

# DD-HDS: A Method for Visualization and Exploration of High-Dimensional Data

Sylvain Lespinats, Michel Verleysen, *Senior Member, IEEE*, Alain Giron, and Bernard Fertil

**Abstract**—Mapping high-dimensional data in a low-dimensional space, for example, for visualization, is a problem of increasingly major concern in data analysis. This paper presents data-driven high-dimensional scaling (DD-HDS), a nonlinear mapping method that follows the line of multidimensional scaling (MDS) approach, based on the preservation of distances between pairs of data. It improves the performance of existing competitors with respect to the representation of high-dimensional data, in two ways. It introduces 1) a specific weighting of distances between data taking into account the concentration of measure phenomenon and 2) a symmetric handling of short distances in the original and output spaces, avoiding false neighbor representations while still allowing some necessary tears in the original distribution. More precisely, the weighting is set according to the effective distribution of distances in the data set, with the exception of a single user-defined parameter setting the tradeoff between local neighborhood preservation and global mapping. The optimization of the stress criterion designed for the mapping is realized by “force-directed placement” (FDP). The mappings of low- and high-dimensional data sets are presented as illustrations of the features and advantages of the proposed algorithm. The weighting function specific to high-dimensional data and the symmetric handling of short distances can be easily incorporated in most distance preservation-based nonlinear dimensionality reduction methods.

**Index Terms**—High-dimensional data, multidimensional scaling (MDS), neighborhood visualization, nonlinear mapping.

## I. INTRODUCTION

**V**ISUALIZATION of high-dimensional data is intended to facilitate the understanding of data sets by preserving some “essential” information. It generally requires the mapping of the data into a low- (usually 2- or 3-) dimensional space. However, high-dimensional data raise unusual problems of analysis, given that some properties of the spaces they live in cannot be extrapolated from our common experience. In particular (notably in the case of Euclidian spaces), we often face the problems of empty space and concentration of measure: When

the number of dimensions is high, the neighborhood of each object is scarcely filled whereas most of the other objects are found in a thin outer shell. Distances between high-dimensional objects are usually very concentrated around their average [1].

Exploration and analysis of high-dimensional data are often made by means of dimension reduction techniques [2], [3]. Since human experience mostly deals with 3-D space (and most data display devices are 2-D), finding a meaningful mapping of high-dimensional data into such low-dimensional spaces is a major issue. Often linear mapping methods do not lead to satisfactory representations. Indeed real data most often show nonlinear relationships that cannot be approximated in a satisfactory way by linear methods. Nonlinear mappings (also called nonlinear methods for dimensionality reduction) offer more flexibility, often at the price of an additional complexity.

In this paper, we propose a nonlinear dimensionality reduction method specifically adapted to high-dimensional data. It follows the line of multidimensional scaling (MDS) methods, based on the preservation of distances between pairs of data [4]. However, it differs from existing methods in two ways. First, it includes a weighting of distances that takes the concentration of measure phenomenon into account (see Section IV-A); this is of primary importance when dealing with high-dimensional data, for which the concept of “small” and “large” distances strongly differs from the traditional view in low-dimensional spaces. Second, existing nonlinear dimensionality reduction methods either favor the preservation of small distances in the original space ([5], for example), at the risk of collapsing far points together in the representation, or favor the preservation of small distances in the output space ([6], for example), allowing sometimes unwanted tears in the original distribution. The method proposed in this paper is symmetric with respect to distances in the original and output spaces (see Section IV-B), leading to better and more intuitive representations, as attested by experiences. Finally, the optimization of the method-specific objective function is performed by force-directed placement (FDP), as an alternative to more traditional gradient-based algorithms (see Section IV-C).

This paper is organized as follows. Section II shows known phenomena occurring in high-dimensional spaces, and how metric transformations can deal with them. Section III briefly reviews dimensionality reduction methods, and highlights the difficulties encountered in these methods with high-dimensional data. Section IV presents our original nonlinear mapping algorithm called data-driven high-dimensional scaling (DD-HDS). Section V defines efficiency measures used in Section VI that presents experimental results and comparisons with existing methods.

Manuscript received December 28, 2005; revised July 12, 2006 and October 5, 2006; accepted December 3, 2006. This work was supported in part by the Action INTER EPST Bio-informatique under 2001 Grant 120910.

S. Lespinats, A. Giron, and B. Fertil are with the UMR INSERM unité 678-Université Pierre et Marie Curie—Paris 6, 75634 Paris, France (e-mail: lespinat@imed.jussieu.fr; giron@imed.jussieu.fr; fertil@imed.jussieu.fr).

M. Verleysen is with the Université catholique de Louvain, B-1348 Louvain-la Neuve, Belgium and with the SAMOS-MATISSE, Université Paris 1 Panthéon-Sorbonne, 75634 Paris Cedex 13, France (e-mail: Verleysen@dice.ucl.ac.be).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNN.2007.891682




Fig. 1. Histogram of distances between uniformly distributed data in a unit cube, according to space dimension. Histograms for dimensions larger than 200 would have the same Gaussian-like shape, but their centers would be shifted to the right proportionally to the square root of the dimension.

## II. DISTANCES IN HIGH-DIMENSIONAL SPACE

Mapping methods based on the comparison of distances between pairs of data in the original and output spaces, such as MDS, need to take care of how distances are measured. Specific distances are often used to measure similarities in the original space; distances in the (low-dimensional) output space are usually measured in a more conventional way. In this paper, we focus on Euclidean and derived metrics (Sections II-A and II-B, respectively) for the representation of data in the original space and on the Euclidean metric for the output space for sake of simplicity. However, most results in this paper may be easily extended to other metrics.

### A. Euclidean Distances in High-Dimensional Space

Several surprising phenomena appear when dealing with high-dimensional data. This fact is known as the “curse of dimensionality” [1], and has strong impact on the validity and performances of data analysis tools [7]. In particular, the Euclidean metric ( $\|x\| = \sqrt{\sum x_i^2}$ ) is known to suffer from the unwanted concentration of measure phenomenon [8]. Let us consider the distances between pairs of uniformly distributed data in an  $n$ -dimensional unit-edge hypercube. It can be shown that the mean of the distances increases with the square root of  $n$ , while the variance remains constant. This is illustrated in Fig. 1. As a result, it is much more difficult to discriminate between *small* and *large* distances in a relative way (e.g., when distances are normalized) in a high-dimensional space. As we will see in a short review in Section III, nonlinear dimensionality reduction tools will thus fail to give more weight to small and/or large distances, as it should be the case for a proper functioning of the algorithms.

The Euclidean norm is not the only one to suffer from the concentration of measure phenomenon. All other Minkowski metrics, Minkowski pseudometrics (fractional norms), Pearson correlation metric, etc., have this characteristic, though at different levels [8]. Transformations of distances may be used to overcome these problems. We briefly review such transformations in Section II-B.

### B. Derived Metrics

Changing data representation to help the mapping procedure and, eventually, to add supervised information is common practice (see [9], for example). Three powerful approaches

are particularly useful as preprocessing when dealing with high-dimensional data.

- 1) Today, kernel methods have a huge attractiveness in data analysis [10]. Kernel methods rely on the principle of first mapping data onto a (usually higher dimensional) space before further processing. In practice, the mapping is not explicitly calculated: a so-called kernel  $k(\cdot, \cdot)$  is used to calculate distances between the mapped data. If the kernel is positive definite, it can be verified that it is indeed the scalar product between data in a transformed space:  $k(x_i, x_j) = \langle \phi(x_i), \phi(x_j) \rangle$  where  $\langle \cdot, \cdot \rangle$  denotes the scalar product and  $\phi(\cdot)$  a mapping to a possibly high-dimensional space. The so-called kernel trick avoids calculating both the mapping and the distances (here scalar products) in the original spaces—only the outputs of kernels have to be evaluated.

By using nonlinear transformations of data, kernel methods succeed in building nonlinear models (e.g., for classification and regression) keeping many advantages of linear tools. Their optimization procedure is also simplified compared to many other nonlinear models.

Kernels are used in the context of nonlinear dimensionality reduction as well. This leads, for example, to the kernel principal component analysis (KPCA) method consisting in applying the linear PCA method in a kernel-induced space [10]–[13]. As kernels are defined as a dot product between nonlinear data transformations, KPCA can be viewed as PCA applied after a transformation of metric. As in MDS methods, the user is faced with the crucial choice of an adequate metric transformation, more precisely, the choice of the kernel [11], [14]. Weinberger *et al.*, for example, use semidefinite programming for manifold learning that provides optimized kernels for high-dimensional data projection [15], [16]. In this context, the high similarity of this approach with fast mixing Markov chains must be pointed out [17]. A variant, described in [10, p. 436], consists in using classical MDS on a distance matrix generated through a kernel function. Taking the exponential of (Euclidean) distances is a possible transformation that enhances the contrast between small and large distances.

- 2) Geodesic metrics (also called curvilinear metrics) may offer interesting data representations when dealing with high-dimensional spaces. They are based on the fact that

if data occupy a nonconvex part of the space, it seems legitimated to force distances to be measured *through* the cloud of data, instead of using a conventional Euclidean distance. Intuitive justifications for the interest of curvilinear distances may be found, for example, in [18] and [19]. The curvilinear distance is measured in a graph the nodes of which are the data themselves (or a reduced set). Close data in the space are linked together to form a connected graph (the way they are connected defines the type of distance; classically, each node is connected to the  $k$  closest nodes). The curvilinear distance between any pair of data  $x_i$  and  $x_j$  is subsequently the sum of Euclidean distances between all pairs of data connected in the shortest path on the graph between  $x_i$  and  $x_j$ . Floyd and Dijkstra's algorithms may be used to compute such distances [20], [21]. Isomap [18] and curvilinear distance analysis (CDA) [19], [22] are methods belonging to this category, the difference resulting in the way distances are weighted: Isomap extends classical MDS [4] while CDA extends the "curvilinear component analysis" [6]. Extensions to these methods have been published, e.g., to allow the possibility for tears in the original distribution, avoiding loops or other closed surfaces to collapse in the mapping (see, for example, [23]).

- 3) A last class of derived metrics discussed here consists in using rank orders. Indeed, on some data (such as genomic signatures [24]), the ranking of neighbors is important, while the distance itself may be less [25]. Switching from distance to rank order may thus reveal interesting data properties. Kruskal's criterion is close to such a transformation [26], [27].

Rank orders are not symmetric: let us note  $X = (x_1, x_2, \dots, x_N)$  the data in the original space; if data  $x_i$  is the  $k$ th neighbor of  $x_j$ ,  $x_j$  is not necessarily the  $k$ th neighbor of  $x_i$ . In order to use the rank as a distance,  $d_{ij}^r$  may be defined as the average of the rank order of  $x_i$  with respect to  $x_j$  and vice-versa. Note that, nevertheless,  $d^r$  must be called a pseudodistance as it does not respect the triangular inequality. Furthermore, this pseudodistance is data set dependent (just as curvilinear distance is), as  $d_{ij}^r$  may be influenced by the addition of a data  $x_k$  ( $k \neq i$  and  $k \neq j$ ) in the data set; to limit this effect, it is possible to normalize the pseudodistances (for example, by their maximum value). By construction, rank orders do not suffer from the concentration of measure phenomenon; however, they are not commonly used in nonlinear dimensionality reduction methods, probably due to the lack of conventional distance properties.

