

# Coherent Point Drift Revisited for Non-rigid Shape Matching and Registration

Aoxiang Fan<sup>1</sup>, Jiayi Ma<sup>1\*</sup>, Xin Tian<sup>1</sup>, Xiaoguang Mei<sup>1</sup>, and Wei Liu<sup>2</sup>

<sup>1</sup> Wuhan University, China    <sup>2</sup> Tencent Data Platform, China

{fanaoxiang, xin.tian}@whu.edu.cn, {jyma2010, meixiaoguang}@gmail.com, wl2223@columbia.edu

## Abstract

In this paper, we explore a new type of extrinsic method to directly align two geometric shapes with point-to-point correspondences in ambient space by recovering a deformation, which allows more continuous and smooth maps to be obtained. Specifically, the classic coherent point drift is revisited and generalizations have been proposed. First, by observing that the deformation model is essentially defined with respect to Euclidean space, we generalize the kernel method to non-Euclidean domains. This generally leads to better results for processing shapes, which are known as two-dimensional manifolds. Second, a generalized probabilistic model is proposed to address the sensibility of coherent point drift method to local optima. Instead of directly optimizing over the objective of coherent point drift, the new model allows to focus on a group of most confident ones, thus improves the robustness of the registration system. Experiments are conducted on multiple public datasets with comparison to state-of-the-art competitors, demonstrating the superiority of our method which is both flexible and efficient to improve the matching accuracy due to our extrinsic alignment objective in ambient space.

## 1. Introduction

Non-rigid shape matching lies at the core of many applications in computer vision and graphics, ranging from statistical shape analysis, information and style transfer to the generation of new shapes among others [68]. Unlike the rigid counterpart, non-rigid shape matching is much more intractable since the matching cannot be exactly modeled by a small number of parameters. As a result, matching extrinsically by determining the deformation in ambient space, *i.e.* 3D Euclidean space, is deemed as a very difficult task and most attention has been given to intrinsic methods.

The group of intrinsic methods is favorable for matching shapes involving complex motions because the intrinsic




Figure 1. A demonstration of our method. The first and second shapes are a matching pair produced by MWP [34], where the result is presented by color transfer (unmatched point in white). Given the imperfect result, our method proceeds by robustly estimating a transformation from the first to the second shape, which results in the third transformed shape. The refined matching result is manifested in the fourth shape (in comparison to the first).

properties or features are invariant to extrinsic changes. In particular, a prominent strategy is to exploit the spectral quantities using eigen-decomposition of the Laplace-Beltrami operator [61], which is in theory invariant to isometric shape deformations. The strategy has proven to be very effective, represented by the framework of functional maps [52] and its numerous follow-up works in recent years [46, 56, 58]. In essence, these methods are to find a well-defined feature space to represent each point in shape, so that Euclidean metrics can be used to establish correspondences or maps. However, some fundamental problems exist in this category. The maps recovered in feature space normally do not have the properties such as continuity or smoothness. Moreover, the transformation of points inevitably means loss of information which inherently prohibits accurate alignment at fine scales, especially in less prominent areas such as flat regions.

In contrast, extrinsic methods have the potential advantages to preserve the desired properties of maps and obtain more accurate alignment, since the objective is based on ambient space. Efforts in this line of work are relatively less compared to the intrinsic methods. In particular, the classic *coherent point drift* (CPD) [51] and its extended variants [4, 31], deformation based on vector field flow [15] and divergence-free field [19] have been proposed in recent years. However, their performances are restricted due to some commonly-known major obstacles. To begin with,

\*Corresponding author

the deformation which requires a high degree of freedom is hard to model in an efficient way. In addition, the optimization can easily get stuck in local optima because of the complexity of deformation.

Our method belongs to the category of extrinsic methods, although intrinsic information may be utilized for optimal performance as input. Our method is both a revisit and a generalization of the classic CPD. The generalization mainly includes two aspects to address the weaknesses of CPD which prevent it from application to general-purpose shape matching. The first aspect is regarding the deformation model. In CPD, a regularized functional framework is utilized, in which the solution is established as a use of the kernel trick. In this paper, we further analyze the direction and generalize the original framework based on Euclidean space to non-Euclidean domains. We show that the kernel method can be extended for shapes with non-Euclidean structures, by connecting the notion of Laplace operator to a non-Euclidean one, *i.e.* Laplace-Beltrami operator. We also show that the new deformation model has a more economic and intuitive alternative expression. The second aspect is regarding the intractability of optimizing the objective of CPD. We observe that this is due to the large search space of the problem. To address this issue, we borrow ideas from the image matching community [26, 44]. In image matching tasks, direct matching of two feature sets is intractable, thus a prevalent approach is to find the most confident correspondences by nearest-neighbor descriptor matching, and geometric models may be found from the set of initial correspondences. To extend the idea to shape matching, we generalize the probabilistic model of CPD to handle such a restrictive set, which also admits a smart and efficient Expectation-Maximization (EM) method. An exemplary demonstration of our method is presented in Fig. 1.

To summarize, the main contributions of this paper include the following three aspects. First, a generalization of the regularized functional framework of CPD to non-Euclidean domains is proposed. Second, a generalization of the probabilistic model of CPD to handle an arbitrary outlier-contaminated initial confident correspondence set is presented. Third, by extensive experimental validation on multiple datasets and experimental analyses on different settings, we demonstrate the effectiveness of our method.

## 2. Related Work

*Metric Space Methods.* In metric spaces, a metric value shall be defined and given for each pair of points, and metric space methods in shape matching attempt to preserve the geodesic distance metric between every pair of points as well as possible. The procedure known as multidimensional scaling [8, 11] embed shapes into an intermediate space, in which the shapes can be seen as rigid and methods like ICP [6] can be used to perform match-

ing. This idea are firstly suggested in [23], followed by a number of works such as [36, 39, 45, 62]. Although this idea is appealing for its simplicity, the error can be substantial. A more straightforward idea is to avoid intermediate simple space, by embedding one shape directly into another. This formulation is firstly developed in [9], and its solver is built on the Gromov-Hausdorff distance framework [49]. A spectral decomposition method is subsequently proposed in [1] to further reduce the complexity. Mathematically, the Quadratic Assignment Problem (QAP) has strong and direct connections to the optimization problem raised by Gromov-Hausdorff distance [32] or its relaxation form Gromov-Wasserstein distance [48], as indicated in [47, 57, 66]. Essentially, many methods adopt a QAP formulation by incorporating geodesic distance preservation into a quadratic objective, and optimizing over permutation matrices [5, 12, 18, 66].

