

Some steps towards a general principle for dimensionality reduction mappings

Barbara Hammer ^{*}; Kerstin Bunte, Michael Biehl [†]

November 2, 2010

Abstract

In the past years, many dimensionality reduction methods have been established which allow to visualize high dimensional data sets. Recently, also formal evaluation schemes have been proposed for data visualization, which allow a quantitative evaluation along general principles. Most techniques provide a mapping of a priorly given finite set of points only, requiring additional steps for out-of-sample extensions. We propose a general view on dimensionality reduction based on the concept of cost functions, and, based on this general principle, extend dimensionality reduction to explicit mappings of the data manifold. This offers the possibility of simple out-of-sample extensions. Further, it opens a way towards a theory of data visualization taking the perspective of its generalization ability to new data points. We demonstrate the approach based in a simple example.

1 Introduction

The amount of electronic data available today doubles approximately every 20 months. At the same time, its complexity and dimensionality increases dramatically due to improved sensor technology, dedicated data formats, and rapidly increasing capabilities to digitally capture different data modalities. As a consequence, data can no longer be inspected manually, rather, automated methods which help humans to quickly scan through massive data volumes are needed. Data visualization relies on the astonishing cognitive capabilities of humans for structure detection in visual images. In this context, the available information and structural characteristics or specifics can be captured almost instantly by humans despite the given number of data points which are represented in the visualization. As a consequence, data visualization and dimensionality reduction play a key role in modern data mining techniques.

^{*}University of Bielefeld, CITEC Center of Excellence, Bielefeld - Germany Email: bhammer@techfak.uni-bielefeld.de

[†]University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science, Nijenborgh 9, Groningen - The Netherlands, Email: k.bunte@rug.nl

A plethora of methods for linear and nonlinear dimensionality reduction has been proposed in the past years, see e.g. [10, 20, 18, 8, 21] for recent overviews. In general, the task is to substitute data points in a high dimensional data manifold by much lower dimensions (ideally two dimensions to obtain a visualization), such that as much information as possible is preserved. Since this problem formulation is ill-posed, a variety of methods can be derived by imposing additional constraints on the visualization task. Spectral dimensionality reduction techniques such as LLE [13], Isomap [16], or Laplacian eigenmaps [2] rely on the spectrum of the neighborhood graph of the data and preserve important properties of this graph. In general, they allow a unique algebraic solution of the corresponding mathematical objective which formalizes the visualization task. Thereby, many methods rely on very simple affinity functions such as Gaussians such that their results are flawed when it comes to effects such as boundaries or different separated manifolds. Using more complex affinities such as present in Isomap [16] or maximum variance unfolding [22] can partially avoid this problem, but at the prize of higher computational costs. Nonlinear methods which do not rely on the spectrum often have the drawback that local optima can easily occur. However, their results can be more appropriate as demonstrated e.g. for SNE [7], t-SNE [18], or elastic embedding [6].

All of these methods, however, map the given data points only and their extension towards novel data points requires additional effort. Essentially, two different ways for out of sample extensions can be found in the literature: either an interpolation takes place, e.g. by fitting a neural network to the data which interpolates the projection mapping. This has the drawback that the mapping is not optimized for the projection task, rather, it interpolates the given (probably faulty) coordinates. Alternatively, novel points can be directly mapped to a position in the projection space which minimizes the underlying cost function of the visualization method, where the coordinates of the priorly given data and their projections are kept fixed. In some cases, an explicit algebraic expression is possible, for complex cost functions, numerical optimization is necessary. Usually, however, the novel coordinates depend on all given data by means of the cost function, which often yields to quadratic effort corresponding to the pairwise affinities of data points captured in the cost function.

In this contribution we take a different point of view and propose a general principle how dimensionality reduction mappings which are optimized for the visualization task can be obtained based on the dimensionality reduction principles as proposed in the literature. For this purpose, a specific form and complexity of the dimensionality reduction mapping is fixed, such as a function stemming from a class which allows universal approximation, e.g. locally linear functions, or a particularly simple function to allow easy interpretability such as a global linear function. Instead of the coordinates of the projected data points, the function parameters are optimized in a second step. A similar mechanism has been proposed in specific settings in the contribution [13], LLE is extended towards a locally linear embedding function, leading to locally linear coordination, in the approach [19] t-SNE is extended towards an embedding given by an encoder networks. We argue that this principle can be generalized towards a

general framework which allows to adapt embedding functions of different complexity according to a given objective as mirrored by the various dimensionality reduction technologies. We exemplarily demonstrate this procedure for a global linear mapping and the visualization cost term of t-SNE.