All these derived metrics are designed to enhance the contrast between small and large distances. However, whatever is the method used for that purpose, all the methods fail to address the specific properties of high-dimensional data. The concentration of distances phenomenon detailed in Section II-A causes that all distances are approximately equal. Transforming these distances by a nonlinear function, such a kernel does not help, unless the transformation is designed specifically to take the

concentration phenomenon into account. This is the goal of the algorithm described in Section IV.

### III. DIMENSION REDUCTION TECHNIQUES

Data  $X = (x_1, x_2, \dots, x_N)$  are defined in the vector space  $E_1$  with associated metric  $m_1$ . Our goal is to represent  $X$  in a vector Euclidian space  $E_2$  of lower dimension than  $N$ .

Mapping data from a high-dimensional space to a low-dimensional one, keeping exactly all distances between the pairs of points in the original and output spaces, is most often impossible whatever is the distance used to measure similarities in the original space. Then, a data representation has to release some constraints according to a specific "point of view." Numerous dimensionality reduction methods have been proposed, with variants in the methodology and in the criterion (the point of view) to optimize. In this section, we briefly present the main aspects of these techniques, before providing some insight about the difficulties encountered when the dimension of the original space is high.

#### A. Approaches for Dimension Reduction

Because keeping exactly all distances between pairs of points unchanged in the representation is most often impossible, all methods emphasize the preservation of some distances or types of distances, therefore, privileging a specific point of view. For example, PCA and classical MDS [4], [28], [29] maximize the variance of the data cloud after projection, under linear projection hypothesis; the resulting representation expresses the overall form of the data set. Locally linear embedding (LLE) [30], Laplacian eigenmaps [31], and Hessian-based locally linear embedding (HLLE) [32] assume that data are located on a manifold, smooth enough to be reasonably well approximated by local linear models; these methods unfold the set of data through local linear projections. The merging of local projections may be optimized afterward [33]. Supervised methods such as discriminant analysis and partial least squares (PLS) regression [34] use a dependent variable (discrete or continuous, respectively) to guide the mapping. Self-organizing maps (SOMs) [35], [36] visualize data on a grid obtained by a topology-preserving vector quantization. A visualization-induced SOM (ViSOM) and a probabilistic regularized SOM (PRSOM) merge SOM and MDS algorithms in order to map data based both on topology and distance preservation [37]–[39]. The generative topographic mapping (GTM) is also inspired from SOM [40]: a lower dimensional manifold is optimized to approach data in the original space (see [41] for discrete data). The Gaussian process latent variable model (GP-LVM) results from a novel probabilistic interpretation of PCA [42]. This nonlinear method is close to KPCA and GTM. Methods such as Sammon's mapping [5], nonlinear MDS [26], [27], [43], curvilinear component analysis (CCA) [6], [44], and isotop (a SOM the nodes of which are positioned by CCA) [45], [46] emphasize on local neighborhood preservation, often at the price of allowing huge deformations of the global shape of the data cloud. Note that the CCA acronym used in this paper according to the literature covering this algorithm does not mean

here canonical correlation analysis. As detailed in Section II, derived distances in the original space may be used, entitling, for example, the use of MDS after a kernel transformation, or the replacement of Euclidean distances by curvilinear distances in MDS and CCA, leading to isomap and CDA, respectively.

The method introduced in Section IV is placed in the context of neighborhood preservation. The goal is to build a method emphasizing on the preservation of small distances, possibly at the price of distortions in large distances. The difference with respect to previously mentioned methods arises from the fact that the specific properties of high-dimensional spaces are taken into account when measuring distances in the original space. Furthermore, the question whether to emphasize on small distances in the original or output space is answered in a symmetric way, offering a compromise between the risk of mapping far points together and the possibility to tear the initial distribution for a better representation.

### B. Sammon's Stress and High-Dimensional Data

Distance preservation methods such as Sammon's mapping, CCA, and CDA minimize the differences between  $d_{ij}$ , the distance between  $x_i$  and  $x_j$  in the original space, and  $d'_{ij}$ , the distance between their representations  $x'_i$  and  $x'_j$  in the output space. Small distances are emphasized in order to preserve local topology (and small distances). An objective criterion  $\zeta = \sum_{i<j} F(|d_{ij} - d'_{ij}|, k(d_{ij}))$  similar to the original Sammon's stress, where  $k(d_{ij})$  is a monotonically decreasing function giving more weight to small distances in the  $F$  criterion is generally used. Sammon's stress [5] is, for example

$$\zeta_{\text{Sammon}} = \frac{1}{\sum_{i<j} d_{ij}} \sum_{i<j} \frac{(d_{ij} - d'_{ij})^2}{d_{ij}} \quad (1)$$

and Kruskal criterion [27] is

$$\zeta_{\text{Kruskal}} = \sqrt{\sum_{i<j} \frac{(d_{ij} - d'_{ij})^2}{d_{ij}^2}}. \quad (2)$$

Stress functions as (1) and (2) are also called error functions, energy functions, or loss functions, depending on the literature.

In both cases, the weighting function is related to the inverse of the distance. In high-dimensional spaces, however, as detailed in Section II-A, all distances tend to be similar. The weighting factor in (1) and (2) does not play its role anymore. Then, the criteria give similar weights to small and large distances, leading to a mapping that mixes global representation and local neighborhood preservations. Such a poor behavior of Sammon's stress was already mentioned by its author, who noted that linearly separable classes in high-dimensional spaces might not be separable in the mapping [5].

### C. Representation With False Neighborhoods or Tears

Distance preservation methods penalize more heavily mismatches in small distances, by weighting the stress criterion with a decreasing function of either the distances in the original space [see (1), for example], or the distances in the output space [see (2), for example]. However, in the first case, it is difficult to tear distributions with loops, leading to the so-called

"false neighborhood" representation: Data far from each other in the original space could be mapped to close points, exactly as PCA "flattens" volumes when the number of principal components used in the projection is not sufficient. As an example, the extreme points of a "C"-shape (2-D space) with two long branches (with respect to their inter-distance) will be projected as neighbors (in a 1-D space), although their distance in the original space is large (compared to distances between neighboring points in the "C"-shape); a good mapping procedure would unroll the "C"-shape instead of flattening it.

In the second case, tears are allowed, with the risk that neighbor points in the original space may be found widely separated in the output space; tears are sometimes necessary, but may lead to wrong interpretations of neighborhoods in the output space. For example, mapping an "O"-shape requires tears to avoid flattening and false neighbors, but the location of tears appears randomly in some methods, or depends more strongly on the density of data in the original space than on the manifold geometry.

False neighbors and tears are limitations of nonlinear dimensionality reduction methods, which cannot be avoided in most cases due to the intrinsic nature of the manifold. Illustrations of tears and false neighbors are provided in Section VI-B, where mappings of open boxes with various algorithms are shown.

However, in most cases, there is no reason to favor *a priori* false neighbors or tears. See [47] for a comprehensive discussion of this problem in terms of similarities and dissimilarities between data. Thus, there is a need for a method that implements a compromise and reduces the risk of both false neighbors and tears.

## IV. DD-HDS: PRINCIPLES RULING THE METHOD

In this section, a new nonlinear dimensionality reduction method is proposed, which addresses the two shortcomings detailed previously: the method implements a weighting of distances specifically adapted to high-dimensional data, and avoids too high risks of both tears and false neighbors through a symmetric weighting approach.

### A. Sigmoid Balancing Function

Based on Fig. 1, it seems obvious that a discrimination between *small* and *large* distances cannot be achieved in high-dimensional spaces by a simple inverse of distances (or inverse of squares of distances) as in Sammon-like (or Kruskal-like) criteria. Thus, we suggest using a weighting function of the form

$$k(d_{ij}) = 1 - \int_{-\infty}^{d_{ij}} f(u, \mu, \sigma) du \quad (3)$$

where  $f(u, \mu, \sigma)$  is the probability density function of a Gaussian variable with mean  $\mu$  and standard deviation (std)  $\sigma$ . An example of such function is shown in Fig. 2, where it can be seen that small distances in their *effective* range will contribute to the objective stress, while large ones will not. Sigmoid-like weighting functions were already proposed by Demartines [6]. Here, the shape of the weighting function is adapted to the effective distribution of the data in the high-dimensional space. Note that using a cumulative Gaussian function as weighting does not assume, in practice, that the distribution of distances is




Fig. 2. Weighting function as implemented by (3) fits the distributions of distances in  $n$ -dimensional spaces (see Fig. 1; histograms from left to right are drawn with  $n = 1, 10, 50$ , and  $200$ , respectively). For  $n = 200$ , the effect of varying parameter  $\lambda$  ( $\lambda = 0.9, 0.7, 0.5, 0.3$ , and  $0.1$ ) is illustrated.

Gaussian. What we are interested in is to discriminate between small and large effective distances in the distribution, with, respectively, a large and a small weighting. Thus, we may observe that the beginning and the end of the decreasing part of the weighting function are located at distances that correspond approximately to small and large effective distances in the distribution, respectively. Besides these characteristics, the exact shape of the weighting is not important, similarly to other choices of weighting functions in other mapping methods. In this case, the cumulative Gaussian function has been chosen because the central limit theorem ensures that distances in high-dimensional spaces will be Gaussian distributed, at least when the marginal distributions are independent identically distributed (i.i.d.). The weighting function is called a sigmoid balancing function, because of its similarity in shape with the sigmoid function, and its balancing role in the weighting of small and large distances.

Of course, mean  $\mu$  and std  $\sigma$  must be chosen in order to adapt to the effective distribution of data. As a rule-of-thumb, it is suggested to take

$$\mu = \frac{\text{mean}}{1 \leq i < j \leq N}(d_{ij}) - 2(1 - \lambda) \frac{\text{std}}{1 \leq i < j \leq N}(d_{ij}) \quad (4)$$

and

$$\sigma = 2\lambda \frac{\text{std}}{1 \leq i < j \leq N}(d_{ij}) \quad (5)$$

where the mean and std are taken over the distribution of distances between all pairs of data in the original space. Thus, the weighting function can make the difference between *small* and *large* distances even if data come from a high-dimensional space. Such data-dependent weighting is similar to the  $p$ -Gaussian kernel proposed in [48] to take into account the effects of dimensionality.  $\lambda$  is a positive user-defined parameter (usually to be taken between 0.1 and 0.9). Section IV-D details how  $\lambda$  is varied during the course of the algorithm, and Section VI-B shows the effect of  $\lambda$  on the resulting mapping in the case of the two open boxes problem. This single parameter allows controlling how large distances are taken into account in the stress objective function, as compared to small distances. Making it vary leads to weighting functions  $k(d_{ij})$ , as shown

in Fig. 2. Of course, for more flexibility or for the fine-tuning of the weighting function,  $\mu$  and  $\sigma$  could be individually considered.

