

Enabling Out-of-Sample Extension in Semi-Supervised Manifold Alignment through Twin Autoencoders

Jake S. Rhodes
Department of Statistics
Brigham Young University
Provo, Utah
rhodes@stat.byu.edu

Adam G. Rustad
Department of Computer Science
Brigham Young University
Provo, Utah
arusty@byu.edu

Marshall S. Nielsen
Department of Statistics
Brigham Young University
Provo, Utah
msn39@byu.edu

Abstract—Manifold alignment aims to find a shared representation, learning inter-domain relationships across multiple domains while retaining intra-domain structure within each domain. Traditional manifold alignment methods lack mechanisms for out-of-sample extension, which requires re-computation of the full embedding alignment when new data are introduced. This limitation reduces their scalability and generalizability to unseen data, posing challenges for real-world applications. To address these issues, we propose an out-of-sample extension that is generalizable to most semi-supervised manifold alignment methods. Our approach leverages a twin autoencoder architecture for multimodal learning, where each autoencoder is trained on a single modality and regularized using a pre-aligned joint embedding. This architecture enables direct out-of-sample extension of points from either modality while preserving a joint geometric structure that can facilitate cross-domain mappings. We validate our approach on bimodal datasets, demonstrating meaningful alignment preservation.

Index Terms—Manifold Alignment, Guided Representation Learning, Out-of-Sample Extension, Regularized Autoencoders

I. INTRODUCTION

In many cases, high-dimensional data can be well represented in far fewer dimensions [1]. For example, labels from high-quality images can be inferred in smaller resolution copies [2], and information from high-dimensional biological systems [3] can be meaningfully extracted. Manifold learning encompasses a set of methodologies to learn or extract a meaningful, low-dimensional representation of high-dimensional data such that relevant attributes of the data are preserved, such as spatial distances (“close” points remain “close”) or features relevant to some supervised task [4]. Additional challenges arise when data is sourced from multiple domains or modalities. In such contexts, manifold alignment can be applied to learn a single, shared manifold or representation of the data from different domains.

The majority of traditional manifold learning methods (and, by extension, manifold alignment methods) map to a set of fixed points but do not have a mechanism to embed points outside of the original training set. That is, explicit out-of-sample extension is not possible. When new data is introduced,

the entire learning algorithm must be repeated, and a new joint manifold, influenced by the new data, is formed. This is problematic due to computational costs and impracticality for real-world applications requiring dynamic updates. Without an out-of-sample extension mechanism, the alignment model cannot generalize to unseen data, reducing the utility of the learned manifold for predictive or inferential tasks.

Kernel methods serve as one approach to address this limitation, such as the Nyström extension [5], or its variants [6]. These methods estimate embeddings for new data points by expressing them as linear combinations of the eigenvectors of a kernel matrix derived from the training data. Although adapted for certain manifold learning methods such as LLE or Isomap [5], these techniques are mainly tailored to eigenvector-based methods and are not suitable for diffusion-based methods or, in particular, manifold alignment methods designed to bridge modalities. Despite their computational advantages, these kernel methods have limited capacity to map nonlinear relationships [7].

An alternative is to train a neural network to regress onto precomputed embeddings derived from a high-dimensional dataset. For example, parametric extensions, such as parametric t -SNE [8] and parametric UMAP [9], have demonstrated the feasibility of incorporating manifold learning into a neural network framework. However, training solely to predict these embeddings often results in solutions that fit the data but fail to capture meaningful patterns or generalize effectively [10]. To address this, multi-task learning [11], [12], such as that done by autoencoders (AEs), has proven effective in improving generalization by simultaneously learning related tasks.

Approaches such as supervised AEs [11], and geometry-regularized AEs [13] have demonstrated the benefits of training networks in simultaneously predicting both targets (embeddings) and inputs (reconstruction), serving as useful methods to improve generalization to unseen data points, preserving geometric structure, and allowing out-of-sample extension in manifold learning tasks [13]. Thus, we leverage regularized AEs to build an extension for the manifold alignment task, learning a function that both preserves the geometric structure

of the extended embedding (the encoder) and also provides a mapping between modalities (the decoder).

