
A Probabilistic Graph Coupling View of Dimension Reduction

Hugues Van Assel
UMPA
ENS Lyon
hugues.van_assel@ens-lyon.fr

Thibault Espinasse
Institut Camille Jordan Lyon 1
Inria Dracula
espinasse@math.univ-lyon1.fr

Julien Chiquet
AgroParisTech
INRAE
julien.chiquet@inrae.fr

Franck Picard
LBMC, CNRS
ENS Lyon
franck.picard@ens-lyon.fr

Abstract

Most popular dimension reduction (DR) methods like t-SNE and UMAP are based on minimizing a cost between input and latent pairwise similarities. Though widely used, these approaches lack clear probabilistic foundations to enable a full understanding of their properties and limitations. To that extent, we introduce a unifying statistical framework based on the coupling of hidden graphs using cross entropy. These graphs induce a Markov random field dependency structure among the observations in both input and latent spaces. We show that existing pairwise similarity DR methods can be retrieved from our framework with particular choices of priors for the graphs. Moreover this reveals that these methods relying on shift-invariant kernels suffer from a statistical degeneracy that explains poor performances in conserving coarse-grain dependencies. New links are drawn with PCA which appears as a non-degenerate graph coupling model.

1 Introduction

Dimensionality reduction (DR) is of central importance when dealing with high-dimensional data [15]. It mitigates the curse of dimensionality, allowing for greater statistical flexibility and less computational complexity. DR also enables visualization that can be of great practical interest for understanding and interpreting the structure of large datasets. Most seminal approaches include Principal Component Analysis (PCA) [31], multidimensional scaling [23] and more broadly kernel eigenmaps methods such as Isomap [4], Laplacian eigenmaps [5] and diffusion maps [12]. These methods share the definition of a pairwise similarity kernel that assigns a high value to close neighbors and the resolution of a spectral problem. They are well understood and unified in the kernel PCA framework [16].

In the past decade, the field has witnessed a major shift with the emergence of a new class of methods. They are also based on pairwise similarities but these are not converted into inner products. Instead, they define pairwise similarity functions in both input and latent spaces and optimize a cost between the two. Among such methods, the Stochastic Neighbor Embedding (SNE) algorithm [18], its heavy-tailed symmetrized version t-SNE [38] or more recent approaches like LargeVis [36] and UMAP [30] are arguably the most used in practice. These will be referred to as *SNE-like* or *neighbor embedding* methods in what follows. They are increasingly popular and now considered as the state-of-art techniques in many fields

[27, 19, 1]. Their popularity is mainly due to their exceptional ability to preserve local structure, *i.e.* close points in the input space have close embeddings, as shown empirically [40]. They also demonstrate impressive performances in identifying clusters [3, 29]. However this is done at the expense of global structure, that these methods struggle in preserving [41, 11] *i.e.* the relative large-scale distances between embedded points do not necessarily correspond to the original ones.

Due to a lack of clear probabilistic foundations, these properties remain mostly empirical. This gap between theory and practice is detrimental as practitioners may rely on strategies that are not optimal for their use case. While recent software developments are making these methods more scalable [8, 33, 28] and further expanding their use, the need for a well-established probabilistic framework is becoming more prominent. In this work we define the generative probabilistic model that encompasses current embedding methods, while establishing new links with the well-established PCA model.

Outline. Consider $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)^\top \in \mathbb{R}^{n \times p}$, an input dataset that consists of n vectors of dimension p . Our task is to embed \mathbf{X} in a lower dimensional space of dimension $q < p$ (typically $q = 2$ for visualization), and we denote by $\mathbf{Z} = (\mathbf{Z}_1, \dots, \mathbf{Z}_n)^\top \in \mathbb{R}^{n \times q}$ the unknown embeddings. The rationale of our framework is to suppose that the observations \mathbf{X} and \mathbf{Z} are structured by two latent graphs with \mathbf{W}_X and \mathbf{W}_Z standing for their n -square weight matrices. As the goal of DR is to preserve the input's structure in the latent space, we propose to find the best low-dimensional representation \mathbf{Z} of \mathbf{X} such that \mathbf{W}_X and \mathbf{W}_Z are close. To build a flexible and robust probabilistic framework, we consider random graphs distributed according to some predefined prior distributions. Our objective is to match the posterior distributions of \mathbf{W}_X and \mathbf{W}_Z . Note that as they share the same dimensionality the latter graphs can be easily compared unlike \mathbf{X} and \mathbf{Z} . The coupling is done with a cross entropy criterion, the minimization of which will be referred to as graph coupling.

In this work, our main contributions are as follows.

- We show that SNE, t-SNE, LargeVis and UMAP are all instances of graph coupling and characterized by different choices of prior for discrete latent structuring graphs (section 3). We demonstrate that such graphs essentially capture conditional independencies among rows through a pairwise Markov Random Field (MRF) model which construction can be found in section 2.
- We uncover the intrinsic probabilistic property explaining why such methods perform poorly on conserving the large scale structure of the data as a consequence of a degeneracy of the MRF when shift invariant kernels are used (theorem 1). Such degeneracy induces the loss of the relative positions of clusters corresponding to the connected components of the posterior latent graphs which distributions are identified (proposition 1). These findings are highlighted by a new initialization of the embeddings that is empirically tested (section 4).
- We show that for Gaussian MRFs, when adapting graph coupling to precision matrices with suitable priors, PCA appears as a natural extension of the coupling problem in its continuous version (theorem 2). Such model does not suffer from the aforementioned degeneracy hence preserves the large-scale structure.

2 Shift-Invariant Pairwise MRF to Model Row Dependencies

We start by defining the distribution of the observations given a graph. The latter takes the form of a pairwise MRF model which as we show is improper (*i.e.* not integrable on $\mathbb{R}^{n \times p}$) when shift-invariant kernels are used. We consider a fixed directed graph $\mathbf{W} \in \mathcal{S}_W$ where:

$$\mathcal{S}_W = \{ \mathbf{W} \in \mathbb{N}^{n \times n} \mid \forall (i, j) \in [n]^2, W_{ii} = 0, W_{ij} \leq n \}$$

Throughout, $(E, \mathcal{B}(E), \lambda_E)$ denotes a measure space where $\mathcal{B}(E)$ is the Borel σ -algebra on E and λ_E is the Lebesgue measure on E .

2.1 Graph Laplacian Null Space

A central element in our construction is the graph Laplacian linear map, defined as follows, where $\mathcal{S}_+^n(\mathbb{R})$ is the set of positive semidefinite matrices.

