SPECRAGE: ROBUST AND GENERALIZABLE MULTI VIEW SPECTRAL REPRESENTATION LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Multi-view representation learning (MvRL) has garnered substantial attention in recent years, driven by the increasing demand for applications that can effectively process and analyze data from multiple sources. In this context, graph Laplacianbased MvRL methods have demonstrated remarkable success in representing multi-view data. However, these methods often struggle with generalization to new data and face challenges with scalability. Moreover, in many practical scenarios, multi-view data is contaminated by noise or outliers. In such cases, modern deeplearning-based MvRL approaches that rely on alignment or contrastive objectives can lead to misleading results, as they may impose incorrect consistency between clear and corrupted data sources. We introduce SpecRaGE, a novel fusion-based framework that integrates the strengths of graph Laplacian methods with the power of deep learning to overcome these challenges. SpecRage uses neural networks to learn parametric mapping that approximates a joint diagonalization of graph Laplacians. This solution bypasses the need for alignment while enabling generalizable and scalable learning of informative and meaningful representations. Moreover, it incorporates a meta-learning fusion module that dynamically adapts to data quality, ensuring robustness against outliers and noisy views. Our extensive experiments demonstrate that SpecRaGE outperforms state-of-the-art methods, particularly in scenarios with data contamination, paving the way for more reliable and efficient multi-view learning. Our code will be made publicly available upon acceptance.

031 032

033

004

010 011

012

013

014

015

016

017

018

019

021

023

024

025

026

027

028

029

1 INTRODUCTION

Multi-view representation learning (MvRL) has become a crucial paradigm in recent years. Its primary objective is to integrate data from multiple sources into a unified representation, which can be used for various tasks such as clustering and classification. The demand for MvRL methods has surged as more applications require analyzing objects or phenomena from diverse perspectives. For instance, streaming platforms rely on the fusion of visual, audio, and textual features to enhance content recommendations, while healthcare systems combine genetic data, imaging, and clinical records to provide a comprehensive view of patient health.

Graph Laplacian methods have demonstrated notable effectiveness in representing both single-view and multi-view data (Ng et al., 2001; Belkin & Niyogi, 2003; Coifman & Lafon, 2006a; Cai et al., 2011; Kumar et al., 2011; Eynard et al., 2012; 2015). This success largely stems from the ability of Laplacian eigenvectors to preserve similarity and capture the underlying cluster structure of the data (Ng et al., 2001; Belkin & Niyogi, 2003). Moreover, by focusing on the eigenvectors corresponding to the smallest eigenvalues, one can uncover the intrinsic low-dimensional manifold structure, as demonstrated in *Laplacian eigenmaps* (Belkin & Niyogi, 2003).

However, in the multi-view setting, traditional graph Laplacian-based methods face two major
limitations: *Generalizability* - the ability to map new test points into the representation space after
the training set has been processed (i.e., out-of-sample extension); *Scalability* – the capability to
process large datasets within a practical time frame efficiently. Current graph Laplacian-based MvRL
approaches often fail to address these two challenges (see discussion in Section 2). These limitations
in generalizability and scalability hinder the full potential of graph Laplacians for representing multi-view data in real-world applications.

054 Another fundamental challenge in the multi-view setting is the presence of noise or outliers in some 055 of the views. For example, in autonomous driving systems that integrate data from multiple sensors, 056 adverse weather conditions can lead to unreliable or noisy data from certain sensors. Most modern 057 deep learning-based MvRL approaches (Andrew et al., 2013; Huang et al., 2019; 2021; Federici et al., 058 2020; Trosten et al., 2021; Xu et al., 2022; Trosten et al., 2023; Wang et al., 2023; Yan et al., 2023) rely on alignment or contrastive objectives that enforces consistency between view-specific representations. These alignment-based methods can become problematic when handling contaminated data, as they 060 may enforce incorrect consistency between clear and degraded data sources. These challenges 061 highlight the pressing need for robust MvRL methods that can effectively handle real-world data 062 imperfections while maintaining the benefits of graph Laplacian-based approaches. 063

064 To address the challenges of Generalizability and Scalability in graph Laplacian-based methods, as well as ensuring *Robustness* in the presence of contaminated views, we propose *SpecRaGE*, a novel 065 fusion-based MvRL framework that combines the strengths of graph Laplacians with the power of 066 deep learning. At its core, SpecRaGE provides a deep-learning solution to the joint diagonalization 067 of Laplacians, by extending SpectralNet (Shaham et al., 2018) to multi-view settings. This joint 068 diagonalization is our key strategy for avoiding the need for alignment between views (see Section 069 4.1 for further discussion). SpecRaGE is inherently scalable, as it is trained on mini-batches in 070 a stochastic manner, allowing it to efficiently process large datasets. Furthermore, the resulting 071 model provides a parametric map that approximates the leading joint eigenvectors of multi-view 072 graph Laplacians. This parametric map enables efficient application to new data, addressing the 073 generalizability challenge of traditional graph Laplacian methods.

Moreover, SpecRaGE incorporates a flexible fusion technique that overcomes the rigid limitations of
 traditional alignment-based methods when dealing with contaminated multi-view data. Specifically,
 SpecRaGE introduces a meta-learning-driven fusion mechanism, which generates sample-specific
 weight vectors, allowing the model to down-weight anomalous or noisy views dynamically.

Our extensive experiments demonstrate that SpecRaGE not only achieves state-of-the-art performance on standard multi-view benchmarks but also significantly outperforms existing methods when dealing with outliers and incomplete views.

The main contributions of this work are: (1) We introduce a generalizable and scalable, graph
 Laplacian-based MvRL framework that extends the power of spectral methods to large-scale multi view data. (2) We propose a meta-learning fusion module that dynamically weights different views,
 providing robust performance in the face of data contamination. (3) We present extensive experimental
 results demonstrating SpecRaGE's superior performance across various benchmarks, particularly in
 scenarios involving outliers and noisy views.

- 880
- 089 090

2 RELATED WORK

091 092

093 Graph Laplacian-based Methods. Various graph Laplacian-based methods (also known as multiview spectral representation learning methods) have been proposed to extract compact and informative 094 representations from multi-view data (Cai et al., 2011; Kumar et al., 2011; Eynard et al., 2012; 2015; 095 Li et al., 2015; Lindenbaum et al., 2020; Yang et al., 2023). These approaches typically aim to learn a 096 fused representation based on multiple graph Laplacians, one for each view. While these methods 097 produce meaningful representations, they often face challenges with generalizability (out-of-sample 098 extension), requiring the recomputation of graph Laplacians to embed new, unseen data into the fused representation space. Some approaches address this issue through out-of-sample extension 100 methods, such as the Nystrom extension (Nyström, 1930) or Geometric Harmonics (Coifman & 101 Lafon, 2006b). However, these techniques were originally developed for single-view data and 102 provide only local extensions, functioning effectively only near existing training points. Additionally, 103 they are computationally intensive and memory-demanding, as they require calculating distances 104 between each new test point and all training points. Furthermore, many graph Laplacian-based MvRL 105 approaches struggle with scalability due to their high computational demands and limited support for sparse matrices, complicating the storage and processing of large Laplacians. These issues with 106 generalizability and scalability hinder the practical application of graph Laplacians for representing 107 multi-view data in real-world settings.

