

# Resonant Deformable Matching: Simultaneous Registration and Reconstruction

John Corring<sup>(✉)</sup> and Anand Rangarajan

Department of Computer and Information Science and Engineering,  
University of Florida, Gainesville, USA  
{corring, anand}@cise.ufl.edu

**Abstract.** In the past decade we have seen the emergence of many efficient algorithms for estimating non-rigid deformations registering a template to target features. Registration of density functions is particularly popular. In contrast to the success enjoyed by the density function representation, we have not seen similar success with the signed distance function representation. Resonant deformable matching (RDM) simultaneously estimates a non-rigid deformation and a set of unknown target normal directions by registering fields comprising signed distance and probability density information. Resonance occurs as the reconstruction estimate comes into agreement with the registered template. We perform experiments probing two problems: point-set registration and normal estimation. RDM compares favorably to top tier point registration and graph algorithms in terms of registration and reconstruction metrics.

## 1 Introduction

Many problems in computer vision require us to determine *correspondences* between similar sets of features. However, we are often faced with scenarios where it is very difficult to even define what a correspondence between two objects should be—no natural map, moreover bijection, may exist at all—often due to mismatched representations. Work focused on determining point correspondences for matching organized features has been abundant, as we highlight below, but there remains a clear need for handling mismatched representations. This work provides a solution for the mismatched case where the template consists of oriented points and the target consists of points, under the assumption that both template and target are drawn from outlines of shapes.

Given a set of point-features, correspondences can be obtained via *registration*. In this approach, sparse feature sets are often first converted into scalar field representations. Then non-rigid matching of the template field with that of the target yields dense point to point correspondences. For example, point features can be converted into a probability density function representation [1–3]. Registration is obtained by deforming the template density onto the target using regularized spatial deformations [4, 5]. Implicit shape representations are not restricted to probability density functions estimated from sets of features.

Implicit representations abound in the literature [6–8]. The signed distance function (SDF) is an example in which the sign encodes interior/exterior properties with the absolute value encoding the distance to the nearest point in the set of curves (surfaces) [9–11]. Contrast this with the unsigned distance function which lacks interior/exterior information. Surprisingly, there is little work on matching template and target SDFs. We address the technical reasons for this now.

The signed distance  $b_S : \mathbb{R}^d \rightarrow \mathbb{R}$  for an open set  $S$  satisfies  $|\nabla b_S| = 1$  :  $b_S|_{\partial S} = 0$  with  $b_S$  continuously differentiable across the set boundary. The first technical problem we encounter in matching SDFs is the choice of a distance measure between them. Far away from the shape boundary (in the *far field*) SDFs take large values. This renders many standard distances useless, like  $L^p$  or  $W^p$ . The second problem we encounter is that SDFs are usually not available in closed form, in sharp contrast to parametric density representations. This implies that closed form distances between SDFs are elusive. Third, note that matching is extremely difficult to perform within the space of SDFs. For  $\phi \in \mathcal{H}$  to maintain the properties of SDFs, we see that  $|(\nabla|_{\phi(x)} b_S) \phi'|_x| = 1$ , implying that  $\phi'|_x \in \mathbb{O}(d)$  for all  $x \in \mathbb{R}^d$ —restricting  $\mathcal{H}$  to rigid transformations. The difficulty of managing this constraint is related to the *reinitialization* problem in level-set methods [12–14] where  $\phi$  is the (instantaneous) motion of an interface represented by a level set function. The complex wave representation [15] of shape is a parametric representation with aspects of the signed distance that avoids these issues. When the target is given as points (from  $\partial S$ ) the problem of estimating the target SDF remains. Estimating an SDF from a sparse set of features is the *very difficult problem* of curve (or surface) **reconstruction** [16–19]. In this paper we use reconstruction in the following sense—consistent assignment of normal vectors to points. Dense normal estimation provides useful data for constructing a surface [20, 21]. The representation that we use allows us to construct surfaces via extraction of a zero level-set, as in Figs. 1 and 2.

Our approach integrates registration and reconstruction from a deformable template viewpoint, hence *Resonant Deformable Matching*. A complex scalar field representation is utilized wherein the squared **magnitude** of the field is proportional to the probability density whereas the **phase** of the field is related to the signed distance (corresponding to implicit curves in 2D and surfaces in 3D). Using a representation with both signed distance and probability information allows us to penalize geometric mismatches in a weighted fashion that preserves the advantages of density fields. During registration, the target signed distance information is *simultaneously* estimated along with the spatial deformation. The advantages of our approach relative to previous work are as follows: (i) our approach employs both probability density and signed distance information for improved registration; (ii) we derive a closed form distance between template and target functions that distinguishes oriented point-sets in the feature space; (iii) reconstruction of the target is achieved during the matching process, which can lead to improved registration; (iv) RDM outperforms competing point and field-based methods on registration and template-based normal estimation.

## 2 Previous Work

Field based methods make a specific choice of representation that is consistent among all point-sets and shapes being matched. Kernel Correlation [1] and `gmmreg` [2] use Parzen-window densities, employing correlation and  $L^2$ -distance objectives respectively. Matching distributions [22] allows singular measures to be matched. A crossover between the density and distance fields in [23] utilized distance transforms yielding a density field which is matched by a geodesic distance. In these works the unifying theme is a field that organizes in terms of *uncertainty*. SDFs, which organize in terms of *geometry*, have also been used for non-rigid registration in [10] where a variational approach leads to a *grid-based* PDE-method that performs distortion on the field. In [24] the signed distance values in the near-field of curves are viewed as random variables with a Gaussian-Kernel based probability model, and the mutual information is minimized using free-form deformations of the domain of the signed distance. Signed-distance registration techniques also have a natural outgrowth towards segmentation through shape-prior approaches [25, 26] to registering shapes to images.

Point-based methods feature explicit estimation of correspondence, possibly in a soft or probabilistic fashion. Coherent Point Drift (CPD) [27] and TPS-RPM [28] are two standard-bearers. TPS-RPM alternates between estimating the (soft) correspondence and a TPS deformation. CPD uses a similar formulation, but also imposes additional constraints (arising from motion coherence theory) on the deformation. RPM-LNS [29] imposes symmetric neighborhood structures to preserve local shape while allowing global deformation.

Graph matching methods have also been employed for point registration [30]. Local and global relations can be encoded in graphs, yielding a powerful structure for correspondence estimation. While graph matching is a computationally hard problem, algorithms for structured graphs and relaxation techniques show promise for point matching [31–33]. When a planar shape is available as a silhouette a cyclic graph emerges and elastic matching can be done quickly [34]. Manifolds induce *Laplace-Beltrami* eigenfunctions [35], providing a canonical basis from which to perform matching from a joint coordinate perspective [36] or a function mapping perspective [37]. These methods all rely on *equivalent organization of source and target*. While organization elevates the richness of the matching techniques available, it also presents a difficulty: these methods require a level footing between template and target. Estimating a graph or mesh from points can be very challenging.

