



3D Shape Analysis Through a Quantum Lens: the Average Mixing Kernel Signature

Luca Cosmo^{1,2,3} · Giorgia Minello³ · Michael Bronstein^{2,4,5} · Emanuele Rodolà¹ · Luca Rossi⁶ · Andrea Torsello³

Received: 16 February 2021 / Accepted: 16 March 2022 / Published online: 11 April 2022
© The Author(s) 2022

Abstract

The Average Mixing Kernel Signature is a novel spectral signature for points on non-rigid three-dimensional shapes. It is based on a quantum exploration process of the shape surface, where the average transition probabilities between the points of the shape are summarised in the finite-time average mixing kernel. A band-filtered spectral analysis of this kernel then yields the AMKS. Crucially, we show that opting for a finite time-evolution allows the signature to account for a mixing of the Laplacian eigenspaces, similar to what is observed in the presence of noise, explaining the increased noise robustness of this signature when compared to alternative signatures. We perform an extensive experimental analysis of the AMKS under a wide range of problem scenarios, evaluating the performance of our descriptor under different sources of noise (vertex jitter and topological), shape representations (mesh and point clouds), as well as when only a partial view of the shape is available. Our experiments show that the AMKS consistently outperforms two of the most widely used spectral signatures, the Heat Kernel Signature and the Wave Kernel Signature, and suggest that the AMKS should be the signature of choice for various compute vision problems, including as input of deep convolutional architectures for shape analysis.

Keywords Shape representation · Shape analysis · Partial Matching · Quantum walks

1 Introduction

Shape descriptors are used in a variety of shape analysis tasks (Aubry et al. 2011a; Gasparetto et al. 2015; Cosmo et al. 2016a, 2017; Rodolà et al. 2017a,b; Vestner et al. 2017; Huang et al. 2018), from surface registration to shape retrieval, and as such they have long been the subject of intense research in the computer vision community. A good descriptor (also known as signature) is able to effectively characterize the local and global geometric information around each point on the surface of a shape while remaining

invariant to rigid transformations, stable against non-rigid ones, and robust to various sources of noise. However, the variety and high complexity of 3D shape representations, as well as the type of transformations they can undergo, make the development of such a descriptor particularly challenging. A further issue is that of potential incompleteness of the shape as a result of occlusions or partial view during the shape acquisition process (Rodolà et al. 2017a). In this case, an ideal signature would still be able to capture the local shape topology while accounting for the partiality induced by the missing parts.

Traditionally, the Heat Kernel Signature (HKS) (Sun et al. 2009) and the Wave Kernel Signature (WKS) (Aubry et al. 2011b) are two of the most widely used descriptors for shape analysis. Both signatures attempt to capture the local shape topology by analysing the eigenvalues and eigenfunctions of the Laplace–Beltrami operator of a shape. They are part of a broader family of descriptors known as spectral descriptors. From the perspective of physics, the HKS characterizes the points of a shape by the way heat dissipates over its surface. The WKS, in contrast, measures the probability of a quantum particle with a certain energy distribution to be located at a given point on the shape. Compared to the HKS, the

Communicated by Adrien Bartoli.

✉ Luca Rossi
luca.rossi@qmul.ac.uk

¹ Sapienza University of Rome, Rome, Italy

² University of Lugano, Lugano, Switzerland

³ Ca' Foscari University of Venice, Venezia, Italy

⁴ Twitter, London, UK

⁵ Imperial College London, London, UK

⁶ Queen Mary University of London, London, UK

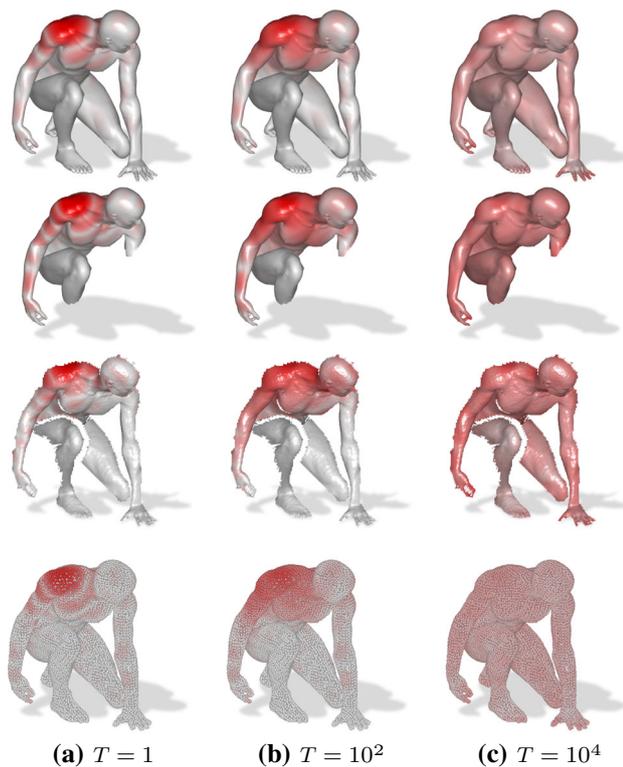


Fig. 1 The average transition probability of a quantum walk starting from one point and spreading to the rest of the shape for increasing values of the stopping time T (left to right), for a full shape (top row), partial shape (second row), range maps (third row), and point cloud (bottom row). More (less) intense shades of red correspond to higher (lower) probability. This is the information contained in the row of the average mixing matrix (Eq. 18) indexed by the starting point. Note the presence of grey bands when $T = 1$ (corresponding to points with zero probability of being visited) which are a result of the interference effects typical of quantum walks

WKS has been shown to exhibit superior feature localization. Indeed, by exploiting interference effects, quantum processes can detect the presence of new structures when compared to classical diffusion processes (Rossi et al. 2012, 2013; Bai et al. 2015; Rossi et al. 2015; Minello et al. 2019).

The Average Mixing Kernel Signature (AMKS) (Rossi et al. 2016; Cosmo et al. 2020) bears similarities with both the HKS and the WKS, as it is based on a quantum mechanical analogue of the heat kernel, known as the average mixing kernel (Cosmo et al. 2020). Like the WKS, the AMKS is based on the idea of letting a quantum system evolve according to the Schrödinger equation, and band-filters are used to compute the signature for a selected set of energies. However, unlike the WKS, the AMKS is not based on the infinite-time average of a single particle evolution, and the band filter is not determined by the initial state. Crucially, the AMKS naturally accounts for the noise-induced mixing of the Laplacian eigenspaces, and it outperforms both the HKS and the WKS in a shape matching task.

In this paper we go beyond the seminal work of Cosmo et al. (2020) and we present a thorough discussion and analysis of the properties and performance of the AMKS. We start by discussing the quantum nature of this signature and we propose a simple yet effective way to modify it in order to cope with shape partiality. We perform an extensive analysis of this signature under a variety of previously untested problem scenarios, ranging from the presence of topological noise to the use of alternative shape representations (point clouds). Our experiments show that the AMKS consistently outperforms both the HKS and the WKS on all these problems, with nearly a 100% performance improvement with respect to these signatures when applied on partial shapes.

The remainder of this paper is organised as follows. Section 2 provides a brief overview of the related work. Section 3 discusses the quantum mechanical background necessary to then introduce the proposed signature in Sect. 4. Section 5 explains the relation between the AMKS and the noise-induced mixing of the Laplacian eigenspaces, while Sect. 6 discusses the computation of the signature. In Sect. 7 we show how the AMKS can be adapted to work on partial shapes, and in Sect. 8 we perform an extensive set of experiments to test the AMKS in a wide range of problem scenarios. Finally, Sect. 9 concludes the paper.

2 Related Work

We divide existing contributions in two main categories: (1) spectral-based descriptors, to which the proposed signature belongs, and (2) learning-based signatures, which have been gaining increasing popularity. While presenting a comprehensive review of spectral and learning-based signatures is beyond the scope of this paper, we refer the reader to the excellent surveys of Masci et al. (2015) and Rostami et al. (2019) for a more in-depth discussion of both categories of signatures.

2.1 Spectral-Based Signatures

As the name suggests, spectral-based signatures attempt to capture the topology of a shape by looking at the eigenfunctions of mesh operators such as the discrete Laplace–Beltrami operator. These signatures can be further separated into two main categories, local and global. Local descriptors are defined on each point of a shape and try to characterise the local structure around these. Global descriptors, on the other hand, are defined on the entire shape and can be often computed by aggregating local descriptors over it.

Shape-DNA (Reuter et al. 2006) is one of the simplest and first examples of spectral descriptor and it corresponds to the truncated sequence of the Laplacian eigenvalues, sorted in

ascending order. This is a global descriptor and, as such, cannot be used for local or partial shape analysis. Global Point Signature (GPS) (Rustamov 2007) is defined as a vector of scaled Laplace–Beltrami eigenfunctions evaluated at a given point on the shape surface. Unfortunately, it has been shown that eigenfunctions corresponding to pairs of eigenvalues with close values can switch, meaning that the GPS of points across different shapes is not easily comparable.

The Heat Kernel Signature (HKS) (Sun et al. 2009) overcomes this problem by taking an exponentially weighted sum of the squares of the eigenfunctions. The underpinning idea is to describe each point of a shape with the amount of heat retained over time. More precisely, the HKS can be seen as performing a low-pass filtering of the spectral information parametrized by the diffusion time t . The result is a signature that is invariant to isometric transformations, robust to small perturbations, and where the HKS of points on different shapes are commensurable.

The HKS is dominated by low frequency information, which corresponds to macroscopic properties of the shape. This means missing out on highly discriminative small-scale information that can be crucial in matching tasks. The Wave Kernel Signature (WKS) (Aubry et al. 2011b) solves this problem by replacing the low-pass filter with a band-pass filter, that better separates the frequency information related to different spatial scales. Stated otherwise, the WKS is parametrised using frequency rather than time (as in the HKS), allowing it to access and control high frequency information as well.

2.2 Learning-Based Signatures

While spectral signatures remain widely used in the shape analysis community due to their efficacy and ease of computation, researchers have recently started turning their attention to learning-based or data-driven descriptors.

Litman and Bronstein (2013) introduced a family of spectral descriptors that generalizes the HKS and the WKS, and proposed a learning scheme to construct the optimal descriptor for a given task. Corman et al. (2014) learn optimal descriptors from a given set of shape correspondences. In particular, they use a functional maps representation where spectral signatures (e.g., HKS or WKS) are used as probe functions that are meant to constrain the degree of deformation between corresponding points.

Fang et al. (2015) use a deep autoencoder to construct a global descriptor for shape retrieval. In this framework, one of the first processing steps is the extraction of the HKS, which is then fed to the autoencoder itself. Masci et al. (2015) extend convolutional neural networks (CNNs) to non-Euclidean manifolds, where a local geodesic patch is defined around each point and a filter is convolved with vertex-level functions (e.g., HKS and WKS). Huang et al. (2018) omit

this preprocessing step by designing a 2D CNN architecture where each point on the shape surface is rendered using multi-scale 2D images.

Finally, while the success of CNNs in computer vision has certainly shifted the interest of the community toward learning-based descriptors, it should be noted that spectral-based signatures are still often used as the first processing step in a modern (deep) learning framework (Huang et al. 2018; Rostami et al. 2019).

3 Quantum Mechanical Background

In this section we introduce the quantum mechanical concepts underpinning our signature, i.e., continuous-time quantum walks and the average mixing kernel.

3.1 Quantum States

Quantum mechanics differs from classical mechanics in that a system can be in a superposition of states, which is represented through a complex linear combination of the classical states of the system. For example, let $S = \{s_1, \dots, s_n\}$ be a (discrete) domain where s_i are the possible classical states, the quantum state over the same domain is represented by a vector of *amplitudes* $a \in \mathbb{C}^n$, such that $a^\dagger a = 1$ and $|a_i|^2$ represents the probability of observing the system in the state s_i . This can be generalized to a continuous domain by substituting a with a function $f : S \rightarrow \mathbb{C}$ such that

$$\int_S f(x)^* f(x) dx = 1. \quad (1)$$

The Dirac notation (also known as *bra-ket* notation) is a standard notation used for describing quantum states. Here a *ket* vector $|a\rangle$ denotes a pure quantum state and is a complex valued column vector (function) of unit Euclidean length, in a Hilbert space. Its conjugate transpose is a *bra* (row) vector, denoted as $\langle a|$. As a result, the inner product between two states $|a\rangle$ and $|b\rangle$ is written $\langle a|b\rangle$, while their outer product is $|a\rangle \langle b|$.

