

Learning Spectral Unions of Partial Deformable 3D Shapes

Luca Moschella¹, Simone Melzi^{1,2}, Luca Cosmo³, Filippo Maggioli¹, Or Litany⁴,
Maks Ovsjanikov⁵, Leonidas Guibas⁶, Emanuele Rodolà¹

¹ Sapienza University of Rome ² University of Milano-Bicocca ³ Ca' Foscari University of Venice ⁴ NVIDIA ⁵ LIX, École Polytechnique, CNRS
⁶ Stanford University

Abstract

*Spectral geometric methods have brought revolutionary changes to the field of geometry processing. Of particular interest is the study of the Laplacian spectrum as a compact, isometry and permutation-invariant representation of a shape. Some recent works show how the intrinsic geometry of a full shape can be recovered from its spectrum, but there are approaches that consider the more challenging problem of recovering the geometry from the spectral information of partial shapes. In this paper, we propose a possible way to fill this gap. We introduce a learning-based method to estimate the Laplacian spectrum of the union of partial non-rigid 3D shapes, without actually computing the 3D geometry of the union or any correspondence between those partial shapes. We do so by operating purely in the spectral domain and by defining the union operation between short sequences of eigenvalues. We show that the approximated union spectrum can be used as-is to reconstruct the complete geometry [MRC*19], perform region localization on a template [RTO*19] and retrieve shapes from a database, generalizing ShapeDNA [RWP06] to work with partialities. Working with eigenvalues allows us to deal with unknown correspondence, different sampling, and different discretizations (point clouds and meshes alike), making this operation especially robust and general. Our approach is data-driven and can generalize to isometric and non-isometric deformations of the surface, as long as these stay within the same semantic class (e.g., human bodies or horses), as well as to partiality artifacts not seen at training time.*

CCS Concepts

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

1. Introduction

Recent progress in spectral geometry processing has brought to significant qualitative leaps that lead to better results in a range of challenging tasks such as deformable shape matching [OBS*12; LRBB17], retrieval [RWP06; BBGO11], style [MRC*19] and pose transfer [KBB*13; YLL*15] among others.

More recently, a great deal of attention has been put on the study of the eigenvalues of the Laplace-Beltrami operator (i.e., the *Laplacian spectrum*) as a compact, isometry and permutation-invariant representation of the input shape. It has been shown that, with the appropriate knowledge on the input domain, this representation contains enough information for localize shape's regions [RTO*19] and even to reconstruct the geometry of the shape [CPR*19; MMRC19]. However, these methods typically operate in a controlled scenario requiring to have access to the *full* geometry of the shape, ignoring the fact that real-world data are riddled with partiality artifacts.

In this paper, we propose a learning-based framework to predict the Laplacian spectrum of the union of two shapes directly from the spectrum of the individual parts. This enables the aforementioned spectral methods to be applied directly in case of par-

tial views of the same shape, without resorting to methods that explicitly fuse the 3D geometry of partial shapes. Indeed, a typical pipeline to combine partial shapes can be very cumbersome, and requires to match the corresponding regions, extract a set of (non-rigid) transformations from the matches, and merge the partial views into a consistent discretization. Each of these steps can be error-prone and difficult to solve, as testified by a wealth of literature on non-rigid shape matching and reconstruction, especially in the case of partial shapes. For example, the mere presence of inconsistent surface sampling can cause problems in most matching pipelines [MMR*19].

Motivated by the excellent results achieved by the methods that exploit the Laplacian spectrum representation, we propose a different perspective. We claim that, in many cases, it is not necessary to have the extrinsic geometry of a target full shape and propose to directly estimate the intrinsic properties of the sought full shape *without* having to materialize its surface geometry. This is done by translating the objective of merging partial shapes from the spatial to a purely *spectral* domain. For each partial surface, our method takes as input the truncated sequence of its Laplacian eigenvalues, which act as a surrogate of the shape geometry, and predicts as output the eigenvalue sequence of the 3D model obtained from the

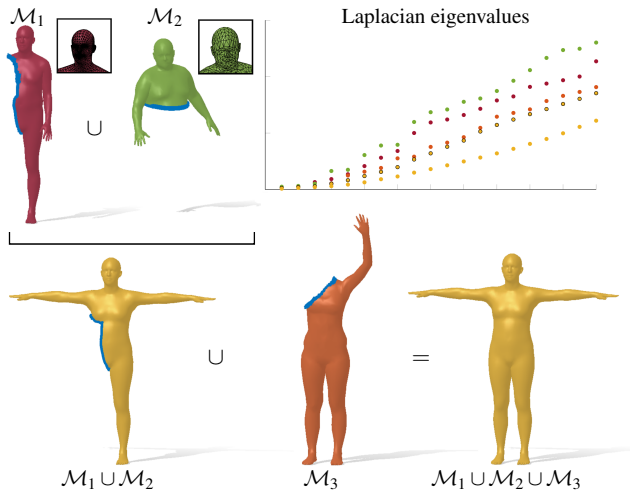


Figure 1: Given a collection of partial deformable shapes $\{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\}$ as input, our method predicts the Laplacian eigenvalues of their union without first having to compute a correspondence or a transformation between the input shapes. The resulting eigenvalues (top right plots, colors correspond to each surface) can be used to reconstruct the final shape if needed, up to isometry/pose (bottom right). In this example, the input shapes have different poses, varying overlap, and different mesh connectivity.

union of the partial surfaces (or an isometric deformation thereof) – but not the 3D model itself, as visualized in Figure 1. This eigenvalues prediction task is in general ill-posed, but can be resolved by means of a data prior, namely by training a deep net on a few hundred examples. The advantages of this approach are numerous, and include associativity (i.e. $A \cup B \cup C = (A \cup B) \cup C$), invariance to deformations and sampling, and generalization to different discrete representations for the input geometry. In a way, this recalls the notion of “homomorphic encryption” in secure computation [RAD79], where the task is to perform calculations on encrypted data without decrypting it first.

Contribution. In this paper, we introduce a learning-based method to estimate the Laplacian spectrum of the union of partial non-rigid 3D shapes, without actually computing the 3D geometry of the union. Sidestepping the reconstruction means that we do not have to commit to one specific 3D embedding in the output (e.g. a specific pose for a human body), but leave this choice to task-specific blocks. Moreover, our method takes advantage of the geometric insight that the spectra can be used not only for single shape recovery and processing (as done in prior works) but also to enable multi-shape operations such as unions. Once a spectrum is predicted, it can be fed *as-is* to any existing spectral pipeline operating with eigenvalues. For example, we can reconstruct the full 3D geometry by using the method in [MRC*19] as an output module. If the geometry is not needed, e.g., for tasks of shape retrieval [RWP06] and region localization [RTO*19], we achieve the same accuracy that can be obtained in the case where the full shape is given.

2. Related work

We discuss two lines of research that are most closely related to our spectral aggregation task: partial non-rigid aggregation of shapes in their extrinsic form (e.g. mesh or point cloud), and spectral analysis of partial shapes.

