

2004s-27

Spectral Dimensionality Reduction

*Yoshua Bengio, Olivier Delalleau,
Nicolas Le Roux, Jean-Francois Paiement,
Pascal Vincent, Marie Ouimet*

Série Scientifique
Scientific Series

Montréal
Mai 2004

© 2004 Yoshua Bengio, Olivier Delalleau, Nicolas Le Roux, Jean-Francois Paiement, Pascal Vincent, Marie Ouimet. Tous droits réservés. *All rights reserved.* Reproduction partielle permise avec citation du document source, incluant la notice ©.

Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source.

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du ministère de la Recherche, de la Science et de la Technologie, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Québec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the Ministère de la Recherche, de la Science et de la Technologie, and grants and research mandates obtained by its research teams.

Les organisations-partenaires / The Partner Organizations

PARTENAIRE MAJEUR

. Ministère du développement économique et régional et de la recherche [MDERR]

PARTENAIRES

. Alcan inc.
. Axa Canada
. Banque du Canada
. Banque Laurentienne du Canada
. Banque Nationale du Canada
. Banque Royale du Canada
. Bell Canada
. BMO Groupe Financier
. Bombardier
. Bourse de Montréal
. Caisse de dépôt et placement du Québec
. Développement des ressources humaines Canada [DRHC]
. Fédération des caisses Desjardins du Québec
. GazMétro
. Hydro-Québec
. Industrie Canada
. Ministère des Finances du Québec
. Pratt & Whitney Canada Inc.
. Raymond Chabot Grant Thornton
. Ville de Montréal

. École Polytechnique de Montréal
. HEC Montréal
. Université Concordia
. Université de Montréal
. Université du Québec à Montréal
. Université Laval
. Université McGill
. Université de Sherbrooke

ASSOCIE A :

. Institut de Finance Mathématique de Montréal (IFM²)
. Laboratoires universitaires Bell Canada
. Réseau de calcul et de modélisation mathématique [RCM²]
. Réseau de centres d'excellence MITACS (Les mathématiques des technologies de l'information et des systèmes complexes)

Les cahiers de la série scientifique (CS) visent à rendre accessibles des résultats de recherche effectuée au CIRANO afin de susciter échanges et commentaires. Ces cahiers sont écrits dans le style des publications scientifiques. Les idées et les opinions émises sont sous l'unique responsabilité des auteurs et ne représentent pas nécessairement les positions du CIRANO ou de ses partenaires.

This paper presents research carried out at CIRANO and aims at encouraging discussion and comment. The observations and viewpoints expressed are the sole responsibility of the authors. They do not necessarily represent positions of CIRANO or its partners.

Spectral Dimensionality Reduction

Yoshua Bengio^{*}, Olivier Delalleau[†], Nicolas Le Roux[‡], Jean-Francois Paiement[§], Pascal Vincent^{**}, Marie Ouimet^{††}

Résumé / Abstract

Dans cet article, nous étudions et développons un cadre unifié pour un certain nombre de méthodes non linéaires de réduction de dimensionnalité, telles que LLE, Isomap, LE (Laplacian Eigenmap) et ACP à noyaux, qui font de la décomposition en valeurs propres (d'où le nom "spectral"). Ce cadre inclut également des méthodes classiques telles que l'ACP et l'échelonnage multidimensionnel métrique (MDS). Il inclut aussi l'étape de transformation de données utilisée dans l'agrégation spectrale. Nous montrons que, dans tous les cas, l'algorithme d'apprentissage estime les fonctions propres principales d'un opérateur qui dépend de la densité inconnue de données et d'un noyau qui n'est pas nécessairement positif semi-défini. Ce cadre aide à généraliser certains modèles pour prédire les coordonnées des exemples hors-échantillons sans avoir à réentraîner le modèle. Il aide également à rendre plus transparent ce que ces algorithmes minimisent sur les données empiriques et donne une notion correspondante d'erreur de généralisation.

Mots clés : modèles non paramétriques, réduction de dimensionnalité non linéaire, modèles à noyau.

In this paper, we study and put under a common framework a number of non-linear dimensionality reduction methods, such as Locally Linear Embedding, Isomap, Laplacian Eigenmaps and kernel PCA, which are based on performing an eigen-decomposition (hence the name 'spectral'). That framework also includes classical methods such as PCA and metric multidimensional scaling (MDS). It also includes the data transformation step used in spectral clustering. We show that in all of these cases the learning algorithm estimates the principal eigenfunctions of an operator that depends on the unknown data density and on a kernel that is not necessarily positive semi-definite. This helps to generalize some of these algorithms so as to predict an embedding for out-of-sample examples without having to retrain the model. It also makes it more transparent what these algorithm are minimizing on the empirical data and gives a corresponding notion of generalization error.

Keywords: *non-parametric models, non-linear dimensionality reduction, kernel models.*

^{*} Université de Montréal et CIRANO, (514) 343-6804, Yoshua.Bengio@cirano.qc.ca.

[†] Olivier Delalleau, Université de Montréal, delallea@iro.umontreal.ca

[‡] Nicolas Le Roux, Université de Montréal, lerouxni@iro.umontreal.ca

[§] Jean-François Paiement, Université de Montréal, paiemeje@iro.umontreal.ca

^{**} Pascal Vincent, Université de Montréal, vincentp@iro.umontreal.ca

^{††} Marie Ouimet, Université de Montréal, ouimema@iro.umontreal.ca

1 Introduction

Unsupervised learning algorithms attempt to extract important characteristics of the unknown data distribution from the given examples. High-density regions are such salient features and they can be described by clustering algorithms (where typically each cluster corresponds to a high density “blob”) or by manifold learning algorithms (which discover high-density low-dimensional surfaces). A more generic description of the density is given by algorithms that estimate the density function.

In the context of supervised learning (each example is associated with a target label) or semi-supervised learning (a few examples are labeled but most are not), manifold learning algorithms can be used as pre-processing methods to perform dimensionality reduction. Each input example is then associated with a low-dimensional representation which corresponds to its estimated coordinates on the manifold. Since the manifold learning can be done without using the target labels, it can be applied on all of the input examples (both those labeled and those unlabeled). If there are many unlabeled examples, it has been shown that they can help to learn a more useful low-dimensional representation of the data (Belkin and Niyogi, 2003b). Dimensionality reduction is an interesting alternative to feature selection. Like feature selection it yields a low-dimensional representation which helps to build lower capacity predictors in order to improve generalization. However, unlike feature selection it may preserve information from all the original input variables. In fact, if the data really lie on a low-dimensional manifold, it may preserve almost all of the original information while representing it in a way that eases learning. For example, manifold learning algorithms such as those described in this paper often have the property of “unfolding” the manifold, i.e. flattening it out, as shown in Figure 1. On the other hand, these techniques being purely unsupervised, they may throw away low variance variations that are highly predictive of the target label. In addition to being useful as a preprocessing step for supervised or semi-supervised learning, linear and non-linear dimensionality reduction is often used for data analysis and visualization, e.g. (Vlachos et al., 2002), since visualizing the projections of the data (two or three dimensions at a time) can help to better understand it.

In the last few years, many unsupervised learning algorithms have been proposed which share the use of an eigen-decomposition for obtaining a lower-dimensional embedding of the data that characterizes a non-linear manifold near which the data would lie: Locally Linear Embedding (LLE) (Roweis and Saul, 2000), Isomap (Tenenbaum, de Silva and Langford, 2000) and Laplacian Eigenmaps (Belkin and Niyogi, 2003a). There are also many variants of spectral clustering (Weiss, 1999; Ng, Jordan and Weiss, 2002), in which such an embedding is an intermediate step before obtaining a clustering of the data that can capture flat, elongated and even curved clusters. The two tasks (manifold learning and clustering) are linked because the clusters that spectral clustering manages to capture can be arbitrary curved manifolds (as long as there is enough data to locally capture the curvature of the manifold): clusters and manifold both are zones of high density. An interesting advantage of the family of manifold learning algorithms described in this paper is that they can easily be applied in the case of non-vectorial data as well

as data for which no vectorial representation is available but for which a similarity function between objects can be computed, as in the MDS (multi-dimensional scaling) algorithms (Cox and Cox, 1994).

There are of course several dimensionality reduction methods that do not fall in the spectral framework described here, but which may have interesting connections nonetheless. For example, the principal curves algorithms (Hastie and Stuetzle, 1989; Kegl and Krzyzak, 2002) have been introduced based on geometric grounds, mostly for 1-dimensional manifolds. Although they optimize a different type of criterion, their spirit is close to that of LLE and Isomap. Another very interesting family of algorithms is the Self-Organizing Map (Kohonen, 1990). With these algorithms, the low dimensional embedding space is discretized (into topologically organized centers) and one learns the coordinates in the raw high-dimensional space of each of these centers. Another neural network like approach to dimensionality reduction is the auto-associative neural network (Rumelhart, Hinton and Williams, 1986; Saund, 1989), in which one trains a multi-layer neural network to predict its input, but forcing the intermediate representation of the hidden units to be a compact code. In section 2.7 we discuss in more detail a family of density estimation algorithms that can be written as mixtures of Gaussians with low-rank covariance matrices, having intimate connections with the LLE and Isomap algorithms.

An interesting question that will not be studied further in this paper is that of selecting the dimensionality of the embedding. This is fundamentally a question of model selection. It could be addressed using traditional model selection methods (such as cross-validation) when the low-dimensional representation is used as input for a supervised learning algorithm. Another approach is that of inferring the dimensionality based on purely unsupervised grounds, using the geometric properties of the empirical data distribution (Kégl, 2003).

1.1 Transduction and Induction

The end result of most inductive machine learning algorithms is a function that minimizes the empirical average of a loss criterion (possibly plus regularization). The function can be applied on new points and for such learning algorithms it is clear that the ideal solution is a function that minimizes the expected loss under the unknown true distribution from which the data was sampled, also known as the generalization error. However, such a characterization was missing for spectral embedding algorithms such as metric Multi-Dimensional Scaling (MDS) (Cox and Cox, 1994), spectral clustering (see (Weiss, 1999) for a review), Laplacian eigenmaps (Belkin and Niyogi, 2003a), Locally Linear Embedding (LLE) (Roweis and Saul, 2000) and Isomap (Tenenbaum, de Silva and Langford, 2000), which are used either for dimensionality reduction or for clustering. As such these algorithms are therefore really transduction algorithms: any test data for which an embedding is desired must be included in the data set on which the algorithm is applied. Their basic form does not provide a function that can be applied to new points, and the notion of generalization error that would be implicitly minimized is not clearly defined.