Definition 1 *The graph Laplacian operator is the map $L: \mathbb{R}_+^{n \times n} \rightarrow \mathcal{S}_+^n(\mathbb{R})$ such that*

$$\text{for } (i, j) \in [n]^2, \quad L(\mathbf{W})_{ij} = \begin{cases} -W_{ij} & \text{if } i \neq j \\ \sum_{k \in [n]} W_{ik} & \text{otherwise.} \end{cases}$$

With an abuse of notation, let $\mathbf{L} = L(\overline{\mathbf{W}})$ where $\overline{\mathbf{W}} = \mathbf{W} + \mathbf{W}^\top$. Let (C_1, \dots, C_R) be a partition of $[n]$ (i.e. the set $\{1, 2, \dots, n\}$) corresponding to the connected components (CCs) of $\overline{\mathbf{W}}$. As well known in spectral graph theory [9], the null space of \mathbf{L} is spanned by the orthonormal vectors $\{\mathbf{U}_r\}_{r \in [R]}$ such that for $r \in [R]$, $\mathbf{U}_r = \left(n_r^{-1/2} \mathbb{1}_{i \in C_r}\right)_{i \in [n]}$ with $n_r = \text{Card}(C_r)$. By the spectral theorem, $\mathbf{U}_{[R]}$ can be completed such that $\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top$ where $\mathbf{U} = (\mathbf{U}_1, \dots, \mathbf{U}_n)$ is orthogonal and $\mathbf{\Lambda} = \text{diag}((\lambda_i)_{i \in [n]})$ with $0 = \lambda_1 = \dots = \lambda_R < \lambda_{R+1} \leq \dots \leq \lambda_n$. In the following, the data is split into two parts: \mathbf{X}_M , the orthogonal projection of \mathbf{X} on $\mathcal{S}_M = (\ker \mathbf{L}) \otimes \mathbb{R}^p$, and \mathbf{X}_C , the projection on $\mathcal{S}_C = (\ker \mathbf{L})^\perp \otimes \mathbb{R}^p$. For $i \in [n]$, $\mathbf{X}_{M,i} = \sum_{r \in [R]} n_r^{-1} \mathbb{1}_{i \in C_r} \sum_{\ell \in C_r} \mathbf{X}_\ell$ hence \mathbf{X}_M stands for the empirical means of \mathbf{X} on CCs, thus modelling the CC positions, while $\mathbf{X}_C = \mathbf{X} - \mathbf{X}_M$ is CC-wise centered, thus modeling the relative positions of the nodes within CCs. We now introduce the probability distribution of these variables.

2.2 Pairwise MRF and Shift-Invariances

In this work, the dependency structure among rows of the data is governed by a graph. The strength of the connection between two nodes is given by a symmetric function $k: \mathbb{R}^p \rightarrow \mathbb{R}_+$. We consider the following pairwise MRF unnormalized density function:

$$f_k: (\mathbf{X}, \mathbf{W}) \mapsto \prod_{(i,j) \in [n]^2} k(\mathbf{X}_i - \mathbf{X}_j)^{W_{ij}}. \quad (1)$$

As we will see shortly, the above is at the heart of DR methods based on pairwise similarities. Note that as k measures the similarity between couples of samples, f_k will take high values if the rows of \mathbf{X} vary smoothly on the graph \mathbf{W} . Thus we can expect \mathbf{X}_i and \mathbf{X}_j to be close if there is an edge between node i and node j in \mathbf{W} . A key remark is that f_k is kept invariant by translating \mathbf{X}_M . Namely for all $\mathbf{X} \in \mathbb{R}^{n \times p}$, $f_k(\mathbf{X}, \mathbf{W}) = f_k(\mathbf{X}_C, \mathbf{W})$. This invariance results in $f_k(\cdot, \mathbf{W})$ being non integrable on $\mathbb{R}^{n \times p}$, as we see with the following example.

Gaussian kernel. For a positive definite matrix $\mathbf{\Sigma} \in \mathcal{S}_{++}^n(\mathbb{R})$, consider the Gaussian kernel $k: \mathbf{x} \mapsto e^{-\frac{1}{2} \|\mathbf{x}\|_{\mathbf{\Sigma}}^2}$ where $\mathbf{\Sigma}$ stands for the covariance among columns. One has:

$$\log f_k(\mathbf{X}, \mathbf{W}) = - \sum_{(i,j) \in [n]^2} W_{ij} \|\mathbf{X}_i - \mathbf{X}_j\|_{\mathbf{\Sigma}}^2 = - \text{tr}(\mathbf{\Sigma}^{-1} \mathbf{X}^\top \mathbf{L} \mathbf{X}) \quad (2)$$

by property of the graph Laplacian (definition [1]). In this case, it is clear that due to the rank deficiency of \mathbf{L} , $f_k(\cdot, \mathbf{W})$ is only $\lambda_{\mathcal{S}_C}$ -integrable. In general DR settings one does not want to rely on Gaussian kernels only. A striking example is the use of the Student kernel in t-SNE [38]. Heavy-tailed kernels appear useful when the dimension of the embeddings is smaller than the intrinsic dimension of the data [20]. Our contribution provides flexibility by extending the previous result to a large class of kernels, as stated in the following theorem.

Theorem 1 *If k is $\lambda_{\mathbb{R}^p}$ -integrable and bounded above $\lambda_{\mathbb{R}^p}$ -almost everywhere then $f_k(\cdot, \mathbf{W})$ is $\lambda_{\mathcal{S}_C}$ -integrable.*

We refer to appendix [A.1] for the proof. We can now define a distribution on $(\mathcal{S}_C, \mathcal{B}(\mathcal{S}_C))$, where $\mathcal{C}_k(\mathbf{W}) = \int f_k(\cdot, \mathbf{W}) d\lambda_{\mathcal{S}_C}$:

$$\mathbb{P}_k(d\mathbf{X}_C | \mathbf{W}) = \mathcal{C}_k(\mathbf{W})^{-1} f_k(\mathbf{X}_C, \mathbf{W}) \lambda_{\mathcal{S}_C}(d\mathbf{X}_C). \quad (3)$$

Remark 1 *Kernels may have node-specific bandwidths τ , set during a pre-processing step, giving $f_k(\mathbf{X}, \mathbf{W}) = \prod_{(i,j)} k((\mathbf{X}_i - \mathbf{X}_j)/\tau_i)^{W_{ij}}$. Note that such bandwidth does not affect the degeneracy of the distribution and theorem [1] still holds.*

Between-Rows Dependency Structure. By symmetry of k , reindexing gives: $f_k(\mathbf{X}, \mathbf{W}) = \prod_{j \in [n]} \prod_{i \in [j]} k(\mathbf{X}_i - \mathbf{X}_j)^{W_{ij}}$. Hence distribution [3] boils down to a pairwise MRF model [10] with respect to the undirected graph $\overline{\mathbf{W}}$, \mathcal{C}_k playing the role of the partition function. Note that since f_k (Equation [1]) trivially factorize according to the cliques of $\overline{\mathbf{W}}$, the Hammersley-Clifford theorem ensures that the rows of \mathbf{X}_C satisfy the local and global Markov properties with respect to $\overline{\mathbf{W}}$.