We propose a twin AE architecture designed to address the challenge of integrating new data from either or both domains while maintaining alignment in the shared embedding space. The architecture consists of two separate AEs, each trained within a single modality that is regularized using a pre-trained, aligned, joint embedding.

Although some methods are capable of handling more than two domains, the majority of experiments and their associated metrics in the literature are currently limited to only two domains. Our problem formulation thus focuses on bimodal problems, that is, manifold alignment restricted to two domains. Where our results focus on only two domains, our approach is capable of handling N domains by training a separate AE for each domain.

II. MANIFOLD ALIGNMENT BACKGROUND

Manifold alignment comprises methodologies intended to uncover a shared, underlying joint manifold representation consisting of disparate data sources. This shared representation, or embedding, can serve as a representation for exploring inter-domain relationships and enhancing subsequent machine learning tasks.

Primary challenges in learning shared representations lie in the heterogeneous nature of the datasets and the availability of known correspondence between points in different domains. For instance, consider a dataset of images paired with corresponding text collected from the Web. Each domain—whether images or text—has a distinct notion of similarity or distance, and some images may lack any associated text altogether. However, shared patterns may still be identifiable across these domains, despite the absence of complete one-to-one correspondence [14].

A fundamental assumption underlying manifold alignment is that the intrinsic structures of the different modalities share a similar manifold structure [15]. Consequently, the latent space is expected to preserve geodesic distance to maintain the integrity of the relationships along the shared manifold. In this space, locally similar instances within a dataset, as well as corresponding instances across datasets, are mapped to similar neighborhoods [16].

Manifold alignment methodologies can be broadly categorized according to the availability of correspondences between observations in different domains. Unsupervised manifold alignment is done without known correspondences [17], [18]. Supervised manifold alignment relies on explicit one-to-one correspondences [19] or complete label information, such as class annotations [20]–[22]. Semi-supervised alignment occupies an intermediate position, where partial correspondence information is available. This may involve multiple views of the same entities, such as various translations of a single source document [23], [24], or shared features across domains [14].

III. RELATED WORKS

In Section III-A, we discuss advances in manifold alignment generally. Outside of neural-network-based approaches, the

described methods lack built-in means to extend to new points. In Section III-B, we discuss approaches for out-of-sample extension in manifold learning, which motivates our approach to extend embedding points in the alignment problem.

A. Semi-Supervised Manifold Alignment Methods

Early work by [23] introduced two methods for semi-supervised manifold alignment: one using predefined embedding coordinates and another modifying Laplacian Eigenmaps to preserve known correspondences in a shared low-dimensional space. We call that latter SSMA and include it in our comparisons.

A similar approach, Joint Laplacian Manifold Alignment (JLMA) [25], builds a joint graph Laplacian using domain-specific and inter-domain similarity matrices. By applying Laplacian Eigenmaps to this joint Laplacian, JLMA produces unified embeddings that respect both intra- and inter-domain relationships.

Manifold Alignment via Procrustes Analysis (MAPA) [24] aligns embeddings by minimizing the distance between known correspondences using Procrustes analysis. MAPA has been shown to outperform SSMA in terms of mapping accuracy.

While many methods rely on Laplacian Eigenmaps, [16] used Diffusion Maps [26] for manifold alignment. This method normalizes similarity graphs and applies Iterative Closest Point [27] and geometric hashing [28] to align embeddings across domains, providing an alternative to eigenmap-based approaches.

[14] introduced the Manifold Alignment Generative Adversarial Network (MAGAN), which employs a correspondence loss to guide alignment between known corresponding points, thus adapted to a semi-supervised alignment setting.

Diffusion Transport Alignment (DTA) [29] combines diffusion processes with optimal transport to align manifolds. By constructing domain-specific similarity matrices and cross-domain transition matrices from known correspondences, DTA provides an alignment that outperforms MAGAN, MAPA, and SSMA in mapping known correspondences.

The recent work in [30] introduced two methods, MASH (Manifold Alignment via Stochastic Hopping) and SPUD (Shortest Paths on the Union of Domains), which integrate inter-domain correspondences into an aligned graph structure. SPUD estimates geodesic distances by learning shortest paths in a combined graph, while MASH uses diffusion-based alignment, offering greater robustness to noise and sparsity in the graph.

