Manifold-aligned Neighbor Embedding

Mohammad Tariqul Islam Jason W. Fleischer

Department of Electrical and Computer Engineering, Princeton University, Princeton, NJ 08544

TL;DR

- We introduce a neighbor embedding framework for manifold alignment.
- We demonstrate the efficacy of the framework using a manifold-aligned version of the uniform manifold approximation and projection algorithm.
- We show that our algorithm can learn an aligned manifold that is visually competitive to embedding of the whole dataset.

Motivation

- Neighbor embedding methods, such as t-SNE and UMAP, break down when points are taken from disparate datasets, as there is no link for pairwise comparison.
- With increasing privacy and proprietary concerns, source data cannot be transferred off-site; the issue of disparate datasets is becoming more and more significant.
- Commmon techniques like Procrustes analysis cannot align point-by-point.

We overcome these issues by iteratively optimizing the neighbor analysis for each dataset and jointly embedding the shared points.

Manifold-aligned Neighbor Embedding (MANE)

Assume the individual n-dimensional local datasets $\mathcal{D}^{(m)} = \{\mathbf{z}_i^{(m)}\}$, where $\mathbf{z}_i \in \mathbb{R}^n$ and $m = 1, 2, 3, \ldots, M$, cannot interact with each other. The index $i = 1, 2, 3, \ldots, N_m$ indexes each data point in the dataset. We construct a seeding dataset $\mathcal{D}^{(0)} = \{\mathbf{z}_i\}$ in order to create extended datasets $D^{(m)} = \mathcal{D}^{(0)} \cup \mathcal{D}^{(m)} = \{\mathbf{x}_i^{(m)}\}$. Without loss of generality, we assume that $\mathbf{x}_0^{(m)} = \mathbf{z}_0, \mathbf{x}_1^{(m)} = \mathbf{z}_1, \ldots, \mathbf{x}_{N_0}^{(m)} = \mathbf{z}_N \in \mathcal{D}^{(0)}$ and $\mathbf{x}_{N_0+1}^{(m)}, \mathbf{x}_{N_0+2}^{(m)}, \ldots, \mathbf{x}_{N_0+N_m}^{(m)} \in \mathcal{D}^{(m)}$. This lets us define the MANE framework.

The high-dimensional pairwise relation is given by

$$p_{i,j}^{(m)} = f_H(d_H(\mathbf{x}_i^{(m)}, \mathbf{x}_j^{(m)}), D^{(m)}) \tag{1}$$

and the low-dimensional pairwise relation is given by

$$q_{i,j}^{(m)} = f_L(d_L(\mathbf{y}_i^{(m)}, \mathbf{y}_j^{(m)}), |D^{(m)}|) \tag{2}$$

Finally the relation between the high-dimensional graphs and their joint low-dimensional embedding is established by optimizing the following problem

$$\begin{aligned} &\min_{|D^{(1)}|,\ldots,|D^{(m)}|} \sum_m \sum_{i,j} l(p_{i,j}^{(m)}, q_{i,j}^{(m)}) \\ &\text{s.t.} \\ &\mathbf{y}_i^{(0)} = \mathbf{y}_i^{(1)} = \cdots = \mathbf{y}_i^{(M)}, \forall i = 1, 2, \ldots N_0, \end{aligned} \tag{3}$$

where, $l(\cdot,\cdot)$ is the loss function.

We devised a UMAP implementaion of MANE. The high-dimensional and the low-dimensional realtions were adapted from the UMAP principles for each of the datasets. Then we define the loss funtion

$$l(p_{i,j}^{(m)}, q_{i,j}^{(m)}) = p_{ij}^{(m)} \log\left(\frac{p_{ij}^{(m)}}{q_{ij}^{(m)}}\right) + \left(1 - p_{ij}^{(m)}\right) \log\left(\frac{1 - p_{ij}^{(m)}}{1 - q_{ij}^{(m)}}\right). \tag{4}$$

This loss funcntion is optimized using negative sampling approach. In each step, we sample one positive edge from one of the graphs and apply the attractive force to it. Then we sample negative edges and apply the repulsive forces to these edges. In both cases, we enforce the constraint in Eq. 3.