The fact that an explicit mapping is obtained instead of coordinates of single points has several benefits: out-of-sample extensions are immediate and reduce to (efficient) function evaluations, whereby the form and complexity of the function can be defined a priori. Approximate inverse mappings can be constructed e.g. by a local linear approximation of the projection and the corresponding pseudoinverse. This way, paths in the projection space can be traced back to paths in the data manifold, shedding some light on the structure of the projection. Since the dimensionality reduction mapping is usually described by a small number of parameters, few data points are sufficient to reliably determine these parameters, i.e. training can be done using a small subset of the data only instead of the full data set. This can dramatically reduce the complexity of the computation since the cost functions often scale at least quadratically with the number of training data. This generalization ability of dimensionality reduction mapping can formally be put into the framework of statistical learning theory. Assuming that a loss function of the dimensionality reduction is fixed, the empirical error of this loss function on a small data set is often already representative for the full error assumed reasonable mappings and loss functions are considered. We will discuss this fact in more detail within this contribution. Further, we will also discuss, in how far this generalization ability of dimensionality reduction mappings can be used to show a formal concept of learnability of dimensionality reduction e.g. based on the reconstruction error of the map.

2 Dimensionality reduction as cost optimization

First, we shortly review some of the most popular dimensionality reduction methods as proposed in the literature. We assume that high dimensional points $X : \{\vec{x}^i \in \mathbb{R}^D\}_{i=1}^n$ are given which should be projected to points $Y : \{\vec{y}^i \in \mathbb{R}^d\}_{i=1}^n$ with $d < D$, ideally $d = 2$ for visualization in the two-dimensional Euclidean plane. Corresponding distances are denoted as $d_{\mathcal{X}}(\vec{x}^i, \vec{x}^j)$ for the original manifold, and $d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)$ for the projection space. Usually, $d_{\mathcal{E}}$ is chosen as the Euclidean distance, while $d_{\mathcal{X}}(\vec{x}^i, \vec{x}^j)$ can be picked arbitrarily (e.g according to Euclidean or geodesic distances in the high dimensional space.)

Multidimensional scaling and extensions

Multidimensional scaling (MDS) [17] or metric MDS, to be more precise, constitutes probably one of the most popular and oldest dimensionality reduction methods. Its goal is to find projections such that the pairwise relations of data are preserved as much as possible as measured in the least squares sense, i.e.

$$E_{\text{MDS}} = \sum_{ij} ((\vec{x}^i)^\top \vec{x}^j - (\vec{y}^i)^\top \vec{y}^j)^2$$

is minimized where, for original MDS, the pairwise relation of data is measured in terms of dot products in the original or projection space, respectively. This formulation has the benefit that an analytical solution is possible in terms of the eigenvectors of the Gram matrix. This objective has later been generalized to explicitly preserve distances, i.e. the cost function becomes

$$E_{\text{MDS}} = \frac{1}{c} \sum_{ij} w_{ij} (d_{\mathcal{X}}(\bar{x}^i, \bar{x}^j) - d_{\mathcal{E}}(\bar{y}^i, \bar{y}^j))^2$$

with Euclidean distances, where the weights w_{ij} can be chosen appropriately, e.g. $w_{ij} = 1$, and c is a normalizing constant [10]. For the popular Sammon mapping, the weights are picked as $w_{ij} = 1/d_{\mathcal{X}}(\bar{x}^i, \bar{x}^j)$, this way putting most emphasis on the preservation of small distances, and c denotes the sum over these distances. In this case, optimization of the cost function usually takes place by means of a gradient descent.

Isomap

Isomap [16] is based on the observation that the Euclidean distance is often not appropriate to describe pairwise relations of data in the original space, rather, the distance should be measured along the data manifold. Therefore, Isomap is based on an approximation of the manifold distance by geodesic distances, i.e. shortest paths lengths in the graph which results if every data point is connected to its nearest neighbors (using either k -neighborhoods or ϵ -balls to define the local neighborhood).

