

Curvature Analysis of Frequency Modulated Manifolds in Dimensionality Reduction

Mijail Guillemand* Armin Iske†

Department of Mathematics, University of Hamburg, Germany

Abstract

Recent advances in the analysis of high-dimensional signal data have triggered an increasing interest in geometry-based methods for nonlinear dimensionality reduction (NDR). In many applications, high-dimensional datasets typically contain redundant information, and NDR methods are important for an efficient analysis of their properties. During the last few years, concepts from differential geometry were used to create a new range of NDR methods. In the construction of such geometry-based strategies, a natural question is to understand their interaction with classical and modern signal processing tools (convolution transforms, Fourier analysis, wavelet functions). In particular, an important task is the analysis of the incurred geometrical deformation when applying signal transforms to the elements of a dataset. In this paper, we propose the concepts of modulation maps and modulation manifolds for the construction of particular datasets relevant in signal processing and NDR. We consider numerical methods for analyzing geometrical properties of the modulation manifolds, with a particular focus on their scalar and mean curvature. In our numerical examples, we apply the resulting geometry-based analysis to simple test cases, where we present geometrical and topological effects of relevance in manifold learning.

Keywords: Nonlinear dimensionality reduction, manifold learning, signal processing, Fourier and wavelet analysis, numerical differential geometry.

1 Introduction

During the last decade, novel concepts for *nonlinear dimensionality reduction* (NDR) have gained significant relevance due to the increasing complexity of new challenging problems in data and signal analysis. In the design of these modern tools, special emphasis is placed on geometrical aspects, where concepts from differential geometry and algebraic topology play an important role [5, 6, 17, 18, 25, 26]. The geometry-based approach of NDR can be viewed as a complementary strategy to statistical oriented methods from machine learning and data mining [14].

*guillemard@math.uni-hamburg.de

†iske@math.uni-hamburg.de

Revised version, last modified: 16th June 2010

To briefly describe the basic problem of NDR and manifold learning, suppose we are given a dataset $X = \{x_i\}_{i=1}^m \subset \mathbb{R}^n$ lying in a high-dimensional Euclidean space, where X is assumed to be sampled from a submanifold \mathcal{M} of \mathbb{R}^n , i.e., $X \subset \mathcal{M} \subset \mathbb{R}^n$. Moreover, we assume that the dimension of \mathcal{M} is much smaller than the dimension of the ambient space, i.e., $\dim(\mathcal{M}) \ll n$. The primary objective of manifold learning is to construct a low-dimensional representation of X which can be used to efficiently visualize and analyze its geometrical properties.

For many examples of datasets $X = \{x_i\}_{i=1}^m \subset \mathbb{R}^n$, each element $x_i \in X$ can be considered as a signal that may be analyzed through a transformation map T , defined via convolution transforms, Fourier analysis, or wavelet functions. Therefore, from a manifold learning perspective, a relevant question is the analysis of the geometrical changes that X goes through when deformed with T resulting in the set $T(X) = \{T(x_i)\}_{i=1}^m$ and, in particular, to study the composition $P \circ T$ of a NDR map P with a signal transform T .

To analyze these problems, we construct a particular class of datasets X sampled from manifolds \mathcal{M} generated by *modulation maps*. The idea of modulation maps is inspired by concepts in frequency modulation techniques from signal transmission in engineering domains. Here, our main interest is to consider these signal processing concepts from a geometrical perspective. To gain further insight into their geometrical properties, we use numerical approximations to construct basic geometric data such as metric and curvature tensors. With these modulation manifolds we can design examples of low-dimensional datasets embedded in high dimensional spaces, which are relevant both in signal processing and dimensionality reduction. A main characteristic of these constructions is that they provide examples (in the same spirit as the Swiss roll dataset) of test cases where classical linear projections, such as principal component analysis (PCA) are outperformed by more recent nonlinear methods, such as isomap, Riemannian normal coordinates (RNC), etc [14, 18, 25].

The contributions of this paper can be summarized as follows. We introduce modulation maps and modulation manifolds as a relevant concept in dimensionality reduction. We provide conditions that justify the terminology used for these objects (as diffeomorphism and manifolds). We consider numerical procedures for computing geometrical quantities, such as metric tensors, Gaussian and mean curvatures. In our numerical examples, we finally illustrate relevant phenomena of modulated manifolds in the context of dimensionality reduction.

The outline of this paper is as follows. In Section 2, basic features of manifold learning and dimensionality reduction are recalled, where we briefly discuss the interaction of dimensionality reduction maps and signal transforms. In Section 3, the concept of modulation manifolds is explained, where one concrete example is provided. This is followed by a discussion on measuring geometric deformations of modulation manifolds using metric and curvature tensors, and the usage of the Laplace-Beltrami operator for computing the mean curvature. In Section 4, numerical examples are provided, where the geometric distortion of selected manifolds (sphere and torus surfaces) is illustrated by combining modulation maps and dimensionality reduction methods. In the numerical examples, we also compare standard PCA techniques with modern (nonlinear) dimensionality reduction strategies. We finally present in Section 4 various geometrical and topological effects that are relevant from both a dimensionality reduction and signal processing viewpoint.

2 Manifold Learning in Dimensionality Reduction

Let $\mathcal{M} \subset \mathbb{R}^n$ denote a smooth compact Riemannian submanifold of a high-dimensional space \mathbb{R}^n , where the dimension of \mathcal{M} is much smaller than that of the ambient space, i.e., $\mathcal{M} \subset \mathbb{R}^n$ with $p = \dim(\mathcal{M}) \ll n$. The primary goal of manifold learning is to construct a low-dimensional representation Ω of \mathcal{M} , with assuming the existence of a diffeomorphism $\mathcal{A} : \Omega \subset \mathbb{R}^d \rightarrow \mathcal{M} \subset \mathbb{R}^n$, $d \ll n$. One relevant task that we wish to address is the analysis of a discrete sample of \mathcal{M} , given by a finite scattered dataset $X = \{x_i\}_{i=1}^m \subset \mathcal{M}$. In this problem, one aims at the construction of a low-dimensional representation $Y = \{y_i\}_{i=1}^m \subset \Omega \subset \mathbb{R}^d$ sharing similar geometrical properties with X . This in turn requires a suitable smooth map $R : \mathcal{M} \rightarrow \Omega$, whose only input is the dataset $X \subset \mathcal{M}$. Due to the (weak) Whitney embedding theorem (which states that any connected smooth p -dimensional manifold can be smoothly embedded in \mathbb{R}^{2p+1}), one basic condition in this problem is $2p + 1 \leq d \leq n$, see [4]. Throughout this paper, we use the term *manifold* to denote a compact smooth connected manifold without boundary being embedded in some Euclidean space.

2.1 Manifold Reconstruction from Discrete Data

Now, a crucial requirement in manifold learning is to ensure conditions under which the finite sampling $X = \{x_i\}_{i=1}^m$ is dense enough for recovering geometrical and topological properties of \mathcal{M} . In the case of topological conditions, new properties have been investigated over the last years [20]. A main concept is the *condition number* $1/\tau$ of the manifold which encodes local and global curvature properties of \mathcal{M} . The condition number can be related to the *medial axis* of \mathcal{M} , which is defined as the closure of the set

$$G = \{x \in \mathbb{R}^n : \exists p, q \in \mathcal{M}, p \neq q, \text{ with } d(x, \mathcal{M}) = \|x - p\| = \|x - q\|\}.$$

By using the medial axis of the manifold, we have

$$\tau = \inf_{p \in \mathcal{M}} d(p, \overline{G}).$$

We also recall that a *deformation retract* is a continuous map $r : U \rightarrow X$ between topological spaces U and X with $X \subset U$, such that the restriction $r|_X$ is the identity [10]. Using these concepts, the following result [20] relates the sampling of the manifold with its homological reconstruction.

Proposition 2.1 (Niyogi, Smale, Weinberger, 2008). *Let \mathcal{M} be a compact Riemannian submanifold of \mathbb{R}^n and $X = \{x_i\}_{i=1}^m \subset \mathbb{R}^n$ a finite $\epsilon/2$ -dense collection in \mathcal{M} , i.e., for each $p \in \mathcal{M}$, there is an $x \in X$ satisfying $\|p - x\|_{\mathbb{R}^n} < \epsilon/2$. Then for any $\epsilon < \sqrt{3/5}\tau$, we have that $U = \bigcup_{x \in X} B_\epsilon(x)$ deformation retracts to \mathcal{M} , and therefore the homology of U equals the homology of \mathcal{M} . ■*

This proposition describes requirements for the discretization of a manifold when reconstructing the homology as a basic topological invariant. But in order to recover geometrical information, we need to consider differential geometric data. Sampling and discretization aspects in differential geometry are an emergent and very active research topics. Basic concepts in this field are the angle defect [2,7] and the usage of the Laplacian

operator when computing the mean and Gaussian curvatures [1, 3, 11, 19]. A series of additional important developments for manifold sampling have also been developed over the last years [22, 23]. Other important topics are generalizations of the curvature concept in Alexandrov spaces or cell-complexes, as discussed in [9, 21].

