

FFT-based Alignment of 2d Closed Curves with Application to Elastic Shape Analysis

Günay Doğan^{1,2}
gunay.dogan@nist.gov

Javier Bernal²
javier.bernal@nist.gov

Charles R. Hagwood²
charles.hagwood@nist.gov

¹ Theiss Research
La Jolla, CA, USA

² National Institute of Standards and
Technology, Gaithersburg, MD, USA

Abstract

For many shape analysis problems in computer vision and scientific imaging (*e.g.*, computational anatomy, morphological cytometry), the ability to align two closed curves in the plane is crucial. In this paper, we concentrate on rigidly aligning pairs of closed curves in the plane. If the curves have the same length and are centered at the origin, the critical steps to an optimal rigid alignment are finding the best rotation for one curve to match the other and redefining the starting point of the rotated curve so that the starting points of the two curves match. Unlike open curves, closed curves do not have fixed starting points, and this introduces an additional degree of freedom in the alignment. Hence the common naive method to find the best rotation and starting point for optimal rigid alignment has $O(N^2)$ time complexity, N the number of nodes per curve. This can be slow for curves with large numbers of nodes. In this paper, we propose a new $O(N \log N)$ algorithm for this problem based on the Fast Fourier Transform. Together with uniform resampling of the curves with respect to arc length, the new algorithm results in an order of magnitude speed-up in our experiments. Additionally, we describe how we can use our new algorithm as part of elastic shape distance computations between closed curves to obtain accurate shape distance values at a fraction of the cost of previous approaches.

1 INTRODUCTION

For many problems in computer vision and scientific imaging, one needs to quantitatively compare boundaries of objects, *e.g.*, organs in computational anatomy, cells in morphological cytometry. Often, the comparison of the boundaries of two objects requires that they be aligned first. Thus, methods for aligning boundaries of 2d objects have been discussed in [2, 3, 8, 9, 10, 12, 14]. Mathematically, the problem is that of aligning pairs of 2d closed curves. In this article, we concentrate on rigidly aligning pairs of 2d closed curves, assumed to be square-integrable. The curves have been processed to match in scale (say, by normalization to unit length) and are centered at the origin of a reference coordinate system. In practice, curves are given as lists of points, each list usually in a counterclockwise order around its curve, with the choice of the first point in a list usually arbitrary. With curves in this form, optimal rigid alignment is achieved by

- 1) *Rotating* one curve, *i.e.*, points that define curve, to match as much as possible the points that define the other curve.
- 2) Among points that define the rotated curve, choosing the best *starting point* to match the starting point of the other curve.

In this paper, we focus on these two factors and propose a new algorithm to find the optimal pair of rotation and starting point. If we do not account for these two factors, we may fail to align even two different versions of the *same* curve. However an additional difficulty could prevent our method from working properly: the discrete sampling of one curve and the placing of the representative nodes often will be incompatible with those of the other, especially if they come from different measured data. As will be described below, this difficulty can usually be alleviated by uniform resampling of the two curves with respect to arc length.

If the starting point is known, the optimal rotation can be found using the Kabsch algorithm [6, 7] in $O(N)$ time, N the number of nodes per curve (see Section 2). Otherwise, a common approach is to loop over all starting point candidates while computing optimal rotation for each, and then choosing the pair of starting point and rotation with optimal alignment. This approach is $O(N^2)$ and can be slow for curves with large numbers of nodes. In this paper, we propose a new fast $O(N \log N)$ method for this problem based on the Fast Fourier Transform (FFT), used together with uniform resampling of the curves with respect to arc length. We then use this method as part of elastic shape distance computations between closed curves and find that we compute accurate shape distances in significantly less time than previous approaches [6, 7].