108 **Deep-learning based Methods.** Modern deep-learning-based MvRL methods attempt to design 109 a loss function that is useful for extracting meaningful representations from multi-view data. One 110 category of methods includes deep extensions of the Canonical Correlation Analysis (CCA) algorithm, 111 such as DCCA, DCCAE, and ℓ_0 -DCCA (Hotelling, 1936; Andrew et al., 2013; Wang et al., 2015; 112 Lindenbaum et al., 2021). These methods utilize deep networks to learn a non-linear mapping for two views, maximizing their correlations. Another set of algorithms relies on information-theory-based 113 metrics (Federici et al., 2020; Lin et al., 2022), aiming to maximize the mutual information between 114 views while minimizing redundant information unique to each view. Contrastive learning methods 115 (Trosten et al., 2021; Xu et al., 2022; Yan et al., 2023; Wang et al., 2023) represent another group 116 of deep learning approaches, utilizing a contrastive alignment objective to achieve view-specific 117 alignment. Deep learning techniques have also been applied to address the multi-view spectral 118 clustering problem (Huang et al., 2019; 2021). These methods are closely related to our work, as they 119 also extend SpectralNet to multi-view settings by incorporating alignment objectives between the 120 spectral embeddings from each view. Despite their promising performance, these approaches rely 121 on some form of alignment between the view-specific representations, making them vulnerable to 122 data contamination. Their underlying assumptions regarding data quality and view consistency may 123 struggle in the presence of noise, outliers, or asymmetric corruption across views, as demonstrated in Section 5.3. 124

125 126

127 128

129

3 PRELIMINARIES

3.1 GRAPH LAPLACIAN AND SPECTRALNET

130 **Graph Laplacian.** Given a dataset of n points x_1, x_2, \ldots, x_n , a positive semi-definite and symmet-131 ric affinity matrix W is an $n \times n$ matrix whose $W_{i,j}$ entry represents the similarity between x_i and x_j . The unnormalized graph Laplacian is defined as L = D - W where D is a diagonal matrix in which 132 the element $D_{i,i} = \sum_{j=1}^{n} W_{i,j}$ correspond to the degrees of the points $x_i, i = 1, ..., n$. The eigen-133 vectors corresponding to the smallest eigenvalues of the Laplacian provide valuable low-dimensional 134 representations, capturing structural information like relationships and similarities between data 135 points. These eigenvectors are widely used in applications such as spectral clustering (Ng et al., 136 2001), dimensionality reduction (Belkin & Niyogi, 2003), graph partitioning (Karypis & Kumar, 137 1998), and image segmentation (Shi & Malik, 2000; Melas-Kyriazi et al., 2022). 138

139

SpectralNet. SpectralNet (Shaham et al., 2018) is a deep-learning model that effectively maps single-view data to the approximate eigenvectors of its Laplacian. This enables the performance of spectral clustering on huge single-view datasets since the loss is amenable to parallelized training. Furthermore, the model can be easily and accurately used to generalize the representation to unseen test data. To learn the eigenvectors, *SpectralNet* minimizes the following Rayleigh-quotient loss: Tr $(Y^T LY)$ s.t. $Y^T Y = I$, where $Y \in \mathbb{R}^{n \times k}$ is the network's output and L is the graph Laplacian.

145 146 147

3.2 JOINT DIAGONALIZATION

Definition 1. A set of diagonalizable matrices $A^{(1)}, A^{(2)}, \ldots, A^{(v)}$ is said to be simultaneously diagonalizable if there exists a single invertible matrix V such that for all $1 \le p \le v$, $V^{-1}A^{(p)}V$ is a diagonal matrix.

Intuitively, *joint diagonalization* (or *simultaneous diagonalization*) seeks to find a common basis by which all matrices could be represented in a diagonal form.

However, exact joint diagonalization can be achieved if and only if $A^{(1)}, A^{(2)}, \ldots, A^{(v)}$ commute. Nevertheless, when commutativity cannot be guaranteed, optimizing a joint diagonality criterion and approximating the solution remains possible. This defines the *approximate joint diagonalization* problem.

Approximate Joint Diagonalization of Laplacians. In terms of graph Laplacians, for a set of graph Laplacians $L^{(1)}, L^{(2)}, \ldots, L^{(v)}$ the objective of *joint diagonalization* is to find a set of orthogonal eigenvectors V such that for each Laplacian $L^{(p)}$, the matrix $V^T L^{(p)} V$ is diagonal and contains the eigenvalues of $L^{(p)}$ on the diagonal. It has been demonstrated in (Eynard et al., 2012; 2015), that it is possible to find an approximate joint diagonalization of Laplacians by solving the following objective: $T_{10}(W^{T}\bar{x}W) = V^{T}W = V$

$$\min_{V \in \mathbb{R}^{n \times k}} \operatorname{Tr} \left(V^T \bar{L} V \right), \quad \text{s.t.} \quad V^T V = I,$$
(1)

where \bar{L} represents a form of average of the graph Laplacians. This average can be computed, for example, as a weighted arithmetic mean: $\bar{L} = \sum_{p=1}^{v} w^{(p)} L^{(p)}$ where $w^{(p)}$ represent the contribution of the *p*-th view.

4 SPECRAGE

166

167 168

169 170

171

183

184

197

198

205 206 207

Problem Statement. Let $X = \{\mathcal{X}^{(p)} \in \mathbb{R}^{n \times d_p}\}_{p=1}^{v}$ be a multi-view dataset where *n* is the number of samples, *v* is the number of views, and d_p denotes the dimensionality of samples within the *p*-th view. MvRL aims to leverage the multi-view information to learn a high-quality unified representation that facilitates downstream tasks such as clustering, manifold learning, or classification.

In this section, we introduce SpecRaGE, our MvRL framework designed to tackle the challenges of generalizability, scalability, and robustness to contaminated views. We begin by outlining the rationale behind our method. Next, we detail how SpecRaGE efficiently learns the approximate joint eigenvectors, facilitating generalization to new data. Finally, we explore the meta-learning fusion module, which enhances robustness in the presence of contaminated views. The overall framework of our method is illustrated in Fig. 1.

4.1 RATIONAL

185 To develop a method that effectively harnesses the compelling properties of graph Laplacians while 186 enhancing robustness against data contamination, it is essential to first eliminate the need for alignment 187 and then implement a mechanism that diminishes the influence of low-quality views. To achieve 188 this, SpecRaGE utilizes joint diagonalization of Laplacians, which fuses information from all views 189 into a unified representation that approximates the joint eigenvectors. By prioritizing the fusion of information over the alignment of view representations, we successfully circumvent the challenges 190 associated with rigid view alignment. Furthermore, to enhance robustness, SpecRaGE incorporates a 191 dynamic fusion module that adjusts the weighting of views based on their quality. This effectively 192 reduces the influence of noisy or outlier views. This fusion-based approach not only leverages the 193 advantages of graph Laplacians but also demonstrates resilience to contaminated data. Crucially, we 194 approximate these joint eigenvectors in a scalable manner using deep learning, facilitating efficient 195 processing of large datasets and enabling generalization to new, unseen data. 196

4.2 GENERALIZABLE AND SCALABLE JOINT EIGENVECTORS APPROXIMATION

Let $F_{\theta} : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \times \cdots \times \mathbb{R}^{d_v} \to \mathbb{R}^k$ be a parametric mapping (i.e., a neural network model) that transforms a multi-view input into the corresponding coordinates in a fused representation. That is, for a multi-view input $\hat{x}_i = \left(x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(v)}\right)$, F_{θ} operates such that $y_i = F_{\theta}(\hat{x}_i)$ where $y_i \in \mathbb{R}^k$ is the corresponding coordinates in the fused representation. To perform a deep joint diagonalization of Laplacians for a batch of m multi-view data points, we propose the following loss:

$$\mathcal{L}_{\theta} = \frac{1}{m^2 v} \sum_{p=1}^{v} \sum_{i,j=1}^{m} W_{i,j}^{(p)} \|y_i - y_j\|_2^2,$$
(2)

where $W^{(p)}$ is an $m \times m$ affinity matrix of the *p*-th view, and $W_{i,j}^{(p)}$ represents an affinity measure between $x_i^{(p)}$ and $x_j^{(p)}$. This loss function encourages similar points (as measured by the affinity matrices) to be close in the fused representation space. As one can observe, this loss can be minimized by mapping all points to the same output (e.g., $F_{\theta}(\hat{x}_i) = y_0$ for all *i*). To avoid such a trivial map, an orthogonality constraint is added to the output, i.e.,

215

$$Y^T Y = I_{k \times k},\tag{3}$$

where Y is an $m \times k$ matrix of the outputs whose *i*-th row is y_i .