### B. Importance Given to a Distance Depends on Original Space and Output Space

Section III-B detailed how existing nonlinear dimensionality reduction methods avoid tears either in the original distribution or false neighborhood representations (see also [49]). It is suggested here to avoid as much as possible both drawbacks, by using a weighting function that is *symmetric* with respect to distances in the original ( $d_{ij}$ ) and output ( $d'_{ij}$ ) spaces: Short distances in the two spaces will be emphasized. The weighting function is subsequently defined by

$$k(\min(d_{ij}, d'_{ij})) = 1 - \int_{-\infty}^{\min(d_{ij}, d'_{ij})} f(u, \mu, \sigma) du. \quad (6)$$

Fig. 3 shows the mismatch level between a distance in the original space and the corresponding distance in the output space, in three different situations: without any weighting, with a weighting using the distance in the original space, and with the weighting given by (6), which clearly shows the symmetry of the weighting with respect to both original and output spaces. The symmetric function prohibits that far points in the original space could be displayed as neighbors in the output space, while still allowing tears when they are necessary to map distributions with closed loops.

Note that a symmetric use of distances in the original and output spaces is made possible by the fact that Sammon-like methods precisely aim at making these distances equal. Thus, there is no scaling problem or risk that the distances in both spaces could not be comparable.

The resulting stress function is given by

$$\varsigma = \sum_{i < j} \left( |d_{ij} - d'_{ij}| \cdot \left( 1 - \int_{-\infty}^{\min(d_{ij}, d'_{ij})} f(u, \mu, \sigma) du \right) \right). \quad (7)$$

Usually, the stress function is related to the square of differences between distances. However, absolute values are used here instead of squares, to avoid giving a too high importance to large distances (often responsible for large differences) in the criteria. Moreover, the formulation of the stress proposed in this

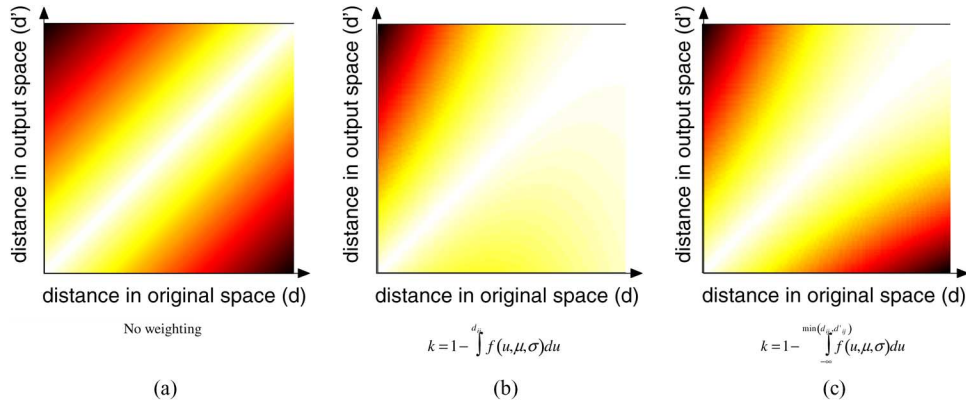


Fig. 3. Stress (weighted mismatch between  $d_{ij}$  and  $d'_{ij}$ ) when (a) no weighting is used, (b) the weighting is based on the distances in the original space, and (c) the symmetric weighting is used (stress increases from light to dark).

paper is consistent with the optimization procedure described in Section IV-C (the spring metaphor).

### C. Optimization by FDP

In general, the optimal position of the data in the output space, resulting from the optimization of stress (7), cannot be obtained analytically. It is necessary to implement a function minimization algorithm with widely recognized robustness and convergence properties. Classically, in the context of dimensionality reduction, one uses the generalized Newton–Raphson algorithm [50], TABU search [51], genetic algorithms [52], [53], simulated annealing [54], or neural networks [6]. To optimize (7) in the context of the proposed method, it is suggested here to use an algorithm based on the FDP.

FDP is an optimization technique for graph visualization introduced by Eades [55]. It compares graphs to spring systems: Nodes are associated to masses and edges to springs between masses [56], [57]. Such system generates forces on the masses, inducing their movement. After a transition phase the system stabilizes; the assumption is made that the final organization corresponds to an acceptable graph representation. The stopping criterion of the algorithm may be a maximum number of iterations, but it has been shown (cf. [56] and [57]) that it is possible to define an energy function whose minimum is attained when the algorithm stabilizes; then, this function may be used to control the convergence and stop the algorithm.

FDP principles are commonly used in graph representation [56]. They are also used for data visualization [58]–[60], as a data set may be seen as a complete graph, the distance matrix defining the edge lengths. A similar approach is also used for the design of printed circuit boards [61], [62]. Although defining objective performance criteria for a mapping must obviously reflect the goal of the user, FDP is known to give satisfying results in mappings for visualization [55]–[57], [60]. One of the main advantages of FDP over other optimization techniques for graph and data visualization is its plasticity: Adding or removing a node or edges rarely induces a strong change in the graph mapping. When new data are added, this makes it possible for the user to keep its intuitive view of the graph and be familiar with the new representation [56], [57]. In addition, FDP makes it possible to escape easily from local minima of the mapping

stress because of the iterative algorithm and the extra random force that accounts for pressure, as detailed in Section IV-E, [58]–[60]. Computational complexity of FDP is  $O(N^2)$  per iteration [58], [59] but can be reduced to  $O(N\sqrt{N})$  [59].

Despite its advantages, FDP is not a central feature of the method presented in this paper. FDP is used as one way to minimize the stress function (7), but other ways could be used as well. In particular, any gradient-based procedure compatible with the high number of parameters (the locations of the data in the mapped space) could be used. Advantages and drawbacks may be found both in the use of FDP and gradient-based optimization procedures. They depend on the number of data, the number of parameters, the complexity of the stress function that results, e.g., from the complexity of the initial manifold, on the need for an easy addition of new data in the mapping, etc.

In the case of the proposed algorithm (FDP), each data  $x_i$  is associated to a node  $x'_i$  in the output space. Each node is linked to all others through springs whose lengths at rest correspond to the distances between nodes in the original space; in this way, FDP places the nodes in the output space keeping all distances as similar as possible to those in the original space.

The stiffness of the springs is adjusted to give more importance to small distances; according to the discussion in Sections IV-A and IV-B, the stiffness is given by (6). The force acting on node  $i$  by node  $j$  is thus given by

$$\begin{aligned} \vec{F}_{x'_i, x'_j} &= (d'_{ij} - d_{ij}) \cdot k(\min(d_{ij}, d'_{ij})) \cdot \vec{u}_{ij} \\ &= (d'_{ij} - d_{ij}) \cdot \left( 1 - \int_{-\infty}^{\min(d_{ij}, d'_{ij})} f(u, \mu, \sigma) du \right) \cdot \vec{u}_{ij} \end{aligned} \quad (8)$$

where  $\vec{u}_{ij}$  is the unitary vector oriented from  $x'_i$  to  $x'_j$ .

At time  $t$ , a node is characterized by its position (noted  $x'_i$ ), its speed (noted  $\vec{v}_i$ ), and its acceleration (noted  $\vec{a}_i$ ).  $\vec{a}_i$  is given by the resultant of forces on the node:  $\vec{a}_i(t) = \sum_{j=1}^N \vec{F}_{x'_j, x'_i} \cdot \Delta t$  being the time increment,  $\vec{v}_i$  is modified according to  $\vec{v}_i(t) = \theta \times \vec{v}_i(t - \Delta t) + \vec{a}_i(t) \times \Delta t$  where  $\theta \in [0, 1]$  is a damping coefficient (here  $\theta = 0.7$ ). The node  $i$  is then moved in the direction of  $\vec{v}_i$  ( $x'_i(t) = x'_i(t - \Delta t) + \vec{v}_i(t) \times \Delta t$ ).

Applying these formulas for acceleration, speed, and position forces the positions  $x'_i$  to converge toward a minimum of

the stress function (7). Indeed the system is relaxed until stability. Stability of each node means that the acceleration of each node, or the sum of the forces applied to it, vanishes. Comparing (7) and (8) shows that this situation is reached when the forces themselves vanish, which results in a minimum of the stress function. However, stability could also be reached when the sum of forces applied to a node vanishes, while forces do not. This situation corresponds to a local minimum of the stress function; Section IV-E describes a stochastic perturbation scheme designed to escape from such a minimum.

The system is relaxed until stability. The level of stability may be measured by the total energy in the system given by

$$E = \frac{1}{2} \sum_{i=1}^N \|\vec{v}_i\|^2 \quad (9)$$

where  $N$  is the number of nodes. When the system becomes stable, the positions of the nodes  $x'_i$  in the output space form a mapping of the original data. Using criterion (9) instead of (7) to measure the level of stability is justified by the fact that (9) can be compared to 0 with a simple, not critical threshold, while (7) never reaches 0; using (7) would mean the development of a strategy based on the empirical derivative, which reveals to be much more critical in practice. Using a stopping criterion such as (9) is also standard practice within FDP framework [56], [57].

Of course, as it is the case in any distance-based dimensionality reduction method, the orientation of the resulting graph has no specific meaning, as any result obtained by rotation or symmetry would be equivalent. The possibility to obtain different mappings when the algorithm is run several times on the same data results from its stochastic character; the only stochastic part of the method is described in Section IV-E.

#### D. Dimensionality Reduction Algorithm

The proposed dimensionality reduction algorithm is detailed here, based on the concepts described previously.

The goal of the algorithm is to find the locations  $x'_i$  of the points in the output space. The unknowns of the optimization procedure are thus these locations. To find them, the stress function given by (7) is minimized. In practice, if the FDP optimization algorithm is used, this is done by moving the points so as to minimize (7) after the computation of acceleration and speed.

The algorithm first builds a global representation based on a limited number of data, and then, iteratively adds subsets of data to refine the mapping at the local level, possibly at the price of a global distortion. Each addition of new data is followed by a learning phase aimed at refining the mapping.

The order of selection of data is obtained because of a procedure that was described by Hastie *et al.* to select adequate “seeds” before a clustering procedure [2]. The advantage of this procedure is that the selected data are guaranteed to spread over the whole domain of the original data; the well-known drawback of Hastie’s clustering procedure, i.e., the need for the complete distance matrix is not relevant here since this matrix is also required for the mapping within the DD-HDS framework. More precisely, the prototypes are selected as follows.

The first prototype is selected as the data for which the sum of distances to all other data is minimum (it is the closest data to the center of gravity).