2 OPTIMAL STARTING POINT AND ROTATION

In this section, we describe our new algorithm to compute optimal starting point and rotation for rigid alignment of closed curves. Initially we have two curves (not necessarily closed) of unit length, centered at the origin, defined by functions $\beta_i : [0, 1] \rightarrow \mathbb{R}^2$, $i = 1, 2$, assumed to be square-integrable, *i.e.*, $\|\beta_i\|_{L^2}^2 = \int_0^1 \|\beta_i(t)\|^2 dt < \infty$, $i = 1, 2$. First we review how the optimal rotation $R = R(\theta) = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}$ is found through angle θ that minimizes mismatch energy

$$E_0(R) = \int_0^1 \|\beta_1(t) - R\beta_2(t)\|^2 dt. \quad (1)$$

For this purpose we first rewrite energy (1) as follows

$$E_0(R) = \|\beta_1\|_{L^2}^2 + \|\beta_2\|_{L^2}^2 - 2 \int_0^1 \beta_1^T(t) R \beta_2(t) dt,$$

where $\|\beta_i\|_{L^2}^2$, $i = 1, 2$, have constant value (note $\|R\beta_2(t)\| = \|\beta_2(t)\|$, since R is a rotation matrix). Then minimizing energy (1) is equivalent to maximizing

$$\int_0^1 \beta_1^T(t) R \beta_2(t) dt = \text{tr}(RA^T),$$

where $\text{tr}(RA^T)$ is the trace of RA^T and A is the 2×2 matrix defined by $A_{kj} = \int_0^1 \beta_{1k}(t) \beta_{2j}(t) dt$, $k, j = 1, 2$. By the Singular Value Decomposition (SVD) theorem $A = USV^T$, where U, V are 2×2 orthogonal matrices, $S = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix}$, $\sigma_1 \geq \sigma_2 \geq 0$. Then $\text{tr}(RA^T) = \text{tr}(RVSU^T) =$

$tr(SU^T RV) = tr(ST)$, where $T = U^T RV$. With $T = \begin{pmatrix} t_{11} & t_{12} \\ t_{21} & t_{22} \end{pmatrix}$, then $tr(ST) = \sigma_1 t_{11} + \sigma_2 t_{22}$. Since U, R, V are orthogonal so is T . Thus $-1 \leq t_{kj} \leq 1, k, j = 1, 2$. If $\det(U)\det(V) > 0$ then $\det(T) = 1$ so that the maximum value of $tr(ST)$ occurs when $t_{11} = t_{22} = 1$ which implies $T = I$, the identity matrix. Thus $U^T RV = T = I$ and $R = UV^T$ is an optimal rotation. If $\det(U)\det(V) < 0$ then $\det(T) = -1$ which implies $t_{11}t_{22} - t_{21}t_{12} = -1$ so that $t_{11}t_{22} \leq 0$. Since $|t_{11}| = |t_{22}|$ then the maximum value of $tr(ST)$ occurs when $t_{11} = 1, t_{22} = -1$ which implies $T = \tilde{S} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. Thus $U^T RV = T = \tilde{S}$ and $R = U\tilde{S}V^T$ is an optimal rotation.

In the case of *closed* curves, we have two such curves of unit length, centered at the origin, defined by periodic functions $\beta_i: \mathbb{R} \rightarrow \mathbb{R}^2, \beta_i(t+1) = \beta_i(t)$ for all values of $t, i = 1, 2$, assumed to be square-integrable over any finite interval of \mathbb{R} . In practice β_1 and β_2 are given as finite lists of points, say N points per curve for some integer $N > 0$, each list in a counterclockwise order around its curve with first and last points the same. At first we parametrize these curves with the discrete uniform parametrization of interval $[0, 1]$ obtained from partitioning $[0, 1]$ into subintervals with endpoints $t_l = (l-1)h, h = 1/(N-1), l = 1, \dots, N$, by defining $\beta_i^l = (\beta_{i1}^l, \beta_{i2}^l)^T$ by $\beta_i^l = \beta_i(t_l), i = 1, 2, l = 1, \dots, N$, where $\beta_i(t_l)$ is the l^{th} point in the list for β_i . We then take advantage of the curves being uniformly parametrized this way to discretize integral (1) using the uniform trapezoidal rule for closed curves:

