



Diffusion representations



Moshe Salhov, Amit Bermanis, Guy Wolf, Amir Averbuch*

School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel

ARTICLE INFO

Article history:

Received 19 November 2015

Received in revised form 2 October 2016

2016

Accepted 7 October 2016

Available online 14 October 2016

Communicated by Gregory Beylkin

Keywords:

Manifold learning

Kernel PCA

Diffusion Maps

Diffusion distance

Distance preservation

ABSTRACT

Diffusion Maps framework is a kernel based method for manifold learning and data analysis that defines diffusion similarities by imposing a Markovian process on the given dataset. Analysis by this process uncovers the intrinsic geometric structures in the data. Recently, it was suggested to replace the standard kernel by a measure-based kernel that incorporates information about the density of the data. Thus, the manifold assumption is replaced by a more general measure-based assumption. The measure-based diffusion kernel incorporates two separate independent representations. The first determines a measure that correlates with a density that represents normal behaviors and patterns in the data. The second consists of the analyzed multidimensional data points.

In this paper, we present a representation framework for data analysis of datasets that is based on a closed-form decomposition of the measure-based kernel. The proposed representation preserves pairwise diffusion distances that does not depend on the data size while being invariant to scale. For a stationary data, no out-of-sample extension is needed for embedding newly arrived data points in the representation space. Several aspects of the presented methodology are demonstrated on analytically generated data.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

Kernel methods constitute of a wide class of algorithms for non-parametric data analysis of massive high dimensional datasets. Typically, a limited set of underlying factors generates the high dimensional observable parameters via non-linear mappings. The non-parametric nature of these methods enables to uncover hidden structures in the data. These methods extend the well known Multi-Dimensional Scaling (MDS) [6,14] method. They are based on an affinity kernel construction that encapsulates the relations (distances, similarities or correlations) among multidimensional data points. Spectral analysis of this kernel provides an efficient representation of the data that simplifies its analysis. Methods such as Isomap [29], LLE [22], Laplacian eigenmaps [1], Hessian eigenmaps [8] and local tangent space alignment [31,33], extend

* Corresponding author. Fax: +972 3 6422020.

E-mail address: amir@math.tau.ac.il (A. Averbuch).

the MDS paradigm by assuming to satisfy the manifold assumption. Under this assumption, the data is assumed to be sampled from a low intrinsic dimensional manifold that captures the dependencies between the observable parameters. The corresponding spectral-based embedding performed by these methods preserves the geometry of the manifold that incorporates the underlying factors in the data.

The diffusion maps (DM) method [5] is a kernel-based method that defines diffusion similarities for data analyzes by imposing a Markovian process over the dataset. It defines a transition probability operator based on local affinities between multidimensional data points. By spectral decomposition of this operator, the data is embedded into a low dimensional Euclidean space, where Euclidean distances represent the diffusion distances in the ambient space. When the data is sampled from a low dimensional manifold, the diffusion paths follow the manifold and the diffusion distances capture its geometry.

DM embedding was utilized for a wide variety of data and pattern analysis techniques. For example it was used to improve audio quality by suppressing transient interference [28]. It was utilized in [25] for detecting and classifying moving vehicles. Additionally, DM was applied to scene classification [12], gene expression analysis [24] and source localization [27]. Furthermore, the DM method can be utilized for fusing different sources of data [16,13].

DM embeddings in both the original version [5,15] and in the measure-based Gaussian correlation (MGC) version [4,3], are obtained by the principal eigenvectors of the corresponding diffusion operator. These eigenvectors represent the long-term behavior of the diffusion process that captures its metastable states [11] as it converges to a unique stationary distribution.

The MGC framework [4,3] enhances the DM method by incorporating information about data distribution in addition to the local distances on which DM is based. This distribution is modeled by a probability measure, which is assumed to quantify the likelihood of data presence over the geometry of the space. The measure and its support in this method replace the manifold assumption. Thus, the diffusion process is accelerated in high density areas in the data rather than being depended solely on the manifold geometry. As shown in [4], the compactness of the associated integral operator enables to achieve dimensionality reduction by utilizing the DM framework.

This MGC construction consists of two independent data points representations. The first represents the domain on which the measure is defined or, equivalently, the support of the measure. The second represents the domain on which the MGC kernel function and the resulting diffusion process are defined. These *measure domain* and the *analyzed domain* may, in some cases, be identical, but separate sets can also be considered by the MGC-based construction. The latter case utilizes a training dataset, which is used as the measure domain to analyze any similar data that is used as the analyzed domain. Furthermore, instead of using the collected data as an analyzed domain, it can be designed as a dictionary or as a grid of representative data points that capture the essential structure of the MGC-based diffusion.

In general, kernel methods can find geometrical meaning in a given data via the application of spectral decomposition. However, this representation changes as additional data points are added to the given dataset. Furthermore, the required computational complexity, which is dictated by spectral decomposition, is $O(n^3)$ that is not feasible for a very large dataset. For example, a segmentation of a medium size image of 512×512 pixels requires a kernel matrix of size $2^{18} \times 2^{18}$. The size of such matrix necessitated about 270 GByte of memory assuming double precision. Spectral decomposition procedure applied to such a matrix is a formidable slow task. Hence, there is a growing need to have more computationally efficient methods that are practical for processing large datasets.

Recently, a method to produce random Fourier features from a given data and a positive kernel was proposed in [21]. The suggested method provides a random mapping of the data such that the inner product of any two mapped points approximates the kernel function with high probability. This scheme utilizes Bochner's theorem [23] that says that any such kernel is a Fourier transform of a uniquely defined probability measure. Later, this work was extended in [17,30] to find explicit features for image classification.

In this paper, we focus on deriving a representation that preserves the diffusion distances between multidimensional data points based on the MGC framework [4,3]. This representation is applicable to process efficiently very large datasets by imposing a Markovian diffusion process to define and represent the non-linear relations between multidimensional data points. It provides a diffusion distance metric that correlates with the intrinsic geometry of the data. The suggested representation and its computational complexity cost per a given data point are invariant to the dataset size.

The paper is organized as follows: Section 2 presents the problem formulation. Section 3 provides an explicit formulation for the diffusion distance. The main results of this paper are established in Sections 4 and 5 that present the suggested low-dimensional embedding of the data and its characterization. Practical demonstration of the proposed methodology is given in Section 6.

2. Problem formulation and mathematical preliminaries

Consider a big dataset $X \subseteq \mathbb{R}^m$ such that for any practical purposes the size of X is considered to be infinite. Without-loss-of-generality, we assume that for all $x \in X$, $\|x\| \leq 1$. Implementation of a kernel method, which uses a full spectral decomposition, becomes impractical when the dataset size is big. Instead, we suggest to represent the given dataset via the density of data points in it using the MGC kernel. In other words, let $q : \mathbb{R}^m \rightarrow [0, 1]$ be the density function of X . We aim to find an explicit embedding function denoted by $f_q : \mathbb{R}^m \rightarrow \mathbb{R}^k$, where k is the embedding dimension such that $m \ll k$. Each member in f_q , which is denoted by $f_q(x)$, $x \in X$, depends on the density function q .

DM provides a multiscale view of the data via a family of geometries that are referred to by *diffusion geometries*. Each geometry is defined by both the associated diffusion metric and the diffusion time parameter t that are linked by $d_\varepsilon^{(t)} : X \times X \rightarrow \mathbb{R}_+$ where ε is a localization parameter. The *diffusion maps* are the associated functions $\Psi^{(t)} : X \rightarrow \mathbb{R}^k$ that embed the data into Euclidean spaces, where the diffusion geometries are preserved such that $\|\Psi^{(t)}(x) - \Psi^{(t)}(y)\| \approx d_\varepsilon^{(t)}(x, y)$, $x, y \in X$.

Given an accuracy requirement $\zeta > 0$, we aim to design an embedding f_q that preserves the diffusion geometry for $t = 1$ such that for all $x, y \in X$,

$$\left| \|f_q(x) - f_q(y)\| - d_\varepsilon^{(1)}(x, y) \right| \leq \zeta. \quad (2.1)$$

We call the embedding f_q the diffusion representation. From the requirement in Eq. (2.1), the Euclidian distance between pairs of representatives approximates the diffusion distance between the corresponding data points over the density of these data points in the MGC-based kernel when $t = 1$. If Eq. (2.1) holds, then f_q preserves the diffusion geometry of the dataset in this sense.