As a natural consequence of providing a unifying framework for these algorithms, we provide an answer to these questions. A loss criterion for spectral embedding algorithms can be defined. It is a reconstruction error that depends on pairs of examples. Minimizing its average value yields the eigenvectors that provide the classical output of these algorithms, i.e. the embeddings. Minimizing its expected value over the true underlying distribution yields the eigenfunctions of a linear operator (called G here) that is defined with a similarity function (a kernel, but not necessarily positive semi-definite) and the data-generating density. When the kernel is positive semi-definite and we work with the empirical density there is a direct correspondence between these algorithms and kernel Principal Components Analysis (PCA) (Schölkopf, Smola and Müller, 1998). Our work is also a direct continuation of previous work (Williams and Seeger, 2000) noting that the Nyström formula and the kernel PCA projection (which are equivalent) represent an approximation of the eigenfunctions of the above linear operator. Previous analysis of the convergence of generalization error of kernel PCA (Shawe-Taylor, Cristianini and Kandola, 2002; Shawe-Taylor and Williams, 2003; Zwald, Bousquet and Blanchard, 2004) also help to justify the view that these methods are estimating the convergent limit of some eigenvectors (at least when the kernel is positive semi-definite). The eigenvectors can then be turned into estimators of eigenfunctions, which can therefore be applied to new points, turning the spectral embedding algorithms into function induction algorithms. The Nyström formula obtained this way is well known (Baker, 1977), and will be given in eq. 2 below. This formula has been used previously for estimating extensions of eigenvectors in Gaussian process regression (Williams and Seeger, 2001), and it was noted (Williams and Seeger, 2000) that it corresponds to the projection of a test point computed with kernel PCA.

In order to extend spectral embedding algorithms such as LLE and Isomap to out-of-sample examples, this paper defines for these spectral embedding algorithms data-dependent kernels K_n that can be applied outside of the training set. See also the independent work (Ham et al., 2003) for a kernel view of LLE and Isomap, but where the kernels are only applied on the training set.

Additional contributions of this paper include a characterization of the empirically estimated eigenfunctions in terms of eigenvectors in the case where the kernel is not positive semi-definite (which is often the case for MDS and Isomap), a convergence theorem linking the Nyström formula to the eigenfunctions of G , as well as experiments on MDS, Isomap, LLE and spectral clustering / Laplacian eigenmaps showing that the Nyström formula for out-of-sample examples is accurate.

1.2 Notation

To simplify the presentation, we will consider the vector-space versions of these algorithms, in which we start from a data set $D = (x_1, \dots, x_n)$ with $x_i \in \mathbb{R}^d$ sampled i.i.d. from an unknown distribution with density p . However, the results in this paper can be readily extended to the case of arbitrary objects, with $p(x)dx$ replaced by $d\mu(x)$ with $\mu(x)$ an appropriate measure, and the only quantity that is required in the algorithms is the similarity or distance between pairs of objects (e.g. similarity $K_n(x_i, x_j)$)

below). See for example the treatment of pairwise measurement data for LLE (Saul and Roweis, 2002) and Isomap (Tenenbaum, de Silva and Langford, 2000), and for MDS in section 2.2.

Below we will use the notation

$$E_x[f(x)] = \int f(x)p(x)dx$$

for averaging over $p(x)$ and

$$\hat{E}_x[f(x)] = \frac{1}{n} \sum_{i=1}^n f(x_i)$$

for averaging over the data in D , i.e. over the empirical distribution denoted $\hat{p}(x)$. We will denote kernels with $K_n(x, y)$ or $\tilde{K}(x, y)$, symmetric functions, not always positive semi-definite, that may depend not only on x and y but also on the data D . The spectral embedding algorithms construct an affinity matrix M , either explicitly through

$$M_{ij} = K_n(x_i, x_j) \tag{1}$$

or implicitly through a procedure that takes the data D and computes M . We denote by v_{ik} the i -th coordinate of the k -th eigenvector of M (sorted in order of decreasing eigenvalues), associated with the eigenvalue ℓ_k . With these notations, the Nyström formula discussed above can be written:

$$f_{k,n}(x) = \frac{\sqrt{n}}{\ell_k} \sum_{i=1}^n v_{ik} K_n(x, x_i) \tag{2}$$

where $f_{k,n}$ is the k -th Nyström estimator with n samples. We will show in section 3 that it estimates the k -th eigenfunction of a linear operator and that it provides an embedding for new example x .

2 Data-Dependent Kernels for Spectral Embedding Algorithms

The first and foremost observation to make is that many algorithm spectral embedding algorithms can be cast in a common framework. The spectral embedding algorithms can be seen to build a $(n \times n)$ similarity matrix M (also called the *Gram* matrix)¹ and compute its principal eigenvectors $v_k = (v_{1k}, \dots, v_{nk})'$ (one entry per exemple) with eigenvalues ℓ_k (sorted by decreasing order). *The embedding associated with the i -th training example is given by the i -th element of the principal eigenvectors, up to some scaling:*

$$(v_{i1}, v_{i2}, \dots, v_{im})' \tag{3}$$

¹For Laplacian Eigenmaps (section 2.4) and LLE (section 2.6), the matrix M discussed here is not the one defined in the original papers on these algorithms, but a transformation of it to reverse the order of eigenvalues, as we see below.

where $m \leq n$ is the desired number of embedding coordinates. For instance, in kernel PCA, MDS and Isomap, v_{ik} is multiplied by $\sqrt{\ell_k}$, and in LLE it is multiplied by \sqrt{n} to obtain the actual embedding coordinates.

In general, we will see that M_{ij} depends not only on (x_i, x_j) but also on the other training examples. Nonetheless, as we show below, it can always be written $M_{ij} = K_n(x_i, x_j)$ where K_n is a “data-dependent” kernel (i.e. it is a function of the n elements of the training set D , and not just of its two arguments). In many algorithms a matrix \tilde{M} is first formed from a simpler, often data-independent kernel (such as the Gaussian kernel), and then transformed into M . We want to think of the entries of M as being generated by applying K_n to the pairs (x_i, x_j) because this will help us to generalize to new examples not in the training set D , and it will help us to think about what happens as n increases.

For each of these methods, by defining a kernel K_n that can be applied outside of the training set, we will be able to generalize the embedding to a new point x , via the Nyström formula (eq. 2 above, and section 3.2). This will only require computations of the form $K_n(x, x_i)$ with x_i a training point.

2.1 Kernel Principal Components Analysis

Kernel PCA is an unsupervised manifold learning technique that maps data points to a new space, generally lower-dimensional (but not necessarily). It generalizes the Principal Components Analysis approach to non-linear transformations using the kernel trick (Schölkopf, Smola and Müller, 1996; Schölkopf, Smola and Müller, 1998; Schölkopf, Burges and Smola, 1999). One considers the data mapped into a “feature space”, a Hilbert space of possibly infinite dimension such that if x is mapped to $\tilde{\phi}(x)$, we have $\langle \tilde{\phi}(x), \tilde{\phi}(y) \rangle = \tilde{K}(x, y)$. Here, \tilde{K} must be a positive (semi)-definite kernel, and is often taken as the Gaussian kernel, i.e.

$$\tilde{K}(x, y) = e^{-\frac{\|x-y\|^2}{\sigma^2}}.$$

The kernel PCA algorithm consists in performing PCA in the feature space: it implicitly finds the leading eigenvectors and eigenvalues of the covariance of the projection $\tilde{\phi}(x)$ of the data. *If the data is centered in feature space* ($\hat{E}_x[\tilde{\phi}(x)] = 0$), the (empirical) feature space covariance matrix is $C = \hat{E}_x[\tilde{\phi}(x)\tilde{\phi}(x)']$. In general, however, the data is not centered, and we need to define a “centered” mapping

$$\phi_n(x) = \tilde{\phi}(x) - \frac{1}{n} \sum_{i=1}^n \tilde{\phi}(x_i)$$

and an associated *data-dependent* kernel K_n such that $K_n(x, y) = \langle \phi_n(x), \phi_n(y) \rangle$, which rewrites:

$$K_n(x, y) = \tilde{K}(x, y) - \hat{E}_{x'}[\tilde{K}(x', y)] - \hat{E}_{y'}[\tilde{K}(x, y')] + \hat{E}_{x'}[\hat{E}_{y'}[\tilde{K}(x', y')]]. \quad (4)$$

The empirical covariance matrix C in “feature space” is thus actually defined by

$$C = \hat{E}_x[\phi_n(x)\phi_n(x)'] \quad (5)$$

with eigenvectors w_k associated with eigenvalues λ_k . As shown in (Schölkopf, Smola and Müller, 1998), this eigen-decomposition of C is related to the one of M (the Gram matrix defined by eq. 1) through $\lambda_k = \ell_k/n$ and

$$w_k = \frac{1}{\sqrt{\ell_k}} \sum_{i=1}^n v_{ik} \phi_n(x_i)$$

where v_k are the eigenvectors of M , associated with eigenvalues ℓ_k . As in PCA, one can then obtain the embedding of a training point x_i by the projection of $\phi_n(x_i)$ on the leading eigenvectors (w_1, \dots, w_m) of C , which yields exactly the embedding of eq. 3, if we multiply v_{ik} by $\sqrt{\ell_k}$.

Note that, as in PCA, we can also compute the projection $P(x) = (P_1(x), \dots, P_m(x))'$ for a new point x , which is written

$$P_k(x) = \langle w_k, \phi_n(x) \rangle = \frac{1}{\sqrt{\ell_k}} \sum_{i=1}^n v_{ik} K_n(x_i, x). \quad (6)$$

This is the key observation that will allow us, in section 3.2, to extend to new points the embedding obtained with other spectral algorithms.

2.2 Multi-Dimensional Scaling

Metric Multi-Dimensional Scaling (MDS) (Cox and Cox, 1994) starts from a notion of distance $d(x, y)$ that is computed between each pair of training examples to fill a matrix $\tilde{M}_{ij} = d^2(x_i, x_j)$. The idea is to find a low-dimensional embedding of the dataset D that preserves the given distances between training points. To do so, the distances are converted to equivalent dot products using the “double-centering” formula, which makes M_{ij} depend not only on (x_i, x_j) but also on all the other examples:

$$M_{ij} = -\frac{1}{2}(\tilde{M}_{ij} - \frac{1}{n}S_i - \frac{1}{n}S_j + \frac{1}{n^2} \sum_k S_k) \quad (7)$$

where $S_i = \sum_{j=1}^n \tilde{M}_{ij}$. The embedding of the example x_i is then given by eq. 3 with v_{ik} multiplied by $\sqrt{\ell_k}$.