Locally linear embedding

Locally linear embedding (LLE) [13] first expresses local topologies by reconstructing a data point by linear combinations of its local neighborhood (denoted by $i \rightarrow j$) in the original space under the constraint that the coefficients sum to one such that translation and rotation invariance is enforced: minimize $\sum_i (\bar{x}^i - \sum_{i \rightarrow j} w_{ij} \bar{x}^j)^2$ with $\sum w_{ij} = 1$. Afterwards, projections are determined such that the local linear relationships are preserved as much as possible in a least squares sense where, again, a normalization of the coefficients leads to a unique optimum of the optimization problem: minimize $\sum_i (\bar{y}^i - \sum_{i \rightarrow j} w_{ij} \bar{y}^j)^2$ such that $\sum \bar{y}^i = 0$ and $\mathbf{Y}^t \mathbf{Y} = \mathbf{n}$, the latter referring to the corresponding matrices.

Laplacian Eigenmaps

Laplacian eigenmaps [2], like LLE and Isomap, start with a local neighborhood graph given by the k nearest neighbors or ϵ -neighborhood, respectively. The connections are weighted with values w_{ij} , e.g. using the heat kernel. Then, projection takes place by picking the eigendirections corresponding to the smallest eigenvalues larger than 0 as computed in the generalized eigenvalue problem given by the corresponding graph Laplacian and the degree matrix of the graph.

This is equivalent to minimizing the embedding objective $\sum_{i \rightarrow j} w_{ij} d_{\mathcal{E}}(\bar{y}^i, \bar{y}^j)^2$ with Euclidean distance, under the constraint $\mathbf{Y}^t D \mathbf{Y} = \mathbf{1}$ and $\mathbf{Y}^t D \bar{\mathbf{1}} = \bar{\mathbf{0}}$, where D is the degree matrix and \mathbf{Y} refers to the matrix of coefficients, to remove scaling factors and translation factors.

Maximum variance unfolding

Maximum variance unfolding (MVU) [22] also first determines a neighborhood graph by taking the k nearest neighbors or ϵ neighborhoods. Afterwards, it finds projections \bar{y}^i such that the variance of the projection is maximized, i.e. $\sum_{i,j} d_{\mathcal{E}}(\bar{y}^i, \bar{y}^j)^2$ is maximum subject to a preservation of neighbors, i.e. $d_{\mathcal{E}}(\bar{y}^i, \bar{y}^j) = d_{\mathcal{X}}(\bar{x}^i, \bar{x}^j)$ for all neighbored points \bar{x}^i and \bar{x}^j , and the normalization $\sum \bar{y}^i = 0$. This can be reformulated as a convex problem by considering the variables $(\bar{y}^i)^\top \bar{y}$ instead. Further, it is not clear that a solution exists due to the constraints, such that possibly slack variables have to be introduced.

Stochastic neighbor embedding

Stochastic neighbor embedding (SNE) [7] defines probabilities

$$p_{j|i} = \frac{\exp\left(\frac{-d_{\mathcal{X}}(\bar{x}^i, \bar{x}^j)^2}{2\sigma_i}\right)}{\sum_{k \neq i} \exp\left(\frac{-d_{\mathcal{X}}(\bar{x}^i, \bar{x}^k)^2}{2\sigma_i}\right)}$$

and

$$q_{j|i} = \frac{\exp(-d_{\mathcal{E}}(\bar{y}^i, \bar{y}^j)^2)}{\sum_{k \neq i} \exp(-d_{\mathcal{E}}(\bar{y}^i, \bar{y}^k)^2)}$$

with Euclidean distances as default. The goal is to optimize the Kullback-Leibler divergence of these two distributions, i.e. the term $E_{\text{SNE}} = -\sum_{i,j} p_{j|i} \log \frac{p_{j|i}}{q_{j|i}}$, where appropriate bandwidths σ_i are determined based on the so-called perplexity which determines the effective number of neighbors of a given data point. Typically, gradient descent is used for the optimization.

T-distributed stochastic neighbor embedding

t-distributed SNE (t-SNE) [18] modifies SNE towards a numerically easier cost function and a distribution in the embedding space which prevents the crowding problem by long tails, student-t. Its cost function is

$$E_{\text{t-SNE}} = \sum_i \sum_j p_{ij} \log \left(\frac{p_{ij}}{q_{ij}} \right)$$

where

$$p_{ij} = \frac{p_{j|i} + p_{i|j}}{2n}$$

symmetrizes the conditional probabilities, n denoting the number of data points, and

$$q_{ij} = \frac{(1 + d_{\mathcal{E}}(\bar{y}^i, \bar{y}^j)/\varsigma)^{-\frac{\varsigma+1}{2}}}{\sum_{k \neq l} (1 + d_{\mathcal{E}}(\bar{y}^k, \bar{y}^l)/\varsigma)^{-\frac{\varsigma+1}{2}}}$$

is given by student-t, and the parameter picked as $\varsigma = -1$, for example. Again, optimization takes place by means of a gradient method.