The rest of this section is dedicated to provide additional details regarding the diffusion geometries that utilize the MGC kernel.

2.1. Diffusion geometries

A family of diffusion geometries of a measurable space (X, μ) with a measure μ is determined by imposing a Markov process over the space. Given a non-negative symmetric kernel function $k_\varepsilon : X \times X \rightarrow \mathbb{R}_+$, then an associated Markov process over the data via the stochastic kernel $p_\varepsilon : X \times X \rightarrow \mathbb{R}_+$ is

$$p_\varepsilon(x, y) \triangleq k_\varepsilon(x, y) / \nu_\varepsilon(x), \quad (2.2)$$

where $\nu_\varepsilon : X \rightarrow \mathbb{R}$ is the local volume function. In a discrete setting, it is called the degree function. In a continuous settings, the local volume function is defined by

$$\nu_\varepsilon(x) \triangleq \int_X k_\varepsilon(x, y) d\mu(y). \tag{2.3}$$

The associated Markovian process over X is defined via the conjugate operator of the integral operator $Pq(x) = \int_X p_\varepsilon(x, y)q(y)d\mu(y)$ that is denoted by P^* . Thus, for any initial probability distribution q_0 over X , $q_1 = P^*q_0$ is the probability distribution over X after a single time step. The probability distribution over X after t time steps is given by the t -th power of P^* . Specifically, if the initial probability measure is concentrated in a specific data point $x \in X$, i.e. $q_0 = \delta(x)$, then the probability distribution after t time steps is $(P^*)^t\delta(x)$, denoted also by $p_\varepsilon^{(t)}(x, \cdot)$. Thus, $p_\varepsilon^{(t)}(x, y)$ is the probability that a random walker, which started his walk in $x \in X$, will end in $y \in X$ after t time steps. Based on this, the t -time diffusion geometry is defined by the distances between probability distributions such that for all $x, y \in X$

$$d_\varepsilon^{(t)}(x, y) \triangleq \left\| p_\varepsilon^{(t)}(x, \cdot) - p_\varepsilon^{(t)}(y, \cdot) \right\|_{L^2(\mathbb{R}^m)}. \tag{2.4}$$

Equations (2.1) and (2.4) suggest that the embedding $f_q(x)$ approximately preserves (for $t = 1$) the distance between probability distributions. Such a family of geometries can be defined for any Markovian process and not necessarily for a diffusion process. It is proved in [5] that under specific conditions, the defined Markovian process approximates the diffusion over a manifold from which the dataset X is sampled. If the Markovian process is ergodic, then it has a unique probability distribution $\hat{\nu}_\varepsilon : X \rightarrow \mathbb{R}_+$, to which it converges independently of its initial distribution, namely, for any $y \in X$, $\hat{\nu}_\varepsilon(y) = \lim_{t \rightarrow \infty} p_\varepsilon^{(t)}(x, y)$, independently of x . This probability measure is an L_1 normalization of the local volume function (Eq. (2.3)), i.e. $\hat{\nu}_\varepsilon(y) = \nu_\varepsilon(y) / \int_X \nu_\varepsilon(y) d\mu(y)$.

2.2. Measure-based Gaussian Correlation kernel

The measure-based Gaussian Correlation (MGC) kernel [4,3] defines the affinities between elements in X , which in this context, it is referred to as the *analyzed domain* via their relations with the reference dataset M that is referred to as the *measure domain*. This framework enables to have a flexible representation of X , as long as M , which in some sense characterizes the data, is sufficiently large. For example, it was shown in [2] that in order to compute the MGC-based diffusion to a large, perhaps infinite dataset, it suffices to sample it to compute the associated DM of the sampled dataset X and to extend it to the rest of the dataset via an out-of-sample extension procedure. The sampling rate, shown in [2], depends on the density of M and on the required accuracy. In this sense, X is considered as a grid for the whole dataset.

Mathematically, for the analyzed domain $X \subset \mathbb{R}^m$ and for the measure domain $M \subset \mathbb{R}^m$ with a density function $q : M \rightarrow \mathbb{R}_+$ defined on the measure domain, the MGC kernel $k_\varepsilon : X \times X \rightarrow \mathbb{R}_+$ is defined as

$$k_\varepsilon(x, y) \triangleq \int_{\mathbb{R}^m} g_m(r; x, \frac{\varepsilon}{2}I_m)g_m(r; y, \frac{\varepsilon}{2}I_m)q(r)dr, \tag{2.5}$$

where I_m is an $m \times m$ unit matrix. For a fixed mean vector $\theta \in \mathbb{R}^m$ and a covariance matrix $\Sigma \in \mathbb{R}^{m \times m}$, $g_m(r; \theta, \Sigma) : \mathbb{R}^m \rightarrow \mathbb{R}_+$ is the normalized Gaussian function given by

$$g_m(r; \theta, \Sigma) \triangleq \frac{1}{(2\pi)^{m/2} |\Sigma|^{1/2}} \exp \left\{ -\frac{1}{2}(r - \theta)^T \Sigma^{-1}(r - \theta) \right\}. \tag{2.6}$$

Since the MGC kernel in Eq. (2.5) is symmetric and positive, it can be utilized to establish a Markov process as was described in Section 2.1. The associated diffusion parameters from Eqs. (2.2), (2.3) and (2.4) are p_ε , ν_ε and $d_\varepsilon^{(t)}$, respectively.

3. Explicit forms for the diffusion distance and stationary distribution

In general, the integral in Eq. (2.5) does not have an explicit form. However, for our purposes, we adopt the Gaussian Mixture Model (GMM), which assumes that the density q is a superposition of normal distributions. Under this assumption, q takes the form

$$q(r) = \sum_{j=1}^n a_j g_m(r; \theta_j, \Sigma_j), \sum_{j=1}^n a_j = 1, \tag{3.1}$$

for appropriate mean vectors θ_j and covariance matrices Σ_j , $j = 1, \dots, n$, (see Eq. (2.6)). Estimating Eq. (3.1) is a generally known problem that has been extensively investigated such as in [7,18] with many published implementations. Such an estimation enables to provide an explicit (closed form) representation of the diffusion geometry in Eq. (2.4).

First, a closed form for the inner product $\langle p_\varepsilon^{(1)}(x, \cdot), p_\varepsilon^{(1)}(z, \cdot) \rangle_{L^2(\mathbb{R}^k)}$, $x, z \in X$, is presented. This inner product closed form enables to get an explicit formulation for the first time step of the MGC-based DM distance $d_\varepsilon^{(1)}(x, z)$. This formalism is established in Theorem 3.1, whose proof appears in Appendix A.

Theorem 3.1. *Assume that the GMM assumption in Eq. (3.1) holds and let $W_{x,z}$ be the inner product $W_{x,z} \triangleq \langle k_\varepsilon(x, \cdot), k_\varepsilon(z, \cdot) \rangle_{L^2(\mathbb{R}^k)}$. Denote $D_j \triangleq (\varepsilon^{-1}I_m + (\varepsilon I_m + 4\Sigma_j)^{-1})^{-1}$ and $c_j(x) \triangleq D_j(\varepsilon^{-1}x + (\varepsilon I_m + 4\Sigma_j)^{-1}(2\theta_j - x))$. Let $\tilde{\Sigma}_j$ be defined by*

$$\tilde{\Sigma}_j \triangleq (\varepsilon/2)I_m + \Sigma_j. \tag{3.2}$$

Then, for any $x, z \in X$, the kernel affinity $k_\varepsilon(x, z)$, the stationary distribution $\nu_\varepsilon(x)$ and the inner product $W_{x,z}$ have explicit forms given by

$$k_\varepsilon(x, z) = \sum_{j=1}^n a_j g_m(x; \theta_j, \tilde{\Sigma}_j) g_m(z; c_j(x), D_j), \tag{3.3}$$

$$\nu_\varepsilon(x) = \sum_{j=1}^n a_j g_m(x; \theta_j, \tilde{\Sigma}_j) \tag{3.4}$$

and

$$W_{x,z} = \sum_{j,i=1}^n a_j a_i g_m(x; \theta_j, \tilde{\Sigma}_j) g_m(z; \theta_i, \tilde{\Sigma}_i) g_m(c_j(x); c_i(z), D_j + D_i). \tag{3.5}$$

Combination of Theorem 3.1 with Eqs. (2.2), (2.3) and (2.4) provides a closed form solution for the first time step ($t = 1$) diffusion metric to be

$$d_\varepsilon^{(1)}(x, y) = \frac{W_{x,x}}{\nu_\varepsilon(x)\nu_\varepsilon(x)} + \frac{W_{z,z}}{\nu_\varepsilon(z)\nu_\varepsilon(z)} - \frac{2W_{x,z}}{\nu_\varepsilon(x)\nu_\varepsilon(z)}. \tag{3.6}$$

Moreover, by combining Eqs. (2.6), (3.1) and (3.4) we get $\|\nu_\varepsilon\|_1 = 1$. Thus, the local volume function is equal to the stationary distribution (Eq. (2.3)) of the MGC-based Markov process.

