





Orthogonalized Fourier Polynomials for Signal Approximation and Transfer

F. Maggioli¹ , S. Melzi¹ , M. Ovsjanikov² , M. M. Bronstein^{3,4,5}  and E. Rodolà¹

¹Sapienza University of Rome, Italy ²LIX, Ecole Polytechnique, IP Paris, France

³Imperial College London, United Kingdom ⁴Twitter, United Kingdom ⁵USI Lugano, Switzerland




Figure 1: Surface approximation using orthogonalized polynomials of increasing order, where order 1 corresponds to the plain Laplacian eigenbasis. All of these approximations are obtained starting from just 5 eigenfunctions. Reconstruction error is encoded by color, growing from white to dark red.

Abstract

We propose a novel approach for the approximation and transfer of signals across 3D shapes. The proposed solution is based on taking pointwise polynomials of the Fourier-like Laplacian eigenbasis, which provides a compact and expressive representation for general signals defined on the surface. Key to our approach is the construction of a new orthonormal basis upon the set of these linearly dependent polynomials. We analyze the properties of this representation, and further provide a complete analysis of the involved parameters. Our technique results in accurate approximation and transfer of various families of signals between near-isometric and non-isometric shapes, even under poor initialization. Our experiments, showcased on a selection of downstream tasks such as filtering and detail transfer, show that our method is more robust to discretization artifacts, deformation and noise as compared to alternative approaches.

CCS Concepts

• **Computing methodologies** → *Shape analysis*; • **Theory of computation** → *Computational geometry*; • **Mathematics of computing** → *Functional analysis*;

1. Introduction and related work

Approximation and transfer of signals between shapes are among the most widely explored tasks in computer vision and graphics, and are at the basis of numerous applications. Common to most approaches is the idea to encode the given surface signal in a basis that allows to represent and transfer it efficiently; among these, approaches based upon the construction of a Fourier-like basis (or rather its surface counterpart [Tau95]) play the lion's share

[Lév06, Rus07, VL08]. The key idea is to project the signal onto a low-dimensional function space, e.g., corresponding to the lowest portion of the frequency band. This yields a well known trade-off between the compactness of the representation and its approximation quality. Efforts have been devoted to strike a balance between these two factors, by resorting to alternative bases or via costly post-processing steps; still, the search for a *compact* basis for representing high level of detail, is an unsolved problem to date.




Figure 2: RGB signal transfer among two non-isometric shapes with different mesh topology. Our approach (ours, ours*) better transfers the surface signal from bison to cow, while using the same amount of information as other existing approaches (eigs [OBCS*12] and prods [NMR*18]). We refer to the experimental section for more details.

A closely related problem to that of signal representation is the need to *transfer* these signals from a source to a target domain. This can often be cast as correspondence problem, where the objective is to find a transformation that acts as a bridge between source and target. This was shown in [OBCS*12] to be equivalent to seeking a coherent set of basis functions for the given pair of shapes; the search for a correspondence is then phrased as the search for a linear map (called functional map) that aligns the basis functions on the source to those on the target. Follow-up works have embraced this view by introducing more stable ways to compute the functional map [NO17, OCB*16, EBC17, RPWO18], by extending the framework to the partial setting [RCB*17, CRM*16], or by constructing new coherent bases explicitly as linear transformations of the Laplacian eigenfunctions [KBB*13, AL19, LRBB17].

The choice of Laplacian eigenfunctions as a reduced basis for representing surface signals is due to their optimality for continuous functions with bounded variation [ABK15]. However, in many real applications such as texture transfer and shape interpolation, this band-limited representation may not provide the necessary accuracy for capturing fine details. To overcome these limitations, two main solutions have been proposed: (i) to design an ad-hoc basis for fixed sets of signals; (ii) to define algorithms for recovering the residual information that is lost in the representation. The former includes wavelets as a localized alternative to the Fourier basis [Zho12, CM06, LDW97, HVG11, Pat18, KMP*20]. Other local constructions include those based on sparse regularization and Hamiltonian operators with step potentials [NVT*14, KGB16, CSBK17, MRCB18]. Specialized bases for piecewise-constant signals and vertex coordinates have been proposed in [Mel19] and [MMM*20] respectively, but these do not generalize well to different function classes. Point (ii) is a more recent trend [MRR*19, ELC20]. The idea is to iteratively seek for bases of increasing dimension starting from an initial alignment between few Laplacian eigenfunctions. The iterative procedure preserves the alignment of the two bases as they increase in dimension, and sidesteps the need for further optimization to get an optimal alignment.

More closely related to ours is the work of Nogneng et al. [NMR*18], where the authors consider the set of pointwise products of the Laplacian eigenfunctions, in addition to the eigenfunctions alone, for representing surface signals more accurately. The main property of these *eigenproducts* is that their alignment can be explicitly and directly derived from the functional map between the standard eigenfunctions; this way, a correct alignment between a few eigenfunctions is automatically extended to the larger set, which includes their products.

Contribution. Our work addresses a key issue of the latter representation, namely that the set containing Laplacian eigenfunctions and eigenproducts is *not* linearly independent in general; as we show in the sequel, it is linearly independent only when very few eigenfunctions are involved. Thus, this set does not provide a unique representation for surface signals. Further, the linear dependence gives rise to instability in the transfer task, which must be handled through additional constraints and pre-processing as shown in [NMR*18]; see Figure 2 for an example. Here we follow a similar idea and use eigenproducts to increase the dimensionality of the basis, and in turn, the quality of the resulting representation. However, differently from [NMR*18], we do not limit our analysis to products of order 2, but we effectively exploit the entire set of “Fourier polynomials” with arbitrary order.

This paper fills the gaps left by [NMR*18] in several ways:

- For the first time, we provide a theoretical analysis on the space spanned by the eigenproducts, including a discussion on the frequency range that they capture;
- We propose the construction of an orthonormal basis on top of the linearly *dependent* set of polynomials, yielding a simpler, more accurate, stable and computationally efficient technique;
- We extend the discussion and empirical evaluation to eigenproducts of order greater than 2.

Our basis applies to several applications, such as detail transfer and spectral filtering, that are impossible to target through the representation proposed in [NMR*18] as we show in the experiments.

2. Background

We model a shapes as a 2-dimensional Riemannian manifolds \mathcal{M} , equipped with the metric tensor g . The positive semi-definite Laplace-Beltrami operator Δ , generalizing the notion of Laplacian to surfaces, can be expressed in local coordinates as (here in the general d -dimensional case):

$$\Delta f = -\frac{1}{\sqrt{\det(g)}} \sum_{i,j=1}^d \frac{\partial}{\partial x_i} \sqrt{\det(g)} g^{i,j} \frac{\partial f}{\partial x_j}, \quad (1)$$

where $g^{i,j} = (g^{-1})_{i,j}$ is the cell (i, j) of the inverse matrix of g .

The operator Δ admits a spectral decomposition into eigenvalues λ_i and associated eigenfunctions ϕ_i . The eigenvalues are non-negative and countably infinite; in this paper, we always assume that the eigenvalues are ordered non-decreasingly, and the associated eigenfunctions follow the same ordering. A standard example is the 1-dimensional manifold representing the real line $[0, T]$. Here, $\Delta = -\frac{\partial^2}{\partial x^2}$, the eigenfunctions are $\phi_k(x) = \sin\left(\frac{2\pi kx}{T}\right)$ and the associated eigenvalues are $\lambda_k = \frac{4\pi^2 k^2}{T^2}$.

Laplacian eigenfunctions span the space $L^2(\mathcal{M})$ of square integrable functions, i.e. the set of functions $f: \mathcal{M} \rightarrow \mathbb{R}$ such that:

$$\langle f, f \rangle_{\mathcal{M}} = \|f\|_{\mathcal{M}}^2 = \int_{\mathcal{M}} f^2(x) dx < \infty. \quad (2)$$

The Laplace-Beltrami operator obeys the Leibniz product rule

with a correction term involving gradients [Cha84]. Namely,

$$\Delta f(x)g(x) = f(x)\Delta g(x) + g(x)\Delta f(x) - 2 \langle \nabla f(x), \nabla g(x) \rangle \quad (3)$$

which, in the case of eigenfunctions, leads to[†]:

$$\Delta \varphi_i(x)\varphi_j(x) = (\lambda_i + \lambda_j)\varphi_i(x)\varphi_j(x) - 2 \langle \nabla \varphi_i(x), \nabla \varphi_j(x) \rangle. \quad (4)$$

A rescaled version of the correction term was empirically used in [SK15] as a descriptor field for shape matching.

An important quantity arising when studying functions on manifolds is their Dirichlet energy, which encodes information about smoothness; in the Fourier setting, this corresponds to the notion of *frequency*. Given a scalar function $f \in L^2(\mathcal{M})$ and its normalized version $u = f/\|f\|_{\mathcal{M}}$, we define the Dirichlet energy of f as

$$\mathcal{E}(f) = \int_{\mathcal{M}} \|\nabla u(x)\|^2 dx = \langle \nabla u, \nabla u \rangle = \langle u, \Delta u \rangle. \quad (5)$$

It can be easily shown that the Dirichlet energy of an eigenfunction φ_k is the associated eigenvalue λ_k .