2.1. Nonrigid shape aggregation

Recovering deformable 3D shapes from partial scans has numerous applications in AR/VR, manufacturing, and robot manipulation. A common setting for this problem is non-rigid registration, where the scans are captured sequentially and exhibit mild inter-frame deformations, and significant overlap. In such cases, template-less methods have been shown to perform well by using general deformation models such as thin-plate splines [BR04; BR07] or as-rigid-as-possible energies [LSP08]. Wand et al. [WJH*07; WAO*09] used dynamic “surfels” to represent the input surfaces, and proposed a statistical model to recover the underlying template shape. Temporal coherence has been used in [TBW*12] to generate dense correspondences from robust landmarks, and in [MFO*07] to reconstruct a space-time surface embedded in 4D. Sharf et al. [SAL*08] incorporated a mass conservation prior to control the plausibility of the reconstructed surface. The “Dynamic Fusion” method of Newcomb et al. [NFS15] and follow-up work [SBC17], demonstrated real-time, template-free non-rigid reconstruction allowing both the object and the camera to move.

More related to our setting are cases where the input set is sparser, and the deformations between the scans can change significantly. In fact, we do *not* assume temporal coherence or an initial alignment. Similar to us, methods designed for these settings usually assume a strong prior on the shape category or even rely on a parametric model. In [ZFY14], a generic human template was used for building a personalized parametric human body model similar to SCAPE [ASK*05]. Chang and Zwicker [CZ09; CZ11] assumed an articulated model and solved for joints and skinning weights. More recently, advances in geometric deep learning for processing point clouds and meshes were used to leverage data-driven priors (e.g. from [BRPB17]) for deformable shape completion and fusion [LBBM18; HIL*20].

2.2. Eigenvalues and partiality

Spectral representations based on the Laplacian are widely used in the analysis of deformable shapes, mainly due to their isometric invariance. Much less attention has been given to the effect of partiality on the spectrum.

Shape correspondence. A first attempt at utilizing the Laplacian eigenfunctions to recover dense correspondences between a partial and a full shape was shown in [RCB*17], building upon the seminal functional maps framework [OBS*12]. This was further extended to matching shapes in the presence of clutter [CRM*16], and to a more efficient fully spectral variant in [LRBB17]. In the context of partial shape aggregation, most relevant is an extension to the multi-part matching algorithm (a.k.a. “non-rigid puzzle”) proposed in [LRB*16]. Recently, deep learning techniques have also been utilizing Laplacian eigenfunctions for matching [LRR*17;

HLR*19; GR20; ALC20; RSO19; APO21]. Replacing eigenfunctions with a basis learned from data was recently proven more robust and therefore applicable to challenging settings including point clouds and partiality [MRMO20].

Reconstruction. Aside from matching, other works have investigated spectral methods for non-rigid completion and registration. In FARM [MMRC20] and its high-resolution variant [MMRC19], a functional maps representation is incorporated into a parametric model-based regression pipeline.

Shape from Spectrum. Most closely related to ours are works that aim directly to recover the shape from its underlying spectrum, also known as the problem of “hearing the shape of a drum” [Kac66]. This procedure was recently studied by Cosmo et al. [CPR*19] in practical rather than purely theoretical settings. Their pipeline, dubbed “isospectralization”, was proven useful in multiple application scenarios, and extended in [MRC*19] by replacing the regularizers of [CPR*19] with a data-driven prior.

In this work, we aim to perform the union of partial deformable shapes from their spectral representation. Our method differs from the ones listed above for two main reasons: i) we consider the shape from spectrum problem in the more challenging setting of partiality, and ii) rather than recovering the geometry of the full shape, we aim to recover its Laplacian spectrum, given the spectrum of two parts. In other words, we introduce the problem of spectral unions of partial shapes and propose an effective solution. In this light our work is related to [SDKG18], which devised a framework to learn fuzzy representations that enable set operations on man-made objects.

3. Proposed method

Let us be given two partial shapes \mathcal{M}_1 and \mathcal{M}_2 , and let $\mathcal{M}_1 \cup \mathcal{M}_2$ denote their non-rigid alignment, as depicted on the left side of Figure 1. We seek an answer to the following question: what can we say about $\mathcal{M}_1 \cup \mathcal{M}_2$, without actually computing this union? More specifically: can we predict the spectrum of $\mathcal{M}_1 \cup \mathcal{M}_2$ without having to solve for the point-to-point correspondence between them?

With no additional priors, the question is ill-posed; for example, there are infinitely many ways in which two sheets of paper can be glued together. In the sequel, we claim that the spectrum of the union can be predicted by coupling Laplacian eigenvalues with a data prior without solving a correspondence or reconstruction problem in the process.

Mathematical preliminaries. We model shapes as Riemannian manifolds \mathcal{M} with boundary $\partial\mathcal{M}$. Each manifold identifies an equivalence class of isometries, and thus has infinitely many embeddings in \mathbb{R}^3 (e.g. changes in pose). Let us be given two manifolds \mathcal{M}_1 and \mathcal{M}_2 , together with a diffeomorphism $\pi: \mathcal{R}_1 \rightarrow \mathcal{R}_2$ between regions $\mathcal{R}_1 \subseteq \mathcal{M}_1$ and $\mathcal{R}_2 \subseteq \mathcal{M}_2$. A third manifold $\mathcal{M}_1 \cup \mathcal{M}_2$ can be obtained by attaching \mathcal{M}_1 to \mathcal{M}_2 over the common region via the

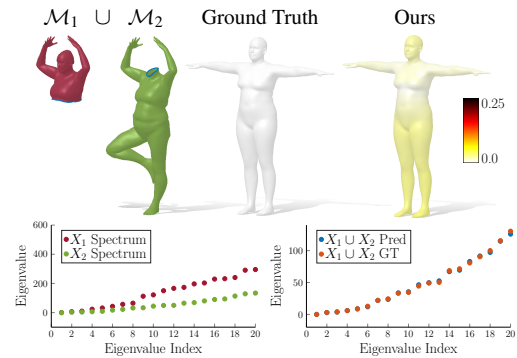


Figure 2: Example of partial shapes whose union entirely covers the full shape. This is the simplest setting that we consider in this paper. Given the spectra of the partial shapes (red and green), we recover the spectrum of their union, and from the spectrum we recover the geometry in standard T-pose using a shape-from-spectrum reconstruction method [MRC*19]. The white shape is recovered from ground-truth eigenvalues; ours is colored with a heatmap, which encodes reconstruction error.

map π (as depicted in the inset figure). We refer to $\mathcal{M}_1 \cup \mathcal{M}_2$ as the *union shape*[†].

On each \mathcal{M} we consider the Laplace-Beltrami operator Δ , extending the notion of Laplace operator from Euclidean geometry to surfaces. This operator admits a spectral decomposition:

$$\Delta\phi_i(x) = \lambda_i\phi_i(x) \quad x \in \text{int}(\mathcal{M}) \quad (1)$$

$$\phi_i(x) = 0 \quad x \in \partial\mathcal{M} \quad (2)$$

into eigenvalues $\lambda_1 < \lambda_2 < \lambda_3 < \dots$ and associated eigenfunctions $\phi_1, \phi_2, \phi_3, \dots$; we adopt homogeneous Dirichlet boundary conditions (2). The set of eigenvalues forms a discrete *spectrum*, which we assume to be ordered non-decreasingly. In this paper, we consider truncated spectra of length k , and introduce the vector-valued function:

$$\boldsymbol{\lambda}: \mathcal{M} \mapsto (\lambda_1, \dots, \lambda_k). \quad (3)$$

In particular, we completely discard the eigenfunctions $\phi_1(x), \phi_2(x), \dots$, which are point-based quantities and thus highly dependent on shape discretization.

