

WESD - Weighted Spectral Distance for Measuring Shape Dissimilarity

Ender Konukoglu, Ben Glocker, Antonio Criminisi and Kilian M. Pohl

Abstract—This article presents a new distance for measuring shape dissimilarity between objects. Recent publications introduced the use of eigenvalues of the Laplace operator as compact shape descriptors. Here, we revisit the eigenvalues to define a proper distance, called Weighted Spectral Distance (WESD), for quantifying shape dissimilarity. The definition of WESD is derived through analysing the heat-trace. This analysis provides the proposed distance an intuitive meaning and mathematically links it to the intrinsic geometry of objects. We analyse the resulting distance definition, present and prove its important theoretical properties. Some of these properties include: i) WESD is defined over the entire sequence of eigenvalues yet it is guaranteed to converge, ii) it is a pseudometric, iii) it is accurately approximated with a finite number of eigenvalues, and iv) it can be mapped to the $[0, 1)$ interval. Lastly, experiments conducted on synthetic and real objects are presented. These experiments highlight the practical benefits of WESD for applications in vision and medical image analysis.

Index Terms—Shape Distance, Spectral Distance, Laplace Operator, Laplace Spectrum, Segmentations, Label Maps, Medical Images

I. INTRODUCTION

Quantifying shape differences between objects is an important task for various areas in computer science, medical imaging and engineering. In manufacturing, for example, one may wish to characterize the difference in shape of two fabricated tools. In radiology, a doctor frequently diagnoses a disease based on anatomical and pathological shape changes over time. In computer vision, discriminative shape models are used for automated object recognition, [1], [2].

In order to define measurements of shape dissimilarity, scientists rely on descriptors of objects that capture information on their geometry [1]. These descriptors can be in the form of parametrized models (e.g. point clouds, surface patches, space curves, medial axis transforms) or in the form of geometric properties (e.g. volume, surface area to volume ratio, curvature maps). Once a descriptor is formulated the distance between two shapes can be defined as the difference between the associated descriptors. The exact definition of the distance however, is a critical issue. In order to define an intuitive and theoretically sound distance, one should ensure that it takes

into account the nature of the descriptor. For instance, the descriptor might be an infinite sequence of positive values, in which case we should be careful not to define a distance that diverges for every non-identical pair of shapes.

Shape descriptors based on the eigensystems of Laplace and Laplace-Beltrami operators, called *spectral signatures*, have recently gained popularity in computational shape analysis [3], [4], [5], [6], [7], [8]. These descriptors leverage the fact that the eigenvalues and the eigenfunctions of Laplace operators contain information on the intrinsic geometry of objects [9], [10], [11]. A visual analogy useful for an intuitive understanding is to think of an object (e.g. in 2D) as the membrane of a drum. In this case the eigenvalues correspond to the fundamental frequencies of vibration of the membrane during percussion, and the eigenfunctions correspond to its fundamental patterns of vibration. Both the eigenvalues and the eigenfunctions depend on the shape of the drum head and thus can be used as shape descriptors for the object.

Despite recent progress by [3], [4], [5], [6], [7], [8], designing meaningful shape distances based on spectral signatures remains challenging. Difficulties arise from the nature of the eigensystems. The eigenfunctions of a shape mostly provide *localized* information on the geometry of small neighborhoods. Aggregating such local information into an overall shape dissimilarity measure is non-trivial. On the other hand, the eigenvalues provide information about the *overall* shape, so they are ideal for defining global distances. However, they form a diverging sequence making it difficult to define a theoretically sound metric. Here, we tackle this latter problem and propose a new shape distance based on the eigenvalues, which is technically sound, intuitive and practically useful.

In the remainder of this section, we first review in further detail the literature on spectral signatures and shape distances related to eigenfunctions and eigenvalues. Then, we provide a brief overview of our new shape distance.

A. Eigenfunctions

The eigenfunctions of an object constitute an infinite set of functions. Each function depends on the shape of the object and is different than the rest of the set. Figure 1 illustrates this for two example objects where a few eigenfunctions are shown. The values these functions attain at each point capture the local geometry around the point, i.e. of its neighborhood. Inspired from this geometric information, methods define *local* shape signatures [4], [5], [6], [12] for each point on an object by evaluating a subset of eigenfunctions at that specific location. Global shape distances are then defined using such local

Corresponding author: E. Konukoglu (e-mail:ender.konukoglu@gmail.com) is with Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital and Harvard Medical School, MA 02129, USA.

B. Glocker and A. Criminisi are with Microsoft Research Cambridge, CB3 0FB, UK.

K.M. Pohl is with University of Pennsylvania, PA 19104, USA.

Authors would like to thank the support in part by Grant Number UL1RR024134 and by the Institute for Translational Medicine and Therapeutics (ITMAT) Transdisciplinary Program.




Fig. 1: Starfish and tarantula. The objects represented as binary maps are shown on the left, followed by the 1st, 2nd, 5th, 20th, and 100th eigenfunction. The values increase from blue (negative) to red (positive) with green being zero.

signatures. Such distance definitions rely on correspondences. These correspondences should hold both in terms of points and the subset of eigenfunctions used in the signatures, a condition hard to satisfy in practice [12]. Explicitly searching for such correspondences leads to expensive algorithms [12], [13], [14], [15], [16]. On the other hand, computing distances between distributions of local signatures obtained by aggregating all the points, as in [5], [6], [17], might implicitly construct false correspondences. In summary, defining a global distance based on local signatures is not an easy task.

Instead of extracting local information from an eigenfunction, one can also think of capturing its global pattern by looking at regions where its values are all positive or all negative. Such regions are called *nodal domains*. Different eigenfunctions induce different patterns and, in turn, have different number of nodal domains, called *nodal counts* [11]. For a given object, the ordered sequence of nodal counts contain information on its overall geometry [18], [19]. Inspired by these observations, authors in [20] used this sequence as a *global* shape signature. They further defined the associated shape distance between two objects as the Euclidean norm of the vector difference between their nodal count sequences. However, it is not intuitively clear what the nodal counts represent. Furthermore, the entire sequence is diverging so that, in practice, one first chooses a finite subset and then computes the distance for that subset. These difficulties make it hard to define an intuitive and sound shape distance based on nodal counts.

B. Eigenvalues

Signatures based on eigenvalues, on the other hand, have a clearer geometric interpretation. The set of eigenvalues contains information on the overall geometry of the object. Specifically, the ordered sequence is analytically related to the intrinsic geometry by the *heat-trace*, [21], [22], [23], [24], [25]. Hence, more intuitive distances can be constructed using the eigenvalues. However, similar to the sequence of nodal counts, the eigenvalue sequence is also divergent. This makes the distance definition theoretically challenging. Inspired by the sequence’s link to the geometry, Reuter et al. in [3], used the smallest N eigenvalues as a shape signature, called *shape-DNA*. As the associated shape distance, the authors proposed the Euclidean norm of the vector difference between the shape-DNAs of objects. Although this is a very good first

attempt, the divergent nature of eigenvalue sequence results in important theoretical limitations for this distance, as also pointed out in [14]. The main problems are i) defining a distance on the entire sequence does not yield a proper metric, ii) the differences between the higher components of two sequences dominate the final distance value, even though these components do not necessarily provide more information on the geometry, and iii) the distance value is sensitive to the choice of the signature size N . These theoretical problems also cause practical drawbacks as we demonstrate later.

This article proposes a new shape distance, called Weighted Spectral Distance (WESD), using the sequence of eigenvalues of the Laplace operator. We derive WESD from the functional relationship between the eigenvalues and the geometric invariants as given by the heat-trace. This derivation provides WESD a clear geometric intuition as a shape distance. It also links WESD to the distance defined by Reuter et al. in [3] as well as to the local signature defined in [5]. The resulting formulation of WESD differs from other previously proposed scores based on eigenvalues, whether in shape analysis or other fields [26], both in its formulation and in the fact that it is defined over the entire sequence. As we will show later, the latter point alleviates the critical importance of the choice of the signature. We furthermore analyse and prove theoretical properties of WESD showing that it does not share some of the fundamental problems the shape distance proposed in [3] has. Specifically, we prove that WESD: i) converges despite the fact that it is defined over the entire eigenvalue sequence, ii) can be mapped to the $[0, 1)$ interval, iii) is accurately approximated with a finite number of eigenvalues and the truncation error has an analytical upper bound and iv) is a pseudometric. These theoretical properties also yield important practical advantages such as being less sensitive to the signature size (truncation parameter) N , providing a principled way of choosing this parameter, providing more stable low-dimensional shape embedding and simplicity in combining with other distances as WESD can be normalised. Applying to synthetic and real objects, we further demonstrate the benefits of WESD in comparison to the other eigenvalue-based distance defined in [3].

The remainder of this article is structured as follows. Section II presents a brief overview of the Laplace operator, the eigenvalue sequence and its role in shape analysis. In Section III we define WESD and derive its theoretical properties. Section IV presents an extensive set of experimental analysis on 2D objects extracted from synthetic binary maps, shape-based retrieval results for 3D objects using the SHREC dataset [27], low dimensional embeddings of real 3D data such as subcortical structures in brain scans and 4D analysis of binary maps extracted from cardiac images.

II. SPECTRUM OF LAPLACE OPERATOR

This section provides a brief background on the Laplace operator, its eigenvalue sequence, called *spectrum*, and its role in shape analysis. We first relate an object’s intrinsic geometry to the spectrum of the corresponding Laplace operator. We then provide some details on the previously proposed shape-

DNA [3] and discuss the associated issues. For further details we refer the reader to [11], [24], [25] and [3].