B. Manifold Learning Extensions

Manifold learning algorithms, such as Diffusion Maps [26], Multidimensional Scaling [31] (MDS), t -SNE [32], and UMAP [33], are designed to uncover low-dimensional representations of high-dimensional data by leveraging eigendecomposition or optimizing latent coordinates based on pairwise similarities. These methods, however, are limited by their inability to generalize to new data points. This arises because they do not parametrize a universal embedding function,

relying instead on the relationships within the training data. As a result, projecting new data points onto the learned manifold is not straightforward.

To address this limitation, methods such as the Nyström extension have been proposed, providing an extension by approximating embeddings for new points using the eigenvectors of a kernel matrix computed on the training set [5]. However, this approach requires recalculating kernel rows for each new data point, leading to computational challenges, particularly for large datasets. Additionally, the quality of the approximations is sensitive to kernel parameters, often necessitating extensive tuning, and must be adapted for specific manifold learning techniques [5]. Other techniques, such as PHATE [3] and landmark MDS [34], approximate embeddings for new points using linear combinations of nearby training data. UMAP also addresses this challenge by initializing new embeddings based on affinities with training points and refining them through gradient descent. While these methods alleviate some computational burdens, they depend on storing either the entire training dataset or a representative subset (e.g., landmarks), which can be memory-intensive.

Neural-network-based approaches such as parametric t -SNE [35] and parametric UMAP [9] employ neural networks to directly optimize an embedding function, enabling the mapping to extend to new inputs. However, they do not offer a straightforward mechanism to approximate the inverse mapping, complicating efforts to assess the fidelity of the embedding.

In contrast, our twin AE-based model inherently supports both forward embedding extensions and inverse mappings. By learning a parametric function, these models eliminate the need to store training data, requiring only the storage of network parameters, making them memory-efficient. This architecture not only scales effectively to large datasets but also facilitates direct evaluation of the embedding’s ability to preserve information, addressing limitations of non-parametric manifold learning techniques.

The general problem of out-of-sample extension in manifold learning naturally extends to manifold alignment. Indeed, many manifold alignment methods are simple extensions of their corresponding manifold learning counterparts [15], [23], [36]. Methods based on neural networks, such as MAGAN, can directly map new points into an embedded space. However, most manifold learning methods face the same out-of-sample problems.

IV. TWIN AUTOENCODER FOR MA EXTENSION

An AE is designed to solve two related problems. First, the model encoder defines a mapping function, f_Z , from the original data space, \mathcal{Z} , to a compressed, lower-dimensional embedding space, $\mathcal{E} \subset R^D$. That is, it learns a function mapping each input to a corresponding representation in a latent space: $f_Z(\mathbf{z}_i) = \mathbf{e}_i$. The decoder, often a mirroring but inverted architecture of the encoder, then maps from the embedded space in an attempt to reconstruct points in the original space: $g_Z(\mathbf{e}_i) = \hat{\mathbf{z}}_i$.

g_Z is learned by minimizing $L(\mathbf{z}_i, \hat{\mathbf{z}}_i)$, where L is typically the mean squared error function. By iteratively learning both functions, the encoder function, f , is guided to retain sufficient information for the decoder to faithfully reconstruct the original data, improving the generalization of the encoder [11]. However, vanilla AEs typically embed into a space that fails to capture the intrinsic data structure, e.g., the embedding is not semantically meaningful [37]. In contrast, regularized AEs (RAEs) can be tailored to guide the encoder to more structurally relevant representations in the embedding space [38], [39].