Remark. Since the Laplacian Δ is invariant to isometries, so is its truncated spectrum encoded in $\boldsymbol{\lambda}$. This means that eigenvalues capture shape information *up to pose*, a fundamental property that is at the basis of our method.

Problem statement. In non-rigid alignment, one is given 3D embeddings (e.g. point clouds) for \mathcal{M}_1 and \mathcal{M}_2 , and must recover a 3D embedding of their union $\mathcal{M}_1 \cup \mathcal{M}_2$. Since, in this setting, \mathcal{M}_1 and \mathcal{M}_2 may undergo wildly different deformations, there is

[†] We keep the mathematical description simple for the sake of clarity. Formally, this operation is called *connected sum*, denoted by $\mathcal{M}_1 \# \mathcal{M}_2$, and is part of the surgery theory of manifolds, see, e.g., [Rol76].

no guarantee that they have the same 3D coordinates on the common region. Therefore, it is not clear how a 3D embedding for $\mathcal{M}_1 \cup \mathcal{M}_2$ should look like.

In our work, we propose to mitigate this problem by switching from a discrete representation of the 3D embedding of $\mathcal{M}_1 \cup \mathcal{M}_2$ to a discrete representation of the entire isometry class, given by $\lambda(\mathcal{M}_1 \cup \mathcal{M}_2)$. Then, we translate the problem of recovering an alignment between 3D embeddings to the estimation of a parametric nonlinear operator $\mathcal{U}_\Theta : \mathbb{R}^k \times \mathbb{R}^k \rightarrow \mathbb{R}^k$, such that:

$$\lambda(\mathcal{M}_1 \cup \mathcal{M}_2) = \mathcal{U}_\Theta(\lambda(\mathcal{M}_1), \lambda(\mathcal{M}_2)). \quad (4)$$

We call \mathcal{U}_Θ the *spectral union* operator, and model it as a deep neural network with learnable parameters Θ . A specific definition for the architecture and the loss are given in Section 4.

Remark. In general, the spectrum of the union shape $\mathcal{M}_1 \cup \mathcal{M}_2$ is *not* simply the union of the spectra of $\mathcal{M}_1, \mathcal{M}_2$. This is only true if \mathcal{M}_1 and \mathcal{M}_2 correspond to disjoint regions of the complete shape (see e.g. [RCB*17, Sec. 3.1]), while in this paper we consider the case in which \mathcal{M}_1 and \mathcal{M}_2 partially overlap, thus making the interactions between the two spectra much more complex.

Difficulty settings and associativity.

Estimating an operator \mathcal{U}_Θ that makes Eq. (4) hold for many different pairs $(\mathcal{M}_1, \mathcal{M}_2)$ is not a simple problem, even with short sequences (in this paper we use $k = 20$). In fact, it is known that Laplacian spectra can vary wildly under partiality perturbations [FM09], and predicting these variations can be difficult.

Based on these observations, we consider two different scenarios with different characteristics:

1. $\mathcal{M}_1 \cup \mathcal{M}_2$ is a complete, watertight shape;
2. $\mathcal{M}_1 \cup \mathcal{M}_2$ is a partial shape itself.

As we demonstrate below, Scenario 1 is simple enough to be solved with a feed-forward network, and generalizes well to unseen data, as shown in Figure 2. Scenario 2 is more difficult, since allowing partiality on the union shape introduces another dimension of variability, as well as more ambiguity on the possible output; see Figure 3 for examples.

Despite being more difficult to solve, the latter scenario lends itself to modeling more complex interactions. In particular, exploiting the associative property of the union, we can compose $m > 2$ partial shapes simply by aggregating pairwise unions:

$$\lambda(\mathcal{M}_1 \cup \mathcal{M}_2 \cup \dots \cup \mathcal{M}_m) = \mathcal{U}_\Theta(\dots(\mathcal{U}_\Theta(\lambda(\mathcal{M}_1), \lambda(\mathcal{M}_2)), \dots), \lambda(\mathcal{M}_m)) \quad (5)$$

See Figure 1 and the inset on the right for an illustration. Note that composing m partial shapes resembles the ‘non-rigid puzzle’ setting seen in [LRB*16], although with a crucial difference: the method of [LRB*16] has access to the complete shape, which is instead unknown to us.

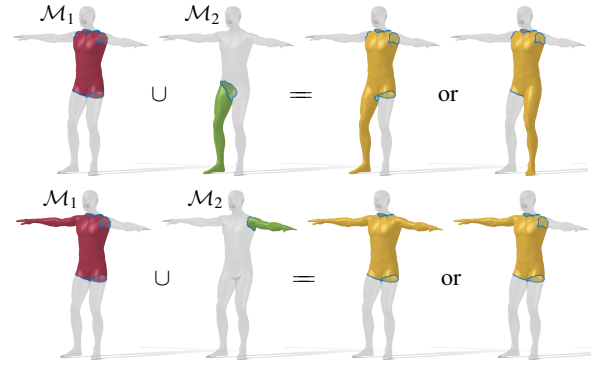
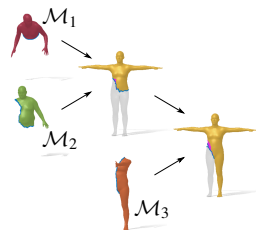


Figure 3: Spectra capture isometry classes, thus there exist ambiguous cases where unions have multiple valid solutions. Top: The two solutions are isometric, hence intrinsically equivalent. Bottom: Since each part is isospectral to its symmetric version, the union of the two spectra can result in three possible solutions (we only show two for simplicity). The semi-transparent full shape is for reference.

4. Network architecture

Our network takes as input two sequences of k eigenvalues, each associated with a partial shape, and outputs a sequence of k eigenvalues, as a prediction of the spectral union. Figure 4 illustrates the neural architecture. It is composed of three main blocks: (1) the projection of the input eigenvalues into a high dimensional space; (2) two transformers, forced to be commutative, to learn the union operation; (3) a dimensionality reduction to decode the spectral union.

Eigenvalue embeddings. The Laplacian eigenvalues of surfaces form a non-decreasing sequence that approximately grows linearly with rate inversely proportional to surface area, a behavior described by Weyl’s asymptotic law [Wey11]. This results in the input eigenvalues hugely varying depending on the area of the partial shape. To guard against network instability we encode the spectra via the offset representation:

$$\text{off}(\lambda_i) = \lambda_i - \lambda_{i-1},$$

with $\text{off}(\lambda_1) = \lambda_1$. This representation has the further advantage of imposing the increasing order constraint on the predicted eigenvalues, by just requiring the non-negativity of the predicted offset sequence.

In practice, the network sees each spectrum as a sequence of length k offsets $\mathbf{A} = (\text{off}(\lambda_1), \dots, \text{off}(\lambda_k)) \in \mathbb{R}^k$, each one is then embedded into a higher-dimensional representation of length $2\ell + 1$, constructed as follows:

$$\text{off}(\lambda_i) \mapsto \left(\vec{\theta}_a^i, \text{off}(\lambda_i) \vec{\theta}_b, \text{off}(\lambda_i) \right)$$

where $\vec{\theta}_a^i$ is a ℓ -dimensional vector acting as a positional encoding for the i -th offset, and $\vec{\theta}_b$ is a linear mapping of the offset to a ℓ -dimensional space. The learnable vectors $\vec{\theta}_a^i$ and $\vec{\theta}_b$, once learned, are independent from the input shapes and eigenvalues.