4. Explicit Diffusion Maps of the analyzed domain

The diffusion distance provides a relation between a pair of data points in the analyzed domain. In this section, we find a representation of any data point in the analyzed domain that preserves the diffusion

distance relation. In the following section we assume that the covariances in the GMM (Eq. (3.1)) are all identical, namely, $\Sigma_j = \Sigma$, $j = 1, \dots, n$, and that the analyzed domain X is a subset of the unit ball in \mathbb{R}^m . The Taylor extension

$$e^\gamma = \sum_{i=0}^{\infty} \frac{\gamma^i}{i!}, \tag{4.1}$$

where γ is a scalar is reformulated with $\gamma = -\frac{1}{2}\|x - y\|^2$ to be

$$e^{-\frac{1}{2}\|x-y\|_2^2} = e^{-\frac{1}{2}\|x\|^2} e^{-\frac{1}{2}\|y\|^2} \sum_{i=0}^{\infty} \frac{(x^T y)^i}{i!}. \tag{4.2}$$

The exponent in Eq. (4.2) is formulated as an inner product such that

$$e^{-\frac{1}{2}\|x-y\|_2^2} = e^{-\frac{1}{2}\|x\|^2} e^{-\frac{1}{2}\|y\|^2} \phi(x)^T \phi(y), \tag{4.3}$$

where $\phi(x)$ is a vector of infinite length whose first term ($i = 1$) is $[\phi(x)]_1 = 1$ and for $i > 1$ we have the generating function $[\phi(x)]_{i+1} = \frac{1}{\sqrt{i+1}} [\phi(x)]_i \otimes x$ where \otimes is the Kronecker product operator. In the following we consider the vector $[\phi(x)]_i$ as a single term of $\phi(x)$ that corresponds to the i th term from the relevant Taylor expansion.

For the diffusion representation, we reformulate the inner product in Eq. (3.5) to be a vector multiplication of two vectors. The first depends only on x and the second depends only on z . Let $h(x)$ be a n -dimensional vector with the entries $[h(x)]_j \triangleq a_j g_m(x; \theta_j, \tilde{\Sigma})$ (Eq. (3.2)). Then, the inner product in Eq. (3.5) is given by the matrix multiplication

$$W_{x,z} = h(x)^T G h(z), \tag{4.4}$$

where the j and k entry in the $n \times n$ matrix G is given by $g_m(c_j(x); c_k(z), D_j + D_k)$. The scalar $[G]_{j,k}$ depends on both x and z . From the assumption on the GMM structure, we have $D_j = D_k = D$. Let $\hat{c}_j(x) \triangleq (2D)^{-\frac{1}{2}} c_j(x)$ then $[G]_{j,k}$ can be reformulated by

$$[G]_{j,k} = \frac{1}{(2\pi)^{m/2} |2D|^{1/2}} \exp \left\{ -\frac{1}{2} (\|\hat{c}_j(x) - \hat{c}_k(z)\|^2) \right\}. \tag{4.5}$$

By reformulating Eq. (4.5) using the inner product in Eq. (4.3), we get

$$[G]_{j,k} = \frac{1}{(2\pi)^{m/2} |2D|^{1/2}} e^{-\frac{1}{2}\|\hat{c}_j(x)\|^2} e^{-\frac{1}{2}\|\hat{c}_k(z)\|^2} \phi(\hat{c}_j(x))^T \phi(\hat{c}_k(z)). \tag{4.6}$$

Hence, the matrix G is the result from an outer-product of the form

$$G = \frac{(2\pi)^{-m/2}}{|2D|^{1/2}} F^T F, \tag{4.7}$$

where

$$F = \left[e^{-\frac{1}{2}\|\hat{c}_1(z)\|^2} \phi(\hat{c}_1(z)), \dots, e^{-\frac{1}{2}\|\hat{c}_n(z)\|^2} \phi(\hat{c}_n(z)) \right]. \tag{4.8}$$

By substituting the outer-product from Eq. (4.8) into Eq. (4.4), we get

$$\frac{h(x)^T Gh(z)}{\nu_\varepsilon(x)\nu_\varepsilon(z)} = f(x)^T f(z), \tag{4.9}$$

where

$$f^T(x) = \frac{(2\pi)^{-m/4}}{|2D|^{1/4}\nu_\varepsilon(x)} h(x)^T F^T \tag{4.10}$$

and the infinite dimensional vector $f(x)$ is the weighted sum of the inner product in the exponent decompositions (Eq. (4.3)).

Lemma 4.1. *The diffusion representations $f(x)$ and $f(z)$ in Eq. (4.10) preserve the diffusion distance in Eq. (3.6).*

Proof. Equation (2.4) together with Eqs. (3.6), (4.4) and (4.9) yield

$$\|f(x) - f(z)\|_{\ell^2(\mathbb{R}^m)} = \|p_\varepsilon(x, \cdot) - p_\varepsilon(z, \cdot)\|_{L^2(\mathbb{R}^m)}. \quad \square \tag{4.11}$$

From Lemma 4.1 we conclude that $f(x)$ is an embedding that preserves the diffusion geometry (for $t = 1$) where $\zeta = 0$ in Eq. (2.1).

5. Truncated diffusion representation

The diffusion representation $f(x)$ in Lemma 4.1 is an infinite size vector. In this section, this vector is truncated and an error estimate for the resulted truncated diffusion distance is provided. The infinite number of entries in $f(x)$ originates from the Gaussian decomposition multiplication in Eq. (4.3). Let e^γ be a Taylor expanded function on the interval $0 \leq \gamma \leq \gamma_{max}$, where γ_{max} is an arbitrary upper limit of γ . Then, the Lagrange remainder of the l th order Taylor (Maclaurin) expansion of e^γ on this interval is given by

$$R_l(\gamma) = \frac{e^b}{(l+1)!} \gamma^{l+1}, \quad 0 \leq b \leq \gamma_{max}. \tag{5.1}$$

Since $e^b \leq e^{\gamma_{max}}$ we have

$$R_l(\gamma) \leq \frac{e^{\gamma_{max}}}{(l+1)!} \gamma^{l+1}. \tag{5.2}$$

Let $f_l(x)$ be the l th order truncated version of $f(x)$ from Eq. (4.10) that is given by,

$$f_l^T(x) = \frac{(2\pi)^{-m/4}}{|2D|^{1/4}\nu_\varepsilon(x)} h(x)^T \hat{F}^T \tag{5.3}$$

where the matrix \hat{F} is given by

$$\hat{F}(x) = \begin{bmatrix} e^{-\frac{1}{2}\|\hat{c}_1(x)\|^2} \hat{\phi}_l(\hat{c}_1(x)) \\ \vdots \\ e^{-\frac{1}{2}\|\hat{c}_n(x)\|^2} \hat{\phi}_l(\hat{c}_n(x)) \end{bmatrix} \tag{5.4}$$

and $\hat{\phi}_l(x)$ is a vector that contains the first l terms from $\phi_l(x)$. Furthermore, let $\phi_{l+1}(x)$ be the vector that contains the truncated terms from the $(l + 1)$ th term in $\phi(x)$ to ∞ . Then, the diffusion distance error is given by

$$\|f(x) - f_l(x)\|_F^2 = \frac{(2\pi)^{-m/2}}{|2D|^{1/2} \nu_\varepsilon(x)^2} \|F_l(x)^T h(x)\|^2, \tag{5.5}$$

where the matrix F_l is given by

$$F_l(x) = \begin{bmatrix} e^{-\frac{1}{2}\|\hat{c}_1(x)\|^2} \phi_{l+1}(\hat{c}_1(x)) \\ \vdots \\ e^{-\frac{1}{2}\|\hat{c}_n(x)\|^2} \phi_{l+1}(\hat{c}_n(x)) \end{bmatrix}. \tag{5.6}$$