A corresponding data-dependent kernel which generates the matrix M is:

$$K_n(a, b) = -\frac{1}{2}(d^2(a, b) - \hat{E}_x[d^2(x, b)] - \hat{E}_{x'}[d^2(a, x')] + \hat{E}_{x, x'}[d^2(x, x')]). \quad (8)$$

2.3 Spectral Clustering

Several variants of spectral clustering have been proposed (Weiss, 1999). They can yield impressively good results where traditional clustering looking for “round blobs” in the data, such as K-means, would fail miserably (see Figure 1). It is based on two main steps: first embedding the data points in a space in which clusters are more “obvious” (using the eigenvectors of a Gram matrix), and then applying a classical clustering algorithm such as K-means, e.g. as in (Ng, Jordan and Weiss, 2002). To construct the spectral clustering affinity matrix M , we first apply a data-independent kernel \tilde{K} such as the Gaussian kernel to each pair of examples: $\tilde{M}_{ij} = \tilde{K}(x_i, x_j)$. The matrix \tilde{M} is then normalized, e.g. using “divisive” normalization (Weiss, 1999; Ng, Jordan and Weiss, 2002)²:

$$M_{ij} = \frac{\tilde{M}_{ij}}{\sqrt{S_i S_j}}. \quad (9)$$

To obtain m clusters, the first m principal eigenvectors of M are computed and K-means is applied on the embedding coordinates after normalizing each embedding vector to have unit norm: the k -th coordinate of the i -th example is $v_{ik} / \sqrt{\sum_{l=1}^m v_{il}^2}$. Note for that for such an embedding to be stable (to keep the same meaning) when n varies, it should be multiplied by \sqrt{n} .

To generalize spectral clustering to out-of-sample points, we will need a kernel that could have generated that matrix M :

$$K_n(a, b) = \frac{1}{n} \frac{\tilde{K}(a, b)}{\sqrt{\hat{E}_x[\tilde{K}(a, x)] \hat{E}_{x'}[\tilde{K}(x', b)]}}. \quad (10)$$

Note that this divisive normalization comes out of the justification of spectral clustering as a relaxed statement of the min-cut problem (Chung, 1997; Spielman and Teng, 1996) (to divide the examples into two groups such as to minimize the sum of the “similarities” between pairs of points straddling the two groups). The additive normalization performed with kernel PCA (eq. 4) makes sense geometrically as a centering in feature space. Both the divisive normalization and the additive normalization procedures make use of a kernel row/column average. It would be interesting to find a similarly pleasing geometric interpretation to the divisive normalization.

2.4 Laplacian Eigenmaps

The Laplacian Eigenmaps method is a recently proposed dimensionality reduction procedure (Belkin and Niyogi, 2003a) that was found to be very successful for *semi-supervised learning*. Several variants have been proposed by the authors and we focus here on the latest one, but they all share the same spirit.

²Better embeddings are usually obtained if we define $S_i = \sum_{j \neq i} \tilde{M}_{ij}$: this alternative normalization can also be cast into the general framework developed here, with a slightly different kernel. Also, one could take the row average instead of the row sum, which seems more natural even if it doesn’t change the embedding.




Figure 1: *Example of the transformation learned as part of spectral clustering. Input data on the left, transformed data on the right. Gray level and cross/circle drawing are only used to show which points get mapped where: the mapping reveals both the clusters and the internal structure of the two manifolds.*

The authors use an approximation of the Laplacian operator such as the Gaussian kernel or the k -nearest-neighbor graph: the *symmetric* matrix \tilde{M} whose element (i, j) is 1 if x_i and x_j are k -nearest-neighbors (x_i is among the k nearest neighbors of x_j or vice versa) and 0 otherwise. Instead of an ordinary eigenproblem, the following generalized eigenproblem is solved:

$$(S - \tilde{M})y_k = \sigma_k S y_k \quad (11)$$

with eigenvalues σ_k , eigenvectors y_k and S the diagonal matrix with elements S_i previously defined (row sums of \tilde{M}). The smallest eigenvalue is left out and the eigenvectors corresponding to the other small eigenvalues are used for the embedding. This is actually the same embedding that is computed with the spectral clustering algorithm from (Shi and Malik, 1997): as noted in (Weiss, 1999) (Normalization Lemma 1), an equivalent result (up to a component-wise scaling of the embedding) can be obtained by considering the principal eigenvectors v_k of the normalized matrix M defined in eq. 9. To fit the common framework for spectral embedding in this paper, we have used the latter formulation. Therefore, the same data-dependent kernel can be defined as for spectral clustering (eq. 10) to generate the matrix M , i.e. spectral clustering just adds a clustering step after a Laplacian Eigenmaps dimensionality reduction.

2.5 Isomap

Isomap (Tenenbaum, de Silva and Langford, 2000) generalizes MDS to non-linear manifolds. It is based on replacing the Euclidean distance by an empirical approximation of the geodesic distance on the manifold. We define the *geodesic distance* $\mathcal{D}(\cdot, \cdot)$ with respect to a data set D , a distance $d(\cdot, \cdot)$ and a neighborhood k as follows:

$$\mathcal{D}(a, b) = \min_{\pi} \sum_{i=1}^{|\pi|} d(\pi_i, \pi_{i+1}) \quad (12)$$

where π is a sequence of points of length $|\pi| = l \geq 2$ with $\pi_1 = a$, $\pi_l = b$, $\pi_i \in D \forall i \in \{2, \dots, l-1\}$ and (π_i, π_{i+1}) are k -nearest-neighbors of each other. The length $|\pi| = l$ is free in the minimization. The Isomap algorithm obtains the normalized matrix M from which the embedding is derived by transforming the raw pairwise distances matrix as follows: (1) compute the matrix $\tilde{M}_{ij} = \mathcal{D}^2(x_i, x_j)$ of squared geodesic distances with respect to the data D and (2) apply to this matrix the distance-to-dot-product transformation (eq. 7), as for MDS. As in MDS, the embedding of x_i is given by eq. 3 with v_{ik} multiplied by $\sqrt{\ell_k}$. Step (1) can be done in $O(n^3)$ operations very easily (e.g. by Floyd’s algorithm), but in (Tenenbaum, de Silva and Langford, 2000) it is suggested to use more efficient algorithms exploiting the sparse structure of the neighborhood graph, such as those presented in (Kumar et al., 1994).

There are several ways to define a kernel that generates M and also generalizes out-of-sample. The solution we have chosen simply computes the geodesic distances without involving the out-of-sample point(s) along the geodesic distance sequence (except for the last distance). This is automatically achieved with the above definition of geodesic distance \mathcal{D} , which only uses the training points to find the shortest path between a and b . The double-centering kernel transformation of eq. 8 can then be applied to obtain K_n , using the geodesic distance \mathcal{D} instead of the MDS distance d .

2.6 Locally Linear Embedding

The Locally Linear Embedding (LLE) algorithm (Roweis and Saul, 2000) looks for an embedding that preserves the local geometry in the neighborhood of each data point. The idea is to find a low-dimensional representation where the reconstruction of a data point from its neighbors is similar to the one in high dimension. First, a sparse matrix of local predictive weights W_{ij} is computed, such that $\sum_j W_{ij} = 1$, $W_{ii} = 0$, $W_{ij} = 0$ if x_j is not a k -nearest-neighbor of x_i and $\|(\sum_j W_{ij}x_j) - x_i\|^2$ is minimized. To find those weights, for a given training point x_i with neighbours (y_{i1}, \dots, y_{ik}) , a local Gram matrix $M^{(i)}$ is computed, such that $M_{rs}^{(i)} = \langle y_{ir} - x_i, y_{is} - x_i \rangle$. To improve the condition of this Gram matrix (to avoid potential issues when solving the linear system below), it is recommended to add a small multiple of the identity matrix:

$$M_{rs}^{(i)} \leftarrow M_{rs}^{(i)} + \delta_{rs} \frac{\Delta^2}{k} Tr(M^{(i)})$$

with Tr the trace operator, and $\Delta^2 \ll 1$. The weights are then obtained by solving the linear system defined by $\sum_r M_{rs}^{(i)} W_{is} = 1$ for all s , then rescaling the W_{is} so that they sum to 1 (Saul and Roweis, 2002).

From the weights W_{ij} , the matrix $\tilde{M} = (I - W)'(I - W)$ is formed. The embedding is obtained from the lowest eigenvectors of \tilde{M} , except for the eigenvector with the smallest eigenvalue, which is uninteresting because it is proportional to $(1, 1, \dots, 1)$ (and its eigenvalue is 0). Since we want to select the principal eigenvectors, we define our normalized matrix by $M = cI - \tilde{M}$ (c being any real number) and ignore the top eigenvector (although one could apply an additive normalization to remove the

components along the $(1, 1, \dots, 1)$ direction). The LLE embedding for x_i is then given by eq. 3 (multiplied by \sqrt{n}), starting at the second eigenvector (since the principal one is constant). If one insists on having a positive semi-definite matrix M , one can take for c the largest eigenvalue of \tilde{M} (note that c only changes the eigenvalues additively and has no influence on the embedding of the training set).

In order to define a kernel K_n generating M , we first denote by $w(x, x_i)$ the weight of x_i in the reconstruction of any point $x \in \mathbb{R}^d$ by its k nearest neighbors in the training set. This is the same reconstruction as above, i.e. the $w(x, x_i)$ are such that they sum to 1, $\|(\sum_i w(x, x_i)x_i) - x\|^2$ is minimized, and $w(x, x_i) = 0$ if x_i is not in the k nearest neighbors of x . If $x = x_j \in D$, we have $w(x, x_i) = \delta_{ij}$. Let us now define a kernel K'_n by $K'_n(x_i, x) = K'_n(x, x_i) = w(x, x_i)$ and $K'_n(x, y) = 0$ when neither x nor y is in the training set D . Let K''_n be such that $K''_n(x_i, x_j) = W_{ij} + W_{ji} - \sum_k W_{ki}W_{kj}$ and $K''_n(x, y) = 0$ when either x or y is not in D . It can be shown that the kernel $K_n = (c - 1)K'_n + K''_n$ is then such that

$$K_n(x_i, x_j) = (c - 1)\delta_{ij} + W_{ij} + W_{ji} - \sum_k W_{ki}W_{kj} = M_{ij}$$

so that it can be used to generate M . There could be other ways to obtain a data-dependent kernel for LLE that can be applied out-of-sample: a justification for using this specific kernel will be given in section 3.1.

