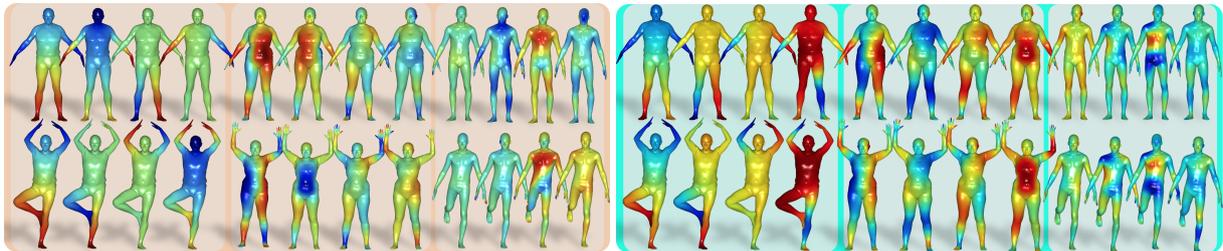


# Shape correspondence by aligning scale-invariant LBO eigenfunctions

A. Bracha, O. Halim, R. Kimmel

Faculty of computer science, the Technion, Israel



**Figure 1:** Four corresponding eigenfunctions textured mapped to the same individual at two different poses (top and bottom), before (pink background) and after (green background) alignment. The first four on the left are eigenfunctions 2-5, followed by 12-15, and 22-25 on the right.

## Abstract

When matching non-rigid shapes, the regular or scale-invariant Laplace-Beltrami Operator (LBO) eigenfunctions could potentially serve as intrinsic descriptors which are invariant to isometric transformations. However, the computed eigenfunctions of two quasi-isometric surfaces could be substantially different. Such discrepancies include sign ambiguities and possible rotations and reflections within subspaces spanned by eigenfunctions that correspond to similar eigenvalues. Thus, without aligning the corresponding eigenspaces it is difficult to use the eigenfunctions as descriptors. Here, we propose to model the relative transformation between the eigenspaces of two quasi-isometric shapes using a band orthogonal matrix, as well as present a framework that aims to estimate this matrix. Estimating this transformation allows us to align the eigenfunctions of one shape with those of the other, that could then be used as intrinsic, consistent, and robust descriptors. To estimate the transformation we use an unsupervised spectral-net framework that uses descriptors given by the eigenfunctions of the scale-invariant version of the LBO. Then, using a spectral training mechanism, we find a band limited orthogonal matrix that aligns the two sets of eigenfunctions.

## 1. Introduction

One of the challenging problems in geometry processing is finding shape correspondence. First attempts to solve the matching problem used point descriptors which can be divided into *extrinsic* and *intrinsic* ones. Extrinsic descriptors are based on geometric properties that describe the embedding of the surface in  $\mathbb{R}^3$ , like the surface normal at a point, the mean curvature, the area encapsulated by spheres of varying radii about a point, etc. These measures are sensitive to poses and postures of articulated objects [STDS14, WBCPS18, ZWL19]. At the other end, intrinsic measures, expressed solely in terms of the metric tensor, are known to be invariant to those natural deformations, that are modeled as isometry to a very good approximation. That observation lead to a

school of thoughts in the field of shape analysis that promotes treating shapes as metric spaces [HRAK19, MS05, BBK06b, BBK06a].

LBO eigenfunctions exhibit an exclusive advantage as intrinsic descriptors, being an orthonormal complete set. As such, they fully determine the spectral transformation, unlike other descriptor functions whose support vanishes at certain spectral domains such as the *heat kernel signatures* [SOG09] and the *wave kernel signatures* [ASC11], that are supported only on the symmetric modes. In [Rus07] it was first argued that the LBO eigenfunctions could serve as descriptors for shape correspondence, yet, the misalignment issue was not addressed directly. In practice, the LBO eigenfunctions are so poorly aligned between two quasi-isometric shapes that they are rarely used as descriptors. The presented framework

targets directly this misalignment concern. It is the first method that solves the eigenfunctions alignment issue so effectively so that a correspondence pipeline built exclusively upon LBO eigenfunctions as dense descriptors achieves on par performance with state-of-the-art correspondence methods. Unlike other alignment methods [SK14] that can resolve only low order eigenfunctions ( $k < 10$ ), the proposed framework is scalable in the number of eigenfunctions and shows a significant improvement of the alignment over a broad range of frequencies, up to  $k = 350$ .

Recent efforts suggest feeding extrinsic descriptors to neural networks in order to learn better features [LRR\*17,HLR\*19,RSO19], provided state-of-the-art results in this field. Still, such networks are tuned to the given data and thus operate well within the domain defined by the training data. The dependency on data can be restrictive, for example, in [LRR\*17] a network was trained to extract descriptors of human bodies as input, while in [HLR\*19] a drop in performance was demonstrated when testing the resulting descriptors on super-hero imaginary figures. Even more importantly, all the mentioned methods are confined to produce a non-linear function of a specific extrinsic descriptor known as SHOT [STDS14], compromising to sub-optimal descriptors. Differently, the proposed method optimizes for the most general intrinsic descriptor. In fact, any intrinsic descriptor function can be represented by the proposed method as a linear combination of the aligned basis functions. Finally, in the era of deep networks, a common approach is to complicate the representation of the solution in-order to achieve better performance over classical models. While this approach has advantages in numerous applications, it sacrifices our ability to explain the solution. Thus, generally, when we have two solvers of the same problem, we prefer the explainable one. Thereby, unlike existing methods that utilize neural networks to construct complex descriptors, we use fundamental intrinsic descriptors and show that, when properly aligned, they are capable of producing the same performance.

To this end, we propose a novel method for shape correspondence based on intrinsic descriptors defined by *scale-invariant LBO* eigenfunctions. Specifically, given two shapes, we first extract their LBO eigenfunctions and the scale invariant LBO (SI-LBO) ones. Next, we jointly align the SI-LBO eigenfunctions in their spectral domain (Fig. 1), modeling the alignment as an operation of a band orthogonal matrix. We then use the aligned eigenfunctions as input descriptors for the functional map framework, where the LBO eigenfunctions are used as a basis of choice [ABK15]. The functional map defines the spectral transformation between the two given manifolds that can be transformed into a point-wise soft-correspondence, see [LRR\*17]. To optimize for the alignment matrix we propose a novel loss defined in terms of the alignment matrix and the point-wise correspondence induced from this matrix, using the described computation pipeline. This loss measures extent to which the aligned SI-LBO eigenfunctions on both shapes are compatible with each other, when mapped by the correspondence that this alignment induces. The resulting optimization process is straightforward, and achieved by back-propagation of the loss gradient to the entries of the alignment matrix. We end up with a well aligned SI-LBO eigenfunctions as well as the point-wise correspondence between the shapes.

