

A Riemannian framework for matching point clouds represented by the Schrödinger distance transform

Yan Deng¹, Anand Rangarajan¹, Stephan Eisenschenk² and Baba C. Vemuri¹

¹Department of CISE, University of Florida, Gainesville, FL 32611, USA

²Department of Neurology, University of Florida, Gainesville, FL 32611, USA

¹{ydeng, anand, vemuri}@cise.ufl.edu ²stephan.eisenschenk@neurology.ufl.edu

Abstract

In this paper, we cast the problem of point cloud matching as a shape matching problem by transforming each of the given point clouds into a shape representation called the Schrödinger distance transform (SDT) representation. This is achieved by solving a static Schrödinger equation instead of the corresponding static Hamilton-Jacobi equation in this setting. The SDT representation is an analytic expression and following the theoretical physics literature, can be normalized to have unit L_2 norm—making it a square-root density, which is identified with a point on a unit Hilbert sphere, whose intrinsic geometry is fully known. The Fisher-Rao metric, a natural metric for the space of densities leads to analytic expressions for the geodesic distance between points on this sphere. In this paper, we use the well known Riemannian framework never before used for point cloud matching, and present a novel matching algorithm. We pose point set matching under rigid and non-rigid transformations in this framework and solve for the transformations using standard nonlinear optimization techniques. Finally, to evaluate the performance of our algorithm—dubbed SDTM—we present several synthetic and real data examples along with extensive comparisons to state-of-the-art techniques. The experiments show that our algorithm outperforms state-of-the-art point set registration algorithms on many quantitative metrics.

1. Introduction

With the advent of new sensing technologies such as the Kinect, fast laser scanners etc., sensing 3D objects to build models and print them is becoming popular. It is not impossible to envision scanning faces in 3D using fast scanners from vantage points, manipulating them in 3D and printing them for possible display as mantle pieces in homes. The problem of matching point clouds in 3D is an often encountered problem in the aforementioned applications.

Similar scenarios arise in 2D when matching feature point clouds sampled from 2D shapes. Within shape matching, two dominant trends have emerged in the past decade—the use of point-set density function representations on the one hand and the deployment of distance transforms (and level sets) for sets of curves embedded in Euclidean space on the other. Recently, the two representations have seen a rapprochement of a certain kind with the development of the Schrödinger distance transform (SDT) [14]—a wave function representation of a point-set with simultaneous and strong relationships to both density functions and distance transforms. In this work, the SDT wave function is a *square-root density* shape representation derived from the given point-clouds since it has a unit norm, thereby enabling the identification of the shape distance with the geodesic length (shortest path) between two shapes situated on the unit Hilbert sphere. This attractive and straightforward geometric metaphor fuels our approach to non-rigid shape and hence point cloud matching. Essentially, after placing both the model and scene point sets represented by the respective SDTs (with unit L_2 -norm) on the unit Hilbert sphere, we perform shape matching by minimizing the geodesic distance between a scene and a model (known in closed form on the unit Hilbert sphere) which is transformed via the action of a regularized warp. Note that knowing the Riemannian geometry of the unit Hilbert sphere fully affords significant computational advantages, hitherto never realized in the context of point cloud matching in published literature. Shape location “jitter” (due to noise etc.) is built in through the use of an uncertainty parameter in the SDT. The geodesic distance between two shapes placed on the unit Hilbert sphere is computed in closed form as a function of the (warped) model and scene locations. Standard nonlinear optimization approaches are then utilized to minimize this objective function. The result, as shown in this paper, is a state of the art point cloud matching method—dubbed SDTM—which outperforms competing methods in matching metrics.