3.2 Quantum Measurements

Another notable difference between classical and quantum systems is the observation process.

First, the observation is statistical in nature, with the wave function, which represents all the information we can have about a quantum system, dictating the probabilities of observing a particular value from a measurement on the system. Further, the statistical nature of the system cannot be reduced to an unknown hidden value, as quantum systems do not satisfy Bell's inequality (Bell 1987), an inequality satisfied by all

classical hidden-value systems. Second, observations change the system state by collapsing the wavefunction on the states compatible with the observed value.

Consider a quantum system in a state $|\psi\rangle$. While in classical mechanics an observable (i.e., a measurable physical quantity) is a real-valued function of the system state, in quantum mechanics an observable is a self-adjoint operator acting on $|\psi\rangle$, such that the possible outcomes of an observation correspond to the eigenvalues of the associated self-adjoint operator, while the eigenvectors (or eigenfunctions if the domain is continuous) form a basis of the states consistent with each outcome.

In this context, observing a state corresponds to performing a projective measurement of $|\psi\rangle$, where the wave function collapse results in an orthogonal projection on the space spanned by the eigenvectors corresponding to the eigenvalue associated with the outcome. Given the observable O , we consider its spectral decomposition $O = \sum_{\lambda} \lambda P_{\lambda}$, where

$$P_{\lambda} = \Phi_{\lambda} \Phi_{\lambda}^{\dagger} = \sum_i \phi_{\lambda i} (\phi_{\lambda i})^{\dagger} \tag{2}$$

is the projector on the subspace spanned by the eigenvectors associated with eigenvalue λ , and $\phi_{\lambda i}$ refers to the i -th eigenvector associated with eigenvalue λ . Then, the outcome λ is observed with probability $p(\lambda) = \langle \psi | P_{\lambda} | \psi \rangle$ and after such observation, the state of the quantum system changes to

$$|\bar{\psi}\rangle = \frac{P_{\lambda} |\psi\rangle}{\|P_{\lambda} |\psi\rangle\|}, \tag{3}$$

where $\| |\psi\rangle \| = \sqrt{\langle \psi | \psi \rangle}$ denotes the norm of $|\psi\rangle$.

3.3 The Schrödinger Equation

We adopt the Schrödinger picture of the dynamics of a quantum system, in which the measurement operators are constant and the wave function varies in time. Here, time is a parameter external to the system and thus this is a non-relativistic model that works in the low-speed limit. According to this picture, in absence of a measurement, the wave function evolves according to the well-known Schrödinger equation

$$\frac{d}{dt} |\psi\rangle = -\frac{i}{\hbar} \mathcal{H} |\psi\rangle. \tag{4}$$

The solution to this equation is given by

$$|\psi(t)\rangle = U(t) |\psi(0)\rangle, \tag{5}$$

where $U(t) = \exp(-\frac{i}{\hbar} t \mathcal{H})$ is a unitary operator. Note that $U(t)$ shares the same eigenvectors of \mathcal{H} , and for each eigenvalue λ of \mathcal{H} there exists a corresponding eigenvalue $e^{-\frac{i}{\hbar} t \lambda}$ of $U(t)$.

3.4 From Classical to Quantum Walk and Diffusion

Let $G = (V, E)$ denote an undirected graph with vertex set V and edge set $E \subseteq V \times V$. Classical and quantum walks on G are both defined as vertex visiting processes where the movements of the walker are constrained by the edge structure of the graph. Crucially, despite some apparent similarities in their definition, the dynamics of continuous-time quantum walks and continuous-time random walks are fundamentally different (Kempe 2003).

Looking at the continuous domain counterparts, given a manifold \mathcal{M} , a diffusion process governed by the heat equation determines how a classical distribution evolves over \mathcal{M} , while a quantum evolution evolves the wave function over the same domain. Again, the two processes share some apparent similarities, but provide fundamentally different dynamics.

3.4.1 Classical Random Walks and the Heat Equation

In the classical case, the state of the walk is described by a time-varying probability distribution, representing the probability of locating the walker at the various vertices of the graph over time. In the case of a continuous diffusion process, the distribution is substituted by a probability density function.

Let \mathcal{M} denote a manifold, the dynamics of the diffusion process follow the heat equation

$$\frac{\partial}{\partial t} f(t) = -\Delta_{\mathcal{M}} f(t), \tag{6}$$

where $f(t) : \mathcal{M} \rightarrow \mathbb{R}$ denotes the heat distribution at time t and $\Delta_{\mathcal{M}}$ is the Laplace–Beltrami operator associated with \mathcal{M} .

When the graph is obtained by discretizing an underlying continuous domain, as in the case of the mesh representation of shapes, and a suitable discretization of the Laplace–Beltrami operator is used, the continuous-time random walk corresponds to an approximation of the heat diffusion process on the shape surface, with the only difference that the continuous-domain heat distribution $f(t)$ is replaced by the discrete state vector $\mathbf{p}(t) \in \mathbb{R}^n$, with n denoting the number of vertices of the graph. For this reason, we will analyze them together and use the formalism interchangeably throughout, with the only caveat that the discretization induces a corresponding dot-product over the states taking into account the variation in sizes of the discretization domains.

Given an initial heat distribution $f(0)$ (or, similarly, an initial state vector $\mathbf{p}(0)$), the heat kernel

$$H(t) = \exp(-t \Delta_{\mathcal{M}}) \tag{7}$$

allows one to determine the heat distribution at time t , i.e.,

$$f(t) = H(t)f(0). \tag{8}$$

Assuming at time $t = 0$ all the heat is concentrated on a single point v of \mathcal{M} , the heat operator kernel (Sun et al. 2009) $H(t)$ determines the amount of heat transferred from v to the other points of \mathcal{M} during the time t . Similarly, in the discrete case of a graph, $H(t)$ is a doubly-stochastic matrix and its entries $H(t) = (h_{u,v})$ measure the probability of a transition from v to u in t units of time.

Most importantly, from $H(t) = e^{-t\Delta_{\mathcal{M}}}$ we see that the heat operator and the Laplace–Beltrami operator share the same eigenfunctions and if λ is an eigenvalue of $\Delta_{\mathcal{M}}$ then $e^{-\lambda t}$ is an eigenvalue of $H(t)$, allowing one to compute the amount of heat (or probability, in the discrete case of random walks) transferred from one point of the surface to the other as

$$k_{uv} = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(u)\phi_i(v), \tag{9}$$

where λ_i and ϕ_i denote the i -th eigenvalue and eigenfunction of $\Delta_{\mathcal{M}}$, respectively.

From this we can clearly see that the process has stationary distributions on the kernel of $\Delta_{\mathcal{M}}$, i.e., the subspace spanned by the eigenfunctions of $\Delta_{\mathcal{M}}$ corresponding to the zero eigenvalue.

3.4.2 Quantum Walks and the Schrödinger Equation

Moving to the quantum realm, continuous-time quantum walks bear many similarities with their classical counterpart, yet they behave profoundly differently. The probability density function $f(t)$ is replaced by the complex valued wave function of the quantum particle moving freely over the graph or the manifold. An important consequence of this is that the state of the walker does not lie in a probability space, but on a complex hyper-sphere, which in turns allows interference to take place (Portugal 2013).

Let \mathcal{M} denote the manifold representing the shape surface. The state of a quantum particle at time t is given by the wave function $|\psi(t)\rangle$, defined as an element of the Hilbert space of functions mapping points in \mathcal{M} to \mathbb{C} , such that $\langle \psi(t) | \psi(t) \rangle = 1$ for all times t .

In the discrete domain, the state of the walker at time t is denoted as

$$|\psi(t)\rangle = \sum_{u \in V} \alpha_u(t) |u\rangle, \tag{10}$$

where $|u\rangle$ denotes the state corresponding to the vertex u and $\alpha_u(t) \in \mathbb{C}$ is the corresponding amplitude. Moreover, we

have that $\alpha_u(t)\alpha_u^*(t)$ gives the probability that at time t the walker is at the vertex u , and thus $\sum_{u \in V} \alpha_u(t)\alpha_u^*(t) = 1$ and $\alpha_u(t)\alpha_u^*(t) \in [0, 1]$, for all $u \in V, t \in \mathbb{R}^+$.

The evolution of the walker is governed by the Schrödinger equation

$$\frac{\partial}{\partial t} |\psi(t)\rangle = -i\Delta_{\mathcal{M}} |\psi(t)\rangle, \tag{11}$$

where the Laplace–Beltrami operator plays the role of a time-independent Hamiltonian. Note that for simplicity here we assume units that guarantee $\hbar = 1$.

The solution to this equation, as we have seen before, is given by

$$|\psi(t)\rangle = \exp(-it\Delta_{\mathcal{M}}) |\psi(0)\rangle. \tag{12}$$

The dynamics of the classical and quantum processes described above, albeit apparently similar in the definition, are fundamentally different. While in the classical case heat dissipates with time and the walker reaches a steady state, in the quantum case the wave function propagates through the domain and the walker state keeps oscillating.

3.5 Average Mixing Matrix

Godsil (2013) introduces the concept of mixing matrix $M_G(t) = (m_{uv})$ as the probability that a quantum walker starting at node v is observed at node u at time t . Given the eigenvalues λ of the Hamiltonian and the associated projectors P_λ , the unitary operator inducing the quantum walk can be rewritten as

$$U(t) = \sum_{\lambda \in \mathfrak{3}} e^{-i\lambda t} P_\lambda, \tag{13}$$

where $\mathfrak{3}$ is the set of unique eigenvalues of the Hamiltonian.

Given Eq. 13, we can write the mixing kernel for the manifold \mathcal{M} in terms of the Schur-Hadamard product (also known as entry-wise product and denoted by the symbol \circ) of the projectors

$$M_{\mathcal{M}}(t) = \sum_{\lambda_1 \in \mathfrak{3}} \sum_{\lambda_2 \in \mathfrak{3}} e^{-i(\lambda_1 - \lambda_2)t} P_{\lambda_1} \circ P_{\lambda_2}. \tag{14}$$

Recall that, while in the classical case the probability distribution induced by a random walk converges to a steady state, this does not happen in the quantum case. However, we will show that we can enforce convergence by taking a time-average even if $U(t)$ is norm-preserving. Let us define the average mixing kernel (Godsil 2013) at time T as

$$\begin{aligned} \widehat{M}_{\mathcal{M};T} &= \frac{1}{T} \int_0^T M_{\mathcal{M}}(t) dt \\ &= \sum_{\lambda_1 \in \mathcal{E}} \sum_{\lambda_2 \in \mathcal{E}} P_{\lambda_1} \circ P_{\lambda_2} \frac{1}{T} \int_0^T e^{-i(\lambda_1 - \lambda_2)t} dt, \end{aligned} \tag{15}$$

which has solution

$$\widehat{M}_{\mathcal{M};T} = \sum_{\lambda_1 \in \mathcal{E}} \sum_{\lambda_2 \in \mathcal{E}} P_{\lambda_1} \circ P_{\lambda_2} \frac{i(1 - e^{iT(\lambda_2 - \lambda_1)})}{T(\lambda_2 - \lambda_1)}. \tag{16}$$

In the limit $T \rightarrow \infty$, Eq. 16 becomes

$$\widehat{M}_{\mathcal{M};\infty} = \sum_{\lambda \in \mathcal{E}} P_{\lambda} \circ P_{\lambda}. \tag{17}$$

Finally, we note that the average mixing kernel can be rewritten as follows, i.e.,

$$\begin{aligned} \widehat{M}_{\mathcal{M};T} &= \sum_{\lambda_1 \in \mathcal{E}} \sum_{\lambda_2 \in \mathcal{E}} P_{\lambda_1} \circ P_{\lambda_2} \frac{i(1 - e^{iT(\lambda_2 - \lambda_1)})}{T(\lambda_2 - \lambda_1)} \\ &= \sum_{\lambda_1 \in \mathcal{E}} \sum_{\lambda_2 \in \mathcal{E}} P_{\lambda_1} \circ P_{\lambda_2} \frac{\sin(T(\lambda_2 - \lambda_1))}{T(\lambda_2 - \lambda_1)} \\ &= \sum_{\lambda_1 \in \mathcal{E}} \sum_{\lambda_2 \in \mathcal{E}} P_{\lambda_1} \circ P_{\lambda_2} \text{sinc}(T(\lambda_2 - \lambda_1)), \end{aligned} \tag{18}$$

where $\text{sinc}(x) = \frac{\sin(x)}{x}$ is the unnormalized sinc function. In Sect. 4 we will show that the noise robustness of our signature is closely related to the rate of mixing of the eigenspaces given by the sinc component when we take a finite time average as in Eq. 18, instead of an infinite time average as in Eq. 17. In Fig. 1 we visualize the entries of the average mixing matrix computed according to Eq. 18 on a mesh of n vertices sampled from an underlying shape, for different choices of the stopping time T .