Point feature organization can be viewed from many perspectives: computational geometric methods [17], psychological gestalt principles [38–40], clustering [41, 42], and level-set methods [18, 19] all organize points in some sense. Shape representations are typically chosen to engender a desired organizational aspect of shapes [43]. Through a multi-valued function or a distributional representation, different aspects of shapes can be embodied in fields that interact predictably [7, 8, 15]. These works provide a spectrum of organizational principles that can be used to temper the difficulty of the point matching problem. **In this work we obtain a reconstruction while matching**, which means that

no target structure needs to be estimated before matching. Few works touting simultaneous matching and reconstruction are currently available [44, 45].

### 3 Complex Wave Mixtures and Signed Distance Functions


In [15], the complex-valued field—or wave representation of shape was introduced. It was used to compress closed curves with PCA (exploiting linearity of the representation) and classify oriented point-sets using reconstruction error as a distance metric. Here, we prove interesting properties of this shape representation for the purposes of curve reconstruction and as a feature function before proceeding to showcasing simultaneous matching and reconstruction. Consider a point-set augmented with directional information at each point. That is, let  $\mathcal{S} = \{(m_a, \nu_a)\}_{a=1}^M$ , where  $\nu_a$  is a normal associated with the point  $m_a$ . We may use  $S$  to denote the set underlying the oriented point-set  $\mathcal{S}$ , with each  $m_a \in \partial S$  and  $\nu_a$  pointing in the outward direction from  $S$ . We refer to this as an *oriented point-set* [15, 20]. The complex field we use extends the standard Gaussian Parzen window density to a square-root of a density by using the normal information, written (unnormalized) as

$$\psi_{\mathcal{S}}(x) = \sum_{a=1}^M e^{-\frac{\|x-m_a\|^2}{2\sigma^2} + i\frac{\nu_a^T(x-m_a)}{\lambda}}. \quad (1)$$

$\lambda$  controls the frequency of the wave: the lower the value of  $\lambda$  the higher the spatial frequency. The wave oscillates along the normal near a point feature but integrates information from different wavefronts in the far field (near and far are a function of  $\sigma, \lambda$ ). The squared magnitude of  $\psi(x)$  encodes probability density information. Zero level-sets of the phase now carry *shape geometry* information.

The mixture in (1) has similarities to the Gabor filter or wavelet—well known to vision researchers and mathematicians [46–48]. This allows us to leverage the mathematical literature to prove useful properties of this representation, such as proof of injectivity below which follows a similar argument for a related Gabor system [48]. Gabor systems are families of *time-frequency* translates of an admissible function. The use of Gabor wavelets for function approximation has been studied in the past [46, 47], but the connection linking the phase of a square-root density estimator to signed distance functions (and static Hamilton-Jacobi equations [49]) is subtle and less well-known.

In contrast to the unsigned distance function, the signed distance is smooth across shape boundaries (providing a reconstruction) with the sign of the distance indicating whether a location is inside or outside the shape. When we fit Parzen window density estimators to a point-set, we can obtain an approximate unsigned distance function at every point. The relation  $G(x) \approx C_R e^{-\frac{R^2(x)}{2\sigma^2}}$  holds (with  $C_R$  being a normalization constant), where the approximate unsigned distance function  $R(x)$  approaches the true distance pointwise as  $\sigma$  decreases toward zero [50].



**Fig. 1.** An example of **surface reconstruction** by RDM. (a) and (b) are the inputs to RDM, the points in (a) are the target and the oriented point-set in (b) is the template. (c) shows the estimated normal vectors from *RDM* and the *true* normal vectors. 99% of the normal vectors are recovered to within  $\pi/4$  angular error. (d) shows the reconstructed surface (the zero level set of the phase of  $\psi$ ) from the true normal vectors and (e) the reconstructed surface from RDM. The protrusions from the ears are due to mis-oriented normals in the high curvature area near the ear lobe.

For oriented point-sets the relation is

$$\psi_S(x) \approx \Psi_S = e^{-\frac{b_S^2(x)}{2\sigma^2} + i\frac{b_S(x)}{\lambda}} \quad (2)$$

where  $b_S(x)$  is the SDF. For a fixed  $S$ , the approximation becomes more accurate as  $\sigma, \lambda \rightarrow 0$ . Note that the magnitude is agnostic to the sign of the distance whereas the phase carries the sign but is modular due to the wrapped nature of the phase. We refer to this as a **modular distance function** since phase unwrapping is required to obtain a global SDF. Note that *we do not require or use phase unwrapping* in this paper. Two key advantages to using the modular distance function in lieu of the signed distance function are: (i) the modulus decays as we approach the far-field, handling the far-field issue mentioned above, (ii) Equation (2) allows us to derive distances in closed form.

### 3.1 Relationship Between Signed Distances and Complex Wave Mixtures

To solidify the claim made in Eq. (2), first note that  $\|\Psi_S\|^2 < \infty$ .  $|\Psi_S(x)|^2 = |\exp\{-\frac{b_S^2(x)}{2\sigma^2} + ib_S(x)/\lambda\}|^2$  is dominated by its concave envelope  $\bar{\Psi}_S$ , which has  $\|\bar{\Psi}_S\|^2 \leq \frac{((2\pi\sigma^2)^{d/2+1}\pi^{d/2} \text{diam}(S)^d)}{\Gamma(\frac{d}{2}+1)}$ , by an application of volumes of revolution.  $\|\Psi_S\| \geq (2\pi\sigma^2)^{d/2}$  as  $d(x, p) > d(x, S)$  for all  $p \in S$ .

Then, note that as  $\sigma \rightarrow 0$  that  $\langle \exp\{-\frac{\|x-m\|^2}{2\sigma^2} + i\frac{v^T(x-m)}{\lambda}\}, \Psi^{\sigma, \lambda} \rangle \rightarrow 0$  whenever  $m \notin \partial S$ . And as  $\lambda \rightarrow 0$  destructive interference causes  $\langle \exp\{-\frac{\|x-m\|^2}{2\sigma^2} + i\frac{v^T(x-m)}{\lambda}\}, \Psi^{\sigma, \lambda} \rangle \rightarrow 0$  by an application of the stationary phase expansion [51]. This means that as  $\sigma, \lambda$  shrink, the only significant coefficients of the Gabor Transform of  $\Psi_S$  come from atoms centered on the boundary, oriented in the outward normal direction. More evidence supporting the substitution of the signed distance by the complex wave mixture is provided in Sect. 4.2.

### 3.2 An Embedding Theorem for Complex Wave Mixtures

In some contexts, invariance of representation is desirable [36, 37]. For the purposes of *deformable matching*, however, having a 1-to-1 mapping between the point features and function representation is a prerequisite for employing distances as objective functions: if a feature function is not injective, it is possible that two non-registered point-sets result in the same feature functions, with zero distance between them. This is precluded in the complex wave representation. Note that this injectivity was not furnished in [15].

