

Incomplete Pivoted QR-based Dimensionality Reduction

Amit Bermanis, Aviv Rotbart, Moshe Salhov and Amir Averbuch
School of Computer Science
Tel Aviv University, Tel Aviv 69978, Israel

May 19, 2015

Abstract

High-dimensional big data appears in many research fields such as image recognition, biology and collaborative filtering. Often, the exploration of such data by classic algorithms is encountered with difficulties due to ‘curse of dimensionality’ phenomenon. Therefore, dimensionality reduction methods are applied to the data prior to its analysis. Many of these methods are based on principal components analysis, which is statistically driven, namely they map the data into a low-dimension subspace that preserves significant statistical properties of the high-dimensional data. As a consequence, such methods do not directly address the geometry of the data, reflected by the mutual distances between multidimensional data point. Thus, operations such as classification, anomaly detection or other machine learning tasks may be affected.

This work provides a dictionary-based framework for geometrically driven data analysis that includes dimensionality reduction, out-of-sample extension and anomaly detection. It embeds high-dimensional data in a low-dimensional subspace. This embedding preserves the original high-dimensional geometry of the data up to a user-defined distortion rate. In addition, it identifies a subset of landmark data points that constitute a dictionary for the analyzed dataset. The dictionary enables to have a natural extension of the low-dimensional embedding to out-of-sample data points, which gives rise to a distortion-based criterion for anomaly detection. The suggested method is demonstrated on synthetic and real-world datasets and achieves good results for classification, anomaly detection and out-of-sample tasks.

Key words: linear dimensionality reduction, incomplete pivoted QR, distortion, out-of-sample extension, user-defined distortion, diffusion maps

1 Introduction

Nowadays, continuous sampling of measurements from sensor systems of real-world processes has generate ever-growing datasets. Analysis of data silos is a fundamental task in many scientific and industrial fields, whose goal is to infer significant information from a collection of observations/measurements that illuminate the underlying phenomenon that generates the observed data. The illuminated observations will assist in tasks such as classification, clustering, forecasting and anomaly detection to name some.

High-dimensional big data analysis is of special interest since multidimensional data points usually reside in a lower dimensional subspace of the ambient space. For example, data clustering requires regions of high density that constitute clusters. In high dimensions, a huge number of data points is required to create high density regions and this number grows exponentially with dimension. High-dimensional data appear in two main forms: parametric and

non-parametric. In the former case, every observed data point consists of many parameters where each corresponds to a single dimension. Typically, parametric data is concerned with the geometry of such data. Non-parametric data is typically originated in artificial geometry ascription to the analyzed data that usually encapsulated in a *kernel* matrix. Analysis of this geometry, through the analysis of the associated kernel, can uncover latent data features alas such geometries may be high-dimensional. Either way, for efficient analysis of high-dimensional data, a dimensionality reduction, which preserves the original high-dimensional geometry, is needed.

Principal Component Analysis (PCA) [29] is a statistically driven linear dimensionality reduction method, originally designates for the analysis of parametric data. It acts on the Gram matrix of the data and embeds the data in a low-dimensional space, whose coordinates are the directions of the high variances of the data that are also known as principal components. PCA can be implemented by the application of the singular value decomposition (SVD) [26] to a data matrix. However, this implementation is computationally expensive. Section 1.4 details the essentials of SVD.

Induced by PCA for linear analysis of parametric data, the kernel PCA approach for non-linear analysis of non-parametric data was introduced in [42]. Essentially, if the kernel matrix is semi-positive definite, then it can be treated as a Gram matrix of the data in the high-dimensional space, which is also referred to as *feature space*. Then, PCA is applied to the kernel matrix in order to embed the data in a low-dimensional space. Some examples for kernel methods, among many others, are local linear embedding (LLE) [40], Isomap [44] and Diffusion Maps (DM) [19]. LLE embeds the data in a low-dimensional space whose geometry represents local linear geometries of the high-dimensional data. The low-dimensional geometry, which is produced by Isomap, preserves the geodesic distances in the original data. DM provides a low-dimensional representation of the diffusion geometry of the data as defined in [14, 19].

In this work, an incomplete pivoted QR-based deterministic method for dimensionality reduction is presented. The method is designed to preserve a high-dimensional Euclidean geometry of parametric data up to a user-specified distortion rate, according to the following definition:

Definition 1.1 (μ -distortion). *Let $(\mathcal{H}, \mathbf{m}_{\mathcal{H}})$ and $(\mathcal{L}, \mathbf{m}_{\mathcal{L}})$ be metric spaces, $\mathcal{A} \subset \mathcal{H}$ and $\mu \geq 0$. A map $\mathbf{F} : \mathcal{A} \rightarrow \mathcal{L}$ is called a μ -distortion of \mathcal{A} if $\sup_{x,y \in \mathcal{A}} |\mathbf{m}_{\mathcal{H}}(x,y) - \mathbf{m}_{\mathcal{L}}(\mathbf{F}(x), \mathbf{F}(y))| \leq \mu$. The space \mathcal{L} is referred to as the μ -embedding space of \mathcal{A} .*

The proposed method is dictionary-based, where the dictionary is chosen from the analyzed dataset \mathcal{A} . The method identifies a Euclidean embedding space, which is spanned by the dictionary members, on which the orthogonal projection of the data provides a user-defined distortion of the original high-dimensional dataset. In that sense, our method is geometrically driven as opposed to PCA. Clearly, there is an interplay between the distortion rate and the dimension of the resulted embedding subspace: the smaller μ is the higher is the embedding's dimension and vice-versa. Our method preserves global patterns of the data and trades local geometry for low-dimensional representation since dense regions in the original data (such as clusters) are more sensitive to distortions than sparse regions such as gaps between clusters.

Additionally to dimensionality reduction, we present two strongly related schemes for out-of-sample extension and anomaly detection, which are naturally stem from the proposed method for dimensionality reduction. Thus, the dimensionality reduction phase, followed by an out-of-sample extension and by an anomaly detection, constitutes a complete framework for semi-supervised learning, where the original (in-sample) dataset \mathcal{A} functions as a training set. In this context, the learning phase is reflected in the extraction of a μ -embedding subspace of \mathcal{A} as defined in Definition 1.1. Out-of-sample data points, whose projection on the representative subspace is of low distortion, are classified as normal, while the rest are classified as abnormal. Therefore, the original dataset \mathcal{A} is considered as normal by definition.

To conclude, the contribution of this work is threefold: first, the suggested method identifies landmark data points (dictionary) that represent the data, as opposed to PCA that lacks this property, and therefore in some sense it is less informative. In matrix decomposition terminology, this is known as the Columns Subset Selection (CSS) problem. Secondly, the presented method requires very low storage budget relatively to PCA consumption. In the worst case, its computational complexity is identical to that of PCA computation. Lastly, the proposed out-of-sample extension and anomaly detection constitute natural extensions to the dimensionality reduction phase.

The rest of this paper is organized as follows: in the rest of this section, we review related works (Section 1.1), describe the used notation and our general approach (Sections 1.2 and 1.3, respectively), and discuss the essentials of PCA-based dimensionality reduction (Section 1.4). Section 2 establishes the theory and the technical tools on which our method is based. It presents the robustness of our method to noise and its relation to matrix approximation, as well. Although the proposed method is designated for parametric data analysis, section 3 describes its utilization for Diffusion Maps (DM), which is a non-parametric analysis method. Section 4 presents the experimental results. Finally, Section 5 concludes the paper and discusses future work.

1.1 Related works

Johnson-Lindenstrauss (JL) Lemma [33] constitutes a basis for many random projection based dimensionality reduction methods [34, 31, 43, 1, 2, 18, 6, 45] to name some. In contrast to the absolute bound in Definition 1.1, JL Lemma guarantees a relative distortion bound of the form $(1 - \varepsilon) \cdot \mathbf{m}_{\mathcal{H}}^2(x, y) \leq \mathbf{m}_{\mathcal{L}}^2(\mathbf{F}(x), \mathbf{F}(y)) \leq (1 + \varepsilon) \cdot \mathbf{m}_{\mathcal{H}}^2(x, y)$ with high probability. From data analysis perspective, the absolute bound may be more useful when the data is comprised of dense clusters separated by sparse regions. In this scenario, the absolute bound guarantees embedding such that intra-cluster distances may be distorted but inter-cluster distances are preserved so that the global high-dimensional geometry is preserved in the μ -embedded space. The relative bound, on the other hand, may produce a low-dimensional space in which clusters are becoming too close to each other.

Another significant branch of dimensionality reduction methods, which is strongly related to the presented work, deals with the CSS problem [24, 22, 11, 10, 9, 35] to name a few. Interpolative decomposition (ID) of a matrix was introduced first for operator compression [17]. It is designed to approximate spectrally linear integral operators. A randomized version of ID is presented in [37]. The class of randomized matrix decomposition algorithms is out of the scope of this paper. CUR decomposition [12, 36] of a given matrix A generalizes the ID in the sense that both subsets of columns and rows of the original matrix are selected to form the matrices C and R , respectively, such that $A \approx CUR$ for a low rank matrix U . In data analysis terms, this decomposition enables us to get sampling of significant data points (rows), as well as significant features (columns). Similar to ID and unlike this work, CUR decompositions are designated to spectrally approximate the original matrix A . Randomized versions of CUR also exist - see [47, 23] and the references therein.

1.2 Notation

In the rest of the paper, the following notation are used: for $k \in \mathbb{N}$, $[k] = \{1, \dots, k\}$. I_n is the $n \times n$ unit matrix. The i -th coordinate of a vector $\mathbf{v} \in \mathbb{R}^n$ is denoted by $\mathbf{v}_{(i)}$. The (i, j) -th entry of a matrix A is $A_{(i,j)}$ and its i -th row and j -th column are $A_{(i,:)}$ and $A_{(:,j)}$, respectively. If A is of size $m \times n$ and $\mathcal{I} = \{i_1, \dots, i_k\} \subset [m]$ and $\mathcal{J} = \{j_1, \dots, j_\ell\} \subset [n]$ are two ordered sets, then $A_{(\mathcal{I},:)}$ is the $k \times n$ matrix B for which $B_{(r,:)} = A_{(i_r,:)}$, $r \in [k]$, $A_{(:,\mathcal{J})}$ is the $m \times \ell$ matrix

C for which $C_{(:,r)} = A_{(:,j_r)}$, $r \in [\ell]$ and $A_{(\mathcal{I},\mathcal{J})}$ is the $k \times \ell$ matrix D , whose (p,q) -th entry is $D_{(p,q)} = A_{(i_p,j_q)}$, $p \in [k], q \in [\ell]$. The transposed matrix of A is denoted by A^* and A^\dagger is the Moore-Penrose pseudo-inverse of A . $\boldsymbol{\pi} : [n] \rightarrow [n]$ denotes permutation, and Π is the associated $n \times n$ permutation matrix such that $\Pi_{(i,j)} = 1$ if $\boldsymbol{\pi}(j) = i$, otherwise $\Pi_{(i,j)} = 0$. For a subspace $\mathcal{S} \subset \mathbb{R}^m$, \mathcal{S}^\perp is the complementary perpendicular subspace of \mathcal{S} in \mathbb{R}^m , and $\mathbf{W}_\mathcal{S}, \mathbf{W}_\mathcal{S}^\perp : \mathbb{R}^m \rightarrow \mathbb{R}^m$ are the corresponding orthogonal projections on these subspaces, respectively. Finally, $\|\mathbf{v}\|$ is the standard Euclidean norm of \mathbf{v} , where $\|\mathbf{v}\|_1 = \sum_{i=1}^n |\mathbf{v}_{(i)}|$ is its ℓ_1 norm.

The explored dataset is $\mathcal{A} = \{\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(n)}\} \subset \mathbb{R}^m$ and the associated data matrix is the $m \times n$ matrix A , whose j -th column is $A_{(:,j)} = \mathbf{a}^{(j)}$, $j \in [n]$.

1.3 General approach

Our approach to achieve a low-rate distortion embedding consists of two main steps:

1. Given a nonnegative distortion parameter μ . An s -dimensional 2μ -embedded subspace $\mathcal{S} \subset \mathbb{R}^m$ is identified, for which the orthogonal projection of any $\mathbf{a} \in \mathcal{A}$ results in an energy loss of at most μ , i.e.

$$\|\mathbf{W}_\mathcal{S}^\perp(\mathbf{a}^{(i)})\| \leq \mu, \quad i \in [n]. \quad (1.1)$$

Here, the nonnegative distortion parameter μ is a user defined input. Clearly, s is a non-increasing function of μ .

2. The subspace \mathcal{S} is orthogonally aligned with \mathbb{R}^s to achieve an s -dimensional representation of \mathcal{A} , i.e. $\mathbf{O}_\mathcal{S} : \mathbb{R}^m \rightarrow \mathbb{R}^s$ is an orthogonal transformation that satisfies

$$\|\mathbf{O}_\mathcal{S}\mathbf{v}\| = \|\mathbf{v}\|, \quad \mathbf{v} \in \mathcal{S}. \quad (1.2)$$

Obviously, such an alignment (which is not unique) does not affect the geometry of the projected set \mathcal{A} on \mathcal{S} .

Application of the above two-stage scheme to \mathcal{A} results in a 2μ -distortion as Lemma 1.1 shows.

Lemma 1.1. *Let $\mathcal{S} \subset \mathbb{R}^m$ be an s -dimensional subspace of \mathbb{R}^m that satisfies Step 1 and let $\mathbf{O}_\mathcal{S} : \mathbb{R}^m \rightarrow \mathbb{R}^s$ be an orthogonal transformation that satisfies Step 2. Then, the s -dimensional map $\mathbf{F}_s : \mathbb{R}^m \rightarrow \mathbb{R}^s$*

$$\mathbf{F}_s \triangleq \mathbf{O}_\mathcal{S} \circ \mathbf{W}_\mathcal{S} \quad (1.3)$$

is a 2μ -distortion of \mathcal{A} .