4 Average Mixing Kernel Signature

The average mixing matrix provides observation probabilities on node j for a particle starting at node i . In doing so, it integrates contributions at all energy levels. In order to construct a signature, we decompose the spectral contributions to the matrix into several bands so as to both limit the influence of the noise-dominated high-frequency components, and create a more descriptive signature as a function of energy level, as first proposed by Aubry et al. (2011b). We do this selectively reducing the spectral components at frequency λ according to an amplitude $f_E(\lambda)$ for a given target

energy level E , obtaining the band-filtered average mixing matrix

$$\text{AMM}(E) = \sum_{\lambda_1, \lambda_2} P_{\lambda_1} \circ P_{\lambda_2} \text{sinc}(T(\lambda_2 - \lambda_1)) f_E(\lambda_1) f_E(\lambda_2), \tag{19}$$

where $f_E(\lambda) = e^{-\frac{(E - \log \lambda)^2}{2\sigma^2}}$, with E being a band parameter as in the WKS.

Quantum-mechanically, this is equivalent to the introduction of a filter or an annihilation process that eliminates from the ensemble walkers at energy-eigenstate μ with rate $1 - f_E(\mu)$, just before the final observation on the vertex basis $|i\rangle$. This is equivalent to an observation on the states images of the projectors P_{μ} and $I - P_{\mu}$, with $\frac{1 - f_E(\mu)}{\langle j | P_{\mu} | j \rangle}$ being the probability that the particle starting at node j is discarded from the ensemble if observed in P_{μ} .

The ensemble at time t of the particle starting at node j after the annihilation is

$$\begin{aligned} &\sum_{\lambda_1, \lambda_2} e^{-i(\lambda_1 - \lambda_2)t} \left(I - \sum_{\mu_1} (1 - f_E(\mu_1)) P_{\mu_1} \right) \left(P_{\lambda_1} |j\rangle \langle j| P_{\lambda_2} \right) \\ &\quad \left(I - \sum_{\mu_2} (1 - f_E(\mu_2)) P_{\mu_2} \right) \\ &= \sum_{\lambda_1, \lambda_2} e^{-i(\lambda_1 - \lambda_2)t} f_E(\lambda_1) f_E(\lambda_2) \left(P_{\lambda_1} |j\rangle \langle j| P_{\lambda_2} \right), \end{aligned} \tag{20}$$

resulting, after time integration, in the band-filtered average mixing matrix (Eq. 19). To obtain the descriptor, we take the self-propagation, that is the probability that a particle in state u is still observed in u after the evolution, as a function of the energy level E , normalized so that the integral over all energies is 1, i.e.,

$$\text{AMKS}_u(E) = \frac{\text{AMM}(E)_{uu}}{\sum_{\lambda_1, \lambda_2} f_E(\lambda_1) f_E(\lambda_2)}. \tag{21}$$

4.1 Relationship with WKS

Our construction bears clear similarities with the WKS, namely the fact that it is based on the evolution of a quantum system under the Schrödinger equation, the time average is taken to move towards a steady-state distribution, and the fact that we adopt the same energy band-filtering over which the signature is defined.

However, there are some stark differences as well. First, band-filtering is not determined by an initial state, like in WKS, but it is posed as a pre-measurement process on the particle ensemble, similar to adding filters in front of a light detector. This is not only cleaner, but physically much easier to implement. A consequence is that we map initial to final states, thus defining a kernel as in HKS, in contrast with WKS

that maps energy bands to point distributions. This potentially offers better localization and opens the possibility to a more descriptive signature.

Mathematically, the way the kernel is defined is different also in the infinite-time limit: while WKS is fundamentally a function of the square of the entries of the Laplacian eigenvectors, our signature results in their *fourth* power. This, in turn, creates a more distinctive descriptor while at the same time flattening small and possibly noise-induced variations. Finally, instead of taking the infinite-time average of the evolution of a single particle, we look at the statistical behavior of the finite time average of an ensemble of particles. This increases the robustness to noise, as it imposes a mixing of the eigenspaces at a rate similar to the effect of noise.

5 Perturbation Analysis

In order to understand the advantage of the use of finite-time averages, we perform a perturbation analysis on the eigen-decomposition of the Laplace–Beltrami operator, showing that first order perturbation introduces a mixing of the eigenspaces similar to what is obtained in the finite-time averaging of the ensemble.

Recall that the stability analysis presented by Aubry et al. (2011b) performs a perturbation analysis on the eigenvalues of the Laplace–Beltrami operator, showing that the ϵ -perturbation $E_k(\epsilon)$ of the energy level (eigenvalue) E_k satisfies

$$\log\left(\frac{E_k(\epsilon)}{E_k}\right) = \epsilon c_k + O(\epsilon^2) \quad (22)$$

where c_k is a constant. This is what justifies the use of the lognormal kernel in the definition of the WKS. It is worth noting that the result does not depend on the signature, but simply quantifies the perturbative variability/indetermination of the energy levels of the Laplace–Beltrami operator. As such, the analysis can be taken as is for any spectral signature gated on the energy distribution. For this reason we are using the same lognormal kernel for our filters, thus ensuring that their analysis still holds for our signature.

Here, however, we want to go beyond the perturbative analysis of the spectrum alone, analyzing the perturbation of the eigenvectors as well, and showing that the effect of a perturbation on the Laplace–Beltrami operator causes a mixing of the eigenspaces that is dampened by a function of the inverse difference between the corresponding eigenvalues. More formally, the eigenvector ϕ_i associated with the eigenvalue λ_i receives a component from the eigenvector ϕ_j that is proportional to $1/(\lambda_i - \lambda_j)$. Note that asymptotically this is the same behavior of the sinc function in our signature, which can then be seen as one sample from the same distri-

bution over the eigenspace of the Laplace–Beltrami operator. Note that, just as in the case of the analysis in Aubry et al. (2011b), the result does not depend on the actual signature produced, but is true for any spectral signature that mixes the eigenvector with a weight that is asymptotically equivalent to $1/(\lambda_i - \lambda_j)$.

Let L be the Laplace–Beltrami operator and assume that we observe a noisy version $\hat{L} = L + \mathcal{E}$ with \mathcal{E} being a suitable additive noise on the operator, and from that an interpolating function $\mathcal{L}(t) = L + t\mathcal{E}$ with $t \in [0, 1]$. Clearly, as $t = 0$ we obtain the noiseless Laplace–Beltrami operator $\mathcal{L}(0) = L$, while for $t = 1$ we have the noisy operator $\mathcal{L}(1) = \hat{L}$. Further, let $\mathcal{L}\Phi = \Phi\Lambda$ be the eigenvalue equation for \mathcal{L} , where Λ is the diagonal matrix of eigenvalues such that $(\Lambda)_{ii} = \lambda_i$, while Φ is the matrix of eigenvectors, so that $\Phi_{\cdot i} = \phi_i$.

Following (Murthy and Haftka 1988), we can write the derivatives at $t = 0$ of the eigenvectors ϕ_i of $\mathcal{L}(t)$ introducing a matrix $B = (b_{ij})$ residing in the tangent space at unity of the group $O(n)$ of orthogonal transformations such that $\Phi' = \Phi B$, where n is the number of vertices of the mesh. Under this representation, we can write the eigenvectors of \hat{L} as a first order approximation of the expansion of \mathcal{L} at 0 through the exponential map of B , obtaining

$$\Phi_{\hat{L}} \approx \Phi \exp(B) \approx \Phi(I + B), \quad (23)$$

where $\Phi_{\hat{L}}$ is the eigenvector matrix of \hat{L} .

Moreover, noting that both L and \mathcal{E} are symmetric, we can solve the perturbation equations for a self-adjoint operator obtaining the following:

$$b_{ij} = \begin{cases} 0 & \text{if } \lambda_i = \lambda_j; \\ \frac{\phi_i^T \mathcal{E} \phi_j}{\lambda_j - \lambda_i} & \text{otherwise.} \end{cases} \quad (24)$$

From this perturbation analysis we see that, at least to the first order, the effect of noise is a mixing into the eigenspace related to eigenvalue λ_i of a component linked to the eigenspace of λ_j at a rate proportional to the reciprocal of the difference in eigenvalues $\frac{1}{\lambda_j - \lambda_i}$, rate that is asymptotically equivalent to the sinc($T(\lambda_j - \lambda_i)$) obtained through finite-time averaging, for T approaching infinity. Hence, we can see the additive mixing of the finite-time averaging to be equivalent to a random projection onto the same portion of subspace induced by the noise process, or equivalently, we can see the additive mixing noise process as a random projection onto the subspace spanned by the additional components in the finite-time averaging.

The end result is a descriptor that already includes and takes into account part of the deformation introduced by the noise process.

5.1 Laplacian Noise

We have seen that adding noise to the Laplacian introduces a mixing of the eigenvectors that decays linearly with the distance between the corresponding eigenvalues, however it is hard to analyze directly the noise in the Laplacian. Here we would like to experimentally characterize the Laplacian noise in terms of noise that is more easily encountered in real world applications: vertex jitter. In particular, we want to show how Gaussian noise on the location of the vertices in a mesh maps to noise on the Laplacian. There are two processes to take into consideration in the mapping from vertices to entries of the Laplacian and that can alter the distribution: the first is the local alteration of the metric, while the second is the discretization of the Laplacian entries, which are non-zero only in correspondence of adjacent vertices in the mesh, at least using the widely adopted first-order FEM discretization (cotangent rule) to create the mesh Laplacian. Both have the possibility of introducing correlations between the entries and severely skewing the distribution.

In Fig. 2 we show a histogram of the difference in the non-zero entries of the Laplacian computed from the original meshes and from the meshes with added Gaussian vertex noise with zero mean and standard deviation 0.01 in each direction. The superimposed dotted line is the density function of the Gaussian noise applied to the vertices. This is accumulated over 800 meshes, specifically 10 noisy versions of each of the 80 meshes from the TOSCA dataset (see Sect. 8.1). As we can see, the noise appears to still be a 0 mean Gaussian with a slightly larger variance.

We further tried to characterize how the variance of the Laplacian noise is related to that of the original vertex noise. To this end, we applied Gaussian vertex noise of increasing variance to the same meshes as in the previous test, and computed the variance of the corresponding Laplacian noise. Figure 3 plots the standard deviation of the vertex noise versus that of the Laplacian noise. As we can see, within the

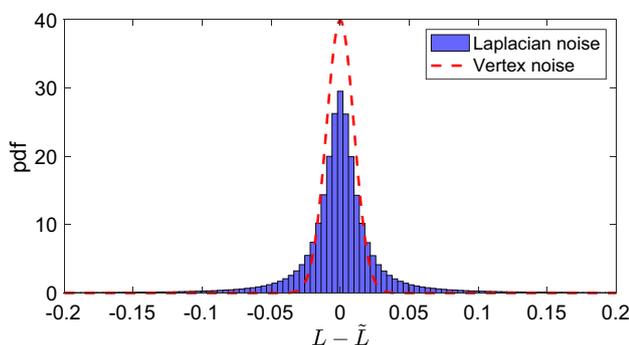


Fig. 2 Histogram of additive noise on the mesh Laplacian, when Gaussian noise with mean 0 and standard deviation 0.01 is added to the mesh vertex positions (red dashed line) (Color figure online)

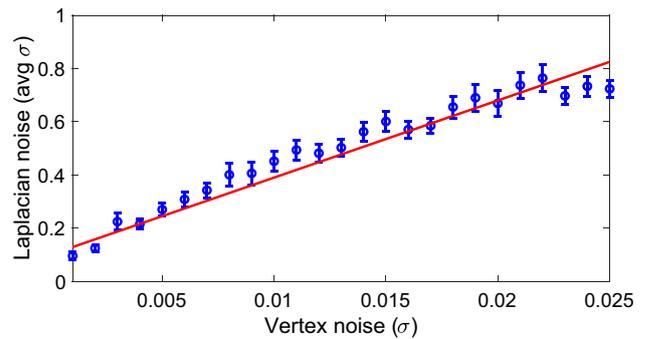


Fig. 3 Average standard deviation (± standard error) of the noise distribution on the mesh Laplacian for increasing values of the standard deviation of additive vertex noise. The red line (with slope ~ 26) highlights the linear relation between the two standard deviations

range of interest the relation is approximately linear with a regression line with a slope of around 26. The observed linear correspondence between vertex noise and Laplacian noise comforts us in the use of Laplacian noise in our perturbation analysis.