2.2 Application Examples

Relevant motivations of our framework are time-frequency representations, where a segmentation of a signal is analyzed with Fourier or wavelet functions. For instance, a typical context is the short term Fourier transform (STFT) for band-limited functions $f \in L^2([0, 1])$, using a window function g ,

$$\mathcal{G}_g f(b, \omega) = \langle f, g_{b, \omega} \rangle = \int_0^1 f(t) \overline{g_{b, \omega}(t)} dt \quad \text{where} \quad g_{b, \omega}(t) = g(t - b) e^{2\pi i \omega t}.$$

We use the core ideas of this framework in order to construct a dataset X_f by considering a segmentation of the domain of f in such a way that small consecutive signal patches are analyzed. More precisely, the set of signal patches can be defined as a dataset of the form

$$X_f = \{x_i\}_{i=1}^m \quad \text{for} \quad x_i = (f(t_{k(i-1)+j}))_{j=0}^{n-1} \in \mathbb{R}^n$$

for $k \in \mathbb{N}$ being a fixed *hop-size*. Here, the regular sampling grid $\{t_\ell\}_{\ell=0}^{km-k+n-1} \subset [0, 1]$ is constructed with considering the Nyquist-Shannon theorem for f . Notice that X_f may be embedded in a very high-dimensional ambient space \mathbb{R}^n , although the dimension of X_f itself may be small. For instance, in audio analysis, for 44kHz signals, $n = 1024$ is commonly used, and therefore customized dimensionality reduction methods could be of main interest. With this particular scheme, the STFT of f can be interpreted as a transformation of the set X_f by taking the (windowed) Fourier transform of each x_i .

In this paper we regard X_f as a geometrical object in the context of manifold learning, a strategy that might provide new information when studying the function f . We remark that the approach taken here essentially differs from traditional nonlinear time-series analysis, in the sense that we don't consider any time-delays and embedding dimensions as used in phase space representations [12].

A second family of examples (similar to the spirit of time-frequency analysis) arises in image processing. One strategy would be to consider a dataset $X_f = \{x_i\}_{i=1}^m$ constructed from a grayscale image $f : [0, 1]^2 \rightarrow [0, 1]$, along with a finite covering of small squares (each containing n pixels) $\{O_i \subset [0, 1]^2\}_{i=1}^m$, centered at pixels positions $\{k_i\}_{i=1}^m \subset [0, 1]^2$. As in the previous situation, when considering band-limited images, the domain $[0, 1]^2$ can be sampled uniformly and the dataset can then be defined as

$$X_f = \{f(O_i) \in \mathbb{R}^n\}_{i=1}^m,$$

where n is the size of the squares O_i , and m denotes the number of pixels k_i . As before, our aim is to analyze the geometry of the image data X_f to gain useful information about the properties of the image f .

As already explained in the outset of the introduction, it is desirable to work with analysis methods that combine signal processing transforms with dimensionality reduction methods. In this case, the basic objects are the manifold \mathcal{M} , the data samples

$X = \{x_i\}_{i=1}^m$ taken from \mathcal{M} , and a diffeomorphism $\mathcal{A} : \Omega \rightarrow \mathcal{M}$, where Ω is the low-dimensional copy of \mathcal{M} to be reconstructed via dimensionality reduction. In this case, the only algorithmic input is the dataset X , but with the assumption that we can reconstruct topological information from \mathcal{M} with X in the spirit of Proposition 2.1. Another basic object in our scheme is a signal processing map $T : \mathcal{M} \rightarrow \mathcal{M}_T$, which may be based on Fourier analysis, wavelet transforms, or convolution filters, together with the resulting set $\mathcal{M}_T = \{T(p), p \in \mathcal{M}\}$ of transformed data. The following diagram describes the basic situation.

$$\begin{array}{ccc} \Omega \subset \mathbb{R}^d & \xrightarrow{\mathcal{A}} & \mathcal{M} \subset \mathbb{R}^n \\ & & \downarrow T \\ \Omega' \subset \mathbb{R}^d & \xleftarrow{P} & \mathcal{M}_T \subset \mathbb{R}^n \end{array}$$

The final component constructs an approximation to Ω , denoted as $\Omega' = P(\mathcal{M}_T)$, by using a dimensionality reduction map P . The characteristics of Ω and Ω' may differ depending on the dimensionality reduction technique, but the main objective is to construct Ω' , so that basic geometrical and topological properties of Ω are recovered. In this paper, we will use modulation maps for the embedding \mathcal{A} to study geometrical and topological effects being incurred by particular dimensionality reduction projections $P : \mathcal{M} \rightarrow \Omega'$.

3 Modulation Maps and Curvature Distortion

This section is devoted to a particular construction of manifolds \mathcal{M} and diffeomorphisms \mathcal{A} based on modulation maps. Modulation techniques are well-known engineering and telecommunication procedures used to transmit information by varying the frequency of a carrier signal ϕ . A main property of these techniques is the simultaneous transmission of different information by using different frequency bands that can be conveniently separated with special convolution filters. Motivated by these ideas, we want to analyze, from a geometrical point of view, a frequency modulation map $\mathcal{A} : \Omega \rightarrow \mathcal{M}$, where \mathcal{M} represents the carrier signals modulated by Ω , which is the information content to be transmitted. Rather than analyzing a specific engineering modulation method, our goal is to use and generalize these concepts in order to construct manifolds and datasets with relevant properties in dimensionality reduction and signal analysis. In our particular situation, the domain Ω is a manifold and so its structural content, transmitted via \mathcal{A} , needs to be extracted from \mathcal{M} .

3.1 Modulation Maps and Modulated Manifolds

We now define the concept of modulation maps and modulation manifolds. Moreover, we provide basic examples together with some properties where these ideas can formally be described. We start with an intuitive explanation and a short motivation using an image processing example. Our current focus on smooth manifolds is an important conceptual step before one may address more general situations (e.g. simplicial complexes, CW-complexes, Alexandrov spaces, etc) that are crucial for many engineering applications.

In the following, we assume \mathcal{M} to be a submanifold of a high dimensional Euclidean space \mathcal{H} , $\mathcal{M} \subset \mathcal{H}$, and Ω to be a parameter space considered as a submanifold of a low

dimensional Euclidean space. The intuition behind the concept of a *modulated manifold* \mathcal{M} is to generate the elements of \mathcal{M} by a family of signal transforms $\{s(\alpha) : \mathcal{H} \rightarrow \mathcal{H}\}_{\alpha \in \Omega}$ and a vector $\phi \in \mathcal{H}$ such that for each $y \in \mathcal{M}$, we have $y = s(\alpha)\phi$ for some $\alpha \in \Omega$. This concept extends the classical notion of modulation by considering maps $s(\alpha) : \mathcal{H} \rightarrow \mathcal{H}$ in order to transform (or *modulate*) a function ϕ for generating the elements of \mathcal{M} . A straightforward generalization is to consider a family of vectors $\{\phi_k\}_{k=1}^d \subset \mathcal{H}$ together with a corresponding family of signal transforms $\{s_k(\alpha) : \mathcal{H} \rightarrow \mathcal{H}\}_{k=1, \alpha \in \Omega}^d$, such that for any $y \in \mathcal{M}$ there is one unique $\alpha \in \Omega$ satisfying

$$y = \sum_{k=1}^d s_k(\alpha)\phi_k.$$

To further explain our motivation, let us return to the analysis of a grayscale image $f : [0, 1]^2 \rightarrow [0, 1]$ using a covering of small squares, or patches, $\{O_i \subset [0, 1]^2\}_{i=1}^m$, each containing n pixels. In this toy example we assume that the corresponding point cloud data $X_f = \{f(O_i) \in \mathbb{R}^n\}_{i=1}^m$ lies in some manifold $\mathcal{M} \subset \mathbb{R}^n$. If an image f is composed of an homogeneous texture, the dataset X_f is a cluster whose elements have similar geometrical characteristics. In a simplified scenario, the idea would be to use a representative patch $\phi \in \mathbb{R}^n$ in order to generate all elements of X_f . The main task is to find a family of transformations (the *modulation maps*) $s(\alpha) : \mathbb{R}^n \rightarrow \mathbb{R}^n$, parametrized by a low dimensional space Ω , such that for any patch $y \in X_f$, there is some $\alpha \in \Omega$ with $y = s(\alpha)\phi$. We remark that several methods in image processing have recently been proposed with a loosely related philosophy (see e.g. the *patch-based texture analysis* as part of classical *texture synthesis* methods [13]).