Embedding Results

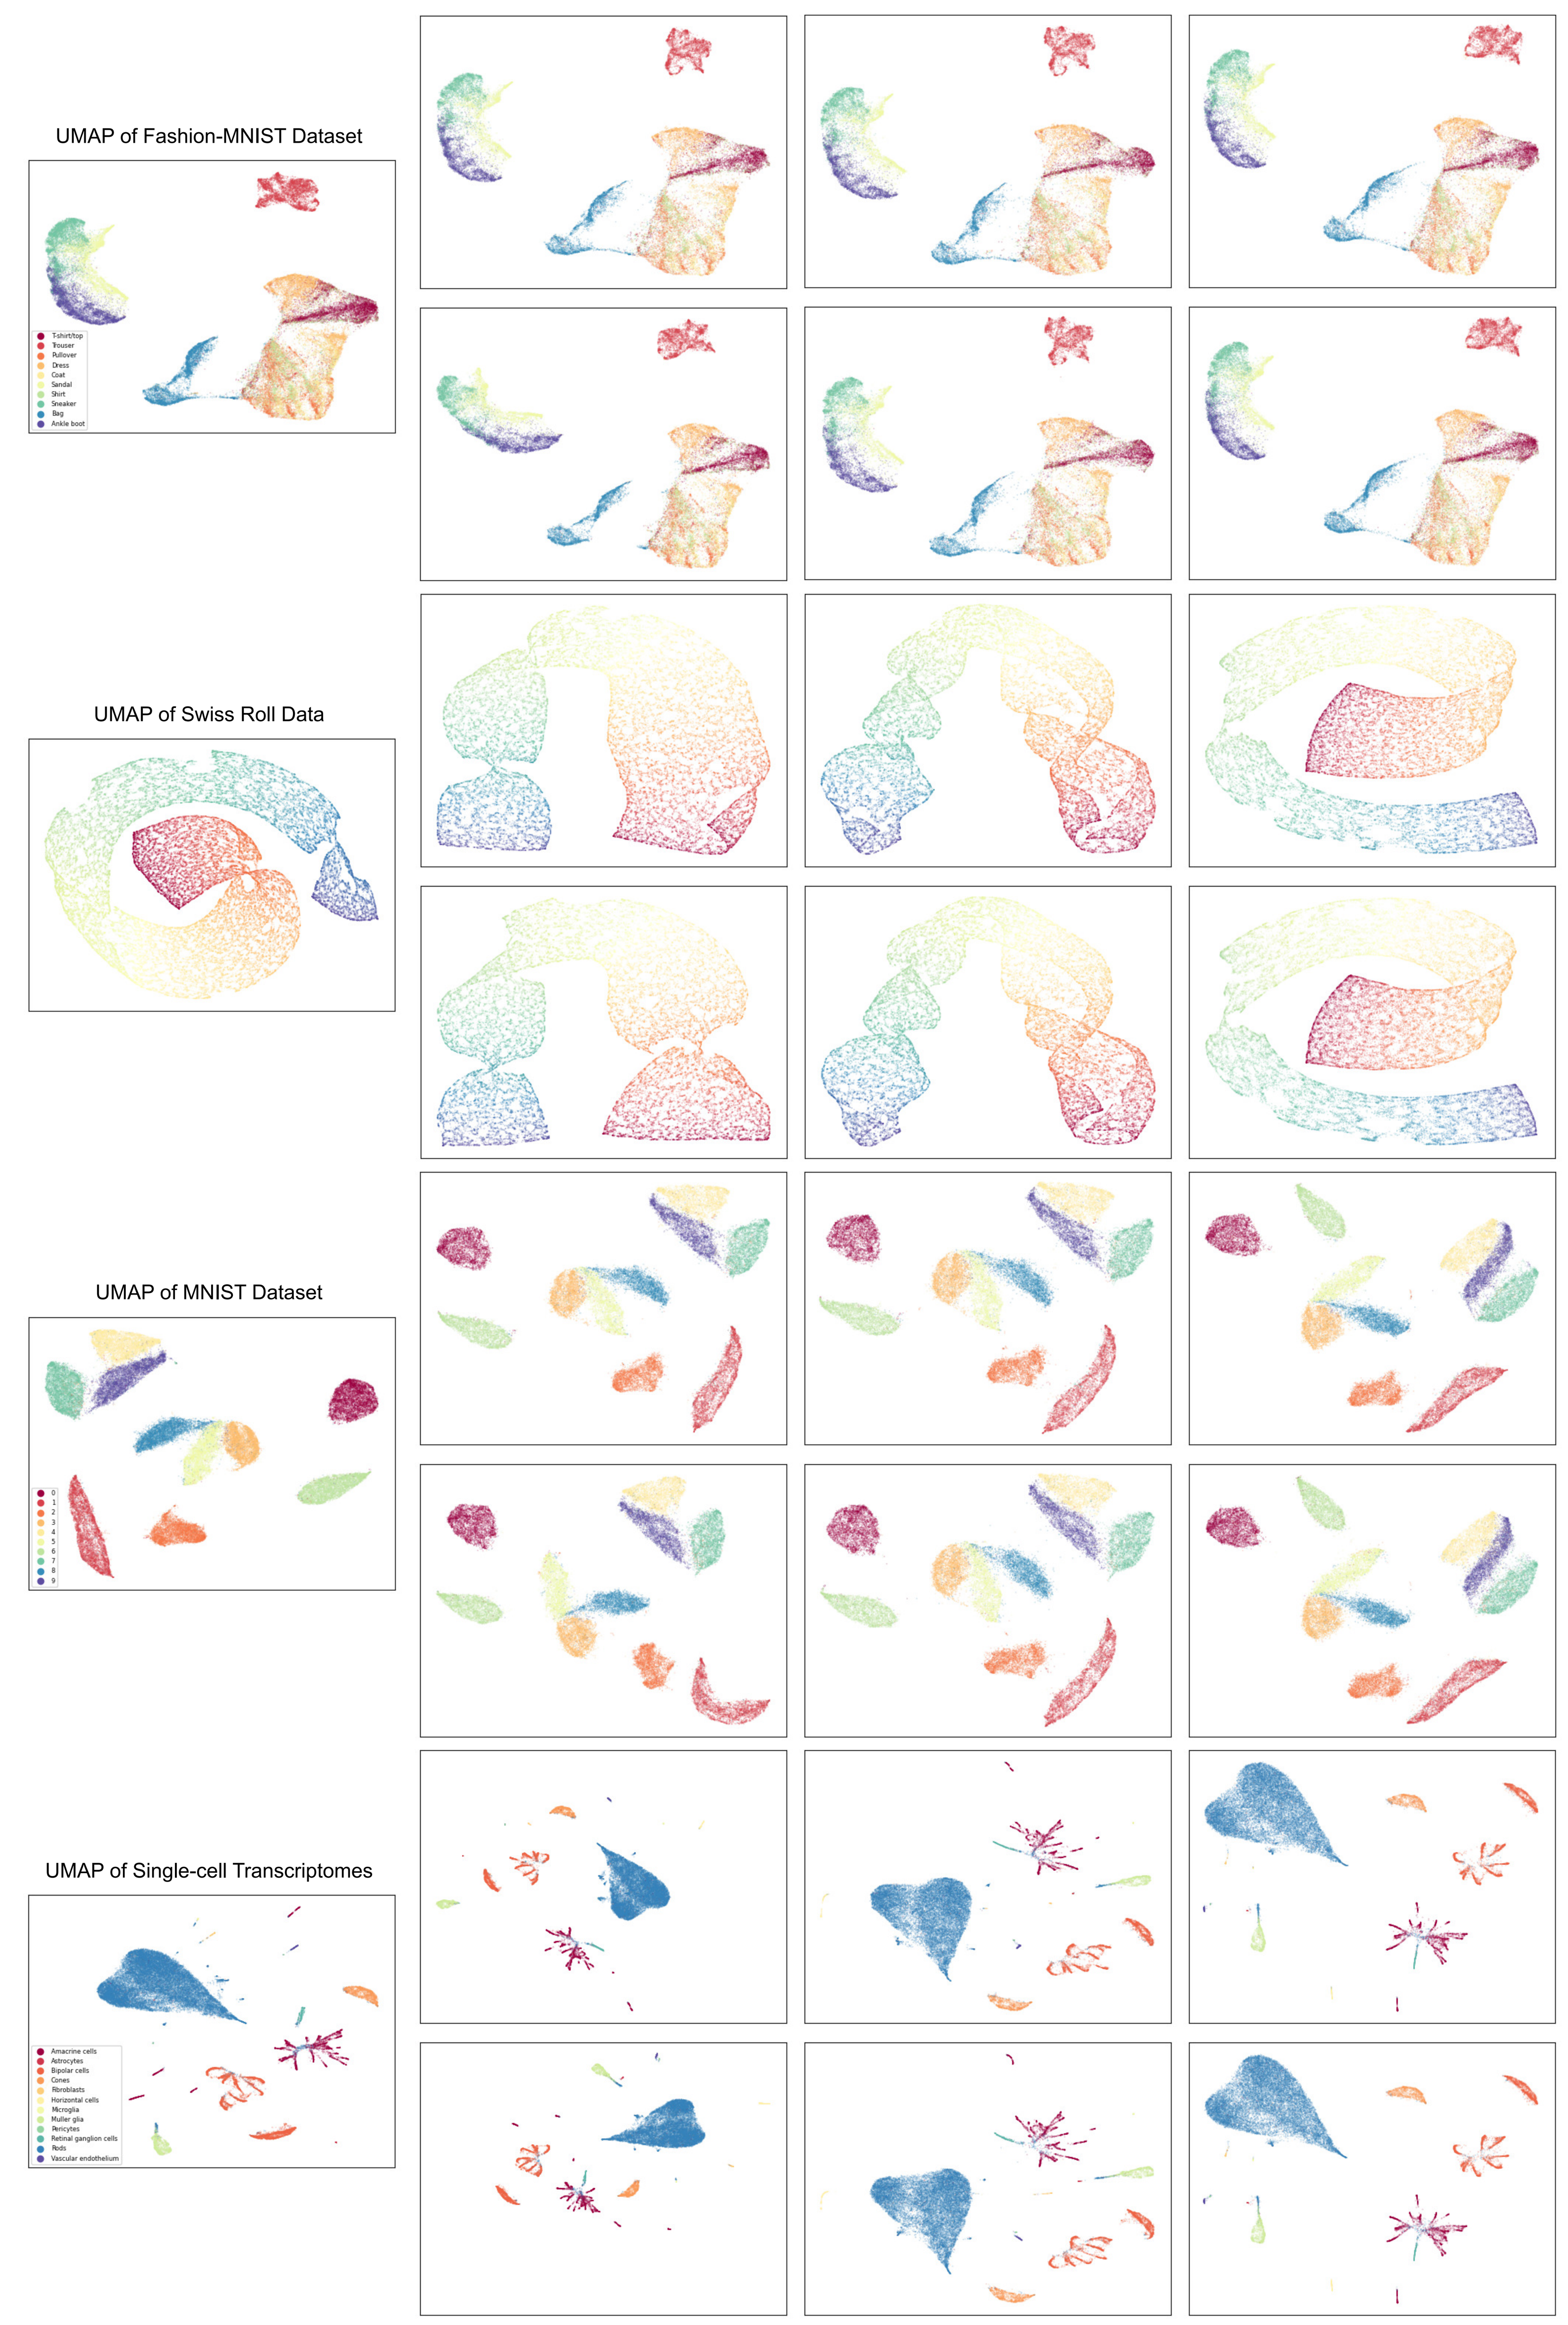

Figure 1. Two-dimensional embeddings of example datasets. For each dataset: (Left) UMAP embeddings. (Right) Top row: embedding of $D^{(1)}$ and bottom row: embedding of $D^{(2)}$ for the individual UMAP, aligned UMAP, and MANE.

Alignment of Shared Data

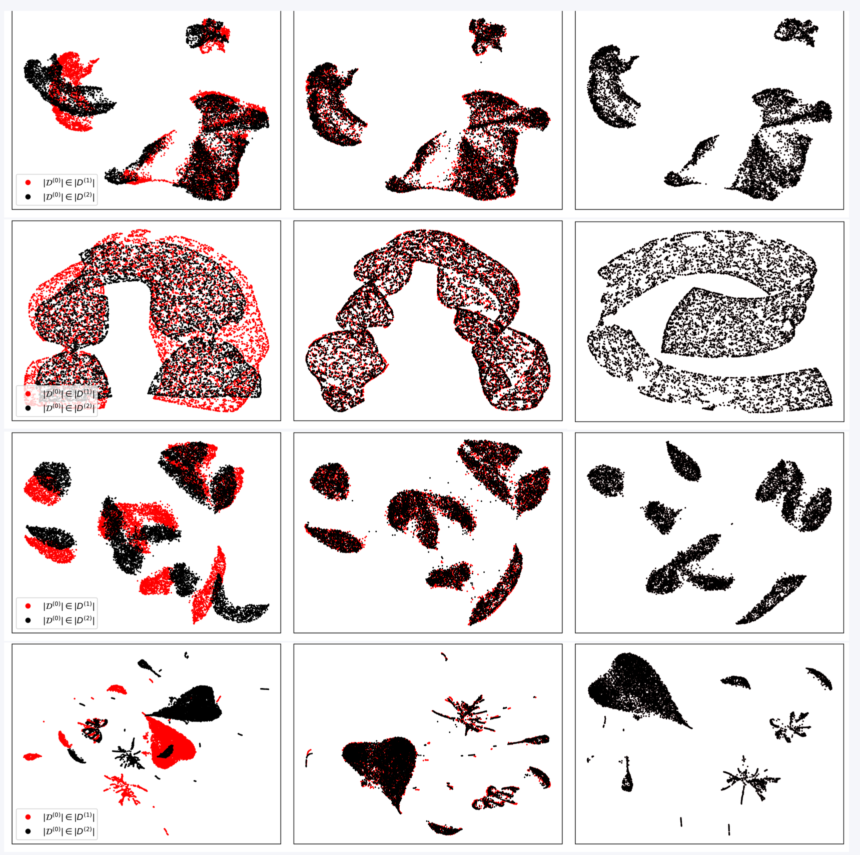

Figure 2. Shared data points from the two-dimensional embeddings of Figure 1. (From left to right) Individual UMAP, aligned UMAP and MANE. MANE shows perfect alignment of shared data.