We denote an object as a closed bounded domain $\Omega \subset \mathbb{R}^d$ with piecewise smooth boundaries. In the case of binary maps, Ω would correspond to the foreground representing the object. For a given Ω , the Laplace operator on this object is defined with respect to a twice differentiable real-valued function f as

$$\Delta_{\Omega} f \triangleq \sum_{i=1}^d \frac{\partial^2}{\partial x_i^2} f, \quad \forall \mathbf{x} \in \Omega$$

where $\mathbf{x} = [x_1, \dots, x_d]$ is the spatial coordinate. The eigenvalues and the eigenfunctions of Δ_{Ω} are defined as the solutions of the Helmholtz equation with Dirichlet type boundary conditions¹, [11],

$$\Delta f + \lambda f = 0 \quad \forall \mathbf{x} \in \Omega, \quad f(\mathbf{x}) = 0, \quad \forall \mathbf{x} \in \partial\Omega,$$

where $\partial\Omega$ denotes the boundary of the object and $\lambda \in \mathbb{R}$ is a scalar. There are infinitely many pairs $\{(\lambda_n, f_n)\}_{n=1}^{\infty}$ satisfying this equation and they form the set of eigenvalues and eigenfunctions respectively. The ordered set of eigenvalues is a positive diverging sequence such that $0 < \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n \leq \dots$. This infinite sequence is called the Dirichlet spectrum of Δ_{Ω} , which we simply refer as the ‘‘spectrum’’. In addition, each component of the spectrum is called a ‘‘mode’’, e.g. λ_n is the called n^{th} mode of the spectrum

The spectrum contains information on the intrinsic geometry of objects. Weyl in [9] showed the first spectrum-geometry link by proving that the asymptotic behavior of the eigenvalues is given as

$$\lambda_n \sim 4\pi^2 \left(\frac{n}{B_d V_{\Omega}} \right)^{2/d}, \quad n \rightarrow \infty,$$

where V_{Ω} is the volume of Ω and B_d is the volume of the unit ball in \mathbb{R}^d . Later works, as [21], [22], [23], [24], extended this result by studying the properties of the Green’s function of the Laplace operator, and showed that a more accurate spectrum-geometry link is given by the *heat-trace*, which in \mathbb{R}^d is given as

$$Z(t) \triangleq \sum_{n=1}^{\infty} e^{-\lambda_n t} = \sum_{s=0}^{\infty} a_{s/2} t^{-d/2+s/2}, \quad t > 0. \quad (1)$$

The coefficients of the polynomial expansion, $a_{s/2}$, are the components carrying the geometric information. These coefficients are given as sums of volume and boundary integrals of some local invariants of the shape, [22], [23], [25]. For instance, as given in [22], the first three coefficients are:

$$\begin{aligned} a_0 &= \frac{1}{(4\pi)^{d/2}} V_{\Omega} \\ a_{1/2} &= -\frac{1}{4(4\pi)^{d/2-1/2}} S_{\Omega}, \\ a_1 &= -\frac{1}{6(4\pi)^{d/2}} \int_{\partial\Omega} \kappa d\Omega, \end{aligned}$$

where S_{Ω} is the surface area (circumference in 2D) and κ is the mean (geodesic) curvature on the boundary of Ω .

¹Other boundary conditions yield different eigensystems. Here we are only interested in the Dirichlet type. Please refer to [11] for the other types.

The functional relationship between the eigenvalue sequence and the coefficients $a_{s/2}$ can be seen in Equation (1). This connection relates the spectrum to the intrinsic geometry, which is the reason why Laplace spectrum is important for the computational study of shapes.

In addition to the spectrum-geometry link, the eigenvalues of the Laplace operator have two other properties which make them useful for shape analysis, [11]. These are: 1) the Laplace operator is invariant to isometric transformations and 2) the spectrum depends continuously on the deformations applied to the boundary of the object. The advantage of the first property is obvious since isometric transformations do not alter the shape. In addition to this, the second property states that there is a continuous link between the differences in eigenvalues and the difference in shape, which makes eigenvalues ideal for measuring shape differences.

Unfortunately, it has also been shown that there exists isospectral non-congruent objects, i.e. objects with different shape but the same spectrum [28]. Therefore, theoretically the Laplace spectrum does not uniquely identify shapes. However, as stated in [3], practically this does not cause a problem mostly because the constructed isospectral non-congruent objects in 2D and 3D are rather extreme examples with nonsmooth boundaries.

The spectral signature, shape-DNA, proposed in [3] is inspired from the properties given above. For a given shape Ω , its shape-DNA is the first N modes of the spectrum of the Laplace operator defined on Ω : $[\lambda_1, \lambda_2, \dots, \lambda_N]$. In addition to the properties the shape-DNA inherits from the eigenvalues, the authors also proposed several normalisations to obtain almost scale invariance². The normalisations used in the experiments in [3], [7], [27], [29] are given as $\lambda_n \rightarrow \lambda_n V_{\Omega}^{2/d}$ and $\lambda_n \rightarrow \lambda_n / \lambda_1$.

In [3], the authors also defined a shape distance based on shape-DNA. Either using the original or its scale invariant version, this distance is given as

$$\rho_{SD}^N(\Omega_{\lambda}, \Omega_{\xi}) \triangleq \left[\sum_{n=1}^N (\lambda_n - \xi_n)^2 \right]^{1/2}, \quad (2)$$

where Ω_{ξ} denotes the object with the spectrum $\{\xi_n\}_{n=1}^{\infty}$. Using $\rho_{SD}^N(\Omega_{\lambda}, \Omega_{\xi})$, the authors were able to distinguish between distinct shapes [27], construct shape manifolds based on the pairwise distances and perform statistical comparisons [7], [29].

However, as also pointed out in [14], due to the diverging nature of the spectrum, ρ_{SD}^N suffers from three essential drawbacks limiting its usability: i) differences at higher modes of the spectrum have higher impacts on the final distance value even though they are not necessarily more informative about the intrinsic geometry, ii) the distance is extremely sensitive to the signature size N , while the choice of this parameter is arbitrary, and iii) the distance cannot be defined over the entire spectrum because it does not yield a proper metric in

²We use the term ‘‘almost’’ because scale invariance is an application dependent concept and the definition of scale difference between arbitrary objects is a mathematically vague notion. A further discussion of scale invariance is outside the scope of this article and we refer the reader to [3].

that case. Therefore, defining a sound and intuitive distance based on the spectrum is still an open question for which we propose a solution in the next section.

III. WEIGHTED SPECTRAL DISTANCE - WESD

This section presents the proposed spectral distance, WESD, the analysis of the heat-trace leading to its definition and its theoretical properties. The structure of presentation aims to separate the definition of the distance, which is essential for its practical implementation, from the details related to its derivation and theoretical properties. In this light, we first present the definitions and mention the associated properties with appropriate references to the following subsections, which contain further details.

We define the Weighted Spectral Distance - WESD - for two closed bounded domains with piecewise smooth boundaries, $\Omega_\lambda, \Omega_\xi \subset \mathbb{R}^d$ as

$$\rho(\Omega_\lambda, \Omega_\xi) \triangleq \left[\sum_{n=1}^{\infty} \left(\frac{|\lambda_n - \xi_n|}{\lambda_n \xi_n} \right)^p \right]^{1/p}, \quad (3)$$

with $p \in \mathbb{R}$ and $p > d/2$. Unlike the distance given in Equation (2), WESD is defined over the entire eigenvalue sequence and the factor p is not fixed to 2. In addition, the difference at each mode contributes to the overall distance proportional to $|\lambda_n - \xi_n|/\lambda_n \xi_n$ instead of $|\lambda_n - \xi_n|$. The additional $\lambda_n \xi_n$ factor (seeming like a simple addition to Equation 2) actually arises from analysing the relation between the n^{th} mode of the spectrum and the heat-trace, which will be presented in Section III-A. This analysis also provides WESD with a geometric intuition. Furthermore, for $p > d/2$ the infinite sum in the definition is guaranteed to converge to a finite value for any pair of shapes. Hence, WESD exists. In addition to its existence, WESD also satisfies the triangular inequality making it a pseudometric. These points are proven in Section III-B. Moreover, the pseudometric WESD has a multi-scale aspect with respect to p . In Section III-C we show that adjusting p controls the sensitivity of WESD with respect shape differences at finer scales, i.e. with respect to geometric differences at local level such as thin protrusions or small bumps. Thus, for higher values of p the distance becomes less sensitive to finer scale differences.

In addition to WESD, we define the normalised score for shape dissimilarity *nWESD* as

$$\bar{\rho}(\Omega_\lambda, \Omega_\xi) \triangleq \frac{\rho(\Omega_\lambda, \Omega_\xi)}{\mathbf{W}(\Omega_\lambda, \Omega_\xi)} \in [0, 1), \quad (4)$$

which maps $\rho(\Omega_\lambda, \Omega_\xi)$ to the $[0, 1)$ interval using the shape-dependent normalisation factor

$$\mathbf{W}(\Omega_\lambda, \Omega_\xi) \triangleq \left\{ C + K \cdot \left[\zeta \left(\frac{2p}{d} \right) - 1 - \left(\frac{1}{2} \right)^{\frac{2p}{d}} \right] \right\}^{\frac{1}{p}}.$$

The factors C and K are the shape based coefficients defined in Corollary 1, and $\zeta(\cdot)$ is the Riemann zeta function [30]. Being confined to $[0, 1)$, *nWESD* allows us to i) compare dissimilarities of different pairs of shapes and ii) easily use the shape dissimilarity in combination with scores quantifying

other type of differences between objects such as volume overlap in case of matching or Jacard's index in case of accuracy assessment.