2.3 Uninformative Model for CC-wise Means

We showed that the MRF [1] is only integrable on \mathcal{S}_C , the definition of which depends on the connectivity structure of \mathbf{W} . As we now demonstrate, the latter MRF can be seen as a limit of proper distributions on $\mathbb{R}^{n \times p}$, see e.g. [35] for a similar construction in the Gaussian case. We introduce the Borel function $f^\varepsilon(\cdot, \mathbf{W}) : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}_+$ for $\varepsilon > 0$ such that for all $\mathbf{X} \in \mathbb{R}^{n \times p}$, $f^\varepsilon(\mathbf{X}, \mathbf{W}) = f^\varepsilon(\mathbf{X}_M, \mathbf{W})$. To allow f^ε to become arbitrarily non-informative, we assume that for all $\mathbf{W} \in \mathcal{S}_W$, $f^\varepsilon(\cdot, \mathbf{W})$ is $\lambda_{\mathcal{S}_M}$ -integrable for all $\varepsilon \in \mathbb{R}_+$ and $f^\varepsilon(\cdot, \mathbf{W}) \xrightarrow{\varepsilon \rightarrow 0} 1$ almost everywhere. We now define the conditional distribution on $(\mathcal{S}_M, \mathcal{B}(\mathcal{S}_M))$ as follows:

$$\mathbb{P}^\varepsilon(d\mathbf{X}_M | \mathbf{W}) = \mathcal{C}^\varepsilon(\mathbf{W})^{-1} f^\varepsilon(\mathbf{X}_M, \mathbf{W}) \lambda_{\mathcal{S}_M}(d\mathbf{X}_M) \quad (4)$$

where $\mathcal{C}^\varepsilon(\mathbf{W}) = \int f^\varepsilon(\cdot, \mathbf{W}) d\lambda_{\mathcal{S}_M}$. With this at hand, the joint conditional is defined as the product measure of [3] and [4] over the row axis, the integrability of which is ensured by the Fubini-Tonelli theorem. In the following we will use the compact notation $\mathcal{C}_k^\varepsilon(\mathbf{W}) = \mathcal{C}_k(\mathbf{W}) \mathcal{C}^\varepsilon(\mathbf{W})$ for the joint normalizing constant.

Remark 2 *At the limit $\varepsilon \rightarrow 0$ the above construction amounts to setting an infinite variance on the distribution of the empirical means of \mathbf{X} on CCs, thus losing the inter-CC structure.*

As an illustration, one can structure the CCs' relative positions according to a Gaussian model with positive definite precision $\varepsilon \boldsymbol{\Theta} \in \mathcal{S}_{++}^R(\mathbb{R})$, as it amounts to choosing $f^\varepsilon : \mathbf{X} \rightarrow \exp(-\frac{\varepsilon}{2} \text{tr}(\boldsymbol{\Sigma}^{-1} \mathbf{X}^\top \mathbf{U}_{[r]} \boldsymbol{\Theta} \mathbf{U}_{[r]}^\top \mathbf{X}))$ such that: $\text{vec}(\mathbf{X}_M) | \boldsymbol{\Theta} \sim \mathcal{N}(\mathbf{0}, (\varepsilon \mathbf{U}_{[r]} \boldsymbol{\Theta} \mathbf{U}_{[r]}^\top)^{-1} \otimes \boldsymbol{\Sigma})$ where \otimes denotes the Kronecker product.

3 Graph Coupling as a Unified Objective for Pairwise Similarity Methods

In this section, we show that neighbor embedding methods can be recovered in the presented framework. They are obtained, for particular choices of graph priors, at the limit $\varepsilon \rightarrow 0$ when f^ε becomes non informative and the CCs' relative positions are lost.

We now turn to the priors for \mathbf{W} . Our methodology is similar to that of constructing conjugate priors for distributions in the exponential family [39], notably we insert the cumulant function $\mathcal{C}_k^\varepsilon$ (i.e. normalizing constant of the conditional) as a multivariate term of the prior. We consider different forms: binary (B), unitary out-degree (D) and n -edges (E), relying on an additional term (Ω) to constraint the topology of the graph. For a matrix \mathbf{A} , A_{i+} denotes $\sum_j A_{ij}$ and A_{++} denotes $\sum_{ij} A_{ij}$. In the following, $\boldsymbol{\pi}$ plays the role of the edge's prior. The latter can be leveraged to incorporate some additional information about the dependency structure, for instance when a network is observed [26].

Definition 2 *Let $\boldsymbol{\pi} \in \mathbb{R}_+^{n \times n}$, $\varepsilon \in \mathbb{R}_+$, $\alpha \in \mathbb{R}$, k satisfies the assumptions of theorem [1] and $\mathcal{P} \in \{B, D, E\}$. For $\mathbf{W} \in \mathcal{S}_W$ we introduce:*

$$\mathbb{P}_{\mathcal{P}, k}^\varepsilon(\mathbf{W}; \boldsymbol{\pi}, \alpha) \propto \mathcal{C}_k^\varepsilon(\mathbf{W})^\alpha \Omega_{\mathcal{P}}(\mathbf{W}) \prod_{(i,j) \in [n]^2} \pi_{ij}^{W_{ij}}$$

where $\Omega_B(\mathbf{W}) = \prod_{ij} \mathbb{1}_{W_{ij} \leq 1}$, $\Omega_D(\mathbf{W}) = \prod_i \mathbb{1}_{W_{i+} = 1}$ and $\Omega_E(\mathbf{W}) = \mathbb{1}_{W_{++} = n} \prod_{ij} (W_{ij}!)^{-1}$.

When $\alpha = 0$, the above no longer depends on ε and k . We will use the compact notation $\mathbb{P}_{\mathcal{P}}(\mathbf{W}; \boldsymbol{\pi}) = \mathbb{P}_{\mathcal{P},k}^{\varepsilon}(\mathbf{W}; \boldsymbol{\pi}, 0)$. Note that by $\mathbf{W} \sim \mathbb{P}_{\mathcal{P}}(\cdot; \boldsymbol{\pi})$ we have the following simple Bernoulli (\mathcal{B}) and multinomial (\mathcal{M}) distributions, where matrix or vector division is to be understood as element-wise.