$$E_0^h(R) = h \sum_{l=1}^{N-1} \|\beta_1^l - R\beta_2^l\|^2. \quad (2)$$

As we did with (1), we can rewrite (2) as follows

$$E_0^h(R) = h \left(\sum_{l=1}^{N-1} (\|\beta_1^l\|^2 + \|\beta_2^l\|^2) - 2 \left(\sum_{l=1}^{N-1} (\beta_1^l)^T R \beta_2^l \right) \right),$$

so that minimizing (2) is equivalent to maximizing

$$\sum_{l=1}^{N-1} (\beta_1^l)^T R \beta_2^l = tr(RA^T),$$

where A is the 2×2 matrix defined by $A_{kj} = \sum_{l=1}^{N-1} \beta_{1k}^l \beta_{2j}^l, k, j = 1, 2$. Then as we did for (1), optimal rotation R for (2) can be computed with the SVD of A or, more precisely, with the Kabsch algorithm (see Algorithm 1) [B, D].

Algorithm 1 Computing optimal R for starting point $t_0 = 0$

Compute $A_{kj} = \sum_{l=1}^{N-1} \beta_{1k}^l \beta_{2j}^l, k, j = 1, 2$.

Compute SVD of A s.t. $A = USV^T$.

If $\det(U)\det(V) > 0$ then $\tilde{S} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ else $\tilde{S} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$.

Return $R = U\tilde{S}V^T$.

Now we find starting point t_0 and rotation R that give optimal rigid alignment of β_1 and β_2 by minimizing mismatch energy

$$E(t_0, R) = \int_0^1 \|\beta_1(t) - R\beta_2(t+t_0)\|^2 dt, \quad t_0 \in [0, 1]. \quad (3)$$

Again with β_1 and β_2 uniformly parametrized as above, for each $m, 1 \leq m \leq N-1$, we define $\beta_2^{\oplus m} = (\beta_{21}^{\oplus m}, \beta_{22}^{\oplus m})^T$, by

$$\beta_2^{\oplus m} = \beta_2(t_l + t_m), \quad l = 1, \dots, N.$$

With $t_0 = t_m$, we then discretize integral (3) using again the uniform trapezoidal rule for closed curves:

$$E^h(t_0, R) = h \sum_{l=1}^{N-1} \|\beta_1^l - R\beta_2^{l\oplus m}\|^2. \quad (4)$$

In addition, for each m , we define 2×2 matrix $A(t_m)$ by

$$A_{kj}(t_m) = \sum_{l=1}^{N-1} \beta_{1k}^l \beta_{2j}^{l\oplus m}, k, j = 1, 2. \quad (5)$$

For $t_0 = t_1, \dots, t_{N-1}$, we could use Algorithm 1 to compute $A = A(t_0)$ and the best rotation $R(t_0) = R$, and then return the pair $(t_0, R(t_0))$ that gives the highest value for $\text{tr}(RA^T)$, *i.e.*, the lowest value for (4). This commonly used approach is $O(N^2)$ as the computation of $A(t_0)$ for a single t_0 is $O(N)$.

With $A_{kj}(t_m)$ as in (5), we propose to compute vectors $A_{kj} = (A_{kj}(t_1), \dots, A_{kj}(t_{N-1}))$, $k, j = 1, 2$, in $O(N \log N)$ time using FFT to accomplish the Discrete Fourier Transform (DFT). For this purpose we define $\tilde{\beta}_1^l = (\tilde{\beta}_{11}^l, \tilde{\beta}_{12}^l)^T$ by $\tilde{\beta}_1^l = \beta_1^{N-l+1}$, $l = 1, \dots, N$, and vectors $\tilde{\beta}_{1k}$ and $\tilde{\beta}_{2j}$ by $\tilde{\beta}_{1k} = (\tilde{\beta}_{1k}^1, \dots, \tilde{\beta}_{1k}^{N-1})$, $\tilde{\beta}_{2j} = (\beta_{2j}^1, \dots, \beta_{2j}^{N-1})$, $k, j = 1, 2$. Given arbitrary vectors x, y of length $N - 1$, we let $\mathbf{DFT}(x)$ and $\mathbf{DFT}^{-1}(y)$ denote the DFT of x and the inverse DFT of y , respectively. With the symbol \cdot indicating component by component multiplication of two vectors, then by the convolution theorem for the DFT we have

