On Bipartite Matching under the RMS Distance

Pankaj K. Agarwal* Jeff M. Phillips†

Abstract

Given two sets A and B of n points each in $\mathbb{R}^2$, we study the problem of computing a matching between A and B that minimizes the root mean square (rms) distance of matched pairs. We can compute an optimal matching in $O(n^{2+\delta})$ time, for any $\delta > 0$, and an ε-approximation in time $O((n/\varepsilon)^{3/2} \log^6 n)$. If the set B is allowed to move rigidly to minimize the rms distance, we can compute a rigid motion of B and a matching in $O((n^4/\varepsilon^{5/2}) \log^6 n)$ time whose cost is within $(1+\varepsilon)$ factor of the optimal one.

1 Introduction

Let A and B be two sets of n points each in $\mathbb{R}^2$. A matching $M \subseteq A \times B$ is a set of n pairs of points so that each point of A or B appears in exactly one pair. We define the cost of a matching M to be

$$\omega(M) = \left[\frac{1}{n} \sum_{(a,b)\in M} ||a-b||^2 \right]^{1/2}.$$

We also define

$$\omega_\infty(M) = \max_{(a,b)\in M} ||a-b||.$$

The *minimum cost matching* of A and B is

$$\mathcal{M}(A,B) = \arg\min_M \omega(M)$$

where the minimum is taken over all matchings of A and B. The *bottleneck matching* of A and B is

$$\mathcal{M}_\infty(A,B) = \arg\min_M \omega_\infty(M).$$

If we allow one of the point sets to translate and rotate, then we define the cost of an optimal matching under rigid motion to be

$$\omega(A,B) = \min_{\substack{t \in \mathbb{R}^2 \\ \rho \in \mathrm{SO}(2)}} \omega(\mathcal{M}(A, \rho(B)+t))$$

where SO(2) is the set of all rotations in $\mathbb{R}^2$. We use $\mathbb{M}(A,B)$ to denote the matching whose cost is $\omega(A,B)$.

The problem of aligning two point sets arises in various areas ranging from structural molecular biology [11] to shape registration [9] to medical imaging [7].

The Hungarian algorithm can be used to compute $\mathcal{M}(A,B)$ in $O(n^3)$ time. No polynomial-time algorithm is known for computing $\mathbb{M}(A,B)$. A popular approach for finding a good alignment between A and B under rigid motion is the so-called iterative closest point (ICP) algorithm [4], which alternates between finding the optimal correspondence between points, and finding a rigid motion of one point set so that the rms distance between the matched points is minimized. However, the correspondence step in many of these algorithms aligns many points of A to one point of B, or vice-versa. One can, of course, use the Hungarian algorithm for the correspondence step.

There has been some work on computing a *Euclidean minimum weight matching* between A and B in which the cost of matching is the average length of an edge. Agarwal *et. al.* [1] developed an $O(n^{2+\delta})$ time algorithm, for any $\delta > 0$, to compute a Euclidean minimum weight matching. Faster approximation algorithms are presented in [10, 3].

Cabello *et. al.* [5] compute the *Earth Mover's Distance* between A and B, where each point has a weight and the Euclidean minimum weight matching is calculated aligning each fractional unit of this weight. When B is allowed to move rigidly and the total weight of A and B are the same, a matching can be computed in $O((n^{7/2}/\varepsilon^{9/2}) \log^6 n)$ time whose cost is within $(1+\varepsilon)$ factor of the optimal cost.

In this paper we present exact and approximation algorithms for computing $\mathcal{M}(A,B)$ and an approximation algorithm for computing $\mathbb{M}(A,B)$. More precisely, we can adapt the algorithm in [1] to compute $\mathcal{M}(A,B)$ in $O(n^{2+\delta})$ time for any $\delta > 0$, and the algorithm in [10] to compute an approximation of $\mathcal{M}(A,B)$ in time $O((n/\varepsilon)^{3/2} \log^6 n)$. Finally, we describe an algorithm to compute a matching of A and B and a rigid motion t, ρ so that the cost of $\mathcal{M}(A, \rho(B)+t)$ is at most $(1+\varepsilon)\omega(A,B)$.

2 Computing $\mathcal{M}(A,B)$

Agarwal *et. al.* describe an $O(n^{2+\delta})$ time algorithm for computing the Euclidean minimum weight matching.