*Functional Maps Methods.* The seminal work of functional maps is originally introduced in [52]. In this formulation and equipped with the spectral decomposition of the Laplace-Beltrami operator, the matching problem reduces to linear optimization over a small functional map matrix, from which dense correspondences can be extracted. This new perspective has encouraged a large volume of works, including partial shape matching [41, 42, 58], additional regularization [37, 54, 55], orientation-preserving operator [56], and alternative functional space representations [38, 72]. Recently, it has been observed that the conversion step from functional maps to dense point-to-point correspondences is also difficult and prone to error [25, 59]. In light of this, a group of methods have been proposed to improve the accuracy and achieved state-of-the-art. These include CPD in the spectral domain [59], kernel density estimation in the product space [70], coarse-to-fine spectral upsampling [46] and its sinkhorn variant [53], among others [25, 34, 56, 73]. In addition, deep neural networks are also investigated for functional maps [17, 28, 33, 40, 60, 63].

*Extrinsic Methods.* The above two categories leverage the intrinsic properties of shapes. A more closely related line of work to ours is the extrinsic methods, which attempts to align shapes in the Euclidean (ambient) space directly. Due to fundamental difficulties of extrinsic shape modeling, the volume of methods are much smaller. In early period, many popular deformation models have been adopted, including Reproducing Kernel Hilbert Space [51], Thin-Plate Splines [13], PCA type representation [43], locally affine model [2], and as-rigid-as-possible model [35]. However, these models pose strong assumptions on the deformation and cannot handle large and uneven deformations ubiquitously existing in shape matching. Time-dependent models are more powerful and permit a broader range of deformations, which is adopted in shape matching in [71] and more recently in [19] on divergence-free fields. However,

the complexity of such methods is prohibitive. Notably, a vector field flow method in [15] refines the functional maps rather than aligns the shape in ambient space. More recently, the smooth shells method [20] provides a hierarchical framework that combines functional maps and as-rigid-as-possible deformation to align shapes in an intrinsic-extrinsic embedding space. Most recently, deformation-based deep learning methods are also emerging [21, 22, 27].

### 3. Methodology

We assume two shapes  $\mathcal{X}$  and  $\mathcal{Y}$  are given, which can be seen as 2D Riemannian manifolds with embeddings in  $\mathbb{R}^3$ . Note that here “embedding” takes a different meaning where  $\mathbb{R}^3$  is the embedding space. In other contexts, embedding also refers to a feature transform from ambient space  $\mathbb{R}^3$  to a high-dimensional embedding space, *e.g.* the process of spectral embedding. Typically, the shapes  $\mathcal{X}$  and  $\mathcal{Y}$  are discretized and represented using triangle meshes. We denote the vertices as point sets  $\mathbf{X} \in \mathbb{R}^{N \times 3}$  and  $\mathbf{Y} \in \mathbb{R}^{M \times 3}$ , respectively. The aim of shape matching is to find the point-to-point mapping  $T : \mathcal{X} \rightarrow \mathcal{Y}$ , *i.e.* to find the (partial) permutation matrix  $\Pi$  between  $\mathbf{X}$  and  $\mathbf{Y}$  in the discrete setting.

#### 3.1. Generalized Functional Deformation

##### 3.1.1 Deformation based on Regularization Theory

The core idea of CPD is to impose a smooth constraint on the displacement function of points in the registration process. To solve the problem, a variational formulation is proposed in [51], which originates from a broader framework, known as Regularized Network [24]. Basically, it investigates a classic machine learning problem to minimize the empirical risk functional plus a regularization term:

$$H[f] = \frac{1}{l} \sum_{i=1}^l c(f(\mathbf{x}_i), y_i) + \frac{\lambda}{2} \|Pf\|^2, \quad (1)$$

where  $l$  is the number of input data pairs  $(\mathbf{x}_i, y_i)$ , and  $c(\cdot)$  is cost function. Practically,  $P$  can be seen as an operator that “extracts” some parts of the function for regularization.

For a detailed discussion to derive the solution of Eq. (1), we refer to the seminal work of [24, 30]. The main conclusion is related to the kernel trick in machine learning, that under rather general conditions the solution can be written in the following form:

$$\begin{aligned} f(\mathbf{x}) &= \sum_{i=1}^l w_i K(\mathbf{x}_i, \mathbf{x}), \\ \|Pf\|^2 &= \mathbf{w}^T \mathbf{K} \mathbf{w}, \end{aligned} \quad (2)$$

where  $w_i$  denotes the weights,  $K(\cdot, \cdot)$  denotes a particular kernel function,  $\mathbf{w} = [w_1, \dots, w_l]^T$  denotes the weight vector, and  $\mathbf{K}$  denotes the kernel matrix  $\mathbf{K}_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$ .

An important perspective that allows us to generalize the framework later is to observe the relation between the regularization operator  $P$  and the kernel function  $K$ . In [65], the

following theorem has been proved which illuminates this relation using the concept of Green’s function [29].

#### Theorem 1 (Green’s Function and Mercer Kernels)

Let  $P$  be a regularization operator, and  $K$  be the Green’s function of self-adjoint operator  $\hat{P}P$  (with  $\hat{P}$  denoting the adjoint operator to  $P$ ). Then  $K$  is a Mercer Kernel that minimizes Eq. (1) with  $P$  as the operator.

Theorem 1 indicates that given the predefined operator, the kernel function can be accordingly derived which characterizes the solution form in Eq. (2).

With the theoretical background established, now we can turn to the idea of CPD. In practice, the regularization term is usually used without specifying the operator, which admits the following form:

$$\phi(f) = \|Pf\|^2 = \langle f, \hat{P}Pf \rangle = \int_{\mathbb{R}^D} |\hat{f}(\omega)|^2 r(\|\omega\|^2) d\omega, \quad (3)$$

where  $\hat{f}(\omega)$  denotes the Fourier transform of  $f$ ,  $r(\|\omega\|^2)$  is a function typically increasing in  $\|\omega\|^2$  to penalize high frequency components. In essence, CPD adopts a Gaussian function, *i.e.*  $r(\|\omega\|^2) = e^{-\beta\|\omega\|^2}$ , to regularize the displacement function, which induces a Gaussian Kernel to formulate the solution in Eq. (2).

##### 3.1.2 Generalized Coherent Point Drift

The original CPD utilizes a regularized framework to model the smooth displacement function  $f \in L^2(\mathbb{R}^D)$ . Clearly, the concepts here are established on the basis of Euclidean space, and the smoothness of function is also defined *w.r.t.* Euclidean space. This poses a major limitation to its range of application, especially when the data are endowed with a non-Euclidean structure, *e.g.* natural shapes.

In current literature of shape analysis, shapes are constantly considered as a low-dimensional manifold embedded in  $\mathbb{R}^3$  and presented as 3D meshes, and an enormous amount of studies making use of the intrinsic properties of shapes have proven to be very successful. Thus it is natural to generalize the idea of CPD to non-Euclidean domains. The generalization is proposed based on several observations and conclusions that have been made in [64]. As we will shortly see, the kernel trick in Euclidean space can be readily generalized to non-Euclidean domains.