Proof. From Eqs. 1.2 and 1.3 we get $\|\mathbf{F}_s(\mathbf{v})\| = \|\mathbf{O}_\mathcal{S} \circ \mathbf{W}_\mathcal{S}(\mathbf{v})\| = \|\mathbf{W}_\mathcal{S}(\mathbf{v})\|$. Since $\mathbf{W}_\mathcal{S}$ is orthogonal projection, then $\|\mathbf{W}_\mathcal{S}(\mathbf{v})\| \leq \|\mathbf{v}\|$. Thus, $0 \leq \|\mathbf{v}\| - \|\mathbf{F}_s(\mathbf{v})\| = \|\mathbf{v}\| - \|\mathbf{W}_\mathcal{S}(\mathbf{v})\| \leq \|\mathbf{v} - \mathbf{W}_\mathcal{S}(\mathbf{v})\| = \|\mathbf{W}_\mathcal{S}^\perp(\mathbf{v})\|$, $\mathbf{v} \in \mathbb{R}^m$. Substituting $\mathbf{v} = \mathbf{a}^{(i)} - \mathbf{a}^{(j)}$, $i, j \in [n]$, yields $0 \leq \|\mathbf{a}^{(i)} - \mathbf{a}^{(j)}\| - \|\mathbf{F}_s(\mathbf{a}^{(i)}) - \mathbf{F}_s(\mathbf{a}^{(j)})\| \leq \|\mathbf{a}^{(i)} - \mathbf{a}^{(j)} - \mathbf{F}_s(\mathbf{a}^{(i)}) + \mathbf{F}_s(\mathbf{a}^{(j)})\| \leq \|\mathbf{W}_\mathcal{S}^\perp(\mathbf{a}^{(i)})\| + \|\mathbf{W}_\mathcal{S}^\perp(\mathbf{a}^{(j)})\| \leq 2\mu$. The last inequality is due to Eq. 1.1. \square

We stress the fact that our goal is to approximate the *geometry* of the dataset \mathcal{A} rather than its members, therefore, we use \mathbf{F}_s rather than $\mathbf{W}_\mathcal{S}$.

1.4 PCA-based dimensionality reduction

A common practice to achieve dimensionality reduction is based on PCA [29] of the (centered) $m \times n$ data matrix A . This method uses a singular value decomposition (SVD) [26] of the data matrix to detect a set of maximum variance orthogonal directions (singular vectors) in \mathbb{R}^m . Projection of the data onto the ρ most significant directions yields the best ρ -dimensional embedding of the data in the mean square error sense. The computational and storage complexities of SVD are $O(\min\{m, n\} \cdot mn)$ and $O(\min\{m, n\}^2)$, respectively.

Numerical methods for the computation of SVD approximation have attracted a growing interest. Out of many methods, we mention here some central ones. In recent years, randomized algorithms for SVD approximation of large matrices have become popular. We refer the reader to [28] and references therein for a review of such methods. An efficient incremental algorithms for computing a thin SVD that considers ρ components is suggested in [13]. The computational complexity of this method is $O(\rho nm)$ for $\rho \leq \min\{m, n\}$. The incremental nature of the algorithm makes it suitable for analysis of dynamic data, where rows/columns are dynamically added and subtracted from the data matrix. Another interesting approach, which reduces the SVD computational cost, is to use matrix sparsification by zeroing out small values in the data matrix. This widely used approach utilizes a sparse eigensolver such as Lanczos to compute the relevant ρ eigen-components [20]. If ρ is small in comparison to the matrix size, then the computational complexity of Lanczos is $O(\max\{m, n\}^2 \cdot \rho)$ [26]. The storage requirements for Lanczos is $O(m)$. Additionally, Lanczos method can be modified to terminate when the smallest estimated eigenvalue is well approximated and its value is lower than a given threshold. More sparsification approaches are given in [46]. Finally, the Nyström extension method [5] provides an additional technique to reduce the SVD computation cost by using a low rank sketch of the data matrix.

Mathematically, suppose that the rank of A is ρ . Let $A = USV^*$ be the (thin) SVD of A , where U and V are $m \times \rho$ and $n \times \rho$ matrices, respectively, whose columns are orthonormal, and S is a diagonal $\rho \times \rho$ matrix, whose diagonal elements are ordered decreasingly $s_1 \geq \dots \geq s_\rho \geq 0$. The columns of U and V are referred to as the left and right singular vectors of A respectively, and the diagonal elements of S as its singular values. Then, for any $k \in [\rho]$ we have $\|A - A_k\| \leq \|A - B\|$ for any orthogonally invariant matrix norm and any $m \times n$ matrix B of rank k or less, where A_k is the k -SVD of A , i.e. $A_k = U_{(:, [k])} S_{([k], [k])} (V_{(:, [k])})^*$. Let \mathcal{U} be the subspace spanned by $U_{(:, [k])}$'s columns, then the k -dimensional embedding $\mathbf{F}_k : \mathbb{R}^m \rightarrow \mathbb{R}^k$, which is defined by $\mathbf{F}_k(\mathbf{v}) \triangleq (U_{(:, [k])})^* \mathbf{v}$, is a composition of the orthogonal map $\mathbf{O}_{\mathcal{U}} : \mathbb{R}^m \rightarrow \mathbb{R}^k$, $\mathbf{O}_{\mathcal{U}}(\mathbf{v}) \triangleq (U_{(:, [k])})^* \mathbf{v}$ and the orthogonal projection $\mathbf{W}_{\mathcal{U}}$, i.e. $\mathbf{F}_k = \mathbf{O}_{\mathcal{U}} \circ \mathbf{W}_{\mathcal{U}}$ (see Eq. 1.3). Lemma 1.2 quantifies the distortion rate of \mathbf{F}_k , applied to \mathcal{A} , with respect to the spectrum of A , which is encapsulated in S .

Lemma 1.2. *The k -dimensional embedding \mathbf{F}_k is a $2s_{k+1}$ -distortion of \mathcal{A} .*

Proof. From the triangular inequality we have $\|A_{(:, i)} - A_{(:, j)}\| \leq \|\mathbf{W}_{\mathcal{U}}^\perp(A_{(:, i)})\| + \|\mathbf{W}_{\mathcal{U}}^\perp(A_{(:, j)})\| + \|\mathbf{W}_{\mathcal{U}}(A_{(:, i)}) - \mathbf{W}_{\mathcal{U}}(A_{(:, j)})\|$. Since $\mathbf{W}_{\mathcal{U}}^\perp = I - \mathbf{W}_{\mathcal{U}}$, and due to A 's SVD, we have $\|\mathbf{W}_{\mathcal{U}}(A_{(:, i)})\| < s_{k+1}$ and $\|\mathbf{W}_{\mathcal{U}}(A_{(:, j)})\| < s_{k+1}$. Moreover, since \mathbf{F}_k is an orthogonal map, we have $\|\mathbf{F}_k(A_{(:, i)} - A_{(:, j)})\| \leq \|A_{(:, i)} - A_{(:, j)}\|$. Thus, $|\|A_{(:, i)} - A_{(:, j)}\| - \|\mathbf{F}_k(A_{(:, i)} - A_{(:, j)})\|| \leq 2s_{k+1}$. \square

A particular case of Lemma 1.2 is when $k = \rho$. Then, \mathbf{F}_k embeds \mathcal{A} accurately in \mathbb{R}^ρ . The computational and storage complexities of the thin SVD are $O(\rho mn)$ and $O(\max\{m, n\}^2)$, respectively. In addition, the principal subspace \mathcal{U} , on which the data is projected, is a mixture of the entire columns set of A which, in terms of data analysis, may be less informative than a dictionary-based subspace.

2 Incomplete Pivoted QR-based Data Analysis

In this section, a QR-based method for data sampling and dimensionality reduction is suggested, as well as consequent out-of-sample and anomaly detection schemes. The method is geometrically driven in the sense that a low-dimensional approximation is constructed to constitute a user-defined distortion of the high-dimensional dataset \mathcal{A} . This is accomplished by using an incomplete pivoted QR decomposition of the data matrix A that is described in section 2.3.

The suggested method incrementally and simultaneously constructs a low-dimensional subspace and projects the data on it. The basis elements for the constructed subspace are chosen from \mathcal{A} . Thus, this method also identifies a subset of representative landmark data points according to the user-defined distortion parameter. The landmarks subset $\mathcal{D} \subset \mathcal{A}$ is referred to as *dictionary*. The dictionary enables both an efficient out-of-sample extension and anomaly detection for any data point $\mathbf{x} \in \mathbb{R}^m \setminus \mathcal{A}$. In this context, \mathcal{A} is referred as a *training set*, and each member is considered as normal. The out-of-sample extension is based only on the geometrical relations between \mathbf{x} and the dictionary members as described in Section 2.5.

Both the computational and the storage costs of the proposed method depend on the dimension of the embedded space. In the worst case, where the dictionary consists of the whole data, these complexities are identical to the corresponding complexities of the thin SVD of the associated data matrix A . Moreover, since the proposed algorithm neither uses the powers of AA^* nor A^*A , as opposed to classical algorithms for SVD computations [26], there is no necessity to store A in the RAM.

There are several methods for practical computation of QR decomposition. Householder [30], Givens rotations [25] and Gram-Schmidt or modified Gram-Schmidt [39] are some typical methods. In [32], an incomplete Gram-Schmidt and incomplete Givens transform are utilized to find an incomplete QR decomposition. Another relevant approach is the Rank Revealing QR (RRQR) method [15]. The RRQR can be used for matrix approximation by proper manipulation of the QR output [16]. The proposed pivoted incomplete QR algorithm is one of many methods to compute a partial orthogonal decomposition [4, 38]. Yet, the theoretical basis for our method is valid for any other version of pivoted incomplete QR algorithm.

Robustness of the proposed method to noise is presented in Section 2.2.1 and the resulted matrix approximation is proved in Section 2.2.2.

2.1 Mathematical preliminaries

QR factorization with columns pivoting [26] of an $m \times n$ matrix A of rank ρ is

$$A\Pi = QR, \quad (2.1)$$

where Π is an $n \times n$ permutation matrix, Q is an $m \times \rho$ matrix whose columns constitute an orthonormal basis for the columns space of A and R is a $\rho \times n$ upper diagonal matrix. This decomposition represents the Gram-Schmidt process applied to A 's columns one-by-one due the order determined by Π . Therefore, for any $k \in [\rho]$ we have

$$\mathcal{A}_{\pi([k])} = \mathcal{Q}_{[k]}, \quad (2.2)$$

and

$$Q_{(:,k)} = \mathbf{W}_{\pi([k-1])}^\perp (A_{(:,\pi(k))}) / \|\mathbf{W}_{\pi([k-1])}^\perp (A_{(:,\pi(k))})\|, \quad (2.3)$$

where $\mathcal{A}_{\mathcal{I}}$ and $\mathcal{Q}_{\mathcal{I}}$ are the subspace spanned by the columns of $A_{(:,\mathcal{I})}$ and $Q_{(:,\mathcal{I})}$, respectively, and $\mathbf{W}_{\pi([k])} : \mathbb{R}^m \rightarrow \mathcal{A}_{\pi([k])}$ is the orthogonal projection on $\mathcal{A}_{\pi([k])}$. Equation 2.2 suggests that for any $k \in [\rho]$

$$\mathbf{W}_{\pi([k])}(\mathbf{v}) = Q_{(:,[k])}(Q_{(:,[k])})^* \mathbf{v}, \quad \mathbf{W}_{\pi([k])}^\perp(\mathbf{v}) = \mathbf{v} - \mathbf{W}_{\pi([k])}(\mathbf{v}), \quad \mathbf{v} \in \mathbb{R}^m. \quad (2.4)$$

The presented dimensionality reduction method is based on the incomplete pivoted QR decomposition of the data matrix A . The criteria for pivoting and incompleteness are based on Lemma 2.1:

Lemma 2.1. *Consider Eq. 2.1. Then, for any $k \in [\rho]$,*

$$R_{(k,k)} = \left\| \mathbf{W}_{\pi([k-1])}^\perp(A_{(:,\pi(k))}) \right\|. \quad (2.5)$$

Proof. Since $\mathbf{W}_{\pi([k-1])}^\perp$ is an orthogonal projection, then $(\mathbf{W}_{\pi([k-1])}^\perp)^* \mathbf{W}_{\pi([k-1])}^\perp = \mathbf{W}_{\pi([k-1])}^\perp$. Therefore,

$$\begin{aligned} \left\| \mathbf{W}_{\pi([k-1])}^\perp(A_{(:,\pi(k))}) \right\|^2 &= (\mathbf{W}_{\pi([k-1])}^\perp(A_{(:,\pi(k))}))^* \mathbf{W}_{\pi([k-1])}^\perp(A_{(:,\pi(k))}) \\ &= (A_{(:,\pi(k))})^* \mathbf{W}_{\pi([k-1])}^\perp(A_{(:,\pi(k))}) \\ &= \left\| \mathbf{W}_{\pi([k-1])}^\perp(A_{(:,\pi(k))}) \right\| (A_{(:,\pi(k))})^* Q_{(:,k]} \\ &= \left\| \mathbf{W}_{\pi([k-1])}^\perp(A_{(:,\pi(k))}) \right\| R_{(k,k)}. \end{aligned}$$

□

Lemma 2.2 stresses the recursive relations between Q 's columns. This relation will be used in Section 2.5.