Definition 3.1 (Modulation maps and modulation manifolds). Let $\{\phi_k\}_{k=1}^d \subset \mathcal{H}$ be a set of vectors in an Euclidean space \mathcal{H} , and $\{s_k : \Omega \rightarrow \mathcal{C}_{\mathcal{H}}(\mathcal{H})\}_{k=1}^d$ a family of smooth maps from a manifold Ω to $\mathcal{C}_{\mathcal{H}}(\mathcal{H})$ (the continuous functions from \mathcal{H} into \mathcal{H}). We say that a manifold $\mathcal{M} \subset \mathcal{H}$ is a $\{\phi_k\}_{k=1}^d$ -**modulated manifold** if

$$\mathcal{M} = \left\{ \sum_{k=1}^d s_k(\alpha)\phi_k, \alpha \in \Omega \right\}.$$

In this case, the map $\mathcal{A} : \Omega \rightarrow \mathcal{M}$, $\alpha \mapsto \sum_{k=1}^d s_k(\alpha)\phi_k$, is called **modulation map**.

This concept can also be described using the language of vector bundles. Recall that a vector bundle (E, Ω, π, F) is defined for a manifold E as a surjective map $\pi : E \rightarrow \Omega$, with a vector space F as a *fiber*, and a manifold *base* Ω such that $\pi^{-1}(\alpha)$ is isomorphic to F for each $\alpha \in \Omega$. Locally, E is homeomorphic to $\Omega \times F$, where E may have a non-trivial global structure. Two typical examples for a vector bundle are the tangent bundle of a manifold and the Möbius band E , whose base Ω is a circle and whose fiber F is a line. *Sections* of a vector bundle $\pi : E \rightarrow \Omega$ are smooth maps $s : \Omega \rightarrow E$ such that $\pi(s(\alpha)) = \alpha$, and their prototypical examples are vector fields on Ω . In our Definition 3.1 we currently use a trivial vector bundle $\pi : E \rightarrow \Omega$ for $E = \Omega \times \mathcal{C}_{\mathcal{H}}(\mathcal{H})$, a fiber $\mathcal{C}_{\mathcal{H}}(\mathcal{H})$, and smooth sections $\{s_k : \Omega \rightarrow E\}_{k=1}^d$.

Example 3.1 (Frequency and scale modulation). A basic example of a modulation map for a manifold $\Omega \subset \mathbb{R}^d$ is a map $\mathcal{A} : \Omega \subset \mathbb{R}^d \rightarrow \mathbb{R}^n$ of the form

$$\mathcal{A}_\alpha(t_i) = \sum_{k=1}^d \phi_k(\alpha_k t_i), \quad \alpha = (\alpha_1, \dots, \alpha_d) \in \Omega, \quad \{t_i\}_{i=1}^n \subset [0, 1],$$

for a finite set of smooth band-limited functions $\{\phi_k \in \mathcal{C}^\infty([0, 1]) \cap \mathcal{B}_{f_s}, \}_{k=1}^d$, with $\mathcal{B}_{f_s} = \{f \in L^2([0, 1]), \text{supp}(\hat{f}) \subseteq [-f_s, f_s]\}$, and a given fixed sampling rate f_s . We use a uniformly spaced finite sampling set $\{t_i\}_{i=1}^n \subset [0, 1]$ (as justified by the Nyquist-Shannon sampling theorem), by considering band-limited functions $\{\phi_k\}_{k=1}^d$. Note that we use the same notation for the band-limited functions ϕ_k and the vector of sampling values $\{\phi_k(t_i)\}_{i=1}^n$, as justified by the *Whittaker-Shannon interpolation formula*. More precisely, as the support of our functions ϕ_k is located in $[0, 1]$, and due to their band-limited property, the interpolation formula allows us to reconstruct each ϕ_k with the finite sampling set $(\phi_k(t_i))_{i=1}^n \in \mathbb{R}^n$, which unambiguously identifies the function $\phi_k : [0, 1] \rightarrow \mathbb{R}$ with a vector $\phi_k \in \mathbb{R}^n$. Note that the maps $s_k(\alpha)$ of Definition 3.1 are in this example given by $s_k(\alpha)\phi_k(t_i) = \phi_k(\alpha_k t_i)$. In other words, we use the (continuous) map $s_k(\alpha) : \mathcal{C}^\infty([0, 1]) \rightarrow \mathcal{C}^\infty([0, 1])$, $f(t) \mapsto f(\alpha_k t)$, as the scaling by factor α_k , the k -th coordinate of vector $\alpha \in \Omega \subset \mathbb{R}^d$. We now consider additional properties to ensure a meaningful definition of frequency modulation, as required in Proposition 3.1.

The first property is the *band separation* for Ω , defined as $B_k \cap B_j = \emptyset$, for all $k \neq j$, with $B_k = \{\alpha_k \in \mathbb{R}, \alpha = (\alpha_1, \dots, \alpha_{k-1}, \alpha_k, \alpha_{k+1}, \dots, \alpha_d) \in \Omega\}$. The intuition behind this assumption is that the coordinates of the manifold Ω are ranged in different and non overlapping regions. This is actually a standard condition in telecommunication engineering, where different non-overlapping frequency bands are used for the transmission of different signals. This property will help ensure the injectivity of \mathcal{A} , as required in Proposition 3.1.

The second property is a *sufficiently dense sampling set* for $\{\phi_k : [0, 1] \rightarrow \mathbb{R}\}_{k=1}^d$ and $\Omega \subset \mathbb{R}^d$, defined as a set $\{t_i\}_{i=1}^n \subset [0, 1]$ that is dense enough for reconstructing each function in $\{\phi_k\}_{k=1}^d$ using the Whittaker-Shannon interpolation formula. Note that these conditions restrict the range of values for Ω , but more general situations could be considered when using different domains for the functions ϕ_k .

Now the following proposition justifies the motivation of Definition 3.1 and the terminology used in the above example.

Proposition 3.1 (Manifold structure of \mathcal{M} and diffeomorphism property of \mathcal{A}). *Suppose that Ω is a submanifold of \mathbb{R}^d with separated bands. Moreover, let $\{t_i\}_{i=1}^n \subset [0, 1]$ be a sufficiently dense sampling set for Ω and a family of smooth band-limited functions $\{\phi_k : [0, 1] \rightarrow \mathbb{R}\}_{k=1}^d$. If for any element $\alpha = (\alpha_1, \dots, \alpha_d) \in \Omega$ the $d \times d$ matrix $\mathfrak{J}_\mathcal{A}(\alpha)$ with entries*

$$\mathfrak{J}_\mathcal{A}(\alpha)_{ks} = \sum_{i=1}^n t_i^2 \phi'_k(\alpha_k t_i) \phi'_s(\alpha_s t_i), \quad 1 \leq k, s \leq d, \quad (1)$$

is invertible, then $\mathcal{M} = \{(\sum_{k=1}^d \phi_k(\alpha_k t_i))_{i=1}^n, \alpha = (\alpha_1, \dots, \alpha_d) \in \Omega\}$ is a submanifold of \mathbb{R}^n , and the map $\mathcal{A} : \Omega \rightarrow \mathcal{M}$, $\mathcal{A}_\alpha(t_i) = \sum_{k=1}^d \phi_k(\alpha_k t_i)$, is a diffeomorphism.

Proof. This follows as a straightforward application of the rank theorem for manifolds. The first step is to verify that the map $\mathcal{A} : \Omega \rightarrow \mathbb{R}^n$, $\mathcal{A}_\alpha(t_i) = \sum_{k=1}^d \phi_k(\alpha_k t_i)$ is an injective immersion. Since Ω is compact we can then conclude that \mathcal{A} is also a smooth embedding (see [16, Proposition 7.4]). With this property, $\mathcal{A} : \Omega \rightarrow \mathbb{R}^n$ is a homeomorphism onto its image $\mathcal{M} = \mathcal{A}(\Omega)$.