A general view

These methods obey one general principle: characteristics of the data \vec{x} are computed and projections \vec{y} are determined such that the corresponding characteristics of the projections are as close to the characteristics of \vec{x} as possible, fulfilling possibly additional constraints or objectives to achieve uniqueness. Thereby, the methods differ in the way how data characteristics are determined and how exactly the similarity of the characteristics is defined and optimized. Table 2 summarizes the properties of the optimization methods under this point of view. Naturally, the methods severely differ with respect to the way in which optimization takes place: in some cases, the characteristics can be directly computed from the data (such as distances), in others, an optimization step is required (such as local linear weights). In some cases, the optimization of the error measure can be done in closed form (such as for Laplacian eigenmaps), in other cases, numerical optimization is necessary (such as for t-SNE).

3 Dimensionality reduction mapping

All dimensionality reduction methods as introduced above give a mapping of points only: $\vec{x}^i \mapsto \vec{y}^i$. Extensions of the map to new data points \vec{x} require a new computation, often the respective coefficients which minimize the objective of dimensionality reduction are determined, keeping all known coefficients fixed. This method has the drawback that additional effort is required if new data points are dealt with. Further, it is not easily possible to formalize and investigate the generalization ability of these mappings, i.e. the question, whether the method works well for future data from the same manifold assumed it works well for the known training set.

These issues can be circumvented if a dimensionality reduction mapping

$$f : \mathcal{X} \rightarrow \mathcal{E}, \vec{x}^i \mapsto \vec{y}^i = f(\vec{x}^i)$$

from the space \mathcal{X} of original data X to the embedding space \mathcal{E} of the projected points Y is computed rather than single coefficients $\vec{x}^i \mapsto \vec{y}^i$ only.

Previous work

In the literature, a few dimensionality reduction technologies provide an explicit mapping of the data: linear methods such as PCA provide an explicit

linear function which optimizes the information loss while projecting [3]. Extensions to nonlinear functions are given by autoencoder networks, which provide a function given by a multilayer feedforward network in such a way that the reconstruction error is minimized when back projecting with another feedforward network [20]. Typically, training takes place by standard back propagation directly minimizing the reconstruction error. Manifold charting starts from locally linear embeddings given by local PCAs and glues these pieces together by minimizing the error on the overlaps [5, 14]. This way, a global embedding mapping is obtained. Topographic maps such as the self-organizing map or generative topographic mapping characterize data in terms of prototypes which are visualized in low dimensions [4, 9]. Due to the clustering, new data can directly be visualized by mapping these data to their closest prototype or its visualization, respectively.

A few dimensionality reduction mappings which give coordinates per default as introduced above have been extended to global dimensionality reduction mappings. Locally linear coordination (LLC) extends LLE in the following way [15]: it is assumed that local linear dimensionality reduction methods are available, such as local PCAs. These are glued together adding affine transformations. These additional parameters are optimized by inserting the resulting points in the LLE cost function and corresponding optimization. Parameterized t-SNE [19] extends t-SNE towards an embedding given by a multilayer neural network. The network parameters are determined using back propagation, where, instead of the mean squared error, the t-SNE cost function is taken as objective.

A general principle

Considering dimensionality reduction as optimization task as formulated in Table 2 allows to simultaneously extend all methods to dimensionality reduction mappings a general way. In a first step, the principled form and complexity of the dimensionality reduction mapping is fixed: a parameterized function

$$f_W : \mathcal{X} \rightarrow \mathcal{E}$$

is chosen with parameters W which have to be determined such that the projections are satisfactory. The form of this function can be given by a linear function, a locally linear function, a feedforward neural network, etc. Then, instead of coefficients \vec{y}^i , the images of the map $f_W(\vec{x}^i)$ are considered and instead of the single coefficients, the map parameters W are optimized. For this purpose characteristics of the data \vec{x}^i can be computed as before. Characteristics of the projected points depend on the parameterized quantities $f_W(\vec{x}^i)$ instead of the coefficients. These terms can be plugged into the corresponding error measure and the parameters W can be determined via optimization taking the same constraints into account as before (or relaxations thereof).