One important issue in defining a distance or a score using the entire eigenvalue sequence is computational limits. In practice we can only compute a finite number of eigenvalues and therefore, can only approximate such distances. Considering this, here we define the finite approximations of WESD and *nWESD* using the smallest N eigenvalues as

$$\rho^N(\Omega_\lambda, \Omega_\xi) \triangleq \left[\sum_{n=1}^N \left(\frac{|\lambda_n - \xi_n|}{\lambda_n \xi_n} \right)^p \right]^{1/p} \quad (5)$$

$$\bar{\rho}^N(\Omega_\lambda, \Omega_\xi) \triangleq \frac{\rho^N(\Omega_\lambda, \Omega_\xi)}{\mathbf{W}(\Omega_\lambda, \Omega_\xi)} \in [0, 1), \quad (6)$$

where N is a truncation parameter. Previous works, such as [3], [5], [17], [20], [26], also define distances based on finite number of modes. However, their view on the distance definition was first to construct finite shape signatures and then to define a distance on the signatures. Therefore, the signature size was a critical component of the definition itself. Furthermore, the effects of the choice of the signature size on the distance values have not been carefully analysed in these works. The view presented here defines the distance directly using the entire sequence without constructing a finite signature. This alleviates the importance of the signature size on the distance. The finite computation given in Equations 5 and 6 are viewed as approximations to the distance and N as the truncation parameter. In this conceptually different setting, unlike previous works, we provide in Section III-D a careful analysis of the choice of N on the spectral distance. Specifically, we prove that $\lim_{N \rightarrow \infty} |\rho(\Omega_\lambda, \Omega_\xi) - \rho^N(\Omega_\lambda, \Omega_\xi)| = 0$ and $\lim_{N \rightarrow \infty} |\bar{\rho}(\Omega_\lambda, \Omega_\xi) - \bar{\rho}^N(\Omega_\lambda, \Omega_\xi)| = 0$. Furthermore, we provide a theoretical upper bound for these errors that shows how fast they decrease in the worst case leading to a principled strategy for choosing N .

Section III-E ends the section by focusing on the invariance of WESD and *nWESD* to global scale (relative size) differences between objects. Specifically, we discuss how an ‘‘approximate’’ scale invariance can be attained for WESD and *nWESD* by following the same strategy proposed in [3].

A. Analysis of the Heat-Trace and Derivation of WESD

We derive WESD by analysing the mathematical link between the spectrum of an object and its geometry. This link is given by the heat-trace defined in Equation (1). Let us consider the heat-trace as a function of both t and the spectrum, $Z(t, \lambda_1, \lambda_2, \dots)$. The main question we answer is how much the $Z(t, \cdot)$ function changes when we change the n^{th} mode of the spectrum from λ_n to ξ_n . Considering the polynomial expansion equivalent to $Z(t, \cdot)$ given in Equation (1), one can see that the change in the value $Z(t, \cdot)$ is directly related to the changes in the coefficients $a_{s/2}$ and so to changes in the integrals over the local invariants. By analysing the influence of the change in the n^{th} mode on $Z(t, \cdot)$, we actually analyse the influence of this change on the integrals over local geometric invariants. Following this line of thought, we

quantify the influence of the change from λ_n to ξ_n on $Z(t, \cdot)$ in terms of λ_n and ξ_n . This can be done by defining

$$\Delta_Z^n \triangleq \int_0^\infty |Z(t, \dots, \lambda_{n-1}, \lambda_n, \lambda_{n+1}, \dots) - Z(t, \dots, \lambda_{n-1}, \xi_n, \lambda_{n+1}, \dots)| dt,$$

which is simply the L_1 -norm of the difference between the functions that is linked to the difference between λ_n and ξ_n . Replacing $Z(t, \cdot)$ with its definition leads to

$$\Delta_Z^n = \int_0^\infty |e^{-\lambda_n t} - e^{-\xi_n t}| dt \quad (7)$$

Without loss of generality let us assume $\xi_n \geq \lambda_n$. Then

$$e^{-\lambda_n t} \geq e^{-\xi_n t} \text{ for } t > 0.$$

We can then evaluate the integral in Equation (7) as

$$\Delta_Z^n = \int_0^\infty e^{-\lambda_n t} - e^{-\xi_n t} dt = \frac{|\lambda_n - \xi_n|}{\lambda_n \xi_n}.$$

Δ_Z^n captures the influence of the difference at the n^{th} mode on $Z(t, \cdot)$. Now, aggregating these influences across all modes leads to the definition of WESD

$$\rho(\Omega_\lambda, \Omega_\xi) = \left[\sum_{n=1}^\infty (\Delta_Z^n)^p \right]^{1/p} = \left[\sum_{n=1}^\infty \left(\frac{|\lambda_n - \xi_n|}{\lambda_n \xi_n} \right)^p \right]^{1/p}.$$

Surprisingly, the formulation of WESD, which results from the analysis presented above, also has very beneficial properties that makes it theoretically sound and useful in practical applications. These properties will be analysed in the following.

Before delving into this analysis though let us make two remarks. The first relates $\rho_{SD}(\cdot, \cdot)$ (Equation (2)) to the analysis of the heat-trace presented above.

Remark 1. *Let us define*

$$\Delta_Z^{n,m} \triangleq \left| \int_0^\infty \frac{d^m}{dt^m} Z(t, \dots, \lambda_{n-1}, \lambda_n, \lambda_{n+1}, \dots) - \frac{d^m}{dt^m} Z(t, \dots, \lambda_{n-1}, \xi_n, \lambda_{n+1}, \dots) dt \right|.$$

One can see that $\Delta_Z^{n,0} = \Delta_Z^n$. Evaluating this integral yields $\Delta_Z^{n,m} = |\lambda_n^{m-1} - \xi_n^{m-1}|$. By setting $m = 2$ $\rho_{SD}(\cdot, \cdot)$ can be derived as follows

$$\rho_{SD}(\Omega_\lambda, \Omega_\xi) = \left[\sum_{n=0}^\infty (\Delta_Z^{n,2})^2 \right]^{1/2}.$$

This relationship not only relates WESD to $\rho_{SD}(\cdot, \cdot)$ but also provides the link between $\rho_{SD}(\cdot, \cdot)$ and the heat-trace. One notices that the functional difference definition used in this remark differs from the previous one used to derive WESD, see Equation 7. This is because ρ_{SD} cannot be derived from the L_1 -distance definition used previously but can be derived from the less ideal definition used in this remark. The existence of an alternative derivation of ρ_{SD} that would start from an appropriate functional difference is an open question.

The second remark notes the link between WESD and *Global Point Signatures* (GPS), a local shape descriptor, presented in [5].

Remark 2. *GPS, as presented in [5], is defined for each point in an object Ω_λ as the infinite series $GPS_{\Omega_\lambda}(\mathbf{x}) \triangleq \{\Phi_{\lambda,n}(\mathbf{x})\} \triangleq \{\lambda_n^{-1/2} f_n(\mathbf{x})\}_{n=1}^\infty$, where $\mathbf{x} \in \Omega_\lambda$ and $f_n(\mathbf{x})$ is the n^{th} eigenfunction. GPS has a connection to WESD arising from the following element-wise integrals*

$$\int_{\Omega_\lambda} \Phi_{\lambda,n}^2(\mathbf{x}) d\mathbf{x} = \int_{\Omega_\lambda} [\lambda_n^{-1/2} f_n(\mathbf{x})]^2 d\mathbf{x} = \lambda_n^{-1},$$

where the equality arises from the fact that eigenfunctions form an orthonormal basis in Ω_λ [11], i.e. $\int_{\Omega_\lambda} f_n(\mathbf{x}) f_m(\mathbf{x}) d\mathbf{x} = \delta(n - m)$ with $\delta(\cdot)$ being the Dirac's delta. Considering this integral, WESD can also be regarded as a distance between GPS' of two objects as

$$\left\{ \sum_{n=1}^\infty \left[\int_{\Omega_\lambda} \Phi_{\lambda,n}^2(\mathbf{x}) d\mathbf{x} - \int_{\Omega_\xi} \Phi_{\xi,n}^2(\mathbf{x}) d\mathbf{x} \right]^p \right\}^{1/p} = \rho(\Omega_\lambda, \Omega_\xi).$$

This link also provides an alternative view on the normalisation factor $\lambda_n^{-1/2}$ used in GPS. In [5] author justifies this normalisation factor by noting that for an object the Green's function can be written as an inner product in the GPS domain, see Section 4 in [5]. This is later used to argue the geometric meaning of GPS as authors point out the use of Green's function in different shape processing tasks. Our link between GPS and WESD provides an alternative view on the normalisation factor as it connects this local signature to the heat-trace $Z(t)$.

B. Existence of the Pseudometric WESD

WESD is defined as the limit of an infinite series as given in Equation (3). For such a distance to be a proper one, actually a pseudometric in this case, the limit of the infinite series should exist for any two spectra. In the case of WESD, this is not evident because it is defined over the entire spectra and each spectrum is a divergent sequence. The first corollary presented below proves that when $p > d/2$ WESD indeed satisfies this condition, i.e. the infinite series converges. The corollary further provides an upper bound for this limit, which is used to construct nWESD. We would like to note that for the ease of presentation, the proofs for all the following corollaries and lemmas are given in Appendix B in the supplemental material.