Other Results

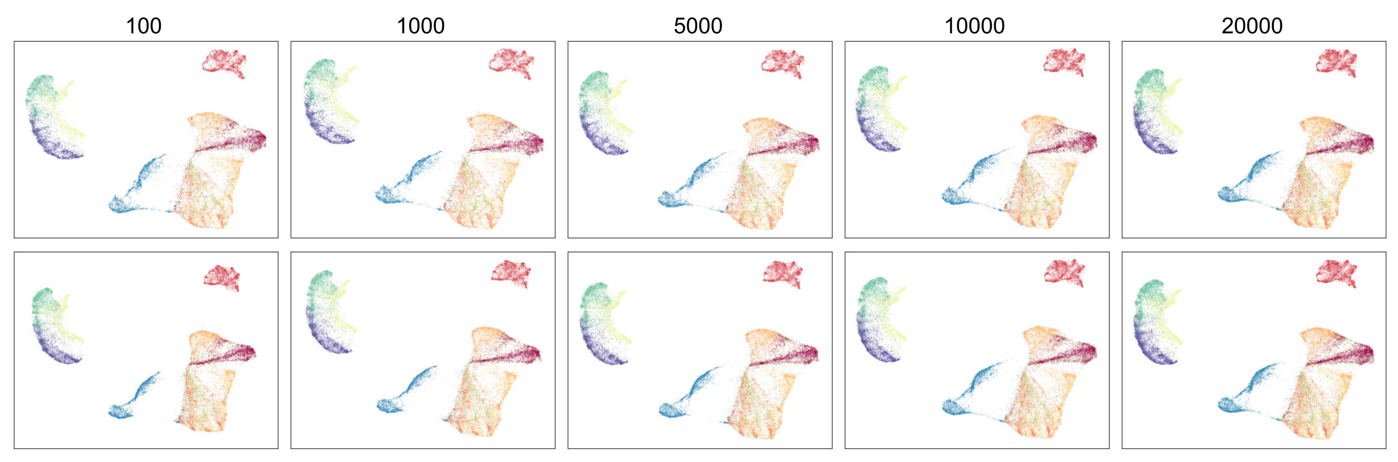

Figure 3. MANE output of Fashion-MNIST data by setting the number of shared points, N_0, to (from left to right) 100, 1000, 5000, 10000, and 20000 respectively. Top row: $|D^{(1)}|$ and bottom row: $|D^{(2)}|$. The fine details are different betwenn $|D^{(1)}|$ and $|D^{(2)}|$ for lower value of N_0.

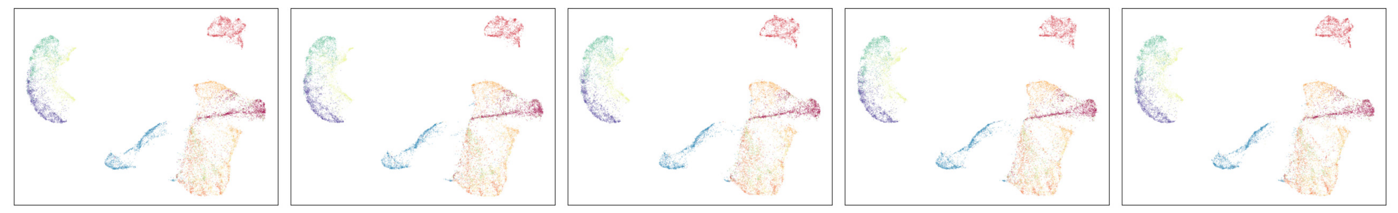

Figure 4. MANE output of Fashion-MNIST data, split into 5 datasets of 14400 data points and 3000 shared points.

References

[1] Yunqian Ma and Yun Fu. *Manifold learning theory and applications*, volume 434. CRC press Boca Raton, 2012.

[2] Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

[3] Chang Wang and Sridhar Mahadevan. Manifold alignment using procrustes analysis. In *Proceedings of the 25th International Conference on Machine Learning*, pages 1120–1127, 2008.