The Frobenius matrix norm and the Euclidean vector norm are compatible [19] such that

$$\|F_l(x)^T h(x)\|_2^2 \leq \|F_l(x)\|_F^2 \|h(x)\|_2^2. \tag{5.7}$$

From Eqs. (2.6) and (3.1) we have $\sum_{j=1}^n a_j = 1$ and $g_m(r; \theta_j, \Sigma_j) \leq 1$, then,

$$\|h(x)\| \leq 1. \tag{5.8}$$

Furthermore, from the Frobenius matrix norm definition we have

$$\|F_l(x)\|_F^2 = \sum_i e^{-\|\hat{c}_i(x)\|^2} \phi_{l+1}(\hat{c}_i(x))^T \phi_{l+1}(\hat{c}_i(x)). \tag{5.9}$$

The product $\phi_{l+1}(\hat{c}_i(x))^T \phi_{l+1}(\hat{c}_i(x))$ is the remainder of the Taylor series. By using Eq. (5.2), we get

$$\phi_{l+1}(\hat{c}_i(x))^T \phi_{l+1}(\hat{c}_i(x)) \leq \frac{e^{\|\hat{c}_i(x)\|^2}}{(l + 1)!} \|\hat{c}_i(x)\|^{2(l+1)}, \tag{5.10}$$

since in this case $\gamma_{max} = \|\hat{c}_i(x)\|^2$. By substituting Eq. (5.10) into Eq. (5.9) and finding the argmax of $\|\hat{c}_i(x)\|$ over x , we get

$$\|F_l(x)\|_F^2 \leq \sum_{i=1}^n \frac{\|\hat{c}_i(x_i^*)\|^{2(l+1)}}{(l + 1)!}, \tag{5.11}$$

where x_i^* is the solution to the Trust Region Problem (TRP) of the form

$$x_i^* = \max_{\|x\| \leq 1} \hat{c}_i(x)^T \hat{c}_i(x) = \max_{\|x\| \leq 1} \|Ax - b\|^2, \tag{5.12}$$

where $A = 2^{-\frac{1}{2}} D^{\frac{1}{2}} (\varepsilon^{-1} I_m - (\varepsilon I_m + 4\Sigma_j)^{-1})$ and $b = 2^{\frac{1}{2}} \theta_j D^{\frac{1}{2}} (\varepsilon I_m + 4\Sigma_j)^{-1}$. TRP in Eq. (5.12) is widely investigated in the literature. Important properties are given in [10,9,20]. In Appendix B, we provide complementary theoretical details about this problem. In the following, we assume that the solution x_i^* was found such that the maximal value of $\hat{c}_i(x)^T \hat{c}_i(x)$ is known.

Lemma 5.1 provides an additional bound for $\|F_l(x)\|_F^2$.

Lemma 5.1. Assume that $F_l(x)$ is defined according to Eq. (5.6). Then, for all $x \in X$

$$\|F_l(x)^T\|_F^2 \leq \sum_{i=1}^n \left[1 - e^{-\|\hat{c}_i(x_i^*)\|^2} \left(\sum_{j=0}^l \frac{\|\hat{c}_i(x_i^*)\|^{2j}}{j!} \right) \right]. \tag{5.13}$$

Proof. The inner product $\phi_{l+1}(\hat{c}_i(x))^T \phi_{l+1}(\hat{c}_i(x))$ is the reminder of the Taylor series of $e^{\|\hat{c}_i(x)\|^2}$ with l terms. Hence,

$$\phi_{l+1}(\hat{c}_i(x))^T \phi_{l+1}(\hat{c}_i(x)) = e^{\|\hat{c}_i(x)\|^2} - \left(\sum_{j=0}^l \frac{\|\hat{c}_i(x)\|^{2j}}{j!} \right). \tag{5.14}$$

Substituting Eq. (5.14) into Eq. (5.9) yields,

$$\|F_l(x)^T\|_F^2 = \sum_i \left(1 - e^{-\|\hat{c}_i(x)\|^2} \sum_{j=0}^l \frac{\|\hat{c}_i(x)\|^{2j}}{j!} \right) = \sum_i \varphi(\|\hat{c}_i(x)\|_2^2), \tag{5.15}$$

where $\varphi(z) \triangleq \left(1 - e^{-z} \sum_{j=0}^l \frac{z^j}{j!} \right)$. The derivative $\frac{d\varphi}{dz}$ is given by

$$\frac{d\varphi}{dz} = e^{-z} \frac{z^l}{l!}. \tag{5.16}$$

From Eq. (5.16) we get that $\frac{d\varphi}{dz} \geq 0$ on the interval $z \in [0, \|\hat{c}_i(x_i^*)\|_2^2]$ (x_i^* was defined in Eq. (5.11)). Hence, $\varphi(z)$ is a monotonic function on the interval and $\|F_l(x)^T\|_F^2$ is bounded by the maximal values on the interval where $\varphi(z)$ is defined by,

$$\|F_l(x)^T\|_F^2 \leq \sum_i \varphi(\|\hat{c}_i(x_i^*)\|_2^2). \quad \square \tag{5.17}$$

Let $G_{b_1}(l) \triangleq \sum_i \frac{\|\hat{c}_i(x_i^*)\|^{2(l+1)}}{(l+1)!}$ be the bound from Eq. (5.11) and let $G_{b_2}(l) \triangleq \sum_i \varphi(\|\hat{c}_i(x_i^*)\|_2^2)$ be the bound from Eq. (5.13). Furthermore, let $B_{l_{max}}$ be the bound on $\|F_l(x)^T\|_F^2 \leq B_{l_{max}}$ that takes into consideration the bounds in Eqs. (5.13) and (5.11) such that

$$B_{l_{max}} = \min(G_{b_1}(l), G_{b_2}(l)). \tag{5.18}$$

By examining the stationary distribution component $\nu_\varepsilon(x)$ in the denominator of Eq. (5.5), we assume that a given a minimal threshold ν_{min} establishes that any $\nu_\varepsilon(x)$, which is smaller than ν_{min} , is negligible. By substituting the bounds in Eqs. (5.8) and (5.18) into Eq. (5.5) and using the definition of ν_{min} we get

$$\|f(x) - f_l(x)\|_2^2 \leq \eta \triangleq \frac{(2\pi)^{-m/2}}{|2D|^{1/2} \nu_{min}^2} B_{l_{max}}. \tag{5.19}$$

By imposing the requirement in Eq. (2.1) for l in the form

$$\|f(x) - f_l(x)\|_2^2 \leq \frac{\zeta^2}{4} \tag{5.20}$$

we find the minimal l_{max} such that

$$\frac{(2\pi)^{-m/2}}{|2D|^{1/2} \nu_{min}^2} B_{l_{max}} \leq \frac{\zeta^2}{4}. \tag{5.21}$$

The number l_{max} of Taylor terms of the Taylor expansion is derived by computing Eq. (5.21) for an increasing integers from 1 until the requirement in Eq. (5.21) is satisfied. We expect that for a practical use, 100 iterations are sufficient.

The error in the estimated embedding is bounded. Hence, the error in the estimated diffusion distance in Eq. (2.1) is also bounded as Lemma 5.2 proposes.