Many previous approaches to shape matching formulate

the problem as point pattern matching wherein one simultaneously solves for the correspondence and shape deformation. Of these methods, the venerable Iterative Closest Point (ICP) method [2] by Besl *et al.* uses the nearest neighbor to assign correspondence. Granger *et al.* [7] formulate point set registration as a maximum likelihood problem by assuming that the posterior probability of the scene data conditioned on the model data is uniformly distributed with the correspondence information as the prior. Eschewing ICP, Chui and Rangarajan’s Robust Point Matching (RPM) algorithm in [5] sets up the correspondence problem as linear assignment and solves for a soft correspondence and deformation using an alternating algorithm. The Euclidean distance measure used in [2, 5, 7] is known to be vulnerable to the outliers. Instead of using the Euclidean metric, Belongie *et al.* in [1] compute a histogram based on the local structure information at each point called the “shape context”. This adds attribute information to the correspondence engine—based on a linear assignment. Another popular approach to point set matching/registration involves formulating the registration problem as a graph matching problem. In [4, 11], graph matching is used to find the point correspondences and the performance depends on the structural topology estimated from the point set. Glaunes *et al.* [6] rewrite the square integrable norm which is used to evaluate the similarity between two point sets as a linear combination of kernel functions. This work was extended in [17] using distribution field representations of surfaces, registered using diffeomorphic non-rigid transformations. The Coherent Point Drift (CPD) method [13] assumes that the warped model is sampled from a Gaussian mixture model with the scene data as the mean vector. This method works under the condition that the model data set is dense and the scene data set is sparse; otherwise, it is prone to failure. In [10, 18], the authors formulate the registration problem as a graph matching problem solved using relaxation labeling. In the former, a binary compatibility coefficient is used while the latter relaxes it to a “soft” non-binary compatibility. Jian and Vemuri [8] proposed a method wherein the L_2 distance between two Gaussian mixture representations of the point sets being matched was used to obtain the deformation, without needing to solve for the correspondence. The kernel correlation method of Tsin and Kanade [15] as well as other methods described in [2, 13] are special cases of the formulation in [8]. Recent work in [12] assumes a normal distribution for the noise between the warped model and the scene and uses the same L_2E estimator as in [8] to optimize the deformation parameters. The correspondence between the model and the scene is calculated iteratively from the shape context feature.

Prior to delving into the details, we now present the organization of the paper: in section 2, we recap the development of the Schrödinger distance transform (SDT) applied

to point sets and obtain the closed form expression for the SDT. The geodesic distance measure between two SDTs is introduced and the matching problem is expressed as a non-linear optimization problem. In section 3, an analytic expression of the geodesic distance measure is derived which is crucial to the optimization. In Section 4, we provide implementation details and validate our algorithm by comparing its performance with three state-of-the-art algorithms. We summarize the present work and discuss future plans in Section 5.

2. A Riemannian framework for matching two point clouds represented by SDTs

In this section, we first briefly recap the development leading to the Schrödinger Distance Transform (SDT). Then, we present closed-form expressions for the square-root density representation of a point-cloud which leads to the analytic expressions for the geodesic distance between the SDTs of the two point clouds in the Riemannian framework, in the following section. Subsequently, we present the shape matching objective function.

2.1. The Schrödinger Distance Transform (SDT) for point-sets

Given a point set $\{x_i \in \mathbb{R}^2, i = 1, 2, \dots, M\}$, the unsigned Euclidean distance transform is defined as $S(x) = \min_{\{x_i\}} \|x - x_i\|, x \in \mathbb{R}^2$. In this (Voronoi) representation, the zero level set of $S(x)$ is the set of Voronoi cell centers, while in general, the zero level-set represents the boundary of the shape from which the point-set has been sampled. The gradient of the distance transform $S(x)$ satisfies the following static Hamilton-Jacobi equation: $\|\nabla S(x)\| = 1$ every where (except at the Voronoi centers where $S(x) = 0$ and at the Voronoi boundaries). The static Schrödinger equation

$$-\tau^2 \nabla^2 \psi(x) + \psi(x) = \psi_0(x) \quad (1)$$

with an uncertainty parameter τ can approximate the above static Hamilton-Jacobi equation producing the exact distance transform as $\tau \rightarrow 0$. The field $\psi_0(x)$ in (1) is highly peaked at the point-set locations. For positive values of τ , the zero level-set obtained from (1) is no longer the set of Voronoi cells and can more closely approximate the shape boundaries (and this depends on the interplay between τ and $\psi_0(x)$). When we constrain $\psi(x)$ to have unit L_2 norm, the mapping from the distance transform $S(x)$ to a square root density is complete.

