

Structured Regularization of Functional Map Computations

Jing Ren¹, Mikhail Panine², Peter Wonka¹, and Maks Ovsjanikov²

¹KAUST, ²LIX, École Polytechnique, CNRS

Abstract

We consider the problem of non-rigid shape matching using the functional map framework. Specifically, we analyze a commonly used approach for regularizing functional maps, which consists in penalizing the failure of the unknown map to commute with the Laplace-Beltrami operators on the source and target shapes. We show that this approach has certain undesirable fundamental theoretical limitations, and can be undefined even for trivial maps in the smooth setting. Instead we propose a novel, theoretically well-justified approach for regularizing functional maps, by using the notion of the resolvent of the Laplacian operator. In addition, we provide a natural one-parameter family of regularizers, that can be easily tuned depending on the expected approximate isometry of the input shape pair. We show on a wide range of shape correspondence scenarios that our novel regularization leads to an improvement in the quality of the estimated functional, and ultimately pointwise correspondences before and after commonly-used refinement techniques.

CCS Concepts

- *Computing methodologies* → *Shape analysis*;

1. Introduction

Shape matching is a fundamental problem in geometry processing and computer graphics, with applications ranging from shape interpolation [HRS*16] to shape exploration [HWG14] and statistical shape analysis [BRLB14].

In this paper, we concentrate on the functional maps framework for computing correspondences between shapes, which has proven to be especially useful when dealing with non-rigid and especially near-isometric shape pairs. This approach, originally introduced in [OBCS*12] and subsequently extended in multiple follow-up works, e.g. [KBB*13, RCB*17, HWG14], is based on optimizing for a linear mapping between function spaces defined on the shapes, which can be conveniently encoded as a matrix. One attractive feature of this approach is that many desirable geometric objectives for the unknown pointwise correspondence, such as preservation of geodesic distances, can be conveniently encoded as constraints on the matrices representing functional maps, and often lead to simple convex optimization problems.

The majority of objectives when optimizing for a functional map between a pair of shapes consist of a) preservation of some pre-computed descriptors and b) a functional map regularization term, based on promoting some desired global map properties. This latter, regularization term is especially important for ensuring the overall global consistency of the computed map.

Common strategies for regularization include commutativity with the Laplacian operator, exploited in the original article [OBCS*12], a mask promoting a slanted diagonal in the case of

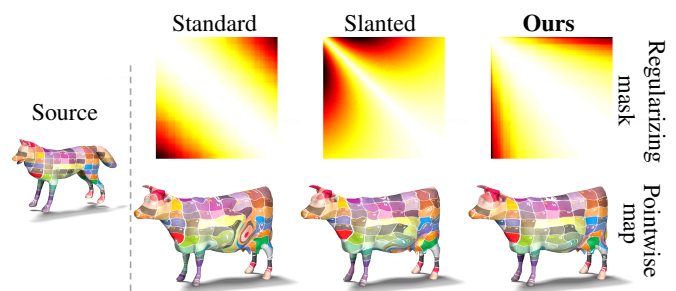


Figure 1: Comparison of different masks. Standard: the Laplacian commutativity operator can be equivalently formulated as a penalty or regularizing mask. Slanted: the weight mask proposed in [RCB*17] designed to promote a slanted structure. Ours: the mask proposed in this paper based on the resolvent of the Laplacian operator. The penalty masks of these three methods are visualized in the first row for an example pair from the SHREC dataset, and the corresponding optimized point-wise maps from the functional map pipeline are shown in the second row.

partial correspondences [RCB*17], or orthonormality of the functional map, which corresponds to local area preservation by the underlying pointwise map [OBCS*12, KBB*13, ROA*13].

Unfortunately, although functional map regularization has been instrumental in obtaining high-quality results on some shape correspondence benchmarks, these regularization terms often lack a theoretical foundation and indeed in some cases result in optimization

objectives which would not be well-defined in the case of smooth surfaces.

In this paper, we introduce a novel way of regularizing functional maps, based on the concept of the resolvent of the Laplacian operator. Unlike the original approach, based on commutativity with the Laplacians, our method has a natural theoretically well-defined analogue in the case of smooth surfaces, and is guaranteed to be bounded even in the full (infinite dimensional) basis case. Moreover, our approach provides a simple one-parameter family of regularizers, that can be tuned depending on the expected approximate deviation from isometry in a given shape pair. Finally, we demonstrate on a wide range of challenging settings that our approach leads to a quantitative and qualitative improvement for the computed functional and eventually pointwise maps, without incurring any additional computational complexity (see Fig. 1 as an example).

2. Related Work

Shape matching in its full generality is an extremely well-studied area of geometry processing and computer graphics, and its full overview is beyond the scope of our paper. Therefore, below we only review the most closely related methods based primarily on the functional maps framework. We refer the interested readers to recent surveys including [VKZHCO11, TCL*13, BCB16] for an in-depth treatment of other shape matching approaches.

Functional Maps. Our approach fits within the functional map framework, which was originally introduced in [OBCS*12] for solving non-rigid shape matching problems, and extended significantly in follow-up works, including [KBB*13, ADK16, KBBV15, RCB*17, EBC17, BDK17] among many others (see also [OCB*17] for an overview). The key observation in these techniques is that it is often easier to estimate correspondences between real-valued functions, rather than points on the shapes. This is because functions have a natural vector space structure, and moreover functional transformations are most often linear, which means that functional maps can be encoded, and thus optimized for, as small matrices in a reduced functional basis.

As observed by several works in this domain, [KBB*13, ROA*13, RCB*17, BDK17] many natural properties on the underlying pointwise correspondences can be expressed as objectives on functional maps. Most notably, this includes: orthonormality of functional maps, which corresponds to the local area-preservation nature of pointwise correspondences [OBCS*12, KBB*13, ROA*13]; preservation of inner products of gradients of functions, which corresponds to conformal maps [ROA*13, BDK17, WLZT18]; preservation of *pointwise products* of functions, which corresponds to functional maps arising from point-to-point correspondences [NO17, NMR*18]; slanted diagonal structure of functional maps, which corresponds to correspondences between partial shapes [RCB*17, LRBB17]. Similarly, several other regularizers have been proposed, including using robust norms and matrix completion techniques [KBB*13, KBBV15], exploiting the relation between functional maps in different directions [ERGB16], the map adjoint [HO17], and powerful cycle-consistency constraints [HWG14] in the context of shape collections, among many

others. More recently constraints on functional maps have been introduced to promote *continuity* of the recovered pointwise correspondence [PSO18], maps between curves defined on shapes [GBKS18], kernel-based techniques aimed at extracting more information from given descriptor constraints [WGBS18], and an approach for incorporating *orientation* information into the functional map inference pipeline [RPWO18] among others.

Among all of these, perhaps the most widely-used building block for regularizing functional maps, introduced in [OBCS*12] and extended in follow-up works, including [WHG13, RCB*17, LRB*16, LRBB17], is based on the commutativity with the Laplacian operators, which implies a diagonal (or slanted diagonal in the case of partial correspondence) structure for functional maps. To promote this structure, the most common method (see also Chapter 2.4 in [OCB*17]) consists in adding an energy to the functional map estimation pipeline, which penalizes the failure of the unknown functional map to commute with the Laplace-Beltrami operators on the source and target shapes. Conveniently, this term still leads to a convex optimization problem during functional map estimation. Moreover, as observed in [OBCS*12], for functional maps arising from pointwise ones, this term is zero if and only if the map preserves geodesic distances exactly.

Unfortunately, although functional map regularization via commutativity with the Laplace-Beltrami operators has been instrumental in obtaining high quality results in challenging benchmarks, the exact properties of this regularization are not well-understood. In particular, as we show below, the commonly-used energy is not bounded in the full (infinite-dimensional) basis in general. Instead, our novel regularizer, although based on a similar underlying principle, overcomes this limitation, and both leads to a theoretically better justified energy, and a practical improvement on a range of challenging datasets.

Optimal Transport. We also note briefly that other commonly-used relaxations for matching problems include those based on optimal transport, e.g. [SPKS16, MCSK*17, VLR*17]. These techniques have recently gained prominence especially due to the computational advances for addressing large-scale transport problems, primarily using the Sinkhorn method with entropic regularization [Cut13, SDGP*15]. Furthermore, other techniques that exploit the formalism of optimal transport, for solving assignment problems include the recent variants of Product Manifold Filter using Gaussian and Heat Kernels [MCSK*17, VLR*17]. Interestingly, these latter methods argue that near-isometric shape matching can be better addressed via preservation of general functional *kernels* rather than preservation of e.g. geodesic distances. Our modification of the regularization term of functional maps can also be seen through a similar light, as we show that a more theoretically justified functional operator leads to an improvement of the overall structural properties of the map, which eventually result in more accurate functional and pointwise maps.

3. Background

Our work is based on the functional map representation and the estimation pipeline. Below we review the basic notions and the main steps for estimating a map between a pair of shapes using

this framework, and refer the interested reader to a recent set of course notes [OCB*17] for a more in-depth discussion.

Basic Pipeline. Given a pair of shapes, S_1, S_2 represented as triangle meshes, and containing, respectively, n_1 and n_2 vertices, the general pipeline for computing a map between them using the functional map representation, consists of the following main steps:

1. Compute a small set of $k_1 \ll n_1$ and $k_2 \ll n_2$ basis functions on each shape. The most common choice consists in using the first k eigenfunctions of the Laplace-Beltrami operator of each shape, although other bases derived from the Hamiltonian operator [CSBK18] and more localized basis functions [NVT*14, MRCB18] have also been used.
2. Compute a set of descriptor functions on each shape, that are expected to be approximately preserved by the unknown map. Store their coefficients in the corresponding bases as columns of matrices A_1, A_2 .
3. Compute the optimal functional map C by solving the following optimization problem:

$$C_{\text{opt}} = \arg \min_{C_{12}} E_{\text{desc}}(C_{12}) + \alpha E_{\text{reg}}(C_{12}) \quad (1)$$

where the first term aims at the descriptor preservation: $E_{\text{desc}}(C_{12}) = \|C_{12}A_1 - A_2\|^2$, whereas the second term regularizes the map by promoting the correctness of its overall structural properties. As mentioned above, the most common approach consists of penalizing the failure of the unknown functional map to commute with the Laplace-Beltrami operators, which can be written as:

$$E_{\text{reg}}(C_{12}) = E_{\text{comm}}(C_{12}) = \|C_{12}\Delta_1 - \Delta_2 C_{12}\|^2 \quad (2)$$

where Δ_1 and Δ_2 are the Laplace-Beltrami operators of the two shapes expressed in the respective bases. Here, and throughout the rest of our paper, unless stated otherwise $\|\cdot\|$ denotes the matrix Frobenius norm.