Lemma 2.2. *For any $k \in [\rho]$, $Q_{(:,k)} = (A_{(:,\pi(k))} - \sum_{i=1}^{k-1} R_{(i,k)} Q_{(:,i)}) / R_{(k,k)}$.*

Proof. According to Eqs. 2.3 and 2.4, and since $R_{(i,j)} = (Q_{(:,i)})^* A_{(:,\pi(j))}$ for any $i \in [\rho], j \in [n]$, we have $\mathbf{W}_{\pi([k-1])}^\perp(A_{(:,\pi(k))}) = \sum_{i=1}^{k-1} R_{(i,k)} Q_{(:,i)}$. Therefore, due to Eq. 2.5, the lemma is proved. □

2.2 Incomplete pivoted QR-based dimensionality reduction - theoretical background

The geometry of A 's (permuted) columns is isomorphic to the geometry of the R 's columns, i.e. for any $i, j \in [n]$, $(A_{(:,\pi(i))})^* A_{(:,\pi(j))} = (R_{(:,i)})^* R_{(:,j)}$. Thus, the upper triangularity of R suggests to embed the dataset \mathcal{A} by an incomplete (truncated) version of R 's rows. Mathematically, following Eq. 1.3, if we set $\mathcal{S} = \mathcal{A}_{\pi([s])}$ and the orthogonal map $\mathbf{O}_{\mathcal{S}} : \mathbb{R}^m \rightarrow \mathbb{R}^s$ is

$$\mathbf{O}_{\mathcal{S}}(\mathbf{v}) \triangleq (Q_{(:,[s])})^* \mathbf{v}, \quad (2.6)$$

then, due to the orthogonality of Q 's columns and Eqs. 2.1, 2.2 and 2.4, the s -dimensional embedding from Eq. 1.3 becomes

$$\mathbf{F}_s(\mathbf{v}) = (Q_{(:,[s])})^* \mathbf{v} \quad (2.7)$$

and specifically,

$$\mathbf{F}_s(A_{(:,\pi(i))}) = R_{([s],i)}, \quad i \in [n]. \quad (2.8)$$

Notice that Eq. 1.2 is satisfied by $\mathbf{O}_{\mathcal{S}}$ from Eq. 2.6. Although this specific choice for $\mathbf{O}_{\mathcal{S}}$ yields $\mathbf{O}_{\mathcal{S}} = \mathbf{F}_s$, this is not always the case since, as aforementioned, $\mathbf{O}_{\mathcal{S}}$ is not unique. For example, in Section 2.3, a different choice of $\mathbf{O}_{\mathcal{S}}$ is presented. The incompleteness of the discussed QR decomposition is reflected in Eq. 2.8, where the s -dimensional embedding is defined via only a partial set of R 's rows. According to the triangularity of R , the geometry of such an embedding is exact on the basis elements of \mathcal{S} which are

$$\mathcal{D} \triangleq \{A_{(:,\pi(1))}, \dots, A_{(:,\pi(s))}\}. \quad (2.9)$$

This set is referred to as the *dictionary* of \mathcal{A} and its elements are referred to as *pivots*. As Lemma 1.1 suggests, a careful choice of Π might result in a low-dimensional distortion \mathbf{F}_s of \mathcal{A} . An algorithm for such a choice is presented in Section 2.3.

We conclude this section with an example that demonstrates the permutation's significance. Let A be the following matrix:

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 20 \end{pmatrix}.$$

Then, in order to achieve a 3-distortion for $\Pi = I_7$, we must have $s = \rho = 7$. On the other hand, one can verify that for any permutation that satisfies $\boldsymbol{\pi}([3]) = \{1, 4, 7\}$, $\|\mathbf{W}_{\boldsymbol{\pi}([3])}^\perp(A_{(:,i)})\| < 1.5$, $i \in [7]$, which according to Lemma 1.1 is a sufficient condition for a 3-dimensional embedding of \mathcal{A} , with distortion rate bounded by 3. Consequently, F_s is an s -dimensional 2μ -distortion of \mathcal{A} , with $s = 3$ and $\mu = 1.5$.

Proposition 2.3, which is a rephrased version of Lemma 1.1 in terms of the pivoted incomplete QR, concludes the above discussion:

Proposition 2.3. *Let $\mu > 0$. If there exist a permutation $\boldsymbol{\pi}$ and $s \in [\rho]$ for which $\|\mathbf{W}_{\boldsymbol{\pi}([s])}^\perp(A_{(:,i)})\| < \mu$ for any $i \in [n]$, then \mathbf{F}_s from Eq. 2.7 is an s -dimensional 2μ -distortion of \mathcal{A} .*

2.2.1 Stability to noise

In real-life, data may be noisy. Therefore, instead of analyzing clean data that is stored in A , a noisy version $\tilde{A} = A + N$ is analyzed where the matrix N represents an additive noise. Proposition 2.4 shows that a distortion by a noisy data is also a distortion of the original clean data, where the error originated by noise is additive.

Proposition 2.4. *Let A be an $m \times n$ data matrix and let $\tilde{A} = A + N$, where N is a noise matrix of the same size of A . Assume that $\|N\| \leq \eta$ for some $\eta \geq 0$, and let $\tilde{\mathbf{F}}_s$ be a 2μ -distortion of \tilde{A} 's columns as defined in Eq. 2.8. Then, $\tilde{\mathbf{F}}_s$ is a $2(\mu + \eta)$ -distortion of A 's columns.*

Proof. Let $\tilde{\mathbf{F}}_s$ be the map defined in Eq. 1.3 with the corresponding elements $\tilde{\mathcal{S}}$, $\mathbf{O}_{\tilde{\mathcal{S}}}$ and $\mathbf{W}_{\tilde{\mathcal{S}}}$. Then, we have

$$\begin{aligned} \left| \|A_{(:,i)} - A_{(:,j)}\| - \left\| \tilde{\mathbf{F}}_s(\tilde{A}_{(:,i)}) - \tilde{\mathbf{F}}_s(\tilde{A}_{(:,j)}) \right\| \right| &= \left| \|A_{(:,i)} - A_{(:,j)}\| - \left\| \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,i)}) - \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,j)}) \right\| \right| \\ &\leq \left\| A_{(:,i)} - A_{(:,j)} - \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,i)}) + \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,j)}) \right\| \\ &\leq \left\| A_{(:,i)} - \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,i)}) \right\| + \left\| A_{(:,j)} - \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,j)}) \right\| \\ &\leq \left\| A_{(:,i)} - \tilde{A}_{(:,i)} \right\| + \left\| \tilde{A}_{(:,i)} - \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,i)}) \right\| \\ &\quad + \left\| A_{(:,j)} - \tilde{A}_{(:,j)} \right\| + \left\| \tilde{A}_{(:,j)} - \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,j)}) \right\| \\ &= \|N_{(:,i)}\| + \left\| \tilde{A}_{(:,i)} - \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,i)}) \right\| \\ &\quad + \|N_{(:,j)}\| + \left\| \tilde{A}_{(:,j)} - \mathbf{W}_{\tilde{\mathcal{S}}}(\tilde{A}_{(:,j)}) \right\| \\ &\leq 2(\mu + \eta). \end{aligned}$$

The first equality is due to Eqs. 1.2 and 1.3 and the last inequality is due to the fact that $\|N_{(:,i)}\| \leq \|N\|$, the proposition's assumption and Eq. 1.1. \square

2.2.2 Matrix approximation error

The notions of low-rank matrix approximation and geometry-preserving dimensionality reduction are different but related. While the s -SVD of a data matrix enables an s -dimensional embedding of the associated data (as was established in Lemma 1.2), it can be shown that the incomplete QR factorization $Q_{(:,[s])}R_{([s],:)}$ of the pivoted version of A (see Eq. 2.1) can be used to form a low-rank approximation of A . The operator norm of an arbitrary matrix M , denoted by $\|M\|_2$, is equal to its maximal singular value. Thus, as was explained in Section 1.4, it measures the maximal linear trend of the data stored in its columns (or rows). Therefore, the operator norm of the difference of two matrices measures the strength of the maximal linear trend of the associated error. Proposition 2.5 provides a bound for the approximation error of $A\Pi$ by its incomplete QR factorization. Its proof uses the Frobenius norm $\|M\|_F \triangleq \sqrt{\sum_{i,j} M_{(i,j)}^2}$ of a matrix and the norms inequality $\|M\|_2 \leq \|M\|_F$ for any matrix M .

Proposition 2.5. *Let μ , π , s and A satisfy the condition of Proposition 2.3, then $\|A\Pi - Q_{(:,[s])}R_{([s],:)}\|_\eta \leq \mu\sqrt{\rho - s}$ for $\eta \in \{2, F\}$.*

Proof. Following Eqs. 2.1 and 2.4, and according to the orthonormality of Q 's columns, we have $\mathbf{W}_{\pi([s])}(A_{(:,i)}) = Q_{(:,[s])}R_{([s],i)}$ for any $i \in [n]$. On the other hand, due to the triangularity of R , $A_{(:,\pi([s]))} = Q_{(:,[s])}R_{([s],[s])}$. Thus, at least s columns from $A - Q_{(:,[s])}R_{(:,[s])}$ are vanishing, and the norms of the rest (mostly) $\rho - s$ are bounded by μ , according to the proposition's assumption. This leads to $\|A - Q_{(:,[s])}R_{(:,[s])}\|_2 \leq \|A - Q_{(:,[s])}R_{(:,[s])}\|_F \leq \mu\sqrt{\rho - s}$. \square

2.3 Incomplete pivoted QR-based (ICPQR) dimensionality reduction: implementation

Algorithm 1 iteratively constructs an incomplete pivoted QR version of the data matrix A to obtain a 2μ -distortion of \mathcal{A} . In its j -th iteration, the algorithm selects the pivot $A_{(:,\pi(j))}$, and projects the dataset \mathcal{A} on $\mathcal{A}_{\pi([j])}$. Based on Step 1 from Section 1.3, the j -th pivot $A_{(:,\pi(j))}$ is chosen to be the element in \mathcal{A} , whose approximation by its orthogonal projection on $\mathcal{A}_{\pi([j-1])}$ is the worst. Thus, the permutation $\pi : [n] \rightarrow [n]$ is determined by the following condition:

$$\pi(j) = \arg \max_{i \in [n] \setminus \pi([j-1])} \|\mathbf{W}_{\pi([j-1])}^\perp(A_{(:,i)})\|. \quad (2.10)$$

Then, the corresponding new column $Q_{(:,j)}$ and row $R_{(j,:)}$ are computed according to Eqs. 2.1, 2.3 and 2.4. Since the columns permutation is updated in every iteration j , the columns of $R_{(:,[j-1])}$ have to be permuted correspondingly. The algorithm terminates when the quantity in Eq. 2.10 is less than μ . Thus s , which is the number of iterations required to provide a 2μ -distortion, is a non-increasing function of μ , bounded from above by ρ that is not known a-priori. When the algorithm ends, the correspondence rule, defined by Eq. 2.8, provides an s -dimensional 2μ -distortion of \mathcal{A} according to Proposition 2.3. For $\mu = 0$, the application of Algorithm 1 to A results in a complete pivoted QR factorization.

Let us make a couple of technical remarks concerning Algorithm 1: 1. In case of limited computational or storage budget, Algorithm 1 can be easily modified to make a limited number of iterations d or, equivalently, to provide a d -dimensional embedding. In this case, the distortion parameter μ is a non-increasing function of d . 2. The dictionary \mathcal{D} is chosen regardless to the data indexing order. This property ensures a relatively sparse dictionary as

Algorithm 1: Incomplete pivoted QR (ICPQR) decomposition

Input : An $m \times n$ data matrix A and a nonnegative distortion parameter μ .

Output: An $m \times s$ matrix Q whose columns are orthonormal, an $s \times n$ upper diagonal matrix R and a permutation π such that the correspondence rule defined by Eq. 2.8 is a 2μ -distortion of A 's columns.

```

1 Initialization: set  $\pi = \text{identity}$ ,  $\Pi = I_n$ ,  $\delta > \mu$ , and  $j = 0$ 
2 while  $\delta > \mu$  do
3   set  $j = j + 1$ 
4   set  $i_j = \arg \max_{i \in [n] \setminus \pi([j-1])} \|\mathbf{W}_{\pi([j-1])}^\perp(A_{(:,i)})\|$  (see Eq. 2.4)
   //consider  $\pi([0]) = \emptyset$ 
5   set  $\Delta = \mathbf{W}_{\pi([j-1])}^\perp(A_{(:,i_j)})$  and  $\delta = \|\Delta\|$ 
6   set  $Q_{(:,j)} = \Delta/\delta$ 
7   switch  $\pi(j) \leftrightarrow \pi(i_j)$ , set  $\Pi = \Pi\Pi_{j \leftrightarrow i_j}$ 
   // $\Pi_{j \leftrightarrow i_j}$  is  $I_n$  with columns  $j, i_j$  swapped
8   set  $R = R\Pi_{j \leftrightarrow i_j}$ 
9   set  $R_{(j,:)} = (Q_{(:,j)})^* \Pi$ 
10 end
11 set  $s = j$ 

```

demonstrated in Section 4.1.4. In order to achieve an optimal¹ coordinates system for the geometry represented by R , an SVD can be utilized. Mathematically, let $R = USV^*$ be the SVD decomposition of R , where U and S are $s \times s$ orthogonal and diagonal matrices, respectively, and V is an $n \times s$ matrix, whose columns are orthonormal. Then, the map $\hat{\mathbf{O}}_{\mathcal{S}} : \mathbb{R}^m \rightarrow \mathbb{R}^s$,

$$\hat{\mathbf{O}}_{\mathcal{S}}(\mathbf{v}) \triangleq (QU)^*\mathbf{v}, \quad \mathbf{v} \in \mathbb{R}^m \quad (2.11)$$

is still isometric on \mathcal{S} , as Step 2 in Section 1.3 requires. Thus, following Eqs. 2.1, 2.8 and 2.11, the map $\hat{\mathbf{F}}_{\mathcal{S}} : \mathbb{R}^m \rightarrow \mathbb{R}^k$ where $\hat{\mathbf{F}}_{\mathcal{S}}(\mathbf{v}) \triangleq \hat{\mathbf{O}}_{\mathcal{S}} \circ \mathbf{W}_{\pi([s])}(\mathbf{v})$, $\mathbf{v} \in \mathbb{R}^m$ is still an s -dimensional 2μ distortion of \mathcal{A} . Moreover, $\hat{\mathbf{F}}_{\mathcal{S}}$ is optimal in the sense that the axes are aligned correspondingly to the variances directions. The computational and storage costs of such an alignment are $O(ns^2)$ and $O(ns)$, respectively. Therefore, the total complexity of Algorithm 1 is not affected by this optional step (see Table 2.4.)