**Theorem 1.**  *$\psi$ . is an injective map from finite sets of oriented points to  $L^2$ . Any metric on  $L^2$  distinguishes oriented point sets under this representation.*

*Proof.* Let  $\mathcal{A} = \{(m_a, \nu_a)\}_{a=1}^A$ ,  $\mathcal{B} = \{(q_b, \omega_b)\}_{b=1}^B$  be distinct oriented point-sets. We will show that  $\psi_{\mathcal{A}} - \psi_{\mathcal{B}}$  is not identically zero. Suppose that  $m_1$  (a location in  $\mathcal{A}$ , with index 1 by reordering) is on the convex hull of  $K = \{m_a\}_{a=1}^A \cup \{q_b\}_{b=1}^B$ . Without loss of generality assume  $m_1 = 0$ . Let  $\mathcal{C} = \mathcal{A} \cup \mathcal{B} \setminus (m_1, \nu_1)$ . Then

$$\psi_{\mathcal{A}} - \psi_{\mathcal{B}} = \exp\left\{-\frac{\|x\|^2}{2\sigma^2}\right\} \left( \exp\left\{i\frac{\nu_1^T x}{\lambda}\right\} + \sum_{(r,\gamma) \in \mathcal{C}} h_{(r,\gamma)}(x) \exp\left\{\frac{x^T r}{\sigma^2}\right\} \right) \quad (3)$$

where each  $h_{(r,\gamma)} = [-1]^{(r,\gamma) \in \mathcal{B}} \exp\left\{-\frac{\|r\|^2}{2\sigma^2} + i\frac{\gamma^T(x-r)}{\lambda}\right\}$ . Since  $m_1$  is on the convex hull, there is a ray  $\{\kappa p\}_{\kappa > 0}$  in the Voronoi cell of  $m_1$  (relative to  $K$ ). So there

is a  $\kappa$  sufficiently large, so that  $\left| \sum_{(r,\gamma) \in \mathcal{A} \cup \mathcal{B} \setminus \{m_1, \nu_1\}} h_{(r,\gamma)}(x) \exp\left\{\frac{\kappa p^T r}{\sigma^2}\right\} \right| < \epsilon/2$ ,

so  $|\psi_{\mathcal{A}}(\kappa p) - \psi_{\mathcal{B}}(\kappa p)|^2 > \exp\left\{-\left\|\frac{\kappa p}{\sigma}\right\|^2\right\}(1 - \epsilon) > 0$ . If an oriented point-set has multiple oriented points with the same location (but distinct normals at these oriented points) we can use the injectivity of the Fourier Transform [52] to show that the sum of trigonometric polynomials (for the duplicated locations) is non-vanishing. Thus, the above argument holds even in that case. If  $d$  is a metric on  $L^2$  it is nonzero on distinct functions, distinguishing oriented point-sets.  $\square$

## 4 Complex Wave Registration and Normal Estimation

In registration, we seek a transformation of the template objects onto the target objects. We denote the transformation of the positions as  $\{\phi(m_a)\}_{a=1}^M$  where  $\phi \in \mathcal{H}$  is an element of the set of non-rigid transformations (assumed to be from the thin-plate spline family in the experiments). We depart from standard registration techniques as in our case the transformation of not only the template centers  $\{m_a\}_{a=1}^M$ , but also the template normals  $\{\nu_a\}_{a=1}^M$ , is carried out under the action of  $\phi$ . The appropriate transformation is the Jacobian,  $\phi'$ , of the deformation  $\phi$ . Note that  $\phi' : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is the derivative of the deformation with respect to the spatial variable, *not the parameters of the deformation*.  $\phi$  acts on  $(m_a, \nu_a)$  by


$$\phi \cdot (m_a, \nu_a) = (\phi(m_a), \phi'|_{m_a} \nu_a), \quad \text{with} \quad (\phi'|_{m_a})_{i,j} = \left. \frac{\partial \phi^{(i)}}{\partial x_j} \right|_{m_a}, \quad (4)$$

and  $\phi \cdot \mathcal{S} = \{\phi \cdot (m_a, \nu_a)\}_{a=1}^N$ . We can write the transformed template as

$$\psi_{\phi \cdot \mathcal{S}}(x) = \sum_{a=1}^M e^{-\frac{\|x - \phi(m_a)\|^2}{2\sigma^2}} e^{i \frac{(\phi' | m_a \nu_a)^T (x - \phi(m_a))}{\lambda}}. \quad (5)$$

Note that the centers and normals have been transformed via the action of the non-rigid deformation  $\phi$  but the location variable  $x$  remains intact. This allows us to define a distance between template and target functions in terms of a feature-space domain integral, which we will minimize w.r.t.  $\phi$ . Note that we actually minimize a regularized version of this distance since large deformations can bring very different point-sets into register. We discuss the specifics of the distance below. First, we address the representational mismatch.

#### 4.1 Introducing Normal Variables for the Target



**Fig. 2.** An example of **curve reconstruction** by RDM. **Left:** The target points are shown in black  $\times$ 's, with the level-sets of the unsigned distance function shown as contours. **Right:** After RDM estimates target normal vectors, the level sets of the phase of  $\psi_{\mathcal{T}(W)}$  are shown. Abutting point-sets make this particular reconstruction problem difficult, because choosing opposite orientations for nearby points goes against most normal estimation regularizers. See also Sect. 5.2.

Oriented point-set matching assumes an additional feature: normal directions for each point of the *template* and the *target*. *Template normals* are estimated offline: a standard approach to estimation involves the fitting of curves and surfaces to the template features followed by a sampling of the curves (or surfaces) into an oriented point-set. We assume template curves (surfaces) do not self-intersect in order to preserve normal uniqueness. This leaves the target normals. To recover a reconstruction of the surface underlying the target point-set, we augment the objective with variables for the target normals  $W = \{\omega_i\}_{i=1}^N$ .

This normal estimation component has **no counterpart** in the density matching literature. Adding these parameters does not increase overfitting of  $\phi$ , since the parameterization of  $\phi$  is independent of the normals. We discuss the effect of this simultaneous fitting and estimating below.

To summarize, first we assume that we are in possession of an oriented template point-set  $\mathcal{S}$ . This template point-set is deformed onto an un-oriented (and un-organized) target point-set  $\mathcal{T}$  via the action of a non-rigid deformation minimizing (8). Since the target point-set is un-organized, we estimate a set of target normals at each point during the matching process, thereby obtaining an oriented target point-set denoted  $\mathcal{T}(W) = \{(q_i, \omega_i)\}_{i=1}^N$ . This simultaneous matching and reconstruction approach is enabled by a closed-form distance measure between template and target complex wave mixtures.

### 4.2 Choosing a Suitable Distance Function

Minimizing  $D(\psi_{\phi,S}, \psi_{\mathcal{T}(W)})$  w.r.t.  $\phi$  and  $W$  is a difficult optimization problem regardless of the choice of  $D$ —symmetries and local minima stand in the way. In the literature, we have seen different choices (geodesic distance on  $S^\infty$  [23], Cauchy-Schwarz [3], Kullback-Leibler [53]) as well as different choices for the Parzen kernel (Gaussian [2], Schrödinger [50]). This cross product space of distances, kernels and algorithms is an active area of research.

