

The Average Mixing Kernel Signature

Luca Cosmo^{*1,2}, Giorgia Minello^{*3},
Michael Bronstein^{2,4,5}, Luca Rossi⁶, and Andrea Torsello³

¹ Sapienza University of Rome, Italy

² University of Lugano, Switzerland

³ Università Ca' Foscari Venezia, Italy

⁴ Imperial College London, United Kingdom

⁵ Twitter, United Kingdom

⁶ Queen Mary University of London, United Kingdom

luca.rossi@qmul.ac.uk

Abstract. We introduce the Average Mixing Kernel Signature (AMKS), a novel signature for points on non-rigid three-dimensional shapes based on the average mixing kernel and continuous-time quantum walks. The average mixing kernel holds information on the average transition probabilities of a quantum walk between each pair of vertices of the mesh until a time T . We define the AMKS by decomposing the spectral contributions of the kernel into several bands, allowing us to limit the influence of noise-dominated high-frequency components and obtain a more descriptive signature. We also show through a perturbation theory analysis of the kernel that choosing a finite stopping time T leads to noise and deformation robustness for the AMKS. We perform an extensive experimental evaluation on two widely used shape matching datasets under varying level of noise, showing that the AMKS outperforms two state-of-the-art descriptors, namely the Heat Kernel Signature (HKS) and the similarly quantum-walk based Wave Kernel Signature (WKS).

Keywords: Shape representation; Shape analysis; Quantum walks.

1 Introduction

A central ingredient in shape analysis, with applications that range from shape correspondence [14, 10, 9, 24, 25, 35] to shape segmentation [1, 16], is the computation of a signature (also known as descriptor or feature descriptor) that effectively characterizes the local and global geometric information around each point on a shape's surface. In this context, an effective descriptor is one that accurately characterizes each point while remaining invariant to rigid transformations, stable against non-rigid ones, and robust to various sources of noise.

Despite a flurry of descriptors introduced in the past years, spectral-based descriptors remain among the most popular type of descriptors [31, 33, 7, 2, 5]. These rely on the analysis of the eigenvalues and eigenfunctions of the Laplace-Beltrami operator. Examples include two widely used state-of-the-art spectral

* Equal contribution.




Fig. 1. The average transition probability of a quantum walk from one point to the rest of the shape for increasing values of the stopping time T (left to right). More (less) intense shades of red correspond to higher (lower) probability. This is the information in the row of the average mixing matrix (Eq. 9) 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.

signatures, the Heat Kernel Signature (HKS) [33] and the Wave Kernel Signature (WKS) [2]. The common idea underpinning these signatures is that of filtering the frequency information, where low and high frequencies capture global and local topological information respectively, with the latter being more discriminative but also easily affected by noise.

The HKS and the WKS are based on analyzing the shape through similarly defined yet profoundly different physical processes, namely heat diffusion and wavefunction evolution. On a discrete domain these correspond to continuous-time random walks and continuous-time quantum walks [17], respectively. Both processes are similar in that they explore the vertices and edges of the mesh representing the discrete approximation of the shape surface. However apparently subtle differences in their mathematical formulation lead to significantly different behaviors, with the emergence of quantum effects resulting in a superior ability to characterize the underlying structure with respect to classical walks. This is not only demonstrated in shape analysis by the WKS being more discriminative than the HKS for shape matching, but also in a variety of other domains, including graph theory [12, 29] and machine learning [27, 28, 3, 22].

In this paper, we propose a novel signature for points on non-rigid three-dimensional shapes based on continuous-time quantum walks. Given a mesh representing the discrete approximation of a shape surface, we define a continuous-time quantum walk over the vertices and edges of the mesh and we let the walker explore the structure until a stopping time T is reached. We compute the average mixing kernel holding the average transition probabilities between each pair of vertices of the mesh. This is the continuous version of the average mixing matrix of Godsil [15] and it's effectively the quantum analogue of the heat kernel underpinning the signature of Sun et al. [33]. We show using perturbation theory that the choice of a finite time T (as opposed to letting the quantum particle explore the shape surface for an infinite time [2]) is crucial to achieve increased robustness against noise and deformation. We then propose to decompose the spectral contributions to the average mixing kernel into several bands. This in turn has the effect of both limiting the influence of noise-dominated high-frequency com-

ponents as well as creating a more descriptive signature as a function of the energy levels. Finally, we use the diagonal information of this band-pass filtered kernel sampled at different temporal evolution steps to compute the proposed Average Mixing Kernel Signature (AMKS).