*Department of Computer Science, Duke University, pankaj@cs.duke.edu
†Department of Computer Science, Duke University, jeffp@cs.duke.edu

Their algorithm basically implements the Hungarian algorithm but exploits geometry to expedite the running time. It uses a dynamic data structure for computing the nearest neighbor of a query point in a weighted point set under the Euclidean distance function. This is the only place where it uses geometry. Their data structure can be adapted to handle squared Euclidean distance without affecting the asymptotic query and update time. Omitting all the details, we conclude the following.

Theorem 1 *Let A and B be two sets of n points each in $\mathbb{R}^2$. $\mathcal{M}(A,B)$ can be computed in time $O(n^{2+\delta})$ for any $\delta > 0$.*

Next we describe how we can adapt the algorithm in [10] to compute, in time $O((n/\varepsilon)^{3/2}\log^6 n)$, a matching M of A and B whose cost is at most $(1+\varepsilon)\mu$, where $\mu = \omega(\mathcal{M}(A,B))$.

We compute $\mathcal{M}_\infty(A,B)$ using the algorithm in [6]. Let $\mu_\infty = \omega_\infty(\mathcal{M}_\infty(A,B))$. A simple calculation shows that $\mu_\infty/\sqrt{n} \le \mu \le \mu_\infty$. Let $d_{\max} = \omega_\infty(\mathcal{M}(A,B))$. Then $\mu_\infty \le d_{\max}$ and $d_{\max}/\sqrt{n} \le \mu \le d_{\max}$. Set $\gamma = \mu_\infty\varepsilon/8n$.

Lemma 2 *Let (a,b) be an edge of $\mathcal{M}(A,B)$, and let $d = ||a-b||$ be its length. Then*

$$(d+\gamma)^2 \le d^2 + 3\varepsilon\cdot\mu^2/8.$$

Proof. Using the fact that $\gamma \le \mu_\infty \le d_{\max}$

$$\begin{aligned} 2d\gamma + \gamma^2 &\le 2d_{\max}\gamma + \gamma^2 \le 3d_{\max}\cdot\gamma \\ &\le 3\mu\sqrt{n}\left(\frac{\mu_\infty\varepsilon}{8n}\right) \le 3\mu(\mu\sqrt{n})\cdot\frac{\varepsilon}{8\sqrt{n}} \\ &\le 3\varepsilon\mu^2/8, \end{aligned}$$

which proves the lemma. □

Let $\mathbb{G} = \{(i\gamma, j\gamma) \mid i,j \in \mathbb{Z}\}$ be a uniform grid. For a point $p \in \mathbb{R}^2$, let $\tilde{p} \in \mathbb{G}$ be the point nearest to p. Let $\tilde{A} = \langle \tilde{a} \mid a \in A\rangle$ and $\tilde{B} = \langle \tilde{b} \mid b \in B\rangle$ be multisets of n points each. Let $\widehat{M} = \{(a,b) \mid (\tilde{a},\tilde{b}) \in \mathcal{M}(\tilde{A},\tilde{B})\}$ be a matching of A and B.

Lemma 3 *$\omega(\widehat{M}) \le (1+\varepsilon)\mu$.*

Proof. Lemma 2 implies that

$$\begin{aligned} &\left[\frac{1}{n}\sum_{(a,b)\in\mathcal{M}(A,B)} ||\tilde{a}-\tilde{b}||^2\right]^{1/2} \\ &\qquad\le \left[\frac{1}{n}\left(n\cdot\mu^2 + \frac{3\varepsilon}{8}\cdot\mu^2\cdot n\right)\right]^{1/2} \\ &\qquad\le (1+3\varepsilon/8)\mu. \end{aligned}$$

Lemma 2 can also be used to show that

$$\begin{aligned} \omega(\widehat{M}) &\le (1+3\varepsilon/8)\cdot\omega(\mathcal{M}(\tilde{A},\tilde{B})) \\ &\le (1+3\varepsilon/8)^2\mu \le (1+\varepsilon)\mu. \end{aligned}$$