2.4 ICPQR reduced cost

Equation 2.8 suggests that Q is not needed for the low rank embedding of \mathcal{A} by \mathbf{F}_s . In this section, we present a more efficient version of Algorithm 1, by which the results in Section 4 were obtained. The algorithm produces no Q and applies no physical permutations.

Consider Eq. 2.1 with π defined by Eq. 2.10, then

$$A = Q\bar{R}, \quad \bar{R} \triangleq R\Pi^*, \quad (2.12)$$

where \bar{R} is no longer triangular. Algorithm 2 is a translated version of Algorithm 1 to this case, where the permutation Π is absorbed in \bar{R} . The low-dimensional embedding from Eq. 2.8 becomes

$$\mathbf{F}_s(A_{(:,i)}) = \bar{R}_{([s],i)}, \quad i \in [n]. \quad (2.13)$$

¹A coordinates system that is determined by principal components.

For the establishment of Algorithm 2, steps 4 and 9 in Algorithm 1, which are dependent on Q , are modified according to Eq. 2.12 to be $(A_{(:,\pi(i))})^* A_{(:,\pi(j))} = (\bar{R}_{(:,pi(i))})^* \bar{R}_{(:,pi(j))} = \sum_{\ell=1}^{\min\{i,j\}} \bar{R}_{(\ell,\pi(i))} \bar{R}_{(\ell,\pi(j))}$, for any $i, j \in [n]$. The upper limit in the sum is due to the upper triangularity of R . Thus, the following recursive relations between the entries of \bar{R} hold:

$$\bar{R}_{(i,\pi(j))} = \begin{cases} \bar{R}_{(i,\pi(i))}^{-1} \left(u_{ij} - \sum_{\ell=1}^{i-1} \bar{R}_{(\ell,\pi(i))} \bar{R}_{(\ell,\pi(j))} \right) & \text{if } i < j \\ \left(u_{ij} - \sum_{\ell=1}^{i-1} (\bar{R}_{(\ell,\pi(i))})^2 \right)^{1/2} & \text{if } i = j \\ 0 & \text{if } i > j \end{cases}, \quad (2.14)$$

where $u_{ij} \triangleq (A_{(:,\pi(i))})^* A_{(:,\pi(j))}$. Moreover, comparing Eq. 2.5 with Eq. 2.14 yields

$$\| \mathbf{W}_{\pi([i-1])}^\perp (A_{(:,\pi(i))}) \| = \left(u_{ii} - \sum_{\ell=1}^{i-1} (\bar{R}_{(\ell,\pi(i))})^2 \right)^{1/2}.$$

Thus, the pivoting criterion from Eq. 2.10 becomes

$$\pi(j) = \arg \max_{i \in [n] \setminus [j-1]} \left(u_{ii} - \sum_{\ell=1}^{i-1} (\bar{R}_{(\ell,\pi(i))})^2 \right)^{1/2}.$$

Algorithm 2: Incomplete pivoted Q-less QR decomposition

Input : An $m \times n$ matrix A and a nonnegative distortion parameter μ .

Output: An $s \times n$ matrix \bar{R} and a permutation π , for which the embedding defined by Eq. 2.13 is a 2μ -distortion of A 's columns.

- 1 Initialization: set $\pi = \text{identity}$, $j = 0$, $\delta > \mu^2$, $\mathbf{y} = \mathbf{0}_n$ (the all zeros vector of length n), and $\mathbf{z} \in \mathbb{R}^n$, for which $\mathbf{z}_{(i)} = \|A_{(:,i)}\|^2$, $i \in [n]$
 - 2 **while** $\delta \geq \mu^2$ **do**
 - 3 set $j = j + 1$
 - 4 set $i_j = \arg \max_{i \in [n] \setminus [j-1]} (\mathbf{z}_{(\pi(i))} - \mathbf{y}_{(\pi(i))})$
 - 5 switch $\pi(j) \leftrightarrow \pi(i_j)$
 - 6 set $\delta = \mathbf{z}_{(\pi(j))} - \mathbf{y}_{(\pi(j))}$
 - 7 for every $i \in [j-1]$ set $\bar{R}_{(j,\pi(i))} = 0$
 - 8 set $\bar{R}_{(j,\pi(j))} = \delta^{1/2}$
 - 9 **for** $i \in [n] \setminus [j]$ **do**
 - 10 set $u_{ji} = (A_{(:,\pi(j))})^* A_{(:,\pi(i))}$
 - 11 $\bar{R}_{(j,\pi(i))} = (u_{ji} - \sum_{\ell=1}^{j-1} \bar{R}_{(\ell,\pi(j))} \bar{R}_{(\ell,\pi(i))}) / \bar{R}_{(j,\pi(j))}$
 - 12 set $\mathbf{y}_{(\pi(i))} = \mathbf{y}_{(\pi(i))} + \bar{R}_{(j,\pi(i))}$
 - 13 **end**
 - 14 **end**
 - 15 set $s = j$
-

The resulted dictionary is the set \mathcal{D} as defined in Eq. 2.9.

Table 2.4 presents the computational and storage complexities of Algorithm 2. The storage of the input matrix A was not taken into account since in the nature of Algorithm 2 there is no need to have its complete storage. For example, the relevant rows and columns of A can be individually computed at each iteration. Therefore, the total storage complexity is smaller than the required storage of the SVD, which is $O(\max\{m, n\}^2)$. The computational complexity of Algorithm 2 depends on μ . In the worst case, when $\mu = 0$ and $s = \rho$, the complexity of Algorithm 2 equals to the complexity of the thin rank- ρ SVD. Otherwise, Algorithm 2 is more efficient than SVD.

Step	Operations	Storage (not including storage of A)
1	$O(mn)$	$O(n)$
4	$O(ns - s^2)$	$O(1)$
7	$O(s^2)$	$O(s^2)$
10	$O(m(ns - s^2))$	$O(m)$
11	$O(s(ns - s^2))$	$O(ns - s^2)$
12	$O(ns - s^2)$	$O(1)$
Total:	$O(mns)$	$O(ns)$

Table 2.1: Computational and storage complexities of Algorithm 2.

2.5 Out-of-sample extension and anomaly detection algorithms

Two fundamental questions may be naturally asked for an out-of-sample data point $\mathbf{x} \in \mathbb{R}^m \setminus \mathcal{A}$: first, is it normal related to the training dataset \mathcal{A} ? and secondly, if it is, then how can the produced low rank embedding be extended to this data point? This section addresses these two questions. Since \mathbf{F}_s (Eq. 2.7) is defined for the entire space of \mathbb{R}^m , it is used to define an out-of-sample extension for the embedding from Eq. 2.13 of any $\mathbf{x} \in \mathbb{R}^m$ and, based on this, to detect anomalies.

2.5.1 Out-of-sample extension

As mentioned above, an out-of-sample extension of the embedding from Eq. 2.13 is defined by Eq. 2.7 for the entire \mathbb{R}^m . As discussed in Section 2.4, since Algorithm 2 produces no Q , \mathbf{F}_s cannot be directly applied to an out-of-sample point $\mathbf{x} \in \mathbb{R}^m$. Therefore, a Q -less tool for calculating the out-of-sample extension, as defined in Eq. 2.7, is provided in Algorithm 3 that is based on the following proposition:

Proposition 2.6. *Let \bar{R} be the $s \times n$ matrix produced by Algorithm 2, $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{f} = \mathbf{F}_s(\mathbf{x}) \in \mathbb{R}^s$ as defined in Eq. 2.7. Then, the following recursive relation holds for the coordinates of \mathbf{f} :*

$$\mathbf{f}_{(j)} = (\bar{R}_{(j,\pi(j))})^{-1} (A_{(:,\pi(j))})^* \mathbf{x} - \sum_{i=1}^{j-1} \bar{R}_{(i,\pi(j))} \mathbf{f}_{(i)}, \quad j \in [s].$$

Proof. According to Eq. 2.7, $\mathbf{f}_{(j)} = (Q_{(:,j)})^* \mathbf{x}$. Thus, due to the recursive relations of Q 's columns, as presented in Lemma 2.2, we have

$$\begin{aligned} \mathbf{f}_{(j)} &= (R_{(j,j)})^{-1} ((A_{(:,\pi(j))})^* \mathbf{x} - \sum_{i=1}^{j-1} R_{(i,j)} (Q_{(:,i)})^* \mathbf{x}) \\ &= (R_{(j,j)})^{-1} ((A_{(:,\pi(j))})^* \mathbf{x} - \sum_{i=1}^{j-1} R_{(i,j)} \mathbf{f}_{(i)}) \\ &= (\bar{R}_{(j,\pi(j))})^{-1} ((A_{(:,\pi(j))})^* \mathbf{x} - \sum_{i=1}^{j-1} \bar{R}_{(i,\pi(j))} \mathbf{f}_{(i)}), \end{aligned}$$

where the last equality is due to Eq. 2.12. □

Since the out-of-sample extension of a data point $\mathbf{x} \in \mathbb{R}^m$ is its orthogonal projection on the s -dimensional dictionary subspace \mathcal{S} , the only required information are the geometric relations between \mathbf{x} and the elements of the dictionary \mathcal{D} (Eq. 2.9), as Proposition 2.6 shows. Algorithm 3 summarized the above.

Algorithm 3: Out-of-sample extension for incomplete pivoted Q-less QR decomposition

Input : Dictionary $\mathcal{D} = \{\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(s)}\}$ (see Eq. 2.9) and $s \times n$ matrix \bar{R} , which are the outputs of Algorithm 2, and vector $\mathbf{v} \in \mathbb{R}^m$.

Output: $\mathbf{F}_s(\mathbf{v})$ as defined by Eq. 2.7.

- 1 Initialization: set $\mathbf{f} \in \mathbb{R}^s$ to be vector of all-zeros, except of the first coordinate,
 $\mathbf{f}_{(1)} = (\bar{R}_{(1, \pi(1))})^{-1}(\mathbf{b}^{(1)})^* \cdot \mathbf{v}$.
 - 2 **for** $j = 2 : s$ **do**
 - 3 | set the j -th coordinate of \mathbf{f} to be $\mathbf{f}_{(j)} = (\bar{R}_{(j, \pi(j))})^{-1}(\mathbf{b}^{(j)})^* \cdot \mathbf{v} - \sum_{i=1}^{j-1} \bar{R}_{(i, \pi(j))} \mathbf{f}_{(i)}$
 - 4 **end**
 - 5 set $\mathbf{F}_s(\mathbf{v}) = \mathbf{f}$
-

2.5.2 Anomaly detection

The 2μ -embedding subspace $\mathcal{S} \subset \mathbb{R}^m$ satisfies $\boldsymbol{\mu}(\mathbf{a}) \leq \mu$ for any $\mathbf{a} \in \mathcal{A}$, where the distortion rate function $\boldsymbol{\mu} : \mathbb{R}^m \rightarrow \mathbb{R}$ is defined to be $\boldsymbol{\mu}(\mathbf{x}) \triangleq \|\mathbf{W}_{\mathcal{S}}^{\perp}(\mathbf{x})\|$ (see discussion in Section 1.3). Once the out-of-sample extension of $\mathbf{F}_s(\mathbf{x})$ was computed by Algorithm 3, the distortion rate of $\mathbf{x} \in \mathbb{R}^m$ can be easily calculated by $\boldsymbol{\mu}^2(\mathbf{x}) = \|\mathbf{x}\|^2 - \|\mathbf{W}_{\mathcal{S}}(\mathbf{x})\|^2 = \|\mathbf{x}\|^2 - \|\mathbf{O}_{\mathcal{S}}\mathbf{W}_{\mathcal{S}}(\mathbf{x})\|^2 = \|\mathbf{x}\|^2 - \|\mathbf{F}_s(\mathbf{x})\|^2$. The first equality is due to the fact that $\mathbf{W}_{\mathcal{S}}$ is orthogonal projection. The second is due to Step 2 in Section 1.3 that suggests that $\|\mathbf{O}_{\mathcal{S}} \circ \mathbf{W}_{\mathcal{S}}(\mathbf{x})\| = \|\mathbf{W}_{\mathcal{S}}(\mathbf{x})\|$ for any $\mathbf{x} \in \mathbb{R}^m$ and the last equality is due to Eq. 2.7. Consequently, we rephrase the distortion rate function to be

$$\boldsymbol{\mu}(\mathbf{x}) \triangleq (\|\mathbf{x}\|^2 - \|\mathbf{F}_s(\mathbf{x})\|^2)^{1/2}. \quad (2.15)$$

Normality of data points by definition 2.1 will serve us in two main forms in the rest of the paper.

Definition 2.1 (κ -normality). *Let $\mathbf{F}_s : \mathcal{A} \rightarrow \mathbb{R}^m$ be an s -dimensional embedding of $\mathcal{A} \subset \mathbb{R}^m$, computed by Algorithm 2, and let $\mathbf{F}_s(\mathbf{x}) \in \mathbb{R}^s$ be its extension to $\mathbf{x} \in \mathbb{R}^m$, produced by Algorithm 3. Then, \mathbf{x} is classified as a κ -normal point relatively to \mathcal{A} if $\boldsymbol{\mu}(\mathbf{x}) \leq \kappa$. Otherwise, it is classified as κ -abnormal.*

Due to Definition 2.1, all the data points in \mathcal{A} are μ -normal. Nevertheless, μ is not necessarily the minimal κ for which \mathcal{A} is classified as κ -normal, as the most strict κ for which \mathcal{A} is still κ -normal, is $\kappa = \mu_{strict}$, where $\mu_{strict} \triangleq \sup_{\mathbf{a} \in \mathcal{A}} \boldsymbol{\mu}(\mathbf{a})$.