Corollary 1. *Let $\Omega_\lambda \subset \mathbb{R}^d$ and $\Omega_\xi \subset \mathbb{R}^d$ be any two closed domains with piecewise smooth boundaries and $\{\lambda\}_{n=1}^\infty$ and $\{\xi\}_{n=1}^\infty$ be their Laplace spectrum. Then the weighted spectral distance*

$$\rho(\Omega_\lambda, \Omega_\xi) = \left[\sum_{n=1}^\infty \left(\frac{|\lambda_n - \xi_n|}{\lambda_n \xi_n} \right)^p \right]^{1/p}$$

converges for $p > \frac{d}{2}$. Furthermore,

$$\rho(\Omega_\lambda, \Omega_\xi) < \left\{ C + K \cdot \left[\zeta \left(\frac{2p}{d} \right) - 1 - \left(\frac{1}{2} \right)^{\frac{2p}{d}} \right] \right\}^{1/p}, \quad (8)$$

where $\zeta(\cdot)$ is the Riemann zeta function and the coefficients C and K are given as

$$C \triangleq \sum_{i=1,2} \left[\frac{d+2}{d \cdot 4\pi^2} \cdot \left(\frac{B_d \hat{V}}{i} \right)^{\frac{2}{d}} - \frac{1}{\mu} \cdot \left(\frac{d}{d+4} \right)^{i-1} \right]^p$$

$$K \triangleq \left[\frac{d+2}{d \cdot 4\pi^2} \cdot \left(B_d \hat{V} \right)^{\frac{2}{d}} - \frac{1}{\mu} \cdot \frac{d}{d+2.64} \right]^p$$

$$\hat{V} \triangleq \max(V(\Omega_\lambda), V(\Omega_\xi)), \quad \mu \triangleq \max(\lambda_1, \xi_1),$$

where $V(\cdot)$ denotes the volume (or area in 2D) of an object.

The Inequality (8) states that WESD has a shape-dependent upper bound. We thus can map the WESD to the $[0, 1)$ interval through normalising it with this upper bound. The nWESD score, given in Equation 4 is constructed based on this strategy. Since its existence is established next we prove that WESD is a pseudometric, i.e. satisfies the other criteria to be a pseudometric, such as the triangle inequality.

Corollary 2. $\rho(\Omega_\lambda, \Omega_\xi)$ is a pseudometric for $d \geq 2$.

We note that WESD is not a metric because the spectrum is invariant to isometries, which is a desirable property for shape analysis. However, in addition to this, the spectrum is also invariant to isospectral non-congruent shapes. This is not desirable but does not cause problems in practice as discussed in Section II and also confirmed in our experiments.

C. On the multi-scale aspect of WESD

The previous section highlighted the role of p on the convergence properties of WESD and therefore on its existence. We now demonstrate that p also provides WESD a multi-scale characteristic. The sensitivity of WESD to the shape differences at finer scales depends on the value of p . Specifically, we show that the higher the value p the less sensitive WESD is to finer scale details and its sensitivity increases as p gets lower.

The multi-scale aspect of WESD arises from the relationship between the Laplace operators and heat diffusion processes [31]. We first present an intuitive summary of this relationship, which is about the multi-scale aspect of $Z(t)$ and t in particular. For a more mathematical treatment we refer the reader to [6]. As stated in [6] and [14], t can be interpreted as the time variable in a heat diffusion process within an object. A useful visual analogy to consider here is the Laplacian smoothing of a surface where t would correspond to the amount of smoothing. Similar to the surface smoothing, as t increases, the local geometric details of an object, such as sharp ridges or steep valleys, lose further their influence on the $Z(t)$ value. As a result $Z(t)$ becomes somewhat insensitive to these local geometric details, i.e. shape details at finer scales. From an alternative view, the value of $Z(t)$ loses its information content with regards to local geometric details. This effect intuitively summarizes the multi-scale characteristic of the heat-trace with respect to t .

Having explained the multi-scale aspect of $Z(t)$, we now analyse how this aspect is reflected upon the eigenvalues. To do so let us define the *influence ratio* $\mathcal{D}(n, t) \triangleq \frac{e^{-\lambda_n t}}{Z(t)}$. This

ratio captures the influence of the n^{th} mode on the heat-trace. In other words, the higher the ratio, the higher the influence of λ_n on the value of $Z(t)$ at that specific t . The following lemma compares the influence ratios of different modes and how this comparison depends on t .

Lemma 1. Let $\Omega_\lambda \subset \mathbb{R}^d$ represent an object with piecewise smooth boundary and $\mathcal{D}(l, t) \triangleq \frac{e^{-\lambda_l t}}{Z(t)}$ be the corresponding influence ratio of mode l at t . Then for any two spectral indices $m > n > 0$

$$\mathcal{D}(n, t) > \mathcal{D}(m, t), \quad \forall t > 0$$

and particularly for two t values such that $t_1 > t_2$

$$\frac{\mathcal{D}(m, t_1)}{\mathcal{D}(n, t_1)} < \frac{\mathcal{D}(m, t_2)}{\mathcal{D}(n, t_2)}.$$

The first inequality of the lemma indicates that the lower modes in the spectrum have more influence on the value of $Z(t)$ than the higher modes. The second inequality shows that the influence of the higher modes become more prominent as t decreases. Considering that for lower t values $Z(t)$ is more informative with regards to shape details at finer scales, Lemma 1 suggests that the higher modes are more important for finer scales than for coarser scales. We illustrate this observation on a synthetic example shown in Figure 2 with the pair (a) + (b) being an example showing coarser shape differences and the pair (a) + (c) showing finer differences. The plots given in Figure 2 (d) and (e) show the corresponding spectral differences observed at modes between 1 and 150. Between (a) and (b) the shape differences are at the coarse level. According to Lemma 1 these differences should show up at the very first modes. On the other hand, between (a) and (c) the differences are at a finer scale and furthermore the objects are very similar at the coarse level. Lemma 1 states that these differences therefore, should show up at higher modes and the differences at the lower modes should be low. Satisfying these expectations, the differences at the first few modes shown in plot (d) have relatively large values compared to the ones in plot (e). Furthermore, the amplitude of the differences at higher modes are generally larger in plot (e) than in plot(d), especially after 100.

In order now to connect these findings to WESD and p let us present the following corollary, which studies the influence of p on the components inside the infinite sum defining the distance.

Corollary 3. Let Ω_λ and Ω_ξ be two objects with piecewise smooth boundaries. Then for any two scalars with $p > d/2$, $q > d/2$, $p \geq q$ and for all n with $|\lambda_n - \xi_n| > 0$ there exists a $M > n$ so that $\forall m \geq M$

$$\frac{\left(\frac{|\lambda_m - \xi_m|}{\lambda_m \xi_m} \right)^p}{\left(\frac{|\lambda_n - \xi_n|}{\lambda_n \xi_n} \right)^p} \leq \frac{\left(\frac{|\lambda_m - \xi_m|}{\lambda_m \xi_m} \right)^q}{\left(\frac{|\lambda_n - \xi_n|}{\lambda_n \xi_n} \right)^q}$$

Thus, the relative contributions of the higher spectral modes on $\rho(\Omega_\lambda, \Omega_\xi)$ with respect to the contributions of the lower modes depend on the value of p . Specifically, the higher spectral modes become more influential as p decreases. Combining this finding with the result of Lemma 1, we follow that as p




Fig. 2: Multi-scale characteristics of different spectral modes:(a), (b) and (c) show three synthetic shapes. In (d) we plot the absolute differences between the corresponding modes of (a) and (b) with respect to the spectral index. In (e) we plot the same difference for the shapes in (a) and (c). The shape difference between (a) and (b), which is at a coarser level, is already captured at the lower spectral modes. The difference between (a) and (c) results in lower differences in lower spectral modes because these objects are more similar at a coarser level. At the higher spectral modes though, the difference between (a) and (c) becomes more prominent since these objects differ more substantially at the finer scales. The plots in (d) and (e) demonstrate that the higher modes for a given object are more important for finer scale shape details.

increases WESD gives less importance to differences at higher spectral modes and therefore becomes less sensitive to the shape differences at finer scales. This provides WESD with a multi-scale aspect with respect to p and also provides us the intuition for choosing a proper value for p .

D. Finite Approximations of WESD and nWESD

One of the important practical questions regarding spectral distances is the number of modes to be included in the calculation of the distance. The computation of eigenvalues and eigenfunctions can be expensive and inaccurate especially for the higher modes. Therefore, spectral distances require the user to set a finite number of modes to be used. This parameter is often referred to as the signature size. Having defined the distance over the entire sequence, we refer to it as the *truncation parameter*. This actually provides a different perspective on the number of modes used to compute the distance. In previous works, such as [3], [5], [20], the value of this parameter, viewed as the signature size, is often set arbitrarily and its effect on the distances have not been carefully analysed. Here, viewing it as a truncation parameter, we study its influence. Specifically, we formulate the difference between using the entire spectra to only using a finite number of modes as an *approximation/truncation error*. So we analyse how this error changes with respect to the truncation parameter. We specifically show in the next corollary that the errors

in approximating WESD and nWESD by the first N modes converges to zero as N increases. Furthermore, we provide an upper bound for both errors as a function of N .