This principle leads to a well defined mathematical objective for the mapping parameters W for every dimensionality reduction method as summarized above, although the way in which optimization takes place is possibly different

as compared to the original method: while numerical methods such as gradient descent can still be used, it is probably no longer possible to find closed form solutions for spectral methods. However, numerical optimization can be used as a default in all cases.

We exemplarily derive formulas for a specific case: a global linear mapping combined with the t-SNE cost function. The suitability of the general principle for different dimensionality reduction cost functions and different parameterizations of the projection mapping will be the subject of future work.

Linear t-SNE Mapping

We derive the formulation in case of a linear hypothesis for the mapping of the high-dimensional data points \vec{x}^l and the t-SNE cost function. The mapping f_W becomes

$$f_W : \vec{x}^l \rightarrow \vec{y}^l = A \cdot \vec{x}^l .$$

The rectangular matrix A defines a linear mapping from $\mathbb{R}^D \rightarrow \mathbb{R}^d$. This matrix can be optimized using a stochastic gradient descent procedure using following gradient of the t-SNE cost function:

$$\begin{aligned} \frac{\partial E_{\text{t-SNE}}}{\partial A} &= \sum_i \sum_j \frac{\partial E_{\text{t-SNE}}}{\partial q_{ij}} \cdot \frac{\partial q_{ij}}{\partial d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)^2} \cdot \frac{\partial d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)^2}{\partial A} \\ &= \frac{\varsigma + 1}{2\varsigma} \sum_i \sum_j (p_{ij} - q_{ji}) \cdot \\ &\quad (1 + d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)/\varsigma)^{-1} \cdot \frac{\partial d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)^2}{\partial A} \end{aligned}$$

with Euclidean distance $d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j) = \|A\vec{x}^i - A\vec{x}^j\|$ follows:

$$\frac{\partial d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)^2}{\partial A} = 2(A\vec{x}^i - A\vec{x}^j)(\vec{x}^i - \vec{x}^j)$$

Hence

$$\begin{aligned} \frac{\partial E_{\text{t-SNE}}}{\partial A} &= \frac{\varsigma + 1}{\varsigma} \sum_i \sum_j (A\vec{x}^i - A\vec{x}^j)(\vec{x}^i - \vec{x}^j) \cdot \\ &\quad (p_{ij} - q_{ji}) \frac{1}{1 + \|A\vec{x}^i - A\vec{x}^j\|^2/\varsigma} . \end{aligned}$$

We test this procedure in comparison to simple PCA on a three dimensional benchmark problem: three Gaussian clusters are stacked together as shown in Fig. 1. Because of the comparably large variance in the direction of the z-coordinate, a PCA mapping projects the data clouds onto each other. In contrast, a linear mapping trained such that the t-SNE cost function of the projections is optimized leads to a much clearer separation of the cluster structure, because it takes into account the preservation of local structures as measured




Figure 1: Simulation Results for a globally linear map trained with PCA and an optimization of the t-SNE costs, respectively. The latter leads to a better separation due to its local nature, which can be formally evaluated referring to the intrusions and extrusions of the mapping.

by the t-SNE cost function. Fig. 1 clearly shows the superiority of the mapping obtained this way, referred to as DiReduct mapping. In addition, the projection is formally evaluated using the error measure as proposed in [11, 12]. Roughly speaking, these rely on the k -intrusions and k -extrusions in the projections, i.e. k -nearest neighbors in the projection, but not the original space, and vice versa. The quality measures refer to the quantities Q which measures the percentage of data which is not k -intrusive or k -extrusive, and B which measures the percentage of k -intrusions minus the percentage of k -extrusions of the map, i.e. it characterizes the behavior of the mapping. Obviously, DiReduct shows a superior quality in particular for small neighborhood ranges since it better preserves local structures of the data. Further, unlike PCA which displays a trend towards intrusions, it is rather neutral in the mapping character, being mildly extrusive for medium sizes of k nearest neighbors.

Alternative more complex mappings such as e.g. locally linear functions as well as an investigation which forms of the mapping and which dimensionality

reduction objectives are suited for the task are the subject of ongoing work. Interestingly, it is possible to integrate additional aspects into the frameworks such as e.g. auxiliary information, by biasing the dimensionality reduction mapping towards the given auxiliary information.