- If $\mathcal{P} = B$, $\forall (i, j) \in [n]^2$, $W_{ij} \stackrel{\perp}{\sim} \mathcal{B}(\pi_{ij}/(1 + \pi_{ij}))$.
- If $\mathcal{P} = D$, $\forall i \in [n]$, $\mathbf{W}_i \stackrel{\perp}{\sim} \mathcal{M}(1, \boldsymbol{\pi}_i/\pi_{i+})$.
- If $\mathcal{P} = E$, $\mathbf{W} \sim \mathcal{M}(n, \boldsymbol{\pi}/\pi_{++})$.

We now show that the posterior distribution of the graph given the observations takes a simple form when the distribution of CC empirical means \mathbf{X}_M diffuses *i.e.* when $\varepsilon \rightarrow 0$ (a proof of the following result can be found in appendix [A.2](#)). In the following, \odot stands for the Hadamard product and \mathcal{D} for the convergence in distribution.

Proposition 1 *Let $\boldsymbol{\pi} \in \mathbb{R}_+^{n \times n}$, k satisfies the assumptions of theorem [1](#) with $\mathbf{K}_X = (k(\mathbf{X}_i - \mathbf{X}_j))_{(i,j) \in [n]^2}$ and $\mathcal{P} \in \{B, D, E\}$. If $\mathbf{W}^{\varepsilon} \sim \mathbb{P}_{\mathcal{P},k}^{\varepsilon}(\cdot; \boldsymbol{\pi}, 1)$ then*

$$\mathbf{W}^{\varepsilon} | \mathbf{X} \xrightarrow[\varepsilon \rightarrow 0]{\mathcal{D}} \mathbb{P}_{\mathcal{P}}(\cdot; \boldsymbol{\pi} \odot \mathbf{K}_X).$$

Remark 3 *For all $\mathbf{W} \in \mathcal{S}_W$, $\mathcal{C}^{\varepsilon}(\mathbf{W})$ diverges as $\varepsilon \rightarrow 0$, hence the graph prior (definition [2](#)) is improper at the limit. This compensates for the uninformative diffuse conditional and allows to retrieve a well-defined tractable posterior limit.*

3.1 Retrieving Well Known DR Methods

We now provide a unified view of neighbor embedding objectives as a coupling between graph posterior distributions. To that extent we derive the cross entropy associated with the various graph priors at hand. In what follows, k_x and k_z satisfy the assumptions of theorem [1](#) and we denote by \mathbf{K}_X and \mathbf{K}_Z the associated kernel matrices on \mathbf{X} and \mathbf{Z} respectively. For both graph priors we consider the parameters $\boldsymbol{\pi} = \mathbf{1}$ and $\alpha = 1$. For $(\mathcal{P}_X, \mathcal{P}_Z) \in \{B, D, E\}^2$, we introduce the cross entropy between the limit posteriors at $\varepsilon \rightarrow 0$,

$$\mathcal{H}_{\mathcal{P}_X, \mathcal{P}_Z} = -\mathbb{E}_{\mathbf{W}_X \sim \mathbb{P}_{\mathcal{P}_X}(\cdot; \mathbf{K}_X)} [\log \mathbb{P}_{\mathcal{P}_Z}(\mathbf{W}_Z = \mathbf{W}_X; \mathbf{K}_Z)]$$

defining a coupling criterion to be optimized with respect to embedding coordinates \mathbf{Z} . We now go through each couple $(\mathcal{P}_X, \mathcal{P}_Z)$ such that $\text{supp}(\mathbb{P}_{\mathcal{P}_X}) \subset \text{supp}(\mathbb{P}_{\mathcal{P}_Z})$ for the cross-entropy to be defined.

SNE. When $\mathcal{P}_X = \mathcal{P}_Z = D$, the probability of the limit posterior graphs factorizes over the nodes and the cross-entropy between limit posteriors takes the form of the objective of SNE [17](#), where for $i \in [n]$, $\mathbf{P}_i^D = \mathbf{K}_{X,i}/K_{X,i+}$ and $\mathbf{Q}_i^D = \mathbf{K}_{Z,i}/K_{Z,i+}$,

$$\mathcal{H}_{D,D} = -\sum_{i \neq j} P_{ij}^D \log Q_{ij}^D.$$

Symmetric-SNE. Choosing $\mathcal{P}_X = D$ and $\mathcal{P}_Z = E$, we define for $(i, j) \in [n]^2$, $\mathbf{Q}_{ij}^E = K_{Z,ij}/K_{Z,++}$ and $\bar{P}_{ij}^D = P_{ij}^D + P_{ji}^D$. The symmetry of \mathbf{Q}^E yields:

$$\mathcal{H}_{D,E} = -\sum_{i \neq j} P_{ij}^D \log Q_{ij}^E = -\sum_{i < j} \bar{P}_{ij}^D \log Q_{ij}^E$$

and the symmetrized objective of t-SNE [38](#) is recovered.

LargeVis. Now choosing $\mathcal{P}_X = D$ and $\mathcal{P}_Z = B$, one can also notice that $\mathbf{Q}^B = (K_{Z,ij}/(1 + K_{Z,ij}))_{(i,j) \in [n]^2}$ is symmetric. With this at hand the limit cross-entropy reads

$$\mathcal{H}_{D,B} = -\sum_{i \neq j} P_{ij}^D \log Q_{ij}^B + (1 - P_{ij}^D) \log(1 - Q_{ij}^B) = -\sum_{i < j} \bar{P}_{ij}^D \log Q_{ij}^B + (2 - \bar{P}_{ij}^D) \log(1 - Q_{ij}^B)$$

which is the objective of LargeVis [36](#).

Table 1: Prior distributions for \mathbf{W}_X and \mathbf{W}_Z associated with the pairwise similarity coupling DR algorithms. Grey-colored boxes are such that the cross-entropy is undefined.