6 Computation of the AMKS

Let us consider the discrete setting in which a shape M is represented as a triangular mesh with n points. We can approximate the Laplace–Beltrami operator (LBO) using the cotangent scheme as $L = A^{-1}S$, where S and A are respectively the stiffness and the mass matrix containing the local area elements of each vertex. We can compute an approximation of the eigenfunctions of M as the solution of the generalized eigenvalue problem $S\Phi = A\Phi\Lambda$, with Λ being the diagonal matrix of eigenvalues. In practice we need to compute just first $d \ll n$ smallest eigenvalues and corresponding eigenvectors, since higher eigenvalues encode finer geometric details of the shape that are mostly dominated by the noise introduced by sampling.

We can rewrite the equation of the AMKS in terms of matrix operations, thus exploiting linear algebra computation capabilities of modern hardware and GPUs. Since we are only interested in the diagonal entries of the matrix $AMKS(E)$, we can expand it as

$$AMKS(E) = \frac{\sum_{\lambda_1, \lambda_2} p_{\lambda_1}^2 \circ p_{\lambda_2}^2 \text{sinc}(T(\lambda_2 - \lambda_1)) f_E(\lambda_1) f_E(\lambda_2)}{\sum_{\lambda_1, \lambda_2} f_E(\lambda_1) f_E(\lambda_2)}, \tag{25}$$

where $p_{\lambda}^2 = \sum_i \phi_{\lambda_i} \circ \phi_{\lambda_i}$ is the diagonal of P_{λ} .

Assuming $k \leq d$ unique eigenvalues, we can define the $k \times k$ matrices $B(E) = (b(E)_{uv})$ and $S = (s_{uv})$ such that

$$s_{uv} = \text{sinc}(T(\lambda_u - \lambda_v))$$

$$b(E)_{uv} = \frac{S_{uv} f_E(\lambda_u) f_E(\lambda_v)}{\sum_{\lambda_1, \lambda_2} f_E(\lambda_u) f_E(\lambda_v)}$$

We can see that

$$B(E) = (\mathbf{1}^\top (\mathbf{f}_E \mathbf{f}_E^\top) \mathbf{1})^{-1} (S \circ \mathbf{f}_E \mathbf{f}_E^\top),$$

with $\mathbf{f}_E = [f_E(\lambda_1) \dots f_E(\lambda_k)]^\top$ being a column vector and $\mathbf{1}$ denoting the column vector of all ones. Let us define $Q = [p_{\lambda_1}^2 | \dots | p_{\lambda_k}^2]$ as the matrix composed by concatenating the column vectors $p_{\lambda_i}^2$. Note that, in the special case in which $k = d$, we have $Q = \Phi^{\circ 2}$ where Φ is the eigenvector matrix and \circ^2 being the elementwise square of a matrix.

The value of the descriptor at energy level E can be computed as

$$\text{AMKS}(E) = (Q \circ (QB)) \mathbf{1}. \tag{26}$$

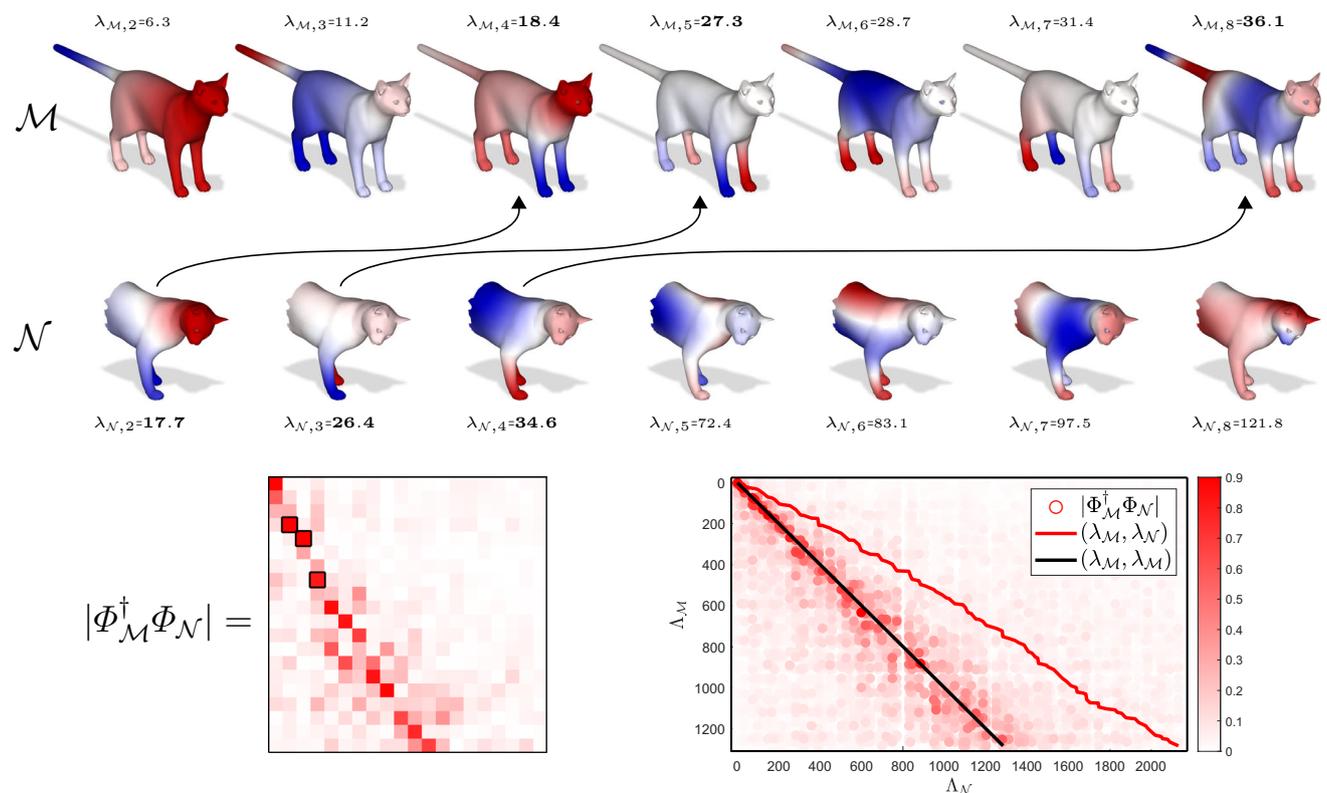


Fig. 4 Example showing the relation between eigenvalues and eigenvectors of a complete shape and a deformed partial version thereof. The **top** figure shows the similarity of the eigenvectors (drawn as a colormap over the surface) corresponding to similar eigenvalues between the two shapes. The **bottom-left** matrix shows the correlation between the eigenvectors of the two shapes according to their index. The 3 black squares correspond to the three arrows of the top figure. The **bottom-**

7 Partiality

Let us consider a setup in which we are given a shape \mathcal{M} and a partial version thereof $\mathcal{N} \subseteq \mathcal{M}$, possibly under a near-isometric transformation.

In Rodolà et al. (2017a), it has been shown that there exists a relation between the eigenbases of \mathcal{M} and \mathcal{N} . In particular, up to the noise introduced by the boundary, there is a one to one correspondence between the eigenfunctions $\Phi_{\mathcal{N}}$ of the partial shape and the eigenfunctions $\Phi_{\mathcal{M}}$ localized on the corresponding part of the full shape. On the other hand, eigenfunctions of the full shape not localized in the corresponding part are not present in the partial shape. The ratio with which eigenvectors of the full shape appear in the partial one roughly follows Weyl’s asymptotic law, and it is directly related to the ratio between the two eigenvalue sequences.

This behavior is visible in the example shown in Fig. 4. Here the first three non-null eigenfunctions of the partial shape \mathcal{N} roughly correspond to the 4th, 5th and 8th eigenfunctions on the full shape \mathcal{M} . Note that the skipped

right scatter plot shows the correlation between the eigenvectors of the two shapes as colored points. The coordinates of each point are equal to the corresponding eigenvalues showing that the eigenvectors of the two shapes have a higher correlation in correspondence of similar eigenvalues (black line) rather than at similar indexes (red line) (Color figure online)

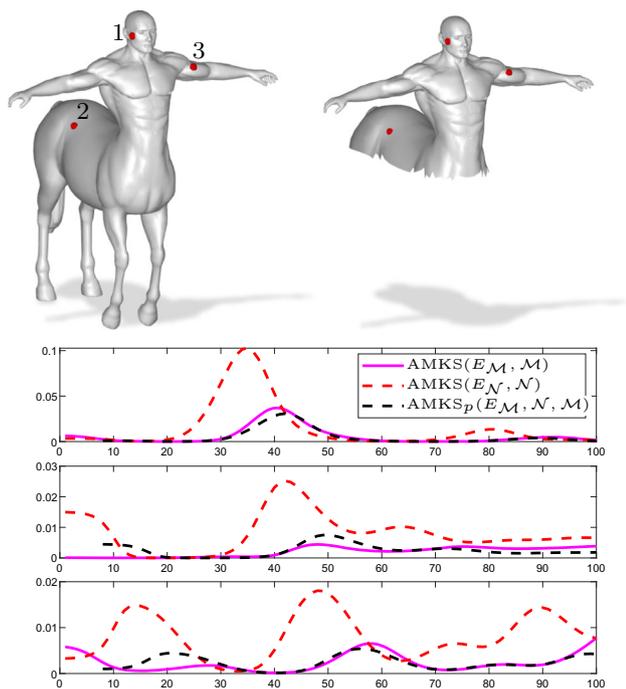


Fig. 5 Comparison between the AMKS descriptors on three corresponding points on a full and a partial shape, before and after applying the proposed correction. From the plot on the bottom, we can see how the corrected descriptor (dashed black line) on the partial shape is much more similar to the corresponding descriptor (magenta line) on the full shape than the original descriptor (dashed red line) (Color figure online)

eigenfunctions on the full shape are indeed mainly localized on the missing part. This correspondence is also visible on the bottom left correlation matrix in terms of dot product between the eigenfunctions of the two shapes. Ideally, the correlation should be one in case of corresponding functions and zero otherwise. As explained in Rodolà et al. (2017a), the noise introduced by the boundary and the non-isometric component of the deformation induces some fuzziness in the correspondence between eigenfunctions which grows with their signal frequency (i.e., the eigenfunction index). It is also interesting to note how similar eigenfunctions appear to correspond to similar eigenvalues on the two shapes. This is further confirmed by the bottom right scatter plot, where we plot the dot product of the eigenvectors of the two shapes as circles with coordinates equal to the corresponding eigenvalue pair. Here the red color indicates higher similarity, showing that the eigenvectors of the two shapes tend to be in correspondence at similar eigenvalues (black line) rather than at the same index on the sequence (red line).

This behavior has to be taken into account when computing the AMKS descriptor for partial shapes. In particular, it affects the choice of the energy levels used to sample the descriptor, which have to be consistent between the full and the partial shape. A logarithmic sampling based on just the eigenvalues of the partial shape considered in isolation is thus

not ideal for the partial case. A simple yet effective solution to this problem is to use the same energy levels $E_{\mathcal{X}}$ computed on the full shape to compute the descriptor also on the partial shape.

Computing the eigenvectors through the generalized eigenvalue problem described in Sect. 6 leads to the identity $\Phi^T A \Phi = I$. Since the norm of each eigenvector depends on the total area of the shape, we have to further scale the descriptors of the partial shape by the squared area ratio between the partial and the full shape. Let $AMKS(E_{\mathcal{M}}, \mathcal{N})$ be the signature computed on \mathcal{N} with energy bands computed on \mathcal{M} . We define the corrected descriptor as:

$$AMKS_p(E_{\mathcal{M}}, \mathcal{N}, \mathcal{M}) = AMKS(E_{\mathcal{M}}, \mathcal{N}) \frac{area(\mathcal{N})^2}{area(\mathcal{M})^2}, \quad (27)$$

where $E_{\mathcal{M}}$ are the energy bands computed on the full shape, and $area(\mathcal{N})$ and $area(\mathcal{M})$ are the surface area of the partial and full shapes, respectively.