The first observation is that the regularization term in Eq. (3) can be equivalently written in terms of Laplace operator as

$$\int_{\mathbb{R}^D} |\hat{f}(\omega)|^2 r(\|\omega\|^2) d\omega = \langle f, r(\Delta)f \rangle, \quad (4)$$

where  $\Delta$  denotes the Laplace operator and  $r(\Delta)$  is the extension of  $r$  to operators simply by applying  $r$  to the spectrum of  $\Delta$

$$\langle f, r(\Delta)f \rangle = \sum_i \langle f, \psi_i \rangle r(\lambda_i) \langle \psi_i, f \rangle, \quad (5)$$

where  $\{(\psi_i, \lambda_i)\}$  is the eigensystem of  $\Delta$ .

**Remark 1.** The importance of Eq. (4) is that it links the Laplace operator to the self-adjoint operator  $\hat{P}P$  as in Eq. (3). Thus by the definition of  $r(\Delta)$  we may determine the form of the kernel function by virtue of Theorem 1.

Now, when we are studying a shape, or equivalently a manifold  $\mathcal{M}$  embedded in  $\mathbb{R}^3$ , the notion of Laplace operator for Euclidean space is not applicable, but there exists a counterpart, *i.e.* the Laplacian-Beltrami operator. Thus, we may define the regularization term in complete analogy to Eq. (4). Note that since the input shapes are a set of discrete samples of the manifold, the Laplacian-Beltrami operator degenerates to a matrix, denoted by  $\mathcal{L}$ . In the spirit of Eq. (5), we may define

$$r(\mathcal{L}) = \sum_i r(\lambda_i) v_i v_i^\top, \quad (6)$$

where  $\{(\lambda_i, v_i)\}$  constitutes the eigensystem of  $\mathcal{L}$ . Under this equation, kernel trick can be readily extended to non-Euclidean domain [64]. Particularly,  $r(\mathcal{L})$  plays the same role of  $\hat{P}P$  of the continuous case, and the kernel matrix in discrete case is exactly the inverse of  $r(\mathcal{L})$  (pseudo-inverse if not invertible), *i.e.* the discrete Green's function as defined in [14]:

$$\mathbf{K} = r(\mathcal{L})^{-1} = \sum_i r^{-1}(\lambda_i) v_i v_i^\top. \quad (7)$$

The different definitions of function  $r(\cdot)$  would result in different regularization operators, and equivalently, different kernels. This completes our generalization since the form of solution is determined by the form of the kernel, in complete analogy to Eq. (2).

**An Alternative Representation.** Analogously to the continuous case, the kernel trick suffers from high computational complexity since  $K \in \mathbb{R}^{N \times N}$ , where  $N$  denotes the number of data points. This is a major limitation for processing shapes since high-resolution shapes can be of tens of thousands points. In this sense, the kernel trick may be inferior to explicitly specifying the form of solution by utilizing the Fourier transform, *i.e.* the eigensystem of the Laplace-Beltrami operator  $\mathcal{L}$ .

Using the intuitive idea of using low-frequency truncated bases, we may express the function of interest as

$$\mathbf{f} = \sum_i w_i v_i = \mathbf{U}\mathbf{w}, \quad (8)$$

where  $\mathbf{f}$  is a vector representing a function over  $\mathcal{M}$  evaluated on sample points.  $\mathbf{U} = [v_1, v_2, \dots, v_k]$  is the matrix of  $k$  truncated bases. The regularization term by applying Eq. (6) now becomes

$$\phi(\mathbf{f}) = \sum_i r(\lambda_i) w_i^2. \quad (9)$$

Practically, we can always use a truncated basis for representing  $\mathbf{f}$ , then the number of variables for representing  $\mathbf{f}$  is

identical to the number of bases used, and independent of the number of data points.

**Remark 2.** The expression of Eq. (8) is a simple and rather intuitive result, which has been used in previous work to model deformation [20]. Our analysis here theoretically connects it to the regularization theory and kernel method, and leads to a new algorithm under the CPD framework.

## 3.2. Generalized Probabilistic Registration

### 3.2.1 Probabilistic Registration Framework

The Gaussian Mixtures Model (GMM) is a classic and prevalent method for the task of registration, where the goal is to simultaneously learn the permutation of points as well as a deformation  $\mathcal{T}(\cdot)$  that aligns two point sets. Following the conventions in [51], we consider the points in  $\mathbf{Y}$  as the GMM centroids and the points in  $\mathbf{X}$  as the data points generated by GMM. For simplicity, equal isotropic covariances  $\sigma^2$  and equal membership probabilities are adopted. To account for outliers, an additional uniform distribution  $\frac{1}{a}$  is used. Consequently, the mixture model takes the form

$$p(\mathbf{x}_n) = \gamma \frac{1}{a} + (1 - \gamma) \sum_{m=1}^M \frac{1}{M} p(\mathbf{x}_n | m), \quad (10)$$

where  $p(\mathbf{x}_n | m) = \frac{1}{(2\pi\sigma^2)^{D/2}} e^{-\frac{\|\mathbf{x}_n - \mathcal{T}(\mathbf{y}_m, \theta)\|^2}{2\sigma^2}}$ ,  $\theta$  denotes the parameters that control the deformation, and  $\gamma$  denotes the weight accounting for outliers.

Adopting the *i.i.d.* data assumption, a Maximum Likelihood Estimate (MLE) objective can be defined utilizing Eq. (10) as  $p(\mathbf{X}) = \prod_{n=1}^N p(\mathbf{x}_n)$ . To circumvent the difficulties of optimizing the MLE objective, a smart EM algorithm can be used. The idea of EM is the iterations between an *E-step* and an *M-step*. The *E-step* is to use current parameters to compute *a posteriori* probability distributions of mixture components

$$p(m | \mathbf{x}_n) = \frac{e^{-\frac{\|\mathbf{x}_n - \mathcal{T}(\mathbf{y}_m, \theta)\|^2}{2\sigma^2}}}{\sum_{k=1}^M e^{-\frac{\|\mathbf{x}_n - \mathcal{T}(\mathbf{y}_k, \theta)\|^2}{2\sigma^2}} + (2\pi\sigma^2)^{D/2} \frac{\gamma}{1-\gamma} \frac{M}{N}}. \quad (11)$$

The *M-step* is to update new parameters by minimizing the expectation of the complete negative log-likelihood function, which, if ignoring independent terms, is written as

$$Q(\theta, \sigma^2) = \sum_{n=1}^N \sum_{m=1}^M p(m | \mathbf{x}_n) \frac{\|\mathbf{x}_n - \mathcal{T}(\mathbf{y}_m, \theta)\|^2}{2\sigma^2} + \frac{N_p D}{2} \log \sigma^2, \quad (12)$$

where  $N_p = \sum_{n=1}^N \sum_{m=1}^M p(m | \mathbf{x}_n)$ .