For the simple case where $\psi_0(x)$ behaves approximately like a set of square-root “delta functions” centered on the point-set locations, the static Schrödinger equation can be

solved in closed-form yielding (in 2D and in 3D)

$$\psi(x) \approx c \sum_{k=1}^M \exp \left\{ -\frac{\|x - x_k\|}{\tau} \right\} \quad (2)$$

where c is chosen to ensure that $\psi(x)$ has unit norm. Note that $\psi(x)$ is not a Gaussian mixture model where the exponent is the square of L_2 distance, instead, it is a square root density with the Euclidean distance in the exponent. As previously mentioned, the square root density $\psi(x)$ is a point located on a unit sphere in the Hilbert space. In this representation, two different shapes become two distinct points on the unit sphere with the geodesic length being the shortest path between them. The cornerstone of the shape matching work presented here is the assumption that minimizing the geodesic length under a suitable (regularized) non-rigid transformation brings the two shapes into correspondence. Once this is established, we envisage extensions to the groupwise setting with applications in atlas-guided segmentation.

2.2. The geodesic distance measure

The distance measure, a geodesic length, is (luckily for us) conceptually simple, analytically elegant and straightforward to derive. The Riemannian geometry of the unit Hilbert sphere is well known, the *Log* and *Exp* maps as well as the geodesic distance between two points on the unit Hilbert sphere are all in closed form. For the sake of space, we simply provide the equations for the geodesic distance on the sphere and refer the reader to [9] for other formulae which are not used here. The geodesic distance between a pair of SDT-derived square-root shape densities is the arc length of the (shorter) segment of a *great circle* with the two shapes as the end points of the segment. Formally, the geodesic distance between two SDTs ψ_1 and ψ_2 can be expressed as

$$\text{dist}(\psi_1, \psi_2) = \|\log_{\psi_1}(\psi_2)\|_{\psi_1} = \cos^{-1}\langle \psi_1, \psi_2 \rangle \quad (3)$$

where the Fisher-Rao metric $\langle \psi_1, \psi_2 \rangle$ is defined as $\int_{\Omega} \psi_1(x)\psi_2(x)dx$. Note that this distance on the unit Hilbert sphere for square root density representations has been widely used in the past for various tasks including shape matching but never for point cloud matching.

2.3. The SDTM objective function

Using the geodesic length given in the last section, the point-set registration problem can be formulated as following: given the model data set M and the scene data set S , we seek optimal transformation parameters \hat{w} which minimize the geodesic distance between the square root density of the warped model $T(M, w)$ and the square root density of the scene data set. Thus, the cost function becomes

$$\hat{w} = \arg \min_w \text{dist}(\psi(T(M, w)), \psi(S)) \quad (4)$$

Despite the dependence of the warped point-set on the shape deformation, the square root density of the warped data set $T(M, w)$ can be written such that $\psi(T(M, w))$ remains a point on the unit Hilbert sphere. Since all normalized square root densities of the warped model and scene are always situated on the unit Hilbert sphere, the optimization follows a path on the sphere allowing us to find a close to optimal deformation \hat{w} (depending on the optimization details). In the process, the geodesic distance between $\psi(T(M, w))$ and $\psi(S)$ is minimized while ensuring that the deformation is appropriately regularized.

We use thin-plate splines (TPS) to model the underlying non-rigid transformation and provide a brief description of the same here. Given a control point set $X = \{x_k \in \mathbb{R}^D\}, k = 1, \dots, K$, the non-rigid mapping is defined as $T(X, w) = Xd + \sum_{k=1}^K \phi(x - x_k)u_k, w = \{d, u\}$, where d is modeled as the affine part and u is a matrix with a D dimensional vector u_k as each row. The parameter matrix u corresponds to the non-affine aspect of the deformation (containing local bending and stretching). In addition to minimizing the geodesic length, we add the regularization (smoothness) term $\lambda \text{trace}[u^T \Phi u]$ to the original cost function. In practice, the re-parametrization described in [5] is required to guarantee that the smoothness term is non-negative definite. In order to avoid non-physical mappings caused by an inappropriate choice of λ , we modify the smoothness term to the Bregman divergence [3] $\lambda \text{trace}[(u - u_o)^T \Phi (u - u_o)]$ where u_o is typically identified with the non-affine TPS parameters from the previous Bregman epoch and this goes to zero when w converges. The final cost function used in our approach is

$$\hat{w} = \underset{w}{\text{argmin}} \left\{ \text{dist}(\psi(T(M, w)), \psi(S)) + \lambda \text{trace}[(u - u_o)^T \Phi (u - u_o)] \right\} \quad (5)$$