An extensive set of experiments shows that the AMKS outperforms state-of-the-art signatures like the HKS and the WKS. Our key findings include:

- when compared to the WKS, we theoretically show that our method increases the robustness to noise by accounting for the noise-induced mixing of the Laplace-Beltrami operator eigenspaces (see the Supplementary Material);
- we show that under Gaussian noise ($\sigma = 0.02$) the proposed AMKS shows an improvement of up to 10% wrt to the WKS, in terms of hit rate;
- when computing the point-wise correspondences between pairs of shapes, we observe an improvement of the geodesic error of up to 4% wrt to the WKS;
- despite the additional need to compute the mixing components wrt the WKS, the AMKS can be fully written in terms of matrix operations and efficiently computed (see the Supplementary Material);
- the computational complexity of the AMKS is linear in the number of vertices of the mesh and energy levels, and quadratic in the dimension of the truncated eigenbase, resulting in an overall computational time that is comparable with that of the WKS.

In [26] Rossi et al. introduced another structural signature based on the average mixing matrix (AMMS) for the analysis of graphs. However, we stress that our work significantly differs from [26] in several key aspects:

- the AMMS is defined on graphs of limited size with long-range connections where using all the spectral range is important. The AMKS is defined on manifold structures where structural information is found at low bandwidths while high bandwidths are dominated by noise. Directly applying the AMMS to our problem would lead to a computationally infeasible, noise-dominated signature where we cannot focus on the energy levels of interest. One way of reducing the computational cost of the AMMS would have been truncating the spectrum, however this would not have led to the quantum-mechanically feasible process one which our framework is based;
- in the AMMS the sampling is done over the time domain, as opposed to the energy domain in the AMKS;
- while a limited empirical analysis suggests that the AMMS is robust to structural noise, we perform the first theoretical analysis of the noise-robustness resulting from the average mixing matrix use, revealing that the mixing factors naturally account for the noise-induced mixing of eigenspaces;
- while there is no analysis of the computational complexity of the AMMS in [26], here we perform a thorough complexity as well as runtime analysis of the AMKS in addition to showing that its computation can be fully expressed in terms of matrix operations;
- finally, as the application domains and communities of interest of the AMMS and AMKS are different, in this paper we perform an entirely separate and more extensive set of experiments.

2 Related work

Spectral signatures aim to characterize a shape by analyzing the eigenvalues and eigenfunctions of the Laplace-Beltrami operator [23, 19]. Based on this key idea, Rustamov [31] introduced the Global Point Signature (GPS), where each point of a mesh is characterized by a vector computed from the scaled eigenfunctions of the Laplace-Beltrami operator evaluated at that point.

While GPS is a global signature and thus lacks the ability to perform partial matching, the Heat Kernel Signature (HKS) of Sun et al. [33] and its scale-invariant version [7] addressed this shortcoming by allowing to capture multi-scale information around each point of a shape. This in turn is achieved by analyzing a heat diffusion process over the shape surface at different time-scales. From a signal processing point of view, the HKS can be seen as performing a low-pass filtering of the spectral information parametrized by the diffusion time t . As a result, its value is dominated by global shape properties, while it fails to accurately capture the highly discriminative small-scale information that is crucial in matching applications.

The Wave Kernel Signature (WKS) [2], on the other hand, solves this issue by acting as a band-pass filter that better separates the frequency information related to different spatial scales. To this end, the authors define a quantum system where a particle explores the shape surface under the Schrödinger equation. This is intimately connected to our approach, however, as explained in Section 4 and in the Supplementary Material, there are some stark differences between the two methods which result in a more accurate and noise resilient descriptor.