231

233

234

236

246 247



Figure 1: SpecRaGE architecture. First, v view-specific representations are obtained from v corresponding neural networks. Subsequently, a fusion is performed. Following the fusion, the fused 232 representation undergoes QR decomposition to ensure orthogonality. Finally, the loss function (Eq. 2 or Eq. 5) is computed where the affinity matrices are constructed by the embeddings of the pre-trained Siamese networks (see discussion in Appendix B). All networks weights are updated by the gradients. 235

237 To satisfy the orthogonality constraint, we follow the same technique from *SpectralNet* and construct 238 an orthogonalization layer that computes the QR decomposition of the network output and returns 239 the orthogonal Q matrix. More specifically, let Y denote the $m \times k$ matrix obtained from the fusion 240 step; the weights of this layer are defined to be the matrix R^{-1} from the QR decomposition of \tilde{Y} . 241 The final orthogonal output is then $Y = \tilde{Y}R^{-1}$. For further details about the training process with 242 the orthogonalization layer, see Appendix F. 243

With some mathematical transitions (see Appendix E), the loss in Eq. 2 can be written in the following 244 Rayleigh quotient form: 245

$$\mathcal{L}_{\theta} = \frac{2}{m^2 v} \operatorname{Tr}\left(Y^T \sum_{p=1}^{v} L^{(p)} Y\right),\tag{4}$$

248 where $L^{(p)}$ is the $m \times m$ graph Laplacian of the p-th view. 249

One can observe that this loss is exactly the arithmetic mean version of the approximate joint 250 diagonalization objective in Eq. 1 where each Laplacian is multiplied with the same weight. Our 251 experiment in Appendix D.1 illustrates that our method effectively approximates the joint eigenvectors 252 of the graph Laplacians. 253

254 The choice of affinity measure plays a crucial role in determining the quality of the generated 255 representations. An appropriate affinity measure can enhance the ability of the model to capture the underlying relationships within the data, while an inadequate one may lead to poor representation 256 quality and misinterpretation of the data structure. In Appendix B, we provide a discussion of the 257 technique we employed to construct the affinity matrices. 258

259 Generalizability and Scalability. Once the framework is trained, it provides a mapping function 260 F_{θ} that transforms each multi-view sample directly into its coordinates in the final unified repre-261 sentation, facilitating efficient generalization for new samples from the same distribution. Notably, 262 all our experiments were conducted using test sets, which illustrates the method's generalizability. 263 Additionally, the stochastic mini-batch training in SpecRaGE avoids computing the full Laplacian 264 eigenvectors, enabling scalability for large datasets. For example, SpecRaGE processed the 1-million-265 sample InfiniteMNIST dataset (see Section 5.1) in about 15 minutes on a GTX-1080 Ti, whereas 266 traditional graph Laplacian methods faced out-of-memory errors or much longer runtimes. The overall running time complexity of SpecRaGE is $O(n(k^2 + mv))$, where k is the output dimension, 267 v is the number of views, and m is the batch size. Since k, v, and m are typically much smaller than 268 n and independent of n, our method exhibits near-linear time complexity. In Appendix C, we present 269 a full time complexity analysis.

4.3 ROBUST META-LEARNING FUSION

272 View-specific Representations. As described above, the loss function in 2 operates on a fused representation $Y \in \mathbb{R}^{m \times k}$, derived from a mini-batch of m multi-view samples. To construct 273 this fused representation, we first need to generate intermediate representations for each individual 274 view. The primary goal of these representations is to embed numeric and categorical features 275 into a common space and ensure that all views are of the same size. To extract the view-specific 276 representations, we introduce an individual neural network, $g_{\theta}^{(p)} : \mathbb{R}^{d_p} \to \mathbb{R}^k$ (e.g., an MLP), for 277 each view. Specifically, given a multi-view input $\hat{x}_i = (x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(v)})$, the intermediate 278 279 representations are $y_i^{(1)}, y_i^{(2)}, \dots, y_i^{(v)}$, where $y_i^{(p)} = g_{\theta}^{(p)} (x_i^{(p)})$. For a batch of m multi-view 280 samples, we obtain v matrices $Y^{(1)}, Y^{(2)}, \ldots, Y^{(v)}$, where $Y^{(p)}$ is an $m \times k$ matrix of the outputs, 281 282 and its *i*-th row corresponds to $y_i^{(p)}$. 283

Merging the view-specific representations $Y^{(1)}, Y^{(2)}, \ldots, Y^{(v)}$ into a unified representation Ypresents a significant challenge. A key difficulty in this process arises from the potential presence of contaminated samples in some views, such as outliers or noisy data. In such cases, it is desirable to give less weight to the low-quality views in order to obtain a more accurate representation. Specifically, we need an approach that can evaluate the quality of each view directly from the data and determine its degree of contribution accordingly.

209

Meta-learning for Fusion. To dynamically weight views based on their quality, we introduce 291 another neural network model that takes a multi-view sample as input and predicts a weights vector 292 w_i . This model functions as a meta-learning model as it learns how to adaptively weight the views 293 based on the data, optimizing the fusion process. The size of the weight vector corresponds to 294 the number of views, and each entry in the vector indicates the importance or contribution of the 295 corresponding view. We apply a softmax on the obtained vector to ensure that the vector sums up to one. The overall process is as follows: First, the view-specific representations $y_i^{(1)}, y_i^{(2)}, \ldots, y_i^{(v)}$ 296 297 are obtained from the multi-view sample \hat{x}_i . Then, the concatenated \hat{x}_i is passed through the metalearning model, generating a weights vector w_i . Subsequently, to get the fused representation y_i , the following weighted sum is performed: $y_i = \sum_{p=1}^{v} w_i^{(p)} \cdot y_i^{(p)}$. Section 5.3 demonstrates that this 298 299 300 fusion approach not only achieves state-of-the-art performance on standard multi-view benchmarks 301 but also significantly outperforms existing methods in handling outliers and noisy views.

To provide meaningful feedback to the meta-learning model, we integrate the view-specific contributions into the loss function. For a batch of size m, weight vectors w_1, w_2, \ldots, w_m are first generated using the meta-learning model. We then compute the mean of these weight vectors, \bar{w} , to capture a global representation of the batch's characteristics. This mean weight vector is embedded into the loss function, guiding the meta-learning model in producing appropriate weights for each view. The resulting loss function is defined as:

308

310

311 312

313 314 315

316 317

318

$$\mathcal{L}_{\theta} = \frac{2}{m^2 v} \operatorname{Tr} \left(Y^T \sum_{p=1}^{v} \bar{w}^{(p)} L^{(p)} Y \right), \tag{5}$$

where $\bar{w}^{(p)}$ denotes the mean contribution of the *p*-th view.

5 EXPERIMENTS

5.1 EXPERIMENTAL SETTINGS

Datasets. We assess the performance of SpecRaGE using five well-studied multi-view datasets.
 Our selection prioritizes datasets that exhibit diversity in the types and number of views, as well as
 the number of classes, as depicted in Table 3 in Appendix A. The datasets are listed as follows: (1)
 BDGP (Cai et al., 2012) contains 2500 images of Drosophila embryos divided into five categories
 with two extracted features. One view has 1750-dimensional visual features, and the other view has
 79-dimensional textual features. (2) Reuters is a multilingual dataset comprising more than 11,000

324 articles across six categories and five languages: English, French, German, Italian, and Spanish. We 325 used a subset of this dataset containing 18,758 samples for our analysis. (3) Caltech20 is a subset of 326 2386 examples derived from the object recognition dataset (Fei-Fei et al., 2004), which comprises 327 20 classes viewed from six different perspectives. The dataset encompasses various features such 328 as Gabor features, wavelet moments, CENTRIST features, histogram of oriented gradients, GIST features, and local binary patterns. (4) Handwritten (Asuncion & Newman, 2007) contains 2,000 329 digital images of handwritten numerals ranging from 0 to 9. This dataset employs two types of 330 descriptors: a 240-dimensional pixel average within 2×3 windows, and a 216-dimensional profile 331 correlations method, serving as two distinct viewpoints. (5) InfiniteMNIST is a large-scale variant 332 of MNIST (LeCun et al., 1998), with 1 million samples. The first view contains the original images, 333 while the second adds Gaussian noise ($\sigma = 0.2$) to images from the same class, as in (Trosten et al., 334 2023). Created similarly to affNIST¹, this dataset applies small random affine transformations to 335 expand the original data. More details about the different datasets can be found in Appendix A. 336