□

We now describe an algorithm for computing an ε-approximation of $\mathcal{M}(\tilde{A},\tilde{B})$ following the approach in [10]. For a briefer exposition let us assume that all points in $\tilde{A}\cup\tilde{B}$ are distinct, though, this assumption is not necessary. We scale $\tilde{A}$ and $\tilde{B}$, in $O(n\log n)$ time, so that the closest pair in $\tilde{A}\cup\tilde{B}$ is at distance 1. After scaling, the length of the longest edge in $\mathcal{M}(\tilde{A},\tilde{B})$ is $d_{\max}\cdot 8n/(\mu_\infty\varepsilon) \le 8n^{3/2}/\varepsilon$. We can assume $\varepsilon > 1/\sqrt{n}$ because otherwise we can simply compute $\mathcal{M}(A,B)$. Hence, the length of the edges in $\mathcal{M}(\tilde{A},\tilde{B})$ are in the range $[1, 8n^2]$. We call all pairs $(\tilde{a},\tilde{b}) \in \tilde{A}\times\tilde{B}$ such that $||a-b|| < 8n^2$ the *interesting pairs*; we ignore the rest.

Let $r \ge \lceil 1/\sqrt{\varepsilon}\rceil$ be an even integer. For $0 \le l < r$, let $u_l = (\cos(2\pi l/r), \sin(2\pi l/r)) \in \mathbb{S}^1$ be a unit vector, and let W_l be the edge formed by u_i and u_{i+1}. $P = \mathrm{conv}(u_0,\ldots,u_{r-1})$ is a regular, centrally symmetric convex r-gon. The Minkowski metric, $d_P(\cdot,\cdot)$, induced by P approximates the Euclidean metric, i.e., for any $a,b\in\mathbb{R}^2$,

$$||a-b|| \le d_P(a,b) \le (1+\varepsilon)||a-b||.$$

As in [10], by setting $k = \log_{1+\varepsilon}(8n^2)$, we compute a family

$$\mathcal{F} = \bigcup_{j\le k}\mathcal{F}_j; \quad \mathcal{F}_j = \{(\tilde{A}_1,\tilde{B}_1),\ldots,(\tilde{A}_u,\tilde{B}_u)\}$$

where (i) $\tilde{A}_i \subseteq \tilde{A}, \tilde{B}_i \subseteq \tilde{B}$, (ii) for every interesting pair $(\tilde{a},\tilde{b}) \in \tilde{A}\times\tilde{B}$, there is exactly one $j \le k$ and exactly one $i \le u$ such that $\tilde{a}\in\tilde{A}_i$ and $\tilde{b}\in\tilde{B}_i$ and $(\tilde{A}_i,\tilde{B}_i) \in \mathcal{F}_j$, and (iii) for any $(\tilde{a},\tilde{b}) \in \tilde{A}_i\times\tilde{B}_i$ and $(\tilde{A}_i,\tilde{B}_i) \in \mathcal{F}_j$, $(d_P(\tilde{a},\tilde{b}))^2 \in \mathcal{I}_j$, where $\mathcal{I}_j = [(1+\varepsilon)^{j-1},(1+\varepsilon)^j]$. For each distance interval $\mathcal{I}_i$ and each wedge W_l, we build in $O(n\log^3 n)$ time a 3-level range tree T_{jl} on B, which stores B as a family $\mathcal{B}_{jl}$ of $O(n\log^3 n)$ canonical subsets, that for a query point $a\in\mathbb{R}^2$ can report all points $b\in B$, as a set $\Phi_{jl}(a)$ of $O(\log^3 n)$ canonical subsets, that satisfy (i) $b \in W_l + a$ and (ii) $d_P(a,b)\in\mathcal{I}_j$, i.e., $b \in \big((1+\varepsilon)^j P + a\big) \setminus \big((1+\varepsilon)^{j-1}P + a\big)$. The first (resp. second) level of T_{jl} is built along direction u_i (resp. u_{i+1}), and the third level is built along $\overrightarrow{u_i u_{i+1}}$. We query the above structure for each $a\in A$. For each $B_i\in\mathcal{B}_{jl}$, we define $A_i = \{a\in A \mid B_i \in \Phi_{jl}(a)\}$.