We conclude this section with a definition of two variants of κ -normality as was defined in Definition 2.1.

Definition 2.2 (Normality and strict normality). *Let $\mathbf{F}_s : \mathcal{A} \rightarrow \mathbb{R}^m$ be an s -dimensional embedding of $\mathcal{A} \subset \mathbb{R}^m$ computed by Algorithm 2, with distortion parameter μ , and let $\mathbf{F}_s(\mathbf{x}) \in \mathbb{R}^s$ be its extension to $\mathbf{x} \in \mathbb{R}^m$ produced by Algorithm 3. Then, \mathbf{x} is classified as a normal data point relatively to \mathcal{A} if $\boldsymbol{\mu}(\mathbf{x}) \leq \mu$, and as a strictly normal point if $\boldsymbol{\mu}(\mathbf{x}) \leq \mu_{strict}$. Otherwise, it is classified as a (strictly) abnormal.*

Obviously, all the data points in \mathcal{A} are strictly-normal and any strictly normal data point is also a normal data point.

3 QR-based Diffusion Maps

Although the QR-based dimensionality reduction method, which was presented in Section 2, is designated for parametric data analysis, this section presents a utilization of our method

for the Diffusion Maps (DM) [19], which is a graph Laplacian based method for analysis of nonparametric data, via exploration of a Markov chain defined on the data. It is mainly utilized for clustering and manifold learning. Typically, application of DM involves a kernel PCA, which is computationally prohibitive for large amount of data.

3.1 DM framework: overview

3.1.1 Diffusion geometry

Let $\mathcal{X} = \{x_1, \dots, x_n\}$ be a dataset and let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a symmetric point-wise positive kernel that defines a connected undirected weighted graph over \mathcal{X} . Then, a Markov process over \mathcal{X} can be defined using $n \times n$ row-stochastic transition probabilities matrix

$$P = D^{-1}K, \quad (3.1)$$

where $K_{(i,j)} = k(x_i, x_j)$, $i, j \in [n]$ and D is a diagonal matrix with the diagonal elements $D_{(i,i)} = \mathbf{d}_{(i)}$, where $\mathbf{d}_{(i)} \triangleq \sum_{j=1}^n K_{(i,j)}$, $i \in [n]$. The vector $\mathbf{d} \in \mathbb{R}^n$ is referred to as the *degrees function* or *degrees vector* of the graph. The associated time-homogeneous Markov chain is defined as follows: for any two time points $t, t_0 \in \mathbb{N}$, $\mathbb{P}(x(t+t_0) = x_j | x(t_0) = x_i) = (P^t)_{(i,j)}$. Assuming that the defined Markov chain is aperiodic (for example, if there is $x \in \mathcal{X}$, for which $k(x, x) > 0$), then it has a unique stationary distribution $\hat{\mathbf{d}} \in \mathbb{R}^n$ which is the steady state of the process, i.e. $\hat{\mathbf{d}}_{(j)} = \lim_{t \rightarrow \infty} (P^t)_{(i,j)}$, regardless the initial point x_i . This steady state is the probability distribution resulted from ℓ_1 normalization of the degrees function \mathbf{d} , i.e.,

$$\hat{\mathbf{d}} = \mathbf{d} / \|\mathbf{d}\|_1. \quad (3.2)$$

The *diffusion distance* in time $t \in \mathbb{N}$ is defined by the metric $\mathbf{D}^{(t)} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ such that

$$\mathbf{D}^{(t)}(x_i, x_j) \triangleq \left\| (P^t)_{(i,:)} - (P^t)_{(j,:)} \right\|_{\ell^2(\hat{\mathbf{d}}^{-1})}, \quad i, j \in [n]. \quad (3.3)$$

By definition, $(P^t)_{(i,:)}$ is the probability distribution over \mathcal{X} after t time steps, where the initial state is x_i . Therefore, the diffusion distance from Eq. 3.3 measures the difference between two propagations along t time steps, one originated in x_i and the other in x_j . Weighing the metric by the inverse of the steady state results in ascribing high weight for similar probabilities on rare states and vice versa.

Due to the above interpretation, the diffusion distances are naturally utilized for multiscale clustering since they uncover the connectivity properties of the graph across time. In [14, 19], it was proved that under some conditions, if \mathcal{X} is sampled from a low intrinsic dimensional manifold then, as n tends to infinity, the Markov chain converges to a diffusion process over that manifold.

3.1.2 Diffusion maps - low rank representation of the diffusion geometry

Diffusion maps [19] are a family of Euclidean representations of the diffusion geometry of \mathcal{X} in different time steps, where the Euclidean distances approximate the diffusion distances in Eq. 3.3. Let $G^{(t)}$ be the $n \times n$ matrix defined by

$$G^{(t)} \triangleq \|\mathbf{d}\|_1^{1/2} D^{-1/2} (P^*)^t, \quad t \in \mathbb{N}. \quad (3.4)$$

Then, due to Eqs. 3.1-3.3, the Euclidean n -dimensional geometry of $G^{(t)}$'s column is isomorphic to the diffusion geometry of the associated data points i.e.,

$$\mathbf{D}^{(t)}(x_i, x_j) = \left\| G^{(t)}_{(:,i)} - G^{(t)}_{(:,j)} \right\|, \quad i, j \in [n]. \quad (3.5)$$

Although embedding of the dataset \mathcal{X} in \mathbb{R}^n by the columns of $G^{(t)}$ preserves the diffusion geometry, it may be ineffective for large n as was explained in Section 1. Therefore, a dimensionality reduction is required.

Since the transition probabilities matrix P (Eq. 3.1) is conjugated to the symmetric matrix $M \triangleq D^{-1/2}KD^{-1/2}$ via the relation $P = D^{-1/2}MD^{1/2}$, P has a complete real eigen-system. Moreover, due to Gershgorin's circle theorem [26] and the fact that P is stochastic, all its eigenvalues are lying in the interval $(-1, 1]$ (the exclusion of -1 is due to the assumption that the chain is aperiodic). Let

$$M = USU^* \quad (3.6)$$

be the eigen-decomposition of M , where U is an orthogonal $n \times n$ matrix and S is a diagonal $n \times n$ matrix, whose diagonal elements s_i are ordered decreasingly due to their modulus $1 = s_1 > |s_2| \geq \dots \geq |s_n| \geq 0$. The first inequality is due to the assumption that the graph is connected. Thus, following Eq. 3.4

$$G^{(t)} = (M^*)^t D^{-1/2} = US^t U^* D^{-1/2}. \quad (3.7)$$

Therefore, due to the orthogonality of U and according to Eq. 3.5

$$\mathbf{D}^{(t)}(x_i, x_j) = \|\mathbf{d}\|_1^{1/2} \|S^t U^* D^{-1/2}(\mathbf{e}^{(i)} - \mathbf{e}^{(j)})\|, \quad i, j \in [n],$$

where $\mathbf{e}^{(i)}$ denotes the i -th standard unit vector in \mathbb{R}^n . As a consequence, the diffusion maps $\Psi^{(t)} : \mathcal{X} \rightarrow \mathbb{R}^n$ are defined as follows:

$$\begin{aligned} \Psi^{(t)}(x_i) &\triangleq \|\mathbf{d}\|_1^{1/2} U^* G^{(t)} \mathbf{e}^{(i)} \\ &= \|\mathbf{d}\|_1^{1/2} S^t U^* D^{-1/2} \mathbf{e}^{(i)} \\ &= \|\mathbf{d}\|_1^{1/2} \mathbf{d}_{(i)}^{-1/2} S^t (U_{(i,:)}^*) \\ &= \hat{\mathbf{d}}_{(i)}^{-1/2} [s_1^t U_{(i,1)}, \dots, s_n^t U_{(i,n)}]^*. \end{aligned} \quad (3.8)$$

Of course, the diffusion maps provide the required embedding of \mathcal{X} in \mathbb{R}^n , since their Euclidean geometry in \mathbb{R}^n is identical to the diffusion geometry of the dataset \mathcal{X} , i.e. $\mathbf{D}^{(t)}(x_i, x_j) = \|\Psi^{(t)}(x_i) - \Psi^{(t)}(x_j)\|$, $i, j \in [n]$. In order to achieve a low-dimensional embedding, the diffusion map from Eq. 3.8 is projected onto its significant principal components according to the decay rate of the spectrum of M^t . Specifically, for a sufficiently small $|s_{k+1}|^t$, the k -dimensional embedding is $\mathbf{T}_k \circ \Psi^{(t)}$, where $\mathbf{T}_k : \mathbb{R}^n \rightarrow \mathbb{R}^k$ is the projection on the first k coordinates. Lemma 3.1 quantifies the distortion resulted by such a projection.

Lemma 3.1. *Let $c = \max_{i \in [n]} \hat{\mathbf{d}}_{(i)}^{-1/2}$ (see Eq. 3.2). Then the k -dimensional embedding $\mathbf{T}_k \circ \Psi^{(t)}$ is a $\sqrt{2}c |s_{k+1}|^t$ -distortion of the n -dimensional diffusion map $\Psi^{(t)}$ from Eq. 3.8.*

Proof. Following Eq. 3.8 we get²

$$\begin{aligned} \left| \|\Psi^{(t)}(x_i) - \Psi^{(t)}(x_j)\| - \|\mathbf{T}_k \circ \Psi^{(t)}(x_i) - \mathbf{T}_k \circ \Psi^{(t)}(x_j)\| \right| &\leq \|\Psi^{(t)}(x_i) - \mathbf{T}_k \circ \Psi^{(t)}(x_i)\| \\ &\quad + \|\Psi^{(t)}(x_j) - \mathbf{T}_k \circ \Psi^{(t)}(x_j)\| \\ &\leq |s_{k+1}|^t (\hat{\mathbf{d}}_{(i)}^{-1/2} + \hat{\mathbf{d}}_{(j)}^{-1/2}) \\ &\leq \sqrt{2}c |s_{k+1}|^t. \end{aligned}$$

□

²Here, for comparison purposes, we use the convention that $\mathbf{T}_k : \mathbb{R}^n \rightarrow \mathbb{R}^n$, is the operator that zeros out the last $n - k$ coordinates.

The distortion bound from Lemma 3.1 is referred to as the *analytic bound*. In many cases, spectral properties of the utilized kernel are known a-priori with no need for its explicit computation. The Gaussian kernel is just one example (see [7, 8]) but not the only. In such cases, only a partial SVD can be calculated to produce the relevant principal components according to the required distortion.

3.2 Efficient ICPQR-based DM framework for data analysis

In this section, we provide a QR-based framework for low-dimensional representation of the DM for a training set \mathcal{X} , its out-of-sample extension and anomaly detection. For this purpose, \mathcal{X} is assumed to be a subset of $\bar{\mathcal{X}}$, on which a symmetric point-wise positive kernel $k : \bar{\mathcal{X}} \times \bar{\mathcal{X}} \rightarrow \mathbb{R}$ is defined.

3.2.1 QR-based low-dimensional embedding

Equation 3.5 suggests that the diffusion geometry is already embodied in the Euclidean geometry of $G^{(t)}$'s columns $\mathcal{G}^{(t)} \triangleq \{(G^{(t)})_{(:,1)}, \dots, (G^{(t)})_{(:,n)}\}$ (see Eq. 3.4). According to Proposition 2.3 and Eq. 3.5, application of Algorithm 2 to $G^{(t)}$ with distortion rate $\mu > 0$ produces an s -dimensional 2μ -distortion $\mathbf{F}_s : \mathcal{G}^{(t)} \rightarrow \mathbb{R}^s$, for which

$$\max_{i,j \in [n]} \|\mathbf{D}^{(t)}(x_i, x_j) - \|\mathbf{h}_s^{(t)}(x_i) - \mathbf{h}_s^{(t)}(x_j)\|\| \leq 2\mu,$$

where $\mathbf{h}_s^{(t)} : \mathcal{X} \rightarrow \mathbb{R}^s$ is defined by

$$\mathbf{h}_s^{(t)}(x_i) \triangleq \mathbf{F}_s((G^{(t)})_{(:,i)}), \quad i \in [n]. \quad (3.9)$$

As was discussed in Section 2.3, the embedding dimension s is not known a-priori and is a non-increasing function of μ . Algorithm 4 summarizes the above.

Algorithm 4: ICPQR-based DM

Input : An $n \times n$ kernel matrix K , time step $t \in \mathbb{N}$ and a nonnegative distortion parameter μ .

Output: An s -dimensional 2μ -distortion $\mathbf{h}_s^{(t)} : \mathcal{X} \rightarrow \mathbb{R}^s$ of $\Psi^{(t)}$, a dictionary $\mathcal{D} \subset \mathcal{G}^{(t)}$ of s elements, an $s \times n$ matrix \bar{R} and a degrees vector $\mathbf{d} \in \mathbb{R}^n$.