The  $(i + 1)$ th prototype is selected as the unselected data giving the best quantization of the data when it is associated with the  $i$  already selected prototypes. The best quantization is defined as the minimum of the quantization error, i.e., the sum of distances between all data and their nearest prototypes.

Once the full data set is sorted as described previously, data are positioned in the output space. It is always possible to place  $p + 1$  data in a  $p$ -dimensional output space, while exactly preserving their distances in the original space. The  $p + 1$  first data (prototypes) are, therefore, first placed in this way. Next, the number of prototypes to map onto the output is increased (multiplied by two, for example) according to the selection order. The FDP algorithm is then used to find a stable configuration considering the new prototypes. Steps of prototype choice, and steps of subset representation are alternated until the total number of data is used. In practice, doubling the number of prototypes after each relaxation step appears to be a good tradeoff between the number of partitions and the relaxation time.

As the stiffness of springs (8) makes use of  $k(\min(d_{ij}, d'_{ij}))$ , it is necessary to choose the value of  $\lambda$  for insertion in (4) and (5). At the beginning of the algorithm, prototypes are far one from another, due to the data ordering procedure. It is thus legitimate to choose a high value for  $\lambda$ , so that large distances influence the mapping. When the number of data is increased, the value of  $\lambda$  is decreased to give more importance to neighborhoods and small distances. The effect of  $\lambda$  is, therefore, similar to the neighborhood parameters in Kohonen’s SOMs. Its final value reflects the user-driven compromise between the efficiency of local representations and a global view of the data cloud or manifold.

For the experiments presented in Section VI, the number of data was doubled and  $\lambda$  was monotonously decreased from 0.9 to 0.1 after each step.

#### E. Pressure Allows Avoiding Local Minima

Forces given by (8) are applied to each node of the graph; the resultant of forces moves the node. The sum of the modules of these forces may be interpreted in terms of “pressure”

$$P_i = \sum_{j=1, j \neq i}^N \|\vec{F}_{x'_i, x'_j}\|. \quad (10)$$

Our use of the “pressure” term is not academic: It does not strictly follow the physical definition of pressure. Nevertheless, it allows differentiating stable nodes because forces are weak (low value of  $P_i$ ), from stable nodes because nonnull forces mutually compensate (high value of  $P_i$ ). In the latter case, the position of the node corresponds to a local minimum of the stress function. In order to escape from this minimum, a supplementary force is added along a random direction (simulating a Brownian movement). Its intensity is function of the local pressure

$$\vec{F}_{\text{Brownian}_i} = \alpha(\text{iteration}) \times \frac{P_i}{N} \times \vec{u}_i \quad (11)$$

where  $\vec{u}_i$  is a unit vector randomly oriented and  $\alpha$  tends toward 0 when the number of iterations increases to allow system relaxation. Equation (11) describes the only stochastic part of the method. It is responsible for the fact that slightly different mappings could result from different runs of the algorithm, as detailed in Section IV-C.

The resulting algorithm is called DD-HDS.

#### F. Computational Complexity

Computational complexities of DD-HDS and other nonlinear mapping algorithms such as CCA and Sammon's mapping are similar. Compared to CCA and Sammon's mapping, the sigmoid balancing function replaces other weighting functions with similar complexity, and the FDP optimization procedure is used as an efficient alternative to gradient-based procedures. The computational complexity of all nonlinear mapping methods is of course larger than the complexity of PCA; the latter relies only on linear algebra computations, while nonlinear mappings require optimization procedures. This difference is essential in terms of computational complexity; the computational load of the optimization procedures themselves is, however, quite impossible to evaluate in practical situations, as it depends dramatically on the content of the initial data set, which determines when the stopping criterion is reached. FDP methods are, however, reputed to be fast and robust [57]; the reader is referred to [56], for a discussion about the computational complexity of FDP.

### V. VISUALIZATION OF THE MAPPING EFFICIENCY

Local and global visualizations may be used in order to explore the efficiency of the mapping with DD-HDS.

Pressure (10) gives valuable information about the efficiency of mapping: The better the placement of a data with respect to its neighbors, the smaller the pressure it undergoes. It must be kept in mind, however, that the pressure depends on  $\lambda$ , i.e., on the size of the effective neighborhood (set by the user).

Criterion (10) may be averaged over all data  $x_i^l$  for a global measure of the representation efficiency. Alternatively, Demartine's  $dy-dx$  diagram [6], [63] may be used to view how distances are preserved in the output space. The principle of this diagram is to plot  $d'_{ij}$  as a function of  $d_{ij}$  for all pairs of data. All pairs for which the distance in the output space is exactly equal to the one in the original space fall on the diagonal of this graph; if short distances are mapped without much distortion, only small departures from the diagonal may be seen close to the origin of the  $dy-dx$  diagram axes.

### VI. EXPERIMENTAL RESULTS

In this section, four databases are mapped on a 2-D space. The two first examples (earth globe and open boxes) are 3-D distributions for which the intrinsic dimension is two (two coordinates are sufficient to describe the location of a point on the distribution).

The two last examples are high-dimensional distributions. The first one is the set used by Tenenbaum [18]: Data are pictures of a virtual face viewed under various angles and illumination conditions. As no preprocessing is used, the dimension of

the data is the number of pixels in the images, i.e., 4096. However, the intrinsic dimensionality is three, as two viewing angles and one illumination angle are sufficient to characterize each image. As in [18], the data are mapped here on a 2-D surface.

The last example is a real-world high-dimensional data set, carrying some information extracted from genome sequences of living species. The intrinsic dimension of these data is not known, as traditional dimension estimation techniques do not result in convincing results. Nevertheless, a 2-D mapping is of interest in order to make it possible browsing the data space.

#### A. Earth Globe

Data to be mapped are 273 large cities around the world. Their distances in the original space are calculated in the 3-D space (left graph of Fig. 4). The mapping by DD-HDS is given in the right graph of Fig. 4. It can be seen that the mapping accounts for the local density of cities. The north hemisphere is properly developed. Continents can be identified. Cities-free areas (such as Pacific Ocean, Antarctic, etc.) are distorted although continuity is preserved in most places.

The grid materializes latitudes and longitudes and shows the deformations resulting from the mapping. The grid was not used during the mapping process. It has been placed on the representation *a posteriori*, through interpolations between mapped cities: Each intersection between latitude and longitude lines were placed in order to best fit its distances with cities in the original space [according to (8)]. Then, points were connected to give a lattice. This interpolation procedure is not specific to DD-HDS and could be used in a similar way with other mapping methods.

This example makes it possible to understand the viewpoint proposed by DD-HDS. Short distances are properly represented, while large ones are not necessarily mapped in a realistic way.

#### B. Open Boxes

Original data live in a 3-D space. They are situated on the sides of two open cubes with open sides pointing toward two different directions (Fig. 5, upper left part). The mapping by DD-HDS is shown to its right. Under these plots, the  $P_i$  criterion (10) for each point (left), and the  $dy-dx$  diagram (right) are shown. Others subplots display mappings achieved by competing methods.

This simulation illustrates on a simple example the advantages of DD-HDS. It correctly develops the two boxes, despite a twist on a large scale. The method effectively combines obvious nonlinear properties with a faithful representation of neighborhoods. Sammons' mapping gives a more expected view, but the lateral faces of the cubes are drastically compressed.

Even if nonlinear mappings are achieved with most methods, it can be seen that more or less intuitive representations are obtained. In the case of SOM, the two clusters are correctly found and neighborhoods are preserved, but the shape of the original objects is not recovered. Methods based on the geodesic concept (isomap and LLE) give two disconnected plots for the two boxes as there is no path available between them [see point 2) in Section II-B]. The black line between the two box representations expresses this segmentation. The impact of long distances



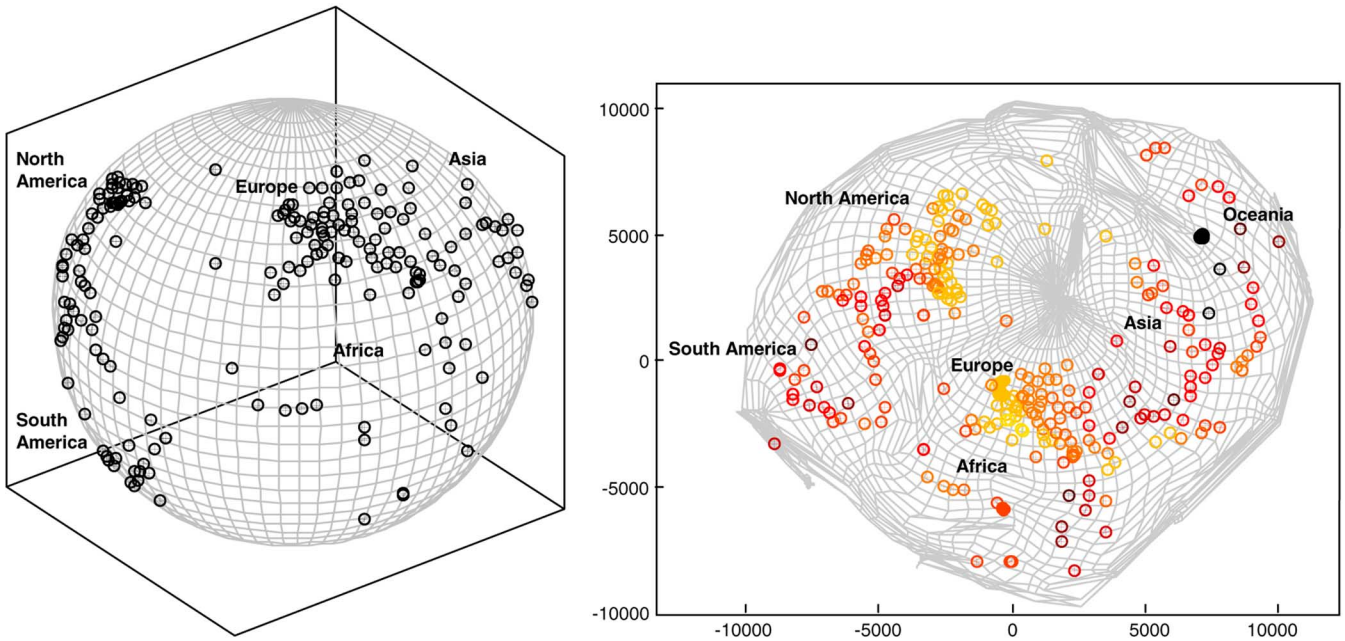


Fig. 4. Mapping of the earth globe (defined by large cities) in a 2-D space. Color (graph on the right) indicates the satisfaction of pairwise distances for the corresponding city [pressure, (10)]. Darkest points indicate highest pressures. For clarity, the part of the grid close to the South Pole is not displayed.