Finally, we formalize the notion of *eigenproduct*.

Definition 2.1 Let $I = \{i_1, \dots, i_n\} \in \mathbb{N}$ be a finite set of indices, possibly containing repeated elements. We define the *eigenproduct* $\varphi_I : \mathcal{M} \rightarrow \mathbb{R}$ to be the scalar function defined as

$$\varphi_I(x) = \prod_{i \in I} \varphi_i(x), \quad (6)$$

where the multiplication is to be taken pointwise. A special case of eigenproduct is when $I = \{i, \dots, i\}$ is a set containing n times the same index. In this case, we define the function $\varphi_I = \varphi_i^n$ an *eigenpower*. Finally, an N -th order *Fourier polynomial* of K eigenfunctions is a linear combination of eigenproducts up to order N involving the first K eigenfunctions (excluding the constant one).

Functional maps. Let $\pi : \mathcal{N} \rightarrow \mathcal{M}$ be a pointwise map between two surfaces \mathcal{N} and \mathcal{M} . A functional map [OBCS*12, OCB*16] is the linear operator $T_F : L^2(\mathcal{M}) \rightarrow L^2(\mathcal{N})$ that maps functions from \mathcal{M} to \mathcal{N} , defined via the pull-back $T_F(f) = f \circ \pi, \forall f \in L^2(\mathcal{M})$. Given two bases $\{\varphi_i\}$ and $\{\psi_j\}$ respectively for $L^2(\mathcal{M})$ and $L^2(\mathcal{N})$, the operator T_F can be encoded in a matrix $\mathbf{C} = (c_{ij})$, computed as:

$$T_F(f) = T_F \left(\sum_i \langle f, \varphi_i \rangle_{\mathcal{M}} \varphi_i \right) = \sum_{ij} \langle f, \varphi_i \rangle_{\mathcal{M}} \underbrace{\langle T_F(\varphi_i), \psi_j \rangle_{\mathcal{N}}}_{c_{ji}} \psi_j \quad (7)$$

In [OBCS*12] the authors propose to truncate the sum on i and j in equation (7). With this choice, the functional map reduces to a small matrix \mathbf{C} that can be efficiently estimated as described in [OBCS*12, NO17] and other related works.

Discretization. In the discrete setting, we represent a Riemannian surface \mathcal{M} as a triangle mesh with n vertices, connected by edges that form a triangle mesh approximating the smooth surface \mathcal{M} . Scalar functions $f : \mathcal{M} \rightarrow \mathbb{R}$ are represented as vectors $\mathbf{f} \in \mathbb{R}^n$. The

[†] Strictly speaking, functions in $L^2(\mathcal{M})$ do not admit a pointwise product; we keep the notation for the sake of simplicity, with the understanding that it remains valid in the proper Sobolev space.




Figure 3: Shapes used for our theoretical results: human ($\sim 7k$ vertices), bunny ($\sim 10k$), cat ($\sim 10k$), and donut ($\sim 20k$).

Laplace-Beltrami operator can be represented as the $n \times n$ matrix $\Delta = \mathbf{A}^{-1}\mathbf{W}$, where matrix \mathbf{A} contains the area elements associated to each vertex, and matrix \mathbf{W} is defined according to the local geometry (see, e.g. [PP93]). The manifold inner product $\langle f, g \rangle_{\mathcal{M}}$ is discretized as the area-weighted scalar product $\mathbf{f}^T \mathbf{A} \mathbf{g}$.

3. Theoretical results

In this section, we present some theoretical results on eigenproducts, together with some interesting implications of their properties. We first examine their frequency distribution, and compare it to the frequency distribution of the eigenfunctions (i.e. their associated eigenvalues). We then discuss a result on the approximation of eigenproducts in the space spanned by the eigenfunctions, providing possible interpretations and implications. We use different shapes to show the generality of our results (see Figure 3).

3.1. Frequency distribution

We now present a result about the distribution of frequencies (i.e. the Dirichlet energies) of eigenproducts.

Theorem 3.1 Let \mathcal{M} be a d -dimensional Riemannian manifold, and Δ be the associated Laplace-Beltrami operator. Then, let I be a set of indices of size N and let φ_I be an eigenproduct. The following relation holds:

$$\mathcal{E}(\varphi_I) \geq \frac{1}{2} \sum_{i \in I} \lambda_i. \quad (8)$$

Furthermore, in the special case of an eigenpower $\varphi_I = \varphi_i^N$, it holds

$$\mathcal{E}(\varphi_i^N) = \frac{N^2}{2N-1} \lambda_i. \quad (9)$$

Proof See Appendix A.

Corollary 3.1.1 Let \mathcal{M} be a 2-dimensional manifold, and Δ be the associated Laplace-Beltrami operator. Fixed $N, K \in \mathbb{N}$, and being Φ the set of N -th order eigenproducts between the first K Laplacian eigenfunctions, we have $\max_{\varphi_I \in \Phi} \{\mathcal{E}(\varphi_I)\} \in \Omega(NK)$.

Proof See Appendix A.

From Theorem 3.1 and Corollary 3.1.1, we can deduce the following. Since, by Weyl's law, for 2-dimensional manifolds the (NK) -th eigenvalue is $\Theta(NK)$, then using the Dirichlet energies of N -th order products between K eigenfunctions allows to express the same frequencies as if we use NK eigenfunctions. Further, within the band of the first NK eigenvalues, each eigenvalue is matched





Figure 4: Each plot shows the relative error $\min\{(\lambda_i - \mathcal{E}(\varphi_I))/\lambda_i\}$ between the i -th Laplacian eigenvalue (where i ranges on the x axis) and the closest frequency of a N -th order eigenproduct. Here NK Laplacian eigenvalues are considered, with $K = 30$ and product order $N = \{1, 2, 3\}$ (left to right). At order 1 the products correspond to the standard eigenfunctions, hence yielding exactly zero error. At increasing order the error stays close to zero, showing that each eigenvalue is approximated by the frequency of a product with $> 99.9\%$ accuracy. Results are averaged on four shapes.

to high accuracy by the Dirichlet energy of an eigenproduct, as we empirically demonstrate in Figure 4.

Intuitively, this means that the eigenproducts up to order N have the same expressive power as the first NK eigenfunctions. Therefore, eigenproducts can be used to represent well band-limited functions within the band of the first NK eigenvalues.

Example. On the real line $\mathcal{M} = [0, T]$, consider the second-order eigenproduct $\tilde{\varphi}(x) = \sin(2\pi x/T) \sin(4\pi x/T)$. Its Dirichlet energy is $\mathcal{E}(\tilde{\varphi}) = 1/\|\tilde{\varphi}\|^2 \int_0^T (\partial/\partial x \tilde{\varphi}(x))^2 dx = (1/\|\tilde{\varphi}\|^2) 5\pi^2/T$. Since the squared norm of $\tilde{\varphi}$ is $T/4$, we get $\mathcal{E}(\tilde{\varphi}) = 20\pi^2/T^2 = 4\pi^2/T^2 + 16\pi^2/T^2 = \lambda_1 + \lambda_2$. See the inset figure for an illustration.



Outside of the band of the first NK eigenvalues, the Dirichlet energy of eigenproducts grows more rapidly than the sum of the eigenvalues; we illustrate this behavior in Figure 5. In fact, the Dirichlet energy of an eigenproduct φ_I can be expressed as (see the proof of Theorem 3.1 in Appendix A for the derivations):

$$\mathcal{E}(\varphi_I) = \frac{1}{2} \sum_{i \in I} \lambda_i + \sum_{i \in I} \int_{\mathcal{M}} \left\| \prod_{\substack{j \in I \\ j \neq i}} \varphi_j(x) \nabla \varphi_i(x) \right\|^2 dx. \quad (10)$$

From Equation (10), we can see how as we add more factors to the eigenproduct, the number of terms grows and, since these terms are all non-negative, the whole energy is increased.

3.2. Approximating eigenproducts

We now investigate the following question, and draw some interesting conclusions that were missing in previous work: Can a given eigenproduct be represented well in the standard eigenbasis?

Let \mathcal{M} be a 2-dimensional manifold and $\{\varphi_i\}$ be the set of its




Figure 5: Dirichlet energies of order-2 products against the sum of the eigenvalues. The former grow much faster in the high portion of the spectrum.