4 Generalization ability

The extension towards dimensionality reduction mappings offers the possibility to learn the mapping based on few randomly selected data points only, since the mapping parameters are already determined by only few points. Depending on the size of the data, this can severely improve the performance of the method, since it reduces the squared complexity to a constant effort. However, an assumption underlying this procedure is that the dimensionality reduction mapping generalizes from few data to new data stemming from the same underlying distribution. That means we have to ensure that the quality measure for all data is good assumed it is good for a given finite subsample used to determine the mapping parameters.

Recently, some work on how dimensionality reduction can be formally evaluated has been proposed [12, 21]. As pointed out in [12], one objective of dimensionality reduction is to preserve the available information as much as possible. In consequence, the possibility to reconstruct the points \vec{x}^i from their projections \vec{y}^i can act as valid evaluation measure. Assuming a dimensionality reduction mapping $f : \mathcal{X} \rightarrow \mathcal{E}$ is given, this results in the reconstruction error

$$E(P) := \int_{\mathcal{X}} \|\vec{x} - f^{-1}(f(\vec{x}))\|^2 P(\vec{x}) d\vec{x}$$

where P defines the probability measure according to which the data \vec{x} are distributed in \mathcal{X} and f^{-1} constitutes an approximate inverse mapping of f , an exact inverse in general not existing. Usually, the full data manifold is not available, but a finite set of samples only. Then, the empirical error can be computed

$$\hat{E}_n(\vec{x}) := \frac{1}{n} \sum_i \|\vec{x}^i - f^{-1}(f(\vec{x}^i))\|^2$$

for given data \vec{x}^i . Note that some dimensionality reduction mappings such as autoencoder networks explicitly optimize this empirical approximation of the costs $E(P)$.

Most dimensionality reduction methods map the data points only, such that neither f nor its approximate inverse are available. Therefore, evaluation measures as proposed in [12, 21] rely on k-neighborhoods in the original and the projection space to approximately capture neighborhood preservation. If a dimensionality reduction mapping is learned, f and its approximate inverse f^{-1} are available. Thus, the evaluation measure $\hat{E}_n(\vec{x})$ can be evaluated. Since the form of f is fixed prior to training, we can specify a function class \mathcal{F} with $f \in \mathcal{F}$ independently of the given training set. Assuming representative vectors \vec{x}^i are

chosen independently and identically distributed according to P the question is whether this quantity allows to limit the real error $E(P)$ we are interested in. As usual, bounds should hold simultaneously for all possible functions in \mathcal{F} to circumvent the problem that the function f is chosen according to the given training data and, thus, the empirical error $\widehat{E}_n(\vec{x})$ is usually small.

This setting can be captured in the classical framework of computational learning theory, as specified e.g. in [1]. We can adapt Theorem 8 from [1] to our setting: We consider a fixed function class

$$\mathcal{F} : \mathcal{X} \rightarrow \mathcal{E}$$

from which the dimensionality reduction mapping is taken. We assume without loss of generality, that the norm of the input data and its reconstructions under mappings $f^{-1} \circ f, f^{-1}$ denoting the approximate inverse of $f \in \mathcal{F}$, are restricted (scaling the data priorly, if necessary), such that the reconstruction error is induced by the squared error, which is a loss function with limited codomain

$$\mathcal{L} : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1], (\vec{x}^i, \vec{x}^j) \mapsto \|\vec{x}^i - \vec{x}^j\|^2$$

Then, as reported in [1] (Theorem 8), assuming i.i.d. data according to P , for any confidence $\delta \in (0, 1)$ and every $f \in \mathcal{F}$ the following holds

$$E(P) \leq \widehat{E}_n(\vec{x}) + R_n(\mathcal{L}_{\mathcal{F}}) + \sqrt{\frac{8 \ln(2/\delta)}{n}}$$

with probability at least $1 - \delta$ where

$$\mathcal{L}_{\mathcal{F}} := \{\vec{x} \mapsto \mathcal{L}(f^{-1}(f(\vec{x})), \vec{x}) \mid f \in \mathcal{F}\}$$

and R_n refers to the so-called Rademacher complexity of the function class. The Rademacher complexity constitutes a quantity which, similar to the Vapnik Chervonenkis dimension, estimates the capacity of a given function class. Assume σ_i are independent identically distributed $\{\pm 1\}$ -valued random variables. The empirical Rademacher complexity of a real valued function class \mathcal{G} is

$$\widehat{R}_n(\mathcal{G}) := \mathbf{E} \left(\sup_{f \in \mathcal{G}} \left| \frac{2}{n} \sum_i \sigma_i f(\vec{x}^i) \right| \mid \text{given } \vec{x}^1, \dots, \vec{x}^n \right)$$

where the expectation is taken over σ_i . It estimates the expected worst case correlation of functions in \mathcal{F} with random ± 1 -valued vectors. The Rademacher complexity denotes the expectation over this quantity taking into account all possible samples \vec{x} .