Finally we set

$$\mathcal{F}_j = \bigcup_l \{(A_i,B_i) \mid B_i \in \mathcal{B}_{jl}\}.$$

By construction $|\mathcal{F}_j| = \sum_{B_i\in\mathcal{B}_{jl}}(|A_i| + |B_i|) = O(1/\sqrt{\varepsilon})\cdot O(n\log^3 n)$. Therefore $|\mathcal{F}| = O(k\cdot(n/\sqrt{\varepsilon})\cdot\log^3 n) = O((n/\varepsilon^{3/2})\cdot\log^4 n)$.

Finally, exploiting the structure of $\mathcal{F}$, Varadarajan and Agarwal [10] show that one can compute in time $O((n/\varepsilon)^{3/2}\log^6 n)$ a matching $\tilde{M}$ of $\tilde{A}$ and $\tilde{B}$ so that

$$\omega(\tilde{M}) \leq (1+\varepsilon)\cdot\omega(\mathcal{M}(\tilde{A},\tilde{B})).$$

Theorem 4 *Let A and B be two sets of n points each in $\mathbb{R}^2$, and let $\varepsilon \geq 0$ be a parameter. We can compute in time $O((n/\varepsilon)^{3/2}\log^6 n)$ a matching of A and B whose cost is at most $(1+\varepsilon)\cdot\omega(\mathcal{M}(A,B))$.*

3 Matching under Rigid Motions

In this section we describe an approximation algorithm for computing $\mathbb{M}(A,B)$. We first consider the translation. Let $\bar{a} = \sum_{a\in A} a/n$ denote the centroid of A, and let $\bar{b}$ denote the centroid of B. Let $Q : A \to B$ be a map. It is well known that

$$\min_{t\in\mathbb{R}^2}\sum_{a\in A}||a-Q(a)-t||^2$$

is attained when $t = \bar{a} - \sum_{a\in A} Q(a)/n$. If Q is bijective, i.e. $\{(a,Q(a)) \mid a \in A\}$ is a matching, then $\sum_{a\in A} Q(a)/n = \bar{b}$. Let $\bar{Q} : A \to B$ be the bijective map corresponding to the matching that attains $\min_{t\in\mathbb{R}^2}\mathcal{M}(A,B+t)$. Then by the above argument,

$$\begin{aligned}\min_t \omega(\mathcal{M}(A,B+t)) &= \left[\frac{1}{n}\sum_{a\in A}||a-\bar{Q}(a)-\bar{a}+\bar{b}||^2\right]^{1/2}\\ &= \omega(\mathcal{M}(A,B+\bar{a}-\bar{b})).\end{aligned}$$

Hence, in order to compute $\mathbb{M}(A,B)$, we first translate B so that the centroids of A and B are aligned by $\bar{a}-\bar{b}$, and then rotate B around their common centroid to minimize the cost of matching under rotation. Note that computing $\mathbb{M}(A,B)$ is now a one-dimensional problem.

With a slight abuse of notation we use B to denote the point set B after it has been translated. We thus have two point sets A and B with a common centroid, say the origin O, and we wish to compute

$$\omega(A,B) = \min_{\rho\in\mathrm{SO}(2)}\omega(\mathcal{M}(A,\rho(B)))$$

and the rotation ρ that attains the minimum. Let

$$\omega_\infty(A,B) = \min_{\rho\in\mathrm{SO}(2)}\omega_\infty(\mathcal{M}(A,\rho(B))).$$

We first compute a 2-approximation of $\omega_\infty(A,B)$ and then use this value to compute an ε-approximation of $\omega(A,B)$ and a matching corresponding to this value. For an angle $\theta \in \mathbb{S}^1$ and for a point $p \in \mathbb{R}^2$, let p_θ be the position of p after being rotated with respect to O by angle θ in the counterclockwise direction. For a point set $X \subseteq \mathbb{R}^2$, let $X_\theta = \{p_\theta \mid p \in X\}$. Note that there is an equivalence between any angle $\theta \in \mathbb{S}^1$ and a particular rotation $\rho \in \mathrm{SO}(2)$, and vice-versa.