As noted independently in (Ham et al., 2003), LLE can be seen as performing kernel PCA with a particular kernel matrix. This identification becomes even more accurate when one notes that getting rid of the constant eigenvector (principal eigenvector of M) is equivalent to the centering operation in feature space required for kernel PCA (Ham et al., 2003).

It is interesting to note a recent descendant of Laplacian Eigenmaps, Isomap and LLE, called Hessian Eigenmaps (Donoho and Grimes, 2003), which considers the limit case of the continuum of the manifold, and replaces the Laplacian in Laplacian Eigenmaps by a Hessian.

2.7 Mixtures of Low-Rank Gaussians

Isomap and LLE are two instances of a larger family of unsupervised learning algorithms which characterize the data distribution by a large set of locally linear low-dimensional patches. A simple example of this type of model is a mixture of Gaussians (centered on each example in the standard non-parametric setting) whose covariance matrices are summarized by a few eigenvectors (i.e. principal directions). The mixture of factor analyzers (Ghahramani and Hinton, 1996) is a parametric version of this type of model, in which the EM algorithm is used to estimate the means and the low-rank covariance matrices. A non-parametric version of the mixture of factor analyzers aimed at capturing manifold structure is the Manifold Parzen Windows algorithm (Vincent and Bengio, 2003), which does not require an iterative algorithm for training. With such models, one can obtain a local low-dimensional representation of examples falling near a Gaussian center, but it may be incomparable to the representation obtained for a nearby Gaussian

center, because the eigenvectors of the covariance matrices of neighboring Gaussians may not be aligned. In order to perform dimensionality reduction from such models, several algorithms have thus been proposed (Teh and Roweis, 2003; Brand, 2003; Verbeek, Roweis and Vlassis, 2004), which look for a global representation that agrees with each local patch. Although these algorithms do not fall into the “spectral manifold learning” family studied in more detailed in this paper, they are very close in spirit.

3 Kernel Eigenfunctions for Induction

With the exception of kernel PCA, the spectral manifold learning algorithms presented in section 2 do not provide us with an immediate way to obtain the embedding for a new point $x \notin D$. However, for some of them, extensions have already been proposed. We briefly review them in section 3.1. In section 3.2, we take advantage of the common framework developed in section 2: each algorithm being associated with a data-dependent kernel K_n generating a Gram matrix M , we can apply the Nyström formula (eq. 2) to obtain the embedding for a new point x .

3.1 Extensions to Spectral Embedding Algorithms

For metric MDS, it is suggested in (Gower, 1968), to solve exactly for the coordinates of the new point that are consistent with its distances to the training points, but in general this requires adding a new dimension. Note also that (Williams, 2001) makes a connection between kernel PCA and metric MDS, remarking that kernel PCA is a form of MDS when the kernel is isotropic. In the following, we will pursue this connection in order to obtain out-of-sample embeddings.

A formula has been proposed (de Silva and Tenenbaum, 2003) to approximate Isomap using only a subset of the examples (the “landmark” points) to compute the eigenvectors. Using the notation of this paper, that formula is

$$e_k(x) = \frac{1}{2\sqrt{\ell_k}} \sum_i v_{ik} (\hat{E}_{x'}[\mathcal{D}^2(x', x_i)] - \mathcal{D}^2(x_i, x)) \quad (13)$$

which is applied to obtain an embedding for the non-landmark examples. One can show (Bengio et al., 2004) that $e_k(x)$ is the Nyström formula when $K_n(x, y)$ is defined as in section 2.5. Landmark Isomap is thus equivalent to performing Isomap on the landmark points only and then predicting the embedding of the other points using the Nyström formula, which is the solution we also propose in what follows.

For LLE, with the notations of section 2.6, an extension suggested in (Saul and Roweis, 2002) is to take for a new point x the embedding $P(x) = (P_1(x), \dots, P_m(x))'$, where

$$P_k(x) = \sum_{i=1}^n P_k(x_i) w(x, x_i).$$

Interestingly, the same embedding can be obtained from the Nyström formula and the kernel K_n defined in section 2.6, when the constant $c \rightarrow +\infty$ (Bengio et al., 2004).

3.2 From Eigenvectors to Eigenfunctions

From the common framework developed in section 2, one can see spectral algorithms as performing a kind of kernel PCA with a specific kernel. (Ng, Jordan and Weiss, 2002) had already noted the link between kernel PCA and spectral clustering. Recently, (Ham et al., 2003) have also shown how Isomap, LLE and Laplacian Eigenmaps can be interpreted as performing a form of kernel PCA. Here, we propose a similar view, extending the framework to allow negative eigenvalues (which may be the case for Isomap). In addition, those papers did not propose to use this link in order to perform function induction, i.e. obtain an embedding for out-of-sample points. Indeed, since there exists a natural extension to new points for kernel PCA (the projection onto the eigenspaces of the covariance matrix, see eq. 6), it is natural to ask whether it makes sense to use such a formula in the more general setting where the kernel may be data-dependent and may have negative eigenvalues.

As noted in (Williams and Seeger, 2000), the kernel PCA projection formula (eq. 6) is proportional to the so-called Nyström formula (Baker, 1977; Williams and Seeger, 2000) (eq. 2), which has been used successfully to “predict” the value of an eigenvector on a new data point, in order to speed-up kernel methods computations by focusing the heavier computations (the eigen-decomposition) on a subset of examples (Williams and Seeger, 2001). The use of this formula can be justified by considering the convergence of eigenvectors and eigenvalues, as the number of examples increases (Baker, 1977; Koltchinskii, 1998; Koltchinskii and Giné, 2000; Williams and Seeger, 2000). In particular, (Shawe-Taylor, Cristianini and Kandola, 2002; Shawe-Taylor and Williams, 2003; Zwald, Bousquet and Blanchard, 2004) give bounds on the kernel PCA convergence error (in the sense of the projection error with respect to the subspace spanned by the eigenvectors), using concentration inequalities.

Based on this kernel PCA convergence results, we conjecture that in the limit, each eigenvector would converge to an eigenfunction for a linear operator (defined below), in the sense that the i -th element of the k -th eigenvector converges to the application of the k -th eigenfunction to x_i . Proposition 2 below formalizes this statement and provides sufficient conditions for convergence of the eigenvectors to eigenfunctions.

In the following we will assume that the (possibly data-dependent) kernel K_n is *bounded* (i.e. $\exists K_{max}, \forall x, y |K_n(x, y)| < K_{max}$) and has a discrete spectrum, i.e. that it can be written as a discrete expansion

$$K_n(x, y) = \sum_{k=1}^{\infty} \alpha_{k,n} \phi_{k,n}(x) \phi_{k,n}(y).$$

Consider the space \mathcal{H}_p of continuous functions f on \mathbb{R}^d that are square integrable as follows:

$$\int f^2(x)p(x)dx < \infty$$

with the data-generating density function $p(x)$. One must note that we actually do not work on functions but on equivalence classes: we say two continuous functions f

and g belong to the same equivalence class (with respect to p) if and only if $\int (f(x) - g(x))^2 p(x) dx = 0$ (if p is strictly positive, then each equivalence class contains only one function).

We will assume that K_n converges uniformly in its arguments (in some probabilistic manner, e.g. almost surely or in probability) to its limit K as $n \rightarrow \infty$. We will associate with each K_n a linear operator G_n and with K a linear operator G , such that for any $f \in \mathcal{H}_p$,

$$G_n f = \frac{1}{n} \sum_{i=1}^n K_n(\cdot, x_i) f(x_i) \quad (14)$$

and

$$G f = \int K(\cdot, y) f(y) p(y) dy \quad (15)$$

which makes sense because we work in a space of functions defined everywhere. Furthermore, as $K_n(\cdot, y)$ and $K(\cdot, y)$ are square-integrable in the sense defined above, for each f and each n , the functions $G_n f$ and $G f$ are square-integrable in the sense defined above. We will show that the Nyström formula (eq. 2) gives the eigenfunctions of G_n (Proposition 1), that their value on the training examples corresponds to the spectral embedding, and that they converge to the eigenfunctions of G (Proposition 2). These results will hold even if K_n has negative eigenvalues.

The eigensystems of interest are thus the following:

$$G f_k = \lambda_k f_k \quad (16)$$

and

$$G_n f_{k,n} = \lambda_{k,n} f_{k,n} \quad (17)$$

where (λ_k, f_k) and $(\lambda_{k,n}, f_{k,n})$ are the corresponding eigenvalues and eigenfunctions.

Note that when eq. 17 is evaluated only at the $x_i \in D$, the set of equations reduces to the eigensystem

$$M v_k = n \lambda_{k,n} v_k.$$

The following proposition gives a more complete characterization of the eigenfunctions of G_n , even in the case where eigenvalues may be negative. The next two propositions formalize the link already made in (Williams and Seeger, 2000) between the Nyström formula and eigenfunctions of G .

Proposition 1 G_n has in its image $m \leq n$ eigenfunctions of the form:

$$f_{k,n}(x) = \frac{\sqrt{n}}{\ell_k} \sum_{i=1}^n v_{ik} K_n(x, x_i) \quad (18)$$

with corresponding non-zero eigenvalues $\lambda_{k,n} = \frac{\ell_k}{n}$, where $v_k = (v_{1k}, \dots, v_{nk})'$ is the k -th eigenvector of the Gram matrix M , associated with the eigenvalue ℓ_k .

For $x_i \in D$ these functions coincide with the corresponding eigenvectors, in the sense that $f_{k,n}(x_i) = \sqrt{n} v_{ik}$.

Proof

First, we show that the $f_{k,n}$ defined by eq. 18 coincide with the eigenvectors of M at $x_i \in D$. For $f_{k,n}$ associated with a non-zero eigenvalue,

$$f_{k,n}(x_i) = \frac{\sqrt{n}}{\ell_k} \sum_{j=1}^n v_{jk} K_n(x_i, x_j) = \frac{\sqrt{n}}{\ell_k} \ell_k v_{ik} = \sqrt{n} v_{ik}. \quad (19)$$

The v_k being orthonormal the $f_{k,n}$ (for different values of k) are therefore different from each other.