In Fig. 5 we show the AMKS descriptor before and after the proposed correction. We can see how the corrected descriptor (dashed black curve) exhibits a much better alignment with the corresponding point descriptor of the complete shape (magenta curve).

8 Experimental Evaluation

In this section we perform a series of quantitative and qualitative analyses to test the performance of our descriptor against two alternative state-of-the-art spectral signatures, the WKS (Aubry et al. 2011b) and the scaled HKS (Sun et al. 2009), which represent two of the most successful and widely used non-learned descriptors for deformable shapes. We start by describing the experimental setup, followed by the application to the classical case of finding the correspondence between two complete shapes. Finally, we analyze our descriptor in the case of part-to-full shape matching.

8.1 Experimental Setup

Datasets The main datasets used for the evaluation of our descriptor are TOSCA (Bronstein et al. 2008) and FAUST (Bogo et al. 2014). The former comprises 8 humanoid and animal shape classes, for a total of 80 meshes of varying resolution (3K to 50K vertices). The FAUST dataset (both the real scans and synthetic version) consists of scanned human shapes in different poses. In total, there are 10 human subjects, each in 10 different poses, for a total of 100 meshes. In the next experiments, we refer to the synthetic version, unless specified otherwise.

Note that both TOSCA and FAUST contain shapes undergoing non-isometric deformations and that in FAUST we also have access to inter-class ground-truth correspondences. In both datasets, all shapes have been remeshed to 5K vertices by iterative edge contractions and re-scaled to unit area. For the *topological noise* experiment we used the SHREC'16 topological noise dataset (Löhner et al. 2016), derived from the KIDS dataset (Rodolà et al. 2014) to include self intersections in the shapes. Finally, for the partial correspondence scenario we used the SHREC'16 partial dataset (Cosmo et al. 2016a) and a synthetic dataset obtained from FAUST (both described in detail in Sect. 8.3).

Distance In order to compare descriptors, we adopted the distance proposed by Aubry et al. in Aubry et al. (2011b), here denoted as L_{KS}^1 . Given two shapes \mathcal{X} and \mathcal{Y} , we measure the distance between two descriptors corresponding to the points $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ as

$$L_{KS}^1(x, y) = \sum_i^d \left| \frac{KS(x)_i - KS(y)_i}{KS(x)_i + KS(y)_i} \right|, \quad (28)$$

where $KS(\cdot)$ is the descriptor vector for a given point.

Metrics The signatures were evaluated using two measures, the Cumulative Match Characteristic (CMC) and the correspondence accuracy (Kim et al. 2011). We use the CMC to estimate the probability of finding a correct correspondence among the p -nearest neighbors in descriptor space. The probability (hit rate) is calculated as the percentage of correct matches among the p -nearest neighbors. We average the hit rate for all points across all pairs of shapes for increasing values of p . The resulting CMC curve is monotonically increasing in p .

The Princeton protocol (Kim et al. 2011), also known as correspondence accuracy, captures the proximity of predicted corresponding points to the ground-truth ones in terms of geodesic distance on the surface. Specifically, we count the percentage of matches that are below a certain distance from the ground-truth correspondence. Given a pair of shapes and an input feature descriptor for a point on one of the shapes, we find the closest point on the other shape in descriptor space. Then we calculate the geodesic distance between its position over the surface and the position of the ground-truth corresponding point. Given a match $(x, y) \in \mathcal{X} \times \mathcal{Y}$ and the ground-truth correspondence (x, y^*) , the normalized geodesic error is $\epsilon(x) = d_{\mathcal{Y}}(y, y^*)/\text{area}(\mathcal{Y})^{0.5}$.

Implementation details We used the cotangent scheme approximation of the Laplace–Beltrami operator. Unless otherwise specified, we considered $k = 100$ eigenvectors of the Laplacian and a descriptor size of $d = 100$ in all methods. The choice of diffusion time will be discussed later. For the settings specific to the HKS and the WKS, we refer the reader to Sun et al. (2009) and Aubry et al. (2011b), respectively.

Our method requires setting the same number of parameters as the HKS and the WKS, with the exception of an additional time parameter. For the energy we use the same range as the WKS (Aubry et al. 2011b).

The time parameter is extensively discussed in the next paragraph, where we show that the AMKS is not particularly sensitive to its value provided that it sits within a certain range.

Choice of time Recall that computing the AMKS requires choosing a time T . At first, this seems to add further complexity to our signature, compared to the HKS and the WKS. For this reason, we decided to perform some tests over two separate subsets of shapes from TOSCA and FAUST in order to evaluate how the choice of the diffusion time affects the performance of our descriptor.

Table 1 shows the hit ratio of our descriptor on the nearest 0.1% of points in descriptor space. In these experiments, we explore how the hit ratio changes for varying levels of noise and choices of the stopping time T . In particular, columns 1 to 11 of Table 1 show the values of the hit ratio for finite values of T , while column 12 shows the value in the limit $T \rightarrow \infty$ (see Eq. 17).

In addition, we compare the results to those obtained with the WKS (column 13), noting that the WKS is time-independent. Finally, we also show the results obtained when using a signature that is a concatenation of the AMKS (for the optimal stopping time T) and the WKS. For the TOSCA dataset, we employed 7 classes (wolf excluded) for a total of 28 pairs, whereas for FAUST we used 10 pairs. Here a pair consists of two shapes: the class reference shape, in canonical pose, and another shape of the same class in a different pose. When adding noise, the reference shape is first rescaled, then Gaussian noise is added to the vertex positions and finally the shape is rescaled again.

It is worth noting that (a) the hit rate of the AMKS is almost always higher than that of the WKS, regardless of the choice of the time T , (b) the optimal times for the two datasets are similar but different, (c) finite stopping times invariably yield better results than taking the limit, and (d) using both the WKS and the AMKS does not increase the performance of the signature, suggesting that the information contained in the AMKS is strictly greater than that contained in the WKS.

Note that computing the AMKS in the time limit $T \rightarrow \infty$ effectively means removing the mixing component, yielding a signature that is equal to the WKS, with the only difference being the use of the fourth power (instead of the square) of the entries of the Laplacian eigenvectors. When compared to the WKS, the results in Table 1 suggest that the fourth power has a stronger smoothing effect close to zero values, resulting in a loss of descriptive power wrt to the WKS in the absence of noise. This in turn highlights the importance of taking into account the mixing between eigenspaces that we observe at finite time averages in the AMKS. While it

Table 1 Average percentage of correspondences \pm standard error on the nearest 0.1% neighbors in the descriptor space for the AMKS (for increasing stopping time T , columns 1–12), the WKS (second to last column), and a concatenation of AMKS and WKS (last column) over a subset of the TOSCA (top) and FAUST (bottom) datasets with added Gaussian noise (η) to the vertex positions (rows)

$\eta \setminus T$	0.005	0.01	0.2	0.5	0.8	1	1.2	1.5	1.8	2	2.5	∞	WKS	AMKS _{(T=0.5)+WKS}
0	36.8 \pm 2.1	36.8 \pm 2.1	37.1 \pm 2.1	36.9 \pm 2.1	36.3 \pm 2.1	35.8 \pm 2.1	35.4 \pm 2.1	34.7 \pm 2.0	34.2 \pm 2.0	33.7 \pm 1.9	32.6 \pm 1.9	27.7 \pm 1.5	33.0 \pm 1.6	36.8 \pm 2.0
0.01	27.9 \pm 1.4	27.9 \pm 1.4	28.2 \pm 1.4	28.3 \pm 1.4	28.0 \pm 1.4	27.7 \pm 1.4	27.3 \pm 1.4	26.9 \pm 1.4	26.5 \pm 1.4	26.2 \pm 1.4	25.5 \pm 1.3	22.1 \pm 1.3	24.9 \pm 1.2	28.4 \pm 1.2
0.015	20.1 \pm 1.1	20.1 \pm 1.1	20.5 \pm 1.1	20.5 \pm 1.1	20.3 \pm 1.0	20.0 \pm 1.0	19.8 \pm 1.0	19.6 \pm 1.0	19.5 \pm 1.0	19.1 \pm 0.9	18.8 \pm 0.9	16.6 \pm 0.8	17.6 \pm 0.8	20.4 \pm 0.9
0.02	14.1 \pm 0.8	14.1 \pm 0.8	14.5 \pm 0.8	14.5 \pm 0.8	14.2 \pm 0.8	14.1 \pm 0.8	14.0 \pm 0.7	13.8 \pm 0.7	13.7 \pm 0.7	13.5 \pm 0.7	13.2 \pm 0.7	11.7 \pm 0.5	11.8 \pm 0.6	14.7 \pm 0.6
$\eta \setminus T$	0.005	0.01	0.2	0.5	0.8	1	1.2	1.5	1.8	2	2.5	∞	WKS	AMKS _{(T=1)+WKS}
0	53.3 \pm 2.4	53.3 \pm 2.4	53.7 \pm 2.4	53.8 \pm 2.5	53.4 \pm 2.5	53.2 \pm 2.4	52.6 \pm 2.4	51.4 \pm 2.4	50.8 \pm 2.4	50.3 \pm 2.4	49.2 \pm 2.4	40.7 \pm 2.5	47.0 \pm 2.5	53.1 \pm 2.5
0.01	33.9 \pm 1.1	33.9 \pm 1.1	34.2 \pm 1.1	34.7 \pm 1.2	35.3 \pm 1.2	35.4 \pm 1.2	35.1 \pm 1.2	34.7 \pm 1.1	34.1 \pm 1.1	33.7 \pm 1.1	32.9 \pm 1.1	24.4 \pm 1.1	29.5 \pm 1.0	35.2 \pm 1.4
0.015	19.3 \pm 0.6	19.3 \pm 0.6	19.6 \pm 0.6	19.9 \pm 0.6	20.4 \pm 0.7	20.4 \pm 0.6	20.3 \pm 0.6	20.1 \pm 0.6	19.6 \pm 0.6	19.3 \pm 0.6	18.7 \pm 0.6	15.1 \pm 0.8	18.1 \pm 0.8	20.7 \pm 0.7
0.02	10.8 \pm 0.4	10.8 \pm 0.4	10.7 \pm 0.4	10.8 \pm 0.5	11.0 \pm 0.4	11.1 \pm 0.5	11.0 \pm 0.5	10.8 \pm 0.5	10.5 \pm 0.5	10.2 \pm 0.5	10.0 \pm 0.4	8.8 \pm 0.5	10.5 \pm 0.4	11.3 \pm 0.5

The best value for each row is highlighted in bold

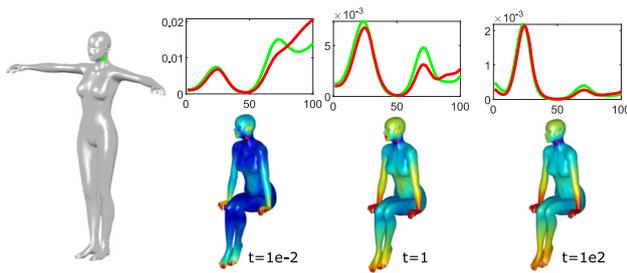


Fig. 6 The figure shows the AMKS of the same point on different shapes, at increasing times $t = 0.01, 1, 100$. Each plot above the sitting shape shows the behavior of two signatures: the green curve belongs to the shape in the reference pose whereas the red curve represents the sitting shape

would be tempting to think that a modified version of the WKS that includes the mixing components may therefore be superior to both the AMKS and the WKS, note that this would make little sense from the point of view of physics and it would have actually a detrimental effect on the signature performance, as we quickly realised during the experimental evaluation. In other words, the mixing alone is not sufficient to justify the increased performance of the AMKS and it is the synergy between the use of the fourth power and finite times that leads to the improved performance and resilience to noise of the AMKS.

As for the optimal value of the stopping time, we observe that the optimal performance of the AMKS in the FAUST dataset is reached at slightly higher times when compared to TOSCA. Indeed, in the TOSCA dataset there is a stronger mixing of eigenvectors as it exhibits smaller deviations from isometry. Note however that the performance of the signature is fairly stable on a relatively wide interval around $T = 0.5$, implying that fine-tuning this parameter is not crucial. This optimal range, in turn, is a consequence of Weyl's law and the eigenvalues normalization that follows from normalizing all shapes to have unit area. Consequently, in order to avoid leaving free parameters and to ensure a fair comparison, in the next experiments we set $T = 0.5$ for both the TOSCA and FAUST datasets. Finally, note that we have run similar experiments on the partial shape datasets and the results were consistent with those of Table 1 and thus are omitted.