$\mathcal{P}_X \backslash \mathcal{P}_Z$	B	D	E
\tilde{B}	UMAP		
D	LARGEVIS	SNE	T-SNE

UMAP. Let us take $\mathcal{P}_X = \mathcal{P}_Z = B$ and consider the symmetric thresholded graph $\tilde{\mathbf{W}}_X = \mathbb{1}_{\mathbf{W}_X + \mathbf{W}_X^\top \geq 1}$. By independence of the edges, $\tilde{W}_{X,ij} \sim \mathcal{B}(\tilde{P}_{ij}^B)$ where $\tilde{P}_{ij}^B = P_{ij}^B + P_{ji}^B - P_{ij}^B P_{ji}^B$ and $\mathbf{P}^B = (K_{X,ij}/(1 + K_{X,ij}))_{(i,j) \in [n]^2}$. Coupling $\tilde{\mathbf{W}}_X$ and \mathbf{W}_Z gives:

$$\mathcal{H}_{\tilde{B},B} = -2 \sum_{i < j} \tilde{P}_{ij}^B \log Q_{ij}^B + (1 - \tilde{P}_{ij}^B) \log (1 - Q_{ij}^B)$$

which is the loss function considered in UMAP [30], the construction of $\tilde{\mathbf{W}}_X$ being borrowed from section 3.1 of the paper.

Remark 4 One can also consider $\mathcal{H}_{E,E}$ but as detailed in [38], this criterion fails at positioning outliers and is therefore not considered. Interestingly, any other feasible combination of the presented priors relates to an existing method.

3.2 Interpretations

As we have seen in section 3.1, SNE-like methods can all be derived from the graph coupling framework. What characterizes each of them is the choice of priors considered for the latent structuring graphs. To the best of our knowledge, the presented framework is the first that manages to unify all these DR algorithms. Such a framework opens many perspectives for improving upon current practices as we discuss in section 4 and section 5. We now focus on a few insights that our work provides about the empirical performances of these methods.

Repulsion & Attraction. Decomposing $\mathcal{H}_{\mathcal{P}_X, \mathcal{P}_Z}$ with Bayes' rule and simplifying constant terms one has the following optimization problem:

$$\min_{\mathbf{Z} \in \mathbb{R}^{n \times q}} - \sum_{(i,j) \in [n]^2} P_{ij}^{P_X} \log k_z(\mathbf{Z}_i - \mathbf{Z}_j) + \log \mathbb{P}(\mathbf{Z}). \quad (5)$$

The first and second terms in eq. (5) respectively summarize the attractive and repulsive forces of the objective. Recall from proposition 1 that \mathbf{P}^{P_X} is the posterior expectation of \mathbf{W}_X . Hence in SNE-like methods, the attractive forces resume to a pairwise MRF log likelihood with respect to a graph posterior expectation given \mathbf{X} . For instance if k_z is the Gaussian kernel, this attractive term reads $\text{tr}(\mathbf{Z}^\top \mathbf{L}^* \mathbf{Z})$ where $\mathbf{L}^* = \mathbb{E}_{\mathbf{W} \sim \mathbb{P}_{P_X}(\cdot; \mathbf{K}_X)}[\mathbf{L}(\mathbf{W})]$, boiling down to the objective of Laplacian eigenmaps [5]. Therefore, for Gaussian MRFs, the attractive forces resume to an unconstrained Laplacian eigenmaps objective. Such link, already noted in [7], is easily unveiled in our framework. Moreover, one can notice that only this attractive term depends on \mathbf{X} as the repulsion is given by the marginal term in (5). The latter reads $\mathbb{P}(\mathbf{Z}) = \sum_{\mathbf{W} \in \mathcal{S}_{\mathbf{W}}} \mathbb{P}(\mathbf{Z}, \mathbf{W})$ with $\mathbb{P}(\mathbf{Z}, \mathbf{W}) \propto f_k(\mathbf{Z}, \mathbf{W}) \Omega_{P_Z}(\mathbf{W})$. Such penalty notably prevents a trivial solution, as $\mathbf{0}$, like any constant vector, is a mode of $f_k(\cdot, \mathbf{W})$ for all \mathbf{W} . Also note that the prior for \mathbf{W}_X only conditions attraction while the prior for \mathbf{W}_Z only affects repulsion. In the present work we focus solely on deciphering the probabilistic model that accounts for neighbor embedding loss functions and refer to [6] for a quantitative study of attraction and repulsion in these methods.

Global Structure Preservation. To gain intuition, consider that \mathbf{W}_X is observed. As we showed in section 2.2, when one relies on shift invariant kernels, the positions of the CC means are taken from a diffuse distribution. Since the above methods are all derived

from the limit posteriors at $\varepsilon \rightarrow 0$, \mathbf{X}_M and \mathbf{Z}_M have no influence on the coupling objective. Hence if two nodes belong to different CCs, their low dimensional pairwise distance will likely not be faithful. We can expect this phenomenon to persist when the expectation on \mathbf{W}_X is considered, especially when clusters are well distinguishable in \mathbf{X} . This observation is central to understand the large scale deficiency of these methods. Note that this happens at the benefit of the local structure which is faithfully represented in low dimension, as discussed in section 1. In the following section we propose to mitigate the global structure deficiency with non-degenerate MRF models.

4 Towards Capturing Large-Scale Dependencies

In this section, we investigate the ability of graph coupling to faithfully represent global structure in low dimension. To gain intuition on the case where the distribution induced by the graph is not degenerate, we consider a proper Gaussian graph coupling model and show its equivalence with PCA. We then provide a new initialization procedure to alleviate the large scale deficiency of graph coupling when degenerate MRFs are used.

4.1 PCA as Graph Coupling

As we argue that the inability of SNE-like methods to reproduce the coarse-grain dependencies of the input in the latent space is due to the degeneracy of the conditional (3), a natural solution would be to consider graphical models that are well defined and integrable on the entire definition spaces of \mathbf{X} and \mathbf{Z} . For simplicity, we consider the Gaussian model and leave the extension to other kernels for future works. Note that in this case integrability translates into the precision matrix being full-rank. As we see with the following, the natural extension of our framework to such models leads to a well-established PCA algorithm. In the following, for a continuous variable Θ_Z , $\mathbb{P}(\Theta_Z = \cdot)$ denotes its density.

Theorem 2 *Let $\nu \geq n$, $\Theta_X \sim \mathcal{W}(\nu, \mathbf{I}_n)$ and $\Theta_Z \sim \mathcal{W}(\nu + p - q, \mathbf{I}_n)$. Assume that Θ_X and Θ_Z structure the rows of respectively \mathbf{X} and \mathbf{Z} such that:*

$$\text{vec}(\mathbf{X})|\Theta_X \sim \mathcal{N}(\mathbf{0}, \Theta_X^{-1} \otimes \mathbf{I}_p), \quad (6)$$

$$\text{vec}(\mathbf{Z})|\Theta_Z \sim \mathcal{N}(\mathbf{0}, \Theta_Z^{-1} \otimes \mathbf{I}_q). \quad (7)$$

Then the solution of the precision coupling problem:

$$\min_{\mathbf{Z} \in \mathbb{R}^{n \times q}} -\mathbb{E}_{\Theta_X|\mathbf{X}} [\log \mathbb{P}(\Theta_Z = \Theta_X|\mathbf{Z})]$$

is a PCA embedding of \mathbf{X} with q components.