Corollary 4. Let $\rho^N(\Omega_\lambda, \Omega_\xi)$ be the truncated approximation of $\rho(\Omega_\lambda, \Omega_\xi)$ based on the first N modes and $\bar{\rho}^N(\Omega_\lambda, \Omega_\xi)$ of $\bar{\rho}(\Omega_\lambda, \Omega_\xi)$. Then $\forall p > d/2$

$$\lim_{N \rightarrow \infty} |\rho - \rho^N| = 0$$

and

$$\lim_{N \rightarrow \infty} |\bar{\rho} - \bar{\rho}^N| = 0.$$

Furthermore, for a given $N \geq 3$ the truncation errors $|\rho - \rho^N|$ and $|\bar{\rho} - \bar{\rho}^N|$ can be bounded by

$$\begin{aligned} |\rho - \rho^N| &< \left\{ C + K \cdot \left[\zeta \left(\frac{2p}{d} \right) - 1 - \left(\frac{1}{2} \right)^{\frac{2p}{d}} \right] \right\}^{\frac{1}{p}} \quad (9) \\ &\quad - \left\{ C + K \cdot \left[\sum_{n=3}^N \left(\frac{1}{n} \right)^{\frac{2p}{d}} \right] \right\}^{\frac{1}{p}} \\ |\bar{\rho} - \bar{\rho}^N| &< 1 - \left\{ \frac{C + K \cdot \left[\sum_{n=3}^N \left(\frac{1}{n} \right)^{\frac{2p}{d}} \right]}{C + K \cdot \left[\zeta \left(\frac{2p}{d} \right) - 1 - \left(\frac{1}{2} \right)^{\frac{2p}{d}} \right]} \right\}^{\frac{1}{p}} \quad (10) \end{aligned}$$

The above corollary has important practical implications. First of all, the sensitivities of ρ^N and $\bar{\rho}^N$ with respect to N decreases as N increases. For any application relying on the shape distances, such as constructing low dimensional embeddings, this reduced sensitivity is particularly important as it provides stability with respect to N both for the distance and for the application using the distance. We note that the opposite is true for ρ_{SD}^N , which is one of the main disadvantages of this distance.

In addition, Corollary 4 can guide the choice for the number of modes N and the norm type p . Specifically, the error upper bounds given in Equations 9 and 10 provide the worst case errors for a given pair of shapes without the need to compute the eigenvalues. So for instance, once a number of modes are computed then based on the distance value obtained so far and the worst case error computed using the upper bounds, one can decide whether to compute more modes or not. Furthermore, these upper bounds are shape-specific as they depend on C and K . One can go one step further and define a shape-independent *residual ratio* for $N \geq 3$ and $p > d/2$ as

$$R(N, p) \triangleq 1 - \left[\frac{\sum_{n=3}^N \left(\frac{1}{n} \right)^{\frac{2p}{d}}}{\zeta \left(\frac{2p}{d} \right) - 1 - \left(\frac{1}{2} \right)^{\frac{2p}{d}}} \right]^{\frac{1}{p}}. \quad (11)$$

that satisfies $R(N, p) > \bar{\rho} - \bar{\rho}^N$, for which the proof is given in Proposition 1 in Appendix B. $R(N, p)$ can be used to select the parameters N and p as it quantifies the quality of the approximation for a given pair of (N, p) in terms of the error upper bounds.

In Figure 3, we plot $R(N, p)$ versus N for different settings of p and $d = 2, 3$. Besides the obvious point that the error upper bound decreases for increasing N we also notice that




Fig. 3: Choosing N : The figures plot the residual ratio $R(N, p)$ versus N for different p values in 2D (left) and in 3D (right). As expected the error upper bound drops with increasing N . The rate of decrease also becomes faster with increasing p . This inverse relation suggests the trade-off between N and the sensitivity of WESD to finer scale shape differences since WESD becomes less sensitive as p increases, see Section III-C.

i) the behavior in 2D and 3D are similar and ii) the rate of decrease of the error upper bound is much faster for higher p . Considering the multi-scale aspect of WESD captured in p , this behavior is interesting. It demonstrates that the choice of p and N are correlated and suggests a trade-off between the rate of decrease of the truncation error and the sensitivity of WESD to shape differences at finer scales. In theory, the choice of these parameters depends on the application and the expected shape differences. If one expects coarse scale differences then choosing a large p and small N might be sufficient. However, if one is interested in finer scale differences then a small p value will be required, which in turn will require a large N value to have a decent approximation. The important aspect of $R(N, p)$ is that it is universal, i.e. it does not depend on the objects. So it can be used in any type of application to choose the parameter pair N, p and to have a rough estimate of the computational costs for computing the distance WESD. We note once again, the specific values should be chosen based on the application and the shapes at hand.

E. Invariance to global scale differences

We end this section studying the impact of global scale differences on WESD and how invariance to such differences can be attained. We would like to note that the notion of global scale in this section refers to the relative size of an object, which is not to be confused with the notion of multi-scale discussed in Section III-C. The spectrum of an object depends on the object's size, i.e. a global scale change alters all the eigenvalues by a constant multiplicative factor [11]. As a result, the global scale difference between two objects contributes to the spectral shape distance WESD. In some applications this contribution might not be desirable, for instance in an object recognition task, where objects in the same category have varying sizes. Therefore, it is a useful property of a shape distance to allow invariance to global scale differences.

Reuter et al. [3] proposed different approximations for normalising the effects of scale differences on the spectrum. In particular, the authors use two different normalisations in their experiments in [7], [27], [29]. Both normalisations directly act

on the eigenvalues. The first one normalises the eigenvalues with respect to the volume (area in 2D or surface area for Riemannian manifolds) and is given as $\lambda_n \rightarrow \lambda_n V_{\Omega_\lambda}^{2/d}$. The second one normalises the eigenvalue with respect to the first eigenvalue in the sequence as $\lambda_n \rightarrow \lambda_n / \lambda_1$. Both of these strategies can be used when computing distances with WESD. Furthermore, since these strategies do not alter the mathematical characteristics of the entire spectrum the theoretical properties of WESD and nWESD hold either way. For our experiments we adopt the first strategy, volume normalisation, using the volume as defined in Euclidean geometry. When using the volume normalised eigenvalues, the only change that applies to the technical details presented so far is \hat{V} in Equation 8 becomes $\hat{V} = 1$. The rest applies directly without any modification.

We would also like to note that estimating the global scale difference between two arbitrary objects is not always a well-posed problem. It is especially hard when the objects are of different category, e.g. an octopus and a submarine. Furthermore, the scale normalisation is application dependent and it might not be desirable for all applications. In Section IV-C2 we present such an example where we analyse the temporal change of the left ventricle shape during a heart cycle. In this case, the volume change is essential for analysing the heart of the same patient so that scale invariance is not appropriate.

IV. EXPERIMENTS

This section presents a variety of experiments on synthetic and real data highlighting the strengths and weaknesses of WESD and nWESD. We start by briefly explaining the details of the numerical implementation of WESD used in the experiments presented here. Then in Section IV-B, the proposed distances are applied to synthetically generated objects demonstrating that

- (i) Ordering objects with respect to their shapes using nWESD results in a visually coherent series (Section IV-B1),
- (ii) WESD is useful for constructing low dimensional embeddings, in particular it yields stable embeddings with respect to the signature size N , (Section IV-B2) and
- (iii) WESD is a suitable distance for shape retrieval, which is shown through experiments on the SHREC dataset [27] (Section IV-B3).

Lastly, in Section IV-C WESD is applied to real objects extracted from 3D medical images. We focus on two examples from a wide variety of applications WESD and nWESD can be beneficial to: population studies of brain structures and analysis of 4D cardiac images.

A. Implementation Details

There are two different aspects in the implementation of WESD: the numerical computation of the Laplace spectra and the parameter settings. First, any numerical method tailored towards computing the eigenvalues of the Laplace operator can be used to compute WESD. Examples of such method are listed in [3], [32]. Our specific implementation represents objects simply as binary images with the foreground defining

Ω . Using the Cartesian grid of the image, it discretizes Δ_Ω through finite difference scheme (see also Chapter 2 of [32]). This step yields a sparse matrix of which we compute the eigenvalues via Arnoldi’s method presented in [33] and implemented in MATLAB[®]. We choose this specific implementation as 1) it is simple 2) it does not introduce any additional parameters and 3) when working with images it avoids any extra preprocessing steps, such as surface extraction or mesh construction.

With regards to the second implementation aspect, we set the parameters N and p empirically. Based on Section III-D, we set $p = 1.5$ or $p = 2.0$ in 2D and $p = 2.0$ in 3D. These values result in a relatively fast diminishing upper bound of the truncation error with respect to N (see Figure 3) while being sensitive to shape differences at finer scales. In both 2D and 3D, we chose $N = 200$ for the number of modes as the truncation error seemed to vanish at that point. Furthermore, in addition to the theoretical considerations on the effects of N and p on WESD given in Section III-D, in Sections IV-B2 and IV-B3 we experimentally study the effects of these parameters on applications using WESD, specifically on constructing low dimensional embeddings and shape retrieval.

B. Synthetic Data

We conduct three experiments: first two are on 2D objects and the last one is on 3D objects. For all of the experiments, we use the scale invariant versions of the spectra obtained by normalising the eigenvalues with the object’s volume as described in Section III-E. As a result the distances WESD and nWESD become “almost” invariant to global scale differences.