4. Convert the functional map C to a point-to-point map, for example using an iterative approach, such as the Iterative Closest Point (ICP) in the spectral embedding, or using other more advanced techniques [RMC15, EBC17].

One of the attractive properties of this pipeline is that the functional map computation in step 3 leads to a simple (convex) least squares optimization with a relatively small number of unknowns, independent of the number of points on the shapes. This step has been further extended e.g. using more powerful descriptor preservation constraints via commutativity [NO17], or using manifold optimization [KGB16] among many others (see also Chapter 3 in [OCB*17]).

4. Functional Map Regularization

Our main goals are to analyze the commonly-used functional map regularization term, to bring attention to some of its theoretical limitations, and to propose a novel regularizer with better theoretical properties, which lead to practical improvements. Therefore, in this in this section, we first consider the standard Laplacian-commutativity term, and then introduce our new regularizer based on the resolvent operator of the Laplacian.

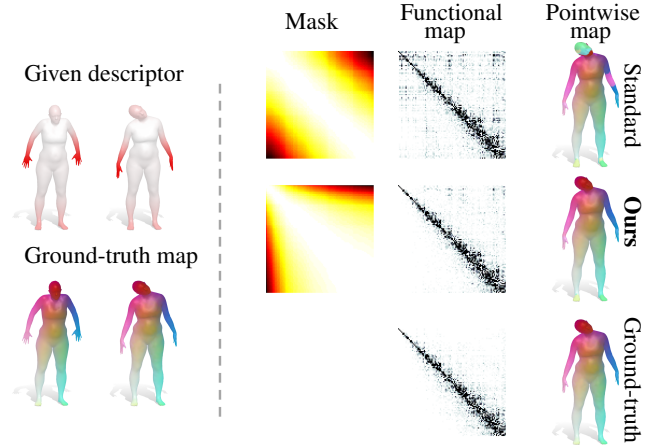


Figure 2: Given a single pair of WKS descriptors, we optimize a 100-by-100 functional map using the standard mask and our resolvent mask and compare to ground truth. The converted point-wise maps are shown on the right. We can see that the functional map stemming from our resolvent mask has less noise than the functional map computed with the standard mask. Also, our resolvent mask leads to a point-wise map with better quality.

4.1. Reformulation of the Laplacian-Commutativity term

As mentioned above, the Laplacian-commutativity term given in Eq. (2) was first introduced to promote approximately isometric correspondences. If the functional map is expressed in the basis given by the eigenfunctions of the Laplace-Beltrami operator, and letting Λ_1 and Λ_2 represent vectors that store the eigenvalues of the Laplacians of shape S_1 and S_2 respectively, this term can be equivalently reformulated as:

$$\begin{aligned} E_{\text{comm}}(C_{12}) &= \|C_{12}\Delta_1 - \Delta_2 C_{12}\|^2 \\ &= \|C_{12}\text{diag}(\Lambda_1) - \text{diag}(\Lambda_2)C_{12}\|^2 \\ &= \|C_{12} \odot (\mathbf{1}_{k_2}\Lambda_1^T) - (\Lambda_2\mathbf{1}_{k_1}^T) \odot C_{12}\|^2 \quad (3) \\ &= \|(\mathbf{1}_{k_2}\Lambda_1^T - \Lambda_2\mathbf{1}_{k_1}^T) \odot C_{12}\|^2 \\ &\triangleq E_{\text{mask}}(C_{12}) = \sum_{ij} [M_{LB}]_{ij} [C_{12}]_{ij}^2, \end{aligned}$$

where \odot is the entry-wise matrix multiplication, $\mathbf{1}_k$ is a k -dim all-ones vector and $[A]_{ij}$ denotes the (i, j) entry of a matrix A .

In other words, the regularization term E_{comm} can be written as a product between the matrix M_{LB} , which we call the penalty mask matrix and the squares of the entries of the unknown functional map C_{12} . In the case of E_{comm} , the entries of M_{LB} are given as $M_{LB}(i, j) = (\Lambda_2(i) - \Lambda_1(j))^2$. Fig. 2 shows a heat map of this matrix (see the first row in the “Mask” column).

Unfortunately, this simple formulation has certain fundamental undesirable properties. In particular, in the case of smooth surfaces, Laplace-Beltrami operators are *unbounded* operators on square-integrable functions [MP49]. Consequently, in general, the energy term $\|C_{12}\Delta_1 - \Delta_2 C_{12}\|^2$ is undefined on smooth surfaces. As a simple example, consider the situation where $C_{12} = Id$, the identity operator, and the two surfaces are scaled versions of each other

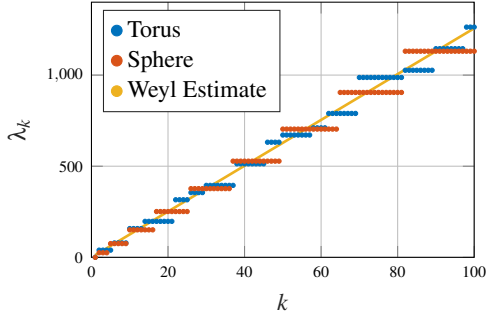


Figure 3: The spectra of a sphere and a square torus of unit area, as well as the corresponding Weyl estimate.

(i.e., $\Delta_2 = c\Delta_1$ for some constant $c \neq 1$). In this case, $\|C_{12}\Delta_1 - \Delta_2 C_{12}\|^2 = |1 - c|^2 \|\Delta_1\|^2$ is infinite.

That being said, the ill-definiteness of $\|C_{12}\Delta_1 - \Delta_2 C_{12}\|^2$ is not a mere question of scale. Recall that, by Weyl’s law [Dod81], large Laplacian eigenvalues of surfaces can be estimated in terms of the surface area S as follows:

$$\lambda_k \sim \frac{4\pi}{S} k. \quad (4)$$

Thus, by rescaling the surfaces such that their areas match guarantees that their eigenvalues have comparable asymptotic growth. This, however, is not sufficient to make $\|C_{12}\Delta_1 - \Delta_2 C_{12}\|^2$ well-defined. This can be seen from another simple, if slightly artificial, example. Consider a sphere (\mathcal{M}_1) and a flat square torus (\mathcal{M}_2), both of unit area. While these surfaces are not diffeomorphic, they are among the few whose Laplace-Beltrami eigenvalues can be explicitly computed (see [Sau06a, Sau06b] among many others), which is why we consider them here. Said spectra, as well as the corresponding Weyl estimate are illustrated in Fig. 3. For simplicity’s sake, we once again pick a functional map of the form $C_{12} = Id$. We illustrate the spectra, as well as the corresponding $\|C_{12}\Delta_1 - \Delta_2 C_{12}\|^2$ on Fig. 4. Notice the divergence of the Laplacian commutator energy. On the same figure, we illustrate its proposed replacement, defined in the next section. Note its rapid convergence.

In addition to being ill-defined in the continuous setting, the Laplacian commutativity mask has another significant problem. Namely, it penalizes the high frequency entries of functional maps in the same way as the low frequency ones, despite the greater instability of the former, and in this way fails to exhibit the funnel-like structure observed in ground-truth functional maps (see e.g. Fig. 2 above or Figure 4 in [OBBS*12]). Recall that a mask serves as a penalty during the optimization of the functional map matrix. Thus, the mask and the ground truth functional map should be complementary, in the sense that for the regions of ground-truth functional maps with smaller (resp. large) values, the mask should add more (resp. less) penalty.

We illustrate this phenomenon in Fig. 5, which shows the average of the squared values of functional maps constructed from the ground-truth correspondences using 250 eigenfunctions over 100 FAUST non-isometric shape pairs. We then compare its structure

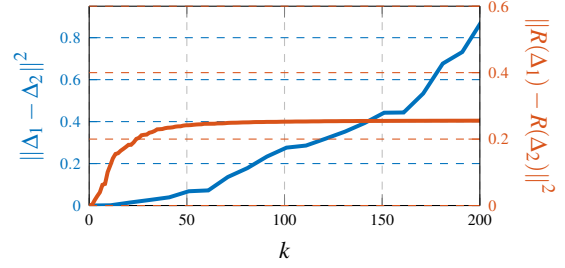


Figure 4: The Laplacian and Resolvent commutator energies for $C_{12} = Id$ computed using the lowest k eigenvalues of the unit area sphere and torus (see Fig. 3). Note the rapid convergence of $\|C_{12}R(\Delta_1) - R(\Delta_2)C_{12}\|^2$ and the divergence of $\|C_{12}\Delta_1 - \Delta_2 C_{12}\|^2$. The spectra are rescaled with respect to the largest eigenvalue with $k = 100$, as described in Eq. (14).

to the Laplacian commutativity mask (labeled “Standard”), averaged over the same shape pairs. We perform the same computation for the heuristic slanted mask introduced in [RCB*17] and, finally, for our proposed resolvent-based mask defined in the next section. Notice that our mask better reproduces the funnel-like structure of the ground-truth maps.

4.2. A novel regularization based on the Resolvent

In this section, we propose an alternative form for the Laplacian commutativity regularizer that overcomes the problems identified in the previous section. The resulting regularizer is better from both the theoretical and practical standpoints.

The first issue of the original Laplacian-commutativity term arises from the fact that the Laplace-Beltrami operator is unbounded. We thus propose to replace it with a meaningful bounded operator, namely its resolvent. Below, we give a brief overview and refer the reader to [RS80] for the detailed functional-analytic underpinnings of the following discussion.

Resolvent. Let $A : \mathcal{D} \rightarrow \mathcal{H}$ be a densely-defined closed operator on some Hilbert space \mathcal{H} with domain $\mathcal{D} \subset \mathcal{H}$. Let $\rho(A) \subset \mathbb{C}$ be the set of all complex numbers μ such that $R_\mu(A) = (A - \mu Id)^{-1}$ is defined and bounded. The set $\rho(A)$ and the operator $R_\mu(A)$ are known as the resolvent set of A and the resolvent (operator) of A at μ , respectively. The resolvent set $\rho(A)$ is the complement of the spectrum of A in the complex plane.