Then for any $x \in \mathbb{R}^d$

$$(G_n f_{k,n})(x) = \frac{1}{n} \sum_{i=1}^n K_n(x, x_i) f_{k,n}(x_i) = \frac{1}{\sqrt{n}} \sum_{i=1}^n K_n(x, x_i) v_{ik} = \frac{\ell_k}{n} f_{k,n}(x) \quad (20)$$

which shows that $f_{k,n}$ is an eigenfunction of G_n with eigenvalue $\lambda_{k,n} = \ell_k/n$. \square

Discussion

The previous result shows that the Nyström formula generalizes the spectral embedding outside of the training set. This means the embedding $P(x)$ for a new point x is given (up to some scaling) by

$$P_k(x) = \frac{f_{k,n}(x)}{\sqrt{n}} = \frac{1}{\ell_k} \sum_{i=1}^n v_{ik} K_n(x, x_i)$$

where the scaling is the same as the one described in section 2 (so that the embedding obtained on the training set is coherent with the one obtained from the eigenvectors, thanks to eq. 19).

However, there could be many possible generalizations. To justify the use of this particular generalization, the following proposition helps to understand the convergence of these functions as n increases. We would like the out-of-sample embedding predictions obtained with the Nyström formula to be somehow close to the asymptotic embedding (the embedding one would obtain as $n \rightarrow \infty$).

Note also that the convergence of eigenvectors to eigenfunctions shown in (Baker, 1977) applies to data x_i which are deterministically chosen to span a domain, whereas here the x_i form a random sample from an unknown distribution.

Proposition 2 *If $K_n = K$ is bounded and not data-dependent, then the eigenfunctions $f_{k,n}$ of G_n associated with non-zero eigenvalues of multiplicity 1 converge to the corresponding eigenfunctions of G (almost surely, and up to the sign).*

For K_n data-dependent but bounded (almost surely, and independently of n) and converging uniformly to K , if the eigen-decomposition of the Gram matrix $(K_n(x_i, x_j))$ converges³ to the eigen-decomposition of the Gram matrix $(K(x_i, x_j))$ then a similar result holds: the eigenfunctions $f_{k,n}$ of G_n associated with non-zero eigenvalues of multiplicity 1 converge to the corresponding eigenfunctions of G (almost surely, and up to the sign).

³the convergences should be almost sure, otherwise the result may hold with a different kind of probabilistic convergence, e.g. in probability.

Proof

In the following, we will denote by $\hat{f} \in \mathcal{H}_{\hat{p}}$ the restriction of a function $f \in \mathcal{H}_p$ to the training set $D = \{x_1, \dots, x_n\}$, and by \hat{G}_n the operator in $\mathcal{H}_{\hat{p}}$ defined as in eq. 14, which has the same eigenvalues and eigenfunctions as G_n (except the eigenfunctions are restricted to D). We start with the case where $K_n = K$. We first take advantage of (Koltchinskii and Giné, 2000), theorem 3.1, that shows that the distance between the eigenvalue spectra of \hat{G}_n and G converges to 0 almost surely. We then use theorem 2.1 from (Koltchinskii, 1998), which is stated as follows. Let K be a symmetric kernel such that $E[|K(X, X)|] < +\infty$ and $E[K^2(X, Y)] < +\infty$ (so that the operator G defined by eq. 15 is Hilbert-Schmidt and K can be written $K(x, y) = \sum_{i \in I} \mu_i \psi_i(x) \psi_i(y)$ with I a discrete set). Suppose that \mathcal{F} is a class of measurable functions such that there exists $F \in \mathcal{H}_p$ verifying $|f(x)| \leq F(x)$ for all $f \in \mathcal{F}$. Moreover, suppose that for all $i \in I$, $\{f \psi_i : f \in \mathcal{F}\} \in GC(p)$, where $GC(p)$ denotes the set of p -Glivenko-Cantelli classes (see, e.g., (van der Vaart and Wellner, 1996)). Then, for all non-zero eigenvalue λ_k

$$\sup_{f, g \in \mathcal{F}} \left| \langle P_k(\hat{G}_n) \hat{f}, \hat{g} \rangle_{\mathcal{H}_{\hat{p}}} - \langle P_k(G) f, g \rangle_{\mathcal{H}_p} \right| \rightarrow 0 \quad (21)$$

almost surely when $n \rightarrow +\infty$, with $P_k(G)$ the projection on the k -th eigenspace of G , and $P_k(\hat{G}_n)$, with probability 1 and for n sufficiently large, the projection on the corresponding eigenspace of \hat{G}_n (for more details see (Koltchinskii, 1998)).

Let us consider the k -th eigenspace of G (of dimension 1 because we have considered eigenvalues of multiplicity 1), i.e. the eigenspace spanned by the eigenfunction f_k : the k -th eigenspace of \hat{G}_n is also 1-dimensional, almost surely (because of the convergence of the spectrum), and spanned by $f_{k,n}$. Let $x \in \mathbb{R}^d$ be any point in the input space, and $\mathcal{F} = \{h_x\}$ with $h_x = K(x, \cdot) \in \mathcal{H}_p$. For any $i \in I$

$$\left| \frac{1}{n} \sum_j h_x(x_j) \psi_i(x_j) - \int h_x(y) \psi_i(y) p(y) dy \right| \rightarrow 0$$

almost surely (thanks to the strong law of large numbers), so that \mathcal{F} verifies the hypothesis needed to apply the theorem above. In addition,

$$\langle P_k(G) h_x, h_x \rangle_{\mathcal{H}_p} = \langle h_x, f_k \rangle_{\mathcal{H}_p}^2 = (G f_k)(x)^2 = \lambda_k^2 f_k(x)^2$$

and similarly, using eq. 20, we have with probability 1 and for n large enough:

$$\langle P_k(\hat{G}_n) \hat{h}_x, \hat{h}_x \rangle_{\mathcal{H}_{\hat{p}}} = \langle \hat{h}_x, \hat{f}_{k,n} \rangle_{\mathcal{H}_{\hat{p}}}^2 = (G_n f_{k,n})(x)^2 = \lambda_{k,n}^2 f_{k,n}(x)^2. \quad (22)$$

The conclusion of the theorem thus tells us that

$$|\lambda_{k,n}^2 f_{k,n}(x)^2 - \lambda_k^2 f_k(x)^2| \rightarrow 0$$

almost surely. Since we have the convergence of the eigenvalues, this implies

$$|f_{k,n}(x)^2 - f_k(x)^2| \rightarrow 0 \quad (23)$$

almost surely, which shows the (simple) convergence of the eigenfunctions, up to the sign. To get the convergence in \mathcal{H}_p , we need to show that $\|(|f_{k,n}| - |f_k|)\|_{\mathcal{H}_p} \rightarrow 0$, i.e.

$$\int g_{k,n}(x)dx \rightarrow 0 \quad (24)$$

with $g_{k,n}(x) = (|f_{k,n}(x)| - |f_k(x)|)^2 p(x)$. We will need to show that both $f_{k,n}$ and f_k are bounded (independently of n). Since f_k is an eigenfunction of G , we have $|\lambda_k f_k(x)| = |(Gf_k)(x)| = |\int K(x, y)f_k(y)p(y)dy| \leq c|\int f_k(y)p(y)dy|$, so that $|f_k(x)| \leq c'_k$. For $f_{k,n}$, we have

$$|\lambda_{k,n} f_{k,n}(x)| = |(G_n f_{k,n})(x)| = \left| \frac{1}{n} \sum_{i=1}^n K(x, x_i) \sqrt{n} v_{ik} \right| \leq \frac{c}{\sqrt{n}} \sum_{i=1}^n |v_{ik}| \leq c$$

because the maximum of $\sum_{i=1}^n |v_{ik}|$ subject to $\|v_k\|^2 = 1$ is \sqrt{n} , so that $|f_{k,n}(x)| \leq c''_k$ almost surely (thanks to the convergence of $\lambda_{k,n}$). Therefore, we have that (almost surely) $g_{k,n}(x) \leq (c'_k + c''_k)^2 p(x)$ which is an integrable function, and from eq. 23, $g_{k,n}(x) \rightarrow 0$ for all x . The theorem of dominated convergence can thus be applied, which proves eq. 24 is true (almost surely), and there is convergence of the eigenfunctions in \mathcal{H}_p .

If K_n is data-dependent but converges, in a way such that the eigen-decomposition of the Gram matrix $(K_n(x_i, x_j))$ converges to the eigen-decomposition of the Gram matrix $(K(x_i, x_j))$, with K the limit of K_n , we want to apply the same reasoning. Because of the convergence of the eigen-decomposition of the Gram matrix, eq. 21 still holds. However, eq. 22 has to be replaced with a limit, because $h_x = K(x, \cdot) \neq K_n(x, \cdot)$. This limit still allows to write eq. 23 (possibly with a different form of probabilistic convergence, depending on the convergence of K_n to K), and the same result is obtained. \square

Discussion

Kernel PCA has already been shown to be a stable and convergent algorithm (Shawe-Taylor, Cristianini and Kandola, 2002; Shawe-Taylor and Williams, 2003; Zwald, Bousquet and Blanchard, 2004). These papers characterize the rate of convergence of the projection error on the subspace spanned by the first m eigenvectors of the feature space covariance matrix. When we perform the PCA or kernel PCA projection on an out-of-sample point, we are taking advantage of the above convergence and stability properties: we trust that a principal eigenvector of the empirical covariance matrix estimates well a corresponding eigenvector of the true covariance matrix. Another justification for applying the Nyström formula outside of the training examples is therefore, as already noted earlier and in (Williams and Seeger, 2000), *in the case where K_n is positive semi-definite*, that it corresponds to the kernel PCA projection (on a corresponding eigenvector of the feature space covariance matrix C).

Clearly, we therefore have with the Nyström formula a method to *generalize spectral embedding algorithms to out-of-sample examples*, whereas the original spectral embedding methods only provide the transformed coordinates of training points (i.e. an embedding of the training points). The experiments described in section 5 show empirically the good generalization of this out-of-sample embedding. Note however that it is not

always clear whether the assumptions needed to apply Proposition 2 are verified or not (especially because of the data-dependency of K_n). This proposition mainly gives an intuition of what a spectral embedding technique is doing (estimating eigenfunctions of a linear operator) in the case of ideal convergence.