$$A_{kj} = \left(\sum_{l=1}^{N-1} \beta_{1k}^l \beta_{2j}^{l\oplus 1}, \dots, \sum_{l=1}^{N-1} \beta_{1k}^l \beta_{2j}^{l\oplus (N-1)} \right) = \mathbf{DFT}^{-1}[\mathbf{DFT}(\tilde{\beta}_{1k}) \cdot \mathbf{DFT}(\beta_{2j})], k, j = 1, 2,$$

which enables us to reduce the computation of the matrix element $A_{kj}(t_0)$ for all t_0 to three $O(N \log N)$ FFT operations (thus a total of twelve for all of $A(t_0)$). Once we have computed $A(t_0)$ for all t_0 this way, we can loop over t_0 candidates, compute for each t_0 the corresponding optimal rotation $R(t_0)$ and then $\text{tr}(R(t_0)A(t_0)^T)$ (instead of the energy $E^h(t_0, R(t_0))$), and return the pair $(t_0, R(t_0))$ that gives the highest $\text{tr}(R(t_0)A(t_0)^T)$ value, *i.e.*, the lowest value for $E^h(t_0, R(t_0))$ in (4). This is summarized in Algorithm 2. There for arbitrary vectors x, y of length $N - 1$, $\mathbf{FFT}(x)$, $\mathbf{IFFT}(y)$ denote $\mathbf{DFT}(x)$, $\mathbf{DFT}^{-1}(y)$, respectively, computed with FFT. Note that the computation of the SVD of $A(t_0)$ has $O(1)$ complexity. In addition, the computation of $\text{tr}(R(t_0)A(t_0)^T)$ has $O(1)$ complexity as opposed to that of $E^h(t_0, R(t_0))$ in (4) that would have $O(N)$ complexity.

Algorithm 2 Fast algorithm for optimal $(t_0, R(t_0))$

Compute $A_{kj} = (A_{kj}(t_1), \dots, A_{kj}(t_{N-1})) = \mathbf{IFFT}[\mathbf{FFT}(\tilde{\beta}_{1k}) \cdot \mathbf{FFT}(\beta_{2j})]$, $k, j = 1, 2$.

for $t_0 = t_1, \dots, t_{N-1}$ **do**

 Compute SVD of $A = A(t_0)$ s.t. $A = USV^T$.

 If $\det(U)\det(V) > 0$ then $\tilde{S} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ else $\tilde{S} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$.

 Set $R = R(t_0) = U\tilde{S}V^T$.

 Compute $\text{tr}(RA^T)$.

end for

Return $(t_0, R(t_0))$ that gives the highest $\text{tr}(RA^T)$.

Above we described the obvious way of parametrizing the curves from their lists of points with the same uniform parametrization of $[0, 1]$. It should be pointed out that for the development of the two algorithms above, any other way of parametrizing the curves uniformly

would have produced the same results. As noted in the introduction, the discrete sampling of one curve and the placing of the representative nodes may not be compatible with those of the other. Accordingly, when rigidly aligning curves with the algorithms above, the obvious way of parametrizing the curves uniformly may not be the most appropriate. Thus, for this purpose, it seems more natural to parametrize each curve by its arc length with the same uniform parametrization of $[0, 1]$. In particular, in the case of similar curves, optimal (t_0, R) values are then produced, as the parametrized curves then match in a way enabling correspondence between the curves at endpoints of subintervals in the parametrization domain $[0, 1]$. We do this in two steps. In the first step, for $i = 1, 2$, we define discrete parameter s_i^l by

$$s_i^l = L_i^l / L_i^N, l = 1, \dots, N, \quad L_i^l = \sum_{m=2}^l \|\beta_i^m - \beta_i^{m-1}\|, l = 2, \dots, N, \quad L_i^1 = 0. \quad (6)$$