The general idea of considering the resolvent of an unbounded operator rather than the operator itself comes from the notion of norm-resolvent convergence, which is used to study the convergence of unbounded self-adjoint operators [RS80]. In that sense, our choice to use the resolvent of the Laplace-Beltrami operator is a natural one. Note that here, closedness is a technical condition used in the definition of the resolvent. In particular, it is satisfied by self-adjoint operators, such as the ones we consider.

Before proceeding further, we slightly generalize our problem. In what follows we will consider powers of the Laplacian rather than the Laplacian itself. Specifically, we will use Δ^γ for some $\gamma >$




Figure 5: The average squared ground-truth functional map and three penalty masks over 100 FAUST non-isometric shape pairs.

0 rather than Δ . As explained below, the parameter γ controls the funnel-like structure of the resulting mask. Thus, this parameter takes on an important role in the numerical tests reported later in this paper.

Now since the Laplace-Beltrami operator Δ is positive and self-adjoint, its spectrum is contained in the non-negative real line. The same also holds for Δ^γ . Thus, we are guaranteed that $R_\mu(\Delta)$ is a well-defined bounded operator for any complex μ not in the non-negative real line.

For our purposes, the resolvent of Δ^γ is expressed as

$$R(\Delta^\gamma) = (\Delta^\gamma - (a + ib)Id)^{-1}, \quad (5)$$

where i is the imaginary unit and $a, b \in \mathbb{R}$. Our proposal is to use the resolvent of the Laplacian instead of the Laplacian in the commutator term. We thus define a new energy term,

$$E_{\text{resolvent}}(C_{12}) = \|C_{12}R(\Delta_1^\gamma) - R(\Delta_2^\gamma)C_{12}\|^2. \quad (6)$$

Before proceeding further, recall that we only need to consider bounded functional maps C_{12} , since functional maps that arise as pullbacks of diffeomorphisms are bounded. This last fact is shown in appendix A, for completeness. Our proposal is motivated by the following result.

Theorem 1 (Bounded Resolvent Commutativity) Let C_{12} be a bounded functional map. Then, in the operator norm,

$$\|C_{12}R(\Delta_1^\gamma) - R(\Delta_2^\gamma)C_{12}\|_{op}^2 < \infty. \quad (7)$$

Proof The result directly follows from the fact that products and linear combinations of bounded operators are bounded. As defined above, $R(\Delta^\gamma)$ is bounded, which concludes the proof. \square

As explained above, there is no analogous guarantee for the commonly used Laplacian commutativity, since the Laplace-Beltrami operator itself is unbounded. In contrast, our resolvent-based commutativity is always well-defined, and leads to a commutator with bounded operator norm, even in the case of smooth surfaces.

Notice that Theorem 1 holds in the operator norm, rather than the more convenient Frobenius norm, which we use to define the energy in Equation (6). Our usage of the Frobenius norm can be justified in two ways. First, a Frobenius norm version of Theorem 1 holds for $\gamma > 1/2$. This is shown in Lemma 2 of Appendix B. Thus, for $\gamma > 1/2$, the usage of the Frobenius norm is perfectly justified in the smooth setting. For lower values of γ , we invoke the fact that, in practice, we work on a finite dimensional vector space. Since all norms are equivalent on finite dimensional vector spaces,

we are free to replace the operator norm with the Frobenius one. In that case, we lose the guarantee that the energy $E_{\text{resolvent}}$ makes sense in the full Laplace-Beltrami basis in the smooth setting. Yet, as shown in the experiments reported below, this does not seem to cause any issue.

Similarly to the usual Laplacian commutativity term E_{comm} , $E_{\text{resolvent}}$ can also be expressed as a mask matrix. This can be done as follows. Notice that $R(\Delta^\gamma)$ is diagonal in the Laplacian eigenbasis. Moreover, if Δ has eigenvalues $\{\lambda_n\}_{n=0}^\infty$, then, by matrix inversion, the eigenvalues $\{r_n\}_{n=0}^\infty$ of $R(\Delta^\gamma)$ will be given by:

$$r_n = \frac{1}{\lambda_n^\gamma - a - ib} = \frac{\lambda_n^\gamma - a}{(\lambda_n^\gamma - a)^2 + b^2} + \frac{b}{(\lambda_n^\gamma - a)^2 + b^2}i. \quad (8)$$

In practice, we choose $a = 0$ and $b = 1$, which yields

$$r_n = \frac{\lambda_n^\gamma}{(\lambda_n^\gamma)^2 + 1} + \frac{1}{(\lambda_n^\gamma)^2 + 1}i. \quad (9)$$

Since the square of the Frobenius norm independently considers the real and imaginary parts of a matrix, we can re-express Equation (6) in terms of two new matrices M_{Re} and M_{Im} , defined below.

$$\begin{aligned} E_{\text{resolvent}}(C_{12}) &= \|M_{Re} \odot C_{12}\|^2 + \|M_{Im} \odot C_{12}\|^2 \\ &= \sum_{ij} [M_{Re}]_{ij}^2 [C_{12}]_{ij}^2 + \sum_{ij} [M_{Im}]_{ij}^2 [C_{12}]_{ij}^2 \\ &= \sum_{ij} \left([M_{Re}]_{ij}^2 + [M_{Im}]_{ij}^2 \right) [C_{12}]_{ij}^2 \end{aligned} \quad (10)$$

The matrices M_{Re} and M_{Im} correspond to the real and imaginary parts of the eigenvalues of the resolvent, respectively. Explicitly, these matrices are given by

$$M_{Re}(i, j) = \frac{\Lambda_2(i)^\gamma}{\Lambda_2(i)^{2\gamma} + 1} - \frac{\Lambda_1(j)^\gamma}{\Lambda_1(j)^{2\gamma} + 1} \quad (11)$$

$$M_{Im}(i, j) = \frac{1}{\Lambda_2(i)^{2\gamma} + 1} - \frac{1}{\Lambda_1(j)^{2\gamma} + 1} \quad (12)$$

These matrices are not quite masks in the sense of Eq. (3). As per Equation (10), the above two matrices can be combined into a single mask

$$M_{res}(i, j) = M_{Re}(i, j)^2 + M_{Im}(i, j)^2. \quad (13)$$

The split of M_{res} into M_{Re} and M_{Im} will be revisited later, when we explore beyond the established theory and consider a mask constructed from weighted combinations of these two matrices.

4.3. Rescaling the Spectra

As it stands now, we observe that in practice the mask defined in equation (13) decays too quickly as the Laplacian eigenvalues grow. In other words, the mask is not sufficiently sensitive to the

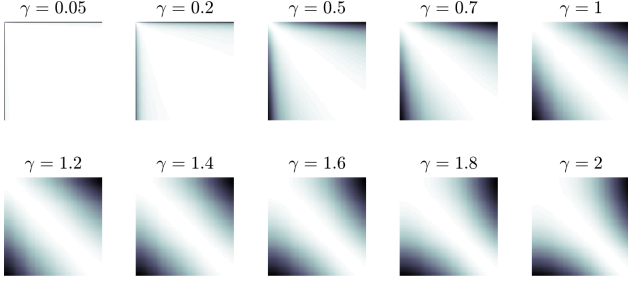


Figure 6: The resolvent mask introduced in Eq. (13) with different γ . The darker the region, the larger the penalty. Here, Δ_1 and Δ_2 are near-isospectral, which explains the approximately symmetric structure of the mask.

higher frequencies. This is due to the scale introduced by the parameter b found in the definition of the resolvent. There are a few equivalent ways of addressing this issue.

Our approach is as follows. We begin by computing the k lowest eigenvalues of the considered Laplacians. Then, both spectra are rescaled according to the rule

$$\Lambda_i \mapsto \frac{\Lambda_i}{\max(\max(\Lambda_1), \max(\Lambda_2))}. \quad (14)$$

In other words, the spectra of the source and target shape are rescaled by the same factor, such that the largest considered eigenvalue (over both spectra) becomes equal to 1. From Equation (8), one can see that this rescaling can be absorbed into a choice of b and a change in the weight of the resolvent energy in the overall energy. Thus, conceptually, the spectral rescaling is equivalent to a choice of the parameter b used in the definition of the resolvent. Alternatively, recall that rescaling the spectrum is equivalent to rescaling the surface. Consequently, this rescaling does not affect the theoretical guarantees of the previous section.

Returning to the previously discussed example of the sphere and torus of unit area (see Fig. 3), we illustrate the resolvent energy computed by the above procedure in Fig. 4. Note the rapid convergence of the resolvent commutator energy (in red) and the divergence of the Laplacian commutator energy (in blue).

4.4. Mask Structure as a Function of γ

Having defined the way to compute the resolvent mask, we are now ready to explore the way in which tuning the parameter γ controls its the funnel-like structure.

Fig. 6 illustrates the resolvent mask for different values of γ . The darker the region, the more penalized the corresponding entry in the functional map C_{12} would be. Notice that the funnel-like structure of the mask changes with different values of γ .

From Fig. 6, we see that the behaviour of the mask as a function of γ can be separated into two regimes. The first corresponds to $\gamma \in (0, 1]$. There, increasing γ results in a narrowing of the funnel-like shape of the mask. Thus, the larger γ is, the more the mask pe-

nalizes functional maps that take eigenfunctions of Δ_1 to eigenfunctions of Δ_2 with distant eigenvalues. Correspondingly, small values of γ are more lax in that regard, allowing for maps between eigenfunctions with quite different eigenvalues. The former choice of γ is appropriate when the shapes under consideration are approximately isometric, as one then expects the eigenfunctions and eigenvalues of both surfaces to be roughly the same. The latter choice is sound for shape pairs that are further away from isometry.

The second regime corresponds to $\gamma > 1$. There, we observe an inversion of the funnel-like structure, with the reverse funnel shape being more and more pronounced as γ increases. This results in a low penalty for maps between low frequency eigenfunctions, which does not respect the shape of the ground-truth maps (see Fig. 5).

Later in this paper, we report empirically obtained optimal values for γ on a benchmark dataset (see Fig. 14).

4.5. Proposed Functional Map Energy

In summary, we propose the following energy to compute a functional map for a pair of shapes, with a set of given descriptors:

$$E(C_{12}) = \alpha_1 E_{\text{desc}} + \alpha_2 E_{\text{mult}} + \alpha_3 E_{\text{orient}} + \alpha_4 E_{\text{resolvent}} \quad (15)$$

where E_{desc} is the descriptor-preserving term defined in Sec. 3, E_{mult} is the descriptor-commutativity term defined as $E_{\text{mult}} = \sum_i \|C_{12} D_{1i}^{\text{mult}} - D_{2i}^{\text{mult}} C_{12}\|^2$ introduced in [NO17], E_{orient} is the orientation-preserving term introduced in [RPWO18]. $E_{\text{resolvent}}$ is the our new term introduced in Eq. (6) and discussed above. As mentioned above, when functional maps are expressed in the Laplace-Beltrami eigenbasis, this term can be written via a penalty using the mask matrix given in Eq. (13).