337

Baselines. We compared the performance of SpecRaGE with seven multi-view representation 338 learning methods including two classic deep methods (DCCA (Andrew et al., 2013) and DCCAE 339 (Wang et al., 2015)) and five state-of-the-art methods (MvSCN (Huang et al., 2019), MIB (Federici 340 et al., 2020), Multi-VAE (Xu et al., 2021), AECoKM (Trosten et al., 2023), and MetaViewer (Wang 341 et al., 2023)). DCCA and DCCAE are deep extensions of classic correlation analysis. MvSCN 342 learns spectral embeddings for each view while aligning view-specific representations with an MSE 343 objective. MIB uses information theory to separate shared and view-specific information. Multi-344 VAE disentangles visual representations into view-common and view-peculiar variables. AECoKM combines autoencoders with contrastive learning for view alignment. MetaViewer learns to fuse 345 representations by observing reconstruction from unified representations to specific views and 346 employs contrastive learning in its improved version. To ensure fairness, we ran each algorithm ten 347 times on the aforementioned datasets using the same backbones, recording the mean and standard 348 deviation of their performance. For clustering, we employed K-means, while Support Vector Machines 349 (SVM) were used for classification. For alignment-based methods, we concatenated the view-specific 350 representations before applying clustering and classification. More details on hyper-parameters and 351 training are in Appendix F. 352

- Evaluation metrics. For clustering, we employed three widely used metrics: Normalized Mutual
 Information (NMI), Accuracy (ACC), and Adjusted Rand Index (ARI). For classification, we used
 Accuracy, Precision, and F-score. Higher values of these metrics indicate superior performance.
- 356 357 358

359

5.2 REPRESENTATION EVALUATION

Clustering Results. In the clustering experiment, we aimed to evaluate SpecRaGE's ability to capture the underlying cluster structure of the data in comparison to the baseline methods. To achieve this, we applied the K-means algorithm to the final representations learned by each method. Table 1 summarizes the clustering results across the five datasets. SpecRaGE achieves top-performing results on most datasets and evaluation metrics. Fig. 2 provides an additional illustration of the efficacy of SpecRaGE in representing the data and capturing its inherent cluster structure.

366

Classification Results. In the classification experiment, we aimed to assess the quality of the
 learned representations by running a SVM on top of the final embeddings produced by each method.
 We chose to use SVM because it does not introduce any additional non-linear transformations to
 the feature spaces, ensuring that the comparison between different algorithms remains unbiased.
 SVM and linear classifiers are widely accepted as standard benchmarks for evaluating learned
 representations and embeddings (See, for example, (Chen et al., 2020; Federici et al., 2020; Bardes et al., 2021; Wang et al., 2023)).

Table 2 presents the classification results across the five datasets. Notably, SpecRaGE consistently
 achieves top performance in most datasets and evaluation metrics, demonstrating its effectiveness in
 generating discriminative representations for classification tasks.

³⁷⁷

¹https://www.cs.toronto.edu/~tijmen/affNIST/

Datasets	Metrics	DCCA	DCCAE	MvSCN	MIB	Multi-VAE	AECoKM	MetaViewer	SpecRaGE
BDGP	ACC NMI ARI	$\begin{array}{c} 75.8 \pm 1.31 \\ 67.9 \pm 1.15 \\ 49.2 \pm 1.31 \end{array}$	$\begin{array}{c} 80.1 \pm \! 1.50 \\ 73.2 \pm \! 1.04 \\ 55.1 \pm \! 1.02 \end{array}$	$\begin{array}{c} 81.9 \pm \!$	$\begin{array}{c} 86.9 \pm \! 0.75 \\ 80.9 \pm \! 0.91 \\ 74.1 \pm \! 1.82 \end{array}$	$\begin{array}{c} 53.8 \pm \! 6.91 \\ 29.9 \pm \! 6.54 \\ 28.9 \pm \! 4.92 \end{array}$	$\begin{array}{c} 76.6 \pm \! 5.33 \\ 68.7 \pm \! 5.89 \\ 54.7 \pm \! 6.34 \end{array}$	$\frac{90.4}{\underline{86.3}} \pm 0.43 \\ \pm 1.05 \\ \underline{89.3} \pm 1.21$	97.6 ±0.53 93.4 ±0.83 94.1 ±1.29
Reuters	ACC NMI ARI	$\begin{array}{c} 47.9 \pm \! 0.93 \\ 26.6 \pm \! 1.32 \\ 12.7 \pm \! 1.44 \end{array}$	$\begin{array}{c} 42.0 \pm 1.23 \\ 20.3 \pm 1.13 \\ 8.5 \pm 1.06 \end{array}$	$\begin{array}{c} 49.5 \pm 2.50 \\ 27.6 \pm 1.32 \\ 23.1 \pm 1.12 \end{array}$	$\frac{\underline{49.5}}{\underline{28.4}} \pm 1.10 \\ \underline{\underline{28.4}} \pm 1.81 \\ \underline{\underline{24.4}} \pm 1.23 \\ \end{array}$	$\begin{array}{c} 36.6 \pm \! 2.34 \\ 21.1 \pm \! 2.20 \\ 12.4 \pm \! 2.34 \end{array}$	$\begin{array}{c} 26.5 \pm \! 0.52 \\ 16.4 \pm \! 1.54 \\ 10.1 \pm \! 2.54 \end{array}$	$\begin{array}{c} 47.8 \pm \! 0.08 \\ 24.2 \pm \! 0.06 \\ 19.19 \pm \! 0.81 \end{array}$	56.7 ±3.22 38.3 ±2.45 29.6 ±2.64
Caltech20	ACC NMI ARI	$\begin{array}{c} 39.7 \pm 1.10 \\ 51.2 \pm 1.21 \\ 23.5 \pm 2.23 \end{array}$	$\begin{array}{c} 36.1 \pm \! 1.50 \\ 52.4 \pm \! 1.71 \\ 25.3 \pm \! 1.14 \end{array}$	$\begin{array}{c} 42.1 \pm 2.59 \\ 53.7 \pm 3.49 \\ 29.7 \pm 3.26 \end{array}$	$\begin{array}{c} 36.1 \pm 1.23 \\ 47.5 \pm 1.55 \\ 23.2 \pm 1.04 \end{array}$	$\begin{array}{c} 32.4 \pm \! 2.25 \\ 46.2 \pm \! 2.82 \\ 24.8 \pm \! 1.79 \end{array}$	$\begin{array}{c} 42.2 \pm 3.65 \\ 57.6 \pm 0.18 \\ 32.2 \pm 1.20 \end{array}$	$\frac{\underline{45.1}}{\underline{60.9}} \pm 1.33 \\ \underline{\underline{60.9}} \pm 2.15 \\ \underline{35.0} \pm 1.24 \\ \end{array}$	$\begin{array}{c} \textbf{50.1} \pm 2.61 \\ \textbf{66.0} \pm 1.72 \\ \textbf{40.1} \pm 3.49 \end{array}$
Handwritten	ACC NMI ARI	$\begin{array}{c} 58.3 \pm 2.10 \\ 70.2 \pm 1.69 \\ 52.2 \pm 1.33 \end{array}$	$\begin{array}{c} 63.8 \pm 1.21 \\ 75.2 \pm 1.34 \\ 60.8 \pm 1.87 \end{array}$	68.3 ± 4.86 70.9 ± 2.90 59.6 ± 4.25	$\begin{array}{c} 67.5 \pm 2.95 \\ 64.7 \pm 2.18 \\ 48.9 \pm 2.56 \end{array}$	$\begin{array}{c} 74.5 \pm 3.34 \\ 70.0 \pm 3.65 \\ 60.5 \pm 3.15 \end{array}$	$\begin{array}{c} 66.4 \pm 2.34 \\ 70.3 \pm 2.70 \\ 61.4 \pm 2.70 \end{array}$	$\frac{\underline{86.3}}{\underline{78.9}} \pm 2.51 \\ \underline{78.9} \pm 1.44 \\ \underline{72.3} \pm 2.94$	91.9 ±3.51 86.5 ±1.74 83.0 ±1.98
InfiniteMNIST	ACC NMI ARI	$\begin{array}{c} 95.6 \pm 1.09 \\ 90.0 \pm 1.39 \\ 90.1 \pm 1.33 \end{array}$	$\begin{array}{c} 95.1 \pm 1.21 \\ 88.9 \pm 1.41 \\ 89.3 \pm 1.01 \end{array}$	$\begin{array}{c} \underline{99.1} \pm 0.28 \\ \underline{97.7} \pm 1.54 \\ \underline{97.9} \pm 0.48 \end{array}$	$\begin{array}{c} 68.6 \pm 3.95 \\ 67.2 \pm 3.15 \\ 66.9 \pm 3.56 \end{array}$	$\begin{array}{c} 98.1 \pm \! 0.50 \\ 96.0 \pm \! 1.20 \\ 96.4 \pm \! 1.05 \end{array}$	99.3 ±0.20 98.1 ±0.42 98.5 ±0.20	$\begin{array}{c} 80.0 \pm 0.22 \\ 72.3 \pm 0.15 \\ 65.2 \pm 0.24 \end{array}$	$\frac{99.1}{97.5} \pm 0.27 \\ \pm 0.54 \\ \underline{97.9} \pm 0.21$

Table 1: Clustering results of all methods on four datasets. The best result in each row is shown inbold and the second-best is underlined.