We use the  $L^2$  distance. The  $L^2$  distance for density function registration was studied in [2] as a specialization of the density power divergence [54]. It strikes a balance between robustness to sampling and computability.  $L^2$  is robust to small Gaussian perturbations in the location parameters:  $\mathbb{E}_\delta [|\psi_S - \psi_{S+\delta}|^2] \rightarrow 0$  as  $\text{var}(\delta) \rightarrow 0$  by Fubini’s theorem [52]. While behavior under resampling is harder to examine theoretically, a certain amount of robustness is borne out in Sect. 5.2. Now, note that if  $\|\psi_S - \psi_{\mathcal{T}}\|^2 < \epsilon$  then  $\|\psi_S/C_S - \psi_{\mathcal{T}}/C_{\mathcal{T}}\|^2 < \epsilon'$  ( $C_S, C_{\mathcal{T}}$  are normalization constants), and so  $1 - \epsilon'/2 < |\langle \psi_S/C_S, \psi_{\mathcal{T}}/C_{\mathcal{T}} \rangle|$ . Continuing the line of reasoning in Sect. 3.1, if we pass to the normalized versions of  $\Psi_S$  and  $\Psi_{\mathcal{T}}$  then we see that the signed distances  $b_S$  and  $b_{\mathcal{T}}$  must be approximately aligned in the near field (of  $S$  and  $\mathcal{T}$ ). Otherwise, destructive interference would cause cancellations in the product field, decreasing the correlation.

We evaluate the squared  $L_2$  distance between the deformed template and target complex wave mixtures, subsequently minimized w.r.t. the unknown matching and normal parameters. The action of the spatial non-rigid deformation results in deformed template points *and* normals. Contrast this to the typical density matching situation in which only the template points are deformed. The squared  $L_2$  distance between the deformed template  $\psi_{\phi,S}$  and target  $\psi_{\mathcal{T}(W)}$ ,  $D(\psi_{\phi,S}, \psi_{\mathcal{T}(W)})$ , is given by

$$\int_{\mathbb{R}^D} \left| \sum_{a=1}^M e^{-\frac{\|x-\phi(m_a)\|^2}{2\sigma^2} + i\frac{\phi \cdot \nu_a^T(x-\phi(m_a))}{\lambda}} - \sum_{b=1}^N e^{-\frac{\|x-q_b\|^2}{2\sigma^2} + i\frac{\omega_b^T(x-q_b)}{\lambda}} \right|^2 dx \tag{6}$$

where the target wave mixture has been specified for the oriented point-set  $\mathcal{T}(W) = \{(q_b, \omega_b)\}_{b=1}^N$ . Note that their cardinalities  $M$  and  $N$  can differ. When evaluating the  $L_2$  distance, we are required to determine the inner product between terms which may differ in their location and frequency (with common scale and frequency parameters  $\sigma$  and  $\lambda$  respectively).

The inner product, denoted  $I_{(m,\nu)}^{(q,\omega)} = \langle \psi_{(m,\nu)}, \psi_{(q,\omega)} \rangle$ , is given by the integral

$$\int_{\mathbb{R}^D} e^{-\frac{\|x-m\|^2 - \|x-q\|^2}{2\sigma^2} + i\frac{\nu^T(x-m) - \omega^T(x-q)}{\lambda}} dx = \frac{e^{-\frac{\|m-q\|^2}{4\sigma^2} - \frac{\sigma^2\|\nu-\omega\|^2}{4\lambda^2} + i\frac{(\nu+\omega)^T(m-q)}{2\lambda}}}{(2\pi\sigma^2)^{\frac{D}{2}}}. \tag{7}$$

If  $m = q$ , then the spatial term goes to 1 and weights the Gaussian corresponding to the frequency term heavily. If  $m \approx q + \delta\omega^\perp$  this weighting is dampened, but we obtain constructive interference provided the normals  $\nu$  and  $\omega$  are aligned. When the normals are not aligned, we get destructive interference. This can either



force the normal estimates in line with the template or influence the template movement, and prevent unnecessary local rotation of the template normals.

The objective function minimized in this work is therefore

$$\mathcal{E}(\phi, W) = D(\psi_{\phi, \mathcal{S}}, \psi_{\mathcal{T}(W)}) + \beta L(\phi). \quad (8)$$

In (8),  $\phi$  and  $W$  are the desired spatial deformation and target normal set respectively. Additionally,  $\beta$  is a regularization parameter and  $L$  a suitable spline regularization (chosen to be the thin plate spline bending energy). Assuming a set of fixed centers  $\{p_b\}_{b=1}^P$  on the template, the thin plate spline [4] maps the location  $x \in \mathbb{R}^D$  to the location  $A(x) + \sum_{b=1}^P \mathbf{C}_b^T K(x - p_b)$  where  $A$  is an affine transformation,  $K$  is the thin-plate spline kernel and  $\{\mathbf{C}_b\}_{b=1}^P$  is the set of spline parameters. The mapping is linear in each  $\mathbf{C}_b$  and  $A$  and therefore so is  $\phi'$ . The regularization term in (8) becomes  $\beta \text{tr}(\mathbf{C}^T \mathbf{K} \mathbf{C})$  where  $\mathbf{C}$  is the  $P \times D$  matrix of spline coefficients, and  $\mathbf{K}_{ij} = K(p_i - p_j)$  is the Gram matrix of the set of control points.

We can characterize the asymptotic behavior of our matching objective. Examining the wave mixture, we see that the wave flattens out as  $\lambda \rightarrow \infty$ —eventually approaching 1. This intuitively results in the Gabor tending to the Gaussian. This is made more precise in the following Proposition, essentially a consequence of the dominated convergence theorem [52].

**Proposition 1.** *Let  $\{m_a, \nu_a\}_{a=1}^M, \{q_b, \omega_b\}_{b=1}^N$  be a pair of oriented point-sets. As  $\lambda \rightarrow \infty$  the objective function (6) converges to*

$$\int_{\mathbb{R}^D} \left| \sum_{a=1}^M e^{-\frac{\|x - m_a\|^2}{2\sigma^2}} - \sum_{i=1}^N e^{-\frac{\|x - q_i\|^2}{2\sigma^2}} \right|^2 dx. \quad (9)$$

with  $\phi \cdot$  acting by restriction to the first coordinate of Eq. (4).