1) *Ordering of Shapes*: For the first experiment we created two synthetic datasets. Each dataset consists of a reference object and random deformations of this reference. These deformed versions are generated by transforming the reference via random deformations of varying magnitude and amount of nonlinearity. As a result the datasets contain objects that are very similar to the reference ones and objects that are substantially different. Figures 4(a) and (b) show some examples from these datasets where the binary images to the very left show the reference objects. In the first dataset, the reference object is a disc. There are a total of 500 random deformations of this reference disc. The first 400 are generated via non-linear deformations while the last 100 are isometric transformations. In the second dataset, the reference is a slightly more complicated object (see Figure 4(b)) and in total there are 400 random transformations of this reference. The first 300 are generated by non-linear deformations and the last 100 produced via isometric transformations. All objects are discretized as binary maps with a size of 200×200 pixels. The numerical computations are performed on these image grids as discussed earlier.

We computed the nWESD scores ($\bar{\rho}^N$ with $p = 1.5$ and $N = 200$) between the reference and the deformed objects in each dataset. Based on these scores, we then ordered the deformed objects according to their similarity in shape to the reference. Figures 4 (c) and (d) show examples of the resulting orderings. We notice that the orderings are visually meaningful




Fig. 4: Shape-based ordering of objects: We generate two artificial datasets each consisting of a reference object and its random deformations. Samples from the datasets are shown in (a) and (b). The binary images to the very left show the reference objects for each dataset. We then ordered all the deformed objects with respect to the nWESD scores between the object and the reference. The graphs in (c) and (d) plot these orderings. Based on visual inspection the ordering is quite reasonable.

, i.e. the further the deformed objects visually deviate from the references, the higher their nWESD score is. Furthermore, all the objects generated via isometric transformations yielded scores close to zero as a result of the invariance of the proposed scores to this type of transformation.

2) *Low Dimensional Embeddings*: In the second experiment we focus on creating low dimensional embeddings. We compare the embeddings constructed by WESD with the ones constructed using ρ_{SD}^N (Equation (2)), the distance proposed in [3]. We do so based on the TOSCA dataset (toolbox for surface comparison and analysis), [34], [35]. This dataset contains binary segmentations of 5 human, 5 centaurs and 5 horses as shown in Figure 5(a). We compute the pairwise affinity matrices between objects via ρ_{SD}^N and WESD (ρ^N with $p = 2.0$). We then apply the ISOMAP algorithm [36] to these matrices, which maps the 15 objects to a 2D plane based on the pairwise shape distances. We repeat this experiment for affinity matrices computed using different number of spectral modes, i.e. $N = 50, 100, 200$, to demonstrate the effect of the signature size (truncation parameter) on both distances.

The plots in Figures 5(b),(d) and (f) present the resulting 2D embeddings of the dataset using ρ_{SD}^N . The embeddings are substantially different for different N . This variation arises due to high sensitivity of ρ_{SD}^N towards the signature size N . Moreover, the embeddings obtained using higher N are less satisfactory in terms of separating the three different




Fig. 5: Low dimensional embeddings: (a) The 15 objects used in this experiment. The graphs plot the 2D embeddings of the objects based on the affinity matrices constructed by ρ_{SD}^N and WESD (ρ^N). Each row presents the results based on different N : 50, 100 and 200 from top to bottom respectively. The structures of the 2D embedding based on ρ_{SD}^N are quite different for different N . WESD however, produces embeddings that are similar. This demonstrates the stability of the embedding with respect to N when WESD is used.

object classes. This is actually expected since the spectral modes with higher indices dominate the value of ρ_{SD}^N even though they are not informative with regards to the overall geometry and thus, negatively impact the outcome. The plots in Figures 5(c), (e) and (g) present the embeddings obtained using WESD. The embeddings obtained at different N are very similar. This shows that the construction of the low dimensional embedding is stable with respect to N when

WESD is used. This is a direct consequence of the convergent behavior of WESD discussed in Sections III-B and III-D. As illustrated by this experiment, this property has very important practical implications.

The experimental analysis presented above requires computing a high number of eigenvalues in advance. Besides this option, through Corollary 4, WESD provides us the opportunity to perform the same analysis without the need to compute eigenvalues in advance. Once N modes are computed, one can use Equation 9 to analyse the stability of the embedding with respect to N and ultimately use this analysis to decide whether N is enough. In the following we perform such an analysis for this experiment.

At three different $N = (100, 150, 200)$, we computed the error upper bounds for each element of the affinity matrix using Equation 9. Let us refer to these bounds as $E^N(\Omega_1, \Omega_2)$, so we can write $\rho(\cdot, \cdot) \in [\rho^N(\cdot, \cdot), \rho^N(\cdot, \cdot) + E^N(\cdot, \cdot)]$. For the analysis we assume $\rho(\cdot, \cdot)$ can lie anywhere in this range, i.e. uniformly distributed. By randomly sampling these ranges for each element, we generate 5000 different affinity matrices and build low-dimensional embeddings for each matrix. In Figure 6 we plot the embeddings obtained with ρ^N , with markers, along with the minimum and the maximum low-dimensional coordinates each object attained during random sampling - indicated with a rectangle around each point. These rectangles show the *maximum* error on the low-dimensional embedding that we might be introducing by truncating the computation of WESD at N . Observing these graphs: i) one can provide a guarantee that adding more modes after $N = 200$ will not change the embedding much, ii) at $N = 150$ one could have stopped computing more modes because the embedding cannot change substantially and iii) at $N = 100$ the embedding theoretically can change so more modes might be necessary. An important point to note is that our experimental findings suggest that the theoretical upper bounds given in Corollary 4 are very conservative. In this experiment, for instance, the embedding does not change much between using $N = 100$ and $N = 200$, i.e. the maximum absolute coordinate change vector for all the objects is $(0.24 \times 10^{-4}, 0.22 \times 10^{-4})$.

3) *Shape-based Retrieval of 3D Objects*: In this last experiment with synthetic data, we focus on the application of *shape-based object retrieval*, i.e. given a test object identifying other “similar” objects within a dataset using shape information. Similarity in this context can be defined in various ways but the definition used here is semantic similarity, meaning that objects that are of the same semantic category (e.g. human bodies, aeroplanes, etc) are similar and objects of different categories are not. Shapes of similar objects have similar traits and properties. Shape distances used for retrieval purposes should be able to capture these traits yielding the lowest values between similar object pairs. Here, WESD’s value for shape-based retrieval is evaluated using the publicly available dataset SHREC presented in [27]³.

SHREC dataset consists of 600 3D non-rigid objects from 30 different categories, i.e. 20 objects per category. Ob-

³Available at <http://www.itl.nist.gov/iad/vug/sharp/contest/2011/NonRigid/>

	NN	FT	ST	E	DCG
WESD (ρ^N) $p=3.15, N=100$	0.9933	0.9020	0.9305	0.6900	0.9706
WESD (ρ^N) $p=2.0, N=100$	0.9933	0.8923	0.9238	0.6824	0.9691
ρ_{SD}^N $N=12, \text{norm1}$	0.9967	0.8896	0.9521	0.6959	0.9748
ρ_{SD}^N $N=12, \text{normA}$	0.9917	0.9153	0.9569	0.7047	0.9783
ρ_{SD}^N $N=15, \text{norm1}$	0.9933	0.8683	0.9431	0.6895	0.9705
MDS-CM-BOF	0.9950	0.9127	0.9691	0.7166	0.9822
SD-GDM-meshSIFT	1.0000	0.9720	0.9901	0.7358	0.9955

(a)




(b) $p = 3.15$ and $N = 100$ (c) $p = 3.15$ (d) $N = 100$

Fig. 7: Shape-based Object Retrieval Results on SHREC Dataset. a) Retrieval scores obtained by WESD for two different sets of N and p values along with the scores obtained by the distance proposed in [3] (values taken from [27].) b) Precision-Recall curves obtained for shape retrieval via WESD for the entire dataset. c) Effect of the signature size N on the retrieval scores obtained by WESD for a fixed $p = 3.15$. d) Effect of the norm type p on the same scores for a fixed $N = 100$.




Fig. 6: The theoretical analysis of the effects of truncating the computation of WESD at N modes on the low-dimensional embeddings: Points (markers) indicate the low-dimensional embeddings obtained by using $\rho^N(\cdot, \cdot)$. The rectangle around a point denotes the theoretical maximum extent that point might move to if infinite number of modes were used to compute the distance, i.e. if $\rho(\cdot, \cdot)$ were used. These bounds are computed using using Equation 9 without the need to compute more eigenvalues than N .

jects from the same category differ with substantial non-linear deformations, which makes retrieval in this dataset challenging. To evaluate the retrieval accuracy of WESD, first each object was converted from its original watertight surface mesh discretization to a 3D binary image using the Iso2mesh software package⁴. Then pairwise shape distances across the entire dataset were computed using WESD and the 600×600 affinity matrix was constructed, where each entry is a pairwise distance. This affinity matrix was then evaluated using the software provided with the dataset³. The evaluation consists of a variety of retrieval accuracy scores such as Nearest Neighbor (NN), First-Tier (FT), Second-Tier (ST), E-Measure (E), Discounted Cumulative Gain (DCG) and Precision-Recall curve. The first two rows of the table in

Figure 7(a) list these scores obtained using WESD for two different settings of the p and N values. Additionally, the next three rows of the same table show the results obtained using ρ_{SD}^N (Equation 2, [3]), as listed in [27]⁵, using two different types of scale normalisation (norm1: normalising with respect to the first eigenvalue, normA: area normalisation, see Section III-E for further details). These accuracy scores show that WESD and ρ_{SD}^N perform very similar in retrieval from the SHREC dataset. The plot given in Figure 7(b) shows the precision-recall curve of WESD ($p = 3.15$ and $N = 100$) for the entire dataset. The curve is very similar to the best curve obtained using ρ_{SD}^N shown in [27]. Once again, this confirms that both distances perform similarly.