Laplacian eigenfunctions. For any function f , let $E_\nu(f)$ be the projection of f onto the first ν Laplacian eigenfunctions and let $R_\nu(f)$ be the L^2 norm of $f - E_\nu(f)$. Namely, the residual:

$$R_\nu(f) = \left\| f - \sum_{i=0}^{\nu} \langle \varphi_i, f \rangle_{\mathcal{M}} \varphi_i \right\|_{L^2}. \quad (11)$$

Affalo et al. [ABK15] proved the upper bound:

$$R_\nu^2(f) \leq \frac{\|\nabla f\|^2}{\lambda_{\nu+1}} \quad \forall f, \quad (12)$$

further showing that the bound can not be tightened by any other sequence of linearly independent functions $\{\varphi_i\} \in L^2(\mathcal{M})$. This yields the optimality of the Laplacian eigenfunctions for representing any function with bounded gradient magnitude. For the special case of f being an eigenproduct φ_I , however, the bound is not very informative. We instead appeal to the following:

Theorem 3.2 [LSS18] Fixed $K, N \in \mathbb{N}$, with $K \gg 1$, for any set of indices $I = \{i_1, \dots, i_N\}$, where each $i_j \leq K$, for any $\nu > K$ and for any $\chi \in \mathbb{N}$, it holds:

$$R_\nu(\varphi_I) \lesssim K^{\frac{N\sigma(2N,2)}{2}} \left(\frac{K}{\nu}\right)^{\frac{\chi}{2}}, \quad (13)$$

where

$$\sigma(p, d) = \max \left\{ \frac{d-1}{2} \left(\frac{1}{2} - \frac{1}{p} \right), d \left(\frac{1}{2} - \frac{1}{p} \right) - \frac{1}{2} \right\}. \quad (14)$$

This result can be read as follows. For large values of K , the products tend to add less and less information to the spanned space, until, eventually, all the products are spanned by a basis of ν eigenfunctions, for every $\nu > K$.

Example. Consider the product $\tilde{\varphi}(x) = \sin(2\pi x) \sin(4\pi x) \sin(6\pi x)$ between the first three non-constant eigenfunctions on $\mathcal{M} = [0, 1]$. If we represent this product in the basis of the first $\nu = 4$ non-constant eigenfunctions, we get the residual $R_4(\tilde{\varphi}) = 3/64$; this residual is already quite small, if we consider that we are using only 1 more eigenfunction for representing the product.

Nevertheless, for small values of K and N it is very rare that an eigenproduct can be expressed exactly as a linear combination of eigenfunctions. Extending the basis with eigenproducts introduces a lot of additional information. Even if they are not all linearly independent, they could expand the dimension of the spanned space by orders of magnitude. In Figure 6 we show examples where using only 20 eigenfunctions, and admitting products up to the 4-th order,

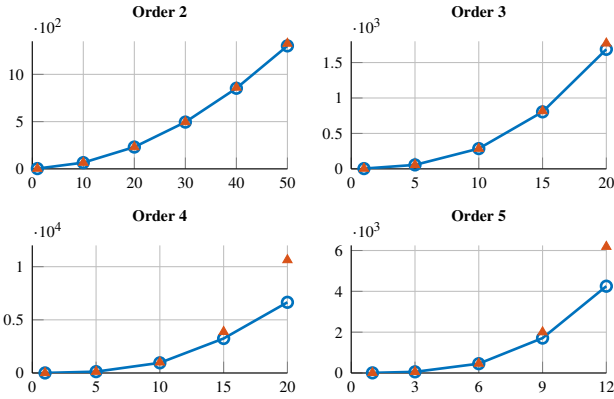


Figure 6: Number of linearly independent eigenproducts (in blue) against the number of eigenfunctions involved; each plot is for a different maximum product order. The total number of eigenproducts (in orange) is plotted for reference. These results are averaged on four different shapes.

generates a basis spanning a functional space with > 3000 dimensions. This kind of rank analysis was also missing in [NMR*18], while the study about the loss of the eigenproducts basis’ full-rank property, reported in Figure 6, gives a heuristic for selecting parameters N and K . According to our theoretical and experimental analysis, we advocate using $K \leq 50, 15, 10$ for, respectively, $N = 2, 3, 4$, since this parameter setup maximizes the size of the basis with respect to the computed products.

3.3. Orthogonalized eigenproducts

In the work of Nogneng et al. [NMR*18], resorting to eigenproducts for representing surface signals also involved solving an optimization problem prone to numerical instability. Here we propose a much simpler alternative that is less empirical, more stable, and provides consistently better results than [NMR*18]. Specifically, consider the set of order- N products of the first K eigenfunctions, spanning a function space $\mathcal{F}(\mathcal{M})$ with some dimension Q . We orthogonalize this set via the Gram-Schmidt (GS) algorithm, and obtain an orthonormal basis for the same space with exactly Q basis functions; see Figures 7 and 8 for examples. Although straightforward, this process offers advantages both in terms of computational effort and numerical stability.

Complexity. In [NMR*18], computing a representation for a given signal in the set of eigenproducts requires computing the SVD decomposition of a $n \times L$ matrix containing all eigenproducts as its columns, where n is the number of vertices and $L = \binom{K+N}{N}$ is the number of N -th order products between the first K eigenfunctions. This SVD decomposition has complexity $\mathcal{O}(n^2L + L^3)$ [GVL13], which for $L \ll n$ (our case) reduces to $\mathcal{O}(n^2L)$. On the other hand, the computational complexity of GS is $\mathcal{O}(nL^2)$, which for $L \ll n$ is much more sustainable than SVD decomposition.

Stability. We demonstrate empirically that our orthogonal basis produces more stable results than [NMR*18]. For the latter method,

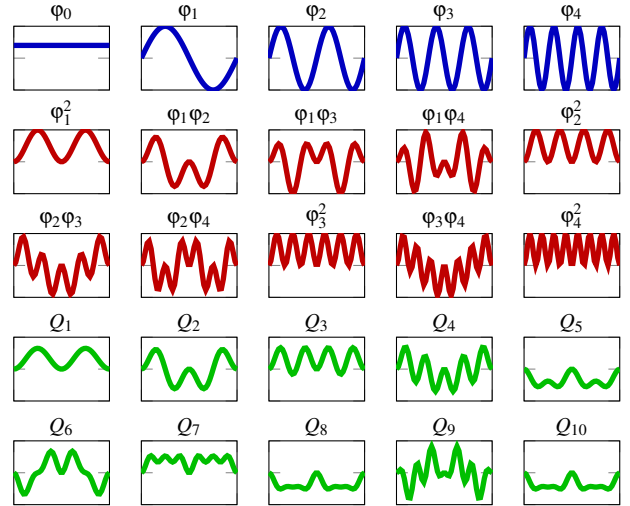


Figure 7: Standard Laplacian eigenfunctions (blue), eigenproducts (red, in lexicographic order) and orthonormalized eigenproducts (green) on the real line $[0, T]$.

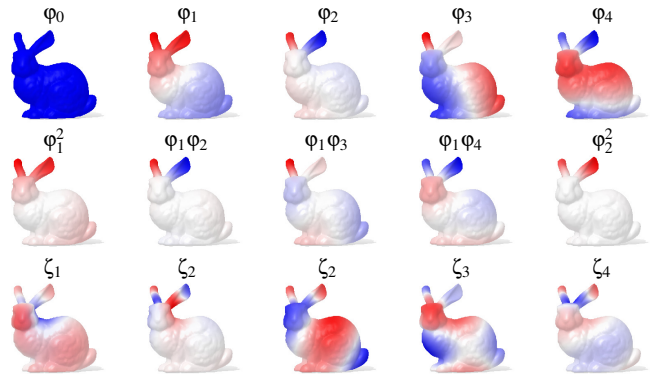


Figure 8: Top to bottom: first five Laplacian eigenfunctions, first five eigenproducts, and the orthogonalized eigenproducts.

numerical inaccuracies occurring in the computation of the representation tend to propagate when the function is mapped to a different domain, producing local scale errors; see Figure 9 and the experimental section for examples.

Transfer. In a similar spirit as [NMR*18], where the authors extend the notion of functional map to eigenproducts, we also provide a way to compute a transfer matrix for our orthogonal bases.