We now highlight the parallels with the previous construction done for neighbor embedding methods. First note that the multivariate Gaussian with full-rank precision is inherently a pairwise MRF [35]. When choosing the Gaussian kernel for neighbor embedding methods, we saw that the graph Laplacian \mathbf{L}_X of \mathbf{W}_X was playing the role of the among-row precision matrix, as we had $\mathbf{X}|\mathbf{W}_X \sim \mathcal{N}(\mathbf{0}, \mathbf{L}_X^{-1} \otimes \mathbf{I}_p)$ (equation 2). Recall that the later always has a null-space which is spanned by the CC indicator vectors of \mathbf{W} (section 2.1). Here, the key difference is that we impose a full-rank constraint on the precision Θ . Concerning the priors, we choose the ones that are conjugate to the conditionals (6) and (7), as previously done when constructing the prior for neighbor embedding methods (definition 2). Hence in the full-rank setting, the prior simply amounts to a Wishart distribution denoted by \mathcal{W} .

The above theorem further highlights the flexibility and generality of the graph coupling framework. Unlike usual constructions of PCA or probabilistic PCA [37], in the above the linear relation between \mathbf{X} and \mathbf{Z} is recovered by solving the graph coupling problem and not explicitly stated beforehand. To the best of our knowledge, it is the first time such a link is uncovered between PCA and SNE-like methods. In contrast with the latter, PCA is well-known for its ability to preserve global structure while being significantly less efficient at identifying clusters [2]. Therefore, as suspected in section 3.2, the degeneracy of the conditional distribution given the graph is key to determine the distance preservation properties of the embeddings. We propose in section 4.2 to combine both graph coupling approaches to strike a balance between global and local structure preservation.

4.2 Hierarchical Graph Coupling

The goal of this section is to show that global structure in SNE-like embeddings can be improved by structuring the CCs' positions. We consider the following hierarchical model for \mathbf{X} , where $\mathcal{P}_X \in \{B, D, E\}$, k_x satisfies the assumptions of theorem 1 and $\nu_X \geq n$:

$$\begin{aligned} \mathbf{W}_X &\sim \mathbb{P}_{\mathcal{P}_X, k_x}^\varepsilon(\cdot; \mathbf{1}, 1), \quad \Theta_X | \mathbf{W}_X \sim \mathcal{W}(\nu_X, \mathbf{I}_R) \\ \mathbf{X}_C | \mathbf{W}_X &\sim \mathbb{P}_{k_x}(\cdot | \mathbf{W}_X), \quad \text{vec}(\mathbf{X}_M) | \Theta_X \sim \mathcal{N}\left(\mathbf{0}, (\varepsilon \mathbf{U}_{[R]} \Theta_X \mathbf{U}_{[R]}^\top)^{-1} \otimes \mathbf{I}_p\right) \end{aligned}$$

where $\mathbf{U}_{[R]}$ are the eigenvectors associated to the Laplacian null-space of $\overline{\mathbf{W}}_X$. Given a graph \mathbf{W}_X , the idea is to structure the CCs' relative positions with a full-rank Gaussian model. The same model is considered for \mathbf{W}_Z , Θ_Z and \mathbf{Z} , choosing $\nu_Z = \nu_X + p - q$ for the Wishart prior to satisfy the assumption of theorem 2. With this in place, we aim at providing a complete coupling objective, matching the pairs (\mathbf{W}_X, Θ_X) and (\mathbf{W}_Z, Θ_Z) . The joint negative cross-entropy can be decomposed as follows:

$$\begin{aligned} \mathbb{E}_{(\mathbf{w}_X, \Theta_X) | \mathbf{X}} [\log \mathbb{P}((\mathbf{W}_Z, \Theta_Z) = (\mathbf{W}_X, \Theta_X) | \mathbf{Z})] \\ = \mathbb{E}_{\mathbf{w}_X | \mathbf{X}} [\log \mathbb{P}(\mathbf{W}_Z = \mathbf{W}_X | \mathbf{Z})] + \end{aligned} \quad (8)$$

$$\mathbb{E}_{(\mathbf{w}_X, \Theta_X) | \mathbf{X}} [\log \mathbb{P}(\Theta_Z = \Theta_X | \mathbf{W}_Z = \mathbf{W}_X, \mathbf{Z})] \quad (9)$$

where (8) is the usual coupling criterion of \mathbf{W}_X and \mathbf{W}_Z capturing intra-CC variability while (9) is a penalty resulting from the Gaussian structure on \mathcal{S}_M . Constructed as such, the above objective allows a trade-of between local and global structure preservation. Following current trends in DR [21], we propose to take care of the global structure first *i.e.* focusing on (9) before (8). The difficulty of dealing with (9) lies in the hierarchical construction of the graph and the Gaussian precision (see fig. 3). We state the following result.

Corollary 1 Let $\mathbf{W}_X \in \mathcal{S}_W$, $\mathbf{L} = L(\overline{\mathbf{W}}_X)$ and $\mathcal{S}_M^q = (\ker \mathbf{L}) \otimes \mathbb{R}^q$, then for all $\varepsilon > 0$, given the above hierarchical model, the solution of the problem:

$$\min_{\mathbf{Z} \in \mathcal{S}_M^q} -\mathbb{E}_{\Theta_X | \mathbf{X}} [\log \mathbb{P}(\Theta_Z = \Theta_X | \mathbf{W}_Z = \mathbf{W}_X, \mathbf{Z})]$$

is a PCA embedding of $\mathbf{U}_{[R]} \mathbf{U}_{[R]}^\top \mathbf{X}$ where $\mathbf{U}_{[R]}$ are the CCs' membership vectors of $\overline{\mathbf{W}}_X$.

Remark 5 Note that while (8) approximates the objective of SNE-like methods when $\varepsilon \rightarrow 0$, the minimizer of (9) given by corollary 1 is stable for all ε .