Learning-based signatures have recently emerged as alternatives to spectral-based signatures, although the latter remain widely used due to their efficacy and ease of computation. Litman and Bronstein [20] introduced a family of spectral descriptors that generalizes the HKS and the WKS and they propose a learning scheme to construct the optimal descriptor for a given task. Corman et al. [8] 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 constraint the degree of deformation between corresponding points.

In [13] Fang et al. 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. [21] extend 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. [16], on the other hand, omit this pre-processing step by designing a 2D convolutional neural network (CNN) architecture where each point on the shape surface is rendered using multi-scale 2D images.

Presenting a comprehensive review of spectral- and learning-based signatures is beyond the scope of this paper, so we refer the reader to the spectral-based signatures review of Masci et al. [21] and the survey of Rostami et al. on data-driven signatures [30]. Finally, we would like to stress that while learning-based descriptors are gaining increasing popularity in the literature, particularly due to

the success of CNNs in computer vision, spectral-based signatures are still often used as the first processing step in a modern (deep) learning framework [16, 30].

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 walks and the Schrödinger equation

Let M denote the manifold representing the shape surface. The state of a quantum particle at time t is given by the wave function $\psi(t)$, defined as an element of the Hilbert space of functions mapping points on the surface of M to \mathbb{C} , such that the squared modulus $|\psi(t)|^2$ is a real number and the integral of this quantity over M is unity. In other words, the squared modulus of the wave function is interpreted as the probability of the quantum walker being located at a certain point on M at time t .

The Dirac notation (also known as *bra-ket* notation) is a standard notation used for describing quantum states, where a *ket* vector $|a\rangle$ denotes a pure quantum state and is a complex valued column vector of unit Euclidean length, in a n -dimensional 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 $|a\rangle\langle b|$. With this notation to hand, on a discrete domain, the state of the walker at time t is denoted as

$$|\psi(t)\rangle = \sum_{u \in V} \alpha_u(t) |u\rangle, \quad (1)$$

where $|u\rangle$ denotes the state corresponding to the vertex u and $\alpha_u(t) \in \mathbb{C}$ is its amplitude. Moreover, we have that $\alpha_u(t)\alpha_u^*(t)$ is 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 $u \in V$, $t \in \mathbb{R}^+$ (where $\alpha_u^*(t)$ is the complex conjugate of $\alpha_u(t)$).

The evolution of the walker is governed by the Schrödinger equation

$$\frac{\partial}{\partial t} |\psi(t)\rangle = -i\Delta_M |\psi(t)\rangle, \quad (2)$$

where the Laplace-Beltrami operator plays the role of a time-independent Hamiltonian, i.e., the operator that accounts for the total energy of the system. The solution to this equation is given by

$$|\psi(t)\rangle = U(t) |\psi(0)\rangle, \quad (3)$$

where $U(t) = e^{-it\Delta_M}$ is a unitary operator. Similarly to the heat operator, $U(t)$ shares the same eigenvectors of Δ_M and for each eigenvalue λ of Δ_M there exists a corresponding eigenvalue of $U(t)$.

Measuring quantities in quantum mechanics 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. In this context, observing a state corresponds to performing a projective measurement of $|\psi\rangle$. Given the observable O and its spectral decomposition $O = \sum_{\lambda} \lambda P_{\lambda}$, where P_{λ} is the projector on the subspace spanned by the eigenvalue λ , the outcome λ is observed with probability $p(\lambda) = \langle \psi | P_{\lambda} | \psi \rangle$. Notably, the state of the quantum system changes after the measurement and becomes $|\bar{\psi}\rangle = P_{\lambda} |\psi\rangle / \|P_{\lambda} |\psi\rangle\|$ where $\| |\psi\rangle \| = \sqrt{\langle \psi | \psi \rangle}$ denotes the norm of the vector $|\psi\rangle$.