Before proceeding to a more extensive evaluation of the proposed energy, in Fig. 2 above, we provide an example of a functional map obtained using this energy and compare it to one obtained using the standard Laplacian commutator regularizer. Fig. 2 also shows the resolvent mask with $\gamma = 0.5$ (second row), compared to the standard mask (first row). The rightmost column shows the quality of pointwise maps, recovered from the functional maps (shown in the second rightmost column) using both our and the standard Laplacian regularizers.

5. Results

We tested the proposed approach using a MATLAB-based implementation, which we adapt from the state-of-the-art functional map approach in [RPWO18]. Here, we first describe the benchmark datasets and the baseline methods.

Datasets. We use the two datasets introduced in [RPWO18]. These consist of shapes from the FAUST [BRLB14] and TOSCA [BBK08] datasets, which were remeshed so that the shapes have different triangulations, and are no longer in one-to-one correspondence, making the matching more challenging and realistic. Specifically, we include 300 FAUST shape pairs and 284 TOSCA pairs for evaluation.

Baselines. To evaluate our new regularizer, we compare to the following two masks:

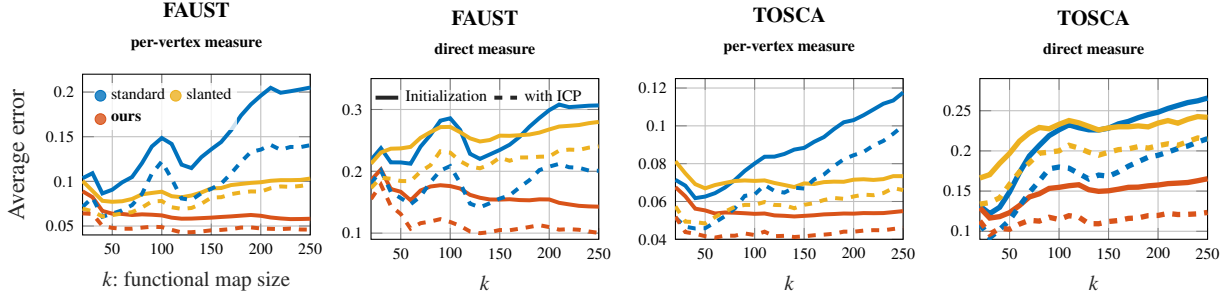


Figure 7: Changing the functional map size. We randomly select 50 FAUST non-isometric pairs and 50 TOSCA isometric pairs. For each of the pairs, we only use one pair of WKS [ASC11] descriptors. We then optimize for a functional map with different size ranging from 20 to 250 using different Laplacian mask terms. The solid lines represent the initialization with different masks and the dashed lines are the ICP-refined results. We report the per-vertex and the direct error measure. We can see that the proposed mask is much more stable than the standard and the slanted mask as the functional map size increases.

- "standard": the standard Laplacian-commutativity mask, which is defined as $M_{LB}(i, j) = (\Delta_2(i) - \Delta_1(j))$ as discussed before.
- "slanted": the heuristic slanted diagonal penalty mask proposed in [RCB*17], which is defined as

$$M(i, j) = \exp\left(-\eta\sqrt{i^2 + j^2}\right) \left\| \frac{\mathbf{n}}{\|\mathbf{n}\|} \times \left((i, j)^T - \mathbf{p}\right) \right\|$$

where $\mathbf{p} = (1, 1)^T$, and $\mathbf{n} = (1, r/k)^T$ is the line direction with slope r/k , where r is the estimated rank of the functional map, and k is the size of the square functional map. This weight matrix is originally applied to partial shape matching in [RCB*17], and we use it as a mask matrix to regularize the functional map. In our tests, η is set to the default value 0.03 as suggested in the original paper. This mask is illustrated on Fig. 5. Note that it exhibits the desired funnel-like structure for the lower part of the spectrum, but not the upper part.

We also compare three different settings:

- **Initialization.** The wave kernel signatures (WKS) [ASC11] are used to construct E_{desc} , E_{mult} , and E_{orient} in Eq. (15). Then we optimize the functional map w.r.t. the energy defined in Eq. (15) with three different masks, namely, the standard, slanted, and our resolvent mask.
- **ICP refinement.** After initialization, we use ICP [OBBS*12] to refine the computed functional maps.
- **BCICP refinement.** After initialization, we use the recently proposed BCICP algorithm [RPWO18] to refine the computed functional and pointwise maps, using the open-source implementation and default parameters provided by the authors.

Measurements. In our experiments, we measured the quality of the functional maps and the recovered point-wise maps:

- **Point-wise maps.** Since most shapes contain left-right symmetries, which are indistinguishable for intrinsic methods, in each dataset, we considered both the ground-truth direct and symmetric correspondences. To measure the accuracy of a computed map, we used the following measures:
 - **per-vertex error:** for each vertex we accept the ground-truth direct and symmetric correspondences and take the minimum

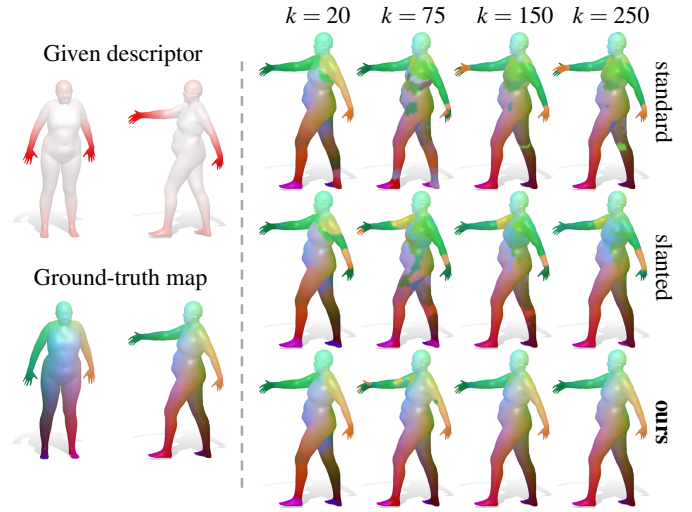


Figure 8: Given one pair of WKS descriptors, as visualized on the left top, we use different Laplacian mask terms to optimize for a functional map with size k . The ground-truth map is visualized on the bottom left. We can see that our mask is much more stable over different size k .

as the error of this vertex. This measure reflects the accuracy of the map regardless of the symmetry.

- **direct error:** we compute the average per-vertex error to the direct ground-truth correspondences only. This measure reflects both the accuracy and the smoothness of the map.
- **Functional maps.** We can evaluate a functional map by the quality of its recovered point-wise map, or by measuring the penalty from a given mask. Specifically, for a given functional map C and a mask matrix M , we can measure the total penalty as $\sum_{ij} [M]_{ij} [C]_{ij}^2$.

5.1. Effect of functional map size

For the functional maps pipeline, a set of corresponding descriptors is given as input. We then optimize a $k \times k$ functional map by minimizing an objective function based on Eq. (15). Therefore, we have to solve for k^2 variables. If k is a smaller value, e.g., $k < 50$, the op-

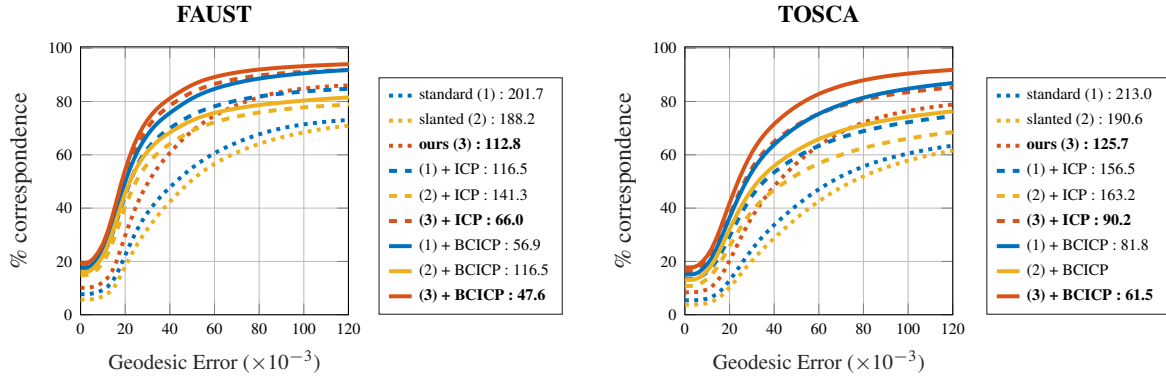


Figure 9: For each dataset, we compare the quality of the functional maps between the standard mask (blue curves), the slanted mask (yellow curves), and our resolvent mask (red curves) with the same set of parameters, where a single pair of descriptors is used to optimize a 100-by-100 functional map. The comparison is made in three different settings: comparing the initialization directly (dotted lines), comparing the initialization with ICP refinement (dashed lines), and with the BCICP refinement (solid lines). Specifically, these curves are measured on 300 FAUST shape pairs and 284 TOSCA pairs. The average direct errors are reported in the legends.

timization problem is easier to solve since the number of variables is small. However, if k is too small, the information that is encoded into the optimized functional map is limited to the low-resolution of the spectrum of the shapes, and hence it will be hard to transfer detailed information. On the other hand, if k is too large, solving the optimization problem is potentially more time-consuming. Even more importantly, this optimization problem can become under-constrained when the number of variables exceeds the constraints stemming from the input descriptors. In this case, we need effective regularizers to regularize the functional maps. In real experiments, the choice of the parameter k is a key hyper parameter.