Here, we introduce a regularized, twin AE architecture to learn a cross-domain embedding function for semi-supervised manifold alignment problems. The two AE models are trained independently of each other but regularized using the same shared embedding space. The combined architecture requires two datasets, $X \subset \mathcal{X}$ and $Y \subset \mathcal{Y}$, and a previously generated aligned embedding generated using an alignment method of choice, E . The sets of points mapped from \mathcal{X} or \mathcal{Y} to the aligned embedding, \mathcal{E} , respectively, are designated as E_X and E_Y . These points serve to guide the regularization of the encoder function. The bottleneck layer of each encoder is guided to map points in each respective domain to the aligned embedding space. To this end, we modify the typical loss function to account for the multi-domain embedding alignment, defining:

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \lambda \mathcal{L}_{\text{align}}, \quad \text{where}$$

$$\mathcal{L}_{\text{recon}} = \frac{1}{n_z} \sum_{i=1}^{n_z} \|\mathbf{z}_i - g_Z(f_Z(\mathbf{z}_i))\|^2, \quad \text{and}$$

$$\mathcal{L}_{\text{align}} = \frac{1}{n_z} \sum_{k=1}^{n_z} \|f_Z(\mathbf{z}_k) - \mathbf{e}_k\|^2$$

Here, λ determines the strength of the influence of the pre-computed aligned embedding. Each autoencoder is trained separately while being regularized using domain-specific points from the aligned embedding. This work primarily focuses on embedding extension, leveraging the decoder to guide representation learning. Future work will refine the decoder’s role in enhancing cross-domain generalization.

V. EXPERIMENTAL SETUP AND RESULTS

A. Data Simulation

To evaluate the AE extension, we simulate multimodal datasets using publicly available datasets from the UCI¹ repository [40]. Five multimodal extension techniques are applied:

- 1) **Random Split:** Features are randomly assigned to domains, with no overlap between domains.
- 2) **Even Split:** Important predictive features are evenly distributed across domains.
- 3) **Skewed Split:** All important predictive features are concentrated in one domain.
- 4) **Additive Noise:** The dataset is duplicated, and Gaussian noise is added to the duplicate.

¹Diabetes, Tic-Tac-Toe, Medical Dataset, Hepatitis, Iris, Audiology, Parkinsons, Seeds, Segmentation, Glass, Heart Disease, Heart Failure, Flare1, Ecoli, Ionosphere, Cancer Data, Hill Valley, and Balance Scale.

MA Method	Correlation \pm SD	Sig. Level
DTA	0.25 \pm 0.12	0.006
JLMA	0.80 \pm 0.32	0.006
MAGAN	0.54 \pm 0.23	0.017
MAPA	0.17 \pm 0.12	0.014
MASH	0.68 \pm 0.33	0.004
MASH-	0.68 \pm 0.33	0.005
NAMA	0.58 \pm 0.27	0.009
SPUD	0.72 \pm 0.29	0.005

TABLE I

In this table, we show a summary of Mantel’s correlations between model embedding and the twin RAE-created embeddings trained on each. The significance level is the correlation coefficient compared to 10,000 random permutations of the embedding distances. The twin RAE embeddings are significantly correlated over all of the models tested, but the correlation is weaker for the DTA and MAPA manifold alignment methods. Scores are averaged across all test datasets and splits.

- 5) **Random Rotation:** A duplicate dataset undergoes a random feature rotation.

B. Evaluation Methods

The evaluation of the extension is two-fold. First, we compare the points embedded by the twin RAE network with those embedded via the full alignment method to assess the fidelity of the function in embedding out-of-sample points. For the comparison, we evaluate the geodesic distance preservation of the embedded points using Mantel’s test [41]. That is, we calculate Mantel’s correlation between the pairwise distance matrices using the extended, embedded points and the embedded points using the full alignment models. The full results are found in Table I. In the table, we see that the correlation between the twin RAE-created embeddings and the test points of the model embeddings is significant across all alignment models, although the correlation is weaker for DTA and MAPA.

Second, we evaluate the integrity of the extended model by comparing supervised tasks trained in the embedding space. To do so, we evaluated the accuracy of random forest and k -nearest neighbor models trained on each of the original domains (baseline scores) and compared them to the accuracy achieved when training on the aligned embedding. To work this approach, an embedding is formed using training points in the set. The twin RAE are then trained on the training set while using the embedded training points to regularize the AE model. The predictive models are trained on the embedded points from the training set and subsequently evaluated on the twin RAE-embedded test points.