The injectivity of \mathcal{A} can be ensured with the *band separation* property of Ω . In order to verify that \mathcal{A} is an immersion, we compute the rank of \mathcal{A} at $\alpha \in \Omega$, defined as the rank of the linear map $\mathcal{A}_* : T_\alpha \Omega \rightarrow T_{\mathcal{A}(\alpha)} \mathbb{R}^n$. For this task, we consider a smooth chart (U, χ) , with $\alpha \in U \subset \Omega$ and $\chi : U \subset \Omega \rightarrow \tilde{U} \subset \mathbb{R}^p$ a smooth map, where $p = \dim(\Omega)$. Now, the derivative \mathcal{A}_* (or pushforward) of \mathcal{A} at $\alpha \in \Omega$ can be described in local coordinates as the Jacobian matrix of $\hat{\mathcal{A}} = \psi \circ \mathcal{A} \circ \chi^{-1} : \chi(U) = \tilde{U} \subset \mathbb{R}^p \rightarrow \psi(\mathbb{R}^n) = \mathbb{R}^n$, at $\chi(\alpha)$ for $\psi = \text{Id}_{\mathbb{R}^n} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, the identity using $(\mathbb{R}^n, \text{Id}_{\mathbb{R}^n})$ as a single chart of the smooth manifold \mathbb{R}^n . The Jacobian matrix of $\hat{\mathcal{A}}$ is the $n \times p$ matrix $J_{\hat{\mathcal{A}}} = J_\psi J_{\mathcal{A}} J_{\chi^{-1}}$, i.e., a product of the $n \times n$ matrix J_ψ , the $n \times d$ matrix $J_{\mathcal{A}}$, and the $d \times p$ matrix $J_{\chi^{-1}}$. As J_ψ is the identity matrix, we have $\text{rank}(J_{\hat{\mathcal{A}}}) = \text{rank}(J_{\mathcal{A}} J_{\chi^{-1}})$, and since $\chi^{-1} : \tilde{U} \subset \mathbb{R}^p \rightarrow U \subset \Omega \subset \mathbb{R}^d$ is a diffeomorphism, we have $\text{rank}(J_{\chi^{-1}}) = p$. Now, if we assume that $\text{rank}(J_{\mathcal{A}}) = d$, with the Sylvester rank inequality we have $\text{rank}(J_{\mathcal{A}}) + \text{rank}(J_{\chi^{-1}}) - d \leq \text{rank}(J_{\mathcal{A}} J_{\chi^{-1}})$. Therefore, due to our assumption $\text{rank}(J_{\mathcal{A}}) = d$, we obtain $\text{rank}(J_{\hat{\mathcal{A}}}) = p$, so that $\mathcal{A} : \Omega \rightarrow \mathbb{R}^n$ is an immersion, since $\text{rank}(\mathcal{A}) = \text{rank}(\mathcal{A}_*) = \dim(\Omega)$. Our assumption $\text{rank}(J_{\mathcal{A}}) = d$ is equivalent to $\det(J_{\mathcal{A}}^T(\alpha) J_{\mathcal{A}}(\alpha)) \neq 0$, for each $\alpha \in \Omega$ and $d \times d$ matrices $J_{\mathcal{A}}^T(\alpha) J_{\mathcal{A}}(\alpha)$. Therefore, for $\mathfrak{J}_{\mathcal{A}}(\alpha) = J_{\mathcal{A}}^T(\alpha) J_{\mathcal{A}}(\alpha)$, we obtain the representation (1).

Using the previous argument we can guarantee that $\mathcal{M} = \mathcal{A}(\Omega)$ is a submanifold of \mathbb{R}^n , if $\mathfrak{J}_{\mathcal{A}}(\alpha)$ is invertible for each $\alpha \in \Omega$. Therefore, the map $\mathcal{A} : \Omega \rightarrow \mathcal{M}$ is a surjective smooth constant-rank map. Finally, we now use an important consequence of the rank theorem for manifolds (see [16, Theorem 7.15]) to conclude that \mathcal{A} is a diffeomorphism between the manifolds Ω and \mathcal{M} , and therefore $\dim(\mathcal{M}) = \dim(\Omega) = p$. ■

Example 3.2 (An explicit case of a frequency modulation map). We now continue the analysis for Example 3.1, using Proposition 3.1, to construct a *frequency modulation map*. A prototypical case is given by trigonometric functions, $\phi(t) = \sin(t)$, where the corresponding modulation map $\mathcal{A} : \Omega \rightarrow \mathbb{R}^n$, with Ω a submanifold of \mathbb{R}^d , is called *frequency modulation map*. A concrete construction is given by $\mathcal{A}^\gamma : \Omega \rightarrow \mathbb{R}^n$ with

$$\mathcal{A}_\alpha^\gamma(t_i) = \sum_{k=1}^d \sin((\alpha_k^0 + \gamma \alpha_k) t_i), \quad \{t_i\}_{i=1}^n \subset [0, 1], \quad (2)$$

for a fixed *bandwidth* parameter $\gamma > 0$, *center frequencies* $\alpha^0 = (\alpha_1^0, \dots, \alpha_d^0) \in \mathbb{R}^d$, and *separated* bands B_k , i.e., $B_k \cap B_j = \emptyset$ for all $k \neq j$, where

$$B_k = \{\alpha_k^0 + \gamma \alpha_k \in \mathbb{R}, \alpha = (\alpha_1, \dots, \alpha_{k-1}, \alpha_k, \alpha_{k+1}, \dots, \alpha_d) \in \Omega\}.$$

This construction can be viewed as the application of an affine transform to the manifold $\Omega \subset \mathbb{R}^d$, with shift α^0 and bandwidth γ , so that the coordinates of the vectors in the resulting set $\alpha^0 + \gamma \Omega = \{\alpha^0 + \gamma \alpha, \alpha \in \Omega\}$ will not share common values. In this case, the frequency content introduced by each coordinate $\alpha \in \Omega$ will not overlap.

Now we need to verify that \mathcal{A}^γ in (2) is a modulation map for specific parameters γ and α^0 . To this end, the following lemma provides only a special case.

Lemma 3.1 (Diffeomorphism property of frequency modulated maps). *Suppose that Ω is a submanifold of \mathbb{R}^d with $\Omega \subset [-1, 1]^d$, and let $\alpha^0 \in \mathbb{R}^d$ be a vector, such that $\Omega_{\alpha^0} = \alpha^0 + \Omega = \{\alpha^0 + \gamma\alpha, \alpha \in \Omega\}$ has separated bands. If the $d \times d$ matrix $\mathfrak{J}(\alpha^0)$ is invertible for $\{t_i = i/n\}_{i=1}^n \subset [0, 1]$, and*

$$\mathfrak{J}(\alpha^0)_{ks} = \sum_{i=1}^n t_i^2 \cos(\alpha_k^0 t_i) \cos(\alpha_s^0 t_i), \quad 1 \leq k, s \leq d,$$

then there exist a bandwidth parameter $\gamma > 0$, such that the matrix $\mathfrak{J}(\beta)_{ks}$ is regular for all $\beta \in \Omega_{\alpha^0}^\gamma = \alpha^0 + \gamma\Omega$.

Proof. This follows from a straightforward application of the Gershgorin circle theorem, where we compare the locations of the eigenvalues of $\mathfrak{J}(\beta)$ with those of the eigenvalues of $\mathfrak{J}(\alpha^0)$. Using the mean value theorem for $\cos(\beta_k t_i)$, with $\beta_k = \alpha_k^0 + \gamma\alpha_k$, we have $\cos((\alpha_k^0 + \gamma\alpha_k)t_i) = \cos(\alpha_k^0 t_i) + \gamma\alpha_k t_i \cos(\eta)$ for $\alpha_k^0 t_i < \eta < \alpha_k^0 t_i + \gamma\alpha_k t_i$. This implies

$$\begin{aligned} \mathfrak{J}(\beta)_{ks} &= \sum_{i=1}^n t_i^2 \cos((\alpha_k^0 + \gamma\alpha_k)t_i) \cos((\alpha_s^0 + \gamma\alpha_s)t_i) \\ &= \sum_{i=1}^n t_i^2 \cos(\alpha_k^0 t_i) \cos(\alpha_s^0 t_i) + \gamma\tau_{ks} \sum_{i=1}^n t_i^2 \\ &= \mathfrak{J}(\alpha^0)_{ks} + \gamma\tau_{ks} \frac{(n+1)(2n+1)}{6n}, \end{aligned}$$

with $t_i = i/n$ and $\tau_{ks} \in \mathbb{R}$ satisfying $|\tau_{ks}| \leq 1$. Using the triangle inequality, we obtain

$$\mathfrak{J}(\beta)_{kk} - \sum_{k \neq s} |\mathfrak{J}(\beta)_{ks}| \geq \mathfrak{J}(\alpha^0)_{kk} - \sum_{k \neq s} |\mathfrak{J}(\alpha^0)_{ks}| - \sum_{i=1}^{d-1} \gamma \frac{(n+1)(2n+1)}{6n}.$$

Therefore, due to the Gershgorin circle theorem, the regularity of the matrices $\mathfrak{J}(\beta)$, for $\beta \in \Omega_{\alpha^0}^\gamma$, can be guaranteed, if $\mathfrak{J}(\alpha^0)_{kk} - \sum_{k \neq s} |\mathfrak{J}(\alpha^0)_{ks}| > \sum_{i=1}^{d-1} \gamma \frac{(n+1)(2n+1)}{6n}$. In particular, this property can be fulfilled for any $\gamma \geq 0$ satisfying

$$\gamma < \gamma_{\max} = \frac{6n \min_{1 \leq i \leq d} (\mathfrak{J}(\alpha^0)_{ii} - \sum_{j \neq i} |\mathfrak{J}(\alpha^0)_{ij}|)}{(d-1)(n+1)(2n+1)}. \quad (3)$$