Note that to facilitate the estimation of deformation, regularization terms are commonly required in addition to Eq. (12). This can be embedded in the GMM probabilistic framework by specifying prior information for deformation, but essentially boils down to minimizing the following objective

$$Q'(\theta, \sigma^2) = Q(\theta, \sigma^2) + \frac{\lambda}{2} \phi(\theta), \quad (13)$$

where  $\phi(\theta)$  is the regularization term.

Now having Eq. (13) in hand as objective, we may discuss how to update the parameters, *i.e.*  $\theta$  and  $\sigma^2$ . Based on our discussion in Sec. 3.1.2, the deformation is defined using displacement function

$$\mathcal{T}(\mathbf{Y}, \mathbf{W}) = \mathbf{Y} + v(\mathbf{Y}) = \mathbf{Y} + \mathbf{U}\mathbf{W}, \quad (14)$$

where  $\mathbf{U} \in \mathbb{R}^{M \times k}$  stacks the truncated  $k$  bases of the Laplacian-Beltrami operator  $\mathcal{L}$ , and  $\mathbf{W} \in \mathbb{R}^{k \times D}$  is the weight matrix to be solved. The regularization term takes the form

$$\phi(\mathbf{W}) = \mathbf{W}^\top \mathbf{R}\mathbf{W}, \quad (15)$$

where  $\mathbf{R}$  is a diagonal matrix and the diagonal elements are the penalizing function of the eigenvalues of the Laplacian-Beltrami operator  $\mathcal{L}$ , *i.e.*  $\mathbf{R}_{ii} = r(\lambda_i)$ . Applying Eq. (14) and Eq.(15) to Eq. (13), we can obtain an analytical solution by solving the following linear system

$$(\mathbf{U}^\top d(\mathbf{P}\mathbf{1})\mathbf{U} + \lambda\sigma^2\mathbf{R})\mathbf{W} = \mathbf{U}^\top \mathbf{P}\mathbf{X} - \mathbf{U}^\top d(\mathbf{P}\mathbf{1})\mathbf{Y}, \quad (16)$$

where  $\mathbf{P}$  has elements  $\mathbf{P}_{mn} = p(m|\mathbf{x}_n)$ ,  $\mathbf{1}$  represents the column vector of all ones, and  $d(\cdot)$  denotes the transform from vector to diagonal matrix. In addition, parameter  $\sigma^2$  can also be derived analytically using simple linear algebra

$$\begin{aligned} \sigma^2 &= \frac{1}{N_p D} \sum_{n=1}^N \sum_{m=1}^M p(m|\mathbf{x}_n) \|\mathbf{x}_n - \mathcal{T}(\mathbf{y}_m, \theta)\|^2 = \\ &= \frac{1}{N_p D} (tr(\mathbf{X}^\top d(\mathbf{P}^\top \mathbf{1})\mathbf{X}) - 2tr((\mathbf{P}\mathbf{X}\mathbf{T}) + tr(\mathbf{T}^\top d(\mathbf{P}\mathbf{1})\mathbf{T})), \end{aligned} \quad (17)$$

where  $\mathbf{T} = \mathbf{Y} + \mathbf{U}\mathbf{W}$ . This completes the whole EM algorithm for GMM-based registration.

### 3.2.2 Generalized Probabilistic Model

The major drawback of the aforementioned GMM framework for registration is its highly non-convex objective, which makes it prone to failures caused by bad initialization. To address this issue, we borrow the idea from image matching [26, 44]. First, we create a set of initial correspondences by nearest matching of feature descriptors, and then estimate transformation from the initial correspondence set. The initial correspondence set contains the most confident correspondences and significantly reduces the search space for registration.

In shape matching, feature descriptors have also been extensively studied in analogy to image matching. For example, SHOT [67] is a frequently used one. In a general sense, the functional maps framework provides a more sophisticated and powerful feature descriptor for shape matching, which greatly enhances spectral embeddings for matching. Thus, we can safely assume that the initial correspondences are given and turn to the next step.

The problem we need to concern for initial correspondence set is the noise, and more importantly, the outlier issue (false correspondences). We next propose a probabilistic framework to address this problem.

Suppose we have obtained a set of putative correspondences  $\mathbf{S} = \{\mathbf{c}_i\}_{i=1}^l$ , where each  $\mathbf{c}_i = (\mathbf{x}_i, \mathbf{y}_i)$ . To identify the inlier set  $\mathcal{I} \subseteq \mathbf{S}$ , we aim to recover from  $\mathbf{S}$  the underlying deformation, *i.e.*  $\mathcal{T} : \mathbb{R}^D \rightarrow \mathbb{R}^D$  in our context, such that  $\mathbf{x}_i = \mathcal{T}(\mathbf{y}_i)$  for  $\mathbf{c}_i \in \mathcal{I}$ . For the inliers, we assume that the noise is isotropic Gaussian with zero mean and variance  $\sigma^2$ . For the outliers, we assume its distribution to be uniform  $\frac{1}{a}$  with  $a$  denoting the volume of this region. We also use  $\gamma$  as the weight to account for outliers. Consequently, the probabilistic model takes the form

$$p(\mathbf{c}_i) = \frac{1-\gamma}{(2\pi\sigma^2)^{D/2}} e^{-\frac{\|\mathbf{x}_i - \mathcal{T}(\mathbf{y}_i)\|^2}{2\sigma^2}} + \gamma \frac{1}{a}. \quad (18)$$

**Remark 3.** We note the resemblance between Eq. (10) and Eq. (18). In essence, our probabilistic model to handle initial correspondences makes a more restricted assumption of the data distribution, *i.e.* the mixture model contains only two components instead of  $M + 1$  in Eq. (10).

Given the assumption of data distribution in Eq. (18), we may now determine the objective as  $p(\mathbf{S}) = \prod_{i=1}^l p(\mathbf{c}_i)$ , using the spirit of MLE. To derive the optimal solution under the MLE objective, the EM algorithm can also be used. In this case, we may also proceed by the iteration of an *E-step* and an *M-step*. The *E-step* to compute *a posteriori* probability distributions of mixture components now read

$$p_i = \frac{(1-\gamma)e^{-\frac{\|\mathbf{x}_i - \mathcal{T}(\mathbf{y}_i)\|^2}{2\sigma^2}}}{(1-\gamma)e^{-\frac{\|\mathbf{x}_i - \mathcal{T}(\mathbf{y}_i)\|^2}{2\sigma^2}} + \gamma \frac{(2\pi\sigma^2)^{D/2}}{a}}. \quad (19)$$

The *M-step* to update new parameters is conducted by minimizing the expectation of the complete negative log-likelihood function as

$$\mathcal{Q}(\theta, \sigma^2) = \sum_{i=1}^l p_i \frac{\|\mathbf{x}_i - \mathcal{T}(\mathbf{y}_i, \theta)\|^2}{2\sigma^2} + \frac{D}{2} \log \sigma^2 \sum_{i=1}^l p_i. \quad (20)$$

Analogously, the regularization term can be included and the final objective is

$$\mathcal{Q}'(\theta, \sigma^2) = \mathcal{Q}(\theta, \sigma^2) + \frac{\lambda}{2} \phi(\theta), \quad (21)$$

where  $\phi(\theta)$  is the regularization term.