- 1 set $\mathbf{d} \in \mathbb{R}^n$, $\mathbf{d}_{(i)} = \sum_{j=1}^n K_{(i,j)}$
 - 2 set the $n \times n$ diagonal matrix D , whose i -th diagonal element is $\mathbf{d}_{(i)}$
 - 3 set the $n \times n$ row stochastic transition probabilities matrix in time t , $P^t = (D^{-1}K)^t$
 - 4 set $G^{(t)} = \|\mathbf{d}\|_1^{1/2} D^{-1/2} (P^t)^*$ (see Eq. 3.4)
 - 5 apply Algorithm 2 to $G^{(t)}$ and μ to get a permutation $\boldsymbol{\pi} : [n] \rightarrow [n]$ and an $s \times n$ matrix \bar{R}
 - 6 define $\mathcal{D} = \{(G^{(t)})_{(:,\boldsymbol{\pi}(1))}, \dots, (G^{(t)})_{(:,\boldsymbol{\pi}(s))}\}$ and $\mathbf{h}_s^{(t)}(x_i) = \bar{R}_{(:,i)}$, $i \in [n]$.
-

The output parameters \bar{R} and \mathbf{d} of Algorithm 4 are needed for the out-of-sample phase, described in Section 3.2.2. In addition, due to Definition 2.2, the anomaly detection in the DM context μ_{strict} takes the form

$$\mu_{strict}^{(t)} = \sup_{i \in [n]} \|(G^{(t)})_{(:,i)} - \mathbf{h}_s^{(t)}(x_i)\|. \quad (3.10)$$

3.2.2 Out-of-sample extension and anomaly detection

Given an out-of-sample data point $x \in \bar{\mathcal{X}} \setminus \mathcal{X}$, the goal of the present section is to extend $\mathbf{h}_s^{(t)}$ from Eq. 3.9 to x . For this purpose, a user-defined probabilities vector $\mathbf{p}^{(t)}(x) \in \mathbb{R}^n$ has to be defined. The i -th entry of $\mathbf{p}^{(t)}(x)$ defines the transition probabilities from x to $x_i \in \mathcal{X}$ in t time steps, i.e. $\mathbf{p}^{(t)}(x)_i = \mathbb{P}(x(t) = x_i | x(0) = x)$, $i \in [n]$. Consequently, consistently with Eq. 3.4, the n dimensional extension of $\mathcal{G}^{(t)}$ to x is defined by $\mathbf{g}^{(t)}(x) \triangleq \|\mathbf{d}\|_1^{1/2} D^{-1/2} \mathbf{p}^{(t)}(x) \in \mathbb{R}^n$. Then, Algorithm 3 is applied to $\mathbf{g}^{(t)}(x)$ to produce an s -dimensional embedding $\mathbf{h}_s^{(t)}(x) = \mathbf{F}_s(\mathbf{g}^{(t)}(x))$.

This scheme is consistent with the low-dimensional embedding scheme, described in Section 3.2.1, in the sense that if the probabilities vector $\mathbf{p}^{(t)}(x)$ equals to an in-sample probabilities vector $(P^*)_{(:,i)}$ for a certain $i \in [n]$, then $\mathbf{h}_s^{(t)}(x) = \mathbf{h}_s^{(t)}(x_i)$. Algorithm 5 summarizes the above.

Algorithm 5: Out-of-sample extension for ICPQR-based DM

Input : A dictionary \mathcal{D} , an $s \times n$ matrix \bar{R} and a degrees vector $\mathbf{d} \in \mathbb{R}^n$, which are the outputs from Algorithm 4 and a transition probabilities vector $\mathbf{p}^{(t)}(x) \in \mathbb{R}^n$.

Output: The extension of $\mathbf{h}_s^{(t)}$ to x , $\mathbf{h}_s^{(t)}(x)$ and the associated distortion rate $\boldsymbol{\mu}(x)$.

- 1 set $\mathbf{g}^{(t)}(x) = \|\mathbf{d}\|_1^{1/2} D^{-1/2} \mathbf{p}^{(t)}(x)$, where D is the diagonal $n \times n$ matrix $\text{diag}(\mathbf{d})$
 - 2 apply Algorithm 3 to \mathcal{D} , \bar{R} and $\mathbf{g}^{(t)}(x)$ to get $\mathbf{h}_s^{(t)}(x) \in \mathbb{R}^s$
 - 3 define $\boldsymbol{\mu}(x) = (\|\mathbf{g}^{(t)}(x)\|^2 - \|\mathbf{h}_s^{(t)}(x)\|^2)^{1/2}$
-

One possibility for the definition of $\mathbf{p}^{(1)}(x)$, which is the first time step transfer probabilities from an out-of-sample data point to \mathcal{X} , is via the kernel function k by

$$\mathbf{p}^{(1)}(x)_{(i)} \triangleq k(x, x_i) / \sum_{j=1}^n k(x, x_j), \quad i \in [n]. \quad (3.11)$$

This definition is consistent with Eq. 3.1. Then, the corresponding transition probabilities vector in time step t can be heuristically defined by $\mathbf{p}^{(t)} = (P^*)^{t-1} \mathbf{p}^{(1)}(x)$. This definition represents a Markovian process for which the new data point x is inaccessible from the dataset \mathcal{X} , and the transition probabilities from x to \mathcal{X} in time-step t are determined by the transition probabilities from x to \mathcal{X} in the first time-step that are represented by $\mathbf{p}^{(1)}(x)$, and the transition probabilities among the elements of \mathcal{X} after $t - 1$ time-steps that are represented by $(P^*)^{t-1}$. Finally, the (strict) normality of an out-of-sample data point x is determined due to Definition 2.2 and Eq. 3.10 by using the distortion rate function $\boldsymbol{\mu}(x)$ from step 3 in Algorithm 5.

4 Experimental Results

This section analyses three different datasets using the proposed methodologies from Sections 2 and 3 that are synthetic and real. Section 4.1 exemplifies the basic notions of geometry preservation, anomaly detection and out-of-sample extension through the application of the QR-based DM to a synthetic dataset as described in Section 3. A comparison with the method proposed in [41] for diffusion geometry preservation is presented in this section as well. A QR-based DM analysis of real data is demonstrated in Section 4.2. The analysis in both of the above examples is based on the corresponding first time step in DM. Finally, Section 4.3 presents a multiclass classification of parametric data, using generalizations of the out-of-sample extension and anomaly detection methods, presented in Section 2.5.

4.1 QR-based DM analysis - toy example

In this section, we present a diffusion-based analysis of a synthetic two dimensional manifold, immersed in a three dimensional Euclidean space. The analyzed dataset $\mathcal{X} \subset \mathbb{R}^3$ consists of $n = 3,000$ data points, uniformly sampled from a Swiss roll, shown in Fig. 4.1.




Figure 4.1: Swiss roll that contains 3,000 uniformly distributed data points. Data points are colored according to their distance from the origin.

The utilized kernel function is the commonly-used Gaussian kernel $k_\varepsilon : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$,

$$k_\varepsilon(x, y) \triangleq e^{-\|x-y\|^2/\varepsilon}, \quad \varepsilon > 0, \quad (4.1)$$

where the norm in the exponent is the standard three dimensional Euclidean norm. The associated $n \times n$ kernel matrix is K_ε whose (i, j) -th entry is

$$(K_\varepsilon)_{(i,j)} \triangleq k_\varepsilon(x_i, x_j), \quad i, j \in [n].$$

The corresponding degrees function is $\mathbf{d}_\varepsilon \in \mathbb{R}^n$, whose i -th entry is $(\mathbf{d}_\varepsilon)_{(i)} \triangleq \sum_{j=1}^n k_\varepsilon(x_i, x_j)$, $i \in [n]$ and D_ε is the $n \times n$ diagonal matrix, whose i -th diagonal entry is $(\mathbf{d}_\varepsilon)_{(i)}$. Based on these, according to Eq. 3.1, the $n \times n$ transition probabilities matrix is defined by

$$P_\varepsilon \triangleq D_\varepsilon^{-1} K_\varepsilon. \quad (4.2)$$

Thus, transition probabilities between close data points are high and low for far points.

Section 4.1.1 addresses the qualitative dependency between the neighborhood parameter ε , and the required embedding's dimensionality. Section 4.1.2 shows a further step of dimensionality reduction, using the optimal coordinates system, as described in Section 2.3. Out-of-sample extension and anomaly detection, as described in Section 3.2.2, are demonstrated in Section 4.1.3. Finally, Section 4.1.4 presents a brief description of the μ -IDM method [41] and compares its performances to the proposed method in this paper.

4.1.1 The dependency between ε and the embedding's dimension s

As was proved in [7], as ε increases, the numerical rank of P_ε decreases and vice versa. Mathematically, let $1 = s_1^{(\varepsilon)} \geq s_2^{(\varepsilon)} \geq \dots \geq s_n^{(\varepsilon)} \geq 0$ be the eigenvalues³ of P_ε . Define $\mathbf{E}_\varepsilon(r) : [0, 1] \rightarrow [0, 1]$, $\mathbf{E}_\varepsilon(r) \triangleq (\sum_{i=1}^t (s_i^{(\varepsilon)})^2 / \sum_{i=1}^n (s_i^{(\varepsilon)})^2)^{1/2}$ to be the energy's portion of P_ε , which is captured by the first t eigenvalues of P_ε , where $r = t/n$ is the corresponding spectrum

³The eigenvalues of P_ε are nonnegative since the Gaussian kernel function k_ε from Eq. 4.1 is positive definite due to Bochner's theorem [48]. Thus, if the data points in \mathcal{X} are all distinct, then k_ε is strictly positive definite and the eigenvalues of P_ε are all positive.

ratio. Then, as ε increases, the number of significant eigen-components decreases as demonstrated in Fig. 4.2. This fact, combined with Lemma 3.1, suggests that when ε decreases, the number of components required to achieve a certain distortion increases as seen in Fig. 4.3.

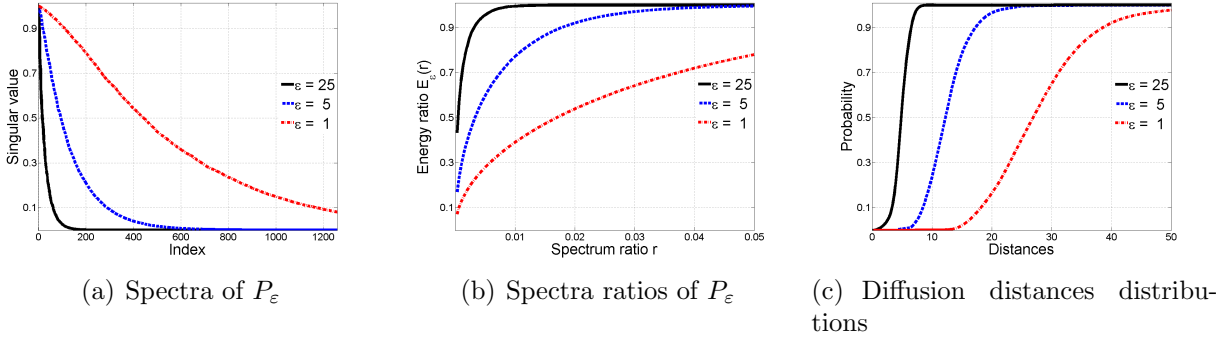


Figure 4.2: Spectral and geometrical views of three diffusion geometries that correspond to the neighborhood parameters $\varepsilon = 1, 5$ and 25 . Figure (a) shows that as ε becomes larger the spectrum decays faster. An immediate consequence is shown in Fig. (b) that shows the relation between the number of significant components and ε . Figure (c) shows the probability distribution of the diffusion distances. It is clear that the use of large ε results in many short diffusion distances and vice-versa.

Figure 4.3 compares between the analytic bound (see Lemma 3.1), the minimal dimension of DM and the QR-based DM dimension that are required to achieve a certain distortion. It also demonstrates the above discussed relation between the neighborhood parameter ε and the dimensionality of the embedding. Thus, for a larger ε a fewer dimensions are required to achieve a certain distortion.




Figure 4.3: Number of components (dimensions- y -axis) required to preserve the Swiss roll diffusion geometry up to a distortion (x -axis) for three different neighborhood sizes. The continuous (green) graph denotes the analytic bound provided by Lemma 3.1, the dashed (red) graph is the QR-based DM dimension produced by Algorithm 2 and the dash-dotted (blue) graph is the minimal DM dimension required to achieve a certain distortion.

4.1.2 Low-dimensional embedding

In this section, a comparison between the classic DM and ICPQR-based DM is presented. Figure 4.4 shows the two-dimensional DM embedding of \mathcal{X} , and a two-dimensional view of an aligned versions of ICPQR-based DM, applied to \mathcal{X} with three different distortion values.




Figure 4.4: Comparison between the two-dimensional DM and the projected aligned ICPQR-based DM of \mathcal{X} with $\varepsilon = 3$, for the first diffusion time-step: (a) two most significant DM coordinates of $\Psi^{(1)}$. (b)-(d): two most significant coordinates of the aligned version of $\mathbf{h}_s^{(1)}$, with (b) $\mu = 0.1$, $s = 1,246$, actual distortion 0.01, (c) $\mu = 1$, $s = 752$, actual distortion 0.23, and (d) $\mu = 5$, $s = 382$, actual distortion 3.91. Data coloring is consistent with Fig. 4.1.

It should be stressed that $\mathbf{h}_s^{(1)}$, which is the output of Algorithm 4, is an s -dimensional 2μ -distortion of the n -dimensional DM $\Psi^{(1)}$ as defined in Definition 1.1. Therefore, it is unlikely that low-dimensional projections (lower than s) will represent similar geometries. Nevertheless, to demonstrate the notion of low-rate distortion, Fig. 4.4 shows a two dimensional view of an aligned version of $\Psi^{(1)}$ with $\Psi^{(1)}$. For that purpose, the DM of \mathcal{X} was explicitly computed. The utilized alignment algorithm is described in Appendix A.

4.1.3 Out of sample extension and anomaly detection

Application of Algorithm 4 to \mathcal{X} with K_ε , $\varepsilon = 3$, $t = 1$ and $\mu = 0.1$ was resulted in an s -dimensional 2μ -distortion of $\Psi^{(1)}$, $\mathbf{h}_s^{(1)} : \mathcal{X} \rightarrow \mathbb{R}^s$ with $s = 1,246$. An out-of-sample extension of $\mathbf{h}_s^{(1)}$ to $\bar{\mathcal{X}}$ is demonstrated in this section, as well as anomaly detection where $\mathcal{X} \subset \mathbb{R}^3$ is a random subset of 10,000 data points that are uniformly sampled from the three dimensional bounding box of \mathcal{X} .