3. The geodesic distance expression

To optimize the cost function 5, a numerical approximation of the Fisher-Rao metric is not necessary. An analytic expression for the geodesic can be obtained which is crucial for gradient-based optimization. Before we can obtain the closed form expression for the geodesic length, we have to engage in some housekeeping. We first have to obtain the normalization term for a square-root SDT density so that it has unit norm:

$$\int_{\Omega} (\psi(x))^2 dx = c \int_{\Omega} \left(\sum_k \exp \left\{ -\frac{\|x - x_k\|}{\tau} \right\} \right)^2 dx = 1 \quad (6)$$

The normalization constant expression will depend on the dimensionality of the point-sets. We first work this out in 2D using Parseval's theorem and 2D circularly symmetric Fourier transforms. The Fourier transform of the integrand

in (2) is

$$\exp\left\{-\frac{\|x - x_k\|}{\tau}\right\} \xrightarrow{\mathcal{F}} \frac{2\pi\frac{1}{\tau}}{(4\pi^2q^2 + \frac{1}{\tau^2})^{\frac{3}{2}}} \exp\{-i2\pi\langle\nu, x_k\rangle\} \quad (7)$$

where ν is the 2D spatial frequency with magnitude q . Using Parseval's theorem, $\int_{\Omega} \psi_1(x)\psi_2^*(x)dx = \int_{\Omega} \Psi_1(\nu)\Psi_2^*(\nu)d\nu$, we can rewrite (6) as

$$c \sum_k \sum_l \int_{\Omega} \frac{4\pi^2\frac{1}{\tau^2}}{(4\pi^2q^2 + \frac{1}{\tau^2})^{\frac{3}{2}}} \exp\{-i2\pi\langle\nu, x_k - x_l\rangle\} d\nu. \quad (8)$$

The integral in (8) can be identified as a Hankel transform. Using standard Hankel transform pairs, we finally obtain the normalization

$$c = \sum_{k,l|k \neq l} \frac{\pi}{4} r_{kl}^2 K_2\left(\frac{r_{kl}}{\tau}\right) + M \frac{\pi h^2}{2} \quad (9)$$

where $r_{kl} = \|x_k - x_l\|_2$ and M is the cardinality of the point-set.

While we seem to have spent an inordinate amount of time deriving a constant c which makes ψ have unit norm, the calculations above actually help set up the analytic geodesic distance expression we seek. Using techniques (Parseval's theorem, Hankel transforms) similar to those above, the inner product between two SDTs $\psi_1(x)$ and $\psi_2(x)$ can be expressed as

$$\langle\psi_1(x), \psi_2(x)\rangle = \int_{\Omega} \psi_1(x)\psi_2(x)dx = \int_{\Omega} \Psi_1(\nu)\Psi_2^*(\nu)d\nu = \frac{1}{\sqrt{c_1c_2}} \sum_k \sum_l \Phi_2(r_{kl}; \tau) \quad (10)$$

where

$$\Phi_2(r; \tau) \stackrel{\text{def}}{=} \begin{cases} \frac{\pi}{4} r^2 K_2\left(\frac{r}{\tau}\right), & r > 0 \\ \frac{\pi h^2}{2}, & r = 0. \end{cases} \quad (11)$$

In (10), $\Psi_{1(2)}(\nu)$ is the 2D Fourier transform of the SDT $\psi_{1(2)}(x)$ and c_1, c_2 are the normalization constants for $\psi_1(x), \psi_2(x)$. Given the inner product (in 2D) between two SDTs, we obtain an analytic expression for the geodesic length in (3).