■

Assuming the hypothesis of the above lemma in combination with Proposition 3.1, we can conclude that $\mathcal{M} = \{(\sum_{k=1}^d \sin((\alpha^0 + \gamma\alpha)t_i))_{i=1}^n \in \mathbb{R}^n, \alpha \in \Omega\}$ is a submanifold of \mathbb{R}^n , and the map $\mathcal{A}^\gamma : \Omega \rightarrow \mathcal{M}$, $\mathcal{A}_\alpha^\gamma(t_i) = \sum_{k=1}^d \sin((\alpha_k^0 + \gamma\alpha_k)t_i)$, is the corresponding diffeomorphism. We remark that it is also straightforward to justify this lemma by using the continuity of the determinant of $\mathfrak{J}(\alpha^0)$. But the bound on γ in (3) allows us to provide a concrete range for the bandwidth parameter γ , where we can ensure the desired diffeomorphism and manifold properties for \mathcal{A} and \mathcal{M} . As a concrete example, consider a sampling rate $n = 256\text{Hz}$, and center frequencies $\alpha^0 = (40\text{Hz}, 60\text{Hz}, 80\text{Hz})$. In this case, the matrix $\mathfrak{J}(\alpha^0)$ is regular for any bandwidth parameter $\gamma < \gamma_{\max}$ with $\gamma_{\max} \approx 0.2454\text{Hz}$. We remark that our numerical simulations suggest that the range for γ , allowing diffeomorphism and manifold properties for \mathcal{A} and \mathcal{M} , is much larger than in the estimate (3). A more detailed analysis on these properties, however, is beyond the aim of the present paper.

3.2 Geometric Deformation with Curvature Measurements

Now we analyze the geometrical deformation between Ω , \mathcal{M} and Ω' , as incurred by the modulation map $\mathcal{A} : \Omega \rightarrow \mathcal{M}$ and the dimensionality reduction transform $P : \mathcal{M} \rightarrow \Omega'$. In order to analyze the geometrical deformation we consider basic curvature concepts, including Gaussian and mean curvature. On the one hand, we first present a very general procedure for computing the scalar curvature using the Riemannian metric as main input. On the other hand, we describe a discrete scheme for computing the mean curvature by using the Laplace-Beltrami operator.

We first recall basic ingredients for computing the scalar curvature of a p -dimensional manifold \mathcal{M} , and refer to [15, 24] for further details. The fundamental tool for describing the geometry of \mathcal{M} is a metric tensor field $g : \mathcal{M} \rightarrow \mathbb{R}^{p \times p}$, defined for local coordinates (x_1, \dots, x_p) as

$$g_{ij}(x) = \langle \partial_i, \partial_j \rangle \quad \text{for } x \in \mathcal{M} \text{ and } 1 \leq i, j \leq p.$$

Here, the partial derivatives ∂_i represent the tangent vectors at each $x \in \mathcal{M}$, and the notation $g_{ij}(x)$ is occasionally used for $g(x)$. With this fundamental building block, all other required structures for defining the scalar curvature are defined, including the Christoffel symbols and curvature tensors. The Christoffel symbols can be described as

$$\Gamma_{ij}^k = \frac{1}{2} \sum_{\ell=1}^p \left(\frac{\partial g_{j\ell}}{\partial x_i} + \frac{\partial g_{i\ell}}{\partial x_j} - \frac{\partial g_{ij}}{\partial x_\ell} \right) g^{\ell k}. \quad (4)$$

Here, the expression g^{ij} denotes the inverse matrix of g_{ij} . An explicit formula for the curvature tensor is given in terms of the Christoffel symbols (the 1,3 curvature tensor) as

$$R^\ell{}_{ijk} = \sum_{h=1}^p (\Gamma_{jk}^h \Gamma_{ih}^\ell - \Gamma_{ik}^h \Gamma_{jh}^\ell) + \frac{\partial \Gamma_{jk}^\ell}{\partial x_i} - \frac{\partial \Gamma_{ik}^\ell}{\partial x_j}. \quad (5)$$

We use the tensor contractions

$$R_{ijkl} = \sum_{h=1}^p R_{ijk}^h g_{\ell h} \quad \text{and} \quad R_{ij} = \sum_{k,\ell=1}^p g^{k\ell} R_{kij\ell} = \sum_{k=1}^p R_{kij}^k \quad (6)$$

for intermediate computations. The scalar curvature is computed together with Gaussian curvature, which for the case of two dimensional manifolds differs by a factor 2, as

$$S = \sum_{i,j=1}^p g^{ij} R_{ij}. \quad (7)$$

In our particular situation, we are considering manifolds embedded in linear spaces (i.e., $\Omega \subset \mathbb{R}^d$ and $\mathcal{M} \subset \mathbb{R}^n$). Our strategy for introducing the concept of curvature distortion is to compare the geometries of Ω and \mathcal{M} generated by their corresponding first fundamental forms, which are particular metrics induced by their ambient space.

The usage of these concepts in discrete settings as meshes for manifolds is a highly nontrivial task, and there are relatively simple situations where convergence properties cannot be guaranteed. A typical example is the *Schwarz lantern* which is a simple mesh

discretization of a cylinder with very poor approximation properties [11]. In our implementation, we consider a straightforward discretization scheme for computing the Gaussian curvature in (7), cf. Algorithm 3.1. But we also consider a rigorous discretization procedure, with good convergence properties, for computing the mean curvature. This can be described as follows.

To analyze the curvature distortion we compute the mean curvature using the Laplace-Beltrami operator $\Delta_{\mathcal{M}}f = \text{div}(\nabla_{\mathcal{M}}f)$, where f are the coordinates of the embedded manifold $\mathcal{M} \subset \mathbb{R}^n$. The Laplace-Beltrami operator is also defined with the metric tensor g_{ij} , its inverse matrix g^{ij} , and its determinant $|g_{ij}|$:

$$\Delta_{\mathcal{M}}f(x) = \frac{1}{\sqrt{|g_{ij}(x)|}} \sum_{k,\ell=1}^p \partial_k \left(\sqrt{|g_{ij}(x)|} g^{ij}(x) \partial_{\ell} f(x) \right) \quad \text{for } x \in \mathcal{M}.$$

Designing accurate and robust methods for discretizing the operator Δ is a very active research topic. A recent method [1] for doing so provides an efficient strategy by using

$$L_K^h f(w) = \frac{1}{4\pi h^2} \sum_{t \in K} \frac{\text{Area}(t)}{\#(t)} \sum_{p \in V(t)} e^{-\frac{\|p-w\|^2}{4h}} (f(p) - f(w)). \quad (8)$$

Here, K is a mesh in \mathbb{R}^3 , the set of vertices is denoted by V , and $f : V \rightarrow \mathbb{R}$. For a face $t \in K$, the number of vertices in t is denoted by $\#(t)$, and $V(t)$ is the set of vertices of t . The parameter h is a positive quantity representing the size of the mesh at each point. As we consider two dimensional manifolds, the mesh $K_{\epsilon,\eta}$ can be described with two variables ϵ, η for controlling the parameter h . We finally remark that the discretization scheme in (8) convergences w.r.t. $\|\cdot\|_{\infty}$ (see [1] for details), i.e.,

$$\lim_{\epsilon, \eta \rightarrow 0} \sup_{K_{\epsilon,\eta}} \|L_{K_{\epsilon,\eta}}^{h(\epsilon,\eta)} f - \Delta_{\mathcal{M}}f\|_{\infty} = 0.$$

3.3 Metric Tensor for Frequency Modulated Manifolds

We now compute a metric tensor in \mathcal{M} in order to analyze the geometrical deformation of Ω as incurred by the application of a modulation map $\mathcal{A} : \Omega \subset \mathbb{R}^d \rightarrow \mathcal{M} \subset \mathbb{R}^n$. The resulting metric tensor can then be used for computing the curvature tensor and the corresponding scalar curvature, which will be used as a measure for the geometric deformation. Our strategy is to consider a parametrization of Ω and to compute the metric tensor generated from the ambient space \mathbb{R}^n . In particular, we use the first fundamental form with respect to the given parametrization. The resulting formula (9) follows by direct computation, as explained in the following proposition, where we compute the first fundamental form of a modulated manifold \mathcal{M} .