### 4.3 Gradient Computation and Optimization Details

We derive the gradient for the TPS parameterization discussed above. The penalty term is easy to differentiate with respect to  $\mathbf{C}$ :

$$\partial_{\mathbf{C}} \beta \text{tr} \mathbf{C}^T \mathbf{K} \mathbf{C} = 2\beta \mathbf{K} \mathbf{C} \quad (10)$$

by differentiating the trace and using the symmetry of  $\mathbf{K}$ . The derivative of the inner product with respect to the parameters  $\mathbf{C}$  is

$$\partial_{\mathbf{C}} I_{(q, \omega)}^{\phi_{\mathbf{C}} \cdot (m, \nu)} = \partial_{\phi_{\mathbf{C}}(m)} I_{(q, \omega)}^{\phi_{\mathbf{C}} \cdot (m, \nu)} \frac{\partial \phi_{\mathbf{C}}(m)}{\partial \mathbf{C}} + \partial_{\phi_{\mathbf{C}} \cdot \nu} I_{(q, \omega)}^{\phi_{\mathbf{C}} \cdot (m, \nu)} \frac{\partial [\phi_{\mathbf{C}} \cdot \nu]}{\partial \mathbf{C}}. \quad (11)$$

Recall that  $\phi_{\mathbf{C}}$  acts on the normal  $\nu$  at point  $m$  by  $\phi_{\mathbf{C}} \cdot \nu = \phi'_{\mathbf{C}}|_m \nu$  [55] where  $\phi'_{\mathbf{C}}|_m$  is the Jacobian at  $m$ . Note that now  $\frac{\partial}{\partial \mathbf{C}} \phi_{\mathbf{C}}$  is the derivative w.r.t. the TPS parameters, *not the spatial variable*. We use  $\partial \cdot$  and  $\frac{\partial}{\partial \cdot}$  interchangeably. Let  $\mathbf{R} \in \mathbb{R}^{P \times N}$  be given by  $\mathbf{R}_{ij} = K(p_i - m_j)$ , the kernel matrix pairing template

and control points. Then  $\frac{\partial}{\partial \mathbf{C}}[\phi_{\mathbf{C}}(m_j)]^a = \mathbf{R}_j e_a$  (the superscript  $a$  indicates the  $a^{\text{th}}$  coordinate,  $e_a \in \mathbb{R}^d$  the  $a^{\text{th}}$  basis row vector) with  $\mathbf{R}_j$  the  $j^{\text{th}}$  column of  $\mathbf{R}$ . Differentiating,

$$\frac{\partial[\phi_{\mathbf{C}}(m_j)]^a}{\partial \mathbf{C}} = \mathbf{R}_j e_a, \quad \left[ \frac{\partial I_{(q,\omega)}^{\phi_{\mathbf{C}} \cdot (m,\nu)}}{\partial \phi_{\mathbf{C}}(m)} \right]^a = \left[ -\frac{\phi_{\mathbf{C}}(m) - q}{2\sigma^2} + i \frac{\phi_{\mathbf{C}} \cdot \nu + \omega}{2\lambda} \right]^a I_{(q,\omega)}^{\phi_{\mathbf{C}} \cdot (m,\nu)}. \quad (12)$$

When applying the entire gradient update (through all points), this is simply an outer product of the derivatives of the inner product and  $\mathbf{R}$ .

The second term in Eq. (11) is not typically seen in the point matching literature. We must differentiate  $\phi'_{\mathbf{C}}|_{m_j} \nu$  with respect to  $\mathbf{C}$ , where  $'$  denotes differentiation with respect to the domain variable. Denote by  $[\mathbf{R}_j']^k$  the matrix of derivatives in the  $k^{\text{th}}$  coordinate of the kernel function at each point in the template set. Then

$$[\phi'_{\mathbf{C}}|_{m_j} \nu_j]^a = [[\mathbf{R}_j']^a \mathbf{C}] \nu_j, \quad \text{and so} \quad \partial_{\mathbf{C}} [\phi'_{\mathbf{C}}|_{m_j} \nu_j]^a = ([\mathbf{R}_j']^a) \nu_j^T,$$

by treating  $\mathbf{C}$  as a scalar-valued form acting on  $([\mathbf{R}_j']^a, \nu)$ . So the second term in Eq. (11) is

$$\frac{\partial I_{(q,\omega)}^{\phi_{\mathbf{C}}(m_j, \nu_j)}}{\partial \mathbf{C}} = I_{(q,\omega)}^{\phi_{\mathbf{C}}(m_j, \nu_j)} \sum_{a=1}^D \left( \left[ -\sigma^2 \frac{\phi_{\mathbf{C}} \cdot \nu_j - \omega}{2\lambda^2} + i \frac{q - m}{2\lambda} \right]^a [\mathbf{R}_j']^a \right) \nu_j^T. \quad (13)$$

The descent direction for the TPS parameters is given by

$$\begin{aligned} \nabla_{\mathbf{C}} D = & 2 \sum_{i=1}^M \sum_{j=1}^M \left[ \partial_{\phi_{\mathbf{C}}(m_i)} I_{(m_j, \nu_j)}^{\phi_{\mathbf{C}}(m_i, \nu_i)} \frac{\partial \phi_{\mathbf{C}}(m_i)}{\partial \mathbf{C}} + \partial_{\phi_{\mathbf{C}} \cdot \nu_i} I_{(m_j, \nu_j)}^{\phi_{\mathbf{C}}(m_i, \nu_i)} \frac{\partial [\phi_{\mathbf{C}} \cdot \nu_i]}{\partial \mathbf{C}} \right] \\ & - 2 \sum_{i=1}^M \sum_{j=1}^N \left[ \partial_{\phi_{\mathbf{C}}(m_i)} I_{(q_j, \omega_j)}^{\phi_{\mathbf{C}}(m_i, \nu_i)} \frac{\partial \phi_{\mathbf{C}}(m_i)}{\partial \mathbf{C}} + \partial_{\phi_{\mathbf{C}} \cdot \nu_i} I_{(q_j, \omega_j)}^{\phi_{\mathbf{C}}(m_i, \nu_i)} \frac{\partial [\phi_{\mathbf{C}} \cdot \nu_i]}{\partial \mathbf{C}} \right]. \end{aligned} \quad (14)$$

To complete the picture, we return to the principal themes of this work—simultaneous registration and reconstruction. Recall that we began by pointing out that there was a paucity of literature on non-rigid SDF matching in comparison to density matching. We zeroed in on the difficulty of estimating SDFs as the main reason. Rather than estimate an SDF for the target point-set with the aid of a deformed template, we chose to estimate target normals as we deformed the template. To do this, we apply the descent direction for each  $\omega_i$  in terms of combinations of  $\partial_{\omega_i} I_{(q_j, \omega_j)}^{\phi_{\mathbf{C}}(m_i, \nu_i)}$  during each round. Further details of the optimization algorithm are provided below. To obtain the signed distance from these normals one may use previously developed methods (e.g. [20, 21]) or use the phase of the resulting wave-function directly (see Figs. 1 and 2). The result is an integrated probability density and SDF approach to simultaneous deformable template matching and multiple curve (or surface) reconstruction.

## 5 Experiments

We compare with the state of the art in density field matching (such as `gmmreg`, abbreviated to GMM) [2], generalized function matching (diffeomorphic measure matching abbreviated DIFF) [22], point-based matching (CPD) [27], and graph-matching (FGM-U) [32]. While other methods [10, 56] are appropriate for further comparison, handling the asymmetry in representation is not possible in their current formulation. The corresponding results are indicated by the appropriate marker and color combinations (see legend in Fig. 5). We investigate the performance of RDM on a variety of datasets and conditions, outlined below.