This representation encodes both the eigenvalue quantity and its

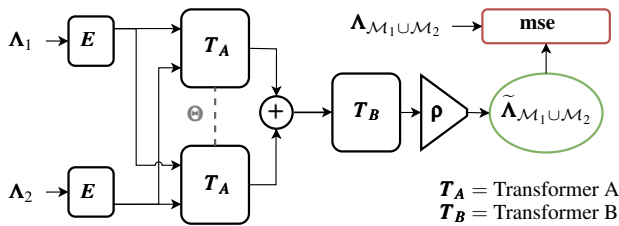


Figure 4: Our neural architecture. Λ_1 and Λ_2 are the input eigenvalues of the partialities, E embeds the eigenvalues into a high dimensional space, the transformer T_A produces a latent representation of the inputs that are summed up to obtain a commutative latent representation of the union, the transformer T_B plus the linear dimensionality reduction ρ decodes this latent representation to obtain the predicted eigenvalues $\hat{\Lambda}_{\mathcal{M}_1 \cup \mathcal{M}_2}$.

position in the sequence, which is a fundamental information for recovering the geometry area.

Symmetric architecture. Given the eigenvalue sequences of Λ_1 and Λ_2 (associated to \mathcal{M}_1 and \mathcal{M}_2 respectively), our neural architecture learns how to perform their union without ever leaving the spectral domain. We further require our model to be commutative, i.e., the result should *not* depend on which pair between (Λ_1, Λ_2) or (Λ_2, Λ_1) is given as input.

We gain this invariance by using a single transformer T_A on the embedded eigenvalues, performing two symmetric operations to obtain a representation of Λ_1 informed about Λ_2 and vice-versa. The two transformed representations are summed together to obtain a commutative representation of the union. We then feed the result to the second transformer T_B , whose task is to decode the union into a representation that can be easily reduced, via a simple linear layer ρ , from the high-dimensional representation back to a sequence of eigenvalues. The whole architecture is illustrated in Figure 4.

The transformers are position-aware neural networks, where the output for each eigenvalue depends on its value and position together with all the other eigenvalues and their positions. It employs an *attention* mechanism to learn relation among eigenvalues.

In the network, the dimensionality of each representation is 32, T_A has 8 heads and 6 layers meanwhile T_B has 8 heads and 3 layers. Thus, ρ reduces the representation dimensionality from 32 to 1. Refer to the supplementary materials for further details.

Training. Our model is trained with a mean squared error loss between the predicted and ground truth spectra. Before entering the loss, the offset representation for the eigenvalues is decoded with a cumulative sum. Experimental results show that penalizing the loss according to the linear increase of the eigenvalues does not yield significant improvements. In the training phase, we augmented the partial regions with small random changes in their surface area. The optimizer used is Adam with a learning rate of $2e-4$ and weight decay of $1e-5$. We use a learning rate scheduler to escape local minima and stabilize the training, in particular the cosine annealing with warm restarts scheduler [LH17], doubling at each restart the number of epochs between restarts. We trained the model for

6741 epochs for a total of 1d 13h 46m on a GeForce RTX 2080 TI, tracking the experiments with [Bie20].

5. Data and evaluation.

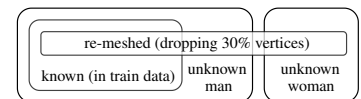
In our experiments we use 3D data from the FAUST [BRLB14] and SURREAL [VRM*17] datasets of deformable human shapes with different identities. This provides us with a total of 50 different identities, each in 10 different poses. To produce partial data, we first extract surface patches of various sizes from the full shapes, and then combine the patches randomly to form two datasets:

- A dataset of ~ 150 partial pairs, where each union covers the entire surface. We test in three different settings depending on the information given at training time: (i) known identity, unknown partiality; (ii) unknown identity, known partiality; (iii) both identity and partiality are unknown. We define an identity as *known* if the training set contains any partiality in any pose of the same shape, and we consider a partiality *known* if the two input partiality types together with their corresponding union are in the training dataset in any shape identity or pose.
- A dataset of ~ 100 partial pairs, whose union does not cover the entire surface. For training, the partial shapes are augmented by enlarging/shrinking the patches randomly. We consider the same three settings as above.

As shown in Figure 3, there are cases in which more than one region on the template is a valid solution to the union problem. Two different ambiguities arise: (a) symmetric counterparts of one or both input partial regions may produce different union regions with different spectra; (b) symmetric union regions are described by the same spectra even though they are localized in different parts of the shape. We remove these ambiguities in the training data by following a minimum union area principle and privileging “left-sided” symmetries exploiting a ground truth symmetry map and labels of the template left side. By this choice, associativity is promoted as we show empirically in the results.

We define two test sets. In TEST A, the pose or the type of partiality have never been seen, but the predicted union may be seen in a different pose or identity at training time. TEST B is more challenging, since the union of the two parts has never been seen at training (neither in a different pose or identity, nor as a union of different partialities). The number of samples in the test datasets is about 15% of all the data available, the remaining data is used for training.

We analyze both TEST A and TEST B scenarios in different settings summarized in the inset Venn diagram.



In Table 1 we report a quantitative analysis of the predictive power of our learning model, according to the mean squared error (mse) and mean absolute error (mae) metrics.

Further, we perform qualitative experiments on different classes, on horses from TOSCA [BBK08] and earphones from PartNet [MZC*19], some examples are in Figure 7, 8 and 9. Through these experiments we prove our method generalization ability to

		mse	mae
TEST A	known man	11.14	2.09
	unknown man	13.25	2.59
	unknown woman	36.92	3.93
	known man re-meshed	29.61	3.31
	unknown man re-meshed	32.67	3.60
	unknown woman re-meshed	62.33	5.23
TEST B	known man	15.41	2.59
	unknown man	24.05	3.60
	unknown woman	64.47	4.99
	known man re-meshed	51.20	4.54
	unknown man re-meshed	75.91	5.90
	unknown woman re-meshed	110.17	6.78

Table 1: Error between the predicted and ground truth eigenvalues in different experimental settings. In each row, “known” denotes an identity included in the training set, “unknown” one not included, and “re-meshed” indicates that the shapes were re-meshed by removing 30% of their vertices before computing their spectrum.

any shape category. Additionally, we successfully trained the neural network on humans and fine-tuned it to work with horses, where the data is scarce, demonstrating that it is possible to perform transfer learning between different shape classes. Details about the data generation and training process for these classes is described in the supplementary material. Sample code and data are available online[†].

6. Applications

We can easily plug our method into existing pipelines that take as input Laplacian eigenvalues. Unique to our approach is that it addresses the scenario in which only partial views of the complete shape are available. We also refer to the Supplementary for further details and results.