This evaluation examines whether the alignment method effectively conveys relevant information from both domains to enhance the supervised task while maintaining the integrity of the predictive model. Using the twin RAE ensures that out-of-sample points are appropriately embedded and do not influence the embedding alignment.

The results in Figure 1 show (top) the random forest out-of-bag accuracies evaluated on the aligned embeddings. These scores are contrasted with the baseline (original data) out-of-bag accuracies in the bottom figure. The scores with much larger differences (MAPA, DTA) correspond to embeddings with much lower Mantel correlations. Most AE-mapped embeddings show nearly identical accuracies compared to the

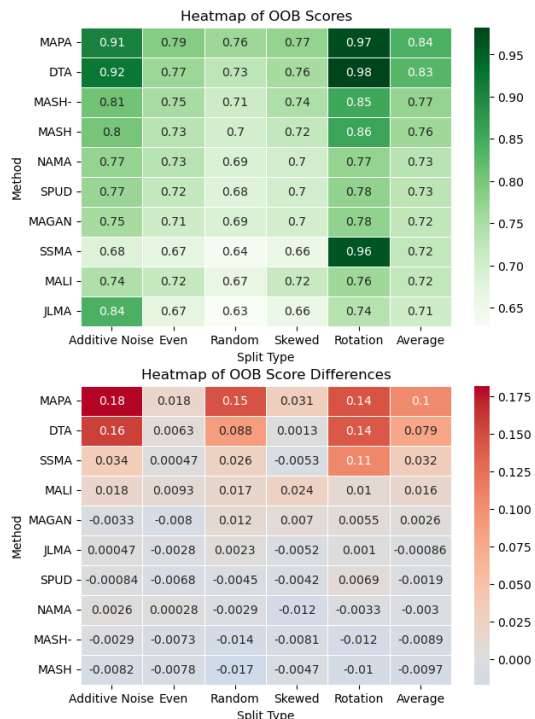


Fig. 1. Here we present heatmaps representing the fidelity of the twin RAE predicted embedding points compared to the embedding created by using the test points in its initialization. In the first graph, random forest out-of-bag scores are created using twin RAE for different manifold alignment methods across five different data splits. In the second graph, the RAE predicted out-of-bag score is compared to the full embedding’s actual out-of-bag score created by the base model. Differences come from five random seeds, 18 test datasets, and 10 manifold alignment methods. The average accuracy for the twin RAE-trained embeddings was found to be 5.46% higher than the model embeddings. The accuracy achieved by the autoencoder is closely comparable to that of the base model.

baselines, suggesting the AE model produces embeddings without losing information related to the supervised task.

VI. CONCLUSION

In this work, we introduced a twin autoencoder architecture to enable out-of-sample extension for semi-supervised manifold alignment. Our approach allows new data points from either domain to be embedded into a pre-aligned joint space without requiring a full recomputation of the alignment model, addressing a key limitation of traditional manifold alignment techniques.

We evaluated our method across multiple bimodal datasets, demonstrating that the twin autoencoder effectively preserves alignment fidelity while maintaining the geometric structure of the shared representation. Experimental results showed significant correlations between the embeddings generated by our approach and those produced by full alignment models, validating the effectiveness of our extension. Furthermore, supervised learning tasks performed on the extended embeddings exhibited comparable predictive accuracy to those trained on the original aligned space, further confirming the utility of our method.