To quantify the stability and the effectiveness of the proposed resolvent mask, we conduct the following test: for each of the test pairs, we only use one pair of corresponding WKS descriptors. We then fix this descriptor pair and optimize for functional maps with different sizes ranging from 20 to 250. We randomly select 50 pairs of FAUST and 50 pairs of TOSCA, and report the average error over the tested shape pairs w.r.t. different functional map size in Fig. 7. We can see that the standard mask fails to regularize the functional map with a large size: the average per-vertex error is three or four times larger than ours. At the same time, the slanted mask has a better performance than the standard one in the per-vertex measure. However, the slanted mask has large direct errors, which suggests that the smoothness is not well preserved. We believe this is due to the fact that the orientation-preserving regularizer starts to fail to disambiguate the symmetry as the functional map size k increases. In this case, the slanted mask cannot help the orientation-preserving operator, while our mask can strengthen the functionality of the orientation-preserving operator and leads to maps with much lower per-vertex and direct error. In summary, even with limited constraints from a single pair of WKS descriptors, increasing the number of variables does not significantly affect the performance of our mask. Fig. 8 shows an illustrative example.

This test shows that our mask is much more stable and can better regularize larger functional maps even in very challenging cases

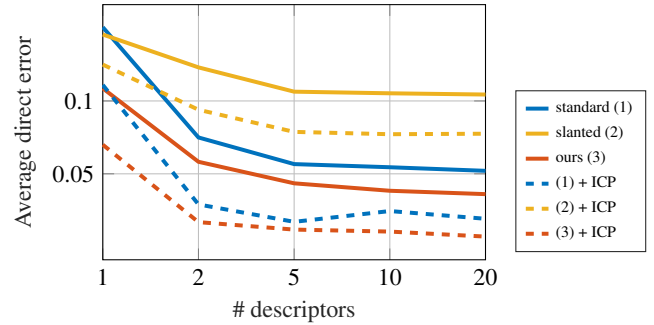


Figure 10: Changing the number of input descriptors. We use different numbers of descriptors to optimize a 100-by-100 functional map with different masks. The results are reported for 300 FAUST shape pairs. As the number of descriptors increases, the results improve for all the masks.

with little input information or constraints. Also, it suggests that with this new mask, we no longer need to tune the parameter k as much as needed by the standard mask to achieve a better result.

5.2. Evaluation on shape matching

In this experiment, we compare our resolvent mask to the standard and the slanted mask on a larger FAUST and TOSCA dataset in three settings: the initialization, with ICP refinement, and with BCICP refinement. To make a fair comparison between different masks, the weights for different terms (the α_i in Eq. (15)) are fixed across different test pairs and different test masks.

In this test, for each test pair, we use one pair of WKS descriptors to optimize a 100-by-100 functional map. As reported in Fig. 9, our mask leads to 40.4%, 43.1%, and 16.3% improvement w.r.t the settings of the initialization, with ICP refinement, and with BCICP refinement respectively, over the best of the standard and the slanted

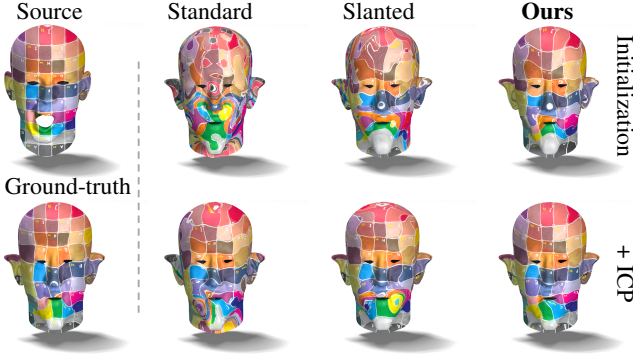


Figure 11: Example. Comparing the quality of the initial maps and after ICP refinement via texture transfer. (One pair of descriptors is used to optimize a 50-by-50 functional map)

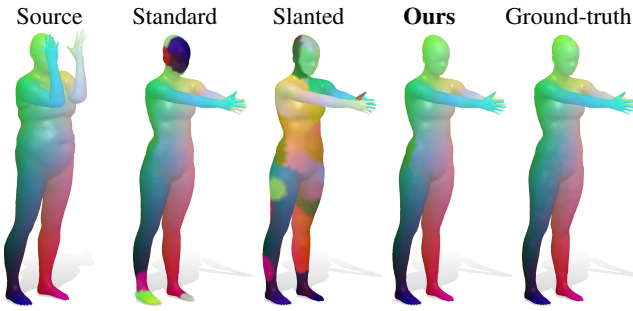


Figure 12: Example. Comparing the quality of the ICP refined maps via color transfer. (One pair of descriptors is used to optimize a 100-by-100 functional map)

mask on the FAUST dataset regarding the average direct error. Similarly, for the TOSCA dataset, ours achieves 34.2%, 42.3%, and 24.8% improvement respectively over the best of the standard and the slanted mask.

Fig. 10 shows a test on the FAUST dataset, where we use different number of carefully curated input descriptors based on WKS [ASC11] using parameters from [RPWO18] to compute 100-by-100 functional maps. As the number of input descriptors increases, more constraints are added to regularize the functional map. Therefore, for all the masks, and all the test settings, the results improve. Observe that when the number of variables is small (as shown in Fig. 7 for small k) or the number of descriptor constraints is large (as shown in Fig. 10 for large descriptor number), all masks perform well since the problem is well-constrained. However, when this is not the case, our resolvent mask can still regularize the functional map better than the other two. For completeness, in Appendix D we also include results with the BCICP refinement. We remark that this refinement is very computationally and memory intensive due, in part, to requiring all-pairs geodesic distances, but can, as such, improve upon even very poor quality maps.

Fig. 11 shows a qualitative example of using a single descriptor pair to optimize for a 100×100 functional map. The first row shows the quality of the initial maps with different masks, and the second



Figure 13: Example. Here we show a challenging pair of a lion and a cat and the results are refined by BCICP. We can see that, when the initial maps are in a low quality and the BCICP refinement fails to improve the initial maps, our new mask still gives a more reasonable map. (Twenty pairs of descriptors are used to optimize a 50-by-50 functional map)

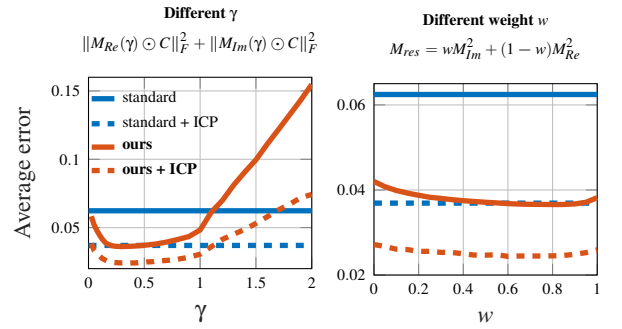


Figure 14: Left: changing the parameter γ (as defined in Eq. (6)); Right: changing the relative weight w between the imaginary part and the real part (see the original definition in Eq. (13)). The results are based on 50 TOSCA pairs.

rows the quality of the corresponding maps refined by ICP. We can see that our resolvent mask outperforms the other two masks. Also the quality of our initial map is close to the ICP refined map, which shows that with our mask, we do not rely on the post-processing refinement as much as the other two. Fig. 12 shows another example of the results refined by ICP. Fig. 13 shows a challenging non-isometric example with results refined after BCICP.

5.3. Analysis of the complex resolvent mask

As shown in Fig. 5, the funnel-shape of our resolvent mask aligns well with the ground-truth functional map, which leads to a better performance over the standard and the slanted mask. To further analyze the properties of our mask, we also conduct the following experiments: We test the range of the parameter γ , and the relative weight between the complex and the real part to construct the resolvent mask. Moreover, we investigate the correlation between the mask penalty added on a functional map and the corresponding recovered point-wise map.

5.3.1. Different parameters for the complex resolvent mask

In this section, we empirically tune the parameters in our resolvent based mask. First, we explore different values of γ . Recall that γ controls the funnel-like structure of the mask. Thus, it is expected that tuning γ can influence the functional map quality.

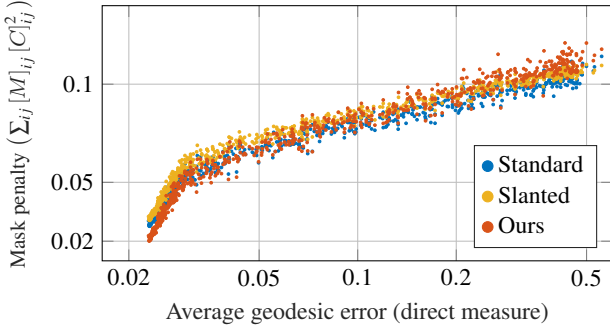


Figure 15: Correlation between the mask penalty on a functional map and the accuracy of the corresponding point-wise map. For a TOSCA isometric pair, we sample 700 functional maps of size 50-by-50 with different quality (normalized to the same scale). We measure the penalty of different masks, as an indicator of the quality of the functional map. We also measure the average geodesic error (w.r.t. the direct measurement) of the recovered point-wise map of the corresponding functional map. The correlation between the functional map penalty and the point-wise map quality is visualized as a scatter plot with 700 samples. Compared to the standard and the slanted mask, our resolvent mask applies a smaller penalty to functional maps with good underlying point-wise maps and more heavily penalizes functional maps with bad underlying point-wise maps.

On Fig. 14, we report results for 100 FAUST pairs, where the result of the standard mask is colored blue, and ours is colored red. As is shown on the left, when γ lies between 0 and 1, our mask always outperforms the standard mask over both the initialization (solid lines) and after the ICP refinement (the dashed lines).

Note that the case $\gamma = 1$ corresponds to the resolvent of the Laplacian. Thus, the fact that our mask outperforms the standard one for $\gamma = 1$ experimentally justifies the usage of the resolvent of the Laplacian, rather than the Laplacian itself. Note also that, as suggested in Sec. 4.4, the resolvent mask performs poorly for $\gamma > 1$.

Finally, we explore the relative contribution of the real and imaginary parts of the resolvent mask in order to analyze the utility of the two components of our mask construction. For this, we analyze the accuracy of the computed pointwise maps when allowing M_{Re} and M_{Im} to have different weights, $(1 - w)$ and w for $w \in [0, 1]$, respectively. As shown in Fig. 14 (right), the convex shape of the red curve (where the weight w changes) demonstrates that both the real part and the imaginary part contribute to the improvement over the standard mask. Note that our mask with any convex combination between the real part and the imaginary parts outperforms the standard mask. In practice we always use the equal weight $w = 0.5$.