In the second step, we strive to approximately parametrize β_1 and β_2 by their arc lengths with the same uniform parametrization of $[0, 1]$ using s_1^l and s_2^l in (6). As we already have parameter and curve node pairs $\{(s_1^l, \beta_1^l)\}_{l=1}^N$ for β_1 , $\{(s_2^l, \beta_2^l)\}_{l=1}^N$ for β_2 , we can interpolate the curves with cubic splines using these parameter and curve node pairs, and obtain new curve nodes β_i^l , $i = 1, 2$, $l = 1, \dots, N$, by evaluating the cubic spline interpolants at $t_l = (l - 1)h$, $h = 1/(N - 1)$, $l = 1, \dots, N$, thus approximately parametrizing each curve by its arc length with the same uniform parametrization of $[0, 1]$. This resampling procedure may introduce some loss in the accuracy of the curve representation (effectively a geometric approximation error), which will have a minor impact on the optimal (t_0, R) computed. We find that the error introduced by resampling is not important for our main application, the computation of elastic shape distances. The algorithm we use for shape distance computation is the iterative algorithm in [5] which improves on an imperfect estimate of (t_0, R) . We could of course use nonuniform parametrizations of $[0, 1]$ to parametrize the curves by their arc lengths, such as those based on curvature [9]. However in that case, a nonuniform FFT would be required.

In Figure 1, using a hippopede curve, we illustrate the impact of sampling and node placement on the computation of optimal (t_0, R) . Top row of Figure 1 includes two different samplings of the curve, one relatively uniform (left-most column) and another one (middle column), which has nodes concentrated in one part of the curve. We are not able to align the two curves well (right-most column) by simply parametrizing them in the obvious way that does not take arc length into consideration. The figure includes resampled versions of the two curves (in the bottom row, same order). As the resampled curves have the same uniform arc length parametrization that ensures good correspondence between the curve nodes, we are able to align the two curves very well. In the figure the first points in the lists of points (initial starting points) are circled in black. The angle of rotation is $\frac{\pi}{3}$.

We did an experimental validation of our algorithm by using different closed curves as test cases. We used five synthetic curves (super-ellipse, hippopede, bumps, limaçon, clover), five (biological) cell boundaries of type A, and five (biological) cell boundaries of type B (see Figure 2). For each of these reference curves, we tried to align a template curve to it. We obtained the template curve by changing the starting point of the reference curve from 0 to 0.25 and rotating the curve by $\frac{\pi}{3}$. We tested with the same uniform arc length parametrization assigned to both curves. We examined the running times and the alignment errors (quantified by Equation (4)) for increasing number of nodes N on the curves. We observed consistent behaviour across different cell examples. Our new $O(N \log N)$ algorithm was much faster than the old $O(N^2)$ algorithm and had exactly the same alignment error values that the old one had. These results are given for the first cell of type B in Figure 3 and Table 1.

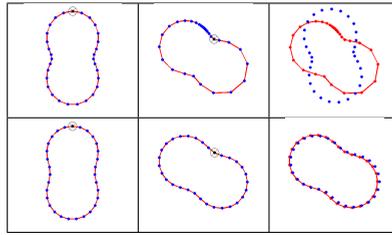


Figure 1: Impact of sampling and node placement in aligning the curves in the 1st and 2nd columns. When the sampling of nodes in the curve is non-uniform (2nd in top row), the alignment fails (3rd in top row). When the sampling is relatively uniform (bottom row), the points in one curve can be easily matched with corresponding nodes in the other curve to ensure successful alignment.

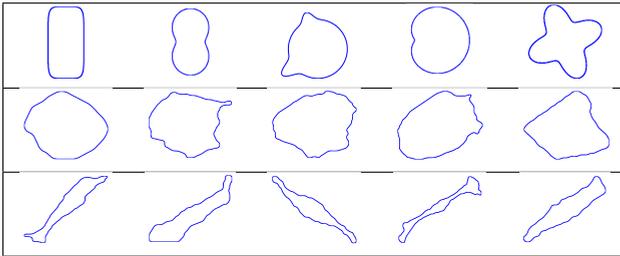


Figure 2: Curve examples used in the experiments: synthetic curves (top row), biological cell boundaries of type A and B (middle and bottom rows).