This result implies that the generalization ability of dimensionality reduction mappings is usually guaranteed since the Gaussian complexity of the class $\mathcal{L}_{\mathcal{F}}$ can be limited for reasonable choices of the mapping function \mathcal{F} . For linear functions for example, bounds on the Rademacher complexity can be derived in the same way as explained in [1].

5 Conclusion

In this contribution, the question how a dimensionality reduction mapping can be inferred rather than coordinates of separated points has been considered. By formulating dimensionality reduction as an optimization problem between structural characteristics, many classical dimensionality reduction techniques can simultaneously be extended towards explicit mappings which depend on a priorly chosen form of the mapping. We have demonstrated the feasibility of this approach in one simple example.

This general view opens the way towards a plethora of alternative mapping possibilities since, in principle, every cost function can be combined with every possible choice of the mapping function. Even more interesting, the framework allows us to consider the generalization ability of dimensionality reduction since an explicit cost function is available in terms of the reconstruction error based on a dimensionality reduction mapping and its approximate inverse. Interestingly, bounds as derived in the context of computational learning theory can directly be transferred to this setting.

The investigation of alternative dimensionality reduction mappings including more global cost functions such as provided by Isomap, and locally non-linear function approximations, as well as the derivative of explicit bounds on its generalization ability will be the subject of future work.

Acknowledgment

This work was supported by the "Nederlandse organisatie voor Wetenschappelijke Onderzoek (NWO)" under project code 612.066.620 and by the "German Science Foundation (DFG)" under grant number HA-2719/4-1.

References

- [1] P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: risk bounds and structural results. *J. Mach. Learn. Res.*, 3:463–482, 2003.
- [2] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15:1373–15396, 2003.
- [3] C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.
- [4] C. M. Bishop and C. K. I. Williams. Gtm: The generative topographic mapping. *Neural Computation*, 10:215–234, 1998.
- [5] M. Brand. Charting a manifold. Technical Report 15, Mitsubishi Electric Research Laboratories (MERL), 2003.
- [6] M. Á. Carreira-perpiñán. The elastic embedding algorithm for dimensionality reduction. In *27th Int. Conf. Machine Learning (ICML 2010)*, pages 167–174, 2010.

- [7] G. Hinton and S. Roweis. Stochastic neighbor embedding. In *Advances in Neural Information Processing Systems 15*, pages 833–840. MIT Press, 2003.
- [8] D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler. Visual analytics: Scope and challenges. In S. Simoff, M. H. Boehlen, and A. Mazeika, editors, *Visual Data Mining: Theory, Techniques and Tools for Visual Analytics*. Springer, 2008. Lecture Notes in Computer Science (LNCS).
- [9] T. Kohonen. *Self-organizing Maps*. Springer, 1995.
- [10] J. Lee and M. Verleysen. *Nonlinear dimensionality reduction*. Springer, 1st edition, 2007.
- [11] J. A. Lee and M. Verleysen. Rank-based quality assessment of nonlinear dimensionality reduction. In *16th European Symposium on Artificial Neural Networks (ESANN)*, pages 49–54, Bruges, Belgium, 2008.
- [12] J. A. Lee and M. Verleysen. Quality assessment of dimensionality reduction: Rank-based criteria. *Neurocomput.*, 72(7-9):1431–1443, 2009.
- [13] S. T. Roweis and L. K. Saul. Nonlinear Dimensionality Reduction by Locally Linear Embedding. *Science*, 290(5500):2323–2326, 2000.
- [14] Y. Teh and S. Roweis. Automatic alignment of local representations. *Advances in Neural Information Processing Systems*, 15:841–848, 2003.
- [15] Y. W. Teh and S. Roweis. Automatic alignment of local representations. In *In Advances in Neural Information Processing Systems 15*, pages 841–848. MIT Press, 2003.
- [16] J. B. Tenenbaum, V. d. Silva, and J. C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290(5500):2319–2323, 2000.
- [17] W. Torgerson. Multidimensional scaling, i: Theory and method. *Psychometrika*, 17:401–419, 1952.
- [18] L. van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, November 2008.
- [19] L. J. P. van der Maaten. Learning a parametric embedding by preserving local structure. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AI-STATS)*, number 5, pages 384–391. JMLR W&CP, 2009.
- [20] L. J. P. van der Maaten, E. O. Postma, and H. J. van den Herik. Dimensionality reduction: A comparative review. Technical Report TiCC-TR 2009-005, Tilburg University, Oct 2009.