Table 2: Classification results of all methods on four datasets. The best result in each row is shown in bold and the second-best is underlined.

Datasets	Metrics	DCCA	DCCAE	MvSCN	MIB	Multi-VAE	AECoKM	MetaViewer	SpecRaGE
BDGP	ACC F-score Precision	$\begin{array}{c} 98.40 \pm \! 0.81 \\ 98.40 \pm \! 0.82 \\ 98.42 \pm \! 0.61 \end{array}$	$\begin{array}{c} 98.65 \pm \! 0.26 \\ 98.50 \pm \! 0.21 \\ 98.63 \pm \! 0.30 \end{array}$	$\begin{array}{c} 98.76 \pm \! 0.10 \\ 98.76 \pm \! 0.10 \\ 98.02 \pm \! 0.15 \end{array}$	$\begin{array}{c} 90.56 \pm \! 1.55 \\ 90.10 \pm \! 0.60 \\ 89.92 \pm \! 1.32 \end{array}$	$\begin{array}{c} 88.87 \pm \! 2.54 \\ 88.87 \pm \! 2.54 \\ 89.07 \pm \! 2.43 \end{array}$	$\begin{array}{c} \underline{98.92} \pm 0.21 \\ \underline{98.92} \pm 0.20 \\ \underline{99.01} \pm 0.20 \end{array}$	$\begin{array}{c} 98.00 \pm 0.11 \\ 98.02 \pm 0.11 \\ 98.59 \pm 0.10 \end{array}$	$\begin{array}{l} \textbf{99.01} \pm 0.30 \\ \textbf{99.00} \pm 0.23 \\ \textbf{99.01} \pm 0.20 \end{array}$
Reuters	ACC F-score Precision	$\begin{array}{c} 74.40 \pm \! 0.92 \\ 74.50 \pm \! 1.10 \\ 74.72 \pm \! 0.92 \end{array}$	$\begin{array}{c} 74.10 \pm \! 0.87 \\ 74.21 \pm \! 0.90 \\ 74.35 \pm \! 1.01 \end{array}$	$\begin{array}{c} \underline{75.52} \pm 0.12 \\ \underline{75.51} \pm 0.12 \\ \underline{75.53} \pm 0.14 \end{array}$	$\begin{array}{c} 71.96 \pm 1.02 \\ 70.78 \pm 0.91 \\ 70.78 \pm 1.10 \end{array}$	$\begin{array}{c} 62.06 \pm 3.40 \\ 59.60 \pm 3.95 \\ 61.31 \pm 3.22 \end{array}$	$\begin{array}{c} 68.01 \pm \! 0.85 \\ 65.21 \pm \! 0.81 \\ 67.24 \pm \! 0.85 \end{array}$	$\begin{array}{c} 59.17 \pm \! 0.10 \\ 51.19 \pm \! 0.09 \\ 56.43 \pm \! 0.12 \end{array}$	76.61 ±1.70 76.60 ±1.77 78.59 ±1.75
Caltech20	ACC F-score Precision	$\begin{array}{c} 72.60 \pm \! 0.51 \\ 40.12 \pm \! 0.50 \\ 46.30 \pm \! 0.44 \end{array}$	$\begin{array}{c} 72.58 \pm \! 0.64 \\ 43.26 \pm \! 0.81 \\ 60.66 \pm \! 0.69 \end{array}$	$\begin{array}{c} 86.54 \pm \! 0.35 \\ 86.75 \pm \! 0.37 \\ 85.32 \pm \! 0.33 \end{array}$	$\begin{array}{c} 73.52 \pm 2.41 \\ 72.52 \pm 2.10 \\ 73.25 \pm 1.93 \end{array}$	$\begin{array}{c} 87.73 \pm \! 0.63 \\ 87.30 \pm \! 0.91 \\ 88.90 \pm \! 1.60 \end{array}$	$\frac{\underline{93.38}}{\underline{93.03}} \pm 1.07 \\ \underline{93.03} \pm 0.97 \\ \underline{93.69} \pm 1.26$	$\begin{array}{c} 92.16 \pm 0.05 \\ 85.72 \pm 0.10 \\ 90.68 \pm 0.15 \end{array}$	$\begin{array}{c} \textbf{95.42} \pm 0.92 \\ \textbf{95.93} \pm 0.93 \\ \textbf{95.44} \pm 0.90 \end{array}$
Handwritten	ACC F-score Precision	$\begin{array}{c} 88.25 \pm \! 0.91 \\ 88.05 \pm \! 1.03 \\ 89.20 \pm \! 0.85 \end{array}$	$\begin{array}{c} 90.01 \pm \! 0.45 \\ 89.92 \pm \! 0.45 \\ 90.48 \pm \! 0.56 \end{array}$	$\begin{array}{c} 96.21 \pm \! 0.21 \\ 96.21 \pm \! 0.21 \\ 96.24 \pm \! 0.21 \end{array}$	$\begin{array}{c} 96.01 \pm 1.04 \\ 96.10 \pm 1.20 \\ 96.03 \pm 0.90 \end{array}$	$\begin{array}{c} 95.36 \pm \! 0.95 \\ 95.36 \pm \! 0.93 \\ 95.38 \pm \! 0.92 \end{array}$	$\begin{array}{c} 96.91 \pm \! 0.65 \\ 96.95 \pm \! 0.63 \\ 96.91 \pm \! 0.34 \end{array}$	$\begin{array}{c} \underline{97.75} \pm 0.21 \\ \underline{97.75} \pm 0.20 \\ \textbf{97.90} \pm 0.19 \end{array}$	$\begin{array}{c} \textbf{97.80} \pm 1.27 \\ \textbf{97.77} \pm 1.27 \\ \underline{97.80} \pm 1.26 \end{array}$
InfiniteMNIST	ACC F-score Precision	$\begin{array}{c} 97.21 \pm \! 0.12 \\ 97.21 \pm \! 0.14 \\ 97.21 \pm \! 0.12 \end{array}$	$\begin{array}{c} 97.60 \pm \! 0.41 \\ 97.60 \pm \! 0.41 \\ 97.60 \pm \! 0.41 \end{array}$	$\begin{array}{c} 99.12 \pm \! 0.10 \\ 99.12 \pm \! 0.10 \\ 99.12 \pm \! 0.10 \end{array}$	$\begin{array}{c} 95.52 \pm \! 0.11 \\ 95.51 \pm \! 0.11 \\ 95.56 \pm \! 0.11 \end{array}$	$\begin{array}{c} 98.75 \pm \! 0.02 \\ 98.75 \pm \! 0.02 \\ 98.78 \pm \! 0.03 \end{array}$	99.50 ±0.09 99.50 ±0.08 99.51 ±0.09	$\begin{array}{c} 95.71 \pm \! 0.10 \\ 95.71 \pm \! 0.10 \\ 95.74 \pm \! 0.15 \end{array}$	$\begin{array}{c} \underline{99.35} \pm 0.04 \\ \underline{99.34} \pm 0.04 \\ \underline{99.35} \pm 0.04 \end{array}$

Visualization. To further demonstrate the effectiveness of the learned unified representation, we utilize the *t*-SNE algorithm on the representation obtained from the validation set during the training. As depicted in Fig. 2, our method successfully separates the data into distinct clusters with increasing training epochs.



Figure 2: Visualization of the unified representation Y during training on the BDGP dataset.

5.3 ROBUSTNESS TO CONTAMINATION EVALUATION

In real-world scenarios, data contamination is a pervasive challenge, whether in the form of noise or anomalous outliers, often introduced by faulty sensors, human errors, or external environmental factors. These types of contamination can severely degrade the performance of multi-view models, making robustness a critical capability. To evaluate SpecRaGE's robustness—a key differentiator of our method—we designed two contamination experiments: one targeting robustness to outliers and the other to Gaussian noise.