An interesting justification for estimating the eigenfunctions of G has been shown in (Williams and Seeger, 2000). When an unknown function f is to be estimated with an approximation g that is a finite linear combination of basis functions, if f is assumed to come from a zero-mean Gaussian process prior with covariance $E_f[f(x)f(y)] = K(x, y)$, then the best choices of basis functions, in terms of expected squared error, are (up to rotation/scaling) the leading eigenfunctions of the linear operator G as defined above.

4 Learning Criterion for the Leading Eigenfunctions

Using an expansion into orthonormal bases (e.g. generalized Fourier decomposition in the case where p is continuous), the best approximation of $K(x, y)$ (in the sense of minimizing expected squared error) using only m terms is the expansion that uses the *first* m eigenfunctions of G (in the order of decreasing eigenvalues):

$$\sum_{k=1}^m \lambda_k f_k(x) f_k(y) \approx K(x, y).$$

This simple observation allows us to define a *loss criterion for spectral embedding algorithms*, something that was lacking up to now for such algorithms. The limit of this loss converges toward an expected loss whose minimization gives rise to the eigenfunctions of G . One could thus conceivably estimate this generalization error using the average of the loss on a test set. That criterion is simply the kernel reconstruction error

$$L_{K_n}(x_i, x_j) = \left(K_n(x_i, x_j) - \sum_{k=1}^m \lambda_{k,n} f_{k,n}(x_i) f_{k,n}(x_j) \right)^2.$$

Proposition 3 *The spectral embedding for a continuous kernel K with discrete spectrum is the solution of a sequential minimization problem, iteratively minimizing for $m = 1, 2, \dots$ the expected value of the loss criterion*

$$L_K(x, y) = K(x, y) - \sum_{k=1}^m \lambda_k f_k(x) f_k(y).$$

Firstly, with $\{(f_k, \lambda_k)\}_{k=1}^{m-1}$ already obtained, one can recursively obtain (λ_m, f_m) by minimizing

$$J_m(\lambda', f') = \int \left(K(x, y) - \lambda' f'(x) f'(y) - \sum_{k=1}^{m-1} \lambda_k f_k(x) f_k(y) \right)^2 p(x) p(y) dx dy \quad (25)$$

where by convention we scale f' such that $\int f'(x)^2 p(x) = 1$ (any other scaling can be transferred into λ').

Secondly, if the same hypothesis on K_n as in Proposition 2 are verified, the Monte-Carlo average of the criterion L_{K_n}

$$\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(K_n(x_i, x_j) - \sum_{k=1}^m \lambda_{k,n} f_{k,n}(x_i) f_{k,n}(x_j) \right)^2$$

converges in probability to the asymptotic expectation of L_K .

Proof

We prove the first part of the proposition concerning the sequential minimization of the loss criterion, which follows from classical linear algebra (Strang, 1980; Kreyszig, 1990). We proceed by induction, assuming that we have already obtained f_1, \dots, f_{m-1} orthogonal eigenfunctions in order of decreasing absolute value of λ_i . We want to prove that (λ', f') that minimizes J_m is (λ_m, f_m) .

Setting $\frac{\partial J_m}{\partial \lambda'} = 0$ yields

$$\lambda' = \langle f', K f' \rangle - \sum_{i=1}^{m-1} \lambda_i \int f_i(x) f_i(y) \cdot f'(x) f'(y) p(x) p(y) dx dy \quad (26)$$

In addition, we have

$$\begin{aligned} J_m(\lambda', f') &= \int (K(x, y) - \sum_{i=1}^{m-1} \lambda_i f_i(x) f_i(y))^2 p(x) p(y) dx dy \\ &- 2 \int \lambda' f'(x) f'(y) \left(K(x, y) - \sum_{i=1}^{m-1} \lambda_i f_i(x) f_i(y) \right) p(x) p(y) dx dy \\ &+ \int (\lambda' f'(x) f'(y))^2 p(x) p(y) dx dy \end{aligned}$$

which gives $J_m(\lambda', f') = C_m - \lambda'^2$ where C_m is a constant, so that λ'^2 should be maximized in order to minimize J_m .

In order to find the minimum of J_m with respect to f' , we write that its Gâteaux differential with any increment h (the limit of $(J_m(\lambda', f' + \alpha h) - J_m(\lambda', f'))/\alpha$ when $\alpha \rightarrow 0$) is 0, which yields the equation:

$$\int K(z, y) f'(y) p(y) dy = \int \lambda' f'(z) f'(y)^2 p(y) dy + \sum_{i=1}^{m-1} \lambda_i \int f_i(z) f_i(y) f'(y) p(y) dy.$$

Using the constraint $\|f'\|^2 = \langle f', f' \rangle = \int f'(y)^2 p(y) dy = 1$, we obtain:

$$(K f')(z) = \lambda' f'(z) + \sum_{i=1}^{m-1} \lambda_i \int f_i(z) f_i(y) f'(y) p(y) dy \quad (27)$$

which rewrites into $Kf' = \lambda' f' + \sum_{i=1}^{m-1} \lambda_i f_i \langle f', f_i \rangle$. Writing Kf' in the basis of all the eigenfunctions, $Kf' = \sum_{i=1}^{\infty} \lambda_i f_i \langle f', f_i \rangle$, we obtain

$$\lambda' f' = \lambda_m f_m \langle f', f_m \rangle + \sum_{i=m+1}^{\infty} \lambda_i f_i \langle f', f_i \rangle.$$

Since the f_i are orthogonal, take the norm and apply Parseval's theorem:

$$\lambda'^2 = \lambda_m^2 \langle f', f_m \rangle^2 + \sum_{i=m+1}^{\infty} \lambda_i^2 \langle f', f_i \rangle^2.$$

If the eigenvalues are distinct, we have $\lambda_m > \lambda_i$ for $i > m$, and the last expression is maximized when $\langle f', f_m \rangle = 1$ and $\langle f', f_i \rangle = 0$ for $i > m$, which proves that $f' = f_m$ is in fact the m -th eigenfunction of the kernel K and thereby $\lambda' = \lambda_m$.

If the eigenvalues are not distinct, then the result can be generalized in the sense that the choice of eigenfunctions is not anymore unique, and the eigenfunctions sharing the same eigenvalue form an orthogonal basis for a subspace. This concludes the proof of the first statement.

To prove the second part (convergence statement), we want to show that the difference between the average cost and the expected asymptotic cost tends toward 0. If we write $\widehat{K}_n(x, y) = \sum_{k=1}^m \lambda_{k,n} f_{k,n}(x) f_{k,n}(y)$ and $\widehat{K}(x, y) = \sum_{k=1}^m \lambda_k f_k(x) f_k(y)$, that difference is

$$\begin{aligned} & \left| \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(K_n(x_i, x_j) - \widehat{K}_n(x_i, x_j) \right)^2 - E_{x,y} \left[\left(K(x, y) - \widehat{K}(x, y) \right)^2 \right] \right| \\ \leq & \left| \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(K(x_i, x_j) - \widehat{K}(x_i, x_j) \right)^2 - E_{x,y} \left[\left(K(x, y) - \widehat{K}(x, y) \right)^2 \right] \right| \\ + & \left| \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(K_n(x_i, x_j) - \widehat{K}_n(x_i, x_j) - K(x_i, x_j) + \widehat{K}(x_i, x_j) \right) \right. \\ & \left. \left(K_n(x_i, x_j) - \widehat{K}_n(x_i, x_j) + K(x_i, x_j) - \widehat{K}(x_i, x_j) \right) \right|. \end{aligned}$$

The kernel and the eigenfunctions being bounded (as shown in the proof of Proposition 2), the second factor in the product (in the second term of the inequality) is bounded by a constant B with probability 1 (because of the $\lambda_{k,n}$ converging almost surely).

Thus, we have with probability 1:

$$\begin{aligned} & \left| \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(K_n(x_i, x_j) - \widehat{K}_n(x_i, x_j) \right)^2 - E_{x,y} \left[\left(K(x, y) - \widehat{K}(x, y) \right)^2 \right] \right| \\ \leq & \left| \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(K(x_i, x_j) - \widehat{K}(x_i, x_j) \right)^2 - E_{x,y} \left[\left(K(x, y) - \widehat{K}(x, y) \right)^2 \right] \right| \end{aligned}$$

$$\begin{aligned}
& + \frac{B}{n^2} \sum_{i=1}^n \sum_{j=1}^n |K_n(x_i, x_j) - K(x_i, x_j)| \\
& + \left| \frac{B}{n^2} \sum_{i=1}^n \sum_{j=1}^n \widehat{K}_n(x_i, x_j) - \widehat{K}(x_i, x_j) \right|
\end{aligned}$$

From the law of large numbers, the first term converges to 0 in probability. The second term converges to 0 because of the uniform convergence of K_n . Now, if we follow the same reasoning as in the proof of Proposition 2 but with $\mathcal{F} = \{h_x, h_y\}$, we get (assuming all eigenvalues $(\lambda_k)_{1 \leq k \leq m}$ are of multiplicity 1)

$$\int (\lambda_{k,n} f_{k,n}(x) f_{k,n}(y) - \lambda_k f_k(x) f_k(y)) p(x) p(y) dx dy \rightarrow 0$$

which, combined with the central limit theorem, shows that the third term also converges to 0 (in probability) as $n \rightarrow +\infty$. We have therefore proved the convergence in probability of the average loss to its asymptotic expectation. \square

Discussion

Note that the empirical criterion is indifferent to the value of the solutions $f_{k,n}$ outside of the training set. Therefore, although the Nyström formula gives a possible solution to the empirical criterion, there are other solutions. Remember that the task we consider is that of estimating the eigenfunctions of G , i.e. approximating a similarity function K where it matters according to the unknown density p . Solutions other than the Nyström formula might also converge to the eigenfunctions of G . For example one could use a non-parametric estimator (such as a neural network) to estimate the eigenfunctions. Even if such a solution does not yield the exact eigenvectors on the training examples (i.e. does not yield the lowest possible error on the training set), it might still be a good solution in terms of generalization, in the sense of good approximation of the eigenfunctions of G . It would be interesting to investigate whether the Nyström formula achieves the fastest possible rate of convergence to the eigenfunctions of G .