5.3.2. Correlation between the mask penalty and map accuracy

To justify that our resolvent mask is a better regularizer than the standard and the slanted ones, we also measure how the mask penalty relates to the accuracy of the corresponding point-wise map shown in Fig. 15. In this experiment, we generate 700 different point-wise maps with different levels of accuracy, then convert

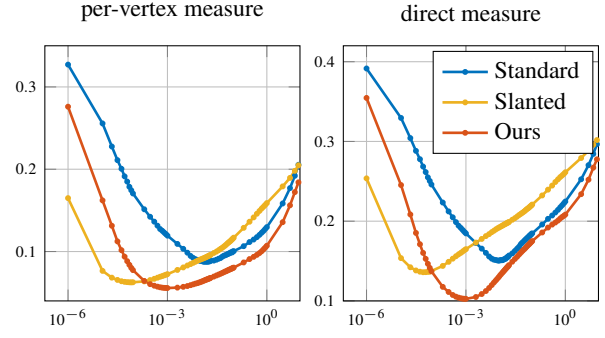


Figure 16: Changing the relative weight α_4 of the mask term. We choose 68 different values for α_4 in the range of 0 and 10 for this test.

Table 1: Average error on 284 TOSCA pairs of different masks, each mask using its own optimal weight from Fig. 16. Specifically, we set α_4 to 10^{-2} , 10^{-4} , 10^{-3} for the standard, slanted, and our mask respectively.

Mask type	Average error ($\times 10^{-3}$)					
	per-vertex measure			direct measure		
	Ini	+ ICP	+ BCICP	Ini	+ ICP	+ BCICP
standard	87.26	65.97	41.22	178.8	130.8	76.87
slanted	64.66	50.76	37.45	166.9	131.5	100.0
ours	55.52	43.82	32.48	124.5	90.6	62.33

them to a functional map representation and measure the penalty added by different masks w.r.t. the average geodesic error computed from the pointwise maps. Each scatter point in Fig. 15 shows such a test sample.

We can observe that, compared to the standard and the slanted mask, the new mask induces a lower penalty on functional maps with a good quality (i.e., smaller average geodesic error), and penalizes a functional map with larger error more heavily. This further confirms that using our resolvent mask is more likely to produce a better functional, and ultimately better pointwise map.

5.4. Parameters

In our tests, we use $\gamma = 0.5$ and $w = 0.5$ to construct our resolvent mask, and use the default value $\eta = 0.03$ to construct the slanted mask. When comparing the three masks in the initialization setting, i.e., to optimize the energy defined in (15), the weights α_i are set to the same values as reported in [RPWO18] but in a relative way (see Appendix D for more details). As for the comparison of the ICP and BCICP refinement settings, the same default parameters and the same number of iterations are used for different masks.

Fig. 16 shows the results of the changing the weight of the mask term, i.e., α_4 , while keeping the rest weights fixed on 50 TOSCA shape pairs. We can see that, for all different choices of the weight, our resolvent mask is always better than the standard mask. When the weight is small enough, it seems the slanted mask can achieve a faster decrease of the error than ours. However, the range of the

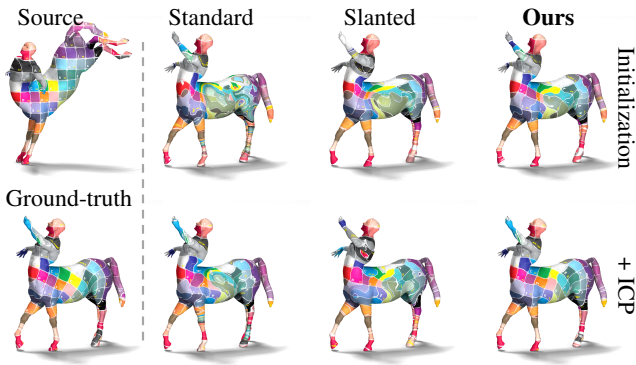


Figure 17: Example. Comparing the quality of the initial maps computed with different masks and after ICP refinement via texture transfer. Each mask uses its own optimal weight from Fig. 16.

effective weight is smaller than ours besides that the construction is purely heuristic and lacks a theoretical justification.

On the other hand, Fig. 16 also suggests that different masks have their own preferred choice of the weight α_4 . Therefore, instead of fixing α_4 , we set it independently w.r.t. the optimal values on a subset reflected in Fig. 16. Specifically, for the standard mask, we set $\alpha_4 = 10^{-2}$, for the slanted mask, we set $\alpha_4 = 10^{-4}$, and for ours, we set $\alpha_4 = 10^{-3}$. The corresponding average errors on the complete TOSCA dataset are reported in Table 1. We can see that, even when the parameters are carefully tuned for the other two masks, our resolvent mask still achieves the best accuracy. Fig. 17 shows a qualitative example.

5.5. Application to non-isometric shape pairs

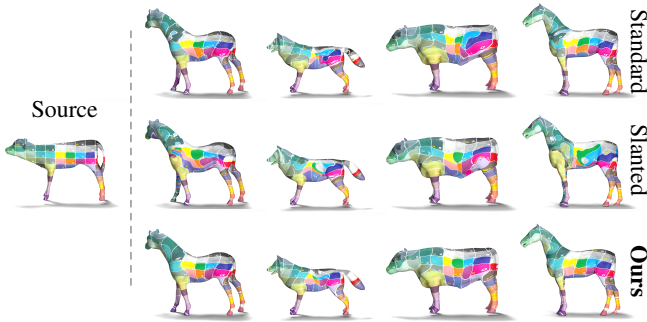


Figure 18: Example. Comparing the quality of the initial maps computed with different masks of four SHREC shape pairs via texture transfer.

To show the usefulness of our resolvent mask on non-isometric shape pairs, we test the 20 FourLeg shapes from the SHREC 2007 dataset [GBP]. Specifically, we use 10 pairs of descriptors constructed from 4 landmarks to optimize a 120-by-120 functional map using the standard, the slanted, and our resolvent mask. For a fair comparison, we use the same parameters as the previous tests: we set $\gamma = 0.5$ for our resolvent mask, and set $\eta = 0.03$ for the slanted mask. Fig. 18 shows a qualitative example. We can see that

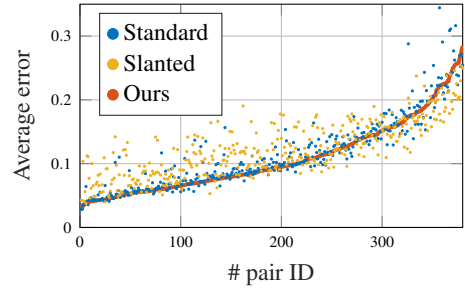


Figure 19: For 380 SHREC shape pairs, we plot the average error for each pair with different masks, where the result of the standard, slanted and our mask is colored in blue, yellow, and red respectively. Therefore, in the case where the red curve is below the blue or the yellow points, our mask leads to a lower error for these pairs.




Figure 20: Failure case. Here we show a challenging shape pair from SHREC, where all the three masks fail to produce a good map.

our resolvent mask gives the best initialization. We also measure the accuracy of the 380 point-wise maps among the 20 shapes on the given landmarks (21 ground-truth landmarks are given for each shape). The average error for the standard, slanted, and our mask is 0.114, 0.122, and 0.107 respectively. Fig. 19 reports the average error for each shape pair. We observe that resolvent mask gives a better result than the standard mask on 68.4% out of 380 pairs, and outperforms the slanted mask on 69.5% pairs. The limited quantitative improvement of our mask over the other two masks is due to: (1) for the shape pairs where our mask significantly outperforms the other two, e.g., as shown in Fig. 18, the map quality is measured on only 21 landmarks, where the average error does not fully reflect our improvement. (2) for some extremely challenging pairs, e.g., a failure case shown in Fig. 20, all the maps from different masks have a poor quality, and thus the “relative improvement” on the average error is not informative. Moreover, we also computed the average error in the case where the mask term is removed from the total energy, i.e., set $\alpha_4 = 0$. In this case, the average error over the complete dataset is 0.112. We can see that, the standard and the slanted mask can have a negative effective in this case, while our resolvent mask still works to improve the map quality to some extent.

6. Conclusion, Limitations and Future Work

In this paper, we proposed a new regularizer, the resolvent Laplacian commutativity, for the functional map pipeline. We first analyzed the limitations of the original Laplacian commutativity term and theoretically justified the effectiveness of our proposed new term. This new regularizer can significantly improve the quality

of the computed functional maps and the corresponding recovered point-wise maps before or after refinement.

However, our method also has several limitations that we would like to overcome in future work. First, our proposed regularizer is well justified on isometric and non-isometric shape pairs, but not on partial shapes, where the ground-truth functional maps can have a different structure. Therefore, it would be interesting to extend the analysis to partial shape pairs. Second, besides the funnel pattern, in our experiments, we also observed the slanted-diagonal structure of the ground-truth functional map of some non-isometric shape pairs as discussed in [RCB*17]. It would be interesting to consider this feature into the mask constructions for further improvement for non-isometric datasets. Thirdly, the role of the (a, b) introduced in Eq. (5) is not well studied, and we would like to leave this exploration as future work. Finally, we believe that it would be interesting to study and potentially apply data-driven techniques to learn the optimal operators for enforcing commutativity across diverse shape pairs.

Acknowledgement The authors would like to thank the anonymous reviewers for their valuable comments and helpful suggestions. Parts of this work were supported by the KAUST OSR Awards No. CRG2017-3426 and CRG2018-3730, a gift from the NVIDIA Corporation, and the ERC Starting Grant No. 758800 (EXPROTEA).