For that purpose, Algorithm 5 was applied to $\bar{\mathcal{X}}$. Beside its three first inputs, which are provided as outputs from Algorithm 4, a transition probabilities vector $\mathbf{p}^{(1)}(x) \in \mathbb{R}^n$ has to be defined for any $x \in \bar{\mathcal{X}}$. In this example, $\mathbf{p}^{(1)}(x)$ was defined consistently with the kernel K_ε by using Eq. 3.11. This definition coincides with the definition of the transition probabilities matrix P_ε in Eq. 4.2 that results in exact extension on \mathcal{X} , i.e. if $x = x_i$ for a certain $i \in [n]$, then $\mathbf{h}_s^{(1)}(x) = \mathbf{h}_s^{(1)}(x_i)$.

The results are shown in Fig. 4.5. Figure 4.5(a) shows a side view of the dataset $\bar{\mathcal{X}}$. Each data point $x \in \bar{\mathcal{X}}$ is colored proportionally to its out-of-sample extension distortion rate $\mu(x)$ (see step 3 in Algorithm 5.) Classification of $\bar{\mathcal{X}}$ to either normal or abnormal classes is shown in Fig. 4.5(b). The normal class $\mathcal{N} \subset \bar{\mathcal{X}}$ is darkly colored and the abnormal class $\bar{\mathcal{N}}$ is brightly colored. The classification was done according to Definition 2.2. Thus, $x \in \bar{\mathcal{X}}$ is classified as normal if its distortion rate satisfies $\mu(x) \leq \mu$. Otherwise, it is classified as abnormal. Lastly, a two dimensional view of the out-of-sample extension of the normal class, namely $\mathbf{h}_s^{(1)}(x)$, $x \in \mathcal{N}$, is shown in Fig. 4.5(c). Each embedded data point is colored in the same color as its nearest neighbor from the embedding of \mathcal{X} by $\mathbf{h}_s^{(1)}(\mathcal{X})$. The shown coordinates system is consistent with the one presented in Fig. 4.4.

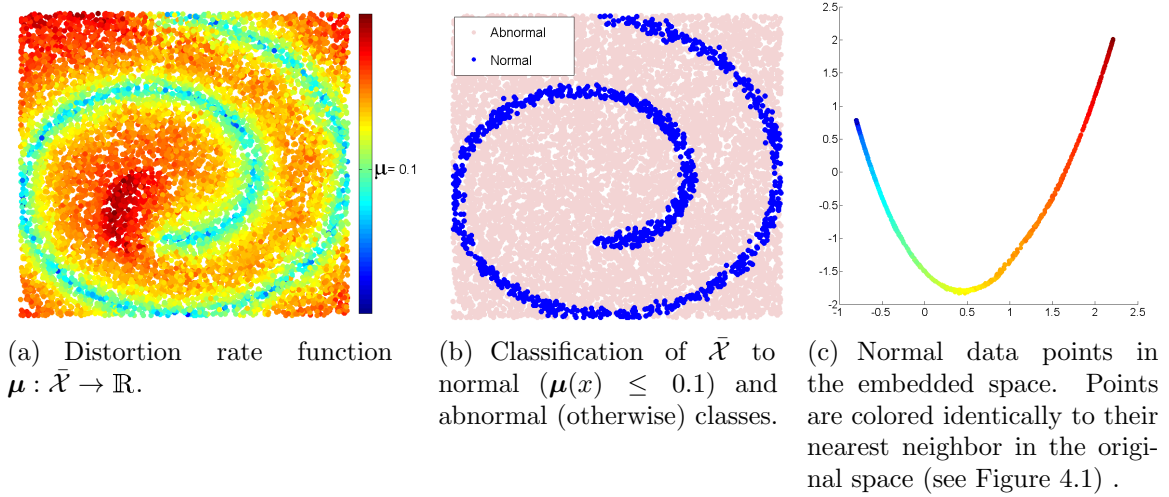


Figure 4.5: 10,000 out-of-sample data points were randomly selected in the bounding box of the original Swiss roll. Left and center: two-dimensional side view of the out-of-sample dataset colored by their distortion rate and classification, respectively. Right: extension of the first two meaningful ICPQR-based DM embedded space coordinates to the normal out-of-sample data points.

4.1.4 Comparison with μ IDM

The μ IDM algorithm in [41] is a dictionary-based method that provides a low rank 2μ -distortion for the first transition time step $\Psi^{(1)}$ of DM. The algorithm incrementally constructs an approximated map by using a single scan of the data. This algorithm is greedy and sensitive to the scan order. Typically, the growth rate of the dictionary is very high at the beginning and decays as time advances. Moreover, the resulted dictionary and the resulted embedding's dimension may be redundant. In each iteration of the μ IDM, a newly processed data point is considered for inclusion in the dictionary that was constructed from previously scanned data points. At the beginning of every iteration, the dictionary elements are already embedded in a low-dimensional space (whose dimension equals to the size of the dictionary), where its geometry is identical to the diffusion geometry of the data restricted to the dictionary. Then, a Nyström-type extension [5] is applied to the scanned data points, based on its affinities with the dictionary elements, in order to approximate the embedding of the newly processed data point. The exact DM of this data point together with the dictionary is then efficiently computed. The geometries of these two embeddings are identical for the dictionary, therefore, at this stage these two geometries are aligned to coincide on the dictionary. The distance of the extended map from the exact map of the examined data point is measured. If this distance is larger than μ then the examined data point is added to the dictionary. The entire computational complexity of this iterative process is lower than the computational complexity of DM. The exact number of required operations depends on the required accuracy and on the dimensionality of the original ambient space.

The presently proposed QR-based DM method considers the entire dataset in each iteration and it is not sensitive to the order of the dataset. Therefore, the resulted dictionary is more sparse as demonstrated in Table 4.1 and in Fig. 4.6.

μ		ICPQR-based DM	μ IDM [41]
0.1	Dictionary size s	1,246	2,305
	Execution time	43 sec.	7 hours
	Actual distortion	0.01	0.03
1	Dictionary size s	752	1,293
	Execution time	27 sec.	71 minutes
	Actual distortion	0.23	0.61
5	Dictionary size s	382	630
	Execution time	17 sec.	15 minutes
	Actual distortion	3.91	4.46
10	Dictionary size s	190	284
	Execution time	9 sec.	4 minutes
	Actual distortion	12.81	13.24

Table 4.1: Comparison between ICPQR-based DM and μ IDM algorithms related to dictionary size, execution time and actual distortion⁴, w.r.t. $\Psi^{(1)}$, for various distortion parameters. Clearly, the actual distortion is bounded by 2μ . Execution times are averaged over 10 runs of the algorithms.

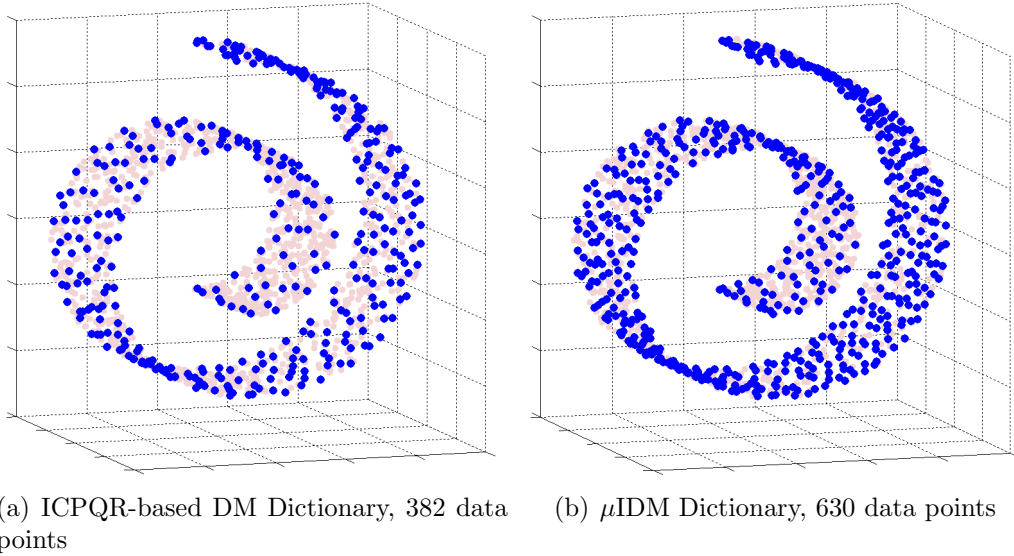


Figure 4.6: Data points admitted to ICPQR-based DM and μ IDM dictionaries (dark points) with distortion parameter $\mu = 5$. The input was given to both algorithms ordered from the inner part of the Swiss roll to the outer with a larger radius.

4.2 QR-based DM analysis - real-world data

This section exemplifies a semi-supervised anomaly detection process applied to a real-world dataset. The examined dataset $\bar{\mathcal{X}} \subset \mathbb{R}^{14}$ is the DARPA dataset [27] consists of $\bar{n} = 12,617$ data points. Each data point is a vector of 14 features that describes a computer network traffic labeled as either normal and legitimate or abnormal that pertains to be an attack and intrusion on the network. The dataset is divided into a training set $\mathcal{X} \subset \bar{\mathcal{X}}$, which contains $n = 6,195$ normally behaving samples, and 5 testing datasets $\{\mathcal{T}_{mon}, \dots, \mathcal{T}_{fri}\}$, which were collected during different days of the week. Each day contains normal and abnormal data points as described in Table 4.2.

⁴An actual distortion of a function f w.r.t. g is $\sup_{x,y \in \mathcal{X}} \|\|f(x) - f(y)\| - \|g(x) - g(y)\|\|$.

First, the training dataset was scaled to place it in the 14-dimensional unit box $[0, 1]^{14}$. Then, the same scaling was applied to the testing datasets $\mathcal{T} = \mathcal{T}_{mon} \cup \dots \cup \mathcal{T}_{fri}$. Application of Algorithm 4 to the training set \mathcal{X} , using the Gaussian kernel from Eq. 4.1 with $\varepsilon = 0.6$, $t = 1$ and $\mu = 10^{-7}$, produced an s -dimensional 2μ distortion of the first time step $\Psi^{(1)}$ of DM of \mathcal{X} , $\mathbf{h}_s^{(1)} : \mathcal{X} \rightarrow \mathbb{R}^s$ with $s = 138$. The neighborhood size parameter ε was chosen to be twice the median of all the mutual distances between the data points in \mathbb{R}^{14} . Such a selection is a common heuristic for determining this parameter in DM context. This concludes the training phase.

As a second stage, for each testing data point $x \in \mathcal{T}$, the probabilities vector $\mathbf{p}^{(1)}(x)$ was defined as a natural extension of the Gaussian kernel from Eq. 4.1 by using Eq. 3.11 with $\varepsilon = 0.6$. Any testing data point, which is distant (relatively to ε) from the training set \mathcal{X} , yields $k_\varepsilon(x, y) \approx 0$, $y \in \mathcal{X}$, and was a-priori classified as abnormal. The subset of distant testing data points is denoted by \mathcal{F} .

Finally, Algorithm 5 was applied to the elements of $\mathcal{T} \setminus \mathcal{F}$, using the previously computed probabilities vectors $\mathbf{p}^{(1)}(x)$, $x \in \mathcal{T}$, to get an extension of $\mathbf{h}_s^{(1)} : \mathcal{T} \setminus \mathcal{F} \rightarrow \mathbb{R}^s$ and a distortion rate $\mu : \mathcal{T} \setminus \mathcal{F} \rightarrow \mathbb{R}$. Then, abnormal data points were detected by following Definition 2.2 with μ_{strict} from Eq. 3.10.

The anomaly detection results are summarized in Table 4.2. Some of them are demonstrated in Fig. 4.7. The rates in the accuracy percentage and in the false alarms columns are related to the original labeling of $\bar{\mathcal{X}}$.

Set	Size	# of anomalies	Accuracy [%]	False Alarms [%]
\mathcal{T}_{mon}	1, 321	1	100	0.68
\mathcal{T}_{tue}	1, 140	53	100	0.53
\mathcal{T}_{wed}	1, 321	16	100	0.08
\mathcal{T}_{thu}	1, 320	24	96	1.74
\mathcal{T}_{fri}	1, 320	18	100	0.15

Table 4.2: Anomaly detection performances.

Figure 4.7 presents three-dimensional views of an aligned version of $\mathbf{h}_s^{(1)}$ of \mathcal{X} , as well as its extension to $\mathcal{T}_{thu} \setminus \mathcal{F}$ with the first three significant coordinates of $\Psi^{(1)}$, the DM of \mathcal{X} in the first time step. As can be seen in Fig 4.7, most of data is located near the training set while the abnormal data points are embedded far away.