3.2 Average mixing matrix

In [15] Godsil 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

$$U(t) = \sum_{\lambda \in \Lambda} e^{-i\lambda t} P_{\lambda}, \quad (4)$$

where Λ is the set of unique eigenvalues. We can then write the mixing kernel in terms of the Schur-Hadamard product (also known as element-wise product) of the projectors

$$M_G(t) = \sum_{\lambda_1 \in \Lambda} \sum_{\lambda_2 \in \Lambda} e^{-i(\lambda_1 - \lambda_2)t} P_{\lambda_1} \circ P_{\lambda_2}. \quad (5)$$

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 can enforce convergence by taking a time-average even if $U(t)$ is norm-preserving. Let us define the average mixing kernel [15] at time T as

$$\widehat{M}_{G;T} = \frac{1}{T} \int_0^T M_G(t) dt = \sum_{\lambda_1 \in \Lambda} \sum_{\lambda_2 \in \Lambda} P_{\lambda_1} \circ P_{\lambda_2} \frac{1}{T} \int_0^T e^{-i(\lambda_1 - \lambda_2)t} dt, \quad (6)$$

which has solution

$$\widehat{M}_{G;T} = \sum_{\lambda_1 \in \Lambda} \sum_{\lambda_2 \in \Lambda} P_{\lambda_1} \circ P_{\lambda_2} \frac{i(1 - e^{iT(\lambda_2 - \lambda_1)})}{T(\lambda_2 - \lambda_1)}. \quad (7)$$

In the limit $T \rightarrow \infty$, Eq. 7 becomes

$$\widehat{M}_{G;\infty} = \sum_{\lambda \in \Lambda} P_{\lambda} \circ P_{\lambda}, \quad (8)$$

where $\tilde{\Lambda}$ is the set of distinct eigenvalues of the Hamiltonian.

Finally, we note that the average mixing kernel can be rewritten as

$$\begin{aligned}
 \widehat{M}_{G;T} &= \sum_{\lambda_1 \in \Lambda} \sum_{\lambda_2 \in \Lambda} P_{\lambda_1} \circ P_{\lambda_2} \frac{i(1 - e^{iT(\lambda_2 - \lambda_1)})}{T(\lambda_2 - \lambda_1)} \\
 &= \sum_{\lambda_1 \in \Lambda} \sum_{\lambda_2 \in \Lambda} P_{\lambda_1} \circ P_{\lambda_2} \frac{i - i \cos(T(\lambda_2 - \lambda_1)) + \sin(T(\lambda_2 - \lambda_1))}{T(\lambda_2 - \lambda_1)} \\
 &= \sum_{\lambda_1 \in \Lambda} \sum_{\lambda_2 \in \Lambda} P_{\lambda_1} \circ P_{\lambda_2} \text{sinc}(T(\lambda_2 - \lambda_1)), \tag{9}
 \end{aligned}$$

where $\text{sinc}(x) = \frac{\sin(x)}{x}$ is the unnormalized sinc function. In Section 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 (Eq. 9) instead of an infinite time one (Eq. 8). Fig. 1 shows the entries of the average mixing matrix computed according to Eq. 9 on a mesh of n vertices sampled from an underlying shape and 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 on 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. [2]. 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

$$\text{AMM}(E) = \sum_{\lambda_1} \sum_{\lambda_2} P_{\lambda_1} \circ P_{\lambda_2} \text{sinc}(T(\lambda_2 - \lambda_1)) f_E(\lambda_1) f_E(\lambda_2), \tag{10}$$

where, following Aubry et al. [2], $f_E(\lambda) = e^{-\frac{(E - \log \lambda)^2}{2\sigma^2}}$, 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 .

Table 1. Average percentage of correspondences \pm standard error on the nearest 0.1% neighbours in the descriptor space for the AMKS (for increasing stopping time T , columns 1 to 13) and the WKS (last column) over a subset of the TOSCA (top) and FAUST (bottom) datasets with added Gaussian noise (η) to the vertex positions (rows). The highest value for each row is highlighted in bold.