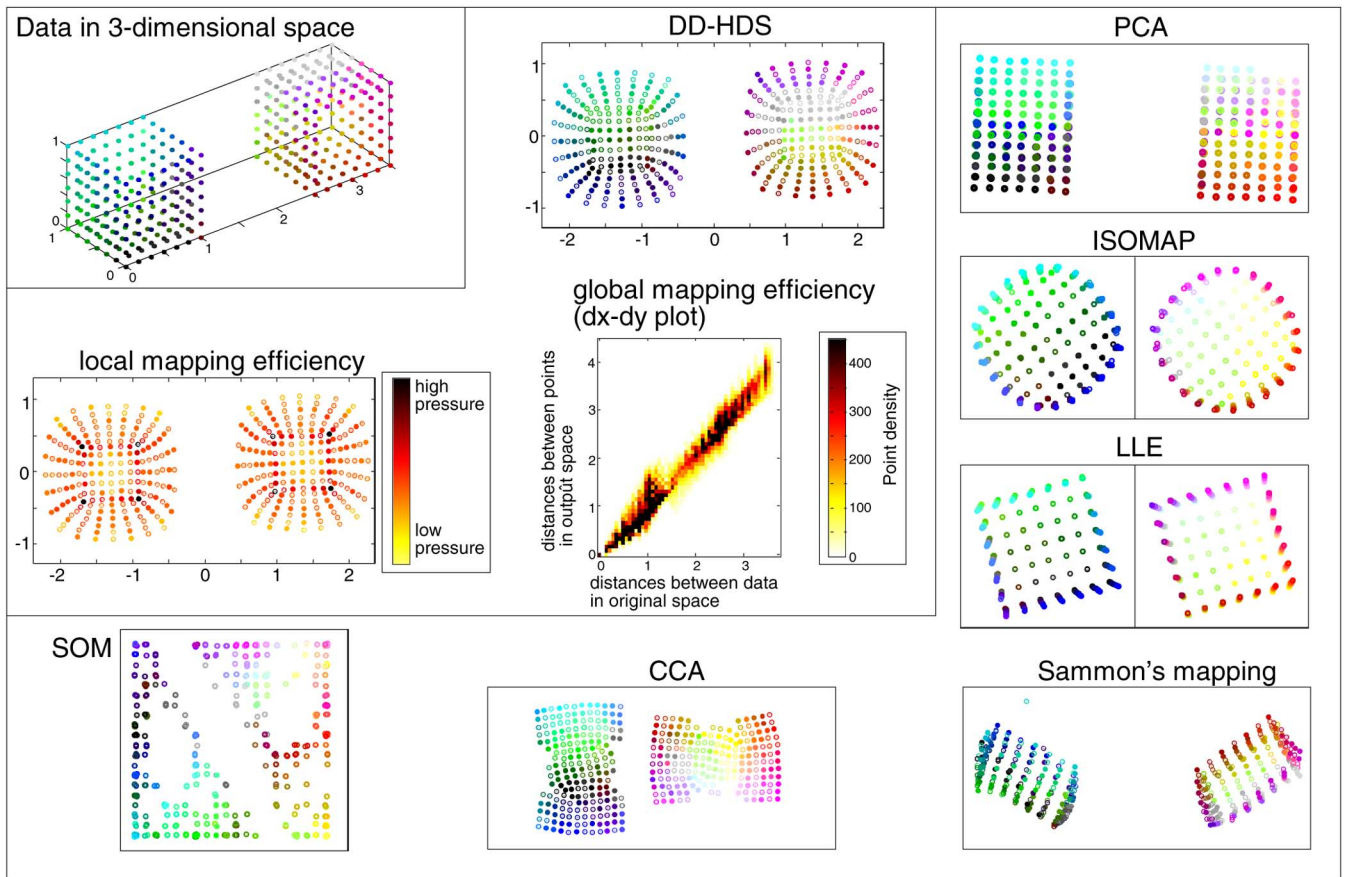


Fig. 5. Mapping of two 3-D open boxes in a 2-D space. The color codes for the position of the data in the original space (except for the local mapping efficiency and global mapping efficiency plots). Upper left—original data (3-D space), upper right—mapping by DD-HDS (2-D space), lower left—pressure (color lut is similar to Fig. 4), and lower right—pairwise distance preservation (color codes for density of distances). Other subplots are self-explanatory. Isomap code used for this simulation is from <http://isomap.stanford.edu/>, SOM is from Matlab neural network toolbox (nnet), and Sammon's mapping and LLE are from <http://www.cis.hut.fi/projects/somtoolbox/>.

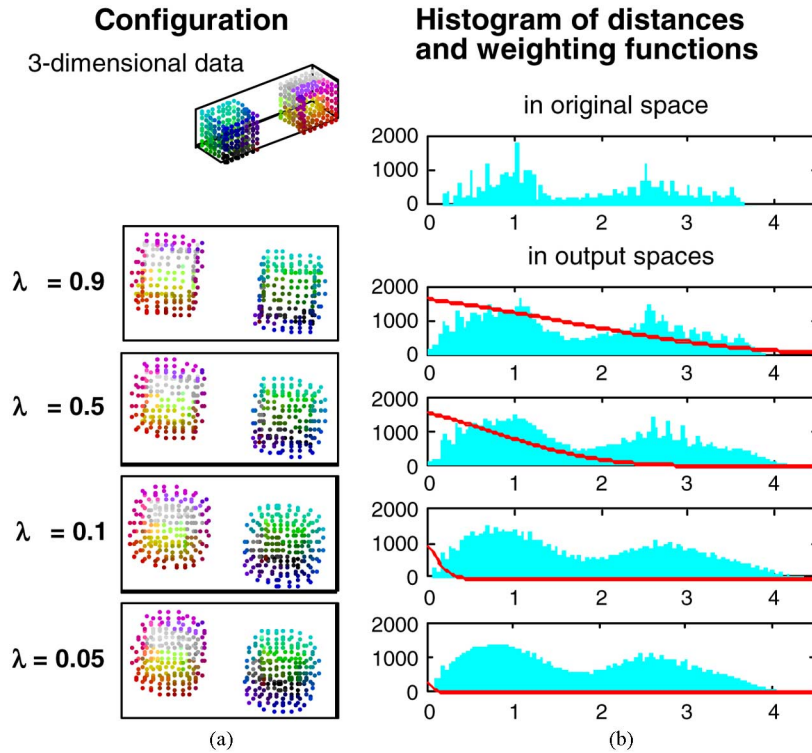


Fig. 6. Mapping of the two boxes data sets according to  $\lambda$ . (a) Color code is the same as in Fig. 5. (b) Associated distances in output space and corresponding weighting functions.

does not allow the development of the sides adjacent to the open side of the boxes.

On the one hand, Sammon's mapping does not generate any tear, but many false neighborhoods (just as PCA does). On the other hand, CCA succeeds in mapping data without any false neighborhood, but some tears can be observed. This result was expected (see Section III-B). Except for the mapping produced by DD-HDS, all other methods show false neighborhoods and/or tears.

The two open boxes data sets can be used to visualize the impact of  $\lambda$  (Fig. 6). High values for  $\lambda$  increase the quality of global structure representations (highest left panels) but neighborhood relations are jeopardized. Low values for  $\lambda$  (0.1, for example) permit a better neighborhood representation, but the overall shape is not guaranteed. Very low values for  $\lambda$  (here 0.05) generate "unreasonable" weighting functions although the mapping may still be found acceptable.

### C. Face Data

The set contains 698 pictures of the same virtual face under different angle and illumination conditions. It has been used by Tenenbaum [18] who tested isomap on high-dimensional data. The original dimension of the data is 4096 (number of pixels), and their intrinsic dimension should be 3 (two viewing angles and one illumination angle). Fig. 7 shows the mappings achieved by PCA, Sammon's mapping, CCA and DD-HDS with Euclidean distance, and Isomap and DD-HDS with geodesic distance (five neighbors have been used to build the grid for geodesic distance). Each resulting map is displayed three times, from top to bottom in each column, but points are colored differently: From top to bottom, they are colored according to the

horizontal viewing angle, the vertical viewing angle, and the illumination angle. Typical examples of images bordered by the color corresponding to the respective angle are shown in the left part of Fig. 7.

The intrinsic dimension of this data set is 3. Actually, most of the tested methods succeed in mapping these data onto a 3-D space. However, maps on Fig. 7 are generated in a 2-D space. This test cannot be perfectly passed: Tears or false neighborhoods are unavoidable. Here, the challenge is to get as less tears as possible while avoiding false neighborhoods. As expected, PCA and Sammon's mappings display high levels of false neighborhood and tears. DD-HDS and CCA with Euclidean distance effectively place points close from one another when their characteristics are similar. This can be seen through the fact that close points have similar colors on the three graphs. Nevertheless, much more exceptions (close points with different colors) are found with CCA. The methods illustrated in the two last columns of Fig. 7 (isomap and DD-HD) use geodesic distance in the original space. It is much easier to observe the continuity of the three intrinsic parameters (the three angles) in these two mappings. The horizontal angle is properly mapped by both methods. The vertical angle is also properly captured by isomap, whereas DD-HDS provides a smooth mapping of light rotation. However, only DD-HDS mapping is such that almost all pairs of close points have similar values (colors) for each of the three angles. Having all three angles similar in a pair of points that are close in the mapping is indeed the necessary condition to have two close points corresponding to close faces (close faces means that all three angles are similar). Although both isomap and DD-HDS lead to close points that do not fulfill these requirements (again,