Proposition 3.2. *Let \mathcal{M} be a manifold constructed from a diffeomorphic modulation map $\mathcal{A} : \Omega \subset \mathbb{R}^d \rightarrow \mathcal{M} \subset \mathbb{R}^n$,*

$$\mathcal{A}_{\alpha}(t_i) = \sum_{k=1}^d \phi_k(\alpha_k t_i) \quad \text{where } \alpha = (\alpha_1, \dots, \alpha_d) \in \Omega \text{ and } \{t_i\}_{i=1}^n \subset [0, 1],$$

and $\{\alpha_j(\theta_1, \dots, \theta_p)\}_{j=1}^d$ be a parametrization of Ω with $p = \dim(\mathcal{M}) = \dim(\Omega)$. Then, the first fundamental form of \mathcal{M} constructed from this parametrization is given by

$$g_{sr} = \sum_{\ell=1}^n t_\ell^2 \sum_{r,q=1}^d \left(\frac{d\phi_r}{dt}(\alpha_r t_\ell) \frac{d\phi_q}{dt}(\alpha_q t_\ell) \frac{\partial \alpha_r}{\partial \theta_s} \frac{\partial \alpha_q}{\partial \theta_r} \right). \quad (9)$$

Proof. This follows as a direct computation of the Jacobian of the composition $\mathcal{A} \circ \alpha$. The Jacobian with respect to parametrization $\alpha_j(\theta_1, \dots, \theta_d)$ of Ω is given by

$$J_{\mathcal{A}} = \left(\frac{\partial \mathcal{A}_\ell}{\partial \theta_i} \right)_{\ell,i} \quad \text{where} \quad \frac{\partial \mathcal{A}_\ell}{\partial \theta_i} = \frac{\partial}{\partial \theta_i} \left(\sum_{j=1}^d \phi_j(\alpha_j t_\ell) \right) = \sum_{j=1}^d \frac{d\phi_j}{dt}(\alpha_j t_\ell) t_\ell \frac{\partial \alpha_j}{\partial \theta_i}.$$

The first fundamental form (metric tensor) of \mathcal{M} is given by

$$\begin{aligned} (J_{\mathcal{A}}^T J_{\mathcal{A}})_{s,r} &= \sum_{\ell=1}^n \left(\sum_{j=1}^d \frac{\partial \phi_j}{\partial t}(\alpha_j t_\ell) t_\ell \frac{\partial \alpha_j}{\partial \theta_s} \right) \left(\sum_{j=1}^d \frac{\partial \phi_j}{\partial t}(\alpha_j t_\ell) t_\ell \frac{\partial \alpha_j}{\partial \theta_r} \right) \\ &= \sum_{\ell=1}^n \sum_{r,q=1}^d \left(\frac{d\phi_r}{dt}(\alpha_r t_\ell) t_\ell \frac{\partial \alpha_r}{\partial \theta_s} \frac{d\phi_q}{dt}(\alpha_q t_\ell) t_\ell \frac{\partial \alpha_q}{\partial \theta_r} \right) \\ &= \sum_{\ell=1}^n t_\ell^2 \sum_{r,q=1}^d \left(\frac{d\phi_p}{dt}(\alpha_r t_\ell) \frac{d\phi_q}{dt}(\alpha_q t_\ell) \frac{\partial \alpha_r}{\partial \theta_s} \frac{\partial \alpha_q}{\partial \theta_r} \right) \end{aligned}$$

As g_{sr} is given by $(J_{\mathcal{A}}^T J_{\mathcal{A}})_{s,r}$, we obtain the resulting equation (9). ■

The expression in equation (9) will play an important role in our following discussion. However, due to its complexity (even for rather simple examples as a sphere or a torus), we prefer to work with a numerical framework for illustrating its properties through the relevant curvature tensors. Moreover, the numerical approach taken provides a flexible scheme that can handle arbitrary two-dimensional frequency modulated manifolds \mathcal{M} that are defined by a finite scattered dataset $X = \{x_i\}_{i=1}^m \subset \mathcal{M}$.

3.4 Numerical Computation of Curvature Tensors

Now we combine the representation (9) with the curvature tensors in (6) and (7) for describing how Ω is geometrically deformed under the mapping $\mathcal{A} : \Omega \rightarrow \mathcal{M}$. In the following computations, we focus on the particular case of two-dimensional manifolds embedded in a three-dimensional space, i.e., $d = 3$, $\dim(\Omega) = 2$, with $\Omega \subset \mathbb{R}^3$. In our computation of the metric tensor in (9), the main inputs are the functions $\{\phi_j\}_{j=1}^3$ and the parametrization $\{\alpha_j(\theta_1, \theta_2)\}_{j=1}^3$ of Ω , which is used to construct the Jacobian components $\frac{\partial \alpha_p}{\partial \theta_s}$ and $\frac{\partial \alpha_q}{\partial \theta_r}$. The following algorithm describes the basic steps for computing the scalar curvature S and the metric tensors g_{ij} of the modulated manifold \mathcal{M} .

Algorithm 3.1. (Curvature and metric tensors of manifolds.)

Input:

- (a) Parametrization $\alpha = (\alpha_j(\theta_1, \theta_2))_{j=1}^3$ of Ω ;
- (b) Functions $\{\phi_j\}_{j=1}^3$ generating the map \mathcal{A} .

- (1) Compute the Jacobian matrices J_α ;
- (2) Compute the metric tensor g_{ij} via equation (9);
- (3) Compute the Christoffel symbols Γ_{ij}^k via equation (4);
- (4) Compute the tensors R_{ijkl} , R_{ij} via equations (6);
- (5) Compute the scalar curvature S via equation (7).

Output: Scalar curvature S of \mathcal{M} .

4 Computational Experiments

We now apply Algorithm 3.1 and the discretization in (8) to compute the Gaussian and mean curvature of frequency modulated manifolds for two different toy examples. The corresponding Matlab code is available at www.math.uni-hamburg.de/home/guillemard/curvature/. In these test cases we use the sphere $\Omega = \mathbb{S}^2$ and the torus $\Omega = \mathbb{T}^2$ to illustrate how the curvature is modified under modulation maps and dimensionality reduction projections.

4.1 Frequency Modulation for a Sphere

In this first numerical example, we use the unit sphere for the parameter space i.e., $\Omega = \mathbb{S}^2 \subset \mathbb{R}^3$, and the modulation map $\mathcal{A} : \mathbb{S}^2 \subset \mathbb{R}^3 \rightarrow \mathcal{M} \subset \mathbb{R}^{256}$ is defined as $\mathcal{A}_\alpha(t_i) = \sum_{k=1}^3 \sin((\alpha_k^0 + \gamma\alpha_k)t_i)$, where

$$\begin{aligned}\alpha_1(u, v) &= \cos(v) \cos(u), \\ \alpha_2(u, v) &= \cos(v) \sin(u), \\ \alpha_3(u, v) &= \sin(v).\end{aligned}$$

Here, we use a finite and regular distribution of values $u \in [0, 2\pi]$, $v \in [0, \pi]$, and $\{t_i\}_{i=1}^{256} \subset [0, 1]$. We apply shifting and scaling to the manifold \mathbb{S}^2 , so that the frequency positions are given by the coordinates $(\alpha_0^1, \alpha_0^2, \alpha_0^3)$ and the scaling factor γ . We use these parameters to obtain a separation of the frequency bands as described in Subsection 3.1. The scaling factor γ gives the spreading of each frequency band (bandwidth). A main observation of the following experiments is that the geometrical deformation depends primarily on the parameter γ . A graphical display of the manifolds \mathbb{S}^2 and \mathcal{M} is presented in Figure 1, where the sphere \mathbb{S}^2 is compared with a three dimensional PCA projection of \mathcal{M} (denoted as $P(\mathcal{M})$). The PCA projection $P(\mathcal{M})$ of $\mathcal{M} \subset \mathbb{R}^n$, $n = 256$, produces a significant geometrical deformation. One objective of the following analysis is to measure this distortion. It can be observed experimentally that an increase in the scale factor γ corresponds to a more pronounced cubic shape deformation, as shown in Figure 1 (b).