## 2. Related Efforts

Non-rigid shape matching techniques can be divided into several classes. A relevant one to our discussion contains those that use intrinsic descriptors for point correspondence and shape matching. In [Rus07], the use of the LBO eigenfunctions weighted by the corresponding eigenvalues was argued to construct isometric invariant *global point signatures* that were then used to classify surfaces via moments of histograms computed via Green functions, also known as *kernels*. Next, in [OMMG10,ASC11] the *heat kernel signature*, and the *wave kernel signature*, were proposed as point descriptors. All these *kernel signatures* were computed using the LBO eigenfunctions, where sign ambiguities were avoided by considering the squared eigenfunctions. Similar kernel signatures can be constructed by using the scale invariant LBO [WWA\*16]. The discrepancies between kernel signatures can be proven to be theoretically bounded by the deformations relating the shapes from which they were extracted from. However, in practice, these bounds are not tight enough leading to numerical errors that make these descriptors inefficient for the task of dense shape correspondence. Also, most kernel signatures are symmetric functions and therefore cannot be used to distinguish symmetric points on the surface.

Another group of methods try to compute shape correspondence by minimizing geodesic distance distortions of embedding one shape into the other, which is intimately related to the Gromov-Hausdorff distance [MS05,BBK08,ADK16,SK17]. These methods often search for a permutation matrix between two sets of sample points that approximate the given surfaces. Bronstein *et al.* [BBK06b] proposed the *generalized multidimensional scaling* (GMDS) to search for the points permutation by gradient descent that minimize a loss based on the Gromov-Hausdorff distance. In [ADK16], the suggested *spectral-GMDS* method exploits the spectral domain to efficiently approximate the permutation matrix, yielding a formulation that can be shown to be related to the *Functional Map* framework [OBCS\*12]. The GMDS evolved into relaxed quadratic assignment problems, see for example [BK17,XLZ20]. In order to find correspondence between two non-isometric surfaces, where the geodesic distance is not preserved, Solomon *et al.* [SPKS16] used the related Gromov-Wasserstein distance [Mém11] with a regularization on the entropy of the correspondence.

Another way to construct robust intrinsic descriptors is to use eigenfunctions of other isometry invariant operators like the eigenfunctions of the inter-geodesics distance matrix [SK17]. The scale invariant metric can also be used for the definition of intrinsic distances, such as the scale invariant diffusion distance [WWJ\*14]. For more variations on intrinsic and extrinsic descriptors we refer to [WS19].

Another family of methods try to align the LBO eigenfunctions for shape matching or shape correspondence. Shtern *et al.* [SK14] considered the eigenfunctions as samples of random variables. They tried to resolve the sign and permutation ambiguities of the eigenfunctions by matching their third order moments. In [EKB\*15], Eynard *et al.* extended the coupled diagonalization idea, put forward in [KBB\*13], and proposed a method to find a mutual eigenfunctions basis of LBOs belonging to different shapes by minimizing a loss that promotes joint-diagonalization. To obtain

the mutual basis, their optimization process requires the existence of corresponding functions on both surfaces, leading to similar difficulties as those encountered in the functional map framework.

The next group of methods include learning descriptors for shape correspondence. Two related frameworks are the FM-Net [LRR\*17] and the unsupervised FM-Net [HLR\*19]. The first is a supervised residual neural network that learns a non-linear function operating on local SHOT descriptors [STDS14] as an input to the functional map framework that computes the desired correspondence between the shapes. Next, Halimi *et al.* converted the FM-Net into an unsupervised learning framework by replacing the supervised loss with the GMDS one. Not long after, [RSO19] suggested to define the unsupervised loss in the spectral domain. See also [GR19] for a recent unsupervised method based on a *Cyclic-loss*, and [YLB\*20] for a different variation on the features processed by the FM-Net. Here, when optimizing for the alignment of the scale invariant eigenfunctions, we use the unsupervised loss [HLR\*19] and the cyclic loss [GR19], as regularization terms to our alignment loss.

Another concept relevant to our discussion is the self-functional map [HK18]. There, Halimi *et al.* proposed a method for shape matching by using descriptors which are defined by a multiplication of SI-LBO eigenfunctions and LBO ones. Theoretically, the proposed framework can be viewed as an intermediate point between the self-functional maps and the unsupervised FM-Net. At one end, we propose to learn the rotation relating the eigenstructures of the two shapes, though there is no deep learning involved. At the other end, we operate within a similar unsupervised Gromov-distance-like GMDS-loss [HLR\*19, BBK06b, ADK16], to obtain accurate dense shape matching.

### 3. Background

**Spectral domain on surfaces.** Here, we lay the groundwork theory for continuous surfaces. However, it worth mentioning that our discrete implementation operates on triangulated surfaces. A two dimensional manifold, also called a surface  $S : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$ , can be used to define a Riemannian manifold  $(S, g)$ , when equipped with a metric tensor  $g$ . The most familiar example is the Euclidean metric tensor

$$g = (g_{ij}) \triangleq (\langle S_i, S_j \rangle). \quad (1)$$

The metric tensor can be used to define differential operators on the manifold such as the *Laplace-Beltrami Operator* (LBO). Given the parametric surface  $S(u, v)$ , an infinitesimal length element on  $S$  is defined by

$$ds^2 = (dudv) \begin{pmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \end{pmatrix} \begin{pmatrix} du \\ dv \end{pmatrix}, \quad (2)$$

where the substitution of the Euclidean metric tensor namely,  $g_{11} = |S_u|^2$ ,  $g_{12} = g_{21} = \langle S_u, S_v \rangle$ ,  $g_{22} = |S_v|^2$ , yields the Euclidean length element. Let  $g \equiv \det(g)$  and  $(g^{ij}) = (g_{ij})^{-1}$ , then, using Einstein summation convention, the LBO can be written as

$$-\Delta_g = \frac{1}{\sqrt{g}} \partial_i \sqrt{g} g^{ij} \partial_j, \quad (3)$$

or explicitly, operating on a function  $f : S \rightarrow \mathbb{R}$ ,

$$-\Delta_g f = \frac{1}{\sqrt{g}} \partial_i \left( \sqrt{g} g^{ij} \partial_j f \right) \quad (4)$$

where  $\partial_i$  is a derivative with respect to the  $i^{\text{th}}$  coordinate,  $u$  or  $v$ , respectively. This intrinsic definition in terms of the metric tensor and its derivatives, makes the LBO invariant to isometries of the surface, i.e. deformations that do not change the metric. In the case of the Euclidean metric defined above, isometries amount to different embedding of the surface in  $\mathbb{R}^3$  without stretching and tearing the surface. LBO is a positive semi-definite operator, therefore the spectral theorem guarantees that the eigenfunctions form an orthonormal basis, that is,

$$\begin{aligned} \Delta_g \phi_i &= \lambda_i \phi_i \\ \langle \phi_i, \phi_j \rangle_g &= \delta_{ij}. \end{aligned} \quad (5)$$