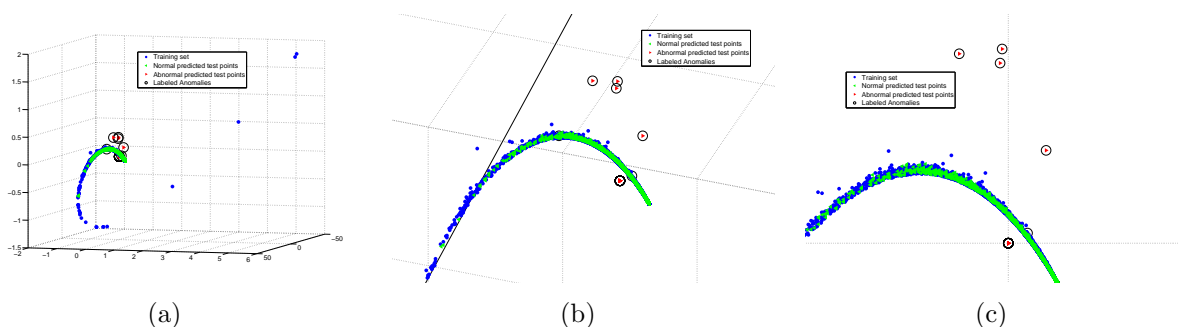


Figure 4.7: Three different angles and scales of an aligned version of a three-dimensional projection of $\mathbf{h}_s^{(1)}$, applied to the training set \mathcal{X} (blue points), its out-of-sample extension to $\mathcal{T}_{thu} \setminus \mathcal{F}$ (strictly-normal data points are green, strictly-abnormal data points are red), and labeled anomalies (black circles). (a) general view. (b) strictly-normal out-of-sample data points are mapped closely to the training set. (c) strictly-abnormal data points.

As in Section 4.1.2, the alignment was done only for visualization purposes using the algorithm in Appendix A.

4.3 Semi-supervised multi-class classification of high dimensional data

In this section, a multi classification process, based on Algorithms 2 and 3, is presented. The analyzed data is the parametric m -dimensional ISOLET dataset [3] $\mathcal{X} \subset \mathbb{R}^m$ where $m = 617$ that contains 7,797 data points in $[-1, 1]^m$. Each data point corresponds to a single human pronunciation of ISOLated LETters. The goal is to classify a testing subset $\mathcal{T} \subset \bar{\mathcal{X}}$, which contains 1,559 letter-pronunciation samples spoken by 30 people, to 26 classes that are based on a training set $\mathcal{X} \subset \bar{\mathcal{X}}$ that contains $n = 6,238$ samples that were already classified to \mathcal{X}_ζ , $\zeta \in \mathcal{I} \triangleq \{A, B, \dots, Z\}$, spoken by 120 different people. For this purpose, Algorithm 2 was applied to each training set \mathcal{X}_ζ with a distortion parameter $\mu = 4.7$, to produce 26 dictionaries $\mathcal{D}_\zeta \subset \mathcal{X}_\zeta$, $\zeta \in \mathcal{I}$. Then, for each of these dictionaries, Algorithm 3 was applied to each element in the testing set $\mathbf{x} \in \mathcal{T}$, to produce its distortion rate $\mu_\zeta(\mathbf{x})$ (see Eq. 2.15). Finally, each of the testing data points was classified to a class whose dictionary described it best, i.e. \mathbf{x} was classified to class ζ_0 , where $\zeta_0 \triangleq \arg \min_{\zeta \in \mathcal{I}} \mu_\zeta(\mathbf{x})$. The parameter μ was determined by taking part of the training set to serve as a validation set. Then, several values of μ were applied on the (reduced) training set, and the validation set was classified based upon these values. The chosen μ was the one that was optimal on the validation set. The results are shown in Fig. 4.8. Out of 1,559 test samples, 92% were classified correctly. The classification of the test data is presented in a confusion matrix in Fig. 4.8. By looking at the shades of the diagonal it can be seen that the majority of the test samples of each class were classified correctly. The sets $\{B, C, D, E, G, P, T, V, Z\}$ and $\{M, N\}$ are the most difficult letters to classify due to high similarity in pronunciation of these letters within each set. The state-of-the-art classification accuracy of this dataset is 96.73%. It was achieved in [21] by using 30-bit error correcting output codes that is based on neural networks. This method is far more complex than the solution proposed in this work.




Figure 4.8: Classification of the test samples in the ISOLET dataset. For each ordered couple (i, j) , the cell at row i and column j is colored according to the number of test samples belonging to class i that were classified by the algorithm to class j . The cells on the diagonal denote correct classification. For each letter, a total of 60 test samples are provided (except for 'M' which is missing one sample due to recording difficulties).

5 Conclusions and Future Works

This work presents a complete framework for linear dimensionality reduction, out-of-sample extension and anomaly detection algorithms for high-dimensional parametric data, which is based on incomplete pivoted QR decomposition of the associated data matrix. The presented method preserves the high-dimensional geometry of the data up to a user-defined distortion parameter. Such a low-dimensional data representation enables further geometrically-based data analysis which, due to the geometry-preservation, is still valid to the original data. The storage complexity of the method is extremely low compared to the classical PCA. In the worst case, its computational complexity is similar to that of the PCA. The method provides a dictionary, which is a subset of landmark data points, that forms a basis for the projection of the entire data to a low dimensional space. Out-of-sample extension and anomalous data point detection become simple tasks once the dictionary is computed. Although the suggested method is designated for parametric data analysis, in some cases it can be adapted to non-parametric data analysis frameworks, as was demonstrated for the DM framework. The stability of our method to perturbations (noise) was proved and its connection to matrix-approximation was presented.

Experimental results show that our method achieves a lower dimensional embedding than PCA achieves for a certain distortion. Moreover, analysis of both synthetic and real-world datasets achieved good performance for dimensionality reduction, out-of-sample extension, anomaly detection and multi-class classification.

Future work includes randomized version of the presented framework to provide a more computationally efficient method for a dictionary-based dimensionality reduction method. In addition, a generalization of our method for non-parametric data analysis methods should be considered, as well as a dynamic framework that enables to cope with datasets that vary across time. Finally, a parallel version of the algorithm, which simultaneously builds dictionaries for subsets of the data and then unifies them, is planned.

A Appendix: Alignment Algorithm

This section details the alignment algorithm that was utilized in Sections 4.1.2 and 4.2, for visualization purposes.

Suppose that A and B are two $m \times n$ matrices of n data points in \mathbb{R}^m . If the sizes of A and B are different, then padding by zeros is performed. Clearly, it will not change the geometry of the columns of these matrices. Let \bar{A} and \bar{B} be the centralized versions of A and B around the columns means of each one of them. Then, the best orthogonal alignment of \bar{B} 's columns with \bar{A} 's columns is provided by the orthogonal matrix $Q = U_A U_B^*$, where U_A and U_B are the left singular vectors of A and B , respectively. Then, the aligned centralized matrix $\tilde{B} = Q * \bar{B}$ is decentralized by A 's columns mean. Obviously, since Q is orthogonal, the geometry of \tilde{B} 's columns is unaffected. Algorithm 6 describes the above.

Algorithm 6: Alignment Algorithm

Input : Two $m \times n$ matrices A and B

Output: An $m \times n$ matrix \tilde{B} whose columns geometry is identical to that of B 's columns, and \tilde{B} 's columns are optimally aligned with A 's columns.

- 1 centralize A : $\bar{A} = A - \mathbf{a}\mathbf{1}_n^*$, where $\mathbf{a} = \sum_{j=1}^n A_{(:,j)}$ and $\mathbf{1}_n \in \mathbb{R}^n$ is the all-ones vector
 - 2 centralize B : $\bar{B} = B - \mathbf{b}\mathbf{1}_n^*$, where $\mathbf{b} = \sum_{j=1}^n B_{(:,j)}$
 - 3 compute the left singular vectors U_A and U_B of \bar{A} and \bar{B} , respectively
 - 4 define $\tilde{B} = U_A U_B^* \bar{B} + \mathbf{a}\mathbf{1}_n^*$
-

Acknowledgments

This research was partially supported by the Israeli Ministry of Science & Technology (Grants No. 3-9096, 3-10898), US-Israel Binational Science Foundation (BSF 2012282) and Blavatnik Computer Science Research Fund.

References

- [1] D. Achlioptas. Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of Computer and System Sciences*, 66(4):671–687, 2003.
- [2] N. Ailon and B. Chazelle. Approximate nearest neighbors and the fast johnson-lindenstrauss transform. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 557–563. ACM, 2006.
- [3] K. Bache and M. Lichman. UCI machine learning repository, 2013.
- [4] Z-Z. Bai, I. S. Duff, and A. J. Wathen. A class of incomplete orthogonal factorization methods. i: Methods and theories. *BIT Numerical Mathematics*, 41(1):53–70, 2001.
- [5] C. T. H. Baker. *The Numerical Treatment of Integral Equations*. Oxford: Clarendon Press, 1977.
- [6] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.

- [7] A. Bermanis, A. Averbuch, and R.R. Coifman. Multiscale data sampling and function extension. *Applied and Computational Harmonic Analysis*, 34(1):15 – 29, 2013.
- [8] A. Bermanis, G. Wolf, and Averbuch. A. Cover-based bounds on the numerical rank of gaussian kernels. *Applied and Computational Harmonic Analysis*, 36(2):302 – 315, 2014.
- [9] C. Boutsidis, M. W. Mahoney, and P. Drineas. Unsupervised feature selection for principal components analysis. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 61–69. ACM, 2008.
- [10] C. Boutsidis, M. W. Mahoney, and P. Drineas. An improved approximation algorithm for the column subset selection problem. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '09*, pages 968–977, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
- [11] C. Boutsidis, J. Sun, and N. Anerousis. Clustered subset selection and its applications on it service metrics. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM '08*, pages 599–608, New York, NY, USA, 2008. ACM.
- [12] C. Boutsidis and D. P. Woodruff. Optimal cur matrix decompositions. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 353–362. ACM, 2014.
- [13] M. Brand. Fast low-rank modifications of the thin singular value decomposition. *Linear Algebra and its Applications*, 415(1):20 – 30, 2006. Special Issue on Large Scale Linear and Nonlinear Eigenvalue Problems.
- [14] P. Brard, G. Besson, and S. Gallot. Embedding riemannian manifolds by their heat kernel. *Geometric and Functional Analysis GAFA*, 4(4):373–398, 1994.
- [15] T. F. Chan. Rank revealing qr factorizations. *Linear Algebra and its Applications*, 88–89(0):67 – 82, 1987.
- [16] T. F. Chan and P. Hansen. Some applications of the rank revealing QR factorization. *SIAM Journal on Scientific and Statistical Computing*, 13(3):727–741, 1992.
- [17] H. Cheng, Z. Gimbutas, P.G. Martinsson, and V. Rokhlin. On the compression of low rank matrices. *SIAM Journal on Scientific Computing*, 26(4):1389–1404, 2005.
- [18] K. L. Clarkson. Tighter bounds for random projections of manifolds. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 39–48. ACM, 2008.
- [19] R.R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
- [20] J.K. Cullum and R.A. Willoughby. *Lanczos Algorithms for Large Symmetric Eigenvalue Computations I: Theory*, volume 41 of *Classics in Applied Mathematics*. SIAM, 2002.
- [21] T.G. Dietterich and G. Bakiri. Error-correcting output codes: A general method for improving multiclass inductive learning programs. In *IN PROCEEDINGS OF AAAI-91*, pages 572–577. AAAI Press, 1991.
- [22] P. Drineas, M. W. Mahoney, and S. Muthukrishnan. Subspace sampling and relative-error matrix approximation: Column-based methods. In *In Proc. of the 10th RANDOM*, pages 316–326, 2006.

- [23] P. Drineas, M. W. Mahoney, and S. Muthukrishnan. Relative-error cur matrix decompositions. *SIAM Journal on Matrix Analysis and Applications*, 30(2):844–881, 2008.
- [24] A. Frieze, R. Kannan, and S. Vempala. Fast monte-carlo algorithms for finding low-rank approximations. *Journal of the ACM (JACM)*, 51(6):1025–1041, 2004.
- [25] W. Givens. Computation of plane unitary rotations transforming a general matrix to triangular form. *Journal of the Society for Industrial and Applied Mathematics*, 6(1):pp. 26–50, 1958.
- [26] G.H. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, fourth edition, 2013.
- [27] I. Graf, R. Lippmann, R. Cunningham, D. Fried, K. Kendall, S. Webster, and M. Zissman. Results of darpa 1998 offline intrusion detection evaluation. In *DARPA PI Meeting*, volume 15, 1998.
- [28] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [29] H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24, 1933.
- [30] A. S. Householder. Unitary triangularization of a nonsymmetric matrix. *J. ACM*, 5(4):339–342, October 1958.
- [31] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613. ACM, 1998.
- [32] A. Jennings and M. Ajiz. Incomplete methods for solving $a^t a x = b$. *SIAM Journal on Scientific and Statistical Computing*, 5(4):978–987, 1984.
- [33] W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- [34] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [35] M. W. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2):123–224, 2011.
- [36] M. W. Mahoney and P. Drineas. Cur matrix decompositions for improved data analysis. *Proceedings of the National Academy of Sciences*, 106(3):697–702, 2009.
- [37] P.G. Martinsson, V. Rokhlin, and M. Tygert. A randomized algorithm for the decomposition of matrices. *Applied and Computational Harmonic Analysis*, 30(1):47–68, 2011.
- [38] A. T. Papadopoulos, I. S. Duff, and A. J. Wathen. A class of incomplete orthogonal factorization methods. ii: implementation and results, 2002.
- [39] J. R. Rice. Experiments on gram-schmidt orthogonalization. *Math. Comp.*, 20:pp. 325–328, 1966.

- [40] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
- [41] M. Salhov, A. Bermanis, G. Wolf, and A. Averbuch. Approximately-isometric diffusion maps. *Applied and Computational Harmonic Analysis*, 38:399–415, 2015.
- [42] B. Schölkopf, A. Smola, E. Smola, and K.R. Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10:1299–1319, 1998.
- [43] L. J. Schulman. Clustering for edge-cost minimization. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 547–555. ACM, 2000.
- [44] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290:2319–2323, 2000.
- [45] S. Vempala. *The random projection method*, volume 65. American Mathematical Soc., 2005.
- [46] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17:395–416, 2007.
- [47] S. Wang and Z. Zhang. Improving cur matrix decomposition and the nyström approximation via adaptive sampling. *The Journal of Machine Learning Research*, 14(1):2729–2769, 2013.
- [48] H. Wendland. *Scattered data approximation*. Cambridge University Press, 2005.