To conclude our investigation on the time parameter, in Fig. 6 we visualize how the AMKS signature of a point changes as the stopping time varies. By increasing the stopping time, the rate at which eigenspaces are mixed decreases, resulting in flatter and less descriptive signatures, while for very short times the signature varies significantly on the high energy bands. Hence, the optimal choice of the diffusion time is a key factor in the trade-off between descriptiveness and sensitivity of the signature.

8.2 Full Shapes Correspondence

In this section we compare the performance of our signature to that of the WKS and the HKS in the context of matching between two complete shapes. We also analyze the descriptors when the surface is represented as a point cloud rather than a triangular mesh, and in the presence of topological noise.

Descriptor robustness We commence by evaluating the robustness of the descriptor when Gaussian noise is added to the coordinates of the mesh vertices. Fig. 7 shows the CMC curve of our method compared to those of the other spectral signatures, for varying levels of Gaussian noise (left to right, $\eta = 0, 0.01, 0.02$). For our signature, we set $T = 0.5$ following the analysis of the previous section. In these plots, the y-axis is the hit rate \pm standard error (the plots were drawn by averaging the hit rate over all the pairs of shapes of each dataset) and the x-axis refers to the p -nearest neighbors (in percentage) in the descriptor space. Our approach clearly outperforms the alternatives, with the difference becoming larger as the noise increases. This is more evident in the TOSCA dataset, as it exhibits smaller deviations from isometry. This shows the noise robustness of our signature.

In Fig. 8 we provide a qualitative comparison of the behaviors of the AMKS and the WKS, showing the distance in descriptor space of a given point from the other points of the shape. We can see that our descriptor is more informative showing a more peaked distribution around the correct match, even in the presence of high levels of Gaussian noise.

Point-wise correspondence In this experiment, we use the three signatures (HKS, WKS, and AMKS) to compute point-wise correspondences between pairs of shapes by taking the nearest neighbor in descriptor space. In Fig. 9 we show the results following the previously described Princeton protocol. As expected, our method shows the largest improvement on the TOSCA and FAUST inter-class datasets, which contain stronger non-isometric deformations.

In Fig. 10 we show some qualitative examples illustrating the normalized geodesic error, where we compare our results with those obtained with the WKS on a selection of shapes from the TOSCA dataset. The lower presence of colored spots for the AMKS suggests a better matching quality.

Topological Noise We use the SHREC16 topological noise dataset (Löhner et al. 2016) to investigate the behavior of our descriptor in the presence of self intersections (topological “gluing”) on the surface representation. Spectral descriptors are known to be particularly sensitive to topological changes in the shape but, as shown in Fig. 11, the AMKS maintains its edge over the HKS and the WKS even in the presence of this type of noise.

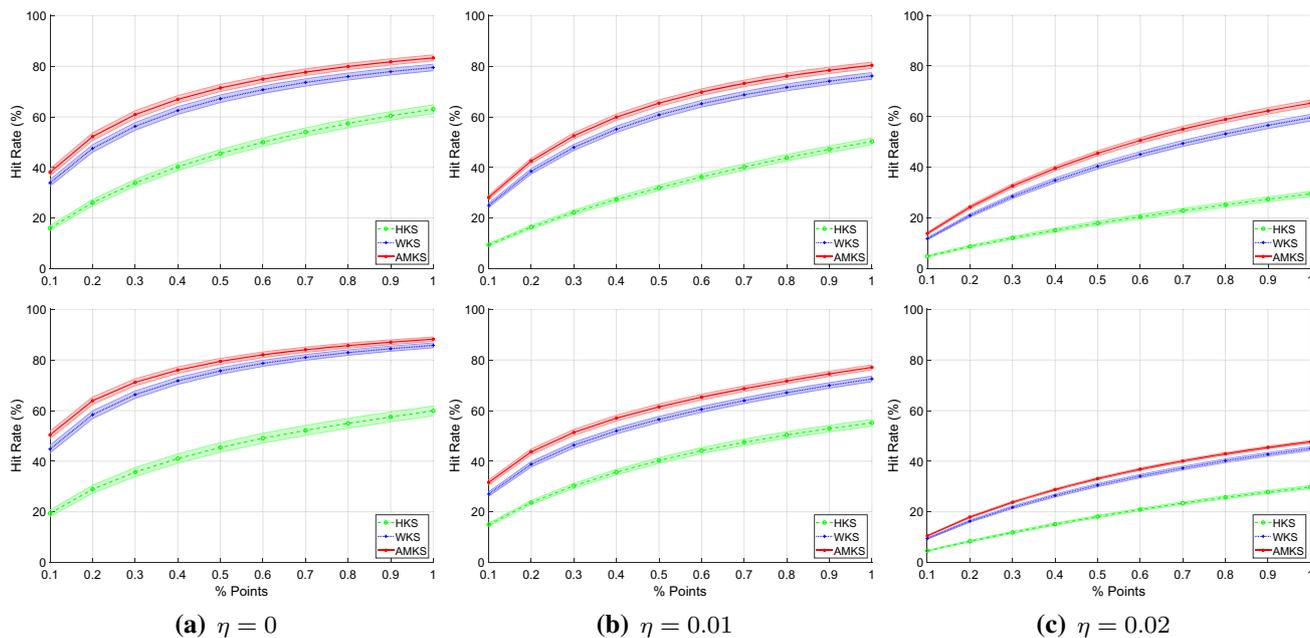


Fig. 7 Quantitative comparison of the matching performance of different descriptors (HKS, WKS, and the proposed AMKS) on the TOSCA dataset (top) and the FAUST dataset (bottom) using the CMC curve \pm

standard error, considering an increasing percentage of nearest points in descriptor space (x -axis) and with increasing level of Gaussian noise (left to right)

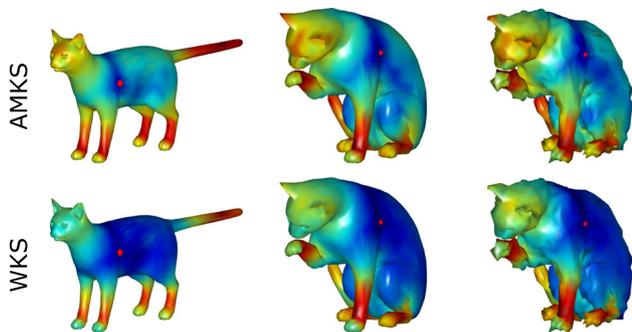


Fig. 8 Distance in the descriptor space between a given point (red dot) and all the other points of the shape, where the blue color corresponds to smaller distances. The *left* figure shows the behavior of our AMKS at increasing noise levels, compared with WKS (Color figure online)

Point clouds Being a spectral descriptor, the AMKS can be applied to any surface representation where a discretization of the Laplace Beltrami Operator can be computed. Point clouds are widely used representations in computer vision for 3D surfaces. We use the shapes in TOSCA and FAUST to generate two datasets of synthetic point clouds with ground-truth correspondences. We sample the surface of each shape on a regular 3D grid resulting in point clouds of approximately 10k points uniformly distributed over the surface. We use the method proposed in Clarenz et al. (2004) to discretize the Laplacian on a point cloud. Figure 12 shows the CMC curves on the point cloud datasets for different amount of Gaussian noise (applied as a displacement along the normal at each point). In both datasets the AMKS outperforms the

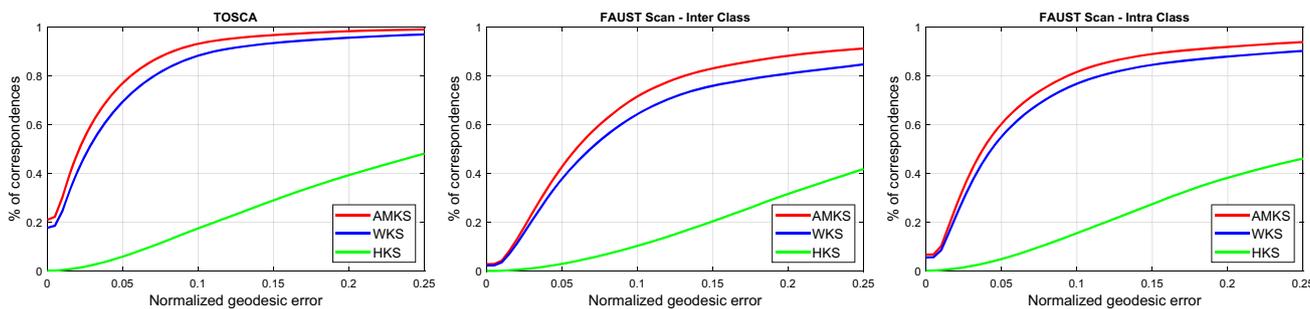


Fig. 9 Quantitative comparison between spectral signatures in terms of the Princeton protocol on the TOSCA (left) and FAUST Scans dataset (center and right). Intra- and inter-class refer to the case where the match is computed only between meshes representing the same and different subjects, respectively

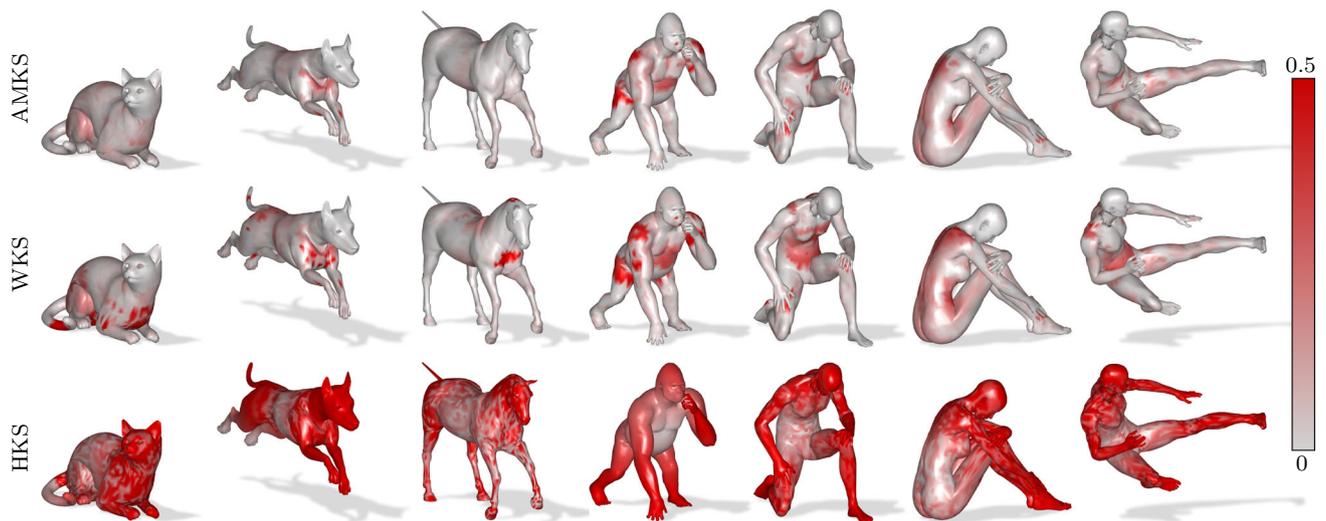


Fig. 10 Qualitative comparison between the WKS and the AMKS on some shapes from the TOSCA dataset. For each point of the surface we show the normalized geodesic error between the closest descriptor on

the corresponding *reference shape* and the ground truth correspondence (or its symmetric one). Overall, the AMKS leads to a better matching quality than the alternative signatures

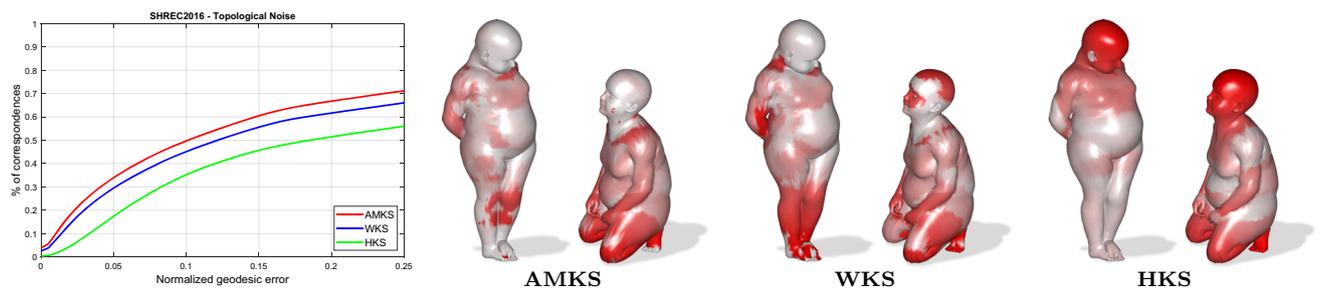


Fig. 11 Left: comparison of CMC curves for different methods on the *SHREC16 Topological Noise* dataset (Lähler et al. 2016). Right: qualitative comparison showing the geodesic distance on the surface between the computed and ground-truth matches

other descriptors, with the margin becoming slightly larger as the noise is increased. This suggests that our descriptor is a good choice in scenarios where the point cloud data is intrinsically noisy due to the acquisition process.