$\eta \backslash T$	0.005	0.01	0.2	0.5	0.8	1	1.2	1.5	1.8	2	2.5	1000	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 2.0	34.7 \pm 2.0	34.2 \pm 2.0	33.7 \pm 1.9	32.6 \pm 1.9	27.8 \pm 1.6	33.0 \pm 1.6
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	21.9 \pm 1.2	24.9 \pm 1.2
0.015	20.1 \pm 1.1	20.1 \pm 1.1	20.5 \pm 1.1	20.5 \pm 1.0	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
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	11.8 \pm 0.6	11.8 \pm 0.6

$\eta \backslash T$	0.005	0.01	0.2	0.5	0.8	1	1.2	1.5	1.8	2	2.5	1000	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 2.5	51.4 \pm 2.4	50.8 \pm 2.4	50.3 \pm 2.4	49.2 \pm 2.5	40.7 \pm 2.5	47.0 \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	24.3 \pm 1.0	29.5 \pm 1.0
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	15.0 \pm 0.8	18.1 \pm 0.8
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 0.4	10.0 \pm 0.4	8.8 \pm 0.4	10.5 \pm 0.4

After annihilation, the ensemble at time t of the particle starting from j is

$$\begin{aligned}
\sum_{\lambda_1} \sum_{\lambda_2} e^{-i(\lambda_1 - \lambda_2)t} & \left(I - \sum_{\mu_1} (1 - f_E(\mu_1)) P_{\mu_1} \right) \left(P_{\lambda_1} |j\rangle \langle j| P_{\lambda_2} \right) \\
& \left(I - \sum_{\mu_2} (1 - f_E(\mu_2)) P_{\mu_2} \right) = \\
& \sum_{\lambda_1} \sum_{\lambda_2} e^{-i(\lambda_1 - \lambda_2)t} f_E(\lambda_1) f_E(\lambda_2) \left(P_{\lambda_1} |j\rangle \langle j| P_{\lambda_2} \right), \quad (11)
\end{aligned}$$

resulting, after time integration, in the band-filtered average mixing matrix (Eq. 10). To obtain the descriptor, we take the self-propagation, i.e., the probability that a particle in state u is still observed in u after the evolution, as a function of the energy level E and normalized so that the integral over all energies is 1, i.e.,

$$AMKS_u(E) = \frac{AMM(E)_{uu}}{\sum_{\lambda_1, \lambda_2} f_E(\lambda_1) f_E(\lambda_2)}. \quad (12)$$

There are some 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 eigenvectors of the Laplace-Beltrami operator, our signature results in their fourth power. This in turn creates a more distinctive




Fig. 2. Quantitative comparison of the match performance of different descriptors (HKS, WKS, and AMKS) on TOSCA (top) and FAUST (bottom) using the CMC curve \pm standard error, considering an increasing percentage of nearest points in the descriptor space (x -axis) and with increasing level of Gaussian noise (left to right).

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 in turn can be shown to increase the robustness to noise as it imposes a mixing of the eigenspaces at a rate similar to the effect of noise (see the Supplementary Material).

5 Experiments

In this section we present quantitative and qualitative analysis testing the performance of our descriptor with respect to two alternative state-of-the-art spectral signatures, the WKS [2] and the scaled HKS [33], as these represent two of the most successful and used non-learned descriptors for deformable shapes. We omit the comparison with the AMMS [26] due to its performance in this domain. Indeed, the AMMS was introduced to work on graphs of limited size and with long-range connections and, as we observed, its poor performance in the shape domain is further exacerbated by the need to truncate the spectrum.

Datasets We evaluated our descriptor on two widely used benchmarks, TOSCA [6] and FAUST [4]. TOSCA comprises 8 humanoid and animals shape classes, for a total of 68 meshes of varying resolution (3K to 50K vertices). FAUST (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




Fig. 3. 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.

total of 100 meshes. In the next experiments, we refer to the synthetic version, unless specified otherwise. In both datasets, all shapes have been remeshed to 5K vertices by iterative pair contractions and re-scaled to unit area.

Setting 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 [33] and [2], respectively. It should be noted that 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 (we refer the reader to [2]). The time parameter is extensively discussed in the next subsection, where we show that the AMKS is not particularly sensitive to its value provided that it sits within a certain range.