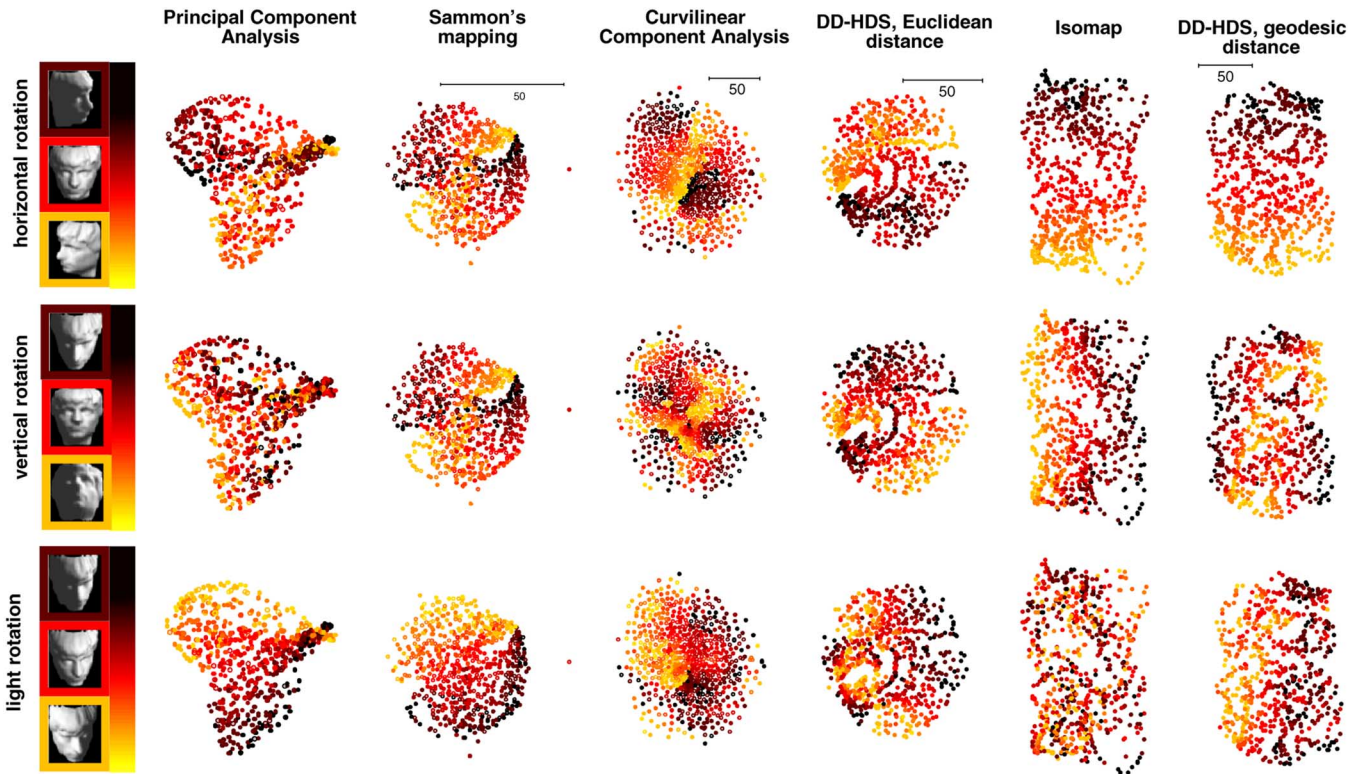


Fig. 7. Representation of face data by PCA, Sammon's mapping, CCA and DD-HDS (four first columns, Euclidean distances), and Isomap and DD-HDS (two last columns, geodesic distances). See Section IV-C for details. Scales of distances are shown for the methods that respect Euclidean distances (Sammon's mapping, CCA, and DD-HDS on Euclidean distances).

tears and false neighbors are unavoidable in this application), Fig. 7 shows that the number of such situations is much lower in DD-HDS: A large number of false neighbors (in isomap) has been replaced by a lower number of tears (in DD-HDS).

#### D. Genomic Signature Issues

Dealing with real data often rises problems that are not encountered with simulated data. In particular, the eventual complexity of real-data distribution in high-dimensional space may strongly reduce mapping efficiency.

Genomic signatures are high-dimensional data resulting from the analysis of DNA sequences in terms of short oligonucleotide frequencies. Within the paradigm of genomic signature, DNA sequences are considered as "texts" built with a four-letter alphabet (nucleotides A, T, C, and G). Short oligonucleotides are small DNA sequences (usually 2–8 nucleotides long). It has been shown that the set of oligonucleotide frequencies calculated from a DNA sequence at least several thousands nucleotides long (the so-called genomic signature) is species specific, i.e., different species have different signatures [24].

Similarities between species allow building a taxonomy of life, usually displayed as a tree (the famous tree of life). Traditional taxonomy is essentially based on macroscopic observations about species. Genomic features may also be used. Along branches of the tree of life, successive refinements lead to an accurate description of species that are the leaves of the tree. For example, "root, cellular organism, Eukaryota,

Fungi/Metazoa group, Metazoa; Eumetazoa, Bilateria, Coelomata, Deuterostomia, Chordata, Craniata, Vertebrata, Gnathostomata, Teleostomi, Euteleostomi, Sarcopterygii, Tetrapoda, Amniota, Mammalia, Theria, Eutheria, Euarchontoglires, Primates, Simiiformes, Catarrhini, Hominoidea, Hominidae, Homo/Pan/Gorilla group, and Homo Sapiens" is the path in the tree to a well known species.<sup>1</sup>

It has been shown that close species in terms of taxonomy have close signatures and vice versa. It is not realistic yet to observe (and corroborate) the full tree of life using the genomic signature as a criterion. However, large partitions (at the bottom of the tree) as well as some specific branches are already properly described with it (see [64]–[66], for examples). It seems, therefore, interesting to check the ability of dimensionality reduction methods to preserve the taxonomic features of genomic signatures.

Depending on the size of examined oligonucleotides, the genomic signature may have a various number of dimensions, ranging usually from 16 to 65 536. For this paper, we focused on 256-dimension signatures that have the most interesting properties with respect to taxonomy. The study presented here concerns 2046 genomic signatures illustrating the diversity of living organisms [697 Eukaryotes (plants, vertebrates, fungus, etc.), 1349 Prokaryotes made up of 1287 bacteria, and 62 archeobacteria].

It has been shown that taxonomic information can be derived from genomic signatures by means of the Euclidean metric

<sup>1</sup>National Center for Biotechnology Information (NCBI): [www.ncbi.nlm.nih.gov/Taxonomy/taxonomyhome.html](http://www.ncbi.nlm.nih.gov/Taxonomy/taxonomyhome.html)




Fig. 8. Weighting functions fitted to the distributions of distances: (a) Euclidean metric. (b) Rank-order pseudometric. (c) Geodesic metric (connectivity=5).

that allows characterizing similarity between them [25], [64]. Euclidean metric and derived “metrics” (rank and geodesic) were successively used within DD-HDS to compare genomic signatures. According to the procedure described previously, the weighting function was fitted to the various distributions of distances (Fig. 8). The interest of an adaptive weighting function is obvious here, considering the diversity in values and shapes of the distance distributions.

Fig. 9 shows the mapping obtained by PCA, KPCA, SOM, Sammon’s mapping, CCA, isomap, and DD-HDS using the three “metrics.” In the upper part of the figure, we are concerned with the ability of the mapping to express segmentation between signatures near the root of the tree of life. The groups the species belong to (namely Eukaryotes, bacteria, and Archebacteria) are recognized by the gray intensity of points. Mappings achieved by PCA and KPCA do not clearly reveal the taxonomic features of the genomic signatures. The high nonlinear correlations between variables are likely responsive for the typical croissant shaped layout. Although Sammon’s mapping makes a better use of the output space, it fails to display the species organization: The projection remains folded, because of the limitations resulting from the Sammon’s stress weighting by distances in the original space, and also probably because of the concentration of measure phenomenon (see Section III-B). Groups are not easily identifiable in the Kohonen map although a general organization of signatures is observable. Isomap allows separating Eukaryotes from bacteria but overlapping remains important. CCA offers an interesting display where groups can be localized and segmented. The paving-like structure may result from the ultimate preservation of short distances between signatures, considering that in high-dimensional Euclidean spaces, even “short” distances are “long.”

DD-HDS also offers mappings where groups of species can be segmented. In addition, DD-HDS provides a sharper separation between groups of species.

Mappings obtained using the rank-order pseudodistance and with isomap are pretty close. The spatial orientation of data follows the nucleotide bias (frequencies of the nucleotides are not necessarily equal over species). It is already known in fact that although the nucleotide bias largely varies between species, it is not linked to taxonomy. The nucleotide bias explains an important part of the overall dispersion of genomic signatures in the high-dimensional space. It is also captured by PCA and KPCA. Mappings obtained from Euclidean distances and geodesic distances are more informative. In particular, substructures in data are observable (see subplots in Fig. 9, lower part, where actinobacteria are highlighted). They correspond to well-identified subgroups of species and are probably the expression of local substructures in the original space.

We would like to point out that the mapping of genomic signatures achieved in this paper would have been better, if more dimensions were allowed for the output space. Our experience with genomic signatures, DD-HDS, and other experimental protocols suggests that the intrinsic dimension of the genomic signatures should be around 7–8.

## VII. CONCLUSION

This paper presents DD-HDS, a mapping method designed to take into account the specificities of high-dimensional data. In particular, it introduces a specific weighting of distances taking into account the concentration of measure phenomenon, and a symmetric handling of short distances in the original and output spaces, avoiding false neighbor representations while still allowing some necessary tears.