Point classification One of the most used methods to train a neural network to perform the task of shape matching is to cast it as a classification task, where corresponding points on different shapes belong to the same class and there are as many classes as the number of points per shape (Rodolà et al. 2014). We used our descriptor as input of two of the state-of-the-art architectures for convolutional deep learning on shapes: FeastNet (Verma et al. 2018) and MoNet (Monti et al. 2017).

FeastNet is a deep neural network based on a graph-convolution operator that dynamically learns a mapping between vertices and filter weights, from features learned by the network. MoNet defines convolution-like operations

as template matching with local intrinsic patches on graphs or manifolds. It replaces the weight functions with Gaussian kernels with learnable parameters, which can be interpreted as a Gaussian Mixture Model (GMM). For both methods we used the code provided by authors on GitHub. For FeastNet we adopted the translation invariant single-scale network architecture proposed in Verma et al. (2018). The first 80 FAUST shapes were used for training and the remaining 20 for test.

Table 2 shows the classification accuracy results after training the networks with different descriptors as input. We compare our descriptor against the WKS, the local rigid descriptor SHOT (Salti et al. 2014) and using the XYZ coordinates of the points as input. The results show how the network is able to exploit the higher descriptiveness of AMKS in order to achieve better classification accuracy. Note that, while training directly on XYZ coordinates

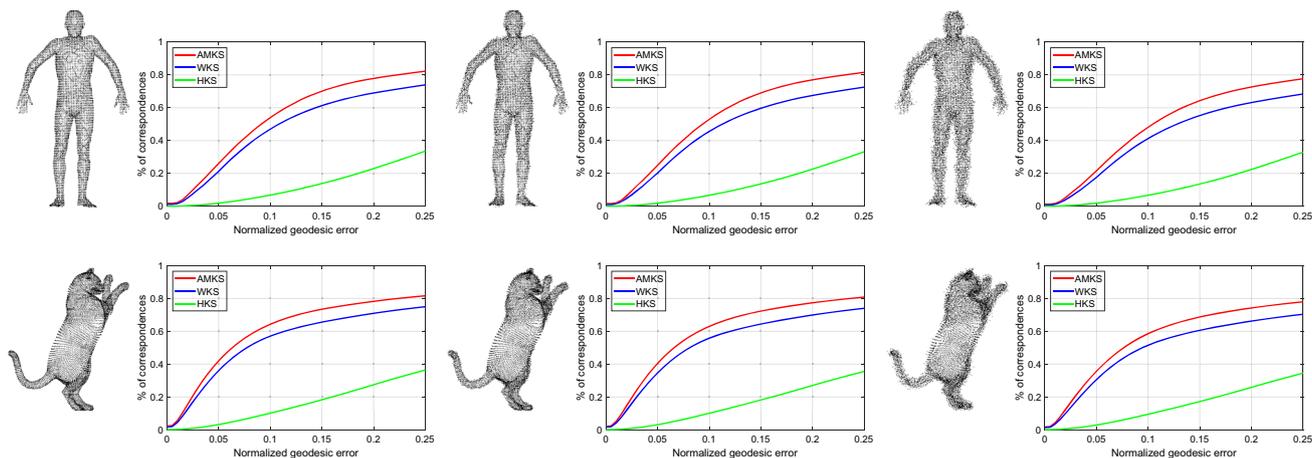


Fig. 12 Quantitative comparison of different methods (HKS, WKS, and the proposed AMKS) on point clouds derived from FAUST (top) and TOSCA (bottom) with increasing levels of Gaussian noise along

the surface normals. Each of the three columns shows a sample point cloud and the corresponding CMC curves for each method at different noise levels: 0.005 (left), 0.02 (middle), 0.05 (right)

Table 2 Classification accuracy of FeastNet and MoNet trained with different input descriptors

	XYZ (%)	ROT(XYZ) (%)	SHOT (%)	WKS (%)	AMKS (%)
FeastNet	86	50	63	80	86
MoNet	88	64	75	94	95

The best result for each architecture is highlighted in bold

allows to achieve a similar accuracy, the performance rapidly decreases if we allow shapes to undergo a random rotation (i.e., ROT(XYZ)) since neither the XYZ coordinates nor the network are intrinsically invariant to this transformation.

It should be noted that this experiment was not meant to provide an exhaustive comparison of our descriptor against learned ones, as the performance of the latter usually comes at the cost of an increased complexity as well as the need of large amounts of training data. Our aim in this experiment

was instead to show that the AMKS can indeed be used as part of standard learning-based matching pipelines.

8.3 Partial Shapes Correspondence

To assess the ability of the proposed descriptor to handle shapes with missing parts we applied it to two datasets containing partial shapes. SHREC16 (Cosmo et al. 2016a) is a dataset of animal and human shapes in different poses containing two types of partiality, neat cuts and holes. In our experiments we consider only the *cuts* sub-dataset, in which the shapes are missing a whole contiguous part of their body. We also test on a more realistic scenario in which the partial shape is given by a synthetically generated range acquisition of a human subject. We refer to this dataset as RANGE. Range maps are obtained by capturing shapes of FAUST with a virtual camera. Some examples of the synthetic point clouds are shown in Fig. 13.



Fig. 13 Left: some examples of the synthetic range maps contained in the RANGE dataset. Right: comparison of CMC curves for the AMKS, its partial version AMKS_p, WKS, its partial version WKS_p, SHOT, and HKS on the RANGE dataset

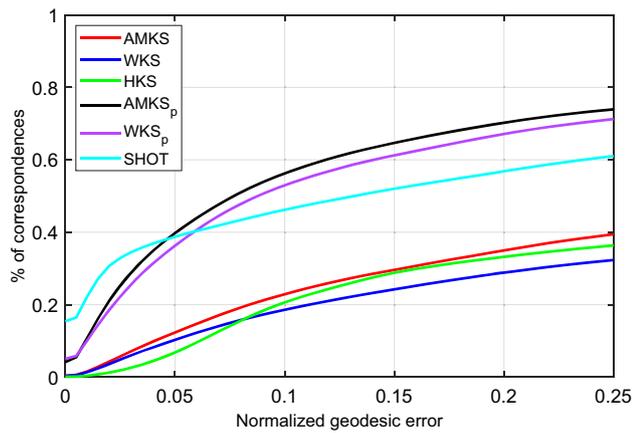


Fig. 14 Comparison of CMC curves of AMKS, its partial version $AMKS_p$, its partial version $AMKS_p$, WKS, its partial version WKS_p , SHOT, and HKS on the cuts sub-dataset of SHREC2016 partial dataset (Cosmo et al. 2016a)

Figures 13 and 14 show the CMC curves for the AMKS adapted to handle partiality ($AMKS_p$), the original formulation of the AMKS, the WKS, and the HKS on the RANGE and SHREC16 datasets, respectively. We also include the results for the SHOT descriptor (Salti et al. 2014) and a version of the WKS adapted to handle partiality similarly to what done for our signature.

As the figures show, correctly handling partiality is essential for the performance of the descriptor as well as for the WKS. However it is also worth noting how the original formulation of the AMKS is still performing better than the

competitors and that the AMKS outperforms the WKS even after both have been modified to cope with partiality.

The superiority of our approach is also visible in the qualitative examples reported in Fig. 15, where we draw over each shape a colormap proportional to the geodesic distance between the correct match and the retrieved one.

8.4 Functional Maps Pipeline

As a further experiment, we evaluated the performance of our descriptor when used in the context of a complete functional maps pipeline. In recent years, approaches based on functional maps (Ovsjanikov et al. 2012; Rodolà et al. 2017a; Cosmo et al. 2016b) have proven to be very effective particularly for non rigid shape, thanks to their implicit shape representation.

For the case where the input shapes are complete, we used the original method proposed in Ovsjanikov et al. (2012). The basic idea is to define a correspondence between two shapes as a linear map between their functional bases defined as the eigenfunctions of their Laplace–Beltrami operators. In this way, finding a correspondence is equal to optimizing over the functional map. Given a large enough set of corresponding functions, i.e., descriptors between the two shapes, this can be solved as a linear system. For the case of partial shapes, we used instead the method of Rodolà et al. (2017a). This is similar to Ovsjanikov et al. (2012), but it leverages the particular slanted diagonal structure of the functional map itself to handle partiality.

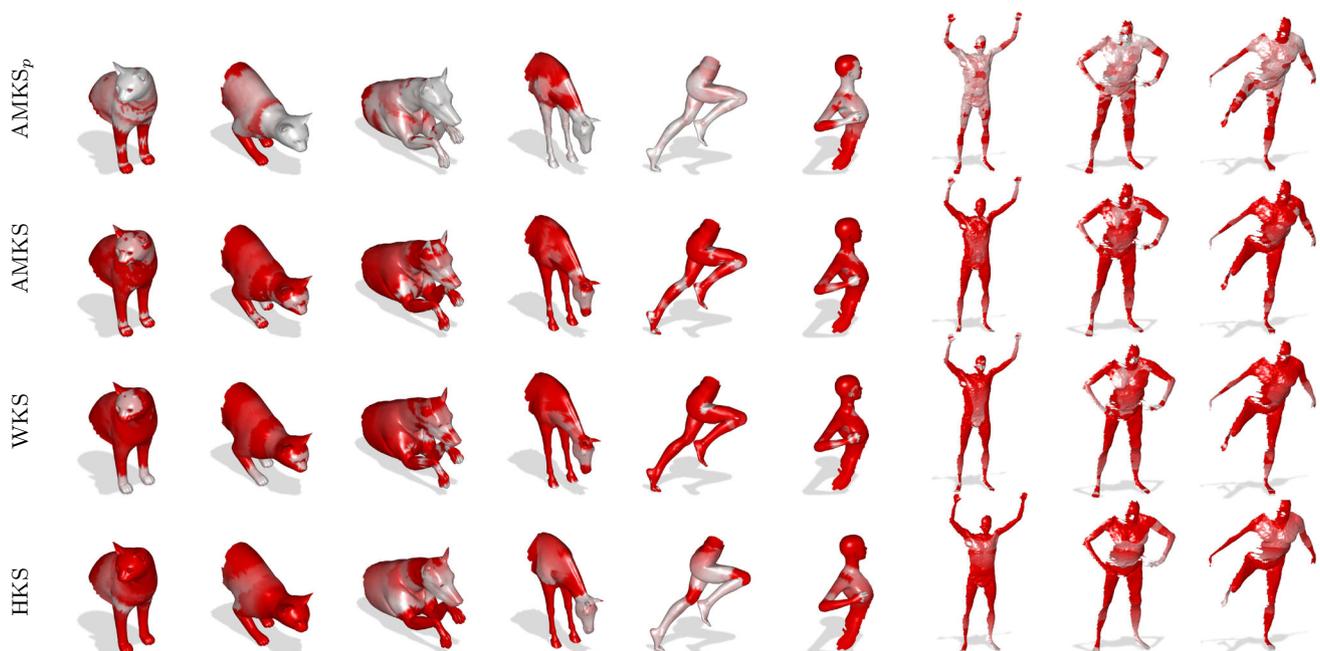


Fig. 15 Qualitative examples of partial shapes with error map (AMKS partial vs AMKS), growing from white to dark red (Color figure online)