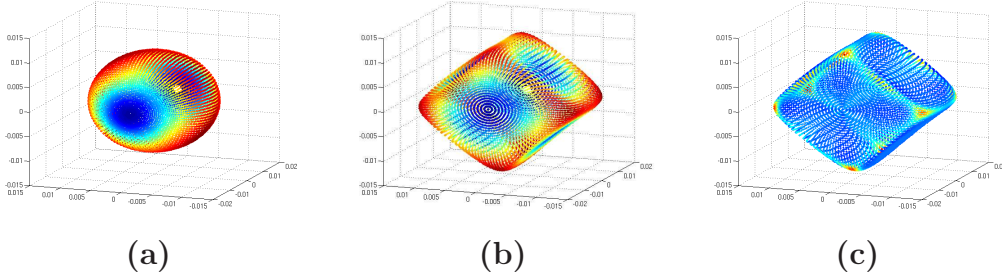


Figure 1: PCA projections of a modulated sphere $\mathcal{M} = \mathcal{A}^\gamma(\Omega) \subset \mathbb{R}^n$, for $\Omega = \mathbb{S}^2 \subset [-1, 1]^3$, with sample frequency $n = 256$, center frequencies $\alpha^0 = (40\text{Hz}, 60\text{Hz}, 80\text{Hz})$, and bandwidth parameter $\gamma = 0.4\text{Hz}$ (a) The sphere $\mathbb{S}^2 \subset \mathbb{R}^3$ with the mean curvature of \mathcal{M} as color code; (b) PCA projection $P_{\text{PCA}}(\mathcal{M}) \subset \mathbb{R}^3$ with its mean curvature in the color code; (c) PCA projection $P_{\text{PCA}}(\mathcal{M}) \subset \mathbb{R}^3$ with its Gaussian curvature in the color code.




Figure 2: Curvatures of the modulated sphere $\mathcal{M} = \mathcal{A}(\Omega)$ and its projection $P_{\text{PCA}}(\mathcal{M})$ (see Figure 1). (a) The mean curvature of $\mathcal{M} \subset \mathbb{R}^{256}$; (b) the mean curvature of $P(\mathcal{M}) \subset \mathbb{R}^3$; (c) the Gaussian curvature of $P(\mathcal{M}) \subset \mathbb{R}^3$.

In order to measure the geometric deformation for this example, we compute the scalar and mean curvature of $\mathcal{M} \subset \mathbb{R}^{256}$ and that of its three-dimensional PCA projection $P(\mathcal{M})$. Figure 2 (a) shows the mean curvature of the manifold \mathcal{M} and the Gaussian curvature of its projection $P(\mathcal{M})$ is displayed in Figure 2 (c). Note that the mean curvature of \mathcal{M} shows some small variations over its surface, indicating only a small deformation of the spherical geometry via the frequency modulation map \mathcal{A} . In contrast, the PCA projection of \mathcal{M} shows significant curvature variations, as presented in Figure 1 (c). Note that there are two sets of four maximal scalar curvature values, corresponding to the corners of the cubic shaped surface shown in Figure 1 (c).

4.2 Frequency Modulation for a Torus

In this second numerical example, we use the torus surface $\Omega = \mathbb{T}^2 \subset \mathbb{R}^3$ in combination with the modulation map $\mathcal{A}_\alpha(t_i) = \sum_{k=1}^3 \sin((\alpha_k^0 + \gamma\alpha_k)t_i)$, and the parametrization

$$\begin{aligned}\alpha_1(u, v) &= (R + r \cos(v)) \cos(u), \\ \alpha_2(u, v) &= (R + r \cos(v)) \sin(u), \\ \alpha_3(u, v) &= r \sin(v).\end{aligned}$$

We use a regular distribution of values $u \in [0, 2\pi]$, $v \in [0, 2\pi]$, and $\{t_i\}_{i=1}^{256} \subset [0, 1]$. As in the previous example, the parameter γ plays a key role in the cubic deformation, cf. Figure 3 (c).

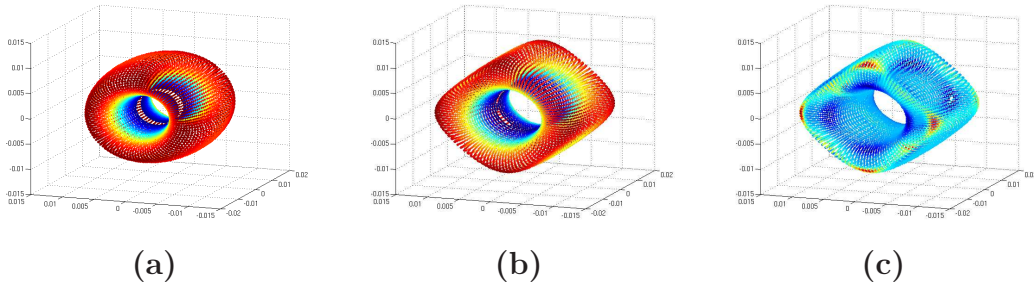


Figure 3: PCA projections of a modulated torus $\mathcal{M} = \mathcal{A}^\gamma(\Omega) \subset \mathbb{R}^n$, for $\Omega = \mathbb{T}^2 \subset [-1, 1]^3$, with sample frequency $n = 256$, center frequencies $\alpha^0 = (40\text{Hz}, 60\text{Hz}, 80\text{Hz})$, and bandwidth parameter $\gamma = 0.3\text{Hz}$ (a) The torus $\mathbb{T}^2 \subset \mathbb{R}^3$ with the mean curvature of \mathcal{M} as color code; (b) PCA projection $P_{\text{PCA}}(\mathcal{M}) \subset \mathbb{R}^3$ with its mean curvature in the color code; (c) PCA projection $P_{\text{PCA}}(\mathcal{M}) \subset \mathbb{R}^3$ with its Gaussian curvature in the color code.




Figure 4: Curvatures of the modulated torus $\mathcal{M} = \mathcal{A}(\Omega)$ and its projection $P_{\text{PCA}}(\mathcal{M})$ (see Figure 3). (a) The mean curvature of $\mathcal{M} \subset \mathbb{R}^{256}$; (b) the mean curvature of $P_{\text{PCA}}(\mathcal{M}) \subset \mathbb{R}^3$; (c) the Gaussian curvature of $P_{\text{PCA}}(\mathcal{M}) \subset \mathbb{R}^3$.

The mean curvature of $\mathcal{M} \subset \mathbb{R}^{256}$ is shown in Figure 4 (a). We observe a typical pattern for the torus geometry: a constant value for the mean curvature along the smaller circle of the torus (depicted with the middle vertical line in Figure 4 (a)), two circles with constant curvature on the top and bottom of the torus (depicted with two vertical lines equidistant to the middle of Figure 4 (a)), and another circle with constant mean curvature along the larger circle of the torus (depicted with the leftmost and rightmost vertical lines of the Figure 4 (a)).

The Gaussian curvature of the PCA projection $P(\mathcal{M})$, shown in Figure 4 (c), illustrates a similar structure, but with a considerable geometrical deformation which includes two sets of four points with maximal scalar curvature, representing the corners of the cubic shaped projection $P(\mathcal{M})$ shown in Figure 3 (c).

We finally remark that, for these toy examples, the task of estimating the frequencies of each $x \in \mathcal{M}$ (and therefore computing Ω) is not a difficult problem. A straightforward

method for doing so is to use the Fourier transform of the signals $x \in \mathcal{M}$, as for instance performed in classical *partial tracking* techniques. But our focus here is to conceptually understand the geometrical effects of dimensionality reduction methods in datasets $\mathcal{M} = \mathcal{A}(\Omega)$. From a signal processing point of view, the problem of estimating Ω becomes much more challenging when considering more complex functions ϕ_k for the modulation map \mathcal{A} . In such situations, the presented geometrical framework will be useful for estimating Ω in more complex test case scenarios.

4.3 Dimensionality Reduction and Topological Effects from Curvature

In our numerical experiments, we observe that an increase of the bandwidth parameter γ amplifies the geometrical distortion of the modulation manifolds $\mathcal{M} \subset \mathbb{R}^n$. In fact, when increasing the bandwidth parameter γ , standard (linear) projection methods, such as PCA, fail to recover the geometry of the manifold. In contrast, recent nonlinear dimensionality reduction methods, such as isomap, achieve to recover (up to a certain bandwidth γ) the geometry of Ω . For the purpose of further illustration, Figure 5 (b) shows how the linear PCA projection destroys the geometrical content of the modulated torus \mathcal{M} , whereas the nonlinear isomap projection achieves to recover the topological and geometrical features of the torus \mathbb{T}^2 reasonably well. This effect is similar to the comparison between the isomap and the PCA method for the classical Swiss Roll dataset example. But with our concept we can generate much more challenging datasets from modulated manifolds to evaluate the performance of NDR techniques.




Figure 5: Dimensionality reduction of a frequency modulated manifold $\mathcal{M} = \mathcal{A}(\Omega) \subset \mathbb{R}^n$, for a torus $\Omega = \mathbb{T}^2 \subset [-1, 1]^3$ using a sampling frequency $n = 256$, with center frequencies $\alpha^0 = (40\text{Hz}, 60\text{Hz}, 80\text{Hz})$, and bandwidth parameter $\gamma = 4\text{Hz}$. (a) $P_{\text{Isomap}}(\mathcal{M})$ isomap projection of \mathcal{M} ; (b) $P_{\text{PCA}}(\mathcal{M})$ PCA projection of \mathcal{M} .