Let us be given a functional map matrix \mathbf{C} between two shapes \mathcal{M} and \mathcal{N} , and let us assume a manifold-independent ordering of the eigenproducts (e.g., a lexicographic ordering on the indices of the factors). We then have an ordered set of P eigenproducts $\tilde{\Phi} = \{\varphi_{I_1}, \dots, \varphi_{I_P}\}$ on \mathcal{M} , and similarly $\tilde{\Psi}$ on \mathcal{N} . After orthonor-




Figure 9: Function transfer example on two pairs from the FAUST [BRLB14] dataset. The adoption of eigenproducts as in [NMR*18] (denoted as prods) yields local scale artifacts (knee and neck in the first row) or loss of high frequency details (arms in the second row). With the standard eigenbasis (eigs), K eigenfunctions are simply not enough to capture the frequency content of the transferred signal. In the bottom row, the second variant of our method (ours*) achieves almost perfect reconstruction.

malization, we get a new set of functions in the form:

$$\zeta_i(x) = \varphi_i(x) - \sum_{j=1}^{i-1} \langle \varphi_i, \zeta_j \rangle_{\mathcal{M}} \zeta_j(x). \quad (15)$$

By Equation (7), matrix \mathbf{C} maps each eigenfunction φ_i on \mathcal{M} to a linear combination $\sum_{j=1}^K c_{j,i} \psi_j$ of eigenfunctions on \mathcal{N} . Assume for now that $T_F(\Phi) = \Psi\mathbf{C}$ is a strict equality. Further, for the sake of simplicity, we limit the exposition to second-order products, but the process can be iterated and generalized to any higher order. Each eigenproduct φ_I on \mathcal{M} is the product between two eigenfunctions φ_i and φ_j . When we consider the mapping induced by \mathbf{C} , we get:

$$\begin{aligned} \varphi_I(x) &= \varphi_i(x)\varphi_j(x) = \left(\sum_{h=1}^K c_{h,i} \psi_h(x) \right) \left(\sum_{p=1}^K c_{p,j} \psi_p(x) \right) = \\ &= \sum_{h,p=1}^K c_{h,i} c_{p,j} \psi_h(x) \psi_p(x). \end{aligned} \quad (16)$$

Thus, the coefficients for transferring products are fully determined by the coefficients for transferring eigenfunctions, as also shown in [NMR*18]. Hence, one can compute a matrix $\tilde{\mathbf{C}}(\mathbf{C})$, depending only on \mathbf{C} , such that $T_F(\tilde{\Phi}) = \tilde{\Psi}\tilde{\mathbf{C}}$.

Going one step further, we extend matrix $\tilde{\mathbf{C}}$ to a new matrix \mathbf{O} that can correctly transfer the orthonormalized basis. We proceed as follows. In general, given a set of vectors $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ with n elements and rank $1 \leq r \leq n$, GS produces an orthogonal basis $\mathbf{Q} = (\mathbf{q}_1, \dots, \mathbf{q}_r)$ with r elements, spanning the same space as \mathbf{B} . A

side-product is the $r \times n$ upper triangular matrix $\mathbf{R} = (r_{i,j})$:

$$\mathbf{R} = (r_{i,j}) = \begin{pmatrix} \langle \mathbf{q}_1, \mathbf{b}_1 \rangle & \langle \mathbf{q}_1, \mathbf{b}_2 \rangle & \dots & \langle \mathbf{q}_1, \mathbf{b}_n \rangle \\ 0 & \langle \mathbf{q}_2, \mathbf{b}_2 \rangle & \dots & \langle \mathbf{q}_2, \mathbf{b}_n \rangle \\ \vdots & \ddots & \dots & \vdots \\ 0 & 0 & \dots & \langle \mathbf{q}_r, \mathbf{b}_n \rangle \end{pmatrix}, \quad (17)$$

such that $\mathbf{QR} = \mathbf{B}$.

Applied to the sets of eigenproducts $\tilde{\Phi}$ and $\tilde{\Psi}$, we get the factorizations $\tilde{\Phi} = \mathbf{Q}_\Phi \mathbf{R}_\Phi$ and $\tilde{\Psi} = \mathbf{Q}_\Psi \mathbf{R}_\Psi$.

Since \mathbf{Q}_Φ and \mathbf{Q}_Ψ span the same space as $\tilde{\Phi}$ and $\tilde{\Psi}$, assuming there exists a meaningful mapping between $\tilde{\Phi}$ and $\tilde{\Psi}$, it makes sense to search for a mapping between \mathbf{Q}_Φ and \mathbf{Q}_Ψ . Thus, we seek for a matrix \mathbf{O} such that:

$$T_F(\mathbf{Q}_\Phi) = \mathbf{Q}_\Psi \mathbf{O}. \quad (18)$$

By the factorization above, we can equivalently rewrite $T_F(\tilde{\Phi}) = \tilde{\Psi}\tilde{\mathbf{C}}$ as $T_F(\mathbf{Q}_\Phi)\mathbf{R}_\Phi = \mathbf{Q}_\Psi\mathbf{R}_\Psi\tilde{\mathbf{C}}$; plugging in Equation (18) and simplifying, we get:

$$\mathbf{O}\mathbf{R}_\Phi = \mathbf{R}_\Psi\tilde{\mathbf{C}}. \quad (19)$$

Since \mathbf{R}_Φ is full row-rank, it has a right inverse $\mathbf{R}_\Phi^\top (\mathbf{R}_\Phi \mathbf{R}_\Phi^\top)^{-1}$. Thus, if we have access to the bases on both manifolds, we can directly compute \mathbf{O} as

$$\mathbf{O} = \mathbf{R}_\Psi \tilde{\mathbf{C}} \mathbf{R}_\Phi^\top (\mathbf{R}_\Phi \mathbf{R}_\Phi^\top)^{-1}. \quad (20)$$

For small enough order N and number of eigenfunctions K , the set of eigenproducts is likely to be full-rank (see Figure 6 for an empirical assessment); in this case, \mathbf{R}_Ψ is a square invertible matrix, leading to:

$$\mathbf{O} = \mathbf{R}_\Psi \tilde{\mathbf{C}} \mathbf{R}_\Phi^{-1}, \quad (21)$$

and since \mathbf{R}_Ψ is upper triangular, its inverse can be computed efficiently.

We also derive the following alternative formula to compute iteratively the columns of \mathbf{O} depending only on the matrices $\tilde{\mathbf{C}}$, \mathbf{R}_Φ , and \mathbf{R}_Ψ :

$$T_F(\zeta_i) = \frac{1}{(\mathbf{R}_\Phi)_{i,i}} \left(\sum_j \underbrace{\left(\sum_{l=1}^i (\mathbf{R}_\Psi)_{j,l} \tilde{c}_{l,i} - \sum_{h=1}^{i-1} (\mathbf{R}_\Phi)_{i,h} o_{j,h} \right)}_{o_{j,i}} \xi_j \right). \quad (22)$$

We report the full derivation of this formula in the supplementary materials. In both Equations (20) and (22), we have an explicit formula to compute \mathbf{O} directly from $\tilde{\mathbf{C}}$, and thus, from \mathbf{C} .

3.4. Implementation details

Despite their simplicity, the analytic solutions we introduced above can be unstable as they rely on the hypothesis that \mathbf{C} maps perfectly the source eigenbasis onto the target, $T_F(\Phi) = \Psi\mathbf{C}$. If the equality is not exact, we observed a quick degradation in accuracy. Moreover, in the iterative process of Equation (22), the error accumulates at each iteration. We highlight this issue in Figure 10 (left), by evaluating the error obtained in the estimation of \mathbf{O} using the formula (20) while we increasingly add white noise to the




Figure 10: Left: Error generated by Equation (20) (red curve) at increasing random noise on the matrix \mathbf{C} . We show the error for the functional map as a reference (blue curve). Right: Error induced by the analytic expression of \mathbf{O} in transferring the orthogonal basis onto the target shape at increasing number of eigenfunctions. Each curve represents a different ratio between the number $K_{\mathcal{N}}$ of eigenfunctions on the target and the number $K_{\mathcal{M}}$ on the source.

coefficients stored in the map \mathbf{C} . We compute the relative error as the Frobenius norm of the difference between $T_F(\Phi)$ and $\Psi\mathbf{C}$ and the difference between $T_F(\mathbf{Q}_\Phi)$ and $\mathbf{Q}_\Psi\mathbf{O}$. The latter difference increases quickly when the former grows. Moreover, the accuracy of the transfer matrix \mathbf{O} decreases as we increase K ; we show this in Figure 10 (right), by plotting the error for different dimensions of the source and target bases. To solve these issues, we propose a different strategy to recover the transfer matrix \mathbf{O} . Being \mathbf{O} a function of $\tilde{\mathbf{C}}$ (and so of \mathbf{C}), we make the educated guess that it is enough to align the first K eigenfunctions to obtain a sufficiently precise transfer matrix. To implement this idea, we proceed as follows:

1. We solve for the best permutation Π as the solution of a nearest neighbor assignment problem between $\tilde{\Psi}$ and $\tilde{\Phi}\tilde{\mathbf{C}}^T$;
2. We estimate \mathbf{O} as the solution in the least-squares sense of $\mathbf{Q}_\Psi\mathbf{O} = \Pi\mathbf{Q}_\Phi$.

This process is equivalent to extracting a point-to-point map from a given functional map as in [OBBS*12], and then converting the point-to-point map to a new matrix with respect to different bases. Since our basis can be completely derived from the first K eigenfunctions (with respect to which the matrix \mathbf{C} is encoded), we expect the estimated Π to contain enough information for estimating an accurate \mathbf{O} . By virtue of our analysis on the space spanned by the set of eigenproducts, we only need a small number of eigenfunctions $K \approx 12$, which ensures an easier optimization for the functional maps. It is worth to stress that this process is more feasible than extracting a point-to-point map directly from the orthogonalized basis, since we extract the map from just K eigenfunctions, rather than searching for a match between two sets of $\mathcal{O}(K^N)$ basis functions. This is even more efficient than fast refinement methods like ZOOMOUT [MRR*19], since we do a one-shot computation, rather than iterating on large bases one function at a time.