- [21] J. Venna, J. Peltonen, K. Nybo, H. Aidos, and S. Kaski. Information retrieval perspective to nonlinear dimensionality reduction for data visualization. *J. Mach. Learn. Res.*, 11:451–490, 2010.
- [22] K. Q. Weinberger and L. K. Saul. An introduction to nonlinear dimensionality reduction by maximum an introduction to nonlinear dimensionality reduction by maximum variance unfolding. *Proceedings of the 21st National Conference on Artificial Intelligence*, 2006.

method	characteristics of data	characteristics of projections	error measure
MDS	Euclidean distance $d_{\mathcal{X}}(\vec{x}^i, \vec{x}^j)$	Euclidean distance $d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)$	minimize weighted least squared error
Isomap	Geodesic distance $d_{\text{geodesic}}(\vec{x}^i, \vec{x}^j)$	Euclidean distance $d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)$	minimize weighted least squared error
LLE	reconstruction weights w_{ij} such that $\sum(\vec{x}^i - \sum_{i \rightarrow j} w_{ij} \vec{x}^j)^2$ is minimum with constraints $\sum_j w_{ij} = 1$	reconstruction weights \tilde{w}_{ij} such that $\sum(\vec{y}^i - \sum_{i \rightarrow j} \tilde{w}_{ij} \vec{y}^j)^2$ is minimum with constraints $\sum \vec{y}^i = 0, \mathbf{Y}^t \mathbf{Y} = \mathbf{n}$	enforce identity $w_{ij} = \tilde{w}_{ij}$
Laplacian eigenmap	negative heat kernel weights $-w_{ij} = \exp(-d_{\mathcal{X}}(\vec{x}^i, \vec{x}^j)^2/t)$ for $i \rightarrow j$	squared Euclidean distance $d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)^2$ for $i \rightarrow j$ with constraints $\mathbf{Y}^t D \mathbf{Y} = \mathbf{1}, \mathbf{Y}^t D \vec{1} = \vec{0}$	maximize correlation
MVU	Euclidean distance $d_{\mathcal{X}}(\vec{x}^i, \vec{x}^j)$ for $i \rightarrow j$	Euclidean distance $d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)$ for $i \rightarrow j$ such that $\sum_{i,j} d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)^2$ is maximum and $\sum_i \vec{y}^i = 0$.	enforce identity (introducing slack variables if necessary)
SNE	probabilities $p_{j i} = \frac{\exp(-d_{\mathcal{X}}(\vec{x}^i, \vec{x}^j)^2/2\sigma_i)}{\sum_{k \neq i} \exp(-d_{\mathcal{X}}(\vec{x}^i, \vec{x}^k)^2/2\sigma_i)}$	probabilities $q_{j i} = \frac{\exp(-d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)^2)}{\sum_{k \neq i} \exp(-d_{\mathcal{E}}(\vec{y}^i, \vec{y}^k)^2)}$	minimize Kullback-Leibler divergences
t-SNE	probabilities $p_{ij} = \frac{p_{j i} + p_{i j}}{2n}$	probabilities $q_{ij} = \frac{(1+d_{\mathcal{E}}(\vec{y}^i, \vec{y}^j)/\varsigma)^{-\frac{\varsigma+1}{2}}}{\sum_{k \neq l} (1+d_{\mathcal{E}}(\vec{y}^k, \vec{y}^l)/\varsigma)^{-\frac{\varsigma+1}{2}}}$	minimize Kullback-Leibler divergence

Table 2: Many dimensionality reduction methods can be put into a general framework: characteristics of the data are extracted. Projections lead to corresponding characteristics depending on the coefficients. These coefficients are determined such that an error measure of the characteristics is minimized, fulfilling probably additional constraints.