For a complete comparison, the table given in Figure 7(a) also reports the results of the two methods that achieved the highest retrieval scores in the SHREC challenge: MDS-CM-BOF ([37]) and SD-GDM-meshSIFT ([38], [39]). Both of these methods construct multi-step pipelines that provide powerful tools for retrieval. However, one has to note that these methods also have added complexities in contrast to WESD and ρ_{SD} , such as: i) the need for a training dataset, ii) the correspondence requirements and the resulting need for remeshing, iii) computation of local key-points for feature computation, iv) high computational cost of geodesic distance matrix, and iv) local feature or point matching to compute a global distance using local features. Nevertheless, their high accuracies demonstrate the advantage of combining different fundamental components to achieve powerful retrieval tools.

Lastly, the graphs shown in Figure 7(c) and (d) provide an analysis of the retrieval results with respect to the parameters N and p . Graphs in Figure 7(c) plot the change of different

⁵We note that for these latter results a slightly different notation is used here than in [27] to conform to the overall notation of this article.

⁴<http://iso2mesh.sourceforge.net/cgi-bin/index.cgi>

retrieval scores with respect to the number of modes used N , i.e. signature size, keeping p fixed at 3.15. Graphs in Figure 7(d) plot the changes with respect to the norm type p keeping N fixed at 100. These graphs show that as N increases the scores seem to increase slowly and then converge. On the other hand, p has a stronger effect on the results than N , particularly on FT, ST and E scores. However, the changes in the scores with respect to changes in N or p are rather small especially compared to the relatively larger fluctuation of the FT score of ρ_{SD}^N with respect to the two sample N values provided in the table in Figure 7(a).

The experiment presented above showed that the retrieval power of WESD is similar to that of the distance ρ_{SD}^N proposed by Reuter et al. [3]. The soundness and theoretical properties of WESD do not come at the expense of lower retrieval power. On the contrary, WESD is able to leverage the descriptive power of the spectra while its properties guarantee that it does not suffer from similar drawbacks as other distances, such as sensitivity to signature size.

C. Real Data

The experiments on real data are conducted on segmentations of 3D structures obtained from magnetic resonance images (MRI). First, we apply WESD to subcortical brain structures. The experiment demonstrates WESD’s capabilities to differentiate categories of objects even in the presence of high intra-class variability. In the second experiment, we focus on temporal analysis of cardiac images. We apply nWESD to delineations of the blood pool of the left ventricle obtained from 3D + time cardiac MRI. The experiment shows that the shape dissimilarity measurements between time points correlates with the dynamic processes of the beating heart.

1) *Clustering Sub-Cortical Structures:* Medical research frequently relies on morphometric studies analysing anatomical shapes from medical images [40]. In this experiment we construct a low dimensional embedding of subcortical structures extracted from Magnetic Resonance Image (MRI) scans of different individuals based on WESD as well as shape-DNA based distance, ρ_{SD}^N , as proposed in [3].

For this experiment, we use the publicly available LPBA40 dataset [41]⁶. The dataset contains manual segmentations of various subcortical structures from MRI brain scans of 40 healthy subjects. Figures 8(a), (b) and (c) show some examples from these structures. There are two main difficulties associated with such datasets. First, the structures have very large intra-class (inter-subject) variability, i.e. the shape of an anatomical structure is often very different across subjects. Second, the segmentations were obtained by manually delineating the 3D objects on successive 2D slices. This creates inconsistencies between segmentations in two successive slices. Such inconsistencies in the end manifest themselves as local artefacts on the object. The protrusion that can be seen on the top of the second hippocampus in Figure 8(a) is an example of such an artefact. These artefacts can influence shape distances negatively.

We select six structures for each patient: left/right caudate nucleus, left/right putamen and left/right hippocampus,

resulting in 240 structures in total. We then create pairwise affinity matrices of the 240 structures first using ρ_{SD}^N with $N = 200$, as proposed in [3], and then WESD (ρ^N with $p = 2$ and $N = 200$). Finally, we use the ISOMAP algorithm [36] to construct 2D embeddings of the structures. Figures 8(d) and (e) show the resulting embeddings. We observe that the embedding obtained via WESD well clusters the data with respect to the anatomical structures. The separation of the clusters for the SD case, however, is more ambiguous, especially between putamen and hippocampus.

The embeddings presented above were obtained by directly using the manual segmentations without any preprocessing. A natural question is how do these embeddings change if the effects of various artefacts are reduced say via surface smoothing. To answer this question, we smooth the surface of the anatomical 3D models and recomputed the embeddings, which are shown in Figures 8(f) and (g). The embedding obtained with ρ_{SD}^N , although to a lesser extent, still suffers from similar ambiguity as in Figure 8(d). The new embedding based on WESD on the other hand, compared to Figure 8(e), even more clearly separates different anatomical structures. However, we also note that this type of preprocessing can also produce undesirable artefacts such as altering the topology of the anatomical object. This is the case for one caudate in Figures 8(f) and (g), which ends up as an outlier that is clearly separated from the other data points. Considering this, the fact that WESD is able to produce visually pleasing embeddings without the need of preprocessing is an advantage.

2) *Analysing Heart Function in 4D MRI:* Four-dimensional imaging of patient anatomy is gaining interest in the medical community. The temporal analysis of anatomical structures is used to extract the characteristics of related dynamic processes, which often indicate certain pathologies [42]. Furthermore, in the recent work [43] authors show that shape information, in addition to volumetric measurements, improve the accuracy of pathology related classification tasks in such dynamical analyses. In this section, we apply nWESD to the shapes of the hearts extracted from four dimensional cardiac images of five different patients. The scan of each patient captures a full cycle of one heartbeat as a series of 20 3D images. Each image shows the left ventricle (LV) at a specific point in the cycle, from which we manually segment the corresponding blood pool. Our reference is the blood pool extracted from the first frame (diastole). We compute the nWESD scores between this reference and all other shapes extracted from the series of images. Here, we do not normalise the eigenvalues with respect to the global scale since size change is an important aspect of the heartbeat dynamics. The graph given in Figure 9 shows the results of these measurements over time across the five patients. The figure also shows some exemplary images and shapes. We observe that the symmetry of the heartbeat along the systolic (as the blood pumps out of the LV pool) and the diastolic phases (as the blood fills in the pool) is well captured with the nWESD score. Furthermore, the end-systolic phase (the time point with the largest distance w.r.t. the reference) is at different time points for different patients, which is to be expected since the different patient scans are not synchronized in time. In summary, WESD well captures the

⁶website:<http://www.loni.ucla.edu/Atlases/LPBA40>




Fig. 8: 2D embedding of subcortical structures: 240 structures (80 caudate nucleus, 80 putamen and 80 hippocampus) are extracted from MR scans of 40 different individuals. (a),(b) and (c) show some example structures from this dataset. Note the high intra-class variability and the artefacts due to finite resolution and manual segmentations. (d) and (e) plot 2D embeddings of these 240 structures obtained based on the affinity matrices computed via ρ_{SD}^N and WESD respectively. These embeddings are computed without any preprocessing applied to the structures. The embedding obtained with WESD distinctly clusters the objects with respect to the anatomical structures. The embedding in (d) however, shows some ambiguities in the separation. Graphs in (f) and (g) plot the similar embeddings obtained after smoothing the surfaces of the structures to remove artefacts. The embedding obtained by ρ_{SD}^N , although better than (d), still suffer from similar problems. The embedding based on WESD on the other hand, now even between better separates the groups.




Fig. 9: Analysing 3D + time (4D) cardiac images: First column shows corresponding 2D slices of a 4D MRI dataset at time points $t = \{0, 6, 12\}$. The second column, 3D shapes extracted at each of the time points. For five patients, we compute the nWESD shape dissimilarity score of the LV blood pool at each time point with respect to its shape at $t = 0$. The graph plots these scores. We note that the proposed shape distance is able to capture the dynamic process of the LV shape changes and furthermore, the symmetry between the two phases of a heart beat: diastole and systole.

dynamics of the beating heart, which is to be expected given the continuous link between the differences in eigenvalues and the difference in shape (see Section II).

V. CONCLUSION

This article proposed WESD, a new spectral shape distance defined over the eigenvalues of the Laplace operator. WESD is a theoretically sound shape metric that is derived from the heat-trace. The theoretical analysis given in this article

presented and proved the properties of WESD related to its existence, computability and multi-scale aspect. The presented experiments showed that the theoretical properties of WESD have many practical advantages over previous works. These experiments further highlighted that WESD is beneficial for various applications.