In a similar way, we may apply Eq. (14) and Eq. (15) to Eq. (21). This leads to a closed-form solution for the deformation parameter matrix  $\mathbf{W}$

$$(\mathbf{U}^\top d(\mathbf{p})\mathbf{U} + \lambda\sigma^2\mathbf{R})\mathbf{W} = \mathbf{U}^\top d(\mathbf{p})\mathbf{X} - \mathbf{U}^\top d(\mathbf{p})\mathbf{Y}, \quad (22)$$

where  $\mathbf{p}$  is the vector of *a posteriori* probability distributions, *i.e.*  $\mathbf{p}_i = p_i$ , and  $\mathbf{X} \in \mathbb{R}^{l \times D}$  and  $\mathbf{Y} \in \mathbb{R}^{l \times D}$  stack  $\mathbf{x}_i$

---

**Algorithm 1** Generalized Coherent Point Drift

---

**Input:** The initial confident correspondence set  $\mathbf{S}$ , two shapes  $\mathcal{X}$  with point set  $\mathbf{X}$  and  $\mathcal{Y}$  with point set  $\mathbf{Y}$  represented as 3D meshes, parameters  $\gamma$ ,  $\lambda$  and  $k$ .

**Output:** Point-to-point correspondences  $\mathcal{C}$ .

- 1: Initialize  $\gamma$ ,  $\lambda$ ,  $\mathbb{T} = \mathbf{0}_{l \times 2}$ ,  $\mathbf{p} = \mathbf{1}$ ;
  - 2: Initialize  $\sigma^2$  by Eq. (23);
  - 3: Compute the smallest first  $k$  eigenvalues and eigenfunctions of the Laplace-Beltrami operator  $\mathcal{L}$  on shape  $\mathcal{Y}$  to form basis Matrix  $\mathbf{U}$  and regularization matrix  $\mathbf{R}$ ;
  - 4: **while**  $Q$  not converge **do**
  - 5:    *E-step:*
  - 6:    Update  $\mathbf{p}$  by Eq. (19);
  - 7:    *M-step:*
  - 8:    Update  $\mathbb{W}$  by Eq. (22),  $\mathbb{T}$  by Eq. (14), and  $\sigma^2$  by Eq. (23);
  - 9: **end while**
  - 10: Initialize  $\mathbf{T}$  by Eq. (14);
  - 11: **while**  $Q$  not converge **do**
  - 12:    *E-step:*
  - 13:    Update  $\mathbf{P}$  by Eq. (11);
  - 14:    *M-step:*
  - 15:    Update  $\mathbf{W}$  by Eq. (16),  $\mathbf{T}$  by Eq. (14), and  $\sigma^2$  by Eq. (17);
  - 16: **end while**
  - 17: Establish point-to-point correspondences  $\mathcal{C}$  by nearest matching between  $\mathbf{X}$  and  $\mathbf{T}$ .
- 

and  $\mathbf{y}_i$  in a correspondence  $\mathbf{c}_i$  respectively. Using simple linear algebra,  $\sigma^2$  has the closed-form solution as

$$\sigma^2 = \frac{\text{tr}((\mathbf{Y} - \mathbb{T})^\top d(\mathbf{p})(\mathbf{Y} - \mathbb{T}))}{D \cdot \text{tr}(d(\mathbf{p}))}, \quad (23)$$

where  $\mathbb{T} = \mathbb{Y} + \mathbf{U}\mathbf{W}$ , and  $\mathbf{W} \in \mathbb{R}^{k \times 3}$ . This completes the whole EM algorithm for deformation estimation using a generalized probabilistic model.

**Algorithmic Summary of Our Method.** Our method assumes that a set of confident initial correspondences are provided, thus we use the generalized probabilistic model to estimate a deformation, and subsequently we use the original GMM-based registration framework to refine the estimation. We name our algorithm as *Generalized Coherent Point Drift* (GCPD) and summarize it in Alg. 1.

### 3.3. Implementation Details

**Computational Complexity.** Assuming the initial correspondences are given, the most time-consuming step is the process of evaluating *a posteriori* probability distributions as in Eq. (11), taking  $O(MN)$  time by straightforward computation. The Fast Gauss Transform is used in CPD to mitigate this issue, which achieves  $O(M + N)$  complexity.

In our method, we utilize a more accurate Figtree method as proposed in [50], which is also of  $O(M + N)$  complexity.

**Initial Confident Correspondences.** Note that a plethora of methods have been proposed for establishing confident correspondences in the literature, based on descriptor matching. In some sense, the Functional Maps methods can also be seen as of this category, which output correspondences based on matching of spectral embeddings. For optimal performance, the output of a Functional Maps method is favored, as will be used in our experiment.

## 4. Experimental Results

### 4.1. Benchmark Comparison

We use three commonly used benchmark datasets, *i.e.* FAUST [7], TOSCA [10] and SCAPE [3] to evaluate the proposed GCPD and other competitor methods.

**FAUST.** FAUST is composed of 10 poses of 10 human subjects with significant variability between different human subjects. Instead of the standard meshes, we use the version provided in [56] where each shape was re-meshed individually. This makes it more challenging and also more realistic, since for real-world scans the sampling of surfaces is generally incompatible. The meshes all have approximately 5k vertices. In our evaluation, we use the 300 matching pairs provided in the dataset, which consists of both isometric and non-isometric pairs.

**TOSCA.** TOSCA consists of 76 shapes in 8 different categories (human and animal shapes) with vertex numbers ranging from 4k to 50k. The high-resolution shapes make this dataset computationally very challenging. In our evaluation, we randomly generate 200 isometric matching pairs.

**SCAPE.** SCAPE contains 71 registered meshes of a particular human subject in different poses. The meshes have identically 12500 vertices. The dataset is a relatively simple one where the meshes are consistent and the resolution is moderate. In our evaluation, we also randomly generate 200 matching pairs for evaluation.