The geodesic distance between two point-sets in 3D can be derived in a similar manner. While some techniques such as Parseval's theorem remain the same, other methods such as the Hankel transform have to be replaced by their 3D counterparts. We omit the calculations which are somewhat ponderous (but straightforward) and merely give the final expression for the inner product between two SDTs $\psi_1(x)$ and $\psi_2(x)$ in 3D:

$$\langle\psi_1(x), \psi_2(x)\rangle = \frac{1}{\sqrt{c_1c_2}} \sum_k \sum_l \Phi_3(r_{kl}; \tau) \quad (12)$$

where

$$\Phi_3(r; \tau) \stackrel{\text{def}}{=} \begin{cases} \frac{\sqrt{2\pi h}}{3} K_{\frac{5}{2}}\left(\frac{r}{h}\right) r^{\frac{5}{2}}, & r > 0 \\ \frac{\pi h^3}{3}, & r = 0. \end{cases} \quad (13)$$

The normalization constants c_1, c_2 here are derived in a manner similar to their 2D siblings.

Now we can use equations (10), (12) as the analytic expressions of the Fisher-Rao metric between two square root densities and the geodesic is merely the inverse cosine of the Fisher-Rao metric. The analytic gradient of the objective function can be easily derived using the relationship $\frac{\partial x^\nu K_\nu x}{\partial(x)} = -x^\nu K_{\nu-1}(x)$. We omit the turgid details here.

4. Optimization Approach

Since the objective function—the geodesic length between two SDTs with a deformation regularization—has well specified analytic expressions, we may use any available gradient-based nonlinear optimization approach for determining a suitable non-rigid shape deformation. After examining available options, we selected the commercial KNITRO[®] package [16] which is very suitable for this problem. Within KNITRO[®], we selected the dense quasi-Newton Broyden Fletcher Goldfarb Shanno (BFGS) approach and in all cases, we terminated the BFGS algorithm after 2000 iterations.

We have not yet mentioned a principled approach to select the two free parameters— λ and τ —in SDTM. The usual approach is cross-validation on a training set followed by testing the generalization performance on a test set. We plan to pursue this aspect of SDTM in the future.

5. Experimental Results

In this section, we present experimental results using our approach—SDTM—on several synthetic and real data sets. We conducted extensive validations against three state-of-the-art point-cloud matching algorithms: Gaussian mixture model-based registration (gmmreg) [8], coherent point drift (CPD) [13], and robust point matching-local neighborhood structures (RPM-LNS) [18]. We were unable to get the code for the topology preserving relaxation labeling (TPRL) algorithm [10] and hence could not use it in our comparisons. We used “Recall”, a widely used measure in the point cloud matching field [8], to gauge the performance. “Recall” is defined as the ratio of true-positive correspondences to the cardinality of the model data within a pre-specified tolerance.

5.1. Rigid Registration

Thirty different shapes from GatorBait_100¹ are used as models for rigid registration. For each model set, we generate the scene data by applying different rotations (to the model set) in the range $[-\frac{\pi}{2}, \frac{\pi}{2}]$ and then use these data as input to gmmreg, CPD and SDTM respectively. The results show that our algorithm and gmmreg recover 100%

¹This dataset is available at

http://www.cise.ufl.edu/~anand/GatorBait_100.tgz.

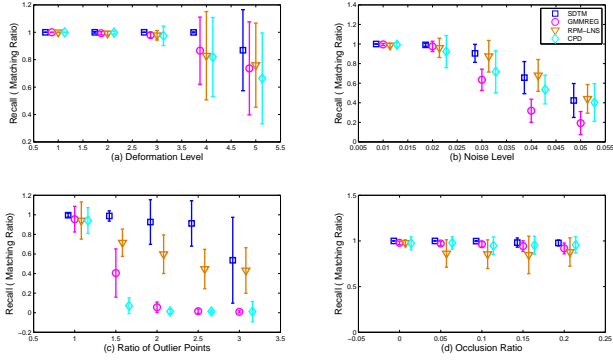


Figure 1. “Recall” for the non-rigid registration algorithms on the fish shape pairs with various degrees of deformation, noise, outliers and occlusion ratios. Legend in (b) is for all four sub-figures. Note that higher the value of “Recall”, better is the method’s performance.

correspondence while CPD converges only for rotations in the range (in radians) $[-1.25, 1.24]$.