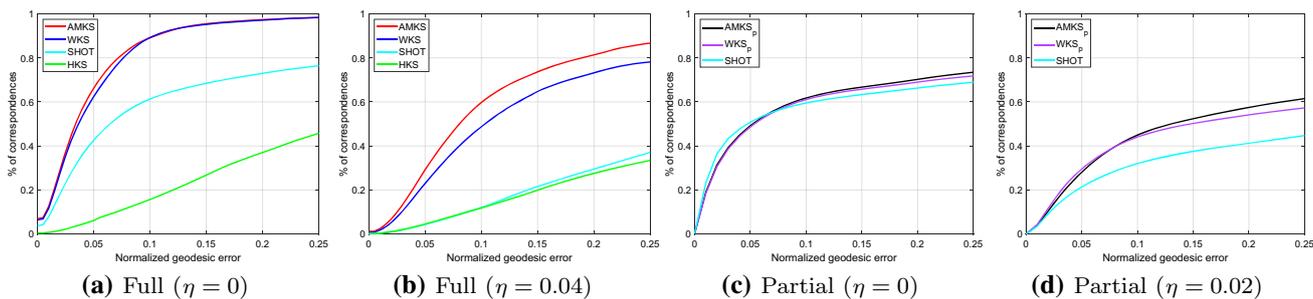


Fig. 16 Comparison of CMC curves for AMKS, WKS, SHOT, and HKS used as input of a functional maps pipeline on **a, b** full and **c, d** partial shapes of the TOSCA dataset. **b, d** refer to the scenarios where noise is added to the dataset (full: $\eta = 0.04$, partial: $\eta = 0.02$)

Figure 16 shows the results obtained when using our descriptor against three alternatives (HKS, WKS, or SHOT) as part of a functional maps pipeline on the full (a) and (c) partial (Cosmo et al. 2016a) shapes of the TOSCA dataset. The various constraints and refinement procedures introduced by the matching pipeline have a clear regularization effect on the final correspondence, resulting in a reduced margin between the AMKS and the WKS, in the case of noiseless shapes. To better appreciate the robustness of AMKS to noise, we then run the same experiment on a noisy version of the two datasets. With the addition of noise, the advantage of using the AMKS, which we have shown to be naturally resilient to noise, becomes apparent, as shown in Fig. 16b, d. Note also that the SHOT descriptor, despite being a good choice for partial shapes, is highly impacted even by mild levels of noise.

8.5 Runtime Analysis

Finally, we analyze the running time needed to compute our descriptor. Without considering the spectral decomposition, the computational complexity of the AMKS is $\mathcal{O}(dnk^2)$, with n being the number of points of the mesh, k the dimension of the truncated eigenbasis, and d the number of energy levels¹. Here we perform an experimental evaluation of the running time needed to compute the AMKS. To this end, we implement our method in MATLAB and we run the code on a desktop workstation with 16GB of RAM and Intel i7 8600 processor. Since the eigendecomposition is a common step in all the compared methods, we ignore it when measuring the running time. Fig. 17 shows the time needed to compute each descriptor, averaged over 100 executions. In the left plot, we fix the number of eigenvalues to 100 and let the number of points vary between 1k and 10k, while in the right plot we keep 5k points and increase the number of eigenvectors from 50 to 500.

As expected, introducing the mixing factor results in a slightly higher running time while the computational com-

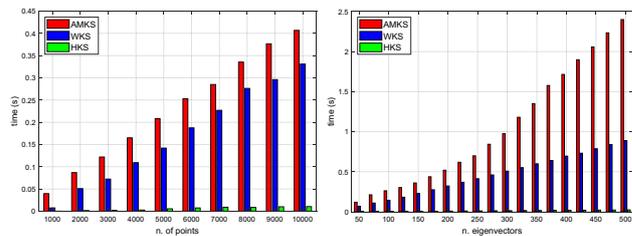


Fig. 17 Comparison of execution times with respect to the number of points (left) and eigenvectors (right)

plexity remains linear in the number of points. On the other hand, we introduce a quadratic complexity with respect to the number of eigenvalues. It is worth noticing, however, that for typical numbers of eigenvectors used in spectral shape retrieval and matching tasks (< 200), the behavior is almost linear and that for a higher number of eigenvectors the main bottleneck would be the spectral decomposition of the Laplace–Beltrami operator.

9 Conclusions

We have proposed a spectral signature for points on non-rigid three-dimensional shapes based on continuous-time quantum walks and the average mixing matrix, holding the transition probabilities between each pair of vertices of a mesh. We have shown, both theoretically and experimentally, that our signature is robust to noise and that it outperforms alternative spectral and non-spectral descriptors under a wide variety of scenarios involving different types of shape representation and noise. We have also proposed a simple yet effective way to cope with shape partiality, and showed that this yields a significant improvement in a partial shape matching task.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence,

¹ The code for the computation of our signature is available on-line at <https://github.com/lcosmo/amks-descriptor>.

unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aubry, M., Schlickewei, U., & Cremers, D. (2011a). Pose-consistent 3d shape segmentation based on a quantum mechanical feature descriptor. In *Joint pattern recognition symposium* (pp. 122–131). Springer.
- Aubry, M., Schlickewei, U., & Cremers, D. (2011b). The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE international conference on computer vision workshops (ICCV Workshops)* (pp. 1626–1633). IEEE.
- Bai, L., Rossi, L., Torsello, A., & Hancock, E. R. (2015). A quantum Jensen–Shannon graph kernel for unattributed graphs. *Pattern Recognition*, *48*(2), 344–355.
- Bell, J. S. (1987). *Speakable and unspeakable in quantum mechanics*. Cambridge: Cambridge University Press.
- Bogo, F., Romero, J., Loper, M., & Black, M. J. (2014). Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 3794–3801).
- Bronstein, A. M., Bronstein, M. M., & Kimmel, R. (2008). *Numerical geometry of non-rigid shapes*. Berlin: Springer.
- Clarenz, U., Rumpf, M., & Telea, A. (2004). Finite elements on point based surfaces. In *Proceedings of the first Eurographics conference on point-based graphics* (pp. 201–211). Eurographics Association.
- Corman, É., Ovsjanikov, M., & Chambolle, A. (2014). Supervised descriptor learning for non-rigid shape matching. In *European conference on computer vision* (pp. 283–298). Springer.
- Cosmo, L., Rodolà, E., Bronstein, M. M., Torsello, A., Cremers, D., & Sahillioglu, Y. (2016a). Shrec'16: Partial matching of deformable shapes. *Proceedings of 3DOR*, *2*(9), 12.
- Cosmo, L., Rodolà, E., Masci, J., Torsello, A., & Bronstein, M. M. (2016b). Matching deformable objects in clutter. In *2016 Fourth international conference on 3D vision (3DV)* (pp. 1–10). IEEE.
- Cosmo, L., Rodolà, E., Albarelli, A., Mémoli, F., & Cremers, D. (2017). Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum*, *36*, 209–221.
- Cosmo, L., Minello, G., Bronstein, M., Rossi, L., & Torsello, A. (2020). The average mixing kernel signature. In *Computer vision—ECCV 2020—16th European conference, Proceedings, Part XX. Lecture Notes in Computer Science* (Vol. 12365, pp. 1–17), Glasgow, UK, August 23–28, 2020. Springer.
- Fang, Y., Xie, J., Dai, G., Wang, M., Zhu, F., Xu, T., & Wong, E. (2015). 3D deep shape descriptor. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2319–2328).
- Gasparetto, A., Minello, G., & Torsello, A. (2015). Non-parametric spectral model for shape retrieval. In *2015 International Conference on 3D Vision* (pp. 344–352). IEEE.
- Godsil, C. (2013). Average mixing of continuous quantum walks. *Journal of Combinatorial Theory, Series A*, *120*(7), 1649–1662.
- Huang, H., Kalogerakis, E., Chaudhuri, S., Ceylan, D., Kim, V. G., & Yumer, E. (2018). Learning local shape descriptors from part correspondences with multiview convolutional networks. *ACM Transactions on Graphics (TOG)*, *37*(1), 6.
- Kempe, Julia. (2003). Quantum random walks: an introductory overview. *Contemporary Physics*, *44*(4), 307–327.
- Kim, V. G., Lipman, Y., & Funkhouser, T. (2011). Blended intrinsic maps. *ACM Transactions on Graphics (TOG)*, *30*, 79.
- Lähner, Z., Rodolà, E., Bronstein, M. M., Cremers, D., Burghard, O., Cosmo, L., Dieckmann, A., Klein, R., & Sahillioglu, Y. (2016). Shrec'16: Matching of deformable shapes with topological noise. In *Proceedings of Eurographics Workshop on 3D Object Retrieval (3DOR)* (Vol. 2).
- Litman, R., & Bronstein, A. M. (2013). Learning spectral descriptors for deformable shape correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *36*(1), 171–180.
- Masci, J., Boscaini, D., Bronstein, M., & Vandergheynst, P. (2015). Geodesic convolutional neural networks on Riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops* (pp. 37–45).
- Minello, G., Rossi, L., & Torsello, A. (2019). Can a quantum walk tell which is which? A study of quantum walk-based graph similarity. *Entropy*, *21*(3), 328.
- Monti, F., Boscaini, D., Masci, J., Rodolà, E., Svoboda, J., & Bronstein, M. M. (2017). Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 5115–5124).
- Murthy, D. V., & Haftka, R. T. (1988). Derivatives of eigenvalues and eigenvectors of a general complex matrix. *International Journal for Numerical Methods in Engineering*, *26*(2), 293–311.
- Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., & Guibas, L. (2012). Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, *31*(4), 1–11.
- Portugal, R. (2013). *Quantum walks and search algorithms*. Berlin: Springer.
- Reuter, M., Wolter, F.-E., & Peinecke, N. (2006). Laplace–Beltrami spectra as ‘shape-dna’ of surfaces and solids. *Computer-Aided Design*, *38*(4), 342–366.
- Rodolà, E., Rota Bulo, S., Windheuser, T., & Vestner, M., & Cremers, D. (2014). Dense non-rigid shape correspondence using random forests. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4177–4184).
- Rodolà, E., Cosmo, L., Bronstein, M. M., Torsello, A., & Cremers, D. (2017a). Partial functional correspondence. *Computer Graphics Forum*, *36*, 222–236.
- Rodolà, E., Cosmo, L., Litany, O., Bronstein, M. M., Bronstein, A. M., Audebert, N., Ben Hamza, A., Boulch, A., Castellani, U., Do, M. N., & Duong, A. D. (2017b). Shrec'17: Deformable shape retrieval with missing parts. In *Proceedings of the Eurographics workshop on 3d object retrieval* (pp. 23–24), Lisbon, Portugal.
- Rossi, L., Torsello, A., & Hancock, E. R. (2012). Approximate axial symmetries from continuous time quantum walks. In *Joint IAPR international workshops on statistical techniques in pattern recognition (SPR) and structural and syntactic pattern recognition (SSPR)* (pp. 144–152). Springer.
- Rossi, L., Torsello, A., Hancock, E. R., & Wilson, R. C. (2013). Characterizing graph symmetries through quantum Jensen–Shannon divergence. *Physical Review E*, *88*(3), 032806.
- Rossi, L., Torsello, A., & Hancock, E. R. (2015). Measuring graph similarity through continuous-time quantum walks and the quantum Jensen–Shannon divergence. *Physical Review E*, *91*(2), 022815.
- Rossi, L., Severini, S., & Torsello, A. (2016). The average mixing matrix signature. In *Joint IAPR international workshops on statistical techniques in pattern recognition (SPR) and structural and syntactic pattern recognition (SSPR)* (pp. 474–484). Springer.
- Rostami, R., Bashiri, F. S., Rostami, B., & Yu, Z. (2019). A survey on data-driven 3D shape descriptors. *Computer Graphics Forum*, *38*, 356–393.
- Rustamov, R. M. (2007). Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the fifth Eurographics symposium on Geometry processing* (pp. 225–233). Eurographics Association.

- Salti, S., Tombari, F., & Di Stefano, L. (2014). Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding*, 125, 251–264.
- Sun, J., Ovsjanikov, M., & Guibas, L. (2009). A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28, 1383–1392.
- Verma, N., Boyer, E., & Verbeek, J. (2018). Feastnet: Feature-steered graph convolutions for 3D shape analysis. In *The IEEE conference on computer vision and pattern recognition (CVPR)*.
- Vestner, M., Lähner, Z., Boyarski, A., Litany, O., Slossberg, R., Remez, T., Rodolà, E., Bronstein, A., Bronstein, M., Kimmel, R., & Cremner, D. (2017). Efficient deformable shape correspondence via kernel matching. In *2017 international conference on 3D vision (3DV)* (pp. 517–526). IEEE.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.