From this observation, we propose a simple heuristic to minimize (9) that consists in computing a PCA embedding of $\mathbb{E}_{\mathbb{P}_{\mathcal{P}_X}(\cdot; \mathbf{K}_X)} [\mathbf{U}_{[R]} \mathbf{U}_{[R]}^\top] \mathbf{X}$. The distribution of the connected components of the posterior of \mathbf{W}_X being intractable, we resort to a Monte-Carlo estimation of the above expectation. The latter procedure called *ccPCA* aims at recovering the inter-CC structure that is filtered by SNE-like methods. *ccPCA* may then be used as initialization for optimizing (8) which is done by running the DR method corresponding to the graph priors at hand (section 3.1). This second step essentially consists in refining the intra-CC structure.

4.3 Experiments with *ccPCA*

Figure 1 shows that a t-SNE embedding of a balanced MNIST dataset of 10000 samples [14] with isotropic Gaussian initialization performs poorly in conserving the relative positions of clusters. As each digit cluster contains approximately 1000 points, with a perplexity of 30, sampling an edge across digit clusters in the graph posterior $\mathbb{P}_{\mathcal{P}_X}(\cdot; \mathbf{K}_X)$ is very unlikely.

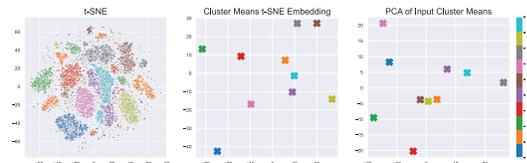


Figure 1: Left: MNIST t-SNE (perp : 30) embeddings initialized with i.i.d $\mathcal{N}(0, 1)$ coordinates. Middle: using these t-SNE embeddings, mean coordinates for each digit are represented. Right: we compute a matrix of mean input coordinates for each of the 10 digits and embed it using PCA. For t-SNE embeddings, the positions of clusters vary across different runs and don't visually match the PCA embeddings of input mean vectors (right plot).

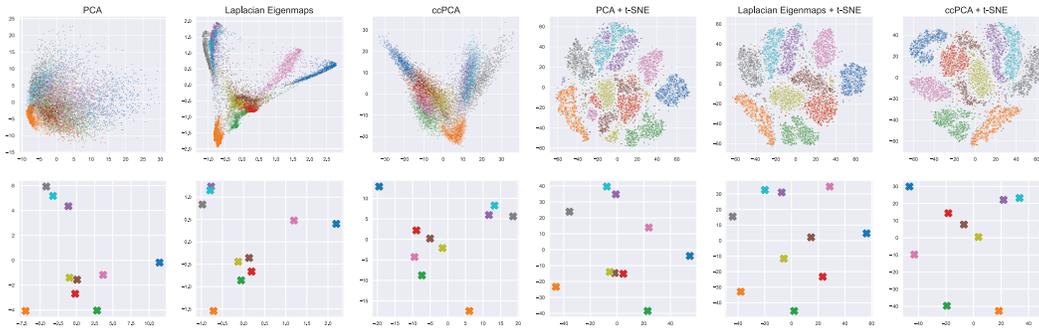


Figure 2: Top: MNIST embeddings produced by PCA, Laplacian eigenmaps, *ccPCA* and finally t-SNE launched after the previous three embeddings to improve the fine-grain structure. Bottom: mean coordinates for each digit using the embeddings of the first row. The color legend is the same as in fig. 1. t-SNE was trained during 1000 iterations using default parameters with the openTSNE implementation [34].

Recall that the perplexity value [38] corresponds to the approximate number of effective neighbors of each point. Hence images of different digits are with very high probability in different CCs of the graph posterior and their CC-wise means are not coupled as discussed in section 3.2. To remedy this in practice, PCA or Laplacian eigenmaps are usually used as initialization [21].

These strategies are tested (fig. 2) together with *ccPCA*. This shows that *ccPCA* manages to retrieve the digits that mostly support the large-scale variability as measured by the peripheral positioning of digits 0 (blue), 2 (green), 6 (pink) and 7 (grey) given by the right side of fig. 1. Other perplexity values for *ccPCA* are explored in appendix B.2 while the experimental setup is detailed in appendix B.1. In appendix B.3, we perform quantitative evaluations of *ccPCA* for both t-SNE and UMAP on various datasets using K-ary neighborhood criteria. We find that using *ccPCA* as initialization is in general more reliable than PCA and Laplacian eigenmaps for preserving global structure using both t-SNE and UMAP.

Compared to PCA, *ccPCA* manages to aggregate points into clusters, thus filtering the intra-cluster variability and focusing solely on the inter-cluster structure. Compared to Laplacian eigenmaps which performs well at identifying clusters but suffers from the same deficiency as t-SNE for positioning them, *ccPCA* retains more of the coarse-grain structure. These observations support our unifying probabilistic framework and the theoretical results about the MRF degeneracy which are the leading contributions of this article. The *ccPCA* initialization appears as a first stepping stone towards more grounded DR methods based on the probabilistic model presented in this article.

5 Conclusion and Perspectives

In this work we shed a new light on most popular DR methods by showing that they can be unified within a common probabilistic model in the form of latent Markov Random Fields Graphs coupled by a cross entropy. The definition of such a model constitutes a major step towards the understanding of common dimension reduction methods, in particular their structure preservation properties as discussed in this article.

Our work offers many perspectives, among which the possibility to enrich the probabilistic model with more suited graph priors. Currently considered priors are simply the ones that are conjugate to the MRFs thus they are mostly designed to yield a tractable coupling objective. However they may not be optimal and could be modified to capture targeted features, *e.g.* communities, in the input data, and give adapted representations in the latent space. The graph coupling approach could also be extended to more general latent structures governing the joint distribution of observations. Finally, the probabilistic model could be leveraged to tackle hyper-parameter calibration, especially kernel bandwidths that have a great influence on the quality of the representations and are currently tuned using heuristics with unclear motivations.

Acknowledgments. The authors would like to thank the anonymous reviewers whose comments and questions helped improve the clarity of this manuscript, as well as Aurélien Garivier, Antoine Barrier and Flosi Poshi for helpful discussions. This work was supported by the Agence Nationale de la Recherche ANR-18-CE45-0023 SingleStatOmics.