References

- [ADK16] AFLALO Y., DUBROVINA A., KIMMEL R.: Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision* 118, 3 (2016), 380–392. 2
- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)* (2011), IEEE, pp. 1626–1633. 7, 9
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer Science & Business Media, 2008. 6
- [BCBB16] BIASOTTI S., CERRI A., BRONSTEIN A., BRONSTEIN M.: Recent trends, applications, and perspectives in 3d shape similarity assessment. In *Computer Graphics Forum* (2016), vol. 35, pp. 87–119. 2
- [BDK17] BURGHARD O., DIECKMANN A., KLEIN R.: Embedding shapes with Green’s functions for global shape matching. *Computers & Graphics* 68 (2017), 1–10. 2
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (Piscataway, NJ, USA, June 2014), IEEE. 1, 6
- [CSBK18] CHOUKROUN Y., SHTERN A., BRONSTEIN A. M., KIMMEL R.: Hamiltonian operator for spectral shape analysis. *IEEE transactions on visualization and computer graphics* (2018). 3
- [Cut13] CUTURI M.: Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems* (2013), pp. 2292–2300. 2
- [Dod81] DODZIUK J.: Eigenvalues of the Laplacian and the heat equation. *The American Mathematical Monthly* 88, 9 (1981), 686–695. 4
- [EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 165–174. 2, 3
- [ERGB16] EYNARD D., RODOLA E., GLASHOFF K., BRONSTEIN M. M.: Coupled functional maps. In *3D Vision (3DV), 2016 Fourth International Conference on* (2016), IEEE, pp. 399–407. 2
- [GBKS18] GEHRE A., BRONSTEIN M., KOBBELT L., SOLOMON J.: Interactive curve constrained functional maps. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 1–12. 2
- [GBP] GIORGI D., BIASOTTI S., PARABOSCHI L.: Shape retrieval contest 2007: Watertight models track. 11
- [HO17] HUANG R., OVSJANIKOV M.: Adjoint map representation for shape analysis and matching. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 151–163. 2
- [HRS*16] HEEREN B., RUMPF M., SCHRÖDER P., WARDETZKY M., WIRTH B.: Splines in the space of shells. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 111–120. 1
- [HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 36. 1, 2
- [KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. In *Computer Graphics Forum* (2013), vol. 32, pp. 439–448. 1, 2
- [KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015), pp. 905–914. 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. 3
- [LRB*16] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D.: Non-rigid puzzles. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 135–143. 2
- [LRBB17] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Fully spectral partial shape matching. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 247–258. 2
- [MCSK*17] MANDAD M., COHEN-STEINER D., KOBBELT L., ALLIEZ P., DESBRUN M.: Variance-minimizing transport plans for inter-surface mapping. *ACM Transactions on Graphics (TOG)* 36 (2017), 14. 2
- [MP49] MINAKSHISUNDARAM S., PLEIJEL A.: Some properties of the eigenfunctions of the Laplace-operator on Riemannian manifolds. *Canadian J. Math* 1, 8 (1949). 3, 14
- [MRCB18] MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M. M.: Localized manifold harmonics for spectral shape analysis. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 20–34. 3
- [NMR*18] NOGNENG D., MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M., OVSJANIKOV M.: Improved functional mappings via product preservation. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 179–190. 2
- [NO17] NOGNENG D., OVSJANIKOV M.: Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum* 36, 2 (2017), 259–267. 2, 3, 6
- [NVT*14] NEUMANN T., VARANASI K., THEOBALT C., MAGNOR M., WACKER M.: Compressed manifold modes for mesh processing. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 35–44. 3
- [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, 2, 4, 7
- [OCB*17] 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 *ACM SIGGRAPH 2017 Courses* (2017), SIGGRAPH ’17, pp. 5:1–5:62. 2, 3

- [PSO18] POULENARD A., SKRABA P., OVSJANIKOV M.: Topological function optimization for continuous shape matching. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 13–25. 2
- [RCB*17] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSSELLO A., CREMERS D.: Partial functional correspondence. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 222–236. 1, 2, 4, 7, 12
- [RMC15] RODOLÀ E., MOELLER M., CREMERS D.: Point-wise map recovery and refinement from functional correspondence. In *Proc. Vision, Modeling and Visualization (VMV)* (2015). 3
- [ROA*13] RUSTAMOV R., OVSJANIKOV M., AZENCOT O., BEN-CHEN M., CHAZAL F., GUIBAS L.: Map-based exploration of intrinsic shape differences and variability. *ACM Transactions on Graphics (TOG)* 32, 4 (July 2013), 72:1–72:12. 1, 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, 6, 7, 9, 10, 14, 15
- [RS80] REED M., SIMON B.: *Methods of Modern Mathematical Physics I: Functional Analysis*, revised and enlarged ed. Academic Press, San Diego, 1980. 4, 13
- [Sau06a] SAUVIGNY F.: *Partial Differential Equations 1: Foundations and Integral Representations*. Springer, Berlin, Heidelberg, 2006. 4
- [Sau06b] SAUVIGNY F.: *Partial Differential Equations 2: Functional Analytic Methods*. Springer, Berlin, Heidelberg, 2006. 4
- [SDGP*15] SOLOMON J., DE GOES F., PEYRÉ G., CUTURI M., BUTSCHER A., NGUYEN A., DU T., GUIBAS L.: Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 66. 2
- [SPKS16] SOLOMON J., PEYRÉ G., KIM V. G., SRA S.: Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 72. 2
- [TCL*13] TAM G. K., CHENG Z.-Q., LAI Y.-K., LANGBEIN F. C., LIU Y., MARSHALL D., MARTIN R. R., SUN X.-F., ROSIN P. L.: Registration of 3D point clouds and meshes: a survey from rigid to nonrigid. *IEEE TVCG* 19, 7 (2013), 1199–1217. 2
- [VKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. In *Computer Graphics Forum* (2011), vol. 30, pp. 1681–1707. 2
- [VLR*17] VESTNER M., LITMAN R., RODOLÀ E., BRONSTEIN A., CREMERS D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Proc. CVPR* (2017), pp. 6681–6690. 2
- [WGBS18] WANG L., GEHRE A., BRONSTEIN M. M., SOLOMON J.: Kernel functional maps. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 27–36. 2
- [WHG13] WANG F., HUANG Q., GUIBAS L. J.: Image co-segmentation via consistent functional maps. In *Proceedings of the IEEE International Conference on Computer Vision* (2013), pp. 849–856. 2
- [WLZT18] WANG Y., LIU B., ZHOU K., TONG Y.: Vector field map representation for near conformal surface correspondence. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 72–83. 2

Appendix A: Pullbacks are Bounded

Lemma 1 Let $\Phi : \mathcal{M} \rightarrow \mathcal{N}$ be a diffeomorphism between connected compact oriented Riemannian manifolds. Then, the associated pullback Φ^* is bounded as a map $L_2(\mathcal{N}) \rightarrow L_2(\mathcal{M})$.

Proof Let $d\mathcal{M}$ and $d\mathcal{N}$ denote the volume forms of \mathcal{M} and \mathcal{N} ,

respectively. Since C^∞ is dense in L_2 , it is enough to show that there exists a constant $B > 0$ such that, for any $f \in C^\infty(\mathcal{N})$,

$$\int_{\mathcal{M}} (\Phi^* f)^2 d\mathcal{M} \leq B \int_{\mathcal{N}} f^2 d\mathcal{N}. \quad (16)$$

By virtue of being a diffeomorphism, Φ is invertible. Using the pullback of Φ^{-1} , we can express the left-hand side of the desired inequality as an integral over \mathcal{N} .

Begin by considering $\left[(\Phi^{-1})^* d\mathcal{M} \right]$, the pullback of the volume form of \mathcal{M} . Since volume forms are top degree forms, there exists $u \in C^\infty(\mathcal{N})$ such that

$$\left[(\Phi^{-1})^* d\mathcal{M} \right] = u d\mathcal{N} \quad (17)$$

Φ is either orientation preserving or orientation reversing. In the former case, $u > 0$. In the latter, $u < 0$. In either case, the left-hand side of the desired inequality can be recast as

$$\begin{aligned} \int_{\mathcal{M}} (\Phi^* f)^2 d\mathcal{M} &= \int_{\mathcal{N}} \left((\Phi^{-1})^* \Phi^* f \right)^2 |u| d\mathcal{N} \\ &= \int_{\mathcal{N}} f^2 |u| d\mathcal{N}. \end{aligned} \quad (18)$$

It remains to bound this expression.

$$\int_{\mathcal{N}} f^2 |u| d\mathcal{N} \leq \sup_{x \in \mathcal{N}} |u(x)| \int_{\mathcal{N}} f^2 d\mathcal{N} \quad (19)$$

Since \mathcal{N} is compact and u is continuous, the supremum is achieved and is finite. This concludes the proof. \square

Appendix B: Bounded Frobenius Norm for $\gamma > 1/2$

In this appendix, we prove a sufficient condition for the Frobenius norm based energy to be defined in the continuous case. We begin by introducing an infinite dimensional analog of the Frobenius norm, which is provided by the Hilbert-Schmidt norm.

Definition 2 (Hilbert-Schmidt Norm) Let $A : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ be a linear operator between Hilbert spaces. Let A^\dagger denote the adjoint of the operator A . Then, the Hilbert-Schmidt norm of A is given by:

$$\|A\|_{HS}^2 = \text{Tr} \left(A^\dagger A \right) = \sum_{i=1}^{\infty} \langle e_i, A^\dagger A e_i \rangle, \quad (20)$$

where $\{e_i\}_{i=1}^{\infty}$ is any orthonormal basis of \mathcal{H}_1 . This norm is also known as the Schatten 2-norm.

The following lemma is the main result of this appendix.

Lemma 2 Let Δ_1 and Δ_2 be Laplacians on compact, connected, oriented surfaces \mathcal{M}_1 and \mathcal{M}_2 , respectively. Let $C_{12} : L_2(\mathcal{M}_1) \rightarrow L_2(\mathcal{M}_2)$ be a bounded operator. If $\gamma > 1/2$, then:

$$\|C_{12} R_\mu(\Delta_1^\gamma) - R_\mu(\Delta_2^\gamma) C_{12}\|_{HS}^2 < \infty, \quad (21)$$

where μ is any complex number not on the non-negative real line.

Proof Operators with finite Hilbert-Schmidt norm are known as operators of Hilbert-Schmidt class. It can be shown (see [RS80]) that linear combinations of Hilbert-Schmidt class operators are of

Hilbert-Schmidt class. Moreover, the product of a bounded operator and a Hilbert-Schmidt class operator is also of Hilbert-Schmidt class. Thus, it is sufficient to show that $R_\mu(\Delta^\gamma)$ has finite Hilbert-Schmidt norm for $\gamma > 1/2$.

Denote the eigenfunctions and eigenvalues of Δ by $\{\psi_k\}_{k=0}^\infty$ and $\{\lambda_k\}_{k=0}^\infty$, respectively. We assume that the eigenvalues are numbered in the usual non-decreasing order.

$R_\mu(\Delta^\gamma)$ is diagonal in the eigenbasis of Δ and has eigenvalues $1/(\lambda_k^\gamma - \mu)$. Thus, the Hilbert-Schmidt norm of $R_\mu(\Delta^\gamma)$ is given by:

$$\begin{aligned} \|R_\mu(\Delta^\gamma)\|_{HS}^2 &= \sum_{k=0}^{\infty} \langle \psi_k, R_\mu(\Delta^\gamma)^\dagger R_\mu(\Delta^\gamma) \psi_k \rangle \\ &= \sum_{k=0}^{\infty} \frac{1}{|\lambda_k^\gamma - \mu|^2} \end{aligned} \quad (22)$$

We will establish the convergence of this series for $\gamma > 1/2$ using the comparison test with the series $\sum_{k=1}^{\infty} \frac{1}{k^p}$. This series is well-known to converge if and only if $p > 1$.