To further investigate the geometrical distortion of modulated manifolds incurred by dimensionality reduction maps, we illustrate how the bandwidth parameter γ affects the topological reconstruction of Ω . To this end, we analyze the evolution of the mean curvature of the manifold $\mathcal{M} = \mathcal{A}^\gamma(\Omega)$ for a range of bandwidth parameters γ (see Figure 6). In this case, we compute for values $\gamma \in [0, 4]$ the maximum mean curvature of \mathcal{M} . We let $\Omega = \mathbb{T}^2 \subset [-1, 1]^3$, and therefore, if the projection $P(\mathcal{M})$ correctly reconstructs the topology of \mathcal{M} , we expect the Betti numbers of $P(\mathcal{M})$ to be $b_0 = 1$ and $b_1 = 2$.

Figure 6 (a) shows an increase of the maximum mean curvature as the bandwidth parameter is amplified. For small values of γ , a standard (linear) dimensionality reduction method, such as PCA, achieves to reproduce the topology of Ω . This can be verified by the persistent homology of $P(\mathcal{M})$, as depicted in the barcode graph of Figure 6 (b), where for $\gamma = 0.2\text{Hz}$, the first two Betti numbers agree with that of a torus structure ($b_0 = 1, b_1 = 2$). But for higher values γ , the geometry of $\mathcal{M} = \mathcal{A}^\gamma(\Omega)$ is more difficult to reconstruct. In fact, the topological structure of the dimensionality reduced datasets $P(\mathcal{M})$ are in this case incorrect. This can be seen in Figure 6 (c), where the persistent homology barcode of $P_{\text{PCA}}(\mathcal{M})$ with $\gamma = 4\text{Hz}$, yields incorrect Betti numbers $b_0 = 3$ and $b_1 = 0$. An example of one data set $P_{\text{PCA}}(\mathcal{M})$ is shown in Figure 5 (b), where $\gamma = 4\text{Hz}$.

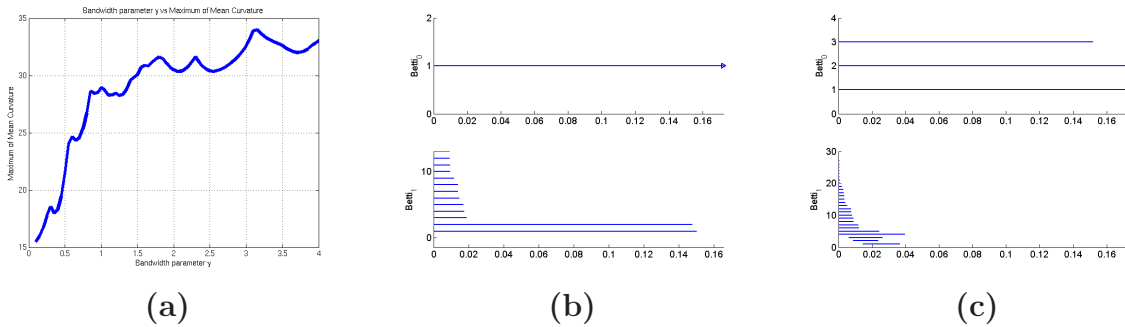


Figure 6: (a) Maximum mean curvature of $\mathcal{M} = \mathcal{A}(\Omega)$ vs bandwidth parameter γ in the interval $[0, 4]$, for $\Omega = \mathbb{T}^2 \subset [-1, 1]^3$, and center frequencies $\alpha^0 = (40\text{Hz}, 60\text{Hz}, 80\text{Hz})$. (b) Betti numbers of $P_{\text{PCA}}(\mathcal{M})$ and bandwidth parameter $\gamma = 0.2\text{Hz}$ ($b_0 = 1, b_1 = 2$), (c) Betti numbers of $P_{\text{PCA}}(\mathcal{M})$ and bandwidth parameter $\gamma = 4\text{Hz}$ ($b_0 = 3, b_1 = 0$).

5 Conclusion and Future Steps

We have introduced the concept of frequency modulation maps and modulation manifolds as relevant objects in manifold learning and dimensionality reduction. We have developed a numerical scheme, Algorithm 3.1, for computing the scalar curvature of a modulation manifold, along with its metric tensor. We applied the algorithm to two different test cases of frequency modulated manifolds. In one test case, we considered the sphere \mathbb{S}^2 , and the other example relies on the torus \mathbb{T}^2 . The numerical examples are illustrating the geometrical distortion incurred by the dimensionality reduction when relying on PCA projections. Moreover, we have shown that the standard linear PCA projection is outperformed by the nonlinear isomap projection, which achieves to recover the geometry of the modulated surface. But the concept of modulation manifold can also be used to construct more challenging datasets for nonlinear dimensionality reduction methods. The findings of this paper provide only a first insight into the nature of frequency modulation maps and modulation manifolds. Weaker structural assumptions on Ω and \mathcal{M} (Alexandrov spaces, cell-complexes, etc) should be considered in future work to cover a larger variety of engineering applications. To this end, the work on persistent homology, discrete Morse theory, or Alexandrov spaces [8, 9, 21] gives a suitable background.

Acknowledgments

We thank the anonymous referees for their useful comments and suggestions. This work is supported by the priority program SPP 1324 of the Deutsche Forschungsgemeinschaft.

References

- [1] M. Belkin, J. Sun, and Y. Wang. Discrete Laplace operator on meshed surfaces. In *Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry*, ACM, 2008, 278–287.
- [2] E.D. Bloch. A characterization of the angle defect and the Euler characteristic in dimension 2. *Discrete and Computational Geometry* **43**(1), 2010, 100–120.
- [3] A.I. Bobenko and Y.B. Suris. *Discrete differential geometry: integrable structure*. Amer. Mathematical Society, 2008.
- [4] D.S. Broomhead and M. Kirby. A new approach to dimensionality reduction: theory and algorithms. *SIAM Journal on Applied Mathematics* **60**(6), 2000, 2114–2142.
- [5] A. Brun, C.-F. Westin, M. Herberthson, and H. Knutsson. Fast manifold learning based on Riemannian normal coordinates. In *Proceedings of the SCIA;05*, Joensuu, Finland, June 2005, 920–929.
- [6] G. Carlsson. Topology and data. *Amer. Mathematical Society* **46**(2), 2009, 255–308.
- [7] J. Cheeger, W. Müller, and R. Schrader. On the curvature of piecewise flat spaces. *Communications in Mathematical Physics* **92**(3), 1984, 405–454.
- [8] R. Forman. *A User’s Guide to Discrete Morse Theory*. Seminaire Lotharingien de Combinatoire, 48:B48c, 2002.
- [9] R. Forman. Bochner’s method for cell complexes and combinatorial Ricci curvature. *Discrete and Computational Geometry* **29**(3), 2003, 323–374.
- [10] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [11] H.C. Hege and K. Polthier. *Visualization and Mathematics III*. Springer, 2003.
- [12] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge University Press, 2004.
- [13] V. Kwatra, A. Schodl, I. Essa, G. Turk, and A. Bobick. Graphcut textures: image and video synthesis using graph cuts. *ACM Transactions on Graphics* **22**(3), 2003, 277–286.
- [14] J.A. Lee and M. Verleysen. *Nonlinear Dimensionality Reduction*. Springer, 2007.
- [15] J.M. Lee. *Riemannian Manifolds: An Introduction to Curvature*. Springer, 1997.
- [16] J.M. Lee. *Introduction to Smooth Manifolds*. Springer, 2003.

- [17] T. Lin and H. Zha. Riemannian manifold learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**, 2008, 796–809.
- [18] T. Lin, H. Zha, and S.U. Lee. *Riemannian Manifold Learning for Nonlinear Dimensionality Reduction*. Lecture Notes in Computer Science, 3951:44, 2006.
- [19] M. Meyer, M. Desbrun, P. Schröder, and A.H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and Mathematics* **3**, 2002, 35–57.
- [20] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete and Computational Geometry* **39**(1), 2008, 419–441.
- [21] Y. Otsu and T. Shioya. The Riemannian structure of Alexandrov spaces. *J. Differential Geom.* **39**(3), 1994, 629–658.
- [22] E. Saucan, E. Appleboim, and Y.Y. Zeevi. *Geometric Sampling of Manifolds for Image Representation and Processing*. Lecture Notes in Computer Science **4485**, 2007, 907–918.
- [23] E. Saucan, E. Appleboim, and Y.Y. Zeevi. Sampling and reconstruction of surfaces and higher dimensional manifolds. *Journal of Mathematical Imaging and Vision* **30**(1), 2008, 105–123.
- [24] G. Walschap. *Metric Structures in Differential Geometry*. Springer, 2004.
- [25] H. Zha and Z. Zhang. Continuum isomap for manifold learning. *Computational Statistics and Data Analysis* **52**(1), 2007, 184–200.
- [26] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.* **33**(2), 2005, 249–274.