**Parameter Configuration:** The following parameters are used in the experiments below unless explicitly stated otherwise. In all of the following experiments the data sets are scaled to  $[0, 1]^2$  before matching. We use a common initial scale parameter ( $\sigma = .1$  for both GMM and RDM) throughout. Two initializations with decreasing  $\sigma$  values are used for both GMM and RDM (DIFF uses 2 reinitializations). CPD estimates scaled progressively during matching, so no reinitialization is performed. For RDM  $\beta = .0075$  for 2d, for GMM  $\beta = .01$ , for CPD  $\beta = 2, \lambda = 2$  (different parameters). For FGM-U Delaunay triangulation was used for graph construction, and 101 iterations at 100 scales were executed for path following. GMM, DIFF, and RDM use MATLAB<sup>®</sup>'s function `fminunc` (set to use a `quasi-newton` solver—BFGS iterations) for optimization. Unless otherwise noted, the error measure is mean Euclidean distance to correspondent.

### 5.1 Synthetic Normal Recovery, Warps, and Oclusions

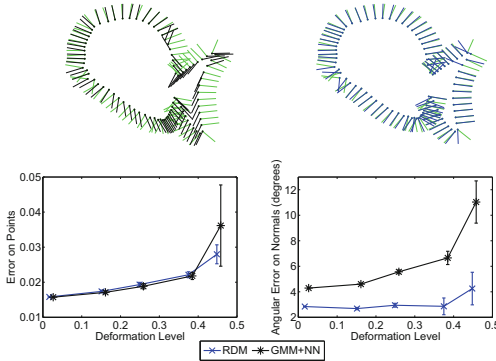
This subsection consists of two sets of experiments. First, we compared our algorithm to a pipeline approach to normal estimation:

1. Register a template point-set to a target by an estimated deformation.
2. Let the deformation act on the normal vectors of the template.
3. Use a nearest-neighbor approach to infer normal vectors onto the corresponding points in the target set.

We used GMM as the matching algorithm. A single curve (a body curve consisting of 80 points) from the multi-curve GatorBait dataset was used. For synthetic deformation, a diffeomorphism is fit to point perturbations by solving a 3d flow problem [5]. Therefore relative normal orientations are preserved. The deformation level corresponds to  $\sup_{x \in \mathbb{R}^D} \|\phi(x) - x\|$  (evaluated on the test points). Results are shown in Fig. 3. In this experiment GMM and RDM only used a single initialization. This experiment shows that a gain in normal recovery is obtained by using RDM instead of imposing template structure *after* matching.

In the second set of experiments, we tested the performance of our algorithm in the 2d and 3d synthetic settings against 4 other methods: CPD, DIFF, FGM-U, and GMM. For 2d we used a point-set consisting of 5 curves (see Fig. 2) from the aforementioned dataset, while for 3d the Stanford bunny and TOSCA datasets were used. We create the target by randomly perturbing points lying along a grid

and solving for a TPS with identity affine component. *No information about the target normal vectors is known beforehand.* After registration, the mean distance to the corresponding point, average error, is recorded. For the occlusion trials, an approximate fixed deformation level is used, .3 in norm for 2d and .15 in 3d. Robustness to outliers and noise is also studied. For these experiments FGM-U was run at 50 scales, due to runtime limitations. One can see from Fig. 4 that FGM-U struggles with nonrigid deformation. A plot showing the percentage of recovered normal orientations is included as well.

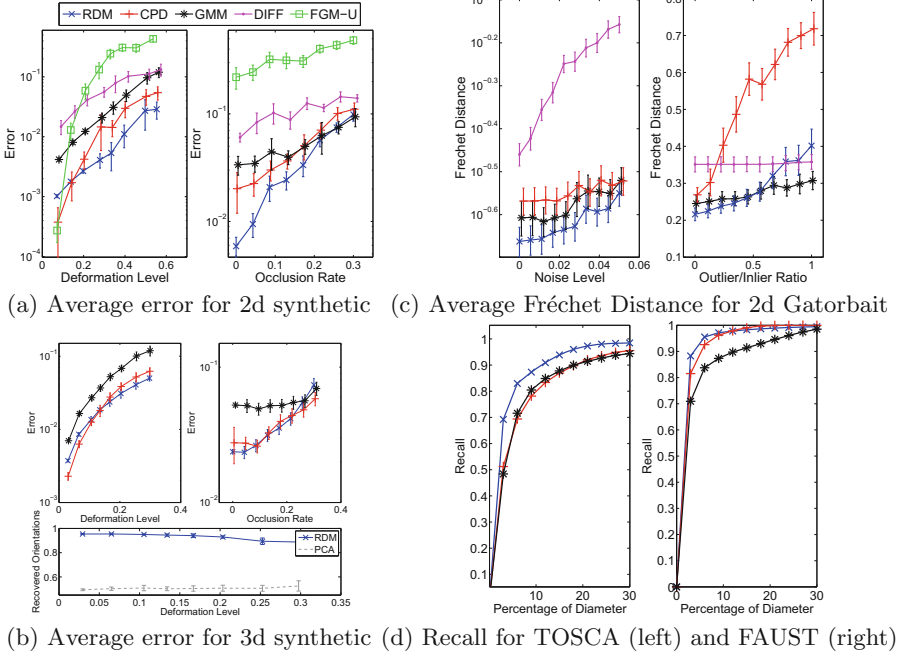


**Fig. 3. Top:** the *recovered normal vectors* for **GMM+NN** (left) and **RDM** (right) and the **True** normals (both) are attached at the points. **Bottom:** Average error and the median angle error between corresponding normal vectors for RDM and GMM+NN. 50 trials per level were performed.

## 5.2 Non-synthetic Matching Experiments

We perform intra-class matching experiments on the TOSCA [57], FAUST [58], and Gatorbait datasets. TOSCA and FAUST represent the 3d performance gauge on real matching experiments. We chose the GatorBait dataset because it has *multiple corresponding parts* which many 2d and 3d point-sets and meshes lack. The same statistic as above—average error to correspondent—is collected for the sets with known correspondence. We present recall (percentage of correct correspondences within a threshold) for matching pose 0 to poses 1, 2, 4, and 5 (smaller deformations) of the FAUST training registrations over all 10 subjects and present comparisons with GMM and CPD. For TOSCA we match the first cat, dog, and gorilla to the remaining poses. We have foregone benchmark comparisons here because in the large deformation regime extrinsic matching is prone to local minima, and we restrict the comparisons to relative performance among other extrinsic matching techniques. We use these datasets as a baseline for comparison with GMM and CPD.