**Scale invariant LBO.** The scale invariant metric was first introduced in [AKR13]. It is defined by the pseudo-metric

$$\tilde{g}_{ij} = |K| g_{ij}, \quad (6)$$

that can be approximated by the metric,

$$\tilde{g}_{ij} = \sqrt{(K^2 + \epsilon^2)} g_{ij}, \quad (7)$$

where  $(g_{ij})$  is the Euclidean metric defined in Eq. (1),  $K$  is the Gaussian curvature, and  $\epsilon$  is a small positive constant preventing the expression from vanishing. Differently from the Euclidean metric, where the distance between two surface points is defined as integrating the locally euclidean geodesic distance, the scale-invariant metric is defined by in Eq. (6). Intuitively, it is realized by infinitesimally scaling the locally Euclidean length by the radius of a circle defined by the surface. One intrinsic choice for such a local scaling factor is  $\rho = |K|^{-\frac{1}{2}}$  the squared root of the Gaussian curvature. The scale-invariant metric  $(\tilde{g}_{ij})$  induces a new LBO introduced by Aflalo *et al.* as the *scale-invariant Laplace-Beltrami Operator* (SI-LBO). The SI-LBO is defined by plugging the the definition of the scale-invariant metric from Eq. (7) into the the LBO definition in Eq. (3), namely,

$$-\Delta_{\tilde{g}} = \frac{1}{\sqrt{\tilde{g}}} \partial_i \sqrt{\tilde{g}} \tilde{g}^{ij} \partial_j. \quad (8)$$

Since the metric  $\tilde{g}$  is scale invariant and invariant to isometries with respect to  $g$ , the SI-LBO is also invariant to these deformations. As before, since it is a positive semi-definite operator, the eigenfunctions form an orthonormal basis. That is,

$$\begin{aligned} \Delta_{\tilde{g}} \tilde{\phi}_i &= \tilde{\lambda}_i \tilde{\phi}_i \\ \langle \tilde{\phi}_i, \tilde{\phi}_j \rangle_{\tilde{g}} &= \delta_{ij}. \end{aligned} \quad (9)$$

**Functional Maps.** Functional maps were first introduced in [OBBS\*12]. There, Ovsjanikov *et al.* showed conceptual equivalence between point-wise correspondence and functional mapping between two different surfaces, and formulated the latter as a linear transformation  $T : \mathcal{F}(\mathcal{X}, \mathbb{R}) \rightarrow \mathcal{F}(\mathcal{Y}, \mathbb{R})$  between the functional spaces defined on the surfaces  $\mathcal{X}$  and  $\mathcal{Y}$ . Specifically, a given real valued function  $f \in \mathcal{F}(\mathcal{X}, \mathbb{R})$  on surface  $\mathcal{X}$ , with the basis representation  $f = \sum_i a_i \phi_i^{\mathcal{X}}$  in a basis  $\{\phi_i^{\mathcal{X}}\}$  that spans  $\mathcal{F}(\mathcal{X}, \mathbb{R})$ , will be

mapped by  $T$  to the surface  $\mathcal{Y}$  as follows,

$$T(f) = \sum_j a_j T(\phi_j^{\mathcal{X}}) = \sum_{i,j} c_{ij} a_j \phi_j^{\mathcal{Y}}, \quad (10)$$

where the translation to the coefficients in the basis  $\{\phi_i^{\mathcal{Y}}\}$  of  $\mathcal{F}(\mathcal{Y}, \mathbb{R})$  is obtained by the functional map  $c_{ij} = \langle \phi_i^{\mathcal{Y}}, T(\phi_j^{\mathcal{X}}) \rangle_{\mathcal{Y}}$ . The functional map  $\mathbf{C} = (c_{ij})$  is derived from a set of linear constraints imposed by a set of corresponding descriptor functions pairs  $\{(f_i, h_i)\}$  on the surfaces  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively. Defining  $\hat{\mathbf{F}}_{ij} = \langle \phi_i^{\mathcal{X}}, f_j \rangle$  and  $\hat{\mathbf{H}}_{ij} = \langle \phi_i^{\mathcal{Y}}, h_j \rangle$ , the matrix  $\mathbf{C}$  can be computed by solving

$$\mathbf{C}\hat{\mathbf{F}} = \hat{\mathbf{H}}, \quad (11)$$

using a least squares method. As stated at [OBBCS\*12] the functional map framework is not constrained to a specific choice of basis, however, in [ABK15], Aflalo *et al.* showed that the LBO eigenfunctions are the optimal  $k$ -dimensional reduced basis and that this selection, which can be regarded as a low pass filter operating on the mapped functions, leads to a smooth correspondence map. Therefore, similar to [HK18], in the proposed method we set the basis functions to the LBO eigenfunctions, denoted by  $\Phi$  and  $\Psi$ , on surfaces  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively. In this paper, we use the SI-LBO eigenfunctions denoted by  $\tilde{\Phi}$  and  $\tilde{\Psi}$ , as descriptors on surfaces  $\mathcal{X}$  and  $\mathcal{Y}$ , and we calculate the functional map by solving (11) with the following matrices

$$\hat{\mathbf{F}}_{ij} = \langle \phi_i, \tilde{\phi}_j \rangle_g, \quad (12)$$

and

$$\hat{\mathbf{H}}_{ij} = \langle \psi_i, \tilde{\psi}_j \rangle_g. \quad (13)$$

As previously discussed, when computed, the eigenfunctions are very poorly aligned undermining the assumption of compatible descriptor functions. Here, we derive an optimization process producing well aligned SI-LBO eigenfunctions. We used the aligned SI-LBO eigenfunctions in the definition of the matrices (12), (13), then, solve for the functional map  $\mathbf{C}$ , and finally extract from  $\mathbf{C}$  the point-wise correspondence.

**Deep functional maps.** Our optimization algorithm for the alignment of the eigenfunctions is based on FM-Net architecture introduced in [LRR\*17]. In their paper, Litany *et al.* suggested to use a neural network in order to find descriptors tuned to be optimal for the functional map framework. Their residual network [HZRS16] based architecture receives a hand-crafted semi-local extrinsic descriptor called SHOT [STDS14], as an input, and outputs a learned descriptor of the same dimension, optimized for the network training task. They trained the network with the following loss function

$$\ell_{\text{sup}}(\mathcal{X}, \mathcal{Y}) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p_{y|x} d_{\mathcal{Y}}^2(y, \pi^*(x)), \quad (14)$$

where  $\mathbf{P}_{i,j} = p_{y(i)|x(j)}$  is a stochastic matrix derived from the functional map  $\mathbf{C}$ , indicating the probability that a vertex  $x(i)$  will be mapped to the vertex  $y(j)$ . The internal sum measures for each point  $x \in \mathcal{X}$ , the expectation of squared geodesic distance between the predicted matched point  $y$ , with probability  $q_{y|x}$ , and the ground-truth mapping  $\pi^*(x)$ . Then, the external sum averages this expectation over the points in surface  $\mathcal{X}$ . The relation between the

functional map  $\mathbf{C}$  and the stochastic matrix  $\mathbf{P}$  is given by

$$\mathbf{P} = \left( |\Psi \mathbf{C} \Phi^T| \wedge \right)^\circ, \quad (15)$$

where  $|\cdot|$  is a pointwise absolute value,  $\wedge$  is a column-wise  $L_2$  normalization and  $\circ$  is a self-Hadamard product. For further details please refer to [HZRS16, HLR\*19]. In the end of the training, the hard correspondence  $\mathbf{\Pi}$  is extracted from the soft correspondence  $\mathbf{P}$  as follows

$$\mathbf{\Pi}_{ij} = \begin{cases} 1 & \text{if } i = \operatorname{argmax}_k \mathbf{P}_{kj} \\ 0 & \text{else} \end{cases}. \quad (16)$$