REFERENCES

- [1] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.290.5500.2323>
- [2] B. Zhu, J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen, "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, no. 7697, pp. 487–492, 2018. [Online]. Available: <https://doi.org/10.1038/nature25988>
- [3] K. R. Moon, D. van Dijk *et al.*, "Visualizing structure and transitions in high-dimensional biological data," *Nat. Biotechnol.*, vol. 37, no. 12, pp. 1482–1492, Dec 2019. [Online]. Available: <https://doi.org/10.1038/s41587-019-0336-3>
- [4] J. S. Rhodes, A. Aumon, S. Morin, M. Girard, C. Laroche, E. Brunet-Ratnasingham, A. Pagliuzza, L. Marchitto, W. Zhang, A. Cutler, F. Grand'Maison, A. Zhou, A. Finzi, N. Chomont, D. E. Kaufmann, S. Zandee, A. Prat, G. Wolf, and K. R. Moon, "Gaining biological insights through supervised data visualization," *bioRxiv*, Jan 2024, preprint. [Online]. Available: <https://www.biorxiv.org/content/10.1101/2023.11.22.568384v1>
- [5] Y. Bengio, J. Paiement *et al.*, "Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering," *NeurIPS*, vol. 16, 2003.
- [6] R. R. Coifman and S. Lafon, "Geometric harmonics: A novel tool for multiscale out-of-sample extension of empirical functions," *Appl. Comput. Harmon. Anal.*, vol. 21, no. 1, pp. 31–52, 2006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1063520306000522>
- [7] Y. Aizenbud, A. Bermanis, and A. Averbuch, "Pca-based out-of-sample extension for dimensionality reduction," *arXiv preprint arXiv:1511.00831*, 2015.
- [8] L. Van Der Maaten, "Learning a parametric embedding by preserving local structure," in *Artificial intelligence and statistics*. PMLR, 2009, pp. 384–391.
- [9] T. Sainburg, L. McInnes, and T. Q. Gentner, "Parametric umap embeddings for representation and semisupervised learning," *Neural Computation*, vol. 33, no. 11, pp. 2881–2907, 2021.
- [10] D. Arpit, S. Jastrzebski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, Y. Bengio *et al.*, "A closer look at memorization in deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 233–242.
- [11] L. Le, A. Patterson, and M. White, "Supervised autoencoders: Improving generalization performance with unsupervised regularizers," in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018.
- [12] Y. Zhang and Q. Yang, "An overview of multi-task learning," *National Science Review*, vol. 5, no. 1, pp. 30–43, 09 2017. [Online]. Available: <https://doi.org/10.1093/nsr/nwx105>
- [13] A. F. Duque, S. Morin, G. Wolf, and K. R. Moon, "Extendable and invertible manifold learning with geometry regularized autoencoders," *2020 IEEE International Conference on Big Data (Big Data)*, pp. 5027–5036, 2020. [Online]. Available: <http://doi.org/10.1109/BigData50022.2020.9378049>
- [14] M. Amodio and S. Krishnaswamy, "Magan: Aligning biological manifolds," in *International Conference on Machine Learning*, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:3303339>
- [15] C. Wang, P. Krafft, and S. Mahadevan, "Manifold alignment," in *Manifold Learning: Theory and Applications*, Y. Ma and Y. Fu, Eds. CRC Press, 2011.
- [16] S. Lafon, Y. Keller, and R. Coifman, "Data fusion and multicue data matching by diffusion maps," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 11, pp. 1784–1797, 2006.
- [17] C. Wang and S. Mahadevan, "Manifold alignment without correspondence," in *International Joint Conference on Artificial Intelligence*, 2009. [Online]. Available: <https://api.semanticscholar.org/CorpusID:59769929>
- [18] J. S. Stanley, S. Gigante, G. Wolf, and S. Krishnaswamy, "Harmonic alignment," in *Proceedings of the 2020 SIAM International Conference on Data Mining (SDM)*, 2020, pp. 316–324. [Online]. Available: <https://epubs.siam.org/doi/pdf/10.1137/1.9781611976236.36>
- [19] O. Lindenbaum, A. Yeredor, M. Salhov, and A. Averbuch, "Multi-view diffusion maps," *Information Fusion*, vol. 55, pp. 127–149, 2020.
- [20] C. Wang and S. Mahadevan, "Heterogeneous domain adaptation using manifold alignment," in *International Joint Conference on Artificial Intelligence*, 2011.