5 Experiments

The theory laid down in the previous sections allows us to unify a number of spectral manifold learning algorithms, and it gives a strong sense that it is possible to transform them into function induction algorithms, with a notion of convergence as $n \rightarrow \infty$ and a notion of generalization error.

Here we want to test one aspect of these theoretical results: does the function induction achieved with the Nyström formula work well? we would like to know if the embedding that it predicts on a new point x is close to the embedding that would have been obtained on x if it had been in the training set. However, how do we evaluate the “error” thus obtained? Our idea is to compare it to the variations in embedding that result from

small perturbations of the training set (such as replacing a subset of the examples by others from the same distribution).

For this purpose we consider splits of the data in three sets, $D = F \cup R_1 \cup R_2$ and training either with $F \cup R_1$ or $F \cup R_2$, comparing the embeddings on F . For each algorithm described in section 2, we apply the following procedure:

1. We choose $F \subset D$ with $m = |F|$ samples. The remaining $n - m$ samples in D/F are split into two equal size subsets R_1 and R_2 . We train (obtain the eigenvectors) over $F \cup R_1$ and $F \cup R_2$ and we calculate the Euclidean distance between the aligned embeddings obtained for each $x_i \in F$. When eigenvalues are close, the estimated eigenvectors are unstable and can rotate in the subspace they span. Thus we estimate an alignment (by linear regression) between the two embeddings using the points in F .
2. For each sample $x_i \in F$, we also train over $\{F \cup R_1\} / \{x_i\}$. We apply the Nyström formula to out-of-sample points to find the predicted embedding of x_i and calculate the Euclidean distance between this embedding and the one obtained when training with $F \cup R_1$, i.e. with x_i in the training set (in this case no alignment is done since the influence of adding a single point is very limited).
3. We calculate the mean **difference** δ (and its standard error) between the distance obtained in step 1 and the one obtained in step 2 for each sample $x_i \in F$, and we repeat this experiment for various sizes of F .

The results obtained for MDS, Isomap, spectral clustering and LLE are shown in Figure 2 for different values of $|R_1|/n$ (i.e the fraction of points exchanged). The vertical axis is δ , the difference between perturbation error and induction error. The horizontal zero line corresponds to no difference between the embedding error due to induction (vs transduction) and the embedding error due to training set perturbation. For values of δ above the zero line, the embedding error due to perturbation is greater than the embedding error due to out-of-sample prediction. Clearly, the in-sample (transduction) vs out-of-sample (induction) difference is of the same order of magnitude as the change in embedding due to exchanging a small fraction of the data (1 to 5%).

Experiments are done over a database of 698 synthetic face images described by 4096 components that is available at <http://isomap.stanford.edu>. Similar results have been obtained over other databases such as Ionosphere⁴ and Swissroll⁵. Each algorithm generates a two-dimensional embedding of the images, following the experiments reported for Isomap. The number of neighbors is 10 for Isomap and LLE, and a Gaussian kernel with a standard deviation of 0.01 is used for spectral clustering / Laplacian eigenmaps. 95% confidence intervals are drawn beside each mean difference of error on the figure.

As expected, the mean difference between the two distances is almost monotonically increasing as the number $|R_1|$ of substituted training samples grows, mostly because

⁴<http://www.ics.uci.edu/~mllearn/MLSummary.html>

⁵<http://www.cs.toronto.edu/~roweis/lle/>

the training set embedding variability increases. We find in most cases that the out-of-sample error is less than or comparable to the training set embedding instability when around 2% of the training examples are substituted randomly.

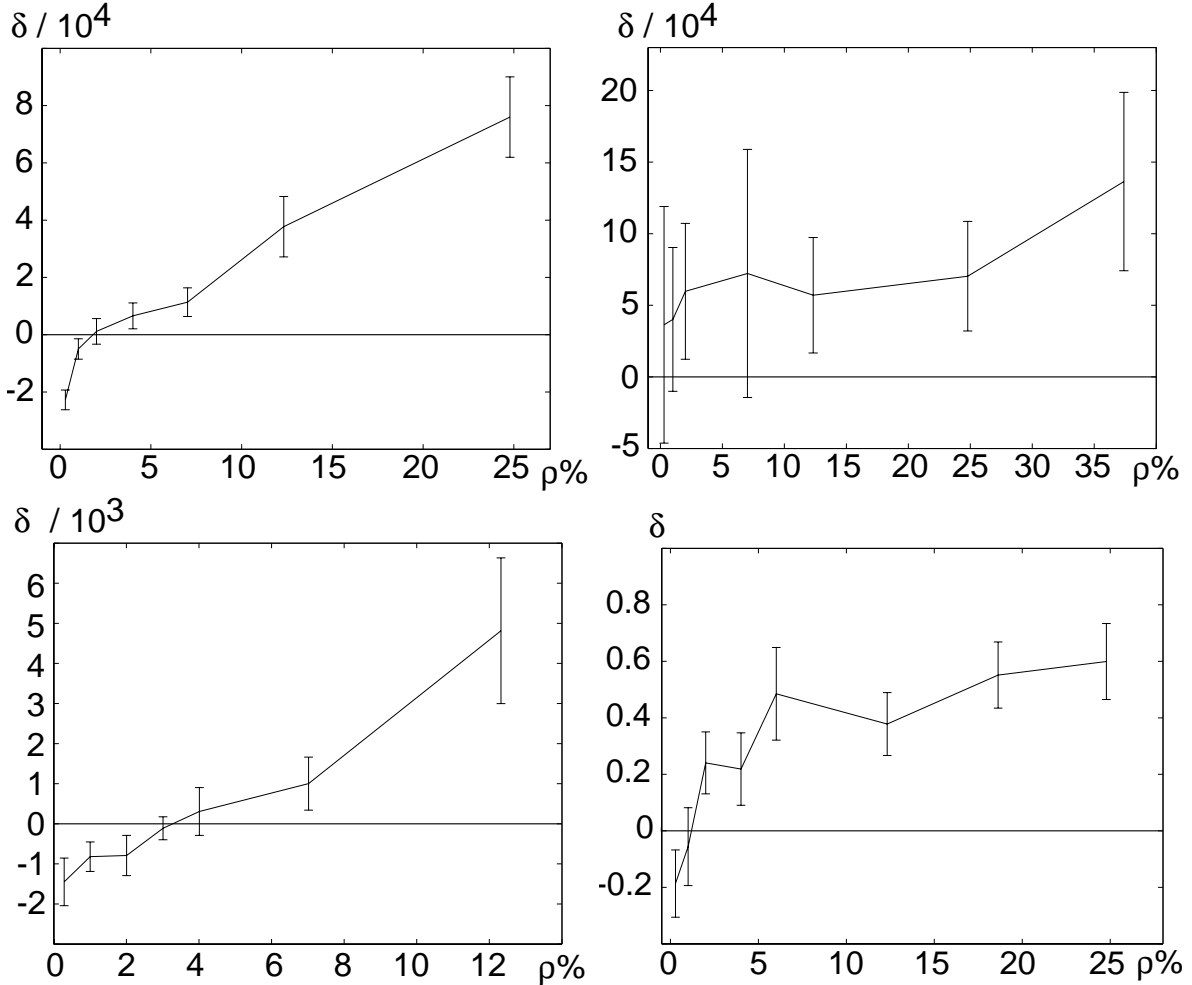


Figure 2: δ (training set variability minus out-of-sample error), w.r.t. ρ (proportion of substituted training samples) on the “Faces” dataset ($n = 698$), obtained with a two-dimensional embedding. Top left: MDS. Top right: spectral clustering or Laplacian eigenmaps. Bottom left: Isomap. Bottom right: LLE. Error bars are 95% confidence intervals. Exchanging about 2% of the training examples has an effect comparable to using the Nyström formula.

6 Conclusion

Manifold learning and dimensionality reduction are powerful machine learning tools for which much progress has been achieved in recent years. This paper sheds light on a

family of such algorithms, involving spectral embedding, which are all based on the eigen-decomposition of a similarity matrix.

Spectral embedding algorithms such as spectral clustering, Isomap, LLE, metric MDS, and Laplacian Eigenmaps are very interesting dimensionality reduction or clustering methods. However, they lacked up to now a notion of generalization that would allow to easily extend the embedding out-of-sample without again solving an eigensystem. This paper has shown with various arguments that the well known Nyström formula can be used for this purpose, and that it thus represents the result of a *function induction* process. These arguments also help us to understand that *these methods do essentially the same thing, but with respect to different kernels: they estimate the eigenfunctions of a linear operator* associated with a kernel and with the underlying distribution of the data. This analysis also shows that these methods are *minimizing an empirical loss*, and that the solutions toward which they converge are the minimizers of a corresponding expected loss, which thus defines what good generalization should mean, for these methods. It shows that these unsupervised learning algorithms can be extended into function induction algorithms. The Nyström formula is a possible extension but it does not exclude other extensions which might be better or worse estimators of the eigenfunctions of the asymptotic linear operator G . When the kernels are positive semi-definite, these methods can also be immediately seen as performing kernel PCA. Note that Isomap generally yields a Gram matrix with negative eigenvalues, and users of MDS, spectral clustering or Laplacian eigenmaps may want to use a kernel that is not guaranteed to be positive semi-definite. The analysis in this paper can still be applied in that case, even though the kernel PCA analogy does not hold anymore. This is important to note because recent work (Laub and Müller, 2003) has shown that the coordinates corresponding to large negative eigenvalues can carry very significant semantics about the underlying objects. In fact, it is proposed in (Laub and Müller, 2003) to perform dimensionality reduction by projecting on the eigenvectors corresponding to the largest eigenvalues **in magnitude** (i.e. irrespective of sign).

In this paper we have provided theorems that provide justification for the Nyström formula in the general case of data-dependent kernels which may not be positive-definite. However, these theorems rely on strong assumptions which may not hold for particular spectral manifold learning algorithms. To help assess the practical validity of the Nyström formula for predicting the embedding of out-of-sample points, we have performed a series of comparative experiments.

The experiments performed here have shown empirically on several data sets that the predicted out-of-sample embedding is generally not far from the one that would be obtained by including the test point in the training set, and that the difference is of the same order as the effect of small perturbations of the training set.

An interesting parallel can be drawn between the spectral embedding algorithms and the view of PCA as finding the principal eigenvectors of a matrix obtained from the data. The present paper parallels for spectral embedding the view of PCA as an estimator of the principal directions of the covariance matrix of the underlying unknown distribution, thus introducing a convenient notion of generalization, relating to an unknown

distribution.