Lemma 5.2. *Let $x, y \in M$ be any two data points. Assume we have a properly trained GMM and the number of Taylor terms l is designed according to Eq. (5.21). Then, the error in the diffusion distance (Eq. (2.1)) is bounded by ζ .*

Proof. Let $f(x)$ and $f(y)$ be the corresponding embedding of x and y , respectively. Furthermore, let $f_l(x)$ and $f_l(y)$ be the corresponding truncated embedding of x and y , respectively. From Eq. (5.20) we have $\|f(x) - f_l(x)\| \leq \frac{\zeta}{2}$ and $\|f(y) - f_l(y)\| \leq \frac{\zeta}{2}$. Therefore, from the triangle inequality

$$\begin{aligned} \|f(x) - f(y)\| &\leq \|f(x) - f_l(x)\| + \|f_l(x) - f_l(y)\| + \|f_l(y) - f(y)\|, \\ &\leq \|f_l(x) - f_l(y)\| + \zeta, \\ \|f_l(x) - f_l(y)\| &\leq \|f(x) - f_l(x)\| + \|f(x) - f(y)\| + \|f_l(y) - f(y)\|, \\ &\leq \|f(x) - f(y)\| + \zeta. \end{aligned}$$

Hence, $\|f(x) - f(y)\| - \zeta \leq \|f_l(x) - f_l(y)\| \leq \|f(x) - f(y)\| + \zeta$. \square

5.1. The computational complexity to compute the diffusion representation

Assume that GMM-based $q(r)$ and l_{max} are designed according to Eq. (5.21). Assume that l_{max} is the length of the representation that is dictated by the dimension of the measurement m such that $k = \sum_{i=1}^{l_{max}} m^{i-1} + 1$. The computation of $\phi(x)$ in Lemma 4.1 has a computational complexity of $O(m^{l_{max}-1}n)$ due to the multiplication $h(x)^T F(x)$ and due to the truncation of $\phi(x)$ that has l_{max} Taylor terms. We also assume that $l_{max} > 2$. Hence, the computational complexity for the computation of the terms in the vector $h(x)$ is negligible compared to the computational complexity of $\phi(x)$.

6. Experimental results

This section presents several examples that demonstrate the principles of the MGC-based closed-form embedding. The first example presents an analysis of a density function for which the stationary distribution is analytically known. The closed-form stationary distribution in this case is compared to the analytical stationary distribution. The second example compares the analytical distance computation with the corresponding distance that is computed via the closed-form embedding defined in Eq. (4.10).

6.1. Example I: a data analysis with an analytically known stationary distribution function

Let the density function $q(r) \in \mathbb{R}^2$ includes two flat squares one above the other with probability $\frac{1}{5}$ to draw samples from the lower square and $\frac{4}{5}$ to draw samples from the upper square. In other words,

$$q(r) = \frac{1}{5} \chi_{[0,1] \times [0,1]}(r) + \frac{4}{5} \chi_{[3,4] \times [3,4]}(r) \tag{6.1}$$

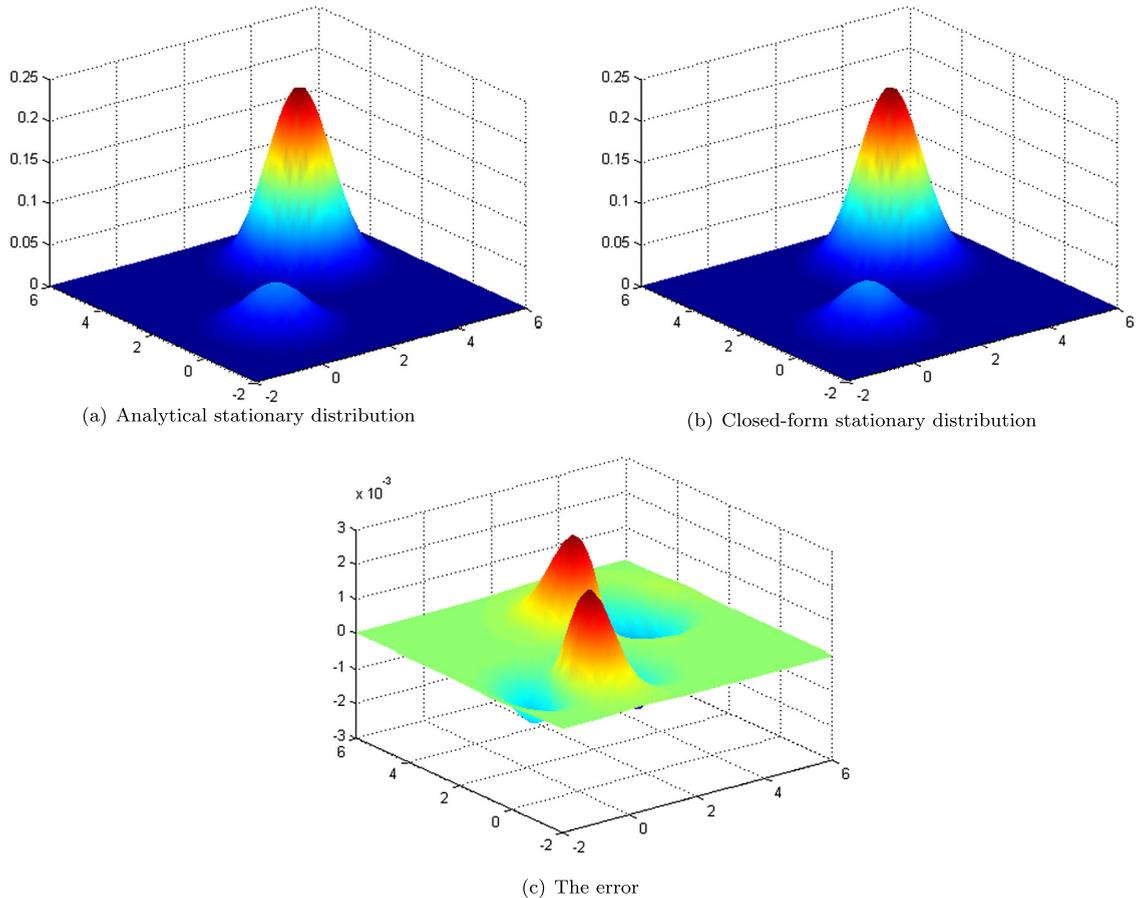


Fig. 6.1. The difference (error) between the closed-form stationary distribution (Eq. (3.4)) and the analytical stationary distribution (Eq. (6.2)). (a) Analytical stationary distribution. (b) Closed-form stationary distribution. (c) The error between (a) and (b).

where $\chi_{[a,b] \times [c,d]}$ is the indicator function for the square $abcd$. Equation (2.3) formulates the stationary distribution computation. Given $\varepsilon = 1$, the integration in Eq. (2.3) is analytically solved by

$$\nu(x_1, x_2) = 0.2H(0, 1, x_1, x_2) + 0.8H(3, 4, x_1, x_2), \quad (6.2)$$

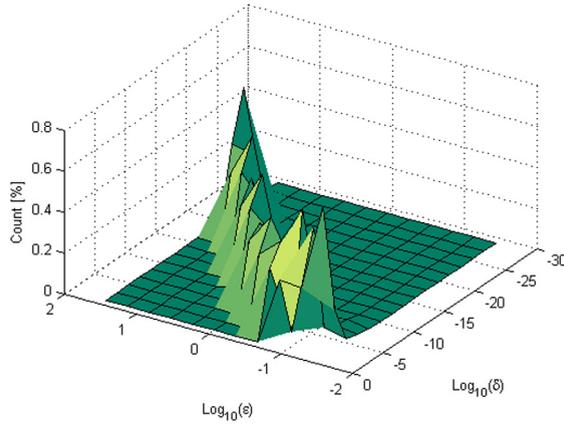
where $H(a, b, x_1, x_2)$ is given by

$$H(a, b, x_1, x_2) \triangleq \frac{1}{4} (\operatorname{erf}(b - x_1) - \operatorname{erf}(a - x_1)) (\operatorname{erf}(b - x_2) - \operatorname{erf}(a - x_2)),$$

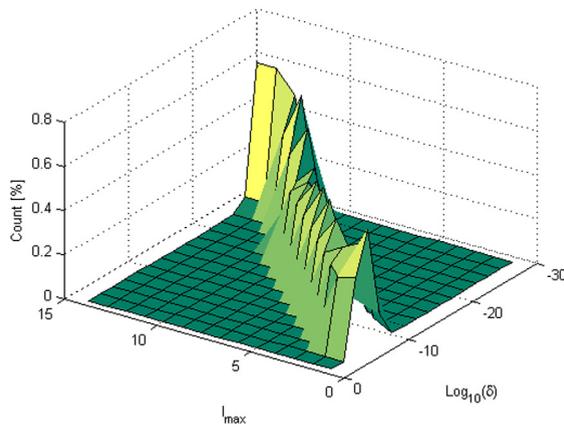
and $\operatorname{erf}(x)$ is the Gauss error function.