**Unsupervised deep functional maps.** Our alignment method leverages the unsupervised loss introduced in [HLR\*19] as a part of the optimization objective. While FM-Net [LRR\*17] is a supervised method, Halimi *et al.* [HLR\*19] used an unsupervised loss based on the generalized multidimensional scaling (GMDS) [BBK06b], that does not require the ground-truth correspondence for training. The unsupervised loss requires only the inter-geodesic distances matrix, where each entry represents the surface distance between two surface points, without the need for point to point matching that is required in the FM-Net model. The suggested *unsupervised* loss measures the inter geodesic distance distortion of the correspondence, that is,

$$\ell_{\text{unsup}}(\mathcal{X}, \mathcal{Y}) = \frac{1}{|\mathcal{X}|^2} \|\mathbf{D}_{\mathcal{X}} - \mathbf{P}^T \mathbf{D}_{\mathcal{Y}} \mathbf{P}\|_{\mathbb{F}}^2. \quad (17)$$

Where  $\mathbf{D}_{\mathcal{X}}$  and  $\mathbf{D}_{\mathcal{Y}}$  are the inter-geodesic distances matrices of the shapes  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively, and  $\mathbf{P}$  is the stochastic matrix derived from the functional map  $\mathbf{C}$ , as in Eq. (15).

#### 4. Method

Our method solves efficiently the misalignment of the SI-LBO eigenfunctions, reclaiming their power as intrinsic independent descriptors that also span any function on the surface. We start by discussing our alignment model, then we define a novel optimization problem to find this alignment, finally we describe the framework used to solve the optimization problem. The final result is 1) two sets of well aligned eigenfunctions and 2) point-wise correspondence obtained by applying these sets as descriptors in the functional map framework.

**Alignment Model.** Let  $\{\tilde{\phi}_i\}$  and  $\{\tilde{\psi}_i\}$  be two set of SI-LBO eigenfunctions, computed on the (quasi-) isometric surfaces  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively, and let  $\tilde{\Phi}, \tilde{\Psi}$  be the matrices containing the eigenfunctions as columns. In practice, the ordered sets of eigenfunctions are a-priori incompatible due to sign ambiguities and possible permutation between adjacent eigenfunctions. While these ambiguities are generally accounted for in previous methods that aims to align the LBO eigenfunctions [SK14, HK18], sign and permutation alone fall short of spanning the whole variety of possible transformations between the sets of eigenfunctions. For example, Fig. 2a demonstrates a case in which the eigenfunctions evaluated on shape  $\mathcal{X}$  are given as a superposition of the eigenfunctions of shape  $\mathcal{Y}$ , using the ground truth mapping between  $\mathcal{X}$  and shape  $\mathcal{Y}$ .

Therefore, it is clear that the alignment transformation should be relaxed to a broader form. The first improvement we suggest is to

(a) The transformation that aligns the 2<sup>nd</sup> and 3<sup>rd</sup> eigenfunction of two different shapes is given by  $\begin{pmatrix} \tilde{\Phi}_2 \\ \tilde{\Phi}_3 \end{pmatrix} \approx \begin{pmatrix} 0.89 & -0.35 \\ 0.4 & -0.91 \end{pmatrix} \begin{pmatrix} \Psi_2 \\ \Psi_3 \end{pmatrix}$ . It represents a rotation in the spectral domain and cannot be captured by changing the signs and/or permutation, which poses a problem for other methods like [SK14].

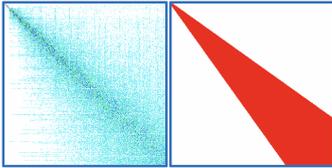
(b) The aligned eigenfunctions after the optimization process. Our algorithm successfully aligned these eigenfunction by finding a suitable rotation in the spectral domain.

**Figure 2:** Subspace alignment of eigenfunctions.

model the alignment as a *band-limited orthogonal transformation*  $\tilde{\mathbf{C}}$ , that admits

$$\mathbf{\Pi}^* \tilde{\Phi} \tilde{\mathbf{C}} = \tilde{\Psi} \quad \text{or alternatively} \quad \mathbf{\Pi}^* \tilde{\Phi} = \tilde{\Psi} \tilde{\mathbf{C}}^T, \quad (18)$$

implying that the eigenfunctions of shape  $\mathcal{X}$  undergo a rotation such that the ground truth correspondence  $\mathbf{\Pi}^* : \mathcal{X} \rightarrow \mathcal{Y}$  maps the rotated eigenfunctions  $\tilde{\Phi} \tilde{\mathbf{C}}$  exactly to  $\tilde{\Psi}$ . We require that  $\tilde{\mathbf{C}}$  is orthogonal since it should operate on a real-valued orthonormal basis  $\tilde{\Phi}$ , and produce one  $(\mathbf{\Pi}^*)^{-1} \tilde{\Psi}$ , assuming the surfaces are isometric. Our motivation for a band restricted form is that eigenfunctions tend to create superposition, as in Fig 2a, only with adjacent eigenfunctions that correspond to similar eigenvalues. Empirically, this coupling tends to increase with frequency, involving more eigenfunctions, thus leading to funnel shape structure of the numerical support of the matrix  $\tilde{\mathbf{C}}$ , as can be seen in Fig. 3. We understand



**Figure 3:** Left: Values of  $\tilde{\mathbf{C}}$  mapping between the eigenfunctions of two SI-LBOs corresponding to two surfaces. Indeed, the matrix exhibits a banded unitary (funnel) structure. Right: The mask by which  $\mathbf{R}$  is restricted. None zero values are allowed only in the red region.

this effect as follows; In [ASC11], Aubry *et al.* analysed the error between corresponding eigenvalues of two quasi-isometric shapes, and found it is proportional to their magnitudes. On the other hand, according to Weyl’s Law applied to 2-manifolds, the asymptotic increment between eigenvalues is constant. Combining the two arguments, as the eigenvalues increase, the ratio between the error and the spectral separation behaves asymptotically as a linear function

of the eigenvalues, and therefore of  $N(\lambda)$ , the index of the eigenvalue. Intuitively, this ratio determines the coupling radius in the spectral domain, which is the maximal difference between the indices of eigenfunctions, that still couple beyond some threshold. Since this radius grows linearly with  $N(\lambda)$ , we model the funnel width as a linear function of  $N(\lambda)$ . The specific equation appears in the next section.

The second improvement in the alignment model is achieved by observing that every orthogonal transformation can be decomposed into a multiplication of two identical unitary transformations, namely,

$$\tilde{\mathbf{C}} = \mathbf{R} \mathbf{R} \quad (19)$$

where  $\mathbf{R}$  represents half of the rotation applied by the original transformation  $\tilde{\mathbf{C}}$ . Thus, we replace Eq. (18) by the requirement for a *band-limited orthogonal transformation*  $\mathbf{R}$ , that admits

$$\mathbf{\Pi}^* \tilde{\Phi} \mathbf{R} = \tilde{\Psi} \mathbf{R}^T \quad (20)$$

Fig. 2b shows that our alignment model can recover the compatible eigenfunctions even under complex misalignment transformations. Empirically, we found that the second model, that apply the alignment on both shapes, is easier to optimize for compared to the model that operate only on a single shape, thus, we adopt it to our method.