**Evaluation Metric.** We use the Princeton benchmark protocol [36] to evaluate the matching accuracy of methods. Specifically, given the ground-truth match  $(x, y^*)$  where  $x \in \mathcal{X}$  and  $y^* \in \mathcal{Y}$ , the error of the calculated match  $(x, y)$  is given by the geodesic distance between  $y$  and  $y^*$  normalized by diameter of  $\mathcal{Y}$ :  $\epsilon(x) = \frac{d_{geo}(y, y^*)}{\sqrt{\text{area}(\mathcal{Y})}}$ . All evaluated correspondences are summarized by a cumulative distribution curve which indicates the final matching performance.

**Competitors and Settings.** The competitors include BIM [36], PMF [70], KernelMatching [69], BCICP [56], SmoothShells [20], MWP [34], ZoomoutSinkhorn [46, 53] and DIR [73]. Among them, we use the provided results in the dataset of [56] for BIM, BCICP, and KernelMatching. We use SHOT [67] matching results to initialize MWP,




Figure 2. Evaluation results of our GCPD and other competing methods on FAUST, TOSCA and SCAPE, in terms of matching accuracy.



Figure 3. Qualitative demonstration using color transfer. For each group of results, the first and last shapes are target and source shapes with ground-truth matching, the second shape represents result of SmoothShells, and the third shape represents result of GCPD.

DIR and ZoomoutSinkhorn. We use the results of MWP to initialize PMF, and use the single-scale version of it.

**GCPD Settings.** In our evaluation, we use the results of MWP to initialize our GCPD for optimal performance. Specifically, for high-resolution shape pairs, we use the output of MWP to produce 5000 sparse matching samples. For low-resolution shape pairs, we directly use all correspondences produced by MWP as initialization. Note that different initialization strategies are analyzed in the next section. We set  $\gamma = 0.1$ ,  $\lambda = 10$  and  $k = 500$  in our experiments. We empirically choose  $r(\lambda_i) = \lambda_i^{1/2}$  for regularization.

**Discussion of Results.** We have provided the comparison results measured by the cumulative distribution curve of matching accuracy in Fig. 2. Some qualitative examples are given in Fig. 3 with comparison to SmoothShells. We also compare the mean runtime of each method on SCAPE dataset and runtime on some representative models from TOSCA in Tab. 1. Note that the runtime includes all pre-processing cost, such as SHOT matching and eigenfunction computation. From the results, we can observe that the recently proposed functional maps variants, such as BCICP, ZoomoutSinkhorn and MWP, significantly outperform previous methods. Among them, the most recent MWP has the best performance, in terms of both matching accuracy and efficiency. This validates our choice for initialization with MWP. Notably, PMF, like our method, also requires initial correspondences. However, PMF fails to improve over MWP input. Among all methods, our GCPD and SmoothShells have the best performance, this is because each method recovers a deformation for matching. Both methods have near-optimal performance, while our method is significantly faster, almost 10 times faster than SmoothShells on SCAPE dataset. This establishes our

Table 1. The runtime comparison (in seconds) of GCPD and state-of-the-arts. We report the average runtime on SCAPE dataset, and runtime of models from TOSCA dataset which reflect different resolutions. SS: SmoothShells; ZS: ZoomoutSinkhorn.

Model	SCAPE	Wolf	Centaur	Horse	Cat	David
Vertices	12500	4344	15768	19248	27894	52565
SS [20]	410.9	149.8	506.2	709.9	986.8	2035.5
ZS [53]	32.8	4.3	43.6	52.6	104.1	583.9
MWP [34]	18.5	2.5	26.3	26.6	58.5	404.1
GCPD	29.4	6.8	81.8	86.1	160.1	813.2

method as the best performer.

**Remark 4.** The DIR method has also been adopted for comparison. However, we only present its results on SCAPE dataset. This is because on TOSCA dataset, the required memory of DIR exceeds the limitation of our machine (64 GB) for high-resolution pairs, due to the need to compute geodesic distance matrix. Also on FAUST, we find that DIR almost fails on each pair, which may be due to the low-resolution of the shapes. Thus although DIR exhibits promising results on SCAPE, it still has some serious limitations in contrast to our method.

## 4.2. Further Analyses

**Ablation Study.** We provide an ablation study of the proposed generalized functional deformation and robust initialization strategy in Fig. 4a. The experiment is conducted on FAUST, which contains both isometric and non-isometric matching pairs. Explanation of settings: To find transformation from  $Y$  to  $X$ , **naive initialization** as used in CPD is to perform matching directly (initialized with  $Y$  itself) with GMM (Sec. 3.2.1). The **proposed initialization**




Figure 4. Further analytical results. (a) Ablation study w.r.t. the generalized functional deformation and robust initialization strategy. (b) Evaluation results of GCPD w.r.t. different initializations. Full line: GCPD refined results; Dashed Line: Initialization methods without GCPD. (c) Evaluation results of GCPD, MWP and ZoomOut for partial shape matching. Full line: cuts dataset; Dashed line: holes dataset.




Figure 5. Qualitative examples showing how our GCPD improves over MWP. For each group, the left two shapes present the matching result of MWP, and the right two shapes present the refined matching result of our GCPD. Note that here the unmatched points are colored as white.

utilizes a different probabilistic model first (Sec. 3.2.2) with initial correspondences (MWP) to obtain  $T$ , then  $T$  is fed into GMM instead of Y. CPD and GCPD simply refers to different functional deformation frameworks. Clearly, the proposed initialization and generalized functional deformation are important for non-rigid shape matching. It is also consistent with the common view for CPD that without a good initialization, the GMM framework fails easily.

**Initialization Choice.** One important aspect of our method is the choice for initialization. We have shown that GCPD initialized by MWP can achieve superior performance, but it remains unclear to what extent the initialization would influence the matching result. Next we provide an experimental analysis of this point. In particular, we additionally use ZoomoutSinkhorn and also plain SHOT matching to provide initial correspondences for our GCPD. The experiments are conducted on FAUST, and the results are presented in Fig. 4b. It can be observed that although GCPD improves the matching accuracy over different initial correspondences, there is a gap between different initialization methods. A better matching result can be expected if the initial correspondence set is better. We also provide some qualitative results involving both isometric and non-isometric pairs in Fig. 5, showing the initial matching result of MWP and how GCPD improves over it.

**Partial Matching.** A prominent advantage of our GCPD method compared to other methods, such as SmoothShells,



Figure 6. Qualitative examples of our GCPD and MWP using color transfer for partial matching. For each group, the first and last columns are target and source shapes with ground-truth matching, the second column represents the results of MWP, and the third column represents the results of our GCPD.

is its flexibility. We make the least of assumptions of the input shapes, and only require a set of initial matches. This property renders our method directly applicable to partial matching. To demonstrate this point, we take the SHREC16 Partial Correspondence benchmark [16] for evaluation. The dataset consists of 8 types of isometric human or animal shapes in different poses with regular “cuts” and irregular “holes”. We test our method by matching each partial shape to the corresponding full shape. For initialization, we still use MWP. The quantitative results are presented in Fig. 4c with comparison to the MWP and ZoomOut [46] method using the tailored partial matching code provided by the authors. The qualitative results are presented in Fig. 6. We can see that our method can enhance the results of MWP even in the partial matching setting.