In the outlier experiment, global anomalies were randomly injected into one view, following the approach from (Han et al., 2022). For the noise experiment, Gaussian noise with $\sigma = 1.2$ was injected into a portion of the samples in one randomly selected view. We conducted both experiments at contamination ratios of 10%, 20%, 30%, and 40%.

These tests were performed on two real-world datasets and a synthetic 2-view version of the scikit-learn Blobs dataset². We chose to include this simple synthetic 2D data in this experiment, as it allows us to more clearly distinguish the effect of the noise or outliers on the model's performance. In these experiments, we report the relative degradation percentage with respect to the uncontaminated baseline, where no contamination is injected.

As shown in Fig 3, SpecRaGE consistently outperforms other methods in terms of robustness, maintaining stable performance even under high levels of contamination. In both the outlier and noisy view experiments, SpecRaGE demonstrates significantly smaller relative degradation percentages across all datasets, even at the highest contamination ratios. This resilience to extreme contamination underscores the method's effectiveness in handling challenging real-world scenarios, where high levels of noise or outliers are often unavoidable.



Figure 3: Performance decrease of methods with various levels of outliers (left column) and noise (right column) across different datasets. Each line represents a different method, showing the relative performance degradation (%) as the percentage of outliers or noise increases.

483 484 485

481

482

²https://scikit-learn.org/stable/modules/generated/sklearn.datasets. make_blobs.html

486 6 CONCLUSION

488 In this work, we introduced SpecRaGE, a novel fusion-based framework for multi-view spectral 489 representation learning that effectively integrates graph Laplacian methods with deep learning 490 techniques. By efficiently learning a parametric map to uncover joint eigenvectors from diverse graph 491 Laplacians, SpecRaGE addresses the challenges of generalizability and scalability, enabling it to handle large datasets while generalizing to new samples. Moreover, SpecRaGE employs a dynamic 492 fusion technique that enhances robustness against outliers and noise in contaminated multi-view data. 493 Extensive experiments validate that SpecRaGE achieves state-of-the-art performance on standard 494 multi-view benchmarks and significantly outperforms existing methods when faced with outliers and 495 incomplete views. These results highlight SpecRaGE's potential to transform multi-view learning in 496 practical applications where data quality is often compromised. 497

497

527

528

Reproducibility Statement. A complete explanation of SpecRaGE's logic is provided in Section 4
 and Algorithm 1. All key details for reproducibility, including the training methodology with the orthogonality constraint, network backbones for both the view-specific and meta-learning networks, hyperparameters, data splits, operating system, and hardware specifications, are available in Appendix 502
 F. Further discussion on the construction of affinity matrices is found in Appendix B.

- 504 505 REFERENCES
- Charu C Aggarwal, Alexander Hinneburg, and Daniel A Keim. On the surprising behavior of distance
 metrics in high dimensional space. In *Database Theory—ICDT 2001: 8th International Conference London, UK, January 4–6, 2001 Proceedings 8.* Springer, 2001.
- Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *International conference on machine learning*, pp. 1247–1255. PMLR, 2013.
- 512 Arthur Asuncion and David Newman. Uci machine learning repository, 2007.
- Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization
 for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- 515
 516
 517
 518
 518
 519
 519
 510
 510
 511
 511
 512
 512
 513
 514
 514
 515
 515
 516
 517
 517
 518
 517
 518
 518
 519
 519
 510
 510
 511
 512
 514
 514
 515
 516
 517
 517
 518
 517
 518
 517
 518
 518
 517
 518
 517
 518
 517
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
 518
- 518 Kevin Beyer, Jonathan Goldstein, Raghu Ramakrishnan, and Uri Shaft. When is "nearest neighbor"
 519 meaningful? In *Database Theory—ICDT'99: 7th International Conference Jerusalem, Israel,*520 *January 10–12, 1999 Proceedings 7.* Springer, 1999.
- Xiao Cai, Feiping Nie, Heng Huang, and Farhad Kamangar. Heterogeneous image feature integration via multi-modal spectral clustering. In *CVPR 2011*, pp. 1977–1984. IEEE, 2011.
- Xiao Cai, Hua Wang, Heng Huang, and Chris Ding. Joint stage recognition and anatomical annotation
 of drosophila gene expression patterns. *Bioinformatics*, pp. i16–i24, 2012.
 - Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607, 2020.
- Ronald R Coifman and Stéphane Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006a.
- Ronald R Coifman and Stéphane Lafon. Geometric harmonics: a novel tool for multiscale out of-sample extension of empirical functions. *Applied and Computational Harmonic Analysis*, 2006b.
- Davide Eynard, Klaus Glashoff, Michael M Bronstein, and Alexander M Bronstein. Multimodal diffusion geometry by joint diagonalization of laplacians. *arXiv preprint arXiv:1209.2295*, 2012.
- Davide Eynard, Artiom Kovnatsky, Michael M Bronstein, Klaus Glashoff, and Alexander M Bronstein.
 Multimodal manifold analysis by simultaneous diagonalization of laplacians. *IEEE transactions* on pattern analysis and machine intelligence, 37(12):2505–2517, 2015.

549

572

577

- Marco Federici, Anjan Dutta, Patrick Forré, Nate Kushman, and Zeynep Akata. Learning robust representations via multi-view information bottleneck. *arXiv preprint arXiv:2002.07017*, 2020.
- Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training
 examples: An incremental bayesian approach tested on 101 object categories. In 2004 conference
 on computer vision and pattern recognition workshop, pp. 178–178. IEEE, 2004.
- Songqiao Han, Xiyang Hu, Hailiang Huang, Minqi Jiang, and Yue Zhao. Adbench: Anomaly
 detection benchmark. *Advances in Neural Information Processing Systems*, 35:32142–32159,
 2022.
- Harold Hotelling. Relations between two sets of variates. *Biometrika*, 28:321–377, 1936.
- Zhenyu Huang, Joey Tianyi Zhou, Xi Peng, Changqing Zhang, Hongyuan Zhu, and Jiancheng Lv.
 Multi-view spectral clustering network. In *IJCAI*, pp. 4, 2019.
- Zhenyu Huang, Joey Tianyi Zhou, Hongyuan Zhu, Changqing Zhang, Jiancheng Lv, and Xi Peng.
 Deep spectral representation learning from multi-view data. *IEEE Transactions on Image Processing*, 30:5352–5362, 2021.
 - George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular
 graphs. *SIAM Journal on scientific Computing*, 1998.
- Gregory Koch, Richard Zemel, Ruslan Salakhutdinov, et al. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*. Lille, 2015.
- Abhishek Kumar, Piyush Rai, and Hal Daume. Co-regularized multi-view spectral clustering.
 Advances in neural information processing systems, 24, 2011.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, pp. 2278–2324, 1998.
- Yeqing Li, Feiping Nie, Heng Huang, and Junzhou Huang. Large-scale multi-view spectral clustering
 via bipartite graph. In *Proceedings of the AAAI conference on artificial intelligence*, 2015.
- Yijie Lin, Yuanbiao Gou, Xiaotian Liu, Jinfeng Bai, Jiancheng Lv, and Xi Peng. Dual contrastive prediction for incomplete multi-view representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 4447–4461, 2022.
- Ofir Lindenbaum, Arie Yeredor, Moshe Salhov, and Amir Averbuch. Multi-view diffusion maps.
 Information Fusion, 55:127–149, 2020.
- Ofir Lindenbaum, Moshe Salhov, Amir Averbuch, and Yuval Kluger. L0-sparse canonical correlation analysis. In *International Conference on Learning Representations*, 2021.
- Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Deep spectral methods:
 A surprisingly strong baseline for unsupervised semantic segmentation and localization. In
 Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8364–8375, 2022.
- Andrew Ng, Michael Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm.
 Advances in neural information processing systems, 14, 2001.
- Evert J Nyström. Über die praktische auflösung von integralgleichungen mit anwendungen auf randwertaufgaben. 1930.
- Uri Shaham, Kelly Stanton, Henry Li, Boaz Nadler, Ronen Basri, and Yuval Kluger. Spectralnet:
 Spectral clustering using deep neural networks. *arXiv preprint arXiv:1801.01587*, 2018.
- Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 2000.
- Daniel J Trosten, Sigurd Lokse, Robert Jenssen, and Michael Kampffmeyer. Reconsidering representation alignment for multi-view clustering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1255–1265, 2021.