Let $\Delta \geq 0$ be a parameter. We describe an algorithm that decides whether $\omega_\infty(A,B) \leq \Delta$. For a pair $(a,b) \in A\times B$, let $\Theta_{ab} \subseteq \mathbb{S}^1$ be the set of angles θ such that $||a-b_\theta|| \leq \Delta$; Θ_{ab} is an angular interval (see Figure 1). Let $\Phi = \langle\phi_0,\ldots,\phi_u\rangle$, $u \leq 2n^2$, be the sequence of endpoints of the intervals Θ_{ab}, $(a,b) \in A\times B$, sorted in counterclockwise direction. For $\theta \in \mathbb{S}^1$, let $G_\theta = \{(a,b) \in A\times B \mid ||a-b_\theta|| \leq \Delta\}$. By construction, G_θ is the same for all angles in a range (ϕ_i,ϕ_{i+1}), which we denote by G_i. Note that $\omega_\infty(A,B) \leq \Delta$ if and only if at least one of the G_i has a perfect matching.

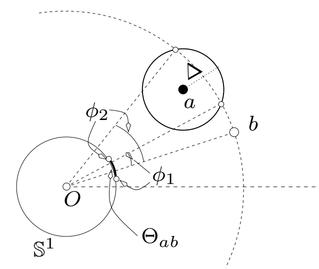

Figure 1: Geometry of interesting angular intervals shown for points a within Δ of b.

By using a disk range searching data structure [2], G_i can be represented implicitly using $O(n^{4/3}\log n)$ edges, and this representation can be computed within the same time bound. We can then compute a matching in G_i in time $O(\sqrt{n}\cdot n^{4/3}\log n) = O(n^{11/6}\log n)$ time [8]. Repeating this step for all $0 \leq i \leq u$, we can determine in $O(n^{23/6}\log n)$ time whether $\omega_\infty(A,B) \leq \Delta$.

Lemma 5 *Let A and B be two sets of n points each in $\mathbb{R}^2$. For a given $\Delta > 0$, we can determine in $O(n^{23/6}\log n)$ time whether $\omega_\infty(A,B) \leq \Delta$.*

Remarks.

1. Since G_i and G_{i+1} differ by a single edge, we can probably improve the running time to $O(n^{10/3}\log^2 n)$.

2. If the distance between two points are computed in L_∞- or L_1-metric, the running time can be improved to $O(n^{7/2}\log^c n)$, which can probably be improved further to $O(n^3\log^c n)$ time.

By performing a binary search and using the above decision procedure we can compute a c-approximation of $\omega_\infty(A,B)$, for any constant c, in time $O(n^{23/6}\log n)$. However, observing that $d_2(x,y) \leq \sqrt{2}\cdot d_\infty(x,y)$ for any $x,y \in \mathbb{R}^2$, and using Remark (2), we can compute a 2-approximation of $\omega_\infty(A,B)$ in time $O(n^{7/2}\mathrm{poly}\log(n))$.

Lemma 6 *Given two sets A and B of n points each in $\mathbb{R}^2$, we can compute a 2-approximation of $\omega_\infty(A,B)$ in time $O(n^{7/2}\text{poly}\log(n))$.*

We are now ready to describe an ε-approximation algorithm for computing $\omega(A,B)$. We compute a quantity $\widehat{\mu}_\infty \le 2\omega_\infty(A,B) \le 2\widehat{\mu}_\infty$, using Lemma 6. Let $\gamma = \widehat{\mu}_\infty\varepsilon/(8n)$. We fix a point $b \in B$. For a point $a \in A$, let $\theta_a = \arg\min_\theta ||a - b_\theta||$, i.e. the ray $\overrightarrow{Ob_{\theta_a}}$ passes through a. Set $\upsilon = 2\sin^{-1}(\gamma/(4||b-O||))$. Let

$$\Psi_a = \left\{\psi_i = \theta_a + i\cdot\upsilon \mid i \in \left[-\lceil 4\widehat{\mu}_\infty\sqrt{n}/\gamma\rceil, \lceil 4\widehat{\mu}_\infty\sqrt{n}/\gamma\rceil\right]\right\},$$

$$\Psi = \bigcup_{a\in A}\Psi_a\,; \qquad |\Psi| = O(n^{5/2}/\varepsilon),$$

as shown in Figure 2. For each $\psi \in \Psi$, we compute in $O((n/\varepsilon)^{3/2}\log^6 n)$ time a matching M_ψ of A and B_ψ such that $\omega(M_\psi) \le (1+\varepsilon/2)\omega(\mathcal{M}(A,B_\psi))$ and return the one with the minimum cost.