## 5. Conclusion

In this paper, we have revisited and generalized the classic CPD method for non-rigid shape matching. First, the deformation model has been generalized to non-Euclidean domains, which makes the method applicable to shapes that involve manifold structures. Second, we have proposed a generalized probabilistic model which can handle outlier-contaminated initial correspondences. This avoids the local optima issue in optimizing the GMM-registration objective. Extensive experiments have been conducted on multiple benchmark datasets which proves the efficacy of the proposed method.



## References

- [1] Yonathan Aflalo, Anastasia Dubrovina, and Ron Kimmel. Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision*, 118(3):380–392, 2016. [2](#)
- [2] Brian Amberg, Sami Romdhani, and Thomas Vetter. Optimal step nonrigid icp algorithms for surface registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007. [2](#)
- [3] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *ACM SIGGRAPH*, pages 408–416. 2005. [6](#)
- [4] Florian Bernard, Luis Salamanca, Johan Thunberg, Alexander Tack, Dennis Jentsch, Hans Lamecker, Stefan Zachow, Frank Hertel, Jorge Goncalves, and Peter Gemmar. Shape-aware surface reconstruction from sparse 3d point-clouds. *Medical image analysis*, 38:77–89, 2017. [1](#)
- [5] Florian Bernard, Christian Theobalt, and Michael Moeller. Ds\*: Tighter lifting-free convex relaxations for quadratic matching problems. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4310–4319, 2018. [2](#)
- [6] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, pages 586–606, 1992. [2](#)
- [7] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014. [6](#)
- [8] Ingwer Borg and Patrick JF Groenen. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2005. [2](#)
- [9] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5):1168–1172, 2006. [2](#)
- [10] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. [6](#)
- [11] Michael M Bronstein, Alexander M Bronstein, Ron Kimmel, and Irad Yavneh. Multigrid multidimensional scaling. *Numerical Linear Algebra with Applications*, 13(2-3):149–171, 2006. [2](#)
- [12] Qifeng Chen and Vladlen Koltun. Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2039–2047, 2015. [2](#)
- [13] Haili Chui and Anand Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2-3):114–141, 2003. [2](#)
- [14] Fan Chung and S-T Yau. Discrete green’s functions. *Journal of Combinatorial Theory, Series A*, 91(1-2):191–214, 2000. [4](#)
- [15] Etienne Corman, Maks Ovsjanikov, and Antonin Chambolle. Continuous matching via vector field flow. In *Computer Graphics Forum*, volume 34, pages 129–139. Wiley Online Library, 2015. [1](#), [3](#)
- [16] Luca Cosmo, Emanuele Rodola, Jonathan Masci, Andrea Torsello, and Michael M Bronstein. Matching deformable objects in clutter. In *Proceedings of the International Conference on 3D Vision*, pages 1–10, 2016. [8](#)
- [17] Nicolas Donati, Abhishek Sharma, and Maks Ovsjanikov. Deep geometric functional maps: Robust feature learning for shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8592–8601, 2020. [2](#)
- [18] Nadav Dym, Haggai Maron, and Yaron Lipman. Ds++: A flexible, scalable and provably tight relaxation for matching problems. *ACM Transactions on Graphics*, 2017. [2](#)
- [19] Marvin Eisenberger, Zorah Löhner, and Daniel Cremers. Divergence-free shape correspondence by deformation. In *Computer Graphics Forum*, volume 38, pages 1–12. Wiley Online Library, 2019. [1](#), [2](#)
- [20] Marvin Eisenberger, Zorah Lahner, and Daniel Cremers. Smooth shells: Multi-scale shape registration with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12265–12274, 2020. [3](#), [4](#), [6](#), [7](#)
- [21] Marvin Eisenberger, David Novotny, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7473–7483, 2021. [3](#)
- [22] Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep shells: Unsupervised shape correspondence with optimal transport. In *Advances in Neural Information Processing Systems*, 2020. [3](#)
- [23] Asi Elad and Ron Kimmel. On bending invariant signatures for surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10):1285–1295, 2003. [2](#)
- [24] Theodoros Evgeniou, Massimiliano Pontil, and Tomaso Poggio. Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13(1):1–50, 2000. [3](#)
- [25] Danielle Ezuz and Mirela Ben-Chen. Deblurring and denoising of maps between shapes. In *Computer Graphics Forum*, volume 36, pages 165–174. Wiley Online Library, 2017. [2](#)
- [26] Aoxiang Fan, Jiayi Ma, Xingyu Jiang, and Haibin Ling. Efficient deterministic search with robust loss functions for geometric model fitting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [2](#), [5](#)
- [27] Wanquan Feng, Juyong Zhang, Hongrui Cai, Haofei Xu, Junhui Hou, and Hujun Bao. Recurrent multi-view alignment network for unsupervised surface registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10297–10307, 2021. [3](#)
- [28] Dvir Ginzburg and Dan Raviv. Cyclic functional mapping: Self-supervised correspondence between non-isometric deformable shapes. In *Proceedings of the European Conference on Computer Vision*, pages 36–52, 2020. [2](#)