- [21] D. Tuia and G. Camps-Valls, "Kernel manifold alignment for domain adaptation," *PLoS One*, vol. 11, no. 2, p. e0148655, Feb. 2016.
- [22] A. F. Duque Correa, M. Lizotte, G. Wolf, and K. R. Moon, "Manifold alignment with label information," in *2023 International Conference on Sampling Theory and Applications (SampTA)*, 2023, pp. 1–6.
- [23] J. Ham, D. Lee, and L. Saul, "Semisupervised alignment of manifolds," in *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, R. G. Cowell and Z. Ghahramani, Eds., vol. R5. PMLR, 06–08 Jan 2005, pp. 120–127, reissued by PMLR on 30 March 2021. [Online]. Available: <https://proceedings.mlr.press/r5/ham05a.html>
- [24] C. Wang and S. Mahadevan, "Manifold alignment using procrustes analysis," in *Proceedings of the 25th International Conference on Machine Learning*, ser. ICML '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 1120–1127. [Online]. Available: <https://doi.org/10.1145/1390156.1390297>
- [25] R. Singh, P. Demetci, G. Bonora, V. Ramani, C. Lee, H. Fang, Z. Duan, X. Deng, J. Shendure, C. Distcheche, and W. S. Noble, "Unsupervised manifold alignment for single-cell multi-omics data," *ACM BCB*, vol. 2020, pp. 1–10, Sep 2020, pubMed-not-MEDLINE.
- [26] R. R. Coifman and S. Lafon, "Diffusion maps," *Appl. Comput. Harmon. Anal.*, vol. 21, no. 1, pp. 5–30, 2006, special Issue: Diffusion Maps and Wavelets. [Online]. Available: <https://doi.org/10.1016/j.acha.2006.04.006>
- [27] A. W. Fitzgibbon, "Robust registration of 2d and 3d point sets," *Image Vis. Comput.*, vol. 21, pp. 1145–1153, 2003. [Online]. Available: <https://api.semanticscholar.org/CorpusID:7576794>
- [28] H. Wolfson and I. Rigoutsos, "Geometric hashing: an overview," *IEEE Computational Science and Engineering*, vol. 4, no. 4, pp. 10–21, 1997.
- [29] A. F. Duque, G. Wolf, and K. R. Moon, "Diffusion transport alignment," in *Advances in Intelligent Data Analysis XXI*, B. Crémilleux, S. Hess, and S. Nijssen, Eds. Cham: Springer Nature Switzerland, 2023, pp. 116–129.
- [30] J. S. Rhodes and A. G. Rustad, "Graph integration for diffusion-based manifold alignment," 2024. [Online]. Available: <https://arxiv.org/abs/2410.22978>
- [31] J. B. Kruskal, "Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis," *Psychometrika*, vol. 29, no. 1, pp. 1–27, Mar 1964. [Online]. Available: <https://doi.org/10.1007/BF02289565>
- [32] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [33] E. Becht, L. McInnes, J. Healy, C.-A. Dutertre, I. W. Kwok, L. G. Ng, F. Ginhoux, and E. W. Newell, "Dimensionality reduction for visualizing single-cell data using umap," *Nature biotechnology*, vol. 37, no. 1, pp. 38–44, 2019.
- [34] V. Silva and J. Tenenbaum, "Global versus local methods in nonlinear dimensionality reduction," *Advances in neural information processing systems*, vol. 15, 2002.
- [35] L. van der Maaten, "Learning a parametric embedding by preserving local structure," in *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, D. van Dyk and M. Welling, Eds., vol. 5. Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA: PMLR, 16–18 Apr 2009, pp. 384–391. [Online]. Available: <https://proceedings.mlr.press/v5/maaten09a.html>
- [36] C. Wang and S. Mahadevan, "Heterogeneous domain adaptation using manifold alignment," in *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Two*, ser. IJCAI'11. AAAI Press, 2011, p. 1541–1546.
- [37] K. Berahmand, F. Daneshfar, E. S. Salehi, Y. Li, and Y. Xu, "Autoencoders and their applications in machine learning: a survey," *Artificial Intelligence Review*, vol. 57, no. 2, p. 28, 2024.
- [38] G. Alain and Y. Bengio, "What regularized auto-encoders learn from the data-generating distribution," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 3563–3593, 2014.
- [39] A. Duque, S. Morin, G. Wolf, and K. Moon, "Geometry regularized autoencoders," *IEEE PAMI*, vol. 45, no. 6, pp. 7381–7394, 2022.
- [40] D. Dua and C. Graff, "Uci machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [41] N. Mantel, "The detection of disease clustering and a generalized regression approach," *Cancer research*, vol. 27, no. 2_Part_1, pp. 209–220, 1967.