Dependency on the first K eigenfunctions. The i -th (for $i > K$) orthogonalized eigenproduct can be explicitly formulated in terms of the first K eigenfunctions. This is not true when we consider a basis of $K^* > K$ Laplacian eigenfunctions. Thus, it is reasonable to think that, starting from a functional map \mathbf{C} for the first K eigen-

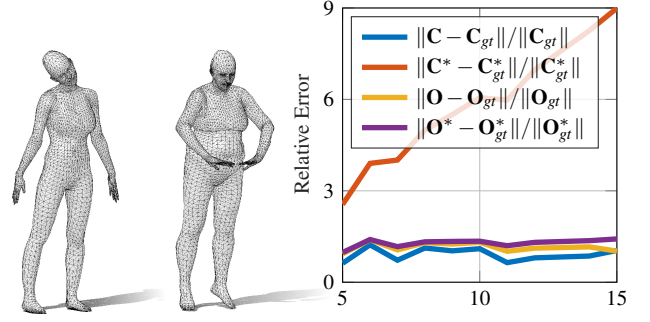


Figure 11: Comparison in the estimation of the transfer matrix directly from the point-to-point map extracted from $\tilde{\mathbf{C}}$. Left: an example pair with different connectivity. Right: error curves vs. size of the initial matrix \mathbf{C} , averaged over 20 random FAUST pairs.

functions, we can handle the orthonormalized products basis better than the standard Laplacian eigenbasis.

To test this claim, we apply the same strategy of computing our transfer matrix \mathbf{O} to estimate a functional map \mathbf{C}^* between eigenfunctions with the same cardinality as our bases. In Figure 11, for products of order $N = 3$ we compare the estimation of the two matrices varying the size K of the initial functional map \mathbf{C} (x -axis). The error (y -axis) is defined as the relative Frobenius norm between the estimated and the ground-truth matrix. We report the error for the initial functional map \mathbf{C} for comparison. In this test, the shapes do not share the same connectivity, as we remesh one of the two to 5k vertices. The errors for \mathbf{O} and \mathbf{O}^* are stable and comparable to the error of \mathbf{C} , while the error for the functional map \mathbf{C}^* increases when K grows. This shows that it is possible to directly estimate \mathbf{O} accurately from the point-to-point map extracted from \mathbf{C} .

Robustness to noise. We add an increasing amount of white noise (x -axis) to an input ground-truth functional map \mathbf{C} , and extract the transfer matrix \mathbf{O} . Then we compare the error in the transfer of the coordinate functions for K eigs, prods, ours and ours*. The curves in the inset figure show that our method is less sensitive to noise. These results are averaged on 20 random FAUST pairs, remeshed as in Figure 11. In Figure 12, we show a qualitative result under maximum noise. The corrupted map generates visible distortion in the transfer for both K eigs and prods, while the proposed method remains stable.

4. Experimental results

We evaluate the performance of our basis in tasks of surface filtering, signal approximation and transfer. For all the tasks we compare with [NMR*18], and also report the results of taking a linear combination of NK eigenfunctions and of $\binom{K+N}{N}$ eigenfunctions.

We also investigate the numerical stability of GS, which may produce non-orthogonal vectors with large, non full-rank input bases, since linearly dependent functions must be discarded from




Figure 12: Comparisons on function transfer with noise.




Figure 13: Detail transfer with our orthogonalized eigenproducts from a textured shape to another. The geometric details on the target are replaced correctly with those of the source shape.

the final set. This is done by thresholding small inner products. We consider (i) a direct implementation with a permissive threshold (10^{-2}) for the identification of linearly dependent vectors (denoted by **Ours**), and (ii) a reiterated variant, with a strict threshold (10^{-9}), which enforces orthogonality (**Ours***). The two approaches may output very different bases, which are shown by experiments to be suitable in different applications. In particular, **Ours** is more numerically stable, therefore more suitable for transferring signals in general (Tables 2 to 4). On the other hand, **Ours*** sacrifices numerical stability for more descriptive power, which can be useful for spectral filtering (Figures 14 and 15), reconstruction (Figure 17), detail transfer (Figure 13), or transfer between very different subjects (Table 6).

The code is available at a public repository on GitHub[‡].

4.1. Detail transfer

In Figure 13, we highlight the representation power of the proposed basis by transferring fine geometric details (encoded as vertex coordinates) from a source shape to a target. For the first pair of shapes, we use products of order $N = 3$ of the first $K_M = 30$ eigenfunctions on the source shape and $K_N = 40$ on the target. For the second pair, we still use products of order $N = 3$, but we limited the number of eigenfunctions to $K_M = 27$ on the source shape and $K_N = 30$ on the target to show the result with a different set of parameters. The standard eigenfunctions cannot represent the details even with NK eigenfunctions, while prods [NMR*18] only approximates the details. Instead, our basis fulfils the task with good accuracy.

[‡] <https://github.com/filthy nobleman/orthogonalized-fourier-polynomial>




Figure 14: Filtering the coefficients of the vertex coordinates represented in the set of eigenproducts (middle) and in our orthogonal basis (right). We use a smoothing filter (blue curve, top row) and a sharpening filter (red curve, bottom row). Eigenproducts are not easy to control and do not provide meaningful variations, while our representation leads to the expected results.




Figure 15: Stability of spectral filtering. From left to right, we gradually increase by ϵ only one coefficient in the spectral representation of the vertex coordinates; we use the same ϵ for both rows. While the reconstruction using eigenproducts degenerates quickly, ours yields a more stable and meaningful deformation.

4.2. Spectral filtering

Since the eigenproducts of [NMR*18] are not linearly independent and do not provide a unique representation for a given signal, it is hard to design operations on the signal coefficients along the lines of spectral filtering approaches such as [VL08]. To demonstrate this, we run a test where we process vertex coordinates with simple filters (Figure 14): a smoothing filter that suppresses the high frequencies, and an enhancement filter that suppresses the lower portion of the spectrum. As a second test, in Figure 15 we continuously change a *single* coefficient in the spectral representation of the vertex coordinates, and then reconstruct the resulting geometry. While our basis can recover a meaningful smooth deformation of the original geometry, the plain eigenproducts rapidly degenerate to a flat shape. All these experiments are performed using products of order 3 on the first 15 eigenfunctions.

4.3. Function approximation

In a discrete setting, computing K Laplacian eigenfunctions boils down to solving an eigenproblem for a sparse symmetric $n \times n$ matrix, which has complexity $\mathcal{O}(Kn^2)$ [Ste02]. This is prohibitive for high resolution meshes and for large K , which is required to capture fine details. Using eigenproducts for the representation as in [NMR*18] can lead to big accuracy improvements, at the cost of increased numerical instability as discussed above. By contrast, our orthonormal basis leads to more stable and accurate results.




Figure 16: Reconstruction of a RGB signal using $K = 100$ eigenfunctions and order-2 products. We compare the reconstruction quality using K eigenfunctions, the eigenproducts of [NMR*18], and our two methods based on the orthogonalized basis.

	K	prods	ours	NK	$\binom{K+N}{N}$
HK k	0.0%	0.0%	0.0%	0.0%	0.0%
HK K	75.8%	18.1%	13.7%	61.1%	0.0%
HKS	1.1%	0.0%	0.0%	0.4%	0.0%
WKS	14.1%	0.0%	0.0%	5.7%	0.1%
Rand	37.5%	33.8%	33.5%	37.2%	33.4%
XYZ	15.2%	2.0%	1.7%	7.6%	0.8%
Ind	29.2%	12.5%	12.2%	23.6%	11.4%
SHOT	66.8%	35.4%	33.8%	58.8%	28.9%
AWFT	12.5%	5.9%	5.8%	10.5%	4.2%

Table 1: Reconstruction error of our method compared to the approach from [NMR*18] (prods) and just taking linear combinations of K eigenfunctions (K). As ideal references, we also report the results obtained via linear combinations of NK and $\binom{K+N}{N}$ eigenfunctions, which are prohibitive to compute on large shapes. Here, we used $K = 30$ and $N = 2$.

In Table 1 we compare with the standard eigenfunctions and the eigenproducts of [NMR*18] in terms of reconstruction error, measured as $(\int_{\mathcal{M}} (f - \tilde{f})^2)^{1/2} / (\int_{\mathcal{M}} f^2)^{1/2}$, where f is the original signal and \tilde{f} is its reconstruction. We consider the same families of functions as in [NMR*18]. Namely, HK k, HK K: the heat kernel between a random point and the rest of the shape approximated using 200 and K eigenfunctions respectively. HKS, WKS: the heat and wave kernel signatures [SOG09, ASC11]. Rand: random functions. XYZ: vertex coordinates. Ind: the binary indicator function of a random region. SHOT, AWFT: local descriptors [TSDS10, MRCB16].