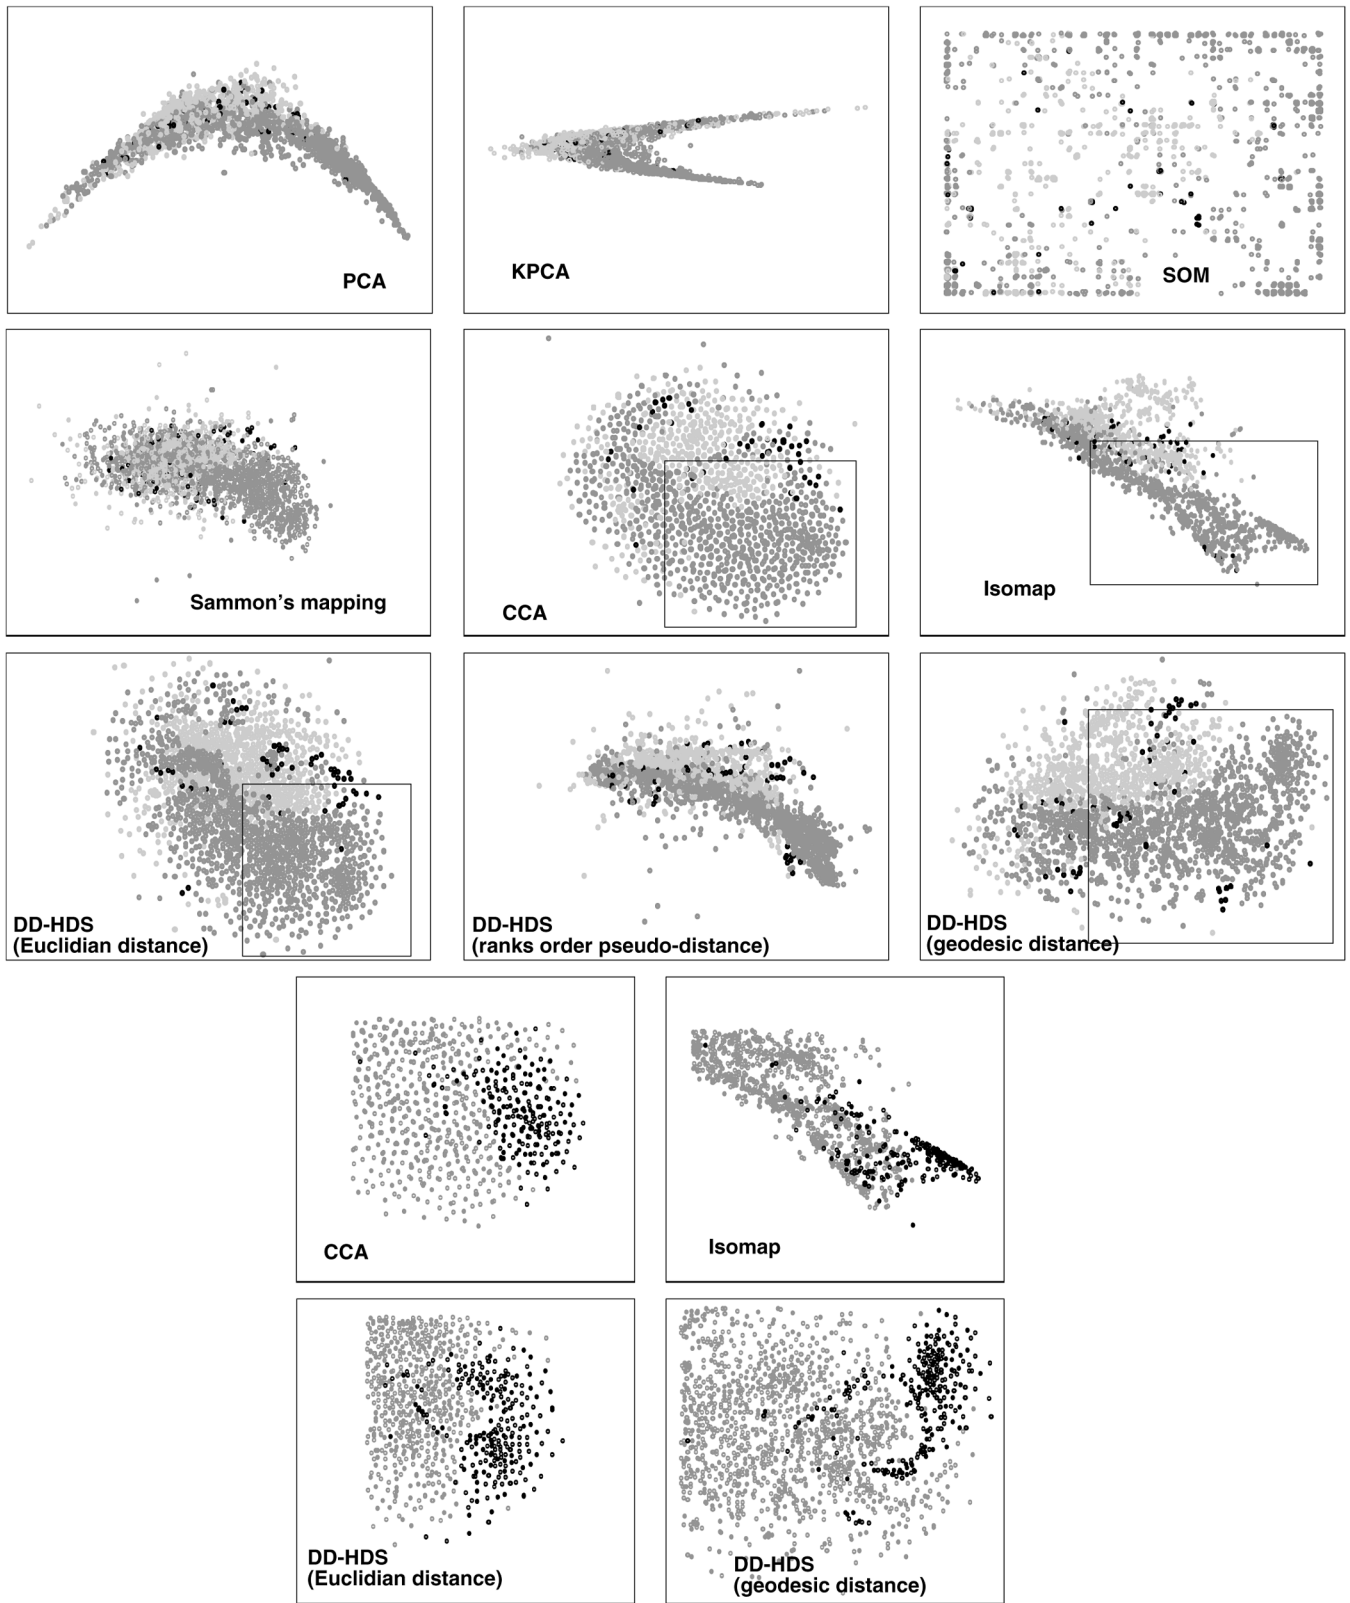


Fig. 9. Mappings of 2046 genomic signatures. Upper panels: signatures as mapped by PCA, KPCA (polynomial kernel of degree 2), SOM (35X55 nodes), Sammon's mapping, isomap (connectivity =5), CCA, and DD-HDS using Euclidean metric, and derived metrics or pseudometrics [rank order and geodesic (connectivity =5)]. Points correspond to species. Grey levels code for taxonomy: light gray = Eukaryotes, dark gray = bacteria, and black = archeobacteria. Lower panels: highlights of the subgroup of actinobacteria for some of the aforementioned mappings (others species are in grey).

Giving an "objective" quantitative evaluation of the efficiency of mapping methods is quite difficult: There is

obviously a subjective part in the low-dimensional mapping of high-dimensional data. Some authors even use the term

aesthetic or pleasant drawing. For the genomic signatures data set at least, we observe structures that correspond to known organizations of species based on the tree of life. The different hierarchical levels of organization, which are available here, may be more or less detected, depending on the methods and parameters. Only experts (biologists in that case) may validate (or invalidate) the mappings. It is believed that the exploration of high-dimensional data must be somehow supervised (i.e., user driven), and, depending on the “point of view,” mappings may be quite different and more or less satisfying. In DD-HDS, a single user-defined parameter allows fixing the compromise between local neighborhood preservation and global mapping; in our experience, this feature turns out very convenient for the interactive exploration of high-dimensional data.

#### ACKNOWLEDGMENT

The authors would like to thank J. A. Lee, Machine Learning Group, Université catholique de Louvain, Belgium, for kindly providing his efficient implementation of the CCA algorithm.

#### REFERENCES

- [1] D. L. Donoho, “High-dimensional data analysis: The curses and blessings of dimensionality,” Los Angeles, CA, 2000, Amer. Math. Soc. Lecture: “Math challenges of the 21st century” [Online]. Available: <http://www-stat.stanford.edu/~donoho/>
- [2] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*. New York: Springer-Verlag, 2001, p. 470.
- [3] W. Basalaj, “Proximity visualization of abstract data,” Ph.D. dissertation, Comput. Lab., Univ. Cambridge, Cambridge, U.K., 2000.
- [4] W. S. Torgerson, “Multidimensional scaling: 1. Theory and method,” *Psychometrika*, vol. 17, pp. 401–419, 1952.
- [5] J. W. Sammon, “A nonlinear mapping for data structure analysis,” *IEEE Trans. Comput.*, vol. C-18, no. 5, pp. 401–409, May 1969.
- [6] P. Demartines and J. Héroult, “Curvilinear component analysis: A self-organizing neural network for nonlinear mapping of data sets,” *IEEE Trans. Neural Netw.*, vol. 8, no. 1, pp. 148–154, Jan. 1997.
- [7] J. Héroult, A. Guérin-Dugué, and P. Villemain, “Searching for the embedded manifolds in high-dimensional data, problems and unsolved questions,” in *Proc. Euro. Symp. Artif. Neural Netw. (ESANN)*, Bruges, Belgium, 2002, pp. 173–184.
- [8] C. C. Aggarwal, A. Hinneburg, and D. A. Keim, “On the surprising behavior of distance metrics in high dimensional space,” in *Lecture Notes in Computer Science*, ser. 1973, J. V. Bussche and V. Vianu, Eds. Berlin, Germany: Springer-Verlag, 2001, pp. 420–434.
- [9] D. Lowe and M. E. Tipping, “Feed-forward neural networks and topographic mappings for exploratory data analysis,” *Neural Comput. Appl.*, vol. 4, pp. 83–95, 1996.
- [10] B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Cambridge, MA: MIT Press, 2002.
- [11] B. Schölkopf, A. J. Smola, and K. R. Müller, “Nonlinear component analysis as a kernel eigenvalue problem,” *Neural Comput.*, vol. 10, pp. 1299–1319, 1998.
- [12] J. Shawe-Taylor and N. Cristianini, *Kernel Methods for Pattern Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [13] W. Wu, D. L. Massarat, and S. de Jong, “The kernel PCA algorithm for wide data. Part II: Fast cross-validation and application in classification of NIR Data,” *Chemometrics Intell. Lab. Syst.*, vol. 37, pp. 271–280, 1997.
- [14] C. K. I. Williams and M. Seeger, “Using the Nyström method to speed up kernel machines,” in *Advances in Neural Information Processing Systems*. Cambridge, MA: MIT Press, 2001, pp. 682–688.
- [15] K. Q. Weinberger, F. Sha, and L. K. Saul, “Learning a kernel matrix for nonlinear dimensionality reduction,” in *Proc. 21st Int. Conf. Mach. Learn. (ICML)*, Banff, AB, Canada, 2004, pp. 839–846.
- [16] K. Q. Weinberger, B. D. Packer, and L. K. Saul, “Nonlinear dimensionality reduction by semidefinite programming and kernel matrix factorization,” in *Proc. 30th Int. Workshop Artif. Intell. Statist.*, Z. Ghahramani and R. Cowell, Eds., Barbados, 2005, pp. 381–388.
- [17] J. Sun, S. Boyd, L. Xiao, and P. Diaconis, “The fastest mixing Markov process on a graph and a connection to a maximum variance unfolding problem,” *SIAM Rev.*, vol. 48, no. 4, pp. 681–699, 2006.
- [18] J. B. Tenenbaum, V. de Silva, and J. C. Langford, “A global geometric framework for nonlinear dimensionality reduction,” *Science*, vol. 290, pp. 2319–2323, 2000.
- [19] J. A. Lee, A. Lendasse, and M. Verleysen, “Nonlinear projection with curvilinear distances: Isomap versus curvilinear distance analysis,” in *Neurocomputing*. New York: Elsevier, 2004, vol. 57, pp. 49–76.
- [20] E. W. Dijkstra, “A note on two problems in connection with graphs,” *Numerisch Mathematik*, vol. 1, pp. 269–271, 1959.
- [21] S. Skiena, *Implementing Discrete Mathematics: Combinatorics and Graph Theory with Mathematica*. Reading, MA: Addison-Wesley, 1990.
- [22] J. A. Lee, N. Lendasse, N. Donckers, and M. Verleysen, “A robust nonlinear projection method,” in *Proc. Euro. Symp. Artif. Neural Netw. (ESANN)*, 2000, pp. 13–20.
- [23] J. A. Lee and M. Verleysen, “Nonlinear dimensionality reduction of data manifolds with essential loops,” *Neurocomput.*, vol. 67, pp. 29–53, 2005.
- [24] P. J. Deschavanne, A. Giron, J. Vilain, G. Fagot, and B. Fertil, “Genomic signature: Characterization and classification of species assessed by chaos game representation of sequences,” *Mol. Biol. Evol.*, vol. 16, pp. 1391–1399, 1999.
- [25] C. Dufraigne, B. Fertil, S. Lespinats, A. Giron, and P. Deschavanne, “Detection and characterization of horizontal transfers in prokaryotes using genomic signature,” *Nucleic Acids Res.* vol. 33, p. e6, 2005 [Online]. Available: <http://nar.oxfordjournals.org/cgi/content/abstract/33/1/e6>
- [26] J. B. Kruskal, “Non-metric multidimensional scaling: A numerical method,” *Psychometrika*, vol. 29, pp. 115–129, 1964.
- [27] J. B. Kruskal, “Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis,” *Psychometrika*, vol. 29, pp. 1–27, 1964.
- [28] I. Borg and P. Groenen, *Modern Multidimensional Scaling*. New York: Springer-Verlag, 1997.
- [29] T. Cox and M. Cox, *Multidimensional Scaling*. London, U.K.: Chapman & Hall, 1994.
- [30] S. T. Roweis and L. K. Saul, “Nonlinear dimensionality reduction by locally linear embedding,” *Science*, vol. 290, pp. 2323–2326, 2000.
- [31] M. Belkin and P. Niyogi, “Laplacian eigenmaps for dimensionality reduction and data representation,” *Neural Comput.*, vol. 15, pp. 1373–1396, 2003.
- [32] D. L. Donoho and C. Grimes, “Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data,” *Proc. Nat. Acad. Sci.*, vol. 100, pp. 5591–5596, 2003.
- [33] S. T. Roweis, L. K. Saul, and G. Hinton, “Global coordination of local linear models,” in *Advances in Neural Information Processing Systems*. Cambridge, MA: MIT Press, 2001, vol. 14, pp. 889–896.
- [34] S. Wold, A. Ruhe, H. Wold, and W. J. Dunn, “The collinearity problem in linear regression. The partial least squares (PLS) approach to generalized inverses,” *SIAM J. Sci. Statist. Comput.*, vol. 5, pp. 735–743, 1984.
- [35] T. Kohonen, “Self-organized formation of topologically correct feature maps,” *Biol. Cybern.*, vol. 43, pp. 59–69, 1982.
- [36] T. Kohonen, *Self-Organizing Maps*, H. K. V. Lotsch, Ed. Heidelberg, Germany: Springer-Verlag, 1997.
- [37] S. Wu and T. W. S. Chow, “PR-SOM: A new visualization method by hybridizing multidimensional scaling and self-organizing map,” *IEEE Trans. Neural Netw.*, vol. 16, no. 6, pp. 1362–1380, Nov. 2005.
- [38] H. Yin, “ViSOM: A novel method for multivariate data projection and structure visualization,” *IEEE Trans. Neural Netw.*, vol. 13, no. 1, pp. 237–243, Jan. 2002.
- [39] H. Yin, “Data visualization and manifold mapping using the ViSOM,” *Neural Netw.*, vol. 15, pp. 1005–1016, 2002.
- [40] C. M. Bishop, M. Svensén, and C. K. I. Williams, “GTM: The generative topographic mapping,” *Neural Comput.*, vol. 10, pp. 215–234, 1998.
- [41] A. Kaban and M. Girolami, “A combined latent class and trait model for the analysis and visualization of discrete data,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 8, pp. 859–872, Aug. 2001.
- [42] N. D. Lawrence, “Probabilistic non-linear principal component analysis with Gaussian process latent variable models,” *J. Mach. Learn. Res.*, vol. 6, pp. 1783–1816, 2005.