Distance We adopted the distance proposed by Aubry et al. in [2], here denoted as L_{KS}^1 . Given two shapes \mathcal{X} and \mathcal{Y} , in order to compare the descriptors for two points $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ we measure their distance as

$$L_{KS}^1(x, y) = \sum_i^d \left| \frac{KS(x)_i - KS(y)_i}{KS(x)_i + KS(y)_i} \right|, \quad (13)$$

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

Metrics We used two measures to evaluate the signatures. First, we made use of the Cumulative Match Characteristic (CMC) measure to estimate the probability of finding a correct correspondence among the p -nearest neighbours in the descriptor space. The probability (hit rate), is calculated as the percentage of correct matches in the p -nearest neighbours. We average the hit rate for all points across all pairs of shapes at increasing values of p . The resulting CMC curve is then a monotonically increasing function of p .

The second measure is the Princeton protocol [18], also known as correspondence accuracy, which captures the proximity of predicted corresponding points




Fig. 4. The two figures show the 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 behaviour of our AMKS at increasing noise levels, compared with WKS. The *right* 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 behaviour of two signatures: the green curve belongs to the shape in reference pose whereas the red curve represents the sitting shape.

to ground-truth ones in terms of geodesic distance on the surface. It counts 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 over 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. In particular, given a match $(x, y) \in \mathcal{X} \times \mathcal{Y}$, and the ground-truth correspondence (x, y^*) , the normalized geodesic error is $\epsilon(x) = d_{\mathcal{Y}}(y, y^*) / \text{area}(\mathcal{Y})^{0.5}$.

5.1 Choice of time

In order to show how the choice of the diffusion time affects the performance of our descriptor, we performed some tests over two subsets of shapes from the two datasets. Table 1 shows the hit ratio of our descriptor on the nearest 0.1% of points in descriptor space. In particular, we explore how the performance changes as noise and time increase. In addition, we compare the results to those obtained with the WKS, noting that the WKS is time-independent. For the TOSCA dataset, we used 7 classes (wolf excluded) for a total of 28 pairs, whereas for FAUST we used 10 pairs. A pair consists of a 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 superior to that of the WKS and b) the optimal times are different in the two datasets. More specifically, in the FAUST dataset the optimal performance of the AMKS 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 (see the perturbation analysis in the Supplementary




Fig. 5. Qualitative comparison between WKS and AMKS on 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). The AMKS improves the overall matching quality.

Material for a discussion of the link between the stopping time T and the noise-induced mixing of the eigenspaces). 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.

5.2 Quantitative Evaluation

We now compare our signature with that of the WKS and the HKS in the context of shape matching following the CMC and Princeton protocols.

Descriptor robustness The first quantitative comparison of the matching performance is shown in Fig. 2. Each plot shows the CMC curve of our method compared to those of the other spectral signatures, for varying levels of Gaussian noise (left to right, $\eta = 0, 0.01, 0.02$) and $T = 0.5$ (time of our signature). Here the y -axis is the hit rate (note that the plots were drawn by averaging the hit rate of all pairs of shapes for each dataset) and the x -axis defines the p -nearest neighbours (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 in turn shows the noise robustness of our signature.

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 neighbour in descriptor space. In Fig. 3 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.

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. We used our descriptor as input to one of the state-of-the-art architectures for deep learning on shapes, FeastNet [34]. We used the code provided by authors, with the translation invariant single-scale network architecture proposed in the paper. The first 80 FAUST shapes have been used for training, the remaining 20 for test.

Table 2 shows the accuracy after training the network with different descriptors as input. We compare against WKS, the local rigid descriptor SHOT [32], and using just XYZ coordinates of points as input. We can see how the network is able to exploit the higher descriptivity of AMKS in order to achieve better classification accuracy. Note that, while training directly on XYZ coordinates allows to achieve the same accuracy, the performance rapidly decreases if we allow shapes to undergo a random rotation (i.e., $\text{ROT}(\text{XYZ})$) since neither the XYZ coordinates nor the network are intrinsically invariant to this transformation.

Table 2. Classification accuracy of FeastNet trained with different input descriptors.