Finally, a better understanding of these methods opens the door to new and potentially much more powerful unsupervised learning algorithms. Several directions remain to be explored:

1. Using a smoother distribution than the empirical distribution to define the linear operator G_n . Intuitively, a distribution that is closer to the true underlying distribution would have a greater chance of yielding better generalization, in the sense of better estimating eigenfunctions of G . This relates to putting priors on certain parameters of the density, e.g. as in (Rosales and Frey, 2003).
2. All of these methods are capturing salient features of the unknown underlying density. Can one use the representation learned through the estimated eigenfunctions in order to construct a good density estimator? Looking at Figure 1 suggests that modeling the density in the transformed space (right hand side) should be much easier (e.g. would require fewer Gaussians in a Gaussian mixture) than in the original space.
3. These transformations discover abstract structures such as clusters and manifolds. It might be possible to learn even more abstract (and less local) structures, starting from these representations. Ultimately, the goal would be to learn higher-level abstractions on top of lower-level abstractions by iterating the unsupervised learning process in multiple “layers”.

Looking for extensions such as these is important because all of the manifold learning algorithms studied here suffer from the following fundamental weakness: they are using mostly the neighbors around each example to capture the local structure of the manifold, i.e. the manifold is seen as a combination of linear patches around each training example. This is very clear in LLE and Isomap, which have a simple geometric interpretation, and it is also clear in non-spectral methods such as Manifold Parzen Windows (Vincent and Bengio, 2003) and other mixtures of factor analyzers (Ghahramani and Hinton, 1996). In low dimension or when the manifold is smooth enough, there may be enough examples locally to characterize the plane tangent to the manifold. However, when the manifold has high curvature with respect to the amount of training data (which can easily be the case, especially with high-dimensional data), it is hopeless to try to capture the local tangent directions based only on local information. Clearly, this is the topic for future work, addressing a fundamental question about generalization in high-dimensional data, and for which the traditional non-parametric approaches may be insufficient.

Acknowledgments

The authors would like to thank Léon Bottou, Christian Léger, Sam Roweis, Yann Le Cun, and Yves Grandvalet for helpful discussions, and the following funding organizations: NSERC, MITACS, IRIS, and the Canada Research Chairs.

References

- Baker, C. (1977). *The numerical treatment of integral equations*. Clarendon Press, Oxford.
- Belkin, M. and Niyogi, P. (2003a). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396.
- Belkin, M. and Niyogi, P. (2003b). Using manifold structure for partially labeled classification. In Becker, S., Thrun, S., and Obermayer, K., editors, *Advances in Neural Information Processing Systems 15*, Cambridge, MA. MIT Press.
- Bengio, Y., Paiement, J., Vincent, P., Delalleau, O., Le Roux, N., and Ouimet, M. (2004). Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. In Thrun, S., Saul, L., and Schölkopf, B., editors, *Advances in Neural Information Processing Systems 16*. MIT Press.
- Brand, M. (2003). Charting a manifold. In Becker, S., Thrun, S., and Obermayer, K., editors, *Advances in Neural Information Processing Systems*, volume 15. MIT Press.
- Chung, F. (1997). Spectral graph theory. In *CBMS Regional Conference Series*, volume 92. American Mathematical Society.
- Cox, T. and Cox, M. (1994). *Multidimensional Scaling*. Chapman & Hall, London.
- de Silva, V. and Tenenbaum, J. (2003). Global versus local methods in nonlinear dimensionality reduction. In Becker, S., Thrun, S., and Obermayer, K., editors, *Advances in Neural Information Processing Systems 15*, pages 705–712, Cambridge, MA. MIT Press.
- Donoho, D. and Grimes, C. (2003). Hessian eigenmaps: new locally linear embedding techniques for high-dimensional data. Technical Report 2003-08, Dept. Statistics, Stanford University.
- Ghahramani, Z. and Hinton, G. (1996). The EM algorithm for mixtures of factor analyzers. Technical Report CRG-TR-96-1, Dpt. of Comp. Sci., Univ. of Toronto.
- Gower, J. (1968). Adding a point to vector diagrams in multivariate analysis. *Biometrika*, 55(3):582–585.
- Ham, J., Lee, D., Mika, S., and Schölkopf, B. (2003). A kernel view of the dimensionality reduction of manifolds. Technical Report TR-110, Max Planck Institute for Biological Cybernetics, Germany.
- Hastie, T. and Stuetzle, W. (1989). Principal curves. *Journal of the American Statistical Association*, 84:502–516.

- Kégl, B. (2003). Intrinsic dimension estimation using packing numbers. In S. Becker, S. T. and Obermayer, K., editors, *Advances in Neural Information Processing Systems 15*, pages 681–688. MIT Press, Cambridge, MA.
- Kégl, B. and Krzyzak, A. (2002). Piecewise linear skeletonization using principal curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(1):59–74.
- Kohonen, T. (1990). The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480.
- Koltchinskii, V. (1998). Asymptotics of spectral projections of some random matrices approximating integral operators. In Eberlein, Hahn, and Talagrand, editors, *Progress in Probability*, volume 43, pages 191–227, Basel. Birkhauser.
- Koltchinskii, V. and Giné, E. (2000). Random matrix approximation of spectra of integral operators. *Bernoulli*, 6(1):113–167.
- Kreyszig, E. (1990). *Introductory Functional Analysis with Applications*. John Wiley & Sons, Inc., New York, NY.
- Kumar, V., Grama, A., Gupta, A., and Karypis, G. (1994). *Introduction to Parallel Computing: Design and Analysis of Algorithms*. Benjamin Cummings, Redwood City, CA.
- Laub, J. and Müller, K.-R. (2003). Feature discovery: unraveling hidden structure in non-metric pairwise data. Technical report, Fraunhofer FIRST.IDA, Germany.
- Ng, A. Y., Jordan, M. I., and Weiss, Y. (2002). On spectral clustering: analysis and an algorithm. In Dietterich, T. G., Becker, S., and Ghahramani, Z., editors, *Advances in Neural Information Processing Systems 14*, Cambridge, MA. MIT Press.
- Rosales, R. and Frey, B. (2003). Learning generative models of affinity matrices. In *Proceedings of the 19th Annual Conference on Uncertainty in Artificial Intelligence (UAI-03)*, pages 485–492, San Francisco, CA. Morgan Kaufmann Publishers.
- Roweis, S. and Saul, L. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326.
- Rumelhart, D., Hinton, G., and Williams, R. (1986). Learning internal representations by error propagation. In Rumelhart, D. and McClelland, J., editors, *Parallel Distributed Processing*, volume 1, chapter 8, pages 318–362. MIT Press, Cambridge.
- Saul, L. and Roweis, S. (2002). Think globally, fit locally: unsupervised learning of low dimensional manifolds. *Journal of Machine Learning Research*, 4:119–155.
- Saund, E. (1989). Dimensionality-reduction using connectionist networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(3):304–314.

- Schölkopf, B., Burges, C. J. C., and Smola, A. J. (1999). *Advances in Kernel Methods — Support Vector Learning*. MIT Press, Cambridge, MA.
- Schölkopf, B., Smola, A., and Müller, K.-R. (1996). Nonlinear component analysis as a kernel eigenvalue problem. Technical Report 44, Max Planck Institute for Biological Cybernetics, Tübingen, Germany.
- Schölkopf, B., Smola, A., and Müller, K.-R. (1998). Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10:1299–1319.
- Shawe-Taylor, J., Cristianini, N., and Kandola, J. (2002). On the concentration of spectral properties. In Dietterich, T., Becker, S., and Ghahramani, Z., editors, *Advances in Neural Information Processing Systems 14*. MIT Press.
- Shawe-Taylor, J. and Williams, C. (2003). The stability of kernel principal components analysis and its relation to the process eigenspectrum. In Becker, S., Thrun, S., and Obermayer, K., editors, *Advances in Neural Information Processing Systems 15*. MIT Press.
- Shi, J. and Malik, J. (1997). Normalized cuts and image segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 731–737.
- Spielman, D. and Teng, S. (1996). Spectral partitioning works: planar graphs and finite element meshes. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*.
- Strang, G. (1980). *Linear Algebra and Its Applications*. Academic Press, New York.
- Teh, Y. W. and Roweis, S. (2003). Automatic alignment of local representations. In Becker, S., Thrun, S., and Obermayer, K., editors, *Advances in Neural Information Processing Systems*, volume 15. MIT Press.
- Tenenbaum, J., de Silva, V., and Langford, J. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323.
- van der Vaart, A. and Wellner, J. (1996). *Weak Convergence and Empirical Processes with applications to Statistics*. Springer, New York.
- Verbeek, J. J., Roweis, S. T., and Vlassis, N. (2004). Non-linear cca and pca by alignment of local models. In Thrun, S., Saul, L., and Schölkopf, B., editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA.
- Vincent, P. and Bengio, Y. (2003). Manifold parzen windows. In Becker, S., Thrun, S., and Obermayer, K., editors, *Advances in Neural Information Processing Systems 15*, Cambridge, MA. MIT Press.
- Vlachos, M., Domeniconi, C., Gunopulos, D., Kollios, G., and Koudas, N. (2002). Non-linear dimensionality reduction techniques for classification and visualization. In *Proc. of 8th SIGKDD*, Edmonton, Canada.

- Weiss, Y. (1999). Segmentation using eigenvectors: a unifying view. In *Proceedings IEEE International Conference on Computer Vision*, pages 975–982.
- Williams, C. (2001). On a connection between kernel pca and metric multidimensional scaling. In Leen, T., Dietterich, T., and Tresp, V., editors, *Advances in Neural Information Processing Systems 13*, pages 675–681. MIT Press.
- Williams, C. and Seeger, M. (2000). The effect of the input density distribution on kernel-based classifiers. In *Proceedings of the Seventeenth International Conference on Machine Learning*. Morgan Kaufmann.
- Williams, C. K. I. and Seeger, M. (2001). Using the Nyström method to speed up kernel machines. In Leen, T., Dietterich, T., and Tresp, V., editors, *Advances in Neural Information Processing Systems 13*, pages 682–688, Cambridge, MA. MIT Press.
- Zwald, L., Bousquet, O., and Blanchard, G. (2004). Statistical properties of kernel principal components analysis. Technical report, submitted.