6.1. Geometry reconstruction

To recover the shape geometry from its predicted union eigenvalues, we use the data-driven method of [MRC*19], which takes eigenvalues as input and directly yields a 3D mesh embedding as output. An example is given in Figure 2, where we compare the geometry recovered from our estimated spectra with the one obtained from the ground truth eigenvalues. For human meshes where the correspondence between their T-Pose and the connectivity adopted in [MRC*19] is known, we can compute the point-wise reconstruction error as the L2 distances between correspondent points. We plot this Euclidean error on the reconstructed surface. White color corresponds to zero error and dark red encodes a larger error. Our spectrum prediction is accurate enough to retain the core geometric information of the original partial shapes, as it can be seen in these examples. For these experiments we used the pre-trained network provided by the authors of [MRC*19] and [MRC*21]. We sampled

[†] Link to the repository <https://github.com/lucmos/spectral-unions>

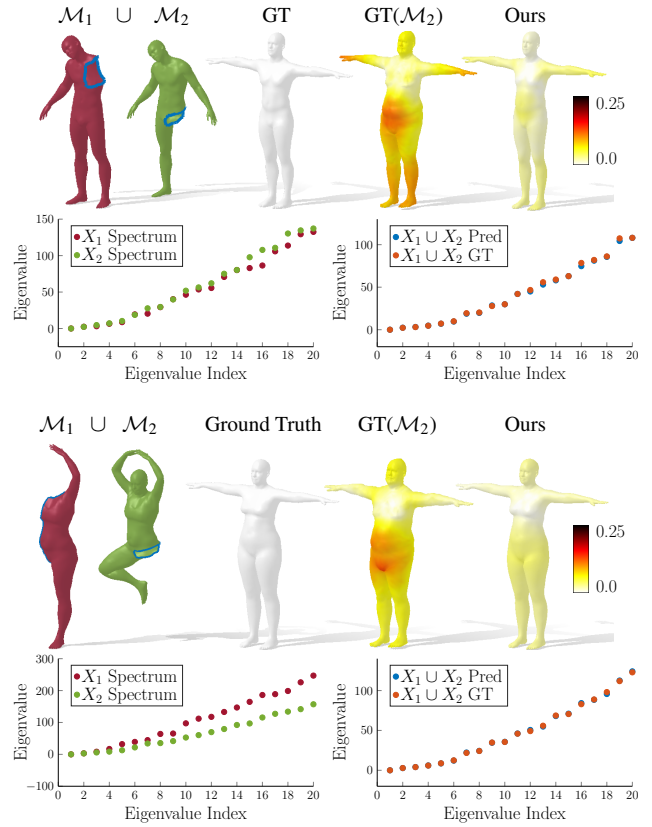


Figure 5: Given two partial shapes as input, we compare the reconstruction obtained by running the method of [MRC*19] only on a partial input (the green shape), yielding the fourth shape, with the reconstruction obtained from our predicted full spectrum, yielding the last shape.

the test shapes outside the training set adopted in these papers. Thus the network is not specifically trained to handle spectra predicted by our pipeline. Moreover, since the spectrum encodes just intrinsic properties (i.e. appearance) of the shape, all the reconstructions of [MRC*19] are in the T-pose.

To emphasize the importance of having an aggregated spectrum, as predicted by our model, in Figure 5 we show the reconstructions obtained with the method of [MRC*19] when using the spectrum of just one of the two partial shapes as input. The result in this case is quite different from what is expected, showing that existing state-of-the-art pipelines are not able to handle partial shapes correctly.

6.2. Region localization

This task, introduced in [RTO*19], consists in locating, on a fixed template, the region corresponding to a given partial shape. To solve this problem we combine the spectral union model introduced in Section 4 with a simple MLP, described in detail in the supplementary materials. The MLP takes as input the predicted eigenvalues of the union, and outputs an indicator function over the vertices

		IoU	Acc.
TEST A	known man	99.28%	99.61%
	unknown man	93.78%	95.83%
	unknown woman	94.19%	96.32%
	known man re-meshed	98.54%	99.06%
	unknown man re-meshed	91.44%	94.08%
	unknown woman re-meshed	93.47%	95.56%
TEST B	known man	97.96%	98.55%
	unknown man	87.58%	92.52%
	unknown woman	96.05%	98.46%
	known man re-meshed	93.04%	97.33%
	unknown man re-meshed	83.69%	91.08%
	unknown woman re-meshed	95.59%	98.43%

Table 2: Intersection over union (IoU) and accuracy in the region localization task, in different experimental settings. Model trained on a single identity, to show generalization.

		IoU	Acc.
TEST A	known man	98.24%	99.09%
	unknown man	96.26%	97.64%
	unknown woman	96.17%	98.04%
	known man re-meshed	97.70%	98.74%
	unknown man re-meshed	95.88%	97.78%
	unknown woman re-meshed	96.04%	97.66%
TEST B	known man	97.43%	99.14%
	unknown man	93.31%	98.23%
	unknown woman	95.74%	98.59%
	known man re-meshed	97.61%	99.11%
	unknown man re-meshed	90.85%	97.63%
	unknown woman re-meshed	96.81%	98.98%

Table 3: Performance when training on six different identities instead of a single identity (compare with Table 2).

of the template. The spectral union operator is not trained to solve the region localization but is used as-is with frozen weights.

In principle, substituting T_B in Figure 4 with the region localization MLP would work if the whole system is trained end-to-end. However, the goal of this work is to perform the union operation in the spectral space. Moreover, if we do not impose the union to be a spectrum, we would not be able to compose the predicted union spectra with another partiality.

In the loss definition, one must take care of the potential ambiguities exemplified in Figure 3; we do so by implementing a symmetry-invariant loss, that does not penalize symmetric solutions. The MLP is trained using the train/test splits described in Section 4, with the difference that we used just 6 different identities in the training phase.

To analyze the prediction quality on this task we adopt two metrics: intersection over union (IoU) of the predicted mask with the ground truth mask, and accuracy, i.e. the ratio of correctly predicted vertices over the full template. We show several qualitative results

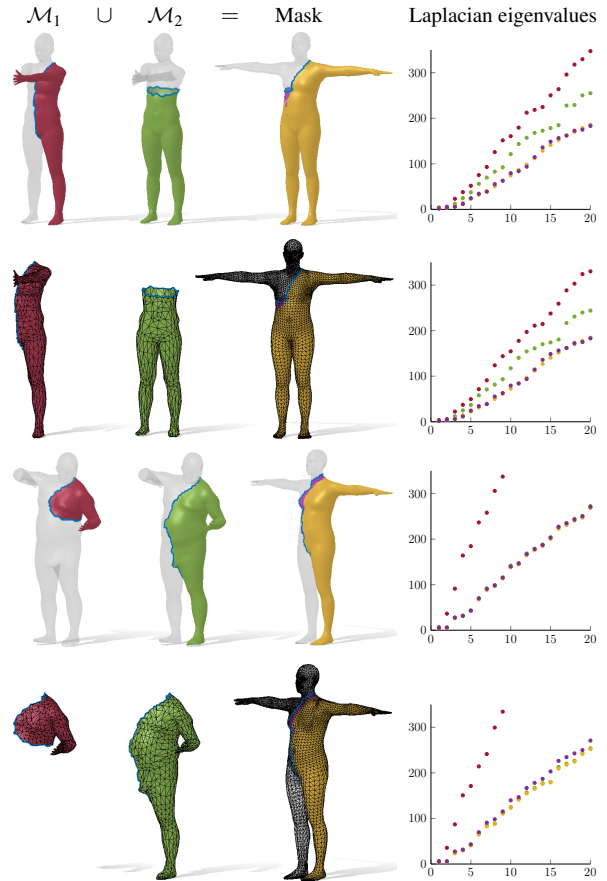


Figure 6: Region localization task, under the effect of different mesh connectivity. Given the eigenvalues of two partial shapes, we correctly predict an indicator function that represents the union of the two over a fixed template.

in Figure 6 and attach an interactive demo in the supplementary materials.