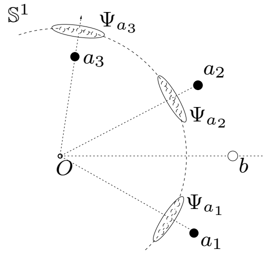

Figure 2: Sampling of $\mathbb{S}^1$ to align a_1, a_2, or a_3 with b. Ψ_{a_1}, Ψ_{a_2}, and Ψ_{a_3} are marked on $\mathbb{S}^1$.

Lemma 7 $\min_{\psi\in\Psi}\omega(A,B_\psi) \le (1+\varepsilon/2)\cdot\omega(A,B)$.

Proof. Observe that $\widehat{\mu}_\infty/(2\sqrt{n}) \le \omega(A,B) \le \widehat{\mu}_\infty$ and that the length of the longest edge in $\mathbb{M}(A,B)$, denoted by $d_{\max}$, is in the range $[\widehat{\mu}_\infty, \widehat{\mu}_\infty\sqrt{n}]$. Consider $(a,b) \in \mathbb{M}(A,B)$. Let $\theta_0 = \arg\min_{\theta\in\mathbb{S}^1}\omega(\mathcal{M}(A,B_\theta))$ be the rotation corresponding to the optimal matching $\mathbb{M}$. Since $||a - b_{\theta_0}|| \le \widehat{\mu}_\infty\sqrt{n}$ and $||b_{\psi_i} - b_{\psi_{i+1}}|| = \gamma/2$, then there must exist a $\psi \in \Psi$ such that $||b_\psi - b_{\theta_0}|| \le \gamma/2$. This condition holds for all $b \in B$ for ψ, thus using Lemma 2 we can bound the total error $\omega(A,B) - \omega(\mathcal{M}(A,B_\psi)) \le \varepsilon/2\cdot\omega(A,B)$. $\square$

We thus compute a $(1+\varepsilon/2)$-approximation matching of A and B_ψ for each $\psi \in \Psi$ and return the matching with the minimum cost. By Lemma 7, it is an ε-approximation of $\omega(A,B)$.

Theorem 8 *Let A and B be two sets of n points each in $\mathbb{R}^2$, and let $\varepsilon > 0$ be a parameter. We can compute in time $O(n^4/\varepsilon^{5/2}\log^6 n)$ time, a matching of A and B (using rigid motion) whose cost is at most $(1+\varepsilon)\cdot\omega(A,B)$.*

References

[1] P. K. Agarwal, A. Efrat, and M. Sharir. Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. *SIAM J. Comput.*, 29:912–953, 2000.

[2] P. K. Agarwal and J. Erikson. Geometric range searching and its relatives.

[3] P. K. Agarwal and K. Varadarajan. A near-linear constant-factor approximation to euclidean bipartite matching? *Proc. 20th Annu. Symp. Comp. Geom.*, pages 247 – 252, 2004.

[4] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Patt. Anal. and Mach. Intel.*, 14(2):239 – 256, 1992.

[5] S. Cabello, P. Giannopoulos, C. Knauer, and G. Rote. Matching point sets with respect to the earth mover's distance. *Proc. 13th Annu. Euro. Symp. Algo.*, 2005.

[6] A. Efrat and A. Itai. Improvements on bottleneck matching and related problems using geometry. *Proc. 12th Annu. Symp. Comp. Geom.*, pages 301–310, 1996.

[7] W. E. Grimson, R. Kikinis, F. A. Jolesz, and P. M. Black. Image-guided surgery. *Sci, Amer.*, 280:62–69, 1999.

[8] J. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matching in bipartite graphs. *SIAM J. of Comput.*, 2:225–231, 1973.

[9] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital michelangelo project: 3D scanning of large statues. *Proc. SIGGRAPH*, pages 131–144, 2000.

[10] K. Varadarajan and P. K. Agarwal. Approximation algorithms for bipartite and non-bipartite matching in the plain. *Proc. 10th Annu. Symp. Disc. Algo.*, pages 805–814, 1999.

[11] T. D. Wu, S. C. Schmidler, T. Hastie, and D. L. Brutlag. Modeling and superposition of multiple protein structures using affine transformations: Analysis of the globins. *Proc. Pac. Symp. on Biocomp.*, pages 507–518, 1998.