REFERENCES

- [1] D. Zhang and G. Lu, "Review of shape representation and description techniques," *Pattern Recognition*, vol. 37, no. 1, pp. 1–19, 2004.
- [2] N. Iyer, S. Jayanti, K. Lou, Y. Kalyanaraman, and K. Ramani, "Three-dimensional shape searching: state-of-the-art review and future trends," *Computer-Aided Design*, vol. 37, pp. 509–30, 2005.
- [3] M. Reuter, F.-E. Wolter, and N. Peinecke, "Laplace-Beltrami spectra as 'Shape-DNA' of surfaces and solids," *Computer-Aided Design*, vol. 38, pp. 342–66, 2006.
- [4] B. Lévy, "Laplace-beltrami eigenfunctions towards an algorithm that 'understands' geometry," in *Shape Modeling and Applications*, 2006.
- [5] R. M. Rustamov, "Laplace-beltrami eigenfunctions for deformation invariant shape representations," in *Eurographics Symposium on Geometry Processing*, 2007.
- [6] J. Sun, M. Ovsjanikov, and L. Guibas, "A concise and provably informative multi-scale signature based on heat diffusion," in *Eurographics Symposium on Geometry Processing*, 2009.
- [7] M. Reuter, F.-E. Wolter, M. Shenton, and M. Niethammer, "Laplace-beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis," *Computer-Aided Design*, vol. 41, pp. 739–55, 2009.
- [8] M. M. Bronstein and A. M. Bronstein, "Shape recognition with spectral distances," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 5, pp. 1065–71, 2011.
- [9] H. Weyl, "Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen," *Mathematische Annalen*, pp. 441–69, 1912.
- [10] M. Kac, "Can one hear the shape of a drum?" *The American Mathematical Monthly*, vol. 73, no. 7, pp. 1–23, 1966.
- [11] R. Courant and D. Hilbert, *Method of Mathematical Physics, vol I*. Interscience Publishers, 1966.
- [12] V. Jain and H. Zhang, "Robust 3d shape correspondence in the spectral domain," in *International Conference Shape Modeling and Applications*, 2006.

- [13] A. M. Bronstein, M. Bronstein, M. Mahmoudi, R. Kimmel, and G. Sapiro, "A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching," *International Journal of Computer Vision*, vol. 89, no. 2,3, pp. 266–286, 2010.
- [14] F. Mémoli, "A spectral notion of gromov-wasserstein distance and related methods," *Applied and Computational Harmonic Analysis*, vol. 30, no. 3, pp. 363–401, 2010.
- [15] M. Ovsjanikov, A. M. Bronstein, M. M. Bronstein, and L. J. Guibas, "Shape google: a computer vision approach to invariant shape retrieval," in *International Conference on Computer Vision Workshops*, 2009, pp. 320–327.
- [16] M. Bronstein, "Scale-invariant heat kernel signatures for non-rigid shape recognition," in *Computer Vision and Pattern Recognition*, 2010.
- [17] R. Lai, Y. Shi, K. Scheibel, S. Fears, R. Woods, A. Toga, and T. Chan, "Metric-induced optimal embedding for intrinsic 3d shape analysis," in *Computer Vision and Pattern Recognition*, 2010.
- [18] S. Gnutzmann, U. Smilansky, and N. Sondergaard, "Resolving isospectral "drums" by counting nodal domains," *Journal of Physics A: Mathematical and General*, vol. 38, no. 41, 2005.
- [19] S. Gnutzmann, P. Karageorge, and U. Smilansky, "Can one count the shape of a drum?" *Physical Review Letters*, vol. 97, no. 9, 2006.
- [20] R. Lai, Y. Shi, I. Dinov, T. Chan, and A. Toga, "Laplace-beltrami nodal counts: A new signature for 3d shape analysis," in *International Symposium on Biomedical Imaging*, 2009.
- [21] A. Pleijel, "A study of certain green's functions with applications in the theory of vibrating membranes," *Arkiv För Matematik*, vol. 2, no. 6, pp. 553–569, 1954.
- [22] H. McKean and I. Singer, "Curvature and the Eigenvalues of the Laplacian," *Journal of Differential Geometry*, vol. 1, pp. 43–69, 1967.
- [23] L. Smith, "The asymptotics of the heat equation for a boundary value problem," *Inventiones Mathematicae*, vol. 63, pp. 467–493, 1981.
- [24] M. Protter, "Can one hear the shape of a drum? Revisited," *SIAM Review*, vol. 29, no. 2, pp. 185 – 197, June 1987.
- [25] D. Vassilevich, "Heat kernel expansion: user's manual," *Physics Reports*, vol. 388, no. 5-6, pp. 279–360, 2003.
- [26] G. Jurman, R. Visintainer, and C. Furlanello, "An introduction to spectral distances in networks (extended version)," Preprint in ArXiv, Oct 2010.
- [27] Z. Lian, A. Godil, B. Bustos, M. Daoudi, J. Hermans, S. Kawamura, Y. Kurita, G. Lavoué, H. V. Nguyen, R. Ohbuchi, Y. Ohkita, Y. Ohishi, F. Porikli, M. Reuter, I. Sipiran, D. Smeets, P. Suetens, H. Tabia, and D. Vandermeulen, "SHREC'11 track: Shape retrieval on non-rigid 3D watertight meshes," in *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on 3D Object Retrieval*, 2011.
- [28] C. Gordon, D. Webb, and S. Wolpert, "Isospectral plane domains and surfaces via riemannian orbifolds," *Inventiones Mathematicae*, vol. 110, no. 1, pp. 1–22, 1992.
- [29] M. Niethammer, M. Reuter, F.-E. Wolter, S. Bouix, N. Peinecke, M.-S. Koo, and M. E. Shenton, "Global medical shape analysis using the Laplace-Beltrami spectrum," in *Medical Image Computing and Computer Assisted Intervention*, 2007.
- [30] E. Whittaker and G. Watson, *A course of modern analysis*. Cambridge Mathematical Library, 1996.
- [31] L. Evans, *Partial Differential Equations*. American Mathematical Society, 1998.
- [32] W. F. Ames, *Numerical Methods for Partial Differential Equations*. Academic Press, 1977.
- [33] W. E. Arnoldi, "The principle of minimized iterations in the solution of the matrix eigenvalue problem," *Quarterly of Applied Mathematics*, vol. 9, no. 17, pp. 17–29, 1954.
- [34] A. Bronstein, A. B. M.M. Bronstein, and R. Kimmel, "Analysis of two-dimensional non-rigid shapes," *International Journal of Computer Vision*, vol. 78, no. 1, pp. 67–77, 2008.
- [35] A. Bronstein, M. Bronstein, and R. Kimmel, *Numerical geometry of non-rigid shapes*. Springer, 2008.
- [36] J. Tenenbaum, V. de Silva, and J. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–23, 2000.
- [37] Z. Lian, A. Godil, X. Sun, and H. Zhang, "Non-rigid 3d shape retrieval using multidimensional scaling and bag-of-features," in *Proceedings of ICIP*, 2010, pp. 3181–84.
- [38] C. Maes, T. Fabry, J. Keustermans, D. Smeets, P. Suetens, and D. Vandermeulen, "Feature detection on 3D face surfaces for pose normalisation and recognition," in *Proceedings of Biometrics: Theory, Applications and Systems*, 2010.
- [39] D. Smeets, T. Fabry, J. Hermans, D. Vandermeulen, and P. Suetens, "Isometric deformation modelling for object recognition," in *Proceedings of Computer Analysis of Images and Patterns*, 2009, pp. 757–65.
- [40] F. L. Bookstein, P. D. Sampson, A. P. Streissguth, and P. D. Connor, "Geometric morphometrics of corpus callosum and subcortical structures in the fetal-alcohol-affected brain," *Teratology*, vol. 64, pp. 4–32, Jul 2001.
- [41] D. Shattuck, M. Mirza, V. Adisetiyo, C. Hojatkashani, G. Salamon, K. Narr, R. Poldrack, R. Bilder, and A. Toga, "Construction of a 3d probabilistic atlas of human cortical structures," *NeuroImage*, vol. 39, no. 3, pp. 1064–80, 2008.
- [42] T. Mansi, I. Voigt, B. Leonardi, X. Pennec, S. Durrleman, M. Sermesant, H. Delingette, A. Taylor, Y. Boudjemline, G. Pongiglione, and N. Ayache, "A statistical model for quantification and prediction of cardiac remodeling: application to tetralogy of fallot," *IEEE Transactions on Medical Imaging*, vol. 30, no. 9, pp. 1605–16, 2011.
- [43] E. Bernardis, E. Konukoglu, Y. Ou, D. N. Metaxas, B. Desjardins, and K. M. Pohl, "Temporal shape analysis via the spectral signature," in *In Proceedings of Medical Image Computing and Computer Assisted Intervention*, 2012.



Ender Konukoglu was born in Istanbul, Turkey, in 1981. He received a Ph.D. degree in computer science specializing in medical image analysis from Université de Nice and INRIA Sophia Antipolis, France. He is now a Research Fellow at Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital / Harvard Medical School. His research interests are medical image analysis, biophysical models, machine learning and applications of partial differential equations.



Ben Glocker was born in Goettingen, Germany, in 1980. He received a Ph.D. degree in medical computer science from Technische Universitaet Muenchen, Germany. He is currently a Postdoctoral Researcher at Microsoft Research Cambridge, UK where he is a member of the Machine Learning and Perception group. His main research focus is on the application of machine learning techniques for medical image analysis.



Antonio Criminisi was born in 1972 in Italy. He received a Degree in Electronics Engineering at the University of Palermo and obtained a PhD in Computer Vision at the University of Oxford. He is currently a Senior Researcher at Microsoft Research Cambridge, UK. His research interests are in the area of medical image analysis, object category recognition, image and video analysis and editing, one-to-one teleconferencing, 3D reconstruction from single and multiple images with application to virtual reality, forensic science and history of art.



Kilian M. Pohl received a M.S. degree from the Department of Mathematics, University of Karlsruhe, and his PhD from the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology. Dr. Pohl is now an Assistant Professor at the Department of Radiology and holds a secondary appointment at the Bioengineering Graduate Group, University of Pennsylvania. His research focuses on creating algorithms for automatically quantifying and generalizing the information latent in medical images.