Given a trained GMM, the stationary distribution is computed by using Eq. (3.4). First, 2000 data points were randomly selected from the distribution in Eq. (6.1). Then, a 48×48 grid was constructed to compute the stationary distribution via Eq. (3.4) based on the analytical solution in Eq. (6.2).

Fig. 6.1(b) presents the stationary distribution (Eq. (3.4)) with a minor distortion compared to the analytical distribution (Eq. (6.2)) presented in Fig. 6.1(a). The error, which is presented in Fig. 6.1(c), is a result from the GMM-based training on a small set of data points. The difference between the closed form stationary distribution and the analytical stationary distribution is in the order of 1% and it is the result from the GMM training error. In the following, we assume that the GMM-based training error is negligible.



(a) Histogram of δ vs. ϵ and $l_{max} = 3$



(b) Histogram of δ vs. l_{max} and $\epsilon = 1$

Fig. 6.2. The histogram of the distance error δ as a function of l_{max} and ϵ . (a) As a function of $\epsilon \in [2^{-5}, 2^5]$ and fixed $l_{max} = 3$. (b) As a function of $l_{max} \in [1, 14]$ and fixed $\epsilon = 1$.

6.2. Example II: estimation of the diffusion distance error originated from the truncated representation (Eq. (5.3))

In this section, we show the diffusion distance error from the proposed truncated representation (Eq. (5.3)) of a specific data. Figs. 6.2 and 6.3 compare between the diffusion distance in Eq. (3.6) and the corresponding distance computed by Eq. (4.11), respectively, using the truncated representations $f_l(x)$ and $f_l(z)$ computed by

$$\delta = \left| \left(\frac{W_{x,x}}{\nu_\epsilon(x)\nu_\epsilon(x)} + \frac{W_{z,z}}{\nu_\epsilon(z)\nu_\epsilon(z)} - \frac{2W_{x,z}}{\nu_\epsilon(x)\nu_\epsilon(z)} \right) - \|f_l(x) - f_l(z)\| \right|^2. \tag{6.3}$$

A data of 3000 data points was randomly selected from the distribution in Eq. (6.1). Furthermore, the bound on the worst-case diffusion distance error η is computed using Eq. (5.19) for the corresponding l_{max} and ϵ where $\nu_{min} = 10^{-3}$.

Fig. 6.2 displays the normalized histogram of δ (Eq. (6.3)) for different values of $\epsilon \in [2^{-5}, 2^5]$ and different values of $l_{max} \in [1, 14]$. Figs. 6.2(a) and 6.2(b) display a reduction of the distance error δ as a function of ϵ and l_{max} , respectively. As ϵ increases, δ is reduced in Fig. 6.2(a). Furthermore, as the number of Taylor terms increases, δ is reduced in Fig. 6.2(b).

Fig. 6.3 compares the bound η to the worst-case of δ from the 3000 randomized data points as a function of $\epsilon \in [2^{-5}, 2^5]$ and as a function of $l_{max} \in [1, 14]$. Fig. 6.3(a) compares between the worst-case of the

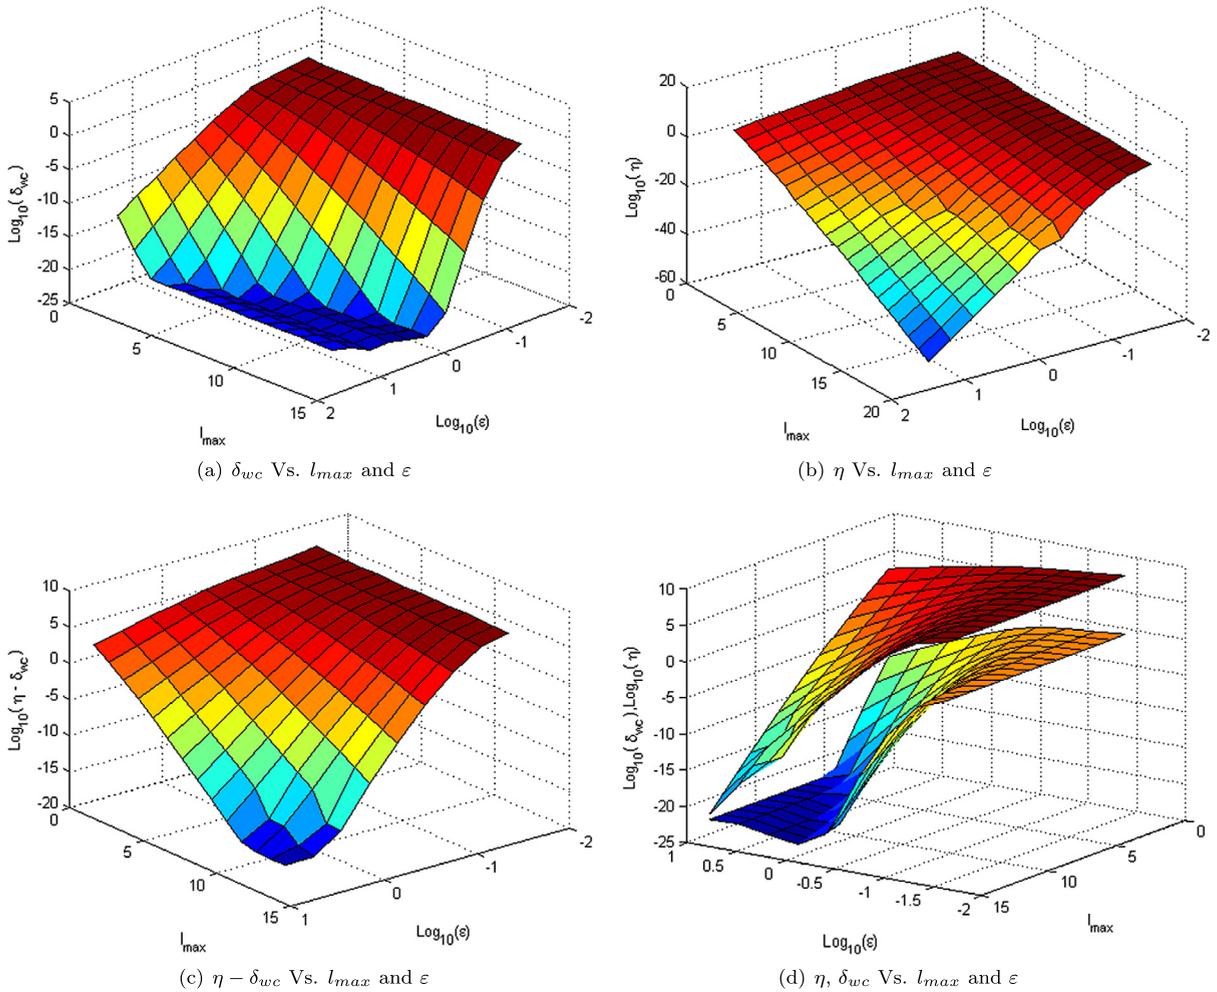


Fig. 6.3. Comparison between the worst-case of the diffusion distance from the data error denoted by δ_{wc} and the corresponding bound η as a function of ϵ and as a function of l_{max} . (a) δ_{wc} as a function of ϵ and l_{max} . (b) η as a function of ϵ and l_{max} . (c) The difference $\eta - \delta_{wc}$ as a function of ϵ and l_{max} . (d) The bound η and the worst-case error δ_{wc} as a function of ϵ and as a function of l_{max} .

diffusion distance error δ_{wc} on the dataset and the corresponding bound η as a function of ϵ and as a function of l_{max} .