Robustness to re-meshing. One key aspect of Laplacian eigenvalues is that they are robust to shape discretization and mesh connectivity. Our model inherits this robustness; see Figure 6, where we highlight the re-meshed inputs by visualizing their surface triangulation. This is supported also by Tables 2 and 3, where the performance on the re-meshed shapes is comparable with the original ones. In these experiments, we test our network with the eigenvalues computed from noisy, re-meshed partial shapes obtained by removing 30% of their vertices with an edge collapse algorithm [GH97].

Generalization to new identities. Our approach generalizes to identities unseen at training time as can be noted in Table 3. To further stress this aspect, we devised an experimental setup in which we used as training set just a single identity. The results of this setup are shown in Table 2.

Generalization to different datasets. In Figure 7 we use partial




Figure 7: Region localization across different datasets. Partial shapes come from datasets not involved in the training.

shapes from other datasets to localize regions on the fixed template. These shapes have different triangulation, vertex density and style, confirming generalization across datasets. More specifically: a shape from TOSCA [BBK08] for humans (first row), one from SMAL [ZKJB17] for the horses (second row) and a camel shape that has a different triangulation and comes from a different class (third row).

Generalization to point clouds. We obtain good results also on point clouds, as shown in Figure 8. For earphones, in the top row, we perform both training and testing on point clouds. In the bottom row, we show that our model trained on human meshes generalizes to point clouds. We compute the Laplacian for point clouds with the method of [SC20].

Associativity. We can compute spectral unions of > 2 partial shapes iteratively as described in Section 3. In Figure 9 we show qualitative results over three parts.

Interpolation. Finally, in Figure 10 we first interpolate the spectra of two partial shapes (in green), and then compute the union of the interpolated spectra with the spectrum of a fixed shape (in red). From each of these unions, we predict a mask on the given template (in yellow). We can see how in the first example (top row) the mask changes smoothly. On the other hand, in the second example it is less obvious how to interpolate the completely missing leg, resulting in an abrupt discontinuity in the predicted mask.

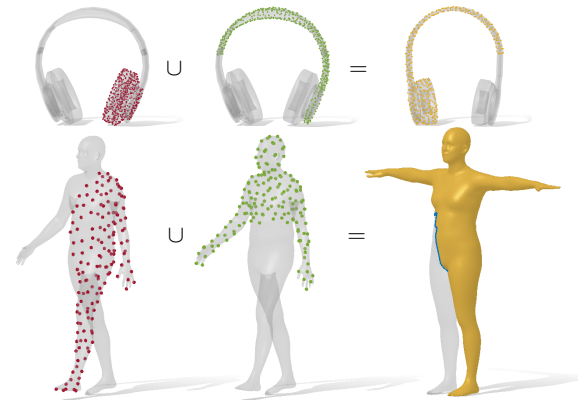


Figure 8: Region localization from partial point cloud spectra. The white mesh is just shown as a visual reference.

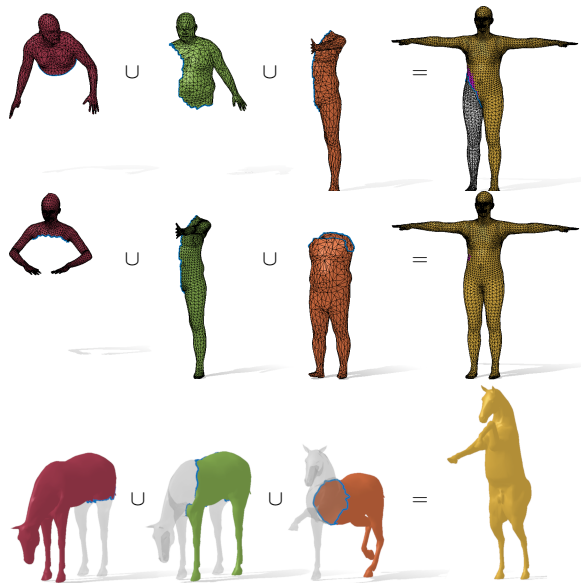


Figure 9: Example of associativity. Note that all the human meshes involved have different connectivity.

6.3. Shape retrieval

This task consists in retrieving a query from a database of shapes that could undergo several deformations. A well-known spectral method to tackle this problem, ShapeDNA [RWP06], adopts the Laplacian spectrum as a shape signature. In the space of these signatures, nearest-neighbor search yields the desired result. However, in order to work correctly, ShapeDNA needs the spectrum of a complete shape; extensions of this signature to the partial case have

	top-1	top-5	top-10
Ours	86.14%	97.75%	99.20%
ShapeDNA	86.59%	96.81%	97.72%

Table 4: Comparisons on the shape retrieval task.

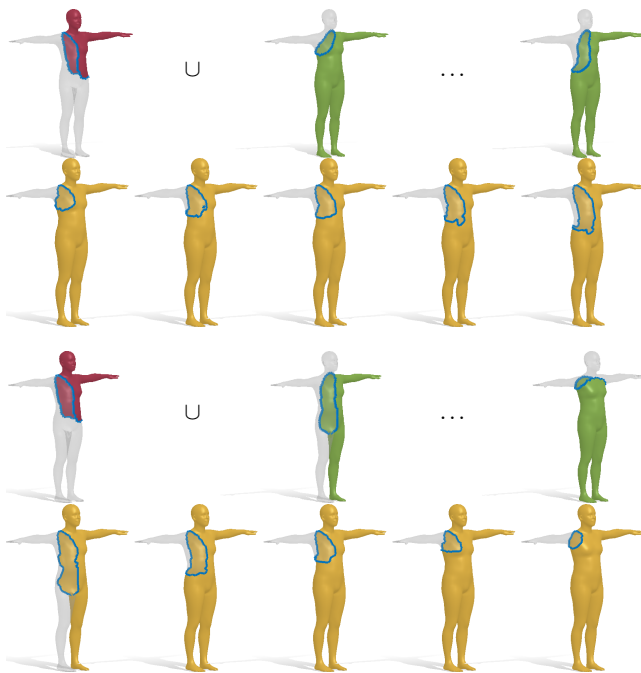


Figure 10: Two examples of linear interpolation of eigenvalues (green shapes), and the resulting predicted masks (in yellow). Please refer to the main text for details.

proven unsuccessful to date [RCL*]. Our method applies directly to this case, since we can estimate the eigenvalues of the unknown complete shape whenever the input query is just a collection of its partial views.

We run our tests on a dataset of 440 complete shapes (44 identities in 10 poses each). For our method, we evaluate 4400 pairs of partial shapes; for each pair we predict the ShapeDNA signature of their union and use it to query the database. We compare it with the accuracy obtained by standard ShapeDNA on each of the 440 complete shapes in the database, which assumes exact knowledge of the union spectra; nevertheless, the identity it retrieves may be wrong due to spectra variations caused by deformations between different poses. We measure the performance using top- k metrics, which count the number of times a shape with the correct identity is in the first k retrieved shapes; we use $k = 1, 5, 10$. The results are reported in Table 4, and show that our predicted eigenvalues are accurate enough to compete with, and even surpass, ShapeDNA for this task. The better performance is due to the robustness of our method to the noise induced by the pose change.

7. Conclusion

We introduced a method to recover the aggregated Laplacian spectrum of a collection of partial deformable shapes, while avoiding the computational burden of computing correspondences or extrinsic alignments. Our method involves a deep net that, given two eigenvalue sequences as input, simply produces another eigenvalue sequence as output. In spite of its apparent simplicity, this method allows to address a number of applications that traditionally require

solving for a correspondence, and retains a comparable quality (in some cases, even higher) to methods that have direct access to the 3D geometry of the full shape.