	XYZ	ROT(XYZ)	SHOT	WKS	AMKS
FeastNet	86%	50%	63%	80%	86%

5.3 Qualitative Evaluation

On the left side of Fig. 4 we contrast the behaviour 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. We employ the same type of visualization in the right side of Fig. 4. Here we can observe the behaviour of AMKS 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.

Finally, in Fig. 5 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 coloured spots for the AMKS suggests a better matching quality.

5.4 Runtime analysis

The computational complexity of our descriptor, without considering the spectral decomposition, is $\mathcal{O}(dnk^2)$ with n being the number of points of the mesh,




Fig. 6. Running times comparison wrt number of points (left) and eigenvectors (right).

k the dimension of the truncated eigenbase, and d the number of energy levels (see the Supplementary Material)⁷. 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. In Fig. 6 we show 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 complexity remains linear in the number of points. On the other hand, we introduce a quadratic complexity with respect to the number of eigenvalues. Note, however, that in the range of number of eigenvectors normally used in spectral shape retrieval and matching tasks (< 200) the behaviour is almost linear and that for a higher number of eigenvectors the main bottleneck is the spectral decomposition of the Laplace-Beltrami operator.

6 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, outperforming both the HKS and the WKS on a shape matching task. Future work will look into the application of our signature on different problems, including partial correspondences [11, 25], point clouds, and shape segmentation.

Acknowledgements Luca Cosmo was supported by the ERC Starting Grant No. 802554 (SPECGEO) and the ERC Consolidator grant No. 724228 (LEMAN).

⁷ <https://github.com/lcosmo/amks-descriptor>.

References

1. Aubry, M., Schlickewei, U., Cremers, D.: Pose-consistent 3d shape segmentation based on a quantum mechanical feature descriptor. In: Joint Pattern Recognition Symposium. pp. 122–131. Springer (2011)
2. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on. pp. 1626–1633. IEEE (2011)
3. Bai, L., Rossi, L., Torsello, A., Hancock, E.R.: A quantum jensen–shannon graph kernel for unattributed graphs. *Pattern Recognition* **48**(2), 344–355 (2015)
4. Bogo, F., Romero, J., Loper, M., Black, M.J.: Faust: Dataset and evaluation for 3d mesh registration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3794–3801 (2014)
5. Boscaini, D., Masci, J., Rodolà, E., Bronstein, M.M., Cremers, D.: Anisotropic diffusion descriptors. In: Computer Graphics Forum. vol. 35, pp. 431–441. Wiley Online Library (2016)
6. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Numerical geometry of non-rigid shapes. Springer Science & Business Media (2008)
7. Bronstein, M.M., Kokkinos, I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 1704–1711. IEEE (2010)
8. Corman, É., Ovsjanikov, M., Chambolle, A.: Supervised descriptor learning for non-rigid shape matching. In: European Conference on Computer Vision. pp. 283–298. Springer (2014)
9. Cosmo, L., Rodola, E., Albarelli, A., Mémoli, F., Cremers, D.: Consistent partial matching of shape collections via sparse modeling. In: Computer Graphics Forum. vol. 36, pp. 209–221. Wiley Online Library (2017)
10. Cosmo, L., Rodolà, E., Bronstein, M.M., Torsello, A., Cremers, D., Sahillioglu, Y.: Shrec’16: Partial matching of deformable shapes. *Proc. 3DOR* **2**(9), 12 (2016)
11. Cosmo, L., Rodola, E., Masci, J., Torsello, A., Bronstein, M.M.: Matching deformable objects in clutter. In: 2016 Fourth International Conference on 3D Vision (3DV). pp. 1–10. IEEE (2016)
12. Emms, D., Severini, S., Wilson, R.C., Hancock, E.R.: Coined quantum walks lift the cospectrality of graphs and trees. *Pattern Recognition* **42**(9), 1988–2002 (2009)
13. Fang, Y., Xie, J., Dai, G., Wang, M., Zhu, F., Xu, T., Wong, E.: 3d deep shape descriptor. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2319–2328 (2015)
14. Gasparetto, A., Minello, G., Torsello, A.: Non-parametric spectral model for shape retrieval. In: 2015 International Conference on 3D Vision. pp. 344–352. IEEE (2015)
15. Godsil, C.: Average mixing of continuous quantum walks. *Journal of Combinatorial Theory, Series A* **120**(7), 1649–1662 (2013)
16. Huang, H., Kalogerakis, E., Chaudhuri, S., Ceylan, D., Kim, V.G., Yumer, E.: Learning local shape descriptors from part correspondences with multiview convolutional networks. *ACM Transactions on Graphics (TOG)* **37**(1), 6 (2018)
17. Kempe, J.: Quantum random walks: an introductory overview. *Contemporary Physics* **44**(4), 307–327 (2003)
18. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. In: *ACM Transactions on Graphics (TOG)*. vol. 30, p. 79. ACM (2011)
19. Levy, B.: Laplace-beltrami eigenfunctions towards an algorithm that “understands” geometry. In: IEEE International Conference on Shape Modeling and Applications 2006 (SMI’06). pp. 13–13. IEEE (2006)