**Optimization objective.** Eq. (20) cannot be solved directly to extract the optimal alignment matrix  $\mathbf{R}$  since we do not assume that we have access to the ground-truth correspondence  $\mathbf{\Pi}^*$  between the shapes  $\mathcal{X}$  and  $\mathcal{Y}$ . In fact the point-wise correspondence is the target of our method, together with the optimal alignment. Therefore, we design an optimization problem that simultaneously solves for both. We define a novel optimization problem based on the alignment equation introduced in (20)

$$\ell_{\text{align}}(\mathbf{R}) = \frac{\|\mathbf{\Pi}(\mathbf{R}) \tilde{\Phi} \mathbf{R} - \tilde{\Psi} \mathbf{R}^T\|_{\mathcal{L}^2(\mathcal{Y},g)}^2}{\|\tilde{\Psi} \mathbf{R}^T\|_{\mathcal{L}^2(\mathcal{Y},g)}^2}, \quad (21)$$

where  $\mathbf{\Pi}(\mathbf{R})$  is no longer the ground-truth correspondence, but rather the hard point-wise correspondence obtained by using  $\tilde{\Phi} \mathbf{R}$  and  $\tilde{\Psi} \mathbf{R}^T$  as descriptors in the FM pipeline, implying it depends on  $\mathbf{R}$ . The specific form of the hard correspondence  $\mathbf{\Pi}(\mathbf{R})$  is according to Eq. (16) in the background section, and we elaborate it explicitly in the next section. While in theory, the correct alignment  $\mathbf{R}^*$  is the global optimum of  $\ell_{\text{align}}$ , in practice it is challenging to reach this optimal value due to mutual dependency of the optimization variables, in the loss definition. Specifically, the evaluation of  $\mathbf{R}$  is unreliable due to the definition of  $\ell_{\text{align}}$  in terms of  $\mathbf{\Pi}(\mathbf{R})$ , which is erroneous in the beginning of the optimization process and vice versa,  $\mathbf{\Pi}(\mathbf{R})$  is evaluated unreliably as long as  $\mathbf{R}$  hasn’t converged. To stabilize the optimization, we introduced two additional terms that operate separately on each of the variables. First, we evaluate the precursor of  $\mathbf{\Pi}(\mathbf{R})$  in the computation pipeline, namely  $\mathbf{P}(\mathbf{R})$ , see Eq. (15) and also the elaboration in the next section. Specifically, the soft correspondence matrix  $\mathbf{P}(\mathbf{R})$  is evaluated, using the unsupervised loss in Equation (17), that is,

$$\ell_{\text{unsup}}(\mathbf{R}) = \frac{1}{|\mathcal{X}|^2} \|\mathbf{D}_{\mathcal{X}} - \mathbf{P}(\mathbf{R})^T \mathbf{D}_{\mathcal{Y}} \mathbf{P}(\mathbf{R})\|_{\mathbb{F}}^2. \quad (22)$$

This loss introduced in [HLR\*19] measures the metric distortion induced by  $\mathbf{P}(\mathbf{R})$ , and will vanish only when  $\mathbf{P}(\mathbf{R})$  approaches an isometric permutation. As opposed to  $\mathbf{\Pi}(\mathbf{R})$ ,  $\mathbf{P}(\mathbf{R})$  doesn't include *argmax* operation, rendering the optimization objective  $\ell_{\text{unsup}}(\mathbf{R})$  smoother. Secondly, we considered the following loss term, evaluating to which extent  $\mathbf{R}$  satisfies the requirement of being orthogonal. To this end, we introduce

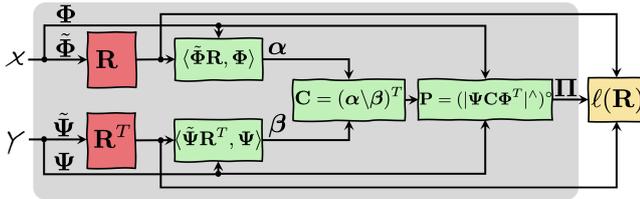
$$\ell_{\text{uni}}(\mathbf{R}) = \frac{1}{k^2} \|\mathbf{R}\mathbf{R}^T - \mathbf{I}\|_F^2, \quad (23)$$

where  $k$  is the number of eigenfunctions and  $\mathbf{I}$  is the identity matrix. However, in practice this loss is extremely non-smooth and we found that it makes the optimization challenging. Therefore, we decided not to use it as a regularization. As a final optimization objective we used

$$\ell(\mathbf{R}) = \ell_{\text{align}}(\mathbf{R}) + \lambda \ell_{\text{unsup}}(\mathbf{R}). \quad (24)$$

Interestingly, while we optimized exclusively on  $\ell(\mathbf{R})$ , we found that  $\ell_{\text{uni}}(\mathbf{R})$  has strong correlation with  $\ell_{\text{align}}(\mathbf{R})$ ; While only monitored,  $\ell_{\text{uni}}(\mathbf{R})$  followed the trend of  $\ell_{\text{align}}(\mathbf{R})$  and eventually reached a very low value. This phenomena is demonstrated in the supplementary. The band restricted support of  $\mathbf{R}$  wasn't penalized for, but rather enforced structurally, as discussed in the next section.

**Optimization framework.** To find the alignment  $\mathbf{R}$  that opti-



**Figure 4:** The proposed numerical solver for the alignment  $\mathbf{R}$ , starting with the input SI-LBO eigenfunctions  $\tilde{\Psi}$ ,  $\tilde{\Phi}$ , that are multiplied by matrix  $\mathbf{R}$  and  $\mathbf{R}^T$ , respectively. Then, the results are used as functional correspondences for the functional map scheme with  $\Psi$  and  $\Phi$  as bases. The output is the dense correspondence matrix  $\mathbf{P}$ . Properly projecting  $\mathbf{P}$  into the permutation  $\mathbf{\Pi}$  we compute the loss function and back-propagate the gradient to optimize for  $\mathbf{R}$ .