5.2. Non-rigid Registration

5.2.1 Experiments on Synthetic Data

Next, we conducted experiments on a 2D fish shape dataset subject to various degrees of deformation, noise, outliers, and occlusion. Fifty fish shapes are used to generate the test cases. The model data set contains 98 points. Twelve points are used as control points for the thin-plate spline in our method while 25 points are used in *gmmreg* as suggested by the authors. As for CPD and RPM-LNS, the entire model set is used as controls for the deformation exactly as was recommended by the authors. Performance accuracy, in terms of the “Recall” measure defined earlier for a tolerance threshold of 0.01, of each of the aforementioned algorithms is depicted in Figure 1. Error bars are also depicted for each method for varying levels of deformation, noise, outliers and occlusion. As evident, SDTM performs either better or almost the same as the best among the competitors.

For each fish dataset, five deformation levels are applied to build the scene data. Our method—SDTM—recovers more than 85% of the correspondences at all the deformation levels, as shown in Figure 1. SDTM clearly performs the best (compared to the other methods) in this scenario and outperforms *gmmreg* (the overall second best results) by 18% in the largest deformation case.

Next, five levels of Gaussian noise with standard deviation σ from 0.01 to 0.05 are used to create the noise corrupted scene data. SDTM outperforms other methods at low and medium noise levels ($\sigma = 0.01, 0.02, 0.03$) and achieves a similar performance to the RPM-LNS method at high noise levels ($\sigma = 0.04, 0.05$).

To test the case of data with outliers, we added 0%, 50%, 100%, 150% and 200% outliers to the scene data




Figure 2. Performance accuracy using the “Recall” measure for non-rigid registration algorithms on the CMU House Sequence (left) and the 3D swagger motion data (right). Legend in (a) is for both sub-figures.

from the moderate deformation case to create the scene data with outliers. The non-rigid registration results in the outlier setting are shown next at the bottom left of Figure 1. As evident, SDTM has the best mean performance compared with the other methods but does show more variability at larger outlier ratios. SDTM outperforms the results from the overall second best method (RPM-LNS) by 55% on the average among cases with varying degree of outliers.

Next, we randomly remove 5%, 10%, 15% and 20% points in the scene data to simulate the case of missing data. The results from application of different algorithms to this type of data are presented at bottom right of Figure 1. Note that SDTM recovers almost 100% of the correspondence in all the cases with lowest variability in comparison to the competition.

5.2.2 Experiments on Real Data

We conducted further quantitative evaluation on SDTM using the CMU house sequence which contains 111 images of a toy house captured from moving viewpoints. Thirty points are chosen in each image with known correspondence. The registration algorithms are applied on image pairs spaced by 20, 40, 60, 80, 100 frames respectively. The results shown on the left of Figure 2 depict performance accuracy plot for the 11 image pairs, using the “Recall” measure with a threshold of 0.03. SDTM shows the best performance in most scenarios excluding the case where the images are spaced by 100 frames while another state of the art method RPM-LNS fails to work on this data set.

We also tested SDTM on a 3D data set—the “swagger” motion capture data from the Advanced Computing Center for the Arts and Design at Ohio State University. Forty two markers are reported in 581 frames. The experiments are conducted on 30 data pairs on frames separated by 50, 100, 150, and 200 frames respectively. The comparison results with *gmmreg* and CPD for which 3D code is available are shown on the right of Figure 2. SDTM yields the best mean value with large variability when the data is spaced by 100, 150, and 200 frames respectively, and *gmmreg* gives similar mean value with small variability in the first case. Overall, our algorithm outperforms the other state of art algorithms in most scenarios. The results suggest that our method can be applied to broad settings and achieves good

performance.