The worst-case error δ_{wc} in Fig. 6.3(a) decreases as l_{max} and ϵ increase. Similarly, the corresponding bound in Fig. 6.3(b) decreases as l_{max} and ϵ increase. The difference between δ_{wc} and the bound η is shown in Figs. 6.3(c) and 6.3(d), respectively. η bounds δ_{wc} from above.

7. Conclusions

The presented methodology provides an alternative to a non-parametric kernel method approach for obtaining data representations via spectral decompositions of a big kernel operator or matrix with finite settings. The presentation of our approach is based on the MGC diffusion kernel [4] and on the resulting measure-based DM embedding obtained by its decomposition. We showed that when the underlying measure is modeled by a GMM, an equivalent embedding, which preserves the diffusion geometry of the data, can be computed without the need to decompose the full kernel. Furthermore, this computation does not depend on the size of the dataset, therefore, it is suitable for modern scenarios where large amounts of data are available for analysis.

The proposed data representation is achieved by finding and decomposing an explicit form of the measure-based diffusion distance in terms of GMM. The decomposition is done by utilizing a Taylor decomposition and by an associated bound for the number of Taylor terms that was formulated for a given measure-based diffusion distance error tolerance requirement. The resulted representation length grows exponentially with the number of Taylor terms, hence, for a practical use, the error tolerance requirement should consider the data dimensionality.

7.1. Future work

Future work will focus on extending the proposed methodology to a more general settings in the following ways below to reduce the representation dimension and to generalize the explicit form for additional kernels.

The utilized generating function for the inner products decomposition is based on the Kronecker product and hence incorporates its structure. This structure can be utilized to reduce the representation length by using a variant of random projection. Thus, the dimensionality of the representation can be reduced.

The proposed methodology is valid for additional Gaussian based kernels such as the standard DM kernel. The formulation of an explicit form for this kernel is straightforward for the above analysis. Hence, similar considerations may lead to an explicit form for the diffusion distance in this case. This research direction may lead to more general explicit forms that can be used to represent datasets.

We provided an explicit form for the inner product associated with the measure-based diffusion distance, therefore, we are able to explicitly and efficiently compute a matrix of pair-wise distances that consider the entire dataset. This matrix can be decomposed via the application of MDS and the proposed representation can be projected to the resulted MDS space that may have a low dimension due to the involved spectral decomposition.

The explicit resulted representation form can be used to find the most distant data points in terms of the diffusion distance by solving a proper optimization problem. The resulted data points can be used to form a dictionary or to provide data points samples for any general use.

The proposed explicit map and the explicit diffusion distance are differentiable, hence, both may be used to characterize data points in terms of change rate in the diffusion maps as a function of a given specific feature. This characterization may be helpful in identifying both the anomaly and the associated features that are the source for anomaly occurrence.

The proposed representation corresponds to $t = 1$. An additional interesting research direction is to generalize the analysis to any $t > 1$ to allow a time diffusion multiscale analysis of datasets.

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), Blavatnik Computer Science Research Fund and Blavatnik ICRC Funds.

Appendix A. Proof of Theorem 3.1

First, we introduce some basic Gaussian distribution related equalities. Denote $\Lambda_{ij} \triangleq (\Sigma_i^{-1} + \Sigma_j^{-1})^{-1}$, $\psi_{ij} \triangleq \Lambda_{ij}(\Sigma_i^{-1}\theta_i + \Sigma_j^{-1}\theta_j)$ and $\Psi_{ij} \triangleq \Sigma_i + \Sigma_j$. In the following, we will use the normal distribution identity

$$g_m(y; \theta_i, \Sigma_i) \cdot g_m(y; \theta_j, \Sigma_j) = g_m(y; \psi_{ij}, \Lambda_{ij}) \cdot g_m(\theta_i; \theta_j, \Psi_{ij}). \tag{A.1}$$

Particularly,

$$g_m^2(y; \theta, \Sigma) = (2\pi)^{-d/2} |2\Sigma|^{-1/2} g_m(y; \theta, (1/2)\Sigma). \tag{A.2}$$

By combining Eq. (A.1) and the fact that the Gaussian function in Eq. (2.6) is normalized to a L^2 unit norm, we have

$$\int g_m(y; \theta_i, \Sigma_i) \cdot g_m(y; \theta_j, \Sigma_j) dy = g_m(\theta_i; \theta_j, \Psi_{ij}). \quad (\text{A.3})$$

Another useful Gaussian relation is due to rescaling of the covariance matrix Σ and the mean vector θ in Eq. (2.6) such that

$$g_m\left(\frac{x+y}{2}; \theta, \Sigma\right) = 4^{d/2} g_m(y+x; 2\theta, 4\Sigma). \quad (\text{A.4})$$

Now, we are ready for the proof of Theorem 3.1.

Proof. An equivalent representation of the MGC kernel in Eq. (2.5), as was proved in [4], is

$$k_\varepsilon(x, y) = g_m(y; x, \varepsilon I_m) \int g_m\left(r; \frac{x+y}{2}, \frac{\varepsilon}{4} I_m\right) q(r) dr. \quad (\text{A.5})$$

By substituting the GMM model (Eq. (3.1)) into Eq. (A.5), we get

$$k_\varepsilon(x, y) = g_m(y; x, \varepsilon I_m) \sum_{j=1}^n a_j \int g_m\left(r; \frac{x+y}{2}, \frac{\varepsilon}{4} I_m\right) \cdot g_m(r; \theta_j, \Sigma_j) dr. \quad (\text{A.6})$$

Thus, from Eq. (A.3), we get

$$k_\varepsilon(x, y) = g_m(y; x, \varepsilon I_m) \sum_{j=1}^n a_j g_m\left(\frac{x+y}{2}; \theta_j, \frac{\varepsilon}{4} I_m + \Sigma_j\right). \quad (\text{A.7})$$

Due to Eqs. (A.1) and (A.4), we get

$$\begin{aligned} g_m(y; x, \varepsilon I_m) g_m\left(\frac{x+y}{2}; \theta_j, \frac{\varepsilon}{4} I_m + \Sigma_j\right) &= 4^{d/2} g_m(y; x, \varepsilon I_m) g_m(y; 2\theta_j - x, \varepsilon I_m + 4\Sigma_j) \\ &= 4^{d/2} g_m(x; 2\theta_j - x, 2\varepsilon I_m + 4\Sigma_j) g_m(y; c_j(x), D_j) \\ &= g_m(x; \theta_j, (\varepsilon/2)I_m + \Sigma_j) g_m(y; c_j(x), D_j) \end{aligned}$$

where $D_j \triangleq (\varepsilon^{-1}I_m + (\varepsilon I_m + 4\Sigma_j)^{-1})^{-1} = \varepsilon(2\varepsilon I_m + 4\Sigma_j)^{-1}(\varepsilon I_m + 4\Sigma_j)$ and $c_j(x) \triangleq D_j(\varepsilon^{-1}x + (\varepsilon I_m + 4\Sigma_j)^{-1}(2\theta_j - x))$. Thus, for $\tilde{\Sigma}_j \triangleq (\varepsilon/2)I_m + \Sigma_j$, we get

$$k_\varepsilon(x, y) = \sum_{j=1}^n a_j g_m(x; \theta_j, \tilde{\Sigma}_j) g_m(y; c_j(x), D_j). \quad (\text{A.8})$$