The GatorBait dataset does not have known correspondences. Furthermore, it consists of nearly abutting curves (see Fig. 2)—organizing points into their




**Fig. 4.** Experimental comparison of RDM, GMM, CPD, DIFF, and FGM. (a) The GatorBait Dataset is deformed as explained in Sect. 5.1. The left plot shows robustness to moderate deformation levels and the right plot shows robustness to occlusion (dropping points in a randomly placed disc). (b) The same experiments are carried out in 3d on the Stanford Bunny. We also report the percentage of normal vectors recovered to within a cone of  $\pi/3$  radians. (c) The Fréchet distances (sum over the parts) between the registered template and the target curves are reported. On the left the target has added noise of the indicated standard deviation and on the right outliers are added. (d) Recall graphs for a subset (See Sect. 5.2) of TOSCA and FAUST.

appropriate curve components is made much harder by the existence of neighborhood points on different curves. The *Fréchet Distance* [59] between corresponding parts in the final registration and the target is recorded. This allows us to measure how accurately *each part* of the template is matched to the target. The first fish species is used as the template and matched to 23 other species. We also perturb the fish with noise and add outliers as uniformly drawn additional points. See Fig. 4 for results. For large 3d datasets, DIFF and FGM were found to be impractical from a runtime perspective (for a runtime comparison see Fig. 5). For the GatorBait dataset, FGM was not competitive.

### 5.3 CMU House Dataset

The CMU House dataset consists of a sequence of image frames and keypoints. The task is to perform point-matching and recover correspondences between

points. From a correspondence standpoint, FGM-U [32] with Delaunay triangulation (FGM-`del`) is the state of the art on this dataset. However, FGM is sensitive to the graph structure—with 2-nearest neighbors FGM’s performance suffers. Should a large set of correspondences be needed, graph matching becomes impractical—even the 30 correspondences here represent significant computational effort for graph matching. RDM would benefit from a denser set of key-points. To initialize normals for RDM, we extract the gradient of the image  $I$  at each of the keypoints in the frames. This is a departure from the usual consideration of ‘normals’—we sample a vector field ( $\nabla I$ ) at discrete points. We have not dealt with this situation explicitly in the text, but syntactically speaking, it is valid. It also represents a case where target structure is already (partially) present—for these experiments we provide RDM the true target normal.



AUC				
Algorithm	Frame 40	60	80	100
FGM- <code>del</code>	1.00	1.00	1.00	1.00
FGM-2NN	1.00	.867	.800	.500
RDM	.931	.871	.857	.833
CPD	.888	.819	.731	.681
GMM	.862	.795	.738	.671
DIFF	.836	.791	.688	.403

Runtimes (in Seconds)				
Algorithm	Frame 40	60	80	100
FGM- <code>del</code>	8.8	11	17	15
FGM-2NN	4.8	6.7	8.3	10
RDM	5.9	5.1	6.5	5.2
CPD	.32	.35	.44	.38
GMM	6.2	7.5	8.8	11
DIFF	12	11	12	15

**Fig. 5.** Recall graphs and area under the curve for the CMU House. For FGM, triangulation yields excellent matches but nearest-neighbor graphs are poor. All experiments run on an AMD X2 B22 with 8 Gb of RAM. We report AUC and runtimes in the table.

## 6 Conclusion and Future Work

Deformable template matching with RDM is done by minimizing the closed form squared  $L^2$  distance between template and target wave mixtures augmented with a standard regularization on the spatial transformation (done in practice via standard nonlinear optimization software implementing quasi-Newton methods). When only the template normals are available at runtime, they can be estimated for the target set during the registration. This provides normal estimates.

In this work we proved injectivity of the representation, derived the gradient term for optimization, showed that RDM can outperform standard normal

transfer by registration, and highlighted the registration accuracy of RDM. We plan to extend RDM to the case where neither point-set is accurately oriented and to evaluate the reconstruction accuracy of RDM against other unsupervised and semi-supervised methods in future work. We are also exploring alternative deformation models for the oriented point-set transformation setting.

## References

1. Tsin, Y., Kanade, T.: A correlation-based approach to robust point set registration. In: Pajdla, T., Matas, J. (eds.) ECCV 2004. LNCS, vol. 3023, pp. 558–569. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-24672-5\\_44](https://doi.org/10.1007/978-3-540-24672-5_44)
2. Jian, B., Vemuri, B.C.: Robust point set registration using Gaussian mixture models. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **33**(8), 1633–1645 (2011)
3. Hasanbelliu, E., Sanchez Giraldo, L., Principe, J.C.: Information theoretic shape matching. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **36**(12), 2436–2451 (2014)
4. Wahba, G.: *Spline Models for Observational Data*, vol. 59. SIAM, Philadelphia (1990)
5. Camion, V., Younes, L.: Geodesic interpolating splines. In: Figueiredo, M., Zerubia, J., Jain, A.K. (eds.) EMMCVPR 2001. LNCS, vol. 2134, pp. 513–527. Springer, Heidelberg (2001). doi:[10.1007/3-540-44745-8\\_34](https://doi.org/10.1007/3-540-44745-8_34)
6. Gorelick, L., Galun, M., Sharon, E., Basri, R., Brandt, A.: Shape representation and classification using the Poisson equation. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **28**(12), 1991–2005 (2006)
7. Charon, N., Trounev, A.: The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM J. Imaging Sci.* **6**(4), 2547–2580 (2013)
8. Guler, R., Tari, S., Unal, G.: Screened Poisson hyperfields for shape coding. *SIAM J. Imaging Sci.* **7**(4), 2558–2590 (2014)
9. Osher, S., Sethian, J.A.: Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. *J. Comput. Phys.* **79**(1), 12–49 (1988)
10. Paragios, N., Rousson, M., Ramesh, V.: Non-rigid registration using distance functions. *Comput. Vis. Image Underst. (CVIU)* **89**(2), 142–165 (2003)
11. Osher, S., Fedkiw, R.: *Level-Set Methods and Dynamic Implicit Surfaces*, vol. 153. Springer Science & Business Media, New York (2006)
12. Sussman, M., Smereka, P., Osher, S.: A level set approach for computing solutions to incompressible two-phase flow. *J. Comput. Phys.* **114**(1), 146–159 (1994)
13. Gomes, J., Faugeras, O.: Reconciling distance functions and level-sets. *J. Vis. Commun. Image Represent.* **11**(2), 209–223 (2000)
14. Estellers, V., Zosso, D., Lai, R., Osher, S., Thiran, J., Bresson, X.: An efficient algorithm for level-set method preserving distance function. *IEEE Trans. Image Process. (TIP)* **21**(12), 4722–4734 (2012)
15. Corring, J., Rangarajan, A.: Shape from phase: an integrated level-set and probability density shape representation. In: International Conference on Pattern Recognition (ICPR), IAPR, pp. 46–51 (2014)
16. Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W.: Surface reconstruction from unorganized points. *Spec. Interest Group Graph. Interact. Techn. (SIGGRAPH)* **26**(2), 71–78 (1992)