Limitations and future directions. Perhaps the main limitation of our method lies in the missing mathematical guarantee that our predicted sequences are actual Laplacian eigenvalues, despite our positive empirical results. We consider enforcing this constraint as an interesting direction of further research. Another interesting area for improvement is the region localization generalization capability, where our current model seems to struggle with out-of-distribution union partialities. We are optimistic that a more diverse and extensive training set would boost the generalization performance. Moreover, we did not consider unprocessed partial single-view or depth scans of physical objects. We expect a drop in performance on such data comparable to other spectral methods. We consider improving the robustness of spectral methods on natural non-pre-processed data as an essential and challenging research direction.

Acknowledgements Parts of this work were supported by the ERC Starting Grants No. 802554 (SPECGEO) and No. 758800 (EX-PROTEA), the ANR AI Chair AIGRETTE, and the SAPIENZA BE-FOR-ERC 2020 Grant (NONLINFMAPS).

References

- [ALC20] AYGÜN, MEHMET, LÄHNER, ZORAH, and CREMERS, DANIEL. “Unsupervised Dense Shape Correspondence using Heat Kernels”. *arXiv preprint arXiv:2010.12682* (2020) 3.
- [APO21] ATTAIKI, SOUHAIB, PAI, GAUTAM, and OVSJANIKOV, MAKS. “Dpfm: Deep partial functional maps”. *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, 175–185 3.
- [ASK*05] ANGUELOV, DRAGOMIR, SRINIVASAN, PRAVEEN, KOLLER, DAPHNE, et al. “SCAPE: shape completion and animation of people”. *ACM SIGGRAPH 2005 Papers*. 2005, 408–416 2.
- [BBGO11] BRONSTEIN, ALEXANDER M, BRONSTEIN, MICHAEL M, GUIBAS, LEONIDAS J, and OVSJANIKOV, MAKS. “Shape google: Geometric words and expressions for invariant shape retrieval”. *ACM Transactions on Graphics (TOG)* 30.1 (2011), 1–20 1.
- [BBK08] BRONSTEIN, A. M., BRONSTEIN, M. M., and KIMMEL, R. *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008 5, 8.
- [Bie20] BIEWALD, LUKAS. *Experiment Tracking with Weights and Biases*. Software available from wandb.com, 2020. URL: <https://www.wandb.com/5>.
- [BR04] BROWN, BENEDICT J and RUSINKIEWICZ, SZYMON. “Non-rigid range-scan alignment using thin-plate splines”. *Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004*. IEEE, 2004, 759–765 2.
- [BR07] BROWN, BENEDICT J and RUSINKIEWICZ, SZYMON. “Global non-rigid alignment of 3-D scans”. *ACM SIGGRAPH 2007 papers*. 2007, 21–es 2.
- [BRLB14] BOGO, FEDERICA, ROMERO, JAVIER, LOPER, MATTHEW, and BLACK, MICHAEL J. “FAUST: Dataset and evaluation for 3D mesh registration”. *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. Piscataway, NJ, USA: IEEE, June 2014 5.
- [BRPB17] BOGO, FEDERICA, ROMERO, JAVIER, PONS-MOLL, GERARD, and BLACK, MICHAEL J. “Dynamic FAUST: Registering Human Bodies in Motion”. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. July 2017 2.