Moving to an orthonormal basis brings additional regularization, leading in turn to an increase in quality. Both our orthogonal basis and the plain eigenproducts produce much better results than those obtained with K or even NK eigenfunctions. The two methods can, in most cases, compete with the approach of using $\binom{K+N}{N}$ eigenfunctions, without incurring in the prohibitive cost of computing a massive number of eigenfunctions (in the order of $\mathcal{O}(K^N)$). Qualitative results on this task are shown in Figures 1 and 17 for the coordinate functions (XYZ), and in Figure 16 for a RGB signal.

4.4. Function transfer

We evaluate the transfer task on several datasets, differing in terms of mesh quality, resolution, regularity, and deformation type (isometry or lack thereof). We consider the same set of functions used for this experiment in [NMR*18]. As a reference, we also report the re-




Figure 17: Coordinate reconstruction of two shapes. With N -th order products of K eigenfunctions, the eigenproducts and the orthogonal basis get better results than NK eigenfunctions, and start to catch some details from the surface. A lower threshold in the orthogonalization process produces a more descriptive basis, in these cases. Here we use $N = 2, K = 50$ for the statue and $N = 3, K = 18$ for the dancing children.

sults obtained with linear combinations of NK standard eigenfunctions. For each method, we compute the functional map matrix \mathbf{C} with the method of [NO17]. As an error measure, we report the normalized error $(\int_{\mathcal{N}} (f_{gt} - \tilde{f})^2)^{1/2} / (\int_{\mathcal{N}} f_{gt}^2)^{1/2}$, where f_{gt} is the ground truth signal on the target \mathcal{N} and \tilde{f} is the transferred counterpart.

Near-isometries (synthetic). For synthetic near-isometric meshes we use 10 random pairs from TOSCA [BBK08]; Table 2 summarizes the comparisons, showing a 5–25% improvement over [NMR*18]. Using third-order products allows to successfully transfer the 46-th eigenfunction by using only 15 eigenfunctions on the source and 18 on the target shape, as shown in Figure 18.

Non-isometries (real humans). Comparisons are favorable also when considering more realistic non-isometric shapes, where we

	K	prods	ours	ours*	NK
HK k	70.3%	64.4%	50.9%	52.0%	17.7%
HK K	78.4%	61.3%	52.8%	55.9%	56.4%
HKS	12.0%	10.4%	5.5%	5.6%	5.9%
WKS	29.3%	22.0%	10.6%	10.8%	17.7%
Rand	50.7%	54.8%	49.9%	50.0%	50.8%
XYZ	40.6%	44.5%	32.4%	33.2%	21.3%
Ind	47.7%	51.2%	27.2%	24.8%	35.9%
SHOT	74.6%	73.3%	65.4%	65.3%	67.0%
AWFT	31.4%	35.4%	23.8%	24.2%	22.9%

Table 2: Transfer error comparison on 10 isometric pairs from TOSCA. The parameters are $N = 3$ and $K_{\mathcal{M}} = K_{\mathcal{N}} = 12$.

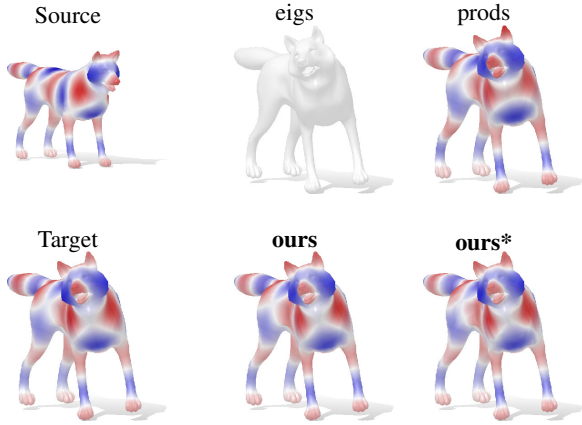


Figure 18: Transferring a high-frequency eigenfunction that is out of the span of the truncated Laplacian eigenbasis. Projecting onto the latter yields a zero signal due to orthogonality, while eigenproducts and our orthogonal basis can transfer it more precisely. Here we used $K_{\mathcal{M}} = 15$ eigenfunctions on the source shape, $K_{\mathcal{N}} = 18$ on the target shape, and $N = 3$ (third-order products).

use 20 random pairs from the FAUST [BRLB14] dataset of real scans. Each pair is simultaneously inter-class and inter-pose. Table 3 summarizes the results. On this dataset, we select K and N as in [NMR*18] for a direct comparison.

	K	prods	ours	ours*	NK
HK k	28.0%	27.4%	26.6%	27.1%	22.5%
HK K	66.7%	51.3%	46.4%	46.4%	56.3%
HKS	8.1%	12.3%	6.6%	6.7%	8.2%
WKS	16.7%	15.5%	11.4%	11.6%	12.3%
Rand	49.9%	51.0%	49.6%	50.5%	50.0%
XYZ	21.3%	21.9%	19.5%	20.0%	19.5%
Ind	36.4%	33.7%	30.1%	30.2%	31.4%
SHOT	76.6%	75.1%	73.7%	76.3%	70.7%
AWFT	15.1%	18.3%	13.4%	13.9%	14.7%

Table 3: Transfer error comparison on 20 inter-class pairs from FAUST. Here $N = 2$, $K_{\mathcal{M}} = 30$ and $K_{\mathcal{N}} = 40$.



Figure 19: RGB signal transfer comparison on a pair from the SHREC'19 dataset. Our approach better transfers the fine details. In this example, the noise in the ground truth correspondence due to wildly different meshings is enough to produce local distortion in the function transfer, but our approach remains stable and produces a smooth correspondence.

Near-isometries (different meshing). We further evaluate the setting where the shapes have a wildly different vertex count and connectivity. For these tests, we take 10 random pairs from SHREC'19 [MMR*19]; the results are reported in Table 4 and highlight how, under a different connectivity, even using a large eigenbasis incurs into issues (NK column), whereas our method still gives stable results. Note that when a mesh in a pair is undersampled with re-

	K	prods	ours	ours*	NK
HK k	27.9%	29.3%	26.2%	26.4%	30.4%
HK K	79.9%	67.4%	62.7%	62.8%	75.8%
HKS	12.6%	17.5%	10.7%	10.8%	13.7%
WKS	25.4%	21.6%	18.5%	18.6%	22.4%
Rand	50.8%	52.8%	50.2%	50.4%	51.5%
XYZ	30.5%	30.2%	29.2%	29.6%	30.8%
Ind	41.0%	38.2%	34.8%	35.7%	36.6%
SHOT	66.3%	66.0%	62.8%	64.7%	64.8%
AWFT	24.6%	28.0%	21.4%	21.6%	25.6%

Table 4: Transfer error comparison on 10 non-isometric pairs (with very different connectivity) from SHREC'19, with $N = 2$, $K_{\mathcal{M}} = 30$ and $K_{\mathcal{N}} = 40$.

spect to the other, the ground-truth correspondence is patchy and low-quality. Figure 19 shows how our approach can still produce a higher quality transfer.

Non-isometries (different animals). We extend our comparison to the recent dataset SHREC'20 [DLR*20], composed by 30 pairs divided in 5 different test-sets (1 partial-to-full and 4 full-to-full shape with 4 different levels of isometry: highest, high, low, lowest). In total, 14 different shapes are involved with different connectivity, topological errors and missing parts. A set of ~ 50 sparse ground truth correspondences are provided for the evaluation. Due to the impossibility of retrieving a dense ground truth, here we measure the average absolute error $|f_{gr} - \hat{f}|$ on the sparse landmarks only. In Table 5, we show that the proposed method transfers functions between these animals better than the competitors. A qualitative visualization is given by the RGB transfer results in Figure 2.

Non-isometries (different semantic class). Finally, we address the difficult setting where the shapes belong to different classes. We

	K	prods	ours	ours*	NK
HK k	12.2%	15.8%	15.2%	13.0%	14.8%
HK K	4.4%	5.9%	5.3%	3.2%	4.8%
HKS	26.6%	30.9%	23.6%	23.6%	26.7%
WKS	19.2%	25.1%	17.7%	17.7%	19.9%
Rand	26.8%	40.4%	27.2%	25.2%	26.9%
XYZ	18.1%	51.5%	11.9%	7.4%	20.4%
Ind	21.0%	26.1%	20.4%	18.6%	22.1%
SHOT	10.4%	15.8%	12.6%	11.3%	10.7%
AWFT	18.8%	30.8%	20.4%	18.3%	19.6%

Table 5: Transfer error comparison for SHREC’20, averaged on the 5 test-sets (30 pairs with different levels of non-isometry, topological errors and missing parts), with $N = 3$, $K_M = 12$, $K_N = 12$. The values are multiplied by 10^2 to help the comparison.

use the MISC dataset [MOR*18] with 10 random pairs in random poses for a woman, a man and a gorilla; a sparse correspondence between $\sim 10\%$ of the vertices is given. The results reported in Table 6 show that, with orthogonalized third-order eigenproducts, we are still able to transfer functions with high accuracy. Having only a sparse ground truth correspondence, we use the same error measure of SHREC’20. We want to stress that, by using this measure, the values shown in Tables 5 and 6 must be interpreted differently from those shown in Tables 2 to 4. The error, here, is not relative, and depending on the availability of landmarks, it can have higher or lower values. However, this does not affect the comparison, since all the methods are compared under the same conditions. In Figure 20, we show a qualitative result for a pair of non-isometric shapes.