Since the λ_k are non-negative and increasing with k towards ∞ , the following inequality holds for all large enough k :

$$\frac{1}{4\lambda_k^{2\gamma}} \leq \frac{1}{|\lambda_k^\gamma - \mu|^2} \leq \frac{2}{\lambda_k^{2\gamma}}. \quad (23)$$

By Weyl's law, there exists a constant $B > 0$ such that $\lambda_k \sim Bk$ for large k [MP49]. Then, the inequality becomes

$$\frac{1}{4B^{2\gamma} k^{2\gamma}} \leq \frac{1}{|\lambda_k^\gamma - \mu|^2} \leq \frac{2}{B^{2\gamma} k^{2\gamma}}. \quad (24)$$

Thus, by the comparison test, the series for $\|R_\mu(\Delta^\gamma)\|_{HS}^2$ converges if and only if $\gamma > 1/2$. This concludes the proof. \square

We conclude this appendix with two remarks on the above result. First, note that we have shown a result slightly stronger than required by the statement of the lemma. In fact, we only needed to show that the series for $\|R_\mu(\Delta^\gamma)\|_{HS}^2$ converges if $\gamma > 1/2$. The "only if" part was optional. We have done this to illustrate that the above proof *strategy* is guaranteed to fail for $\gamma \leq 1/2$. Specifically, it is no longer sufficient to assume that C_{12} is merely a bounded operator. One also cannot simply assume C_{12} to be of Hilbert-Schmidt class, as this rules out the important case of $C_{12} = Id$, which is bounded, but not Hilbert-Schmidt.

As a final remark, note that an analogous proof strategy can be applied to the Schatten p -norm, which can be seen as the l_p generalization of the Hilbert-Schmidt norm. There, the k^{th} term of the series would be $1/|\lambda_k^\gamma - \mu|^p$ and convergence would be guaranteed for $\gamma > 1/p$. Thus, one can find p large enough so that $R_\mu(\Delta^\gamma)$ has a well-defined Schatten p -norm for any given $\gamma > 0$.

Appendix C: Comparison to heat mask

Besides our resolvent-based commutativity term, another natural bounded option would be to use the commutativity with the *heat* operators, which are also bounded linear functional operators, which can act as regularizers on functional maps. This would lead

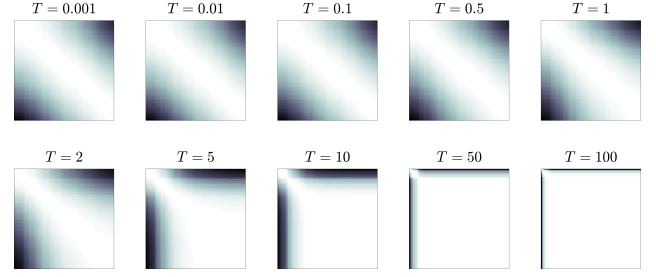


Figure 21: Visualization of the heat mask with different time-scale parameter T .

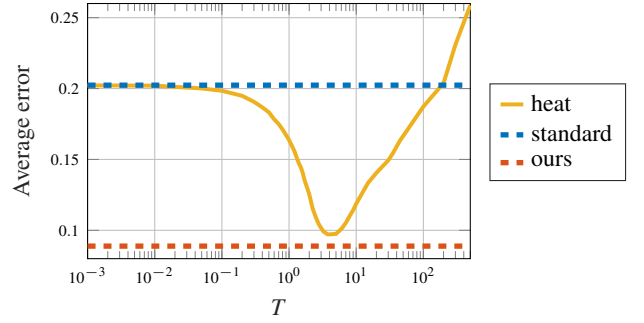


Figure 22: Changing the time-scale parameter T of the heat mask. Here we test the performance of the heat mask (yellow curve) with different T and report the average direct error of 50 random FAUST isometric shape pairs. The average error of the standard mask (blue dashed line) and the average error of our resolvent mask (red dashed line) are included for comparison.

to the following mask matrix:

$$M_{heat}(i, j) = \exp(-T\Lambda_2(i)) - \exp(-T\Lambda_1(j)), \quad (25)$$

where T is the scalar time parameter. Fig. 21 visualizes the heat mask with different values of T on an isometric shape pair. In Fig. 22 we compared the performance of the heat mask with different values of T ranging from 10^{-3} to 500 on FAUST isometric shape pairs. We can see that, as a bounded operator, the heat mask has a better performance than the standard mask. When $T = 5$, it gives the best performance among the tested values. The heat mask with $T = 5$ (in Fig. 21) has a similar funnel structure, but still does not achieve the quality of the results we obtain with the mask based on the resolvent operators.

Appendix D: Change the descriptor size

Fig. 23 shows the average direct error of 300 FAUST shapes with different number of input descriptors. The descriptors are used to optimize a 100-by-100 functional map for each test pair. The results with ICP/BCICP refinement are also included.

We also compare to the results obtained directly using the code and the dataset provided by the authors of [RPWO18] (see Table 2). The reproduced results that are reported in Table 2 are consistent with the values reported in the Table 1-2 ("WKS + directOp + BCICP") in the paper, and Table 2-3 ("WKS + directOp") in the

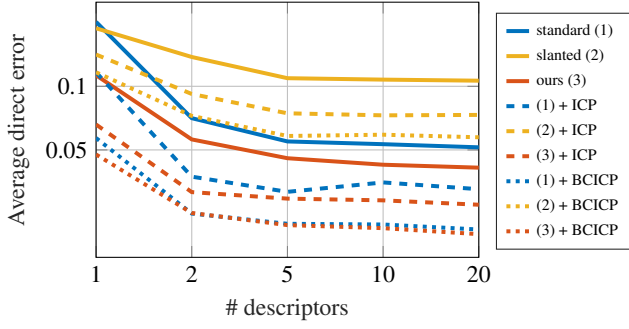


Figure 23: The average geodesic error v.s. the number of descriptors used for functional map estimation on the FAUST dataset. Note that our mask leads to better initialization, which results in better maps, even after the ICP and BCICP refinement.

Table 2: We also compare our results to the exact setting of [RPWO18] on the same set of shape pairs of FAUST dataset. The average direct geodesic error over 300 shape pairs are reported and our mask leads to better results.

	Avg. direct error ($\times 10^{-3}$)		
	Ini	+ICP	+BCICP
[RPWO18]	175.5	121.2	58.2
Ours	72.7	53.5	42.0

supplementary materials of [RPWO18]. Note that [RPWO18] split the dataset into isometric and non-isometric categories, and here we report the results altogether.

Recall that the total energy to optimize is $E(C_{12}) = \alpha_1 E_{\text{desc}} + \alpha_2 E_{\text{mult}} + \alpha_3 E_{\text{orient}} + \alpha_4 E_{\text{mask}}$. We would like to highlight the fact that the parameters in [RPWO18] are not set in the same way as in our work. In particular, in [RPWO18] the standard Laplacian mask is used, and the weight α_i are set to fixed values α_i^* . While in our test, we used the proposed resolvent Laplacian mask, and the weight α_i are set to $\alpha_i^* / E_i(C_{\text{ini}})$, where $E_i(C_{\text{ini}})$ is the corresponding energy term acting on the initial functional map C_{ini} .

Our approach allows a better control over the relative contribution of the different terms in the energy and we observed that it typically works better in practice as well. Specifically, in our setting, the relative weight of each term $\alpha_i E_i(C)$ are fixed across different shape pairs. In this case, if we change the mask construction, we can conclude that the improvement indeed come from our proposed resolvent mask. However, in the comparison to the exact setting of [RPWO18], since only α_i^* is fixed over different test pairs, the relative weight of each term $\alpha_i^* E_i(C)$ can have different scale since different test pairs may have different scale of eigenvalues, descriptors and etc. Therefore, the improvement from Table 2 is not obtained in a well controlled setting, since the improvement can also come from the change of the relative weight of different energy terms as well as from our resolvent mask.

Thus, in addition to the new Table 2, we also emphasize that in Fig. 23 we provide a more fair and controlled comparison to [RPWO18] in which 10 pairs of descriptors are used. Note that, in all the tests across the paper, we used the above discussed way to

fix the relative weight of the mask term w.r.t. the rest terms to make sure the improvement indeed comes from our new resolvent mask.

Appendix E: Stability under remeshing and refinement

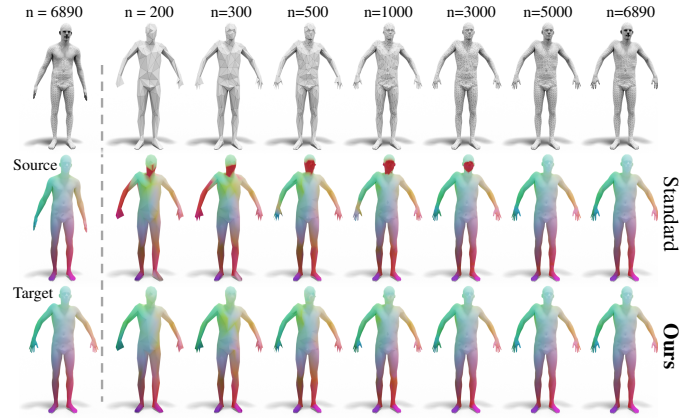


Figure 24: First row: The target shape is fixed, while the source shape is remeshed and downsampled to resolution ranging from 200 to 5000 (the original target shape without remeshing is shown in the last column); Second row: we use the standard mask to optimize a 100-by-100 functional map with 3 descriptors. The recovered pointwise maps are visualized on the corresponding target shape; Third row: similar to the second row but using our resolvent mask.

Fig. 24 illustrate the stability of our resolvent mask under remeshing and refinement. Specifically, we fix the source shape, and remesh and downsample the target shape to different resolutions ranging from 200 to 5000. Note that the original source shape (first row, first column) and the original target shape (first row, last column) have the same triangulation. The downsampled meshes (using QSlm) are shown in the first row with the number of vertices reported in the above.

We then compute a 100-by-100 functional map between the source shape and the downsampled target shape using the standard Laplacian mask (the corresponding point maps are shown in the second row) and our resolvent mask (in the third row). We can see that, our resolvent mask is more stable than the standard mask across different mesh resolution and irregular/inconsistent triangulation.