mizes the loss (24), we employ the core Functional Map module of [LRR\*17] in a shallow network designed as a numerical solver, see Fig. 4. network optimizes over the the entries in the matrix  $\mathbf{R}$ , using back-propagation to estimate the loss gradient with respect to  $\mathbf{R}$  and to update the entries every optimization step. To account for the band restricted structure of  $\mathbf{R}$ , we enforced the funnel shape by multiplying  $\mathbf{R}$  element-wise with a funnel shaped binary mask, as can be seen in Fig. 3 (right). We modeled the support of the mask matrix using a linearly increasing band width

$$W_i = \left\lceil \frac{i}{a} \right\rceil + b, \quad (25)$$

where  $W_i$  is the width of the  $i$ 'th row, and  $b$  and  $a$  are hyper-parameters. Next, we describe the architecture of the optimizer. Given two input shapes  $\mathcal{X}$  and  $\mathcal{Y}$ , we extract the SI-LBO eigenfunctions of each shape, denoted by  $\tilde{\Psi}$  and  $\tilde{\Phi}$ , respectively. The

first shape's eigenfunctions,  $\tilde{\Phi}$ , are then multiplied by a band unitary matrix  $\mathbf{R}$ , and the second shape's eigenfunctions,  $\tilde{\Psi}$ , are multiplied by the transpose of the same matrix  $\mathbf{R}^T$ . We use the aligned bases  $\tilde{\Phi}\mathbf{R}$  and  $\tilde{\Phi}\mathbf{R}^T$  as descriptors for the functional map scheme to extract the functional correspondence  $\mathbf{C}$ , similarly to equation (11). We define the projection of of the descriptors on the LBO eigenbasis as

$$\begin{aligned} \alpha &\equiv \langle \tilde{\Phi}\mathbf{R}, \Phi \rangle_{g_{\mathcal{X}}} \in \mathbb{R}^{d \times k} \\ \beta &\equiv \langle \tilde{\Psi}\mathbf{R}^T, \Psi \rangle_{g_{\mathcal{Y}}} \in \mathbb{R}^{d \times k}, \end{aligned} \quad (26)$$

where  $d$  is the dimension of the descriptor (= number of SI-LBO eigenfunctions),  $k$  is the number of LBO eigenfunctions, and finally,  $g_{\mathcal{X}}$  and  $g_{\mathcal{Y}}$  are the Euclidean metric tensors on surfaces  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively, by which the corresponding LBO's were constructed (note that the projected scale invariant descriptors,  $\alpha$  and  $\beta$ , are in fact the aligned *self functional maps* of each shape, as introduced in [HK18]). Then, we calculate the functional map  $\mathbf{C}^T = \alpha \setminus \beta$ . Here,  $\setminus$  denotes the least squares solver. Then, using Eq. (15), we construct a soft-correspondence stochastic matrix  $\mathbf{P}$ , similarly to [HLR\*19, LRR\*17],

$$\mathbf{P} = (|\Psi \mathbf{C} \Phi^T|^w)^o. \quad (27)$$

The hard correspondence  $\mathbf{\Pi}$  is then extracted from the soft correspondence  $\mathbf{P}$  as follows. Each column of the matrix  $\mathbf{P}$  is associated with a vertex from  $\mathcal{X}$ , and the index of the maximal value along each column represents the corresponding vertex in  $\mathcal{Y}$ . Therefore, we approximate the permutation matrix  $\mathbf{\Pi}$ , by

$$\Pi_{ij} = \begin{cases} 1 & \text{if } i = \operatorname{argmax}_k P_{kj} \\ 0 & \text{else} \end{cases}. \quad (28)$$

Finally both  $\mathbf{\Pi}$  and  $\mathbf{R}$  (as well as  $\mathbf{R}^T$ ) are used in the loss definition (24), and the convergence of this loss produces simultaneously the optimal alignment  $\mathbf{R}$ , and the point-wise correspondence  $\mathbf{\Pi}$ .

## 5. Implementation

**Implementation considerations.** The proposed framework was implemented in Tensorflow [ABC\*16], while the data pre-processing in Matlab. It is important to stress that there is no deep learning involved, though utilizing the optimization mechanism of TensorFlow has proven to be really useful. In our experiments, we used 352 SI-LBO eigenfunctions as descriptors, the same number used in the [HLR\*19, LRR\*17]. The parameter  $\lambda$  in the optimization objective (24) was empirically set to 200. The hyper-parameters of the mask model in Eq. (25) were set to  $a = 2.5$  and  $b = 2$ . We initialized  $\mathbf{R}$  as a diagonal matrix, where each diagonal element was randomly set to 1 or  $-1$ . Since the optimal  $\mathbf{R}$  has a band unitary structure, this initialization gives the optimization process a head start. We used 1000 iterations for the optimization process. When evaluating the gradient of term (21), we considered only the gradient of the numerator with respect to the alignment matrix  $\mathbf{R}$ , while the denominator was only used for dynamic scaling of the gradient. Our algorithm consumes approximately one minute for building the computational graph in Tensorflow, and six minutes for the optimization process, that was performed on Nvidia

GTX 2080 GPU. This time does not include hyper-parameter tuning; As described below, our algorithm uses a hyper-parameter  $\xi$  for the definition of the SI-LBO operator. For more details, see the above section and the supplementary material.

**Pre-processing.** We evaluated our method on FAUST [BRLB14] dataset. For each model, we calculated the LBO eigenfunctions and the eigenfunctions of the generalized SI-LBO, derived from a metric tensor of the general form

$$\tilde{g}_{ij} = |K|^\xi g_{ij}. \quad (29)$$

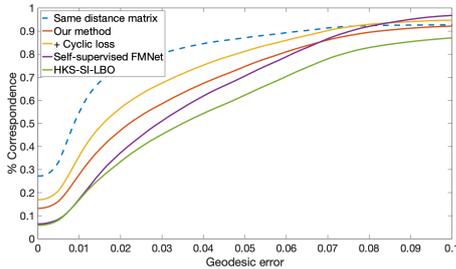
The discretization of generalized SI-LBO operator is given by

$$\tilde{L} = (|K|^\xi A)^{-1} W, \quad (30)$$

which, for simplicity, will be referred to as SI-LBO. Here,  $K$  is the Gaussian curvature,  $W$  is the cotangent matrix and  $A$  is the vertex area diagonal matrix. The scalar  $\xi \in (0, 1]$  determines a metric interpolating between the regular and the scale invariant one, was chosen to be the one providing the best match for each pair of shapes, see supplementary.

## 6. Experiments

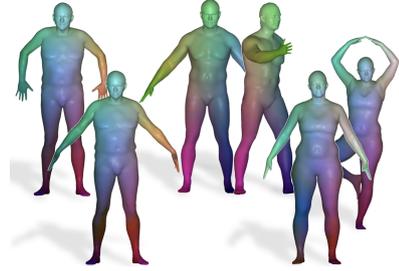
**Simultaneous correspondence and alignment.** We show correspondence results on FAUST [BRLB14] dataset, reporting the correspondence error on 100 pairs; For each subject, we evaluate the correspondence between the first pose and the other poses. We compare our method with [HLR\*19], running it in the self-



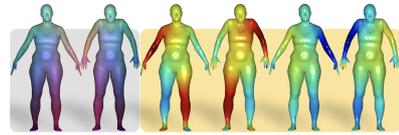
**Figure 5:** Comparison of the matching quality of the different methods proposed in the paper, and the self-supervised FMNet. We see the potential quality in the case of perfect isometry, which explains the improvement of adding the cyclic loss.

supervised operation mode, on each of the pairs. The correspondence results are reported in Fig. 5, showing similar performance of the two methods. Our method is more accurate than the self-supervised network, as can be seen in the domain of small geodesic radius, while it also suffers slightly more of global mismatches, scoring a bit lower for large geodesic radius. We emphasize that our descriptors are the aligned SI-LBO eigenfunction, while the descriptors of the self-supervised method are the output of a deep network. Fig. 1 demonstrates qualitatively the aligned sets of eigenfunctions for different ranges of eigenvalues. As opposed to previous alignment methods, like [SK14], that only showed alignment of the few first eigenfunctions ( $1 \leq k \leq 4$ ), our alignment is scalable in the number of eigenfunctions and we show consistent alignment up to high frequencies. Fig. 9 compares the alignment error of

our method to the unaligned SI-LBO eigenfunctions. We observe a significant improvement of the alignment in a very broad range of frequencies. Finally, Fig. 6 shows the qualitative correspondence results.