594 595 596	Daniel J Trosten, Sigurd Løkse, Robert Jenssen, and Michael C Kampffmeyer. On the effects of self-supervision and contrastive alignment in deep multi-view clustering. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pp. 23976–23985, 2023.
597 598 599 600	Ren Wang, Haoliang Sun, Yuling Ma, Xiaoming Xi, and Yilong Yin. Metaviewer: Towards a unified multi-view representation. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pp. 11590–11599, 2023.
601 602	Weiran Wang, Raman Arora, Karen Livescu, and Jeff Bilmes. On deep multi-view representation learning. In <i>International conference on machine learning</i> , pp. 1083–1092. PMLR, 2015.
603 604 605 606 607	Jie Xu, Yazhou Ren, Huayi Tang, Xiaorong Pu, Xiaofeng Zhu, Ming Zeng, and Lifang He. Multi-vae: Learning disentangled view-common and view-peculiar visual representations for multi-view clustering. In <i>Proceedings of the IEEE/CVF international conference on computer vision</i> , pp. 9234–9243, 2021.
608 609 610	Jie Xu, Huayi Tang, Yazhou Ren, Liang Peng, Xiaofeng Zhu, and Lifang He. Multi-level feature learning for contrastive multi-view clustering. In <i>Proceedings of the IEEE/CVF conference on computer vision and pattern recognition</i> , pp. 16051–16060, 2022.
611 612 613	Weiqing Yan, Yuanyang Zhang, Chenlei Lv, Chang Tang, Guanghui Yue, Liang Liao, and Weisi Lin. Gcfagg: Global and cross-view feature aggregation for multi-view clustering. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 19863–19872, 2023.
615 616 617	Junchen Yang, Ofir Lindenbaum, Yuval Kluger, and Ariel Jaffe. Multi-modal differentiable un- supervised feature selection. In <i>Uncertainty in Artificial Intelligence</i> , pp. 2400–2410. PMLR, 2023.
618	
600	
620	
621	
622	
623	
624	
625	
626	
627	
628	
629	
630	
631	
632	
634	
634	
636	
637	
638	
639	
640	
641	
6/12	
6/13	
643	
645	
646	
647	

A DATASETS CHARACTERISTICS

 In Table 3, we provide additional information regarding the sample size, number of views, and dimensions of the various datasets.

Table 3: Characteristics of the datasets in our experiments

Dataset	#Samples	#Classes	#Views	#Dimensions
BDGP	2500	5	2	1750; 79
Reuters	18,758	6	5	21531; 24892; 34251; 15506; 11547
Caltech20	2386	20	6	48; 40; 254; 1984; 512; 928
Handwritten	2000	10	2	240; 216
InfiniteMNIST	1,000,000	10	2	784;784

B SELF-SUPERVISED AFFINITY LEARNING.

To perform joint diagonalization of Laplacians, it is necessary to know how to construct an affinity matrix for each view. A widely used approach is to use a Gaussian kernel with a specific scale parameter $\sigma > 0$. For instance, for the *p*-th view, the affinity matrix is defined as follows:

$$W_{i,j}^{(p)} = \begin{cases} \exp\left(-\frac{\|x_i^{(p)} - x_j^{(p)}\|^2}{2\sigma^2}\right), & x_j^{(p)} \text{ is one of the } l \text{ nearest neighbors of } x_i^{(p)}, \\ 0, & \text{otherwise.} \end{cases}$$
(6)

However, Euclidean distance may offer a limited measure of similarity, particularly for highdimensional data (Beyer et al., 1999; Aggarwal et al., 2001). Therefore, instead of directly computing the Euclidean distance between x_i and x_j , we choose to train a SiameseNet (Koch et al., 2015) for each view p, denoted as $h_{\theta_{\text{siamese}}}^{(p)}$. This replaces the Euclidean distance $||x_i^{(p)} - x_j^{(p)}||^2$ in Eq. 6 with $||z_i^{(p)} - z_j^{(p)}||^2$, where $z_i^{(p)} = h_{\theta_{\text{siamese}}}^{(p)} \left(x_i^{(p)}\right)$.

Training a SiameseNet typically requires positive and negative pairs. When labels are provided, pairs from the same class are considered positive, while those from different classes are negative. In our unsupervised setup, we determine positive pairs based on small Euclidean distances between points $x_i^{(p)}$ and $x_j^{(p)}$, and negative pairs otherwise. Specifically, positive pairs are formed from the *l* nearest neighbors of each point, while negative pairs are generated from points that are not included in the nearest neighbors set. The parameters *l* and σ are hyper-parameters, discussed in Appendix F.

The training of the SiameseNets serves as a pre-processing step. Once the SiameseNets are trained,
we use them to construct batch affinity matrices for each view during the training. As in (Shaham
et al., 2018), we empirically found (see the ablation study in Appendix D.2 for the experiments) that
in various datasets, using SiameseNet for the affinities improves the quality of the representations.

702 C TIME COMPLEXITY ANALYSIS

Algo	rithm 1 SpecRaGE
1:]	Input: Multi-view data $\hat{x}_1, \hat{x}_2, \dots, \hat{x}_n$, output dimension k, batch size m, number of epochs T
2:	Dutput: Fused representation $y_1, \ldots, y_n = Y \in \mathbb{R}^{n \times k}$ - the approximate joint eigenvectors.
3: f	for each epoch $t \in \{1, 2, \dots, T\}$ do
4:	for each mini-batch of size m do
5:	Obtain view-specific representations $Y^{(1)}, Y^{(2)}, \dots, Y^{(v)}$
6:	Construct $m \times m$ affinity matrices $W^{(1)}, W^{(2)}, \dots, W^{(v)}$
7:	Compute the corresponding graph Laplacians $L^{(1)}, L^{(2)}, \ldots, L^{(v)}$
8:	Generate m weights vectors w_1, w_2, \ldots, w_m using the meta-learning model
9: 10	Obtain the fused representation Y using the weights vectors And a structure $\tilde{V} D^{-1}$
10:	Apply orthogonality constraint using QK decomposition: $Y = Y R^{-1}$
11.	Compute loss in 5
13:	end for
14: (end for
15: I	Forward propagate $\hat{x}_1, \hat{x}_2, \dots, \hat{x}_n$ and obtain <i>n</i> the outputs $y_1, y_2, \dots, y_n \in \mathbb{R}^k$
T 1	
Ine	time complexity of the algorithm in 1 can be analyzed as follows:
Give	n that:
	• The size of the networks and the number of enochs are constant
	Retch size is m
	• The number of views is v.
	• The total number of samples is <i>n</i> .
	• The output dimension is k.
The 1	running time breakdown per batch is:
	• line 5: $O(mv)$.
	• lines 6-7: $O(m^2 v)$
	$= \lim_{n \to \infty} \Theta(n + 0).$
	• Time 8: $O(mv)$.
	• line 9: $O(mv)$.
	• line 10: $O(mk^2)$.
	• line 11: $O(mv)$.
	• line 12: $O(m^2 v)$.
The	Overall complexity:
	• Per batch: $O(4mv + mk^2 + 2m^2v)$
	$= \operatorname{Peremetry} O(\frac{n}{2} + \frac{n}{2} + \frac{n}{2}$
	• Per epoch: $O(\frac{m}{m} \cdot (4mv + mk^2 + 2m^2v)) = O(n(k^2 + v(2m + 4))) = O(n(k^2 + mv)).$
Now	since m, v , and k are much smaller than n , this method presents almost linear running time
com	plexity.
-	
D	FURTHER EXPERIMENTS
D.1	APPROXIMATION OF THE JOINT EIGENVECTORS

We begin by demonstrating that our unified representation approximates the joint eigenvectors. To show this approximation, we compute the Grassmann distance between the subspace of SpecRaGE's



Figure 4: Grassmann distance as a function of the number of epochs.