17. Edelsbrunner, H.: Shape reconstruction with Delaunay complex. In: Lucchesi, C.L., Moura, A.V. (eds.) LATIN 1998. LNCS, vol. 1380, pp. 119–132. Springer, Heidelberg (1998). doi:[10.1007/BFb0054315](https://doi.org/10.1007/BFb0054315)
18. Zhao, H.K., Osher, S., Merriman, B., Kang, M.: Implicit and nonparametric shape reconstruction from unorganized data using a variational level-set method. *Comput. Vis. Image Underst. (CVIU)* **80**(3), 295–314 (2000)
19. Mullen, P., de Goes, F., Desbrun, M., Cohen-Steiner, D., Alliez, P.: Signing the unsigned: Robust surface reconstruction from raw pointsets. *Comput. Graph. Forum* **29**(5), 1733–1741 (2010)
20. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. *Proceedings of the Fourth Eurographics Symposium on Geometry Processing (SGP)*, vol. 7, pp. 61–70 (2006)
21. Manson, J., Petrova, G., Schaefer, S.: Streaming surface reconstruction using wavelets. *Comput. Graph. Forum* **27**(5), 1411–1420 (2008)
22. Glaunes, J., Trounev, A., Younes, L.: Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub-manifolds matching. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, pp. 712–718. IEEE (2004)
23. Deng, Y., Rangarajan, A., Eisenschenk, S., Vemuri, B.C.: A Riemannian framework for matching point clouds represented by the Schrödinger distance transform. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3756–3761. IEEE (2014)
24. Huang, X., Paragios, N., Metaxas, D.N.: Shape registration in implicit spaces using information theory and free form deformations. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **28**(8), 1303–1318 (2006)
25. Leventon, M.E., Grimson, W.E.L., Faugeras, O.: Statistical shape influence in geodesic active contours. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 316–323. IEEE (2000)
26. Rousson, M., Paragios, N.: Shape priors for level set representations. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) *ECCV 2002*. LNCS, vol. 2351, pp. 78–92. Springer, Heidelberg (2002). doi:[10.1007/3-540-47967-8\\_6](https://doi.org/10.1007/3-540-47967-8_6)
27. Myronenko, A., Song, X.: Point set registration: coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **32**(12), 2262–2275 (2010)
28. Chui, H., Rangarajan, A.: A new point matching algorithm for non-rigid registration. *Comput. Vis. Image Underst. (CVIU)* **89**(2), 114–141 (2003)
29. Zheng, Y., Doermann, D.: Robust point matching for non-rigid shapes by preserving local neighborhood structures. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **28**(4), 643–649 (2006)
30. Gold, S., Rangarajan, A.: A graduated assignment algorithm for graph matching. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **18**(4), 377–388 (1996)
31. Zhou, F., de la Torre, F.: Factorized graph matching. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 127–134. IEEE (2012)
32. Zhou, F., de la Torre, F.: Deformable graph matching. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2922–2929. IEEE (2013)
33. Kezurer, I., Kovalsky, S.Z., Basri, R., Lipman, Y.: Tight relaxation of quadratic matching. *Comput. Graph. Forum* **34**(5), 115–128 (2015)
34. Schmidt, F.R., Farin, D., Cremers, D.: Fast matching of planar shapes in sub-cubic runtime. In: *IEEE International Conference on Computer Vision (ICCV)*, pp. 1–6. IEEE (2007)



35. Lévy, B.: Laplace-Beltrami eigenfunctions: towards an algorithm that “understands” geometry. In: IEEE International Conference on Shape Modeling and Applications, pp. 13–21. IEEE (2006)
36. Kovnatsky, A., Bronstein, M.M., Bronstein, A.M., Glashoff, K., Kimmel, R.: Coupled quasi-harmonic bases. *Comput. Graph. Forum* **32**(2), 439–448 (2013)
37. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph. (TOG)* **31**(4), 30:1–30:11 (2012)
38. Guy, G., Medioni, G.: Inferring global perceptual contours from local features. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 786–787. IEEE (1993)
39. Mordohai, P., Medioni, G.: Tensor voting: a perceptual organization approach to computer vision and machine learning. *Synth. Lect. Image Video Multimedia Process.* **2**(1), 1–136 (2006)
40. Boyer, K.L., Sarkar, S.: *Perceptual Organization for Artificial Vision Systems*. vol. 546. Springer Science & Business Media (2012)
41. Pauly, M., Gross, M.H., Kobbelt, L.: Efficient simplification of point-sampled surfaces. In: IEEE Visualization, pp. 163–170 (2002)
42. Zelnik-Manor, L., Perona, P.: Self-tuning spectral clustering. In: *Advances in Neural Information Processing Systems (NIPS)*, pp. 1601–1608 (2004)
43. Kimia, B.B., Tannenbaum, A.R., Zucker, S.W.: Shapes, shocks, and deformations I: the components of two-dimensional shape and the reaction-diffusion space. *Int. J. Comput. Vis. (IJCV)* **15**(3), 189–224 (1995)
44. Mjolsness, E., Gindi, G., Anandan, P.: Optimization in model matching and perceptual organization. *Neural Comput.* **1**(2), 218–229 (1989)
45. Cho, M., Lee, K.M.: Progressive graph matching: Making a move of graphs via probabilistic voting. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 398–405. IEEE (2012)
46. Daubechies, I.: The wavelet transform, time-frequency localization and signal analysis. *IEEE Trans. Inf. Theory (IT)* **36**(5), 961–1005 (1990)
47. Lee, T.S.: Image representation using 2D Gabor wavelets. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **18**(10), 959–971 (1996)
48. Heil, C., Ramanathan, J., Topiwala, P.: Linear independence of time-frequency translates. *Proc. Am. Math. Soc. (AMS)* **124**(9), 2787–2795 (1996)
49. Aubert, G., Kornprobst, P.: *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*. Applied Mathematical Sciences, vol. 147. Springer (2006)
50. Sethi, M., Rangarajan, A., Gurumoorthy, K.: The Schrödinger distance transform (SDT) for point-sets and curves. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 198–205. IEEE (2012)
51. Wong, R.: *Asymptotic Approximations of Integrals*. Society for Industrial and Applied Mathematics (2001)
52. Folland, G.B.: *Real analysis*, 2nd edn. Pure and Applied Mathematics (New York). John Wiley & Sons Inc., New York (1999). Modern techniques and their applications, A Wiley-Interscience Publication
53. Wang, Y., Woods, K., McClain, M.: Information-theoretic matching of two point sets. *IEEE Trans. Image Process. (TIP)* **11**(8), 868–872 (2002)
54. Basu, A., Harris, I.R., Hjort, N.L., Jones, M.: Robust and efficient estimation by minimising a density power divergence. *Biometrika* **85**(3), 549–559 (1998)
55. Spivak, M.: *A comprehensive introduction to differential geometry*, 3rd edn, vol. 1–5. Publish or Perish (1999)

56. Siddiqi, K., Shokoufandeh, A., Dickinson, S.J., Zucker, S.W.: Shock graphs and shape matching. *Int. J. Comput. Vis. (IJCV)* **35**(1), 13–32 (1999)
57. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: *Numerical Geometry of Non-Rigid Shapes*. Monographs in Computer Science. Springer (2009)
58. Bogo, F., Romero, J., Loper, M., Black, M.J.: FAUST: Dataset and evaluation for 3D mesh registration. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3794–3801. IEEE, June 2014
59. Veltkamp, R.C.: Shape matching: similarity measures and algorithms. In: *IEEE International Conference on Shape Modeling and Applications (SMI)*, pp. 188–197. IEEE (2001)