**Figure 6:** Three couples of matching human body shapes at different poses. Similar colors correspond to matched points.



**Figure 7:** A case where the method fails to properly align the eigenfunctions due to symmetry of the optimization objective. In left gray is the correspondence produced by the algorithm. Here, the correspondence of the body is reflected (right-left). The right box shows the aligned eigenfunctions, in which the anti-symmetric eigenfunctions were reflected.

**Analysis.** In this experimental part, we aim to study the core limitations of the proposed method, in order to improve the algorithm. Our first concern in the source of the global mismatches that degrades the performance in Fig. 5. Analysing the qualitative correspondence results, we observed that in a small fraction of the examples the algorithm converges to the reflection of the correct mapping, being also an isometry. In this case, the optimal alignment produces two anti-symmetric eigenfunctions that are related by reflection. This phenomena is demonstrated in Fig. 7. The second issue we would like to explore, is how the deviation from perfect isometry influences our alignment algorithm. We note that this deviation influences twice; First, by perturbing the SI-LBO eigenfunctions, introducing misalignment. Secondly, by slightly misleading our alignment algorithm, through the  $\ell_{\text{unsup}}(\mathbf{R})$  loss term, which is based on *perfect* isometry. To analyse the second effect we artificially inject the distance matrix of shape  $\mathcal{X}$  into the distance matrix of  $\mathcal{Y}$ , ending up with  $D_{\mathcal{X}} = D_{\mathcal{Y}}$ . This way, we can examine our optimization algorithm in the artificial setting of perfect isometry. We show correspondence and alignment results in Fig. 5 and Fig. 9, respectively. We observe a significant improvement of both the alignment and correspondence results. By neutralizing the effect of deviation between the axiomatic model and the real data, the eigenfunctions are better aligned, as can be seen in Fig. 8, showing a failure case of our method which is aligned correctly in the absence of this deviation.

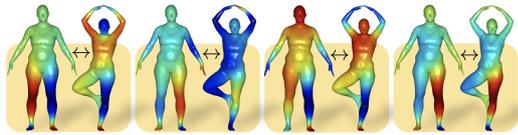
**Refinement with cycle-consistency loss.** Here, we use the analysis performed in previous section to improve the basic algorithm. Specifically, we apply the conclusion that the deviation from perfect isometry has significant influence on our optimization process. Therefore, we would like to use additional reliable priors to compensate for the limitation of this approximation. As such strong reliable prior, we consider the cycle consistency loss, introduced in [GR19]. We add the loss term

$$\ell_{\text{cyc}}(\mathbf{R}) = \frac{1}{|\mathcal{X}|^2} \|\mathbf{D}_{\mathcal{X}} - \mathbf{P}(\mathbf{R})^T \tilde{\mathbf{P}}(\mathbf{R})^T \mathbf{D}_{\mathcal{X}} \tilde{\mathbf{P}}(\mathbf{R}) \mathbf{P}(\mathbf{R})\|_{\text{F}}^2 \quad (31)$$

as additional regularization to our optimization objective, where  $\mathbf{P}$  is given by Eq. (27) and  $\tilde{\mathbf{P}}$  is obtained similar to  $\mathbf{P}$ , by replacing the roles of  $\mathcal{X}$  and  $\mathcal{Y}$ . Note, that this loss must be added to  $\ell_{\text{unsup}}(\mathbf{R})$ , rather than replace it, as bijection is a too permissive requirement. There are many undesired bijections which are not isometric at all. Specifically, the optimization objective we used is

$$\ell(\mathbf{R}) = \ell_{\text{align}}(\mathbf{R}) + \lambda \frac{\ell_{\text{unsup}}(\mathbf{R}) + \ell_{\text{cyc}}(\mathbf{R})}{2}, \quad (32)$$

replacing  $\ell_{\text{unsup}}(\mathbf{R})$  by the average of  $\ell_{\text{unsup}}(\mathbf{R})$  and  $\ell_{\text{cyc}}(\mathbf{R})$ , using the same value for  $\lambda$  as in the original objective. We show alignment and correspondence results in Fig. 9 and Fig. 5, respectively. We conclude that the requirement for cycle-consistency resulted in significant improvement of the alignment algorithm. As can be seen in Fig. 8, the eigenfunctions are aligned better in some of the cases, that failed with the original optimization objective.

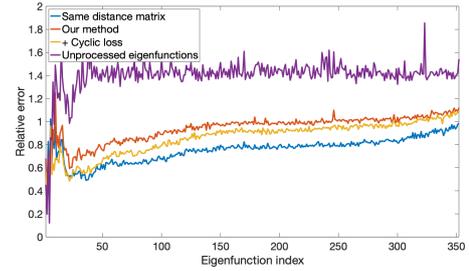


**Figure 8:** The improvement of the alignment assuming isometry is preserved between poses. Left to right: the left couple displays the 7<sup>th</sup> scale-invariant eigenfunctions of the two shapes, the next couple shows the aligned eigenfunctions using our algorithm, next couple is the alignment output in the case where the inter-geodesic distances are forced to be identical. That is, the inter-geodesic distances of the left surface were used also as those of right surface, using the given ground truth correspondence. Finally, the right most couple shows the alignment result for the cycle loss (without restricting the distances to be the same).

**A modified scale-invariant LBO.** The relation between the Gaussian curvature and the heat kernel signature (HKS) [SOG09] implies that the SI-LBO operator could be computed alternatively, in terms of the HKS. This specific construction is described in the supplementary material. Fig. 5 shows correspondence results, obtained by aligning the eigenfunctions of the modified SI-LBO operator.

## 7. Conclusions

We introduced a new method for shape correspondence, facilitating the aligned scale-invariant LBO eigenfunctions, as a complete orthogonal set of intrinsic descriptors. We suggested to model this alignment as an orthogonal matrix, restricted to linearly increasing



**Figure 9:** The relative alignment error as a function of the eigenfunction index, defined as  $err(k) = \|\text{col}_k(\mathbf{\Pi}^* \tilde{\Psi} \mathbf{R}^T - \tilde{\Phi} \mathbf{R})\|_{\mathcal{L}^2(\mathcal{Y},g)} / \|\text{col}_k(\tilde{\Psi} \mathbf{R}^T)\|_{\mathcal{L}^2(\mathcal{Y},g)}$ . The proposed method improves the alignment for all eigenfunctions compared to the original eigenfunctions.

band support, that represents the coupling radius of the eigenfunctions. To extract the alignment matrix, we optimized a novel alignment loss, using the *generalized multidimensional scaling* based unsupervised loss, and the *cycle-consistency loss*, as regularization terms. For the optimization process we constructed an effective numerical solver, based on the core module of FM-Net, stripping away the deep network components. By this, we regained the effectiveness of the scale-invariant LBO eigenfunctions as competitive descriptors, and achieved state-of-the-art correspondence results. We analysed the proposed algorithm to understand its essential limitations, and guide future improvements. We evidenced that the deviation from isometry could be compensated for, by the additional requirement for cycle-consistency. Also, we observed that the reflection symmetry of the optimization objective imposes a limitation. Our conclusion is that by preventing the occurrences of reflected eigenfunctions in the aligned sets, the correspondence improves. We believe that this can be achieved by further requiring the alignment of the triple products, defined by the gradients of the eigenfunctions and the surface normal, as done in [SK14]. Also, in the future, we would like to efficiently extend the current framework to a set of shapes.