6. Summary and Discussion

We motivated this paper at the start with the observation that previous work in the theoretical physics literature allows us to transition from Hamilton-Jacobi scalar fields to Schrödinger distance transforms (SDTs) which have unit L_2 norm. This identification permits us to treat Schrödinger distance transforms as square-root densities. Based on this identification, we represented point-sets as square-root densities with the geometry allowing us to place them at unique locations on a unit Hilbert sphere. The problem of deformable shape matching requires us to construct a distance measure between shapes which when minimized brings the point-sets into correspondence. The unit Hilbert sphere geometry naturally leads to a geodesic length minimization principle. Furthermore, the form of the Schrödinger wave function enables the geodesic length to be expressed (using Parseval's theorem and the Fourier transform) as a compact analytic expression. Nonlinear optimization of the geodesic length under the action of regularized spline-based deformations completes the picture. The result is SDTM, a novel method which is a contender for the state of the art point-set matching when evaluated against many quantitative metrics.

Possible directions for future work follow from the technical underpinnings of the present work and the empirical studies conducted. We have used a real-valued Schrödinger wave function—a fundamental limitation of the present work. In physics, the phase of the wave function is connected to the action whereas the magnitude is related to a probability density function. This clearly suggests that we can embed topology information in the phase of the wave function while continuing to represent point-set location information in the magnitude—an exciting direction for future work. More mundane extensions—shape atlases and dictionaries for example—represent low hanging fruit which can be immediately procured. Finally, we can impose constraints on the optimization path which may lead to intermediate points retaining semantic content.

Acknowledgement

This research has been in part supported by the grants NIH NS066340 to Dr. Vemuri and the Wayne Densch Epilepsy Research Fund to Dr. Eisenschenk.

References

[1] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(4):509–522, 2002. 2

- [2] P. J. Besl and N. D. McKay. A method for registration of 3D shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992. 2
- [3] L. M. Bregman. The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex programming. *USSR Comp. Math. and Math. Phys.*, 7(3):200–217, 1967. 3
- [4] M. Carcassoni and E. R. Hancock. Spectral correspondence for point pattern matching. *Pattern Recognition*, 36(1):193–204, 2003. 2
- [5] H. Chui and A. Rangarajan. A new algorithm for non-rigid point matching. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2:44–51, 2000. 2, 3
- [6] J. Glaunes, A. Trouvé, and L. Younes. 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)*, pages 712–718, 2004. 2
- [7] S. Granger and X. Pennec. Multi-scale EM-ICP: A fast and robust approach for surface registration. In *European Conference on Computer Vision (ECCV)*, pages 418–432, 2002. 2
- [8] B. Jian and B. C. Vemuri. Robust point set registration using Gaussian mixture models. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 33(8):1633–1645, 2011. 2, 4
- [9] S. H. Joshi, E. Klassen, A. Srivastava, and I. Jermyn. A novel representation for Riemannian analysis of elastic curves in \mathbb{R}^n . In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–7, 2007. 3
- [10] J.-H. Lee and C.-H. Won. Topology preserving relaxation labeling for nonrigid point matching. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(2):427–432, 2011. 2, 4
- [11] B. Luo and E. R. Hancock. A unified framework for alignment and correspondence. *Computer Vision and Image Understanding*, 92:26–55, 2003. 2
- [12] J. Ma, J. Zhao, J. Tian, Z. Tu, and A. L. Yuille. Robust estimation of nonrigid transformation for point set registration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2147–2154, 2013. 2
- [13] A. Myronenko and X. Song. Point-set registration: Coherent point drift. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, 2010. 2, 4
- [14] M. Sethi, A. Rangarajan, and K. S. Gurumoorthy. The schrödinger distance transform (SDT) for point-sets and curves. In *CVPR*, pages 198–205, 2012. 1
- [15] Y. Tsin and T. Kanade. A correlation-based approach to robust point set registration. In *ECCV (3)*, pages 558–569, 2004. 2
- [16] R. A. Waltz, J. L. Morales, J. Nocedal, and D. Orban. An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Mathematical Programming*, 107(3):391–408, 2006. 4
- [17] L. Younes. *Shapes and diffeomorphisms*. Springer, 2010. 2
- [18] Y. Zheng and D. S. Doermann. Robust point matching for nonrigid shapes by preserving local neighborhood structures. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(4):643–649, 2006. 2, 4