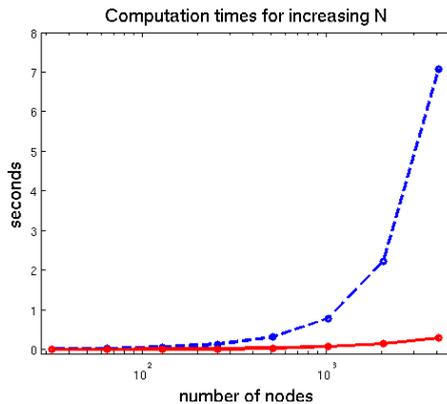


Figure 3: Comparison of timings on an example curve, between the slow $O(N^2)$ alignment algorithm (blue) and the fast $O(N \log N)$ alignment algorithm proposed (red).

	N=64	128	256	512	1024	2048	4096
Slow	.027	.06	.14	.31	.79	2.2	7.1
Fast	.005	.01	.02	.04	.07	.15	.30
Factor	5.4x	6.9x	7.4x	7.5x	11x	15x	24x
Error	2e-2	1e-2	6e-3	3e-3	1.5e-3	7e-4	4e-4

Table 1: Timings (seconds) and errors for alignment algorithms. The two algorithms produce same alignment errors.

3 ALGORITHM FOR ELASTIC SHAPE DISTANCES

In [13], Srivastava et al. proposed an elastic shape distance between closed curves of unit length. In order to represent the shape of one such curve β , $\beta : [0, 1] \rightarrow \mathbb{R}^2$, they defined a function called the square root velocity (SRV) function of β by $q(t) = \dot{\beta}(t) / \|\dot{\beta}(t)\|^{1/2}$, $t \in [0, 1]$, and an orbit $[q] = \{\sqrt{\gamma}Rq(\gamma) \mid \gamma \in \Gamma, R \in SO(2)\}$ representing the shape of β (encoding invariance with respect to rotation and reparameterization by R and γ , respectively; translation invariance is taken care of automatically as $\frac{d}{dt}(\beta(t) + c) = \dot{\beta}(t)$ for any constant planar point c). Here $SO(2)$ is the set of all rotations in \mathbb{R}^2 and Γ is the set of all diffeomorphisms of $[0, 1]$ into $[0, 1]$, with $\gamma(0) = 0$, $\gamma(1) = 1$, for γ in Γ . We note that β of unit length implies $\|q\|_{L^2} = 1$ which in turn implies q is square-integrable. In this setting, the shape distance between closed curves β_1 and β_2 of unit length is the distance between the corresponding square-integrable optimally-matching SRV functions $q_1^*(t) := q_1(t)$ and $q_2^*(t) = \sqrt{\gamma(t)}Rq_2(t_0 + \gamma(t))$, assumed to be centered at the origin, where the triple (t_0, R, γ) is the global minimizer of mismatch energy

$$E(t_0, R, \gamma) = \int_0^1 \|q_1(t) - \sqrt{\gamma(t)}Rq_2(t_0 + \gamma(t))\|^2 dt. \quad (7)$$

Srivastava et al. minimize (7) by looping over all t_0 candidates, computing optimal R with Kabsch algorithm for each t_0 , and then using dynamic programming to get best γ for each (t_0, R) pair [10, 11, 13]. This algorithm, which we call Approach 1, is computationally expensive as it is $O(N^3)$, N the number of nodes per curve. Recently, an $O(N^2)$ iterative algorithm was proposed in [5] of observed subquadratic almost linear time complexity, to optimize (7), which we call Approach 2. It starts from an initial triple of (t_0, R, γ) ($(0, Id, \gamma(t) = t)$ in [5]), and then updates (t_0, R, γ) as it alternates between optimizations with respect to (t_0, R) and γ , until (7) is minimized. Here, assuming without loss of generality that β_1 and β_2 are square-integrable and centered at the origin, we propose to use the FFT-based optimal rigid alignment algorithm of the previous section applied on $\beta_1(t)$ and $R\beta_2(t + t_0)$, together with uniform resampling of β_1 and β_2 with respect to arc length, to compute optimal (t_0^*, R^*) of (3) and then start the algorithm in [5] with $(t_0^*, R^*, \gamma(t) = t)$. Additionally, we propose to modify the algorithm in [5], *i.e.*, Approach 2, as follows. The energy in (7) is reformulated as