## 8. Acknowledgements

This research was partially supported by the Israel Ministry of Science and Technology grant number 3-14719 and the Technion Hiroshi Fujiwara Cyber Security Research Center and the Israel Cyber Directorate. Oshri Halimi is a recipient of the Jabotinsky Fellowship for doctoral students, awarded by the Israel Ministry of Science and Technology.

## References

- [ABC\*16] ABADI M., BARHAM P., CHEN J., CHEN Z., DAVIS A., DEAN J., DEVIN M., GHAWAT S., IRVING G., ISARD M., ET AL.: Tensorflow: A system for large-scale machine learning. In *12th Symposium on Operating Systems Design and Implementation* (2016), pp. 265–283. 6
- [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

- [ADK16] AFLALO Y., DUBROVINA A., KIMMEL R.: Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision* 118, 3 (2016), 380–392. 2, 3
- [AKR13] AFLALO Y., KIMMEL R., RAVIV D.: Scale invariant geometry for nonrigid shapes. *SIAM Journal on Imaging Sciences* 6, 3 (2013), 1579–1597. 3
- [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. 1, 2, 5
- [BBBK08] BRONSTEIN A. M., BRONSTEIN M. M., BRUCKSTEIN A. M., KIMMEL R.: Analysis of two-dimensional non-rigid shapes. *International Journal of Computer Vision* 78, 1 (Jun 2008), 67–88. 2
- [BBK06a] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Expression-invariant representations of faces. *IEEE Transactions on Image Processing* 16, 1 (2006), 188–197. 1
- [BBK06b] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences* 103, 5 (2006), 1168–1172. 1, 2, 3, 4
- [BK17] BURGHARD O., KLEIN R.: Efficient lifted relaxations of the quadratic assignment problem. In *Proceedings of the conference on Vision, Modeling and Visualization* (2017), Eurographics Association, pp. 119–128. 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. 7
- [EKB\*15] EYNARD D., KOVNATSKY A., BRONSTEIN M. M., GLASHOFF K., BRONSTEIN A. M.: Multimodal manifold analysis by simultaneous diagonalization of laplacians. *IEEE transactions on pattern analysis and machine intelligence* 37, 12 (2015), 2505–2517. 2
- [GR19] GINZBURG D., RAVIV D.: Cyclic functional mapping: Self-supervised correspondence between non-isometric deformable shapes. *arXiv preprint arXiv:1912.01249* (2019). 3, 8
- [HK18] HALIMI O., KIMMEL R.: Self functional maps. In *2018 International Conference on 3D Vision (3DV)* (2018), IEEE, pp. 710–718. 3, 4, 6
- [HLR\*19] HALIMI O., LITANY O., RODOLA E., BRONSTEIN A. M., KIMMEL R.: Unsupervised learning of dense shape correspondence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 4370–4379. 2, 3, 4, 6, 7
- [HRAK19] HALIMI O., RAVIV D., AFLALO Y., KIMMEL R.: Computable invariants for curves and surfaces. In *Handbook of Numerical Analysis*, vol. 20. Elsevier, 2019, pp. 273–314. 1
- [HZRS16] HE K., ZHANG X., REN S., SUN J.: Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778. 4
- [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, Wiley Online Library, pp. 439–448. 2
- [LRR\*17] LITANY O., REMEZ T., RODOLA E., BRONSTEIN A., BRONSTEIN M.: Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE International Conference on Computer Vision* (2017), pp. 5659–5667. 2, 3, 4, 6
- [Mém11] MÉMOLI F.: Gromov–wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics* 11, 4 (2011), 417–487. 2
- [MS05] MÉMOLI F., SAPIRO G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics* 5, 3 (2005), 313–347. 1, 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), 1–11. 2, 3, 4
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 1555–1564. 2
- [RSO19] ROUFOSSE J.-M., SHARMA A., OVSJANIKOV M.: Unsupervised deep learning for structured shape matching. In *Proceedings of the IEEE international conference on computer vision* (2019), pp. 1617–1627. 2, 3
- [Rus07] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the fifth Eurographics symposium on Geometry processing* (2007), Eurographics Association, pp. 225–233. 1, 2
- [SK14] SHTERN A., KIMMEL R.: Matching the lbo eigenspace of non-rigid shapes via high order statistics. *Axioms* 3, 3 (2014), 300–319. 2, 4, 5, 7, 8
- [SK17] SHAMAI G., KIMMEL R.: Geodesic distance descriptors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2017), pp. 6410–6418. 2
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum* (2009), vol. 28, Wiley Online Library, pp. 1383–1392. 1, 8
- [SPKS16] SOLOMON J., PEYRÉ G., KIM V. G., SRA S.: Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–13. 2
- [STDS14] SALTI S., TOMBARI F., DI STEFANO L.: SHOT: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding* 125 (2014), 251–264. 1, 2, 3, 4
- [WBCPS18] WANG Y., BEN-CHEN M., POLTEROVICH I., SOLOMON J.: Steklov spectral geometry for extrinsic shape analysis. *ACM Transactions on Graphics (TOG)* 38, 1 (2018), 1–21. 1
- [WS19] WANG Y., SOLOMON J.: Intrinsic and extrinsic operators for shape analysis. In *Handbook of Numerical Analysis*, vol. 20. Elsevier, 2019, pp. 41–115. 2
- [WWA\*16] WANG K., WU Z., ALI S., ZHAO J., JIA T., SHUI W., ZHOU M.: *Scale-Invariant Heat Kernel Mapping for Shape Analysis*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016, pp. 74–90. 2
- [WJW\*14] WANG K., WU Z., JIA T., ALI S., ZHAO J., YANG G., ZHOU M.: A scale-invariant diffusion distance for non-rigid shape analysis. In *2014 International Conference on Virtual Reality and Visualization* (2014), pp. 290–295. 2
- [XLZ20] XIANG R., LAI R., ZHAO H.: Efficient and robust shape correspondence via sparsity-enforced quadratic assignment. *arXiv preprint arXiv:2003.08680* (2020). 2
- [YLB\*20] YANG Z., LITANY O., BIRDAL T., SRIDHAR S., GUIBAS L.: Continuous geodesic convolutions for learning on 3d shapes. *arXiv preprint arXiv:2002.02506* (2020). 3
- [ZWL19] ZENG H., WANG Q., LIU J.: Multi-feature fusion based on multi-view feature and 3d shape feature for non-rigid 3d model retrieval. *IEEE Access* 7 (2019), 41584–41595. 1