20. Litman, R., Bronstein, A.M.: Learning spectral descriptors for deformable shape correspondence. *IEEE transactions on pattern analysis and machine intelligence* **36**(1), 171–180 (2013)
21. Masci, J., Boscaini, D., Bronstein, M., Vandergheynst, P.: Geodesic convolutional neural networks on riemannian manifolds. In: *Proceedings of the IEEE international conference on computer vision workshops*. pp. 37–45 (2015)
22. Minello, G., Rossi, L., Torsello, A.: Can a quantum walk tell which is which? a study of quantum walk-based graph similarity. *Entropy* **21**(3), 328 (2019)
23. Reuter, M., Wolter, F.E., Peinecke, N.: Laplace-spectra as fingerprints for shape matching. In: *Proceedings of the 2005 ACM symposium on Solid and physical modeling*. pp. 101–106. ACM (2005)
24. Rodola, E., Cosmo, L., Litany, O., Bronstein, M., Bronstein, A., Audebert, N., Hamza, A.B., Boulch, A., Castellani, U., Do16, M., et al.: Shrec’17: Deformable shape retrieval with missing parts. In: *Proceedings of the Eurographics Workshop on 3D Object Retrieval, Lisbon, Portugal*. pp. 23–24 (2017)
25. Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. In: *Computer Graphics Forum*. vol. 36, pp. 222–236. Wiley Online Library (2017)
26. Rossi, L., Severini, S., Torsello, A.: 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 (2016)
27. Rossi, L., Torsello, A., Hancock, E.R.: A continuous-time quantum walk kernel for unattributed graphs. In: *International Workshop on Graph-Based Representations in Pattern Recognition*. pp. 101–110. Springer (2013)
28. Rossi, L., Torsello, A., Hancock, E.R.: Measuring graph similarity through continuous-time quantum walks and the quantum jensen-shannon divergence. *Physical Review E* **91**(2), 022815 (2015)
29. Rossi, L., Torsello, A., Hancock, E.R., Wilson, R.C.: Characterizing graph symmetries through quantum jensen-shannon divergence. *Physical Review E* **88**(3), 032806 (2013)
30. Rostami, R., Bashiri, F.S., Rostami, B., Yu, Z.: A survey on data-driven 3d shape descriptors. In: *Computer Graphics Forum*. vol. 38, pp. 356–393. Wiley Online Library (2019)
31. Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: *Proceedings of the fifth Eurographics symposium on Geometry processing*. pp. 225–233. Eurographics Association (2007)
32. Salti, S., Tombari, F., di Stefano, L.: Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding* **125**, 251–264 (2014)
33. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: *Computer graphics forum*. vol. 28, pp. 1383–1392. Wiley Online Library (2009)
34. Verma, N., Boyer, E., Verbeek, J.: Feastnet: Feature-steered graph convolutions for 3d shape analysis. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018)
35. Vestner, M., Löhner, Z., Boyarski, A., Litany, O., Slossberg, R., Remez, T., Rodola, E., Bronstein, A., Bronstein, M., Kimmel, R., et al.: Efficient deformable shape correspondence via kernel matching. In: *2017 International Conference on 3D Vision (3DV)*. pp. 517–526. IEEE (2017)