Thus, an immediate consequence from Eq. (2.3) is $\nu_\varepsilon(x) = \sum_{j=1}^n a_j g_m(x; \theta_j, \tilde{\Sigma}_j)$. Finally, since

$$k_\varepsilon(x, y) \cdot k_\varepsilon(z, y) = \sum_{j,i=1}^n g_m(x; \theta_j, \tilde{\Sigma}_j) g_m(y; c_j(x), D_j) g_m(z; \theta_i, \tilde{\Sigma}_i) g_m(y; c_i(z), D_i), \quad (\text{A.9})$$

then by utilizing Eq. (A.3), we get

$$\begin{aligned} \langle k_\varepsilon(x, \cdot), k_\varepsilon(z, \cdot) \rangle_{L^2(\mathbb{R}^d)} &= \int k_\varepsilon(x, y)k_\varepsilon(z, y)dy \\ &= \sum_{j,i=1}^n a_j a_i g_m(x; \theta_j, \tilde{\Sigma}_j)g_m(z; \theta_i, \tilde{\Sigma}_i)g_m(c_j(x); c_k(z), D_j + D_i). \quad \square \end{aligned}$$

Appendix B. The Trust Region optimization problem (TRP)

The trust region problem (or subproblem) is an optimization problem of the form

$$w_{opt} = \arg \min_{\|w\|_2 \leq \rho} \{w^* B w + 2w^* u\}, \tag{B.1}$$

where $B \in R^{m \times m}$ is a symmetric matrix and $u \in R^m$. In the diffusion representation case, we have

$$w_{opt} = \arg \max_{\|w\|_2 \leq \rho} \|Aw + b\|_2^2 = \arg \min_{\|w\|_2 \leq \rho} \{w^* (-A^T A) w - 2w^* A^T b - b^T b\}. \tag{B.2}$$

Define the Lagrange function

$$L(w, \tau) = w^* (-A^T A) w - 2w^T A^T b - b^T b + \tau (w^* w - \rho^2). \tag{B.3}$$

Let $B = A^T A$ and $u = A^T b$. A necessary and sufficient conditions for w to be an optimal solution for the TRP is that there exists $\tau_{opt} \geq 0$ such that the following are satisfied [26]

- $w_{opt}^T w_{opt} \leq \rho^2$ (Feasibility),
- $(\tau_{opt} I - B) w_{opt} = u$ (Stationarity),
- $\tau_{opt} (w_{opt}^* w_{opt} - \rho^2) = 0$ (Complementary Slackness),
- $\tau_{opt} I - B \succeq 0$ (2nd order necessary conditions).

Furthermore, when $\tau_{opt} I - B > 0$, then the optimal solution w_{opt} is unique. Depending on the values of A , b and ρ , different ways of solving the TRP need to be considered. Different cases may occur and are referred in the literature as the easy case and the hard case. The hard case or near hard case is what causes numerical difficulties in solving the TRP. For the easy case, the optimization problem in Eq. (B.1) has a single solution. However, for the hard case or near hard case, this optimization problem had multiple solutions that maximizes the optimization objective. For more details see [32].

References

- [1] M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation, *Neural Comput.* 15 (6) (2003) 1373–1396.
- [2] A. Bermanis, M. Salhov, G. Wolf, A. Averbuch, Measure-based diffusion grid construction and high-dimensional data discretization, *Appl. Comput. Harmon. Anal.* 40 (2016) 207–228.
- [3] A. Bermanis, G. Wolf, A. Averbuch, Measure-based diffusion kernel methods, in: *SampTA 2013: 10th International Conference on Sampling Theory and Applications*, Bremen, Germany, 2013.
- [4] A. Bermanis, G. Wolf, A. Averbuch, Diffusion-based kernel methods on Euclidean metric measure spaces, *Appl. Comput. Harmon. Anal.* 41 (1) (2016) 190–213.
- [5] R.R. Coifman, S. Lafon, Diffusion maps, *Appl. Comput. Harmon. Anal.* 21 (1) (2006) 5–30.
- [6] T. Cox, M. Cox, *Multidimensional Scaling*, Chapman and Hall, London, UK, 1994.
- [7] N.E. Day, Estimating the components of a mixture of normal distributions, *Biometrika* 56 (3) (1969) 463–474.
- [8] D.L. Donoho, C. Grimes, Hessian eigenmaps: new locally linear embedding techniques for high dimensional data, *Proc. Natl. Acad. Sci. USA* 100 (May 2003) 5591–5596.
- [9] D.M. Gay, Computing optimal locally constrained steps, *SIAM J. Sci. Stat. Comput.* 2 (2) (1981) 186–197.
- [10] S.M. Goldfeld, R.E. Quandt, H.F. Trotter, Maximization by quadratic hill-climbing, *Econometrica* 34 (2) (1966) 541–551.
- [11] W. Huisinga, *Metastability of Markovian Systems*, PhD thesis, Freie Universität Berlin, 2001.
- [12] L. Jingen, Y. Yang, M. Shah, Learning semantic visual vocabularies using diffusion distance, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2009, CVPR 2009, 2009, pp. 461–468.

- [13] Y. Keller, R.R. Coifman, S. Lafon, S.W. Zucker, Audio-visual group recognition using diffusion maps, *IEEE Trans. Signal Process.* 58 (1) (2010) 403–413.
- [14] J.B. Kruskal, Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis, *Psychometrika* 29 (1964) 1–27.
- [15] S. Lafon, Diffusion Maps and Geometric Harmonics, PhD thesis, Yale University, May 2004.
- [16] S. Lafon, Y. Keller, R.R. Coifman, Data fusion and multicue data matching by diffusion maps, *IEEE Trans. Pattern Anal. Mach. Intell.* 28 (11) (2006) 1784–1797.
- [17] S. Maji, A.C. Berg, Max-margin additive classifiers for detection, in: *IEEE 12th International Conference on Computer Vision 2009*, Sept 2009, pp. 40–47.
- [18] G.J. McLachlan, K.E. Basford, *Mixture Models: Inference and Applications to Clustering*, Marcel Dekker, New York, 1988.
- [19] C.D. Meyer (Ed.), *Matrix Analysis and Applied Linear Algebra*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [20] J.J. More, D.C. Sorensen, Computing a trust region step, *SIAM J. Sci. Stat. Comput.* 4 (3) (1983) 553–572.
- [21] A. Rahimi, B. Recht, Random features for large-scale kernel machines, in: *Neural Information Processing Systems*, 2007.
- [22] S.T. Roweis, L.K. Saul, Nonlinear dimensionality reduction by locally linear embedding, *Science* 290 (December 2000) 2323–2326.
- [23] W. Rudin, *Fourier Analysis on Groups*, Interscience Tracts in Pure and Applied Mathematics, Interscience Publishers, New York, London, 1962.
- [24] X. Rui, S. Damelin, D.C. Wunsch, Applications of diffusion maps in gene expression data-based cancer diagnosis analysis, in: *29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2007, EMBS 2007*, 2007, pp. 4613–4616.
- [25] A. Schclar, A. Averbuch, N. Rabin, V. Zheludev, K. Hochman, A diffusion framework for detection of moving vehicles, *Digit. Signal Process.* 20 (1) (2010) 111–122.
- [26] D.C. Sorensen, Newton’s method with a model trust region modification, *SIAM J. Numer. Anal.* 19 (2) (1982) 409–429.
- [27] R. Talmon, I. Cohen, S. Gannot, Supervised source localization using diffusion kernels, in: *2011 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, WASPAA, 2011*, pp. 245–248.
- [28] R. Talmon, I. Cohen, S. Gannot, Single-channel transient interference suppression with diffusion maps, *IEEE Trans. Audio, Speech Language Process.* 21 (1–2) (2013) 132–144.
- [29] J.B. Tenenbaum, V. de Silva, J.C. Langford, A global geometric framework for nonlinear dimensionality reduction, *Science* 290 (5500) (2000) 2319–2323.
- [30] A. Vedaldi, A. Zisserman, Efficient additive kernels via explicit feature maps, *IEEE Trans. Pattern Anal. Mach. Intell.* 34 (3) (March 2012) 480–492.
- [31] G. Yang, X. Xu, J. Zhang, Manifold alignment via local tangent space alignment, in: *International Conference on Computer Science and Software Engineering*, December 2008.
- [32] Ya-xiang Yuan, Recent advances in trust region algorithms, *Math. Program.* 151 (1) (2015) 249–281.
- [33] Z. Zhang, H. Zha, *Principal Manifolds and Nonlinear Dimension Reduction via Local Tangent Space Alignment*, Technical Report CSE-02-019, Department of Computer Science and Engineering, Pennsylvania State University, 2002.