771 output and that of the true eigenvectors of the matrix $\sum_{p=1}^{v} L^{(p)}$, referred to as the joint eigenvectors. 772 The squared Grassmann distance measures the sum of squared sines of the angles between two 773 k-dimensional subspaces, yielding values within the [0, k] range. To demonstrate this approximation, 774 we used the BDGP dataset and the scikit-learn 2D Blobs dataset, where the second view is generated 775 by applying a random rotation transformation to the first view. 776

In the BDGP dataset, our model has an output dimension of 5, so the maximum distance between 777 the subspace of SpecRaGE's output and that of the joint eigenvectors is also 5. As shown in Fig. 778 4, this distance significantly decreases early in training and stabilizes around 0.9, indicating that 779 SpecRaGE effectively approximates the joint eigenvectors. For the Blobs dataset, the Grassmann distance quickly approaches zero, indicating that the subspace of SpecRaGE output is extremely 781 close, if not identical, to the subspace spanned by the exact true joint eigenvectors. 782

783 784

808

809

769

770

ABLATION STUDY WITH SIAMESENETS D.2

785 In Eq. 6 we show how we construct the Gaussian kernel affinities, essential for our primary joint 786 diagonalization objective. To ensure these Gaussian kernels effectively encode the similarities 787 between the data points, we opt to train a SiameseNet on each view. After the SiameseNets are trained, we compute the Euclidean distance on the Siamese's output instead of computing it on the 788 feature space. 789

790 The goal of this experiment is to demonstrate that using the Siamese-based affinity significantly 791 improves the quality of the final unified representation, implying that Siamese-based affinity captures more accurately the similarities between the points. In Table 4, it is evident that replacing the 793 Euclidean distance computed on the feature space with the Euclidean distance computed on Siamese representations (referred to as Siamese distance) leads to significant improvements in clustering 794 results. The "Euclidean distance" column presents the clustering results where Euclidean distance is 795 utilized for each view's affinity $(W^{(p)})$, while the "Siamese distance" column showcases the results 796 nity.

797	whe	ere Siamese dista	ance is employed f	for each view's affini
798		Tak	le 4. ACC eluctor	in a magnita with and
799		140	at ACC cluster	ing results with and
800			Detecets	Euclidean distance
801			Datasets	
802			BDGP	82.2 ± 5.68
803			Handwritten	77.9 ± 1.06
804			InfiniteMNIST	86.5 ± 0.10
805				
806				
807	Е	EQUIVALEN	CE OF THE LO	SS FUNCTIONS

d without the SiameseNets.

Siamese distance

 97.4 ± 0.53

 92.0 ± 3.51 99.2 ± 0.27

EQUIVALENCE OF THE LOSS FUNCTIONS E

In Section 4.2, we claim that the loss function in Eq. 2 is equivalent to the joint-diagonalization loss in Eq. 4, that is:

811
812
813

$$\sum_{i,j=1}^{m} W_{i,j} \|y_i - y_j\|_2^2 = 2 \operatorname{Tr} \left(Y^T \sum_{p=1}^{v} L^{(p)} Y \right)$$

Here, we provide a proof for this equation.

Proof. As demonstrated in (Belkin & Niyogi, 2003) the following relation holds:

818 819 820

821 822

823 824

825

827

828 829 830

831 832

833 834

835

836

838

817

810

814

815 816

$$\sum_{i,j=1}^{m} W_{i,j} \|y_i - y_j\|_2^2 = 2 \operatorname{Tr} \left(Y^T L Y \right)$$

Therefore, we can derive the following expression:

 $\sum_{p=1}^{v} \sum_{i,j=1}^{m} W_{i,j}^{(p)} \|y_i - y_j\|_2^2 = 2 \sum_{p=1}^{v} \operatorname{Tr}\left(Y^T L^{(p)} Y\right)$ $= 2\operatorname{Tr}\left(\sum_{p=1}^{v} Y^{T} L^{(p)} Y\right) = 2\operatorname{Tr}\left(Y^{T} \sum_{p=1}^{v} L^{(p)} Y\right)$

F **TECHNICAL DETAILS**

For fairness, we run each of the compared algorithms ten times on the above datasets, recording both the mean and standard deviation of their performance. The same backbones are employed across all methods and datasets. Specifically, we used an MLP with hidden layers of sizes 1024, 1024, and 512 837 for all view-specific networks $g_{\theta}^{(p)}$ across all datasets and methods. The meta-learning model also uses an MLP backbone with three hidden layers, each containing 100 units. 839

Additionally, for certain datasets such as BDGP, InfiniteMNIST, and Reuters, we initially embedded 840 the raw features of each view $\mathcal{X}^{(p)}$ using a pre-trained Autoencoder (AE) with hidden layers of 841 sizes 512, 512, and 2048. This pre-trained AE is used to obtain a lower-dimensional input, typically 842 containing less nuisance information, as shown in (Shaham et al., 2018). 843

844 The size of the output of our model (k), is determined by the number of categories of the data. In the 845 case of coupled view methods like DCCA, DCCAE, and MIB, we present results based on the two 846 best-performing views. For alignment-based methods, the final unified representations are produced 847 through concatenation. Subsequently, clustering and classification tasks are conducted using K-means and SVM classifiers, respectively. Training typically took 35-50 epochs for each dataset. 848

849 **Training with Orthogonalization Layer.** To train the model with the orthogonalization layer, we 850 adopt a technique similar to that used in *SpectralNet* (Shaham et al., 2018). This technique employs a 851 coordinate descent training approach comprising two main optimization steps: the "orthogonalization 852 step" and the "gradient step". Each step involves processing a different mini-batch. 853

854 During the orthogonalization step, we forward a mini-batch through the model and compute the 855 QR decomposition of the fused representation to update the weights of the orthogonalization layer. In the gradient step, we pass another mini-batch through the model and use the orthogonalization 856 weights from the preceding orthogonalization step to orthogonal the output. Following this, we 857 compute the loss and update the network weights via backpropagation, while keeping the weights of 858 the orthogonalization layer unchanged. 859

860 After training the model, all weights, including those of the orthogonal layer, are fixed. Consistent with observations from *SpectralNet*, we empirically found that when employing large mini-batches, 861 the orthogonalization layer can also approximately orthogonalize the output of other mini-batches 862 towards the end of training. For instance, in the Blobs dataset, when a random batch of size 1024 863 passes through the model with the fixed orthogonalization layer (i.e., after training), the resulting

output Y, exhibits approximate orthogonality. This is evident in $Y^T Y$, where the average deviation of off-diagonal elements from 0 is merely 0.04.

Hyper-parameters. In Table 5, we provide a breakdown of the hyperparameters utilized for the various datasets. The #Neighbors parameter corresponds to the number of nearest neighbors (l) used for each data point in the Gaussian kernel, as outlined in Appendix B. Specifically, we select the same number of neighbors for each view $\mathcal{X}^{(p)}$, determined via hyperparameter tuning. The scale parameter (σ) is also utilized for the Gaussian kernel and was chosen for each view $\mathcal{X}^{(p)}$ as the median of the distances from any point in $\mathcal{X}^{(p)}$ to its *l*-nearest neighbors (across all points in the view), resulting in a global scale. The temperature parameter is applied in the softmax function used on the weight vector generated by the meta-learning model. Using a temperature greater than 1 in the softmax function helps to smooth the weight distribution across the views, reducing the absolute dominance of any single view. This is particularly useful when we want to ensure that all views contribute to the final representation, even if some might be slightly less informative.

The initial learning rate (LR) was uniformly set to 10^{-3} for all datasets, with a decay policy in place. This decay policy is contingent on monitoring the validation loss. If the validation loss fails to improve over 10 epochs, the LR is multiplied by 0.1. Furthermore, if the LR decreases to 10^{-8} , training is terminated. Adam optimizer is used for training.

Table 5: Hyper-parameters.

Hyper-params	BDGP	Reuters	Caltech20	Handwritten	InfiniteMNIST
LR	10^{-3}	10^{-3}	10^{-3}	10^{-3}	10^{-3}
Batch size	1024	2048	1024	1024	1024
#Neighbors (l)	22	18	18	22	30
scale (σ)	Global; Median				
Softmax temperture	250	500	250	250	250000

Data Split. For each dataset, we initially divide it into an 80% training set and a 20% testing set. Subsequently, for training, we further divide the training set into a 90% training subset and a 10% validation subset

OS and Hardware. The training procedures were executed on Rocky Linux 9.3, utilizing Nvidia GPUs including GeForce GTX 1080 Ti and A100 80GB PCIe.