- [CPR*19] COSMO, LUCA, PANINE, MIKHAIL, RAMPINI, ARIANNA, et al. "Isospectralization, or how to hear shape, style, and correspondence". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2019, 7529–7538 **1, 3**.
- [CRM*16] COSMO, LUCA, RODOLÀ, EMANUELE, MASCI, JONATHAN, et al. "Matching deformable objects in clutter". *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE. 2016, 1–10 **2**.
- [CZ09] CHANG, WILL and ZWICKER, MATTHIAS. "Range scan registration using reduced deformable models". *Computer Graphics Forum*. Vol. 28. 2. Wiley Online Library. 2009, 447–456 **2**.
- [CZ11] CHANG, WILL and ZWICKER, MATTHIAS. "Global registration of dynamic range scans for articulated model reconstruction". *ACM Transactions on Graphics (TOG)* 30.3 (2011), 1–15 **2**.
- [FM09] FILOCHE, MARCEL and MAYBORODA, SVITLANA. "Strong localization induced by one clamped point in thin plate vibrations". *Physical review letters* 103.25 (2009), 254301 **4**.
- [GH97] GARLAND, MICHAEL and HECKBERT, PAUL S. "Surface simplification using quadric error metrics". *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 1997, 209–216 **7**.
- [GR20] GINZBURG, DVIR and RAVIV, DAN. "Cyclic functional mapping: Self-supervised correspondence between non-isometric deformable shapes". *European Conference on Computer Vision*. Springer. 2020, 36–52 **3**.
- [HIL*20] HALIMI, OSHRI, IMANUEL, IDO, LITANY, OR, et al. *The Whole Is Greater Than the Sum of Its Nonrigid Parts*. 2020. arXiv: 2001.09650 [cs.CV] **2**.
- [HLR*19] HALIMI, OSHRI, LITANY, OR, RODOLÀ, EMANUELE, et al. "Unsupervised learning of dense shape correspondence". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2019, 4370–4379 **3**.
- [Kac66] KAC, MARK. "Can one hear the shape of a drum?": *The american mathematical monthly* 73.4P2 (1966), 1–23 **3**.
- [KBB*13] KOVNATSKY, ARTIOM, BRONSTEIN, MICHAEL M, BRONSTEIN, ALEXANDER M, et al. "Coupled quasi-harmonic bases". *Computer Graphics Forum*. Vol. 32. 2pt4. Wiley Online Library. 2013, 439–448 **1**.
- [LBBM18] LITANY, OR, BRONSTEIN, ALEX, BRONSTEIN, MICHAEL, and MAKADIA, AMEESH. "Deformable shape completion with graph convolutional autoencoders". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, 1886–1895 **2**.
- [LH17] LOSHCILOV, ILYA and HUTTER, FRANK. *SGDR: Stochastic Gradient Descent with Warm Restarts*. 2017. arXiv: 1608.03983 [cs.LG] **5**.
- [LRB*16] LITANY, OR, RODOLÀ, EMANUELE, BRONSTEIN, ALEXANDER M, et al. "Non-rigid puzzles". *Computer Graphics Forum*. Vol. 35. 5. Wiley Online Library. 2016, 135–143 **2, 4**.
- [LRBB17] LITANY, OR, RODOLÀ, EMANUELE, BRONSTEIN, ALEXANDER M, and BRONSTEIN, MICHAEL M. "Fully spectral partial shape matching". *Computer Graphics Forum*. Vol. 36. 2. Wiley Online Library. 2017, 247–258 **1, 2**.
- [LRR*17] LITANY, OR, REMEZ, TAL, RODOLÀ, EMANUELE, et al. "Deep functional maps: Structured prediction for dense shape correspondence". *Proceedings of the IEEE International Conference on Computer Vision*. 2017, 5659–5667 **2**.
- [LSP08] LI, HAO, SUMNER, ROBERT W, and PAULY, MARK. "Global correspondence optimization for non-rigid registration of depth scans". *Computer graphics forum*. Vol. 27. 5. Wiley Online Library. 2008, 1421–1430 **2**.
- [MFO*07] MITRA, NILOY J, FLÖRY, SIMON, OVSJANIKOV, MAKS, et al. "Dynamic geometry registration". *Symposium on geometry processing*. 2007, 173–182 **2**.
- [MMR*19] MELZI, SIMONE, MARIN, RICCARDO, RODOLÀ, EMANUELE, et al. "Matching humans with different connectivity". (2019) **1**.
- [MMRC19] MARIN, RICCARDO, MELZI, SIMONE, RODOLÀ, EMANUELE, and CASTELLANI, UMBERTO. "High-Resolution Augmentation for Automatic Template-Based Matching of Human Models". *2019 International Conference on 3D Vision (3DV)*. IEEE. 2019, 230–239 **1, 3**.
- [MMRC20] MARIN, RICCARDO, MELZI, SIMONE, RODOLÀ, EMANUELE, and CASTELLANI, UMBERTO. "Farm: Functional automatic registration method for 3d human bodies". *Computer Graphics Forum*. Vol. 39. 1. Wiley Online Library. 2020, 160–173 **3**.
- [MRC*19] MARIN, RICCARDO, RAMPINI, ARIANNA, CASTELLANI, UMBERTO, et al. "Instant recovery of shape from spectrum via latent space connections". *2019 International Conference on 3D Vision (3DV)*. IEEE. 2019, 37–46 **1–3, 6**.
- [MRC*21] MARIN, R, RAMPINI, A, CASTELLANI, U, et al. "Spectral Shape Recovery and Analysis Via Data-driven Connections". *International Journal of Computer Vision* 129.10 (2021), 2745–2760 **6**.
- [MRMO20] MARIN, RICCARDO, RAKOTOSAONA, MARIE-JULIE, MELZI, SIMONE, and OVSJANIKOV, MAKS. "Correspondence learning via linearly-invariant embedding". *Advances in Neural Information Processing Systems* 33 (2020) **3**.
- [MZC*19] MO, KAICHUN, ZHU, SHILIN, CHANG, ANGEL X., et al. "PartNet: A Large-Scale Benchmark for Fine-Grained and Hierarchical Part-Level 3D Object Understanding". *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019 **5**.
- [NFS15] NEWCOMBE, RICHARD A, FOX, DIETER, and SEITZ, STEVEN M. "Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015, 343–352 **2**.
- [OBS*12] OVSJANIKOV, MAKS, BEN-CHEN, MIRELA, SOLOMON, JUSTIN, et al. "Functional maps: a flexible representation of maps between shapes". *ACM Transactions on Graphics (TOG)* 31.4 (2012), 1–11 **1, 2**.
- [RAD79] RIVEST, RONALD L, ADLEMAN, LEN, and DERTOUZOS, MICHAEL L. "ON DATA BANKS AND PRIVACY HOMOMORPHISMS". *FOUNDATIONS OF SECURE COMPUTATIONS* (1979), 171 **2**.
- [RCB*17] RODOLÀ, EMANUELE, COSMO, LUCA, BRONSTEIN, MICHAEL M, et al. "Partial functional correspondence". *Computer Graphics Forum*. Vol. 36. 1. Wiley Online Library. 2017, 222–236 **2, 4**.
- [RCL*] RODOLÀ, E, COSMO, L, LITANY, O, et al. "SHREC'17: Deformable Shape Retrieval with Missing Parts". () **9**.
- [Rol76] ROLFSEN, D. *Knots and Links*. Mathematics lecture series. Publish or Perish, 1976. ISBN: 9780914098164 **3**.
- [RSO19] ROUFOSSE, JEAN-MICHEL, SHARMA, ABHISHEK, and OVSJANIKOV, MAKS. "Unsupervised deep learning for structured shape matching". *Proceedings of the IEEE international conference on computer vision*. 2019, 1617–1627 **3**.
- [RTO*19] RAMPINI, ARIANNA, TALLINI, IRENE, OVSJANIKOV, MAKS, et al. "Correspondence-free region localization for partial shape similarity via hamiltonian spectrum alignment". *2019 International Conference on 3D Vision (3DV)*. IEEE. 2019, 37–46 **1, 2, 6**.
- [RWP06] REUTER, MARTIN, WOLTER, FRANZ-ERICH, and PEINECKE, NIKLAS. "Laplace–Beltrami spectra as 'Shape-DNA' of surfaces and solids". *Computer-Aided Design* 38.4 (2006). Symposium on Solid and Physical Modeling 2005, 342–366 **1, 2, 8**.
- [SAL*08] SHARF, ANDREI, ALCANTARA, DAN A, LEWINER, THOMAS, et al. "Space-time surface reconstruction using incompressible flow". *ACM Transactions on Graphics (TOG)* 27.5 (2008), 1–10 **2**.
- [SBCI17] SLAVCHEVA, MIROSLAVA, BAUST, MAXIMILIAN, CREMERS, DANIEL, and ILIC, SLOBODAN. "Killingfusion: Non-rigid 3d reconstruction without correspondences". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, 1386–1395 **2**.

- [SC20] SHARP, NICHOLAS and CRANE, KEENAN. “A Laplacian for Non-manifold Triangle Meshes”. *Computer Graphics Forum (SGP)* 39.5 (2020) 8.
- [SDKG18] SUNG, MINHYUK, DUBROVINA, ANASTASIA, KIM, VLADIMIR G, and GUIBAS, LEONIDAS. “Learning fuzzy set representations of partial shapes on dual embedding spaces”. *Computer Graphics Forum*. Vol. 37. 5. Wiley Online Library. 2018, 71–81 3.
- [TBW*12] TEVS, ART, BERNER, ALEXANDER, WAND, MICHAEL, et al. “Animation cartography—intrinsic reconstruction of shape and motion”. *ACM Transactions on Graphics (TOG)* 31.2 (2012), 1–15 2.
- [VRM*17] VAROL, GÜL, ROMERO, JAVIER, MARTIN, XAVIER, et al. “Learning from Synthetic Humans”. *CVPR*. 2017 5.
- [WAO*09] WAND, MICHAEL, ADAMS, BART, OVSJANIKOV, MAKSIM, et al. “Efficient reconstruction of nonrigid shape and motion from real-time 3D scanner data”. *ACM Transactions on Graphics (TOG)* 28.2 (2009), 1–15 2.
- [Wey11] WEYL, HERMANN. “Über die asymptotische Verteilung der Eigenwerte”. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1911 (1911), 110–117 4.
- [WJH*07] WAND, MICHAEL, JENKE, PHILIPP, HUANG, QIXING, et al. “Reconstruction of deforming geometry from time-varying point clouds”. *Symposium on Geometry processing*. 2007, 49–58 2.
- [YLL*15] YIN, MENGXIAO, LI, GUIQING, LU, HUINA, et al. “Spectral pose transfer”. *Computer Aided Geometric Design* 35 (2015), 82–94 1.
- [ZFY14] ZHANG, QING, FU, BO, YE, MAO, and YANG, RUIGANG. “Quality dynamic human body modeling using a single low-cost depth camera”. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2014, 676–683 2.
- [ZKJB17] ZUFFI, SILVIA, KANAZAWA, ANGIOO, JACOBS, DAVID, and BLACK, MICHAEL J. “3D Menagerie: Modeling the 3D Shape and Pose of Animals”. *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. July 2017 8.