$$E(t_0, R, \gamma) = \int_0^1 \|Rq_1(t + t_0) - \sqrt{\gamma(t)}q_2(\gamma(t))\|^2 dt \quad (8)$$

and each optimization of (8) with respect to (t_0, R) is then carried out using in the same way as above the FFT-based optimal rigid alignment algorithm of the previous section (obviously without arc-length uniform resampling of any of the functions involved) applied now on square-integrable functions $q_1^*(t) := Rq_1(t + t_0)$ and $q_2^*(t) = \sqrt{\gamma(t)}q_2(\gamma(t))$, instead of

<i>distance</i>		N=64	128	256	512	1024
Approach 1	γ^1	.0184	.0409	.0362	.0361	N/A
	γ^2	.3344	.3373	.3445	.3441	N/A
Approach 2	$\gamma^{1,2}$.0037	0	0	0	0
<i>timing</i>		N=64	128	256	512	1024
Approach 1	$\gamma^{1,2}$	1.45	11.5	94	756	N/A
Approach 2 (no FFT)	γ^1	0.20	0.20	0.33	0.63	1.75
	γ^2	0.34	0.29	0.37	0.70	1.68
Approach 2 (FFT)	γ^1	0.15	0.14	0.19	0.35	0.85
	γ^2	0.16	0.19	0.26	0.42	0.96

Table 2: Average shape distances and timings (in seconds) for the synthetic curves shown in Figure 2, using Approach 1, Approach 2 without FFT speed-up, Approach 2 with it.

$\beta_1(t)$ and $R\beta_2(t+t_0)$, with q_1^* and q_2^* assumed to be centered at the origin. Accordingly, as mentioned above, Approach 2 is initialized by pre-alignment, *i.e.*, starts from the initial triple $(t_0^*, R^*, \gamma(t) = t)$, and then updates (t_0, R, γ) as it alternates between optimizations with respect to (t_0, R) (using the FFT) and γ (using fast dynamic programming and nonlinear constrained optimization as described in [9]), until (8) is minimized. We found that modifying the algorithm in [9] with the FFT this way results in superb performance in terms of computation times and computed minimizers (t_0, R, γ) . To evaluate our algorithm, we performed two sets of experiments, first with synthetic curves, then with (biological) cell boundaries.

In the first set of experiments, we aimed to evaluate both accuracy and scalability with respect to increasing N . For this, we took a synthetic curve β_1 , changed starting point from 0 to 0.25, rotated by $\frac{\pi}{3}$, and reparameterized β_2 by $\gamma^1(t) = t + 0.025 \sin(4\pi t)$ and then $\gamma^2(t) = t + 1.6t^2(t-1)^2$ (which induce mildly and strongly nonuniform distributions of nodes, respectively). The theoretical shape distance between β_1 and β_2 is zero. We computed the shape distance using Approach 1, Approach 2 initialized by pre-alignment, but no FFT in the actual optimization, then again Approach 2 now with the FFT used in the SRV function optimization as well, for all of the five synthetic curves in Figure 2 with increasing N . The average shape distances and computation times are given in Table 2. Approach 2 computed zero for almost all distances, whereas Approach 1 computed significant nonzero values. Moreover, the computation times of Approach 1 increased dramatically as we increased N , whereas the computation time of Approach 2 grew very slowly, and was less than 1 s even for curves with $N = 1024$ nodes.