References

- [1] Friedrich Anders, Cristina Chiappini, Basilio Xavier Santiago, Gal Matijević, Anna B Queiroz, Matthias Steinmetz, and Guillaume Guiglion. Dissecting stellar chemical abundance space with t-sne. *Astronomy & Astrophysics*, 619:A125, 2018.
- [2] Farzana Anowar, Samira Sadaoui, and Bassant Selim. Conceptual and empirical comparison of dimensionality reduction algorithms (PCA, KPCA, LDA, MDS, SVD, LLE, ISOMAP, LE, ICA, t-SNE). *Computer Science Review*, 40:100378, 2021.
- [3] Sanjeev Arora, Wei Hu, and Pravesh K Kothari. An analysis of the t-sne algorithm for data visualization. In *Conference On Learning Theory*, pages 1455–1462. PMLR, 2018.
- [4] Mukund Balasubramanian, Eric L Schwartz, Joshua B Tenenbaum, Vin de Silva, and John C Langford. The isomap algorithm and topological stability. *Science*, 295(5552):7–7, 2002.
- [5] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- [6] Jan Niklas Böhm, Philipp Berens, and Dmitry Kobak. A unifying perspective on neighbor embeddings along the attraction-repulsion spectrum. *arXiv preprint arXiv:2007.08902*, 2020.
- [7] Miguel A Carreira-Perpinán. The elastic embedding algorithm for dimensionality reduction. In *ICML*, volume 10, pages 167–174. Citeseer, 2010.
- [8] David M Chan, Roshan Rao, Forrest Huang, and John F Canny. t-sne-cuda: Gpu-accelerated t-sne and its applications to modern data. In *2018 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, pages 330–338. IEEE, 2018.
- [9] F. R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.
- [10] Peter Clifford. Markov random fields in statistics. *Disorder in physical systems: A volume in honour of John M. Hammersley*, pages 19–32, 1990.
- [11] Andy Coenen and Adam Pearce. Understanding umap. *Google PAIR*, 2019.
- [12] Ronald R Coifman and Stéphane Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.
- [13] Gabor Csardi, Tamas Nepusz, et al. The igraph software package for complex network research. *InterJournal, complex systems*, 1695(5):1–9, 2006.
- [14] Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- [15] David L Donoho. High-dimensional data analysis: The curses and blessings of dimensionality. *AMS math challenges lecture*, 1(2000):32, 2000.
- [16] Jihun Ham, Daniel D Lee, Sebastian Mika, and Bernhard Schölkopf. A kernel view of the dimensionality reduction of manifolds. In *Proceedings of the twenty-first international conference on Machine learning*, page 47, 2004.
- [17] Geoffrey Hinton and Sam T Roweis. Stochastic neighbor embedding. In *NIPS*, volume 15, pages 833–840. Citeseer, 2002.

- [18] Geoffrey E Hinton and Sam T. Roweis. Stochastic neighbor embedding. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 857–864. MIT Press, 2003.
- [19] Dmitry Kobak and Philipp Berens. The art of using t-sne for single-cell transcriptomics. *Nature communications*, 10(1):1–14, 2019.
- [20] Dmitry Kobak, George Linderman, Stefan Steinerberger, Yuval Kluger, and Philipp Berens. Heavy-tailed kernels reveal a finer cluster structure in t-sne visualisations. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 124–139. Springer, 2019.
- [21] Dmitry Kobak and George C Linderman. Initialization is critical for preserving global data structure in both t-sne and umap. *Nature biotechnology*, 39(2):156–157, 2021.
- [22] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 and cifar-100 datasets. URL: <https://www.cs.toronto.edu/kriz/cifar.html>, 6(1):1, 2009.
- [23] Joseph B Kruskal. *Multidimensional scaling*. Number 11. Sage, 1978.
- [24] Nadia S Kurd, Zhaoren He, Tiani L Louis, J Justin Milner, Kyla D Omilusik, Wenhao Jin, Matthew S Tsai, Christella E Widjaja, Jad N Kanbar, Jocelyn G Olvera, et al. Early precursors and molecular determinants of tissue-resident memory cd8+ t lymphocytes revealed by single-cell rna sequencing. *Science immunology*, 5(47):eaaz6894, 2020.
- [25] John A Lee, Diego H Peluffo-Ordóñez, and Michel Verleysen. Multi-scale similarities in stochastic neighbour embedding: Reducing dimensionality while preserving both local and global structure. *Neurocomputing*, 169:246–261, 2015.
- [26] Tianxi Li, Cheng Qian, Elizaveta Levina, and Ji Zhu. High-dimensional gaussian graphical models on network-linked data. *J. Mach. Learn. Res.*, 21:74–1, 2020.
- [27] Wentian Li, Jane E Cerise, Yaning Yang, and Henry Han. Application of t-sne to human genetic data. *Journal of bioinformatics and computational biology*, 15(04):1750017, 2017.
- [28] George C Linderman, Manas Rachh, Jeremy G Hoskins, Stefan Steinerberger, and Yuval Kluger. Fast interpolation-based t-sne for improved visualization of single-cell rna-seq data. *Nature methods*, 16(3):243–245, 2019.
- [29] George C Linderman and Stefan Steinerberger. Clustering with t-sne, provably. *SIAM Journal on Mathematics of Data Science*, 1(2):313–332, 2019.
- [30] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [31] Karl Pearson. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science*, 2(11):559–572, 1901.
- [32] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- [33] Nicola Pezzotti, Julian Thijssen, Alexander Mordvintsev, Thomas Höllt, Baldur Van Lew, Boudewijn PF Lelieveldt, Elmar Eisemann, and Anna Vilanova. Gpgpu linear complexity t-sne optimization. *IEEE transactions on visualization and computer graphics*, 26(1):1172–1181, 2019.
- [34] Pavlin G Poličar, Martin Stražar, and Blaž Zupan. opentsne: a modular python library for t-sne dimensionality reduction and embedding. *BioRxiv*, page 731877, 2019.
- [35] Havard Rue and Leonhard Held. *Gaussian Markov random fields: theory and applications*. CRC press, 2005.

- [36] Jian Tang, Jingzhou Liu, Ming Zhang, and Qiaozhu Mei. Visualizing large-scale and high-dimensional data. In *Proceedings of the 25th international conference on world wide web*, pages 287–297, 2016.
- [37] Michael E Tipping and Christopher M Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- [38] L.J.P. van der Maaten and G.E. Hinton. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9(nov):2579–2605, 2008. Pagination: 27.
- [39] Martin J Wainwright and Michael Irwin Jordan. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.
- [40] Yingfan Wang, Haiyang Huang, Cynthia Rudin, and Yaron Shaposhnik. Understanding how dimension reduction tools work: an empirical approach to deciphering t-sne, umap, trimap, and pacmap for data visualization. *J Mach. Learn. Res.*, 22:1–73, 2021.
- [41] Martin Wattenberg, Fernanda Viégas, and Ian Johnson. How to use t-sne effectively. *Distill*, 1(10):e2, 2016.
- [42] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [No]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes] In appendix [A](#).
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] In appendix [B](#).
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] In appendix [B](#).
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In appendix [B](#).
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [Yes]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]