- [29] Federico Girosi, Michael Jones, and Tomaso Poggio. Priors stabilizers and basis functions: From regularization to radial, tensor and additive splines. 1993. 3
- [30] Federico Girosi, Michael Jones, and Tomaso Poggio. Regularization theory and neural networks architectures. *Neural Computation*, 7(2):219–269, 1995. 3
- [31] Vladislav Golyanik, Bertram Taetz, Gerd Reis, and Didier Stricker. Extended coherent point drift algorithm with correspondence priors and optimal subsampling. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016. 1
- [32] Mikhail Gromov. *Metric structures for Riemannian and non-Riemannian spaces*. Springer Science & Business Media, 2007. 2
- [33] Oshri Halimi, Or Litany, Emanuele Rodola, Alex M Bronstein, and Ron Kimmel. Unsupervised learning of dense shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4370–4379, 2019. 2
- [34] Ling Hu, Qinsong Li, Shengjun Liu, and Xinru Liu. Efficient deformable shape correspondence via multiscale spectral manifold wavelets preservation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14536–14545, 2021. 1, 2, 6, 7
- [35] Qi-Xing Huang, Bart Adams, Martin Wicke, and Leonidas J Guibas. Non-rigid registration under isometric deformations. In *Computer Graphics Forum*, volume 27, pages 1449–1457. Wiley Online Library, 2008. 2
- [36] Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *ACM Transactions on Graphics*, 30(4):1–12, 2011. 2, 6
- [37] Artiom Kovnatsky, Michael M Bronstein, Xavier Bresson, and Pierre Vandergheynst. Functional correspondence by matrix completion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 905–914, 2015. 2
- [38] Artiom Kovnatsky, Michael M Bronstein, Alexander M Bronstein, Klaus Glashoff, and Ron Kimmel. Coupled quasi-harmonic bases. In *Computer Graphics Forum*, volume 32, pages 439–448. Wiley Online Library, 2013. 2
- [39] Yaron Lipman and Thomas Funkhouser. Möbius voting for surface correspondence. *ACM Transactions on Graphics*, 28(3):1–12, 2009. 2
- [40] Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5659–5667, 2017. 2
- [41] Or Litany, Emanuele Rodolà, Alexander M Bronstein, and Michael M Bronstein. Fully spectral partial shape matching. In *Computer Graphics Forum*, volume 36, pages 247–258. Wiley Online Library, 2017. 2
- [42] Or Litany, Emanuele Rodolà, Alexander M Bronstein, Michael M Bronstein, and Daniel Cremers. Non-rigid puzzles. In *Computer Graphics Forum*, volume 35, pages 135–143. Wiley Online Library, 2016. 2
- [43] Marcel Lüthi, Thomas Gerig, Christoph Jud, and Thomas Vetter. Gaussian process morphable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(8):1860–1873, 2017. 2
- [44] Jiayi Ma, Xingyu Jiang, Aoxiang Fan, Junjun Jiang, and Junchi Yan. Image matching from handcrafted to deep features: A survey. *International Journal of Computer Vision*, 129(1):23–79, 2021. 2, 5
- [45] Diana Mateus, Radu Horaud, David Knossow, Fabio Cuzzolin, and Edmond Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. 2
- [46] Simone Melzi, Jing Ren, Emanuele Rodola, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics*, 38(6):155:1–155:14, 2019. 1, 2, 6, 8
- [47] Facundo Mémoli. Spectral gromov-wasserstein distances for shape matching. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 256–263, 2009. 2
- [48] Facundo Mémoli. Gromov–wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics*, 11(4):417–487, 2011. 2
- [49] Facundo Mémoli and Guillermo Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5(3):313–347, 2005. 2
- [50] Vlad I Morariu, Balaji Vasani Srinivasan, Vikas C Raykar, Ramani Duraiswami, and Larry S Davis. Automatic online tuning for fast gaussian summation. In *Advances in Neural Information Processing Systems*, pages 1113–1120, 2008. 6
- [51] Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, 2010. 1, 2, 3, 4
- [52] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics*, 31(4):1–11, 2012. 1, 2
- [53] Gautam Pai, Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Fast sinkhorn filters: Using matrix scaling for non-rigid shape correspondence with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 384–393, 2021. 2, 6, 7
- [54] Jonathan Pokrass, Alexander M Bronstein, Michael M Bronstein, Pablo Sprechmann, and Guillermo Sapiro. Sparse modeling of intrinsic correspondences. In *Computer Graphics Forum*, volume 32, pages 459–468. Wiley Online Library, 2013. 2
- [55] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured regularization of functional map computations. In *Computer Graphics Forum*, volume 38, pages 39–53. Wiley Online Library, 2019. 2
- [56] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics*, 37(6):1–16, 2018. 1, 2, 6

- [57] Emanuele Rodola, Alex M Bronstein, Andrea Albarelli, Filippo Bergamasco, and Andrea Torsello. A game-theoretic approach to deformable shape matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 182–189, 2012. [2](#)
- [58] Emanuele Rodolà, Luca Cosmo, Michael M Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. In *Computer Graphics Forum*, volume 36, pages 222–236. Wiley Online Library, 2017. [1](#), [2](#)
- [59] Emanuele Rodolà, Michael Moeller, and Daniel Cremers. Point-wise map recovery and refinement from functional correspondence. In *Proc. Vision, Modeling and Visualization (VMV)*, 2015. [2](#)
- [60] Jean-Michel Roufosse, Abhishek Sharma, and Maks Ovsjanikov. Unsupervised deep learning for structured shape matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1617–1627, 2019. [2](#)
- [61] Raif M Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the Eurographics Symposium on Geometry Processing*, pages 225–233, 2007. [1](#)
- [62] Raif M Rustamov, Maks Ovsjanikov, Omri Azencot, Mirela Ben-Chen, Frédéric Chazal, and Leonidas Guibas. Map-based exploration of intrinsic shape differences and variability. *ACM Transactions on Graphics*, 32(4):1–12, 2013. [2](#)
- [63] Abhishek Sharma and Maks Ovsjanikov. Weakly supervised deep functional maps for shape matching. *Advances in Neural Information Processing Systems*, 33, 2020. [2](#)
- [64] Alexander J Smola and Risi Kondor. Kernels and regularization on graphs. In *Learning Theory and Kernel Machines*, pages 144–158. 2003. [3](#), [4](#)
- [65] Alex J Smola, Bernhard Schölkopf, and Klaus-Robert Müller. The connection between regularization operators and support vector kernels. *Neural Networks*, 11(4):637–649, 1998. [3](#)
- [66] Justin Solomon, Gabriel Peyré, Vladimir G Kim, and Suvrit Sra. Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics*, 35(4):1–13, 2016. [2](#)
- [67] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *Proceedings of the European Conference on Computer Vision*, pages 356–369. Springer, 2010. [5](#), [6](#)
- [68] Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. In *Computer Graphics Forum*, volume 30, pages 1681–1707. Wiley Online Library, 2011. [1](#)
- [69] Matthias Vestner, Zorah Löhner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodola, Alex Bronstein, Michael Bronstein, Ron Kimmel, et al. Efficient deformable shape correspondence via kernel matching. In *Proceedings of the International Conference on 3D Vision*, pages 517–526, 2017. [6](#)
- [70] Matthias Vestner, Roei Litman, Emanuele Rodola, Alex Bronstein, and Daniel Cremers. Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3327–3336, 2017. [2](#), [6](#)
- [71] Wolfram Von Funck, Holger Theisel, and Hans-Peter Seidel. Vector field based shape deformations. *ACM Transactions on Graphics*, 25(3):1118–1125, 2006. [2](#)
- [72] Yuexuan Wang, B Liu, K Zhou, and Yu Tong. Vector field map representation for near conformal surface correspondence. In *Computer Graphics Forum*, volume 37, pages 72–83. Wiley Online Library, 2018. [2](#)
- [73] Rui Xiang, Rongjie Lai, and Hongkai Zhao. A dual iterative refinement method for non-rigid shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15930–15939, 2021. [2](#), [6](#)