In the second set of experiments, we computed matrix of pairwise distances for the 10 cell boundaries, each uniformly sampled at $N = 256$. Computing the 10×10 distance matrix took 9218 s with Approach 1, 125 s with Approach 2 and FFT-based optimization (74X faster). The role of estimating t_0 and R was not as apparent for this curve set. Still Approach 2 computed smaller distances than Approach 1 for 61 of the 100 pairs, indicating better minimization of energy (7). In fact, in only 6 of the remaining 39, Approach 1 computed a significantly smaller distance than Approach 2 (difference ≥ 0.05). If we look at a more relevant subset of the cells, say type B curves, for which optimizing rotation plays a role, we see that Approach 1 computed smaller distances in only 6 of the 25 pairs (see Table 3).

0.0/0.0	.406/.373	.363/.343	.316/.303	.351/.378
.386/.409	0.0/0.0	.430/.414	.358/.337	.376/.365
.343/.372	.434/.421	0.0/0.0	.375/.344	.317/.295
.314/.333	.367/.363	.384/.395	0.0/0.0	.327/.316
.349/.345	.376/.372	.306/.309	.328/.305	0.0/0.0

Table 3: Matrix of pairwise shape distances of type B cells. The first and second values of a pair computed by Approach 1 and Approach 2 with FFT speed-up, respectively. The instances when Approach 1 computed smaller distances than Approach 2 shown in bold.

4 CONCLUSIONS

In this paper, we propose a fast curve rigid alignment algorithm of $O(N \log N)$ time complexity based on the FFT, N the number of nodes per curve. Given two closed curves, together with uniform resampling of the curves with respect to arc length, our algorithm computes the optimal starting point t_0 and rotation R that must be applied on one curve in order to rigidly align it with the other curve. This is the main contribution of this paper. Additionally, we describe how this fast rigid alignment algorithm can be used to initialize the iterative algorithm in [5] for computing elastic shape distances, and how it can be used for the (t_0, R) optimization step of the same algorithm. The resulting new algorithm computes then accurate shape distances at a fraction of the cost of the original one, and scales very well to curves with large numbers of nodes.

References

- [1] Code from Statistical Shape Analysis and Modeling Group, Florida State University. <http://ssamg.stat.fsu.edu/downloads/ClosedCurves2D3D.zip>. Accessed: 2014-06-20.
- [2] N. J. Ayache and O. D. Faugeras. Hyper: A new approach for the recognition and positioning of two-dimensional objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):44–54, 1986.
- [3] I. Cohen, N. J. Ayache, and P. Sulger. Tracking points on deformable objects using curvature information. In *Proceedings European Conference on Computer Vision*, pages 136–144, 1993.
- [4] M. Cui, J. Femiani, J. Hu, P. Wonka, and A. Razdan. Curve matching for open 2d curves. *Pattern Recognition Letters*, 30(1):1–10, 2009.
- [5] G. Doğan, J. Bernal, and C.R. Hagwood. A fast algorithm for elastic shape distances between closed planar curves. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2015.
- [6] W. Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics*, 32(5):922–923, 1976.
- [7] W. Kabsch. A discussion of the solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics*, 34(5):827–828, 1978.

-
- [8] R. Larsen. L1 generalized procrustes 2d shape alignment. *Journal of Mathematical Imaging and Vision*, 31(2-3):189–194, 2008.
- [9] H. Li, T. Shen, and X. Huang. Approximately global optimization for robust alignment of generalized shapes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(6):1116–1131, June 2011.
- [10] W. Mio, A. Srivastava, and S. Joshi. On shape of plane elastic curves. *International Journal of Computer Vision*, 73(3):307–324, 2007.
- [11] J. T. Schwartz and M. Sharir. Identification of partially obscured objects in two and three dimensions by matching noisy characteristic curves. *Int'l J. Robotics Research*, 6(2):29–44, 1987.
- [12] T. B. Sebastian, P. N. Klein, and B. B. Kimia. On aligning curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(1):116–125, 2003.
- [13] A. Srivastava, E. Klassen, S.H. Joshi, and I.H. Jermyn. Shape analysis of elastic curves in Euclidean spaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(7):1415–1428, 2011.
- [14] S. Umeyama. Parameterized point pattern matching and its application to recognition of object families. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(2):136–144, 1993.