- [43] K. V. Mardia, J. T. Kent, and J. M. Bibby, *Multivariate Analysis*. London, U.K.: Academic Press, 1979.
- [44] J. Héroult, A. Oliva, and A. Guérin-Dugué, "Scene categorization by curvilinear component analysis of low frequency spectra," in *Proc. Euro. Symp. Artif. Neural Netw. (ESANN)*, Bruges, Belgium, 1997, pp. 91–96.
- [45] J. A. Lee and M. Verleysen, "Nonlinear projection with the isotop method," in *Lecture Notes in Computer Science*, ser. 2415, J. R. Dorronsoro, Ed. Berlin, Germany: Springer-Verlag, 2002, pp. 933–938.
- [46] J. A. Lee, C. Archambeau, and M. Verleysen, "Locally linear embedding versus isotop," in *Proc. Euro. Symp. Artif. Neural Netw. (ESANN)*, Bruges, Belgium, 2003, pp. 527–534.
- [47] N. D. Lawrence and J. Quiñero-Candela, "Local distance preservation in the GP-LVM through back constraints," in *Proc. 23rd Int. Conf. Mach. Learn. (ICML)*, W. Cohen and A. Moore, Eds., Pittsburgh, PA, 2006, pp. 513–520.
- [48] D. Francois, V. Wertz, and M. Verleysen, "About the locality of kernels in high-dimensional spaces," in *Proc. Int. Symp. Appl. Stochastic Models Data Anal. (ASMDA)*, Brest, France, 2005, pp. 238–245.
- [49] G. Hinton and S. T. Roweis, "Stochastic neighbor embedding," in *Proc. Neural Inf. Process. Syst. (NIPS)*, Vancouver, BC, Canada, 2002, pp. 833–840.
- [50] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge, U.K.: Cambridge Univ. Press, 1992.
- [51] F. Glover and M. Laguna, , C. R. Reeves, Ed., *Tabu Search*. New York: McGraw-Hill, 1995.
- [52] D. E. Goldberg, *Genetic algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley, 1989.
- [53] C. R. Reeves, , C. R. Reeves, Ed., *Genetic Algorithms*. New York: McGraw-Hill, 1995.
- [54] K. A. Dowsland, *Simulated Annealing*, C. R. Reeves, Ed. New York: McGraw-Hill, 1995.
- [55] P. Eades, "A heuristic for graph drawing," in *Proc. 13th Manitoba Conf. Numer. Math. Comput. (Congressus Numerantium)*, D. S. Meek and G. H. J. V. Rees, Eds., Winnipeg, MB, Canada, 1984, vol. 42, pp. 149–160.
- [56] T. Fruchterman and E. Reingold, "Graph drawing by force-directed placement," *Software-Practice Exp.*, vol. 21, pp. 1129–1164, 1991.
- [57] G. D. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis, *Graph Drawing: Algorithms for the Visualization of Graphs*. Englewood Cliffs, NJ: Prentice-Hall, 1999.
- [58] M. Chalmers, "A linear iteration time layout algorithm for visualizing high-dimensional data," in *Proc. 7th Conf. Visualization*, R. Yagel and G. M. Nielson, Eds., San Francisco/Los Alamitos, CA, 1996, pp. 127–132.
- [59] A. Morrison, G. Ross, and M. Chalmers, "Fast Multidimensional Scaling through Sampling, Springs and Interpolation," *Inf. Visualization*, vol. 2, pp. 68–77, 2003.
- [60] J. X. Li, "Visualization of high-dimensional data with relational perspective map," *Inf. Visualization*, vol. 3, pp. 49–59, 2004.
- [61] A. Frick, A. Ludwig, and H. Mehldau, "A fast adaptive layout algorithm for undirected graphs," in *Proc. Int. Workshop Graph Drawing*, R. Tamassia and I. G. Tollis, Eds., Berlin, Germany, 1994, pp. 388–403.
- [62] N. Quinn and M. Breuer, "A force-directed component placement procedure for printed circuit boards," *IEEE Trans. Circuits Syst.*, vol. CAS-26, no. 6, pp. 377–388, Jun. 1979.
- [63] P. Demartines, "Mesures d'organisation du réseau de Kohonen," presented at the Congrès Satellite du Congrès Européen de Mathématiques: Aspects Théoriques des Réseaux de Neurones, Paris, France, 1992, unpublished.
- [64] S. V. Edwards, B. Fertil, A. Giron, and P. J. Deschavanne, "A genomic schism in birds revealed by phylogenetic analysis of DNA strings," *Syst. Biol.*, vol. 51, pp. 599–613, 2002.
- [65] C. Chapus, C. Dufraigne, S. Edwards, A. Giron, B. Fertil, and P. J. Deschavanne, "Exploration of phylogenetic data using a global sequence analysis method," *BMC Evol. Biol.* vol. 5, no. 63, 2005 [Online]. Available: <http://www.biomedcentral.com/1471-2148/5/63>

- [66] Y. L. Yap, X. W. Zhang, and A. Danchin, "Relationship of SARS-CoV to other pathogenic RNA viruses explored by tetranucleotide usage profiling," *BMC Bioinf.* vol. 4, no. 43, 2003 [Online]. Available: <http://www.biomedcentral.com/1471-2105/4/43>



**Sylvain Lespinats** received the M.S. degree in biomathematics from Denis Diderot University, Paris VII, France, in 2002 and the Ph.D. degree in biomathematics from Pierre and Marie Curie University, Paris VI, France, in 2006.

Currently, he is a Postdoctoral Researcher at INSERM U722, Paris, France. His research interests are focused on the exploration of high-dimensional data with special application to the genomic signatures.



**Michel Verleysen** (S'87–M'92–SM'04) was born in 1965 in Belgium. He received the M.S. and Ph.D. degrees in electrical engineering from the Université catholique de Louvain, Louvain-la Neuve, Belgium, in 1987 and 1992, respectively.

He was an Invited Professor at the Ecole Polytechnique Fédérale de Lausanne (E.P.F.L.), Switzerland, in 1992, at the Université d'Evry Val d'Essonne, France, in 2001, and at the Université Paris I Panthéon-Sorbonne, France, in 2002, 2003, and 2004.


Currently, he is a Research Director of the Belgian Fonds National de la Recherche Scientifique (F.N.R.S.) and Lecturer at the Université catholique de Louvain. His research interests include artificial neural networks, self-organization, time-series forecasting, nonlinear statistics, adaptive signal processing, and biomedical systems.

Dr. Verleysen is the Editor-in-Chief of the *Neural Processing Letters* and an Associate Editor of the IEEE TRANSACTIONS ON NEURAL NETWORKS. He is the Chairman of the European Symposium on Artificial Neural Networks (ESANN) and member of the editorial board and program committee of several journals and conferences on neural networks and learning.



**Alain Giron** received the M.S. degree in artificial intelligence from Pierre et Marie Curie University, Paris VI, France, in 1985 and the Ph.D. degree in systems control from the Compiègne University of Technology, France, in 1990.

Currently, he is a Research Engineer at INSERM unit U678, Paris, France. His research interests include genomic signature and biostatistics.



**Bernard Fertil** received the Ph.D. degree in physics from Paris XI University, Paris, France, in 1975 and in living sciences from Paris VI University, Paris, France, in 1984.

Currently, he is a Research Director at the CNRS institute, heading Group 4 of INSERM unit 678 in Paris, France. He is the author or coauthor of more than 120 scientific papers. His research interests include data mining, image analysis, modeling, with specific applications to medical decision making, bioinformatics, and radiobiology.