	K	prods	ours	ours*	NK
HK k	11.0%	10.6%	12.0%	9.2%	8.7%
HK K	3.1%	3.7%	2.9%	2.1%	3.2%
HKS	5.8%	6.2%	5.4%	5.4%	9.6%
WKS	11.0%	8.2%	7.3%	7.3%	9.3%
Rand	25.4%	28.6%	25.2%	25.4%	25.9%
XYZ	20.1%	21.1%	19.5%	19.2%	18.1%
Ind	10.9%	10.1%	8.0%	6.9%	11.3%
SHOT	0.2%	0.2%	0.2%	0.2%	0.2%
AWFT	5.9%	7.8%	6.5%	6.3%	5.6%

Table 6: Transfer error comparison on 10 strongly non-isometric pairs from MISC. In this setting, $N = 3$ and $K_M = K_N = 12$.

5. Conclusions

In this paper, we proposed a new orthonormal basis based on the pointwise products of the eigenfunctions of the Laplace-Beltrami operator. We provided a theoretical analysis of the properties carried by this basis, and assessed its practical value in different settings and applications. In particular, orthogonalization yields a notable improvement over prior work [NMR*18] in terms of computational complexity and expressive power.

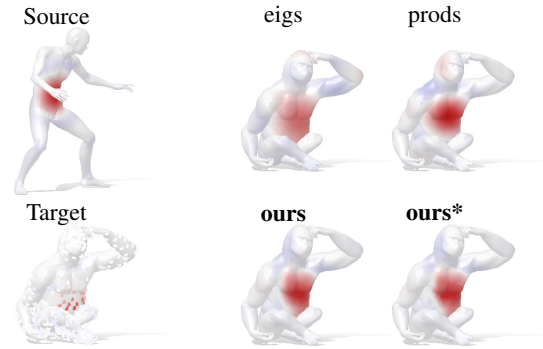


Figure 20: Example of function transfer on the MISC dataset. The ground truth on the target is only shown for the given landmarks.

Limitations and future work. Perhaps the main limitation of our approach lies in the pipeline for function transfer, where we strongly depend on the quality of the given functional map. If the method for building this map is not stable, the behavior of the orthogonal basis can be unpredictable. Furthermore, currently we need to rely on a point-to-point map in order to compute the transfer matrix, since our analytical derivations are not stable. Directly estimating this matrix *without* resorting to point-to-point conversion, is an important direction that we intend to explore. Another possible avenue for future research is to extend our construction to accommodate different definitions of products, for a more efficient and accurate representation for surface signals. Finally, we will explore the adoption of our basis in existing pipelines, such as ZOOMOUT [MRR*19], which are basis-agnostic and allow to extend this work to a number of other tasks.

Acknowledgments FM, SM and ER are supported by the ERC Starting Grant No. 802554 (SPECGEO). Parts of this work were supported by the ERC Starting Grant No. 758800 (EXPROTEA) and the ANR AI Chair AIGRETTE.

References

- [ABK15] AFLALO Y., BREZIS H., KIMMEL R.: On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences* 8, 2 (2015), 1141–1160. 2, 4
- [AL19] AZENCOT O., LAI R.: Shape analysis via functional map construction and bases pursuit, 2019. [arXiv:1909.13200](https://arxiv.org/abs/1909.13200). 2
- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV Workshops* (2011), IEEE, pp. 1626–1633. 9
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, New York, NY, 2008. 9
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR* (Columbus, Ohio, 2014), IEEE, pp. 3794–3801. 6, 10
- [Cha84] CHAVEL I.: *Eigenvalues in Riemannian geometry*. Academic Press, 1984. 3
- [CM06] COIFMAN R. R., MAGGIONI M.: Diffusion wavelets. *Applied and Computational Harmonic Analysis* 21, 1 (2006), 53–94. 2

- [CRM*16] COSMO L., RODOLÀ E., MASCI J., TORSSELLO A., BRONSTEIN M.: Matching deformable objects in clutter. In *Proc. 3D Vision (3DV)* (2016), pp. 1–10. 2
- [CSBK17] CHOUKROUN Y., SHTERN A., BRONSTEIN A., KIMMEL R.: Hamiltonian operator for spectral shape analysis. *TVCG* 26 (2017). 2
- [DLR*20] DYKE R. M., LAI Y.-K., ROSIN P. L., ZAPPALÀ S., DYKES S., GUO D., LI K., MARIN R., MELZI S., YANG J.: SHREC'20: Shape correspondence with non-isometric deformations. *Computers & Graphics* 92 (2020), 28–43. 10
- [EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. *Computer Graphics Forum* 36, 5 (2017), 165–174. 2
- [ELC20] EISENBERGER M., LAHNER Z., CREMERS D.: Smooth shells: Multi-scale shape registration with functional maps. In *Proc. CVPR* (June 2020). 2
- [GVL13] GOLUB G., VAN LOAN C.: *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. JHU Press, 2013. 5
- [HVG11] HAMMOND D. K., VANDERGHEYNST P., GRIBONVAL R.: Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis* 30, 2 (2011), 129–150. 2
- [KBB*13] KOVNATSKY A., BRONSTEIN M., BRONSTEIN A., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Computer Graphics Forum* 32, 2pt4 (2013), 439–448. 2
- [KGB16] KOVNATSKY A., GLASHOFF K., BRONSTEIN M. M.: Madmm: a generic algorithm for non-smooth optimization on manifolds. In *Proc. ECCV* (2016), Springer, pp. 680–696. 2
- [KMP*20] KIRGO M., MELZI S., PATANÈ G., RODOLÀ E., OVSJANIKOV M.: Wavelet-based heat kernel derivatives: Towards informative localized shape analysis. *Computer Graphics Forum* (2020). first online Nov. 2020. 2
- [LDW97] LOUNSBERY M., DEROSE T. D., WARREN J.: Multiresolution analysis for surfaces of arbitrary topological type. *ACM Trans. on Graphics* 16, 1 (1997), 34–73. 2
- [Lév06] LÉVY B.: Laplace-Beltrami eigenfunctions towards an algorithm that understands geometry. In *Proc. SMI* (2006), pp. 13–25. 1
- [LRBB17] LITANY O., RODOLÀ E., BRONSTEIN A., BRONSTEIN M.: Fully spectral partial shape matching. *Computer Graphics Forum* 36, 2 (2017), 247–258. 2
- [LSS18] LU J., SOGGE C. D., STEINERBERGER S.: Approximating pointwise products of laplacian eigenfunctions, 2018. [arXiv:1811.10447](https://arxiv.org/abs/1811.10447). 4
- [Mel19] MELZI S.: Sparse representation of step functions on manifolds. *Computers & Graphics* 82 (2019), 117 – 128. 2
- [MMM*20] MELZI S., MARIN R., MUSONI P., BARDON F., TARINI M., CASTELLANI U.: Intrinsic/extrinsic embedding for functional remeshing of 3D shapes. *CAG* 88 (2020), 1 – 12. 2
- [MMR*19] MELZI S., MARIN R., RODOLÀ E., CASTELLANI U., REN J., POULENARD A., WONKA P., OVSJANIKOV M.: SHREC 2019: Matching Humans with Different Connectivity. In *Eurographics Workshop on 3D Object Retrieval* (2019), The Eurographics Association. 10
- [MOR*18] MELZI S., OVSJANIKOV M., ROFFO G., CRISTANI M., CASTELLANI U.: Discrete time evolution process descriptor for shape analysis and matching. *ACM Transactions on Graphics (TOG)* 37, 1 (Jan. 2018), 4:1–4:18. 11
- [MRCB16] MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M.: Shape analysis with anisotropic windowed fourier transform. In *International Conference on 3D Vision (3DV)* (2016). 9
- [MRCB18] MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M.: Localized manifold harmonics for spectral shape analysis. *Computer Graphics Forum* 37, 6 (2018), 20–34. 2
- [MRR*19] MELZI S., REN J., RODOLÀ E., SHARMA A., WONKA P., OVSJANIKOV M.: Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)* 38, 6 (Nov. 2019), 155:1–155:14. 2, 7, 11
- [NMR*18] NOGNENG D., MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M., OVSJANIKOV M.: Improved functional mappings via product preservation. *Computer Graphics Forum* 37, 2 (2018), 179–190. 2, 5, 6, 7, 8, 9, 10, 11
- [NO17] NOGNENG D., OVSJANIKOV M.: Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum* 36, 2 (2017), 259–267. 2, 3, 9
- [NVT*14] NEUMANN T., VARANASI K., THEOBALT C., MAGNOR M., WACKER M.: Compressed manifold modes for mesh processing. *Computer Graphics Forum* 33, 5 (2014), 35–44. 2
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 30:1–30:11. 2, 3, 7
- [OCB*16] OVSJANIKOV M., CORMAN E., BRONSTEIN M., RODOLÀ E., BEN-CHEN M., GUIBAS L., CHAZAL F., BRONSTEIN A.: Computing and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses* (New York, NY, USA, 2016), ACM, pp. 9:1–9:60. 2, 3
- [Pat18] PATANÈ G.: Laplacian spectral basis functions. *Computer-Aided Geometric Design* 65 (2018), 31 – 47. 2
- [PP93] PINKALL U., POLTHIER K.: Computing Discrete Minimal Surfaces and their Conjugates. *Exp. Math.* 2, 1 (1993), 15–36. 3
- [RCB*17] RODOLÀ E., COSMO L., BRONSTEIN M., TORSSELLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum* 36, 1 (2017), 222–236. 2
- [RPWO18] REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (TOG)* 37, 6 (2018). 2
- [Rus07] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP* (2007), Eurographics Association, pp. 225–233. 1
- [SK15] SHTERN A., KIMMEL R.: Spectral gradient fields embedding for nonrigid shape matching. *CVIU* 140 (2015), 21–29. 3
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. *Computer graphics forum* 28, 5 (2009), 1383–1392. 9
- [Ste02] STEWART G. W.: A krylov–schur algorithm for large eigenproblems. *SIAM J. Matrix Anal. Appl.* 23, 3 (2002), 601–614. 8
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *ACM SIGGRAPH* (1995), pp. 351–358. 1
- [TSDS10] TOMBARI F., SALTI S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *Proc. ECCV* (2010), Springer, pp. 356–369. 9
- [VL08] VALLET B., LÉVY B.: Spectral geometry processing with manifold harmonics. *Computer Graphics Forum* 27, 2 (2008), 251–260. 1, 8
- [Zho12] ZHONG M.: *Harmonic Shape Analysis: From Fourier to Wavelets*. Master’s thesis, Stony Brook University, 2012. 2

Appendix A: Proofs

Laplacian of a product

We want to show that, given a function $f(x) = \prod_{i=1}^n f_i(x)$, the application of the Laplace-Beltrami operator results in

$$\Delta f(x) = \sum_{i=1}^n \prod_{j=1, j \neq i}^n f_j(x) \Delta f_i(x) - \sum_{i,j=1}^n \prod_{h=1, h \neq i,j}^n f_h(x) \langle \nabla f_i(x), \nabla f_j(x) \rangle \quad (23)$$

We show it by induction, with the base case given by Equation (3) for $n = 2$. Assuming this true up to $n - 1$, for $g = \prod_{i=1}^{n-1} f_i$ we have

$$\Delta f(x) = f_n(x)\Delta g(x) + g(x)\Delta f_n(x) - 2\langle \nabla f_n(x), \nabla g(x) \rangle \quad (24)$$

The first term expands to

$$\sum_{i=1}^{n-1} \prod_{j=1, j \neq i}^n f_j(x)\Delta f_i(x) - \sum_{i,j=1}^{n-1} \prod_{h=1, h \neq i,j}^n f_h(x) \langle \nabla f_i(x), \nabla f_j(x) \rangle, \quad (25)$$

while the second term is $\prod_{i=1}^{n-1} f_i(x)\Delta f_n(x)$, and the third term is

$$-2 \sum_{j=1}^{n-1} \prod_{h=1, h \neq j,n}^n f_h(x) \langle \nabla f_j(x), \nabla f_n(x) \rangle \quad (26)$$

By adding these all together, we finally get back Equation (23).

Proof of Theorem 3.1

Proof For simplicity, we assume $\|\varphi_I\|^2 = 1$. Also, to keep the equations more readable, we will abuse of our notation. When we iterate over multiple indices and we indicate $i \neq j$ we do not mean that index i must be necessarily different from index j (since eigenproducts can involve the same eigenfunction multiple times), but that we cannot pick the i -th index two times.

By observing that

$$\Delta \varphi_I(x) = \sum_{i \in I} \lambda_i \varphi_I(x) - \sum_{i,j \in I} \prod_{h \in I, h \neq i,j} \varphi_h(x) \langle \nabla \varphi_i(x), \nabla \varphi_j(x) \rangle \quad (27)$$

and by knowing $\langle \nabla f, \nabla g \rangle_{T(\mathcal{M})} = \langle f, \Delta g \rangle_{\mathcal{M}}$, we easily get

$$\mathcal{E}(\varphi_I) = \sum_{i \in I} \lambda_i - \sum_{i,j \in I} \int_{\mathcal{M}} \prod_{h \in I, h \neq i} \varphi_h(x) \prod_{p \in I, p \neq j} \varphi_p(x) \langle \nabla \varphi_i(x), \nabla \varphi_j(x) \rangle dx \quad (28)$$

Now, let us consider the following terms:

$$\sum_{i \in I} \int_{\mathcal{M}} \prod_{j \in I, j \neq i} \varphi_j^2(x) \langle \nabla \varphi_i(x), \nabla \varphi_i(x) \rangle dx. \quad (29)$$

Since the integrands are always non-negative, the whole sum is non-negative. Hence, we can give a lower bound to Equation (28):

$$\mathcal{E}(\varphi_I) \geq \sum_{i \in I} \lambda_i - \sum_{i,j \in I} \int_{\mathcal{M}} \prod_{h \in I, h \neq i} \varphi_h(x) \prod_{p \in I, p \neq j} \varphi_p(x) \langle \nabla \varphi_i(x), \nabla \varphi_j(x) \rangle dx \quad (30)$$

The second term is the Dirichlet energy of φ_I . In fact, we have:

$$\begin{aligned} \mathcal{E}(\varphi_I) &= \langle \nabla \varphi_I, \nabla \varphi_I \rangle_{T(\mathcal{M})} = \\ &= \sum_{i,j \in I} \left\langle \prod_{h \in I, h \neq i} \varphi_h \nabla \varphi_i, \prod_{p \in I, p \neq j} \varphi_p \nabla \varphi_j \right\rangle_{T(\mathcal{M})} = \\ &= \sum_{i,j \in I} \int_{\mathcal{M}} \prod_{h \in I, h \neq i} \varphi_h(x) \prod_{p \in I, p \neq j} \varphi_p(x) \langle \nabla \varphi_i(x), \nabla \varphi_j(x) \rangle dx \end{aligned} \quad (31)$$

By plugging it into Equation (30), we get:

$$\mathcal{E}(\varphi_I) \geq \sum_{i \in I} \lambda_i - \mathcal{E}(\varphi_I), \quad (32)$$

which results in the claim.

Consider now the special case of $\varphi_I = \varphi_I^n$. In this case, all the integrals in Equation (28) are equal and can be written as

$$\int_{\mathcal{M}} \langle \varphi_I^{n-1}(x) \nabla \varphi_I(x), \varphi_I^{n-1}(x) \nabla \varphi_I(x) \rangle dx \quad (33)$$

which, recalling $\frac{1}{\alpha} \nabla f^\alpha(x) = f^{\alpha-1}(x) \nabla f(x)$, gives

$$\int_{\mathcal{M}} \left\| \varphi_I^{n-1}(x) \nabla \varphi_I(x) \right\|^2 dx = \frac{1}{n^2} \int_{\mathcal{M}} \left\| \nabla \varphi_I^n(x) \right\|^2 dx = \frac{1}{n^2} \mathcal{E}(\varphi_I^n) \quad (34)$$

Since in Equation (28) there is a sum over $i \neq j$, we have a total of $n(n-1)$ of these terms. Hence, we are left with

$$\mathcal{E}(\varphi_I^n) = n\lambda_i - \frac{n-1}{n} \mathcal{E}(\varphi_I^n) \quad (35)$$

which, with simple algebraic manipulations, leads to

$$\mathcal{E}(\varphi_I^n) = \frac{n^2}{2n-1} \lambda_i. \quad (36)$$

□

Proof of Corollary 3.1.1

Proof We know that the set of eigenproducts contains all the eigenpowers, and in particular $\varphi_K^N \in \Phi$. Thus, the following holds:

$$\max_{\varphi \in \Phi} \{\mathcal{E}(\varphi)\} \geq \mathcal{E}(\varphi_K^N) = \frac{N^2}{2N-1} \lambda_K. \quad (37)$$

By Weyl's asymptotic law for 2-dimensional manifolds we have $\lambda_K \in \Theta(K)$, therefore $\mathcal{E}(\varphi_K^N) \in \Theta(NK)$. Since the maximum Dirichlet energy is lower bounded by this value, it must be that $\max_{\varphi \in \Phi} \{\mathcal{E}(\varphi)\} \in \Omega(NK)$. □