

Minimum Partial-Matching and Hausdorff RMS-Distance Under Translation: Combinatorics and Algorithms

Rinat Ben-Avraham^{1,*}, Matthias Henze^{2,**}, Rafel Jaume^{2,***}, Balázs Keszegh^{3,†}, Orit E. Raz^{1,‡}, Micha Sharir^{1,§}, and Igor Tubis¹

¹ Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel, rinatba@gmail.com, {oritraz,michas}@post.tau.ac.il, mrtubis@gmail.com

² Institut für Informatik, Freie Universität Berlin, Berlin, Germany, matthias.henze@fu-berlin.de, jaume@mi.fu-berlin.de

³ Alfréd Rényi Institute of Mathematics, Hungarian Academy of Sciences, Budapest, Hungary, keszegh@renyi.hu

Abstract. We consider the RMS-distance (sum of squared distances between pairs of points) under translation between two point sets in the plane. In the Hausdorff setup, each point is paired to its nearest neighbor in the other set. We develop algorithms for finding a local minimum in near-linear time on the line, and in nearly quadratic time in the plane. These improve substantially the worst-case behavior of the popular ICP heuristics for solving this problem. In the partial-matching setup, each point in the smaller set is matched to a distinct point in the bigger set. Although the problem is not known to be polynomial, we establish several structural properties of the underlying subdivision of the plane and derive improved bounds on its complexity. In addition, we show how to compute a local minimum of the partial-matching RMS-distance under translation, in polynomial time.

Keywords: partial matching, Hausdorff RMS-distance, polyhedral subdivision, local minimum

* Supported by Grant 2012/229 from the U.S.-Israel Binational Science Foundation.

** Supported by ESF EUROCORES programme EuroGIGA-VORONOI, (DFG): RO 2338/5-1.

*** Supported by La-Caixa and the DAAD.

† Supported by Hungarian National Science Fund (OTKA), under grant PD 108406, NN 102029 (EUROGIGA project GraDR 10-EuroGIGA-OP-003), NK 78439, by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences and by the DAAD.

‡ Supported by Grant 892/13 from the Israel Science Foundation.

§ Supported by Grant 2012/229 from the U.S.-Israel Binational Science Foundation, by Grant 892/13 from the Israel Science Foundation, by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11), and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

1 Introduction

Let A and B be two finite sets of points in the plane, of respective cardinalities n and m . We are interested in measuring the similarity between A and B , under a suitable proximity measure. We consider two such measures where the proximity is the sum of the squared distances between pairs of points. In the first, we assume that $n > m$ and we want to match all the points of B (a specific pattern that we want to identify), in a one-to-one manner, to a subset of A (a larger picture that “hides” the pattern) of size $|B|$. This is motivated by situations where we want a one-to-one matching between A and B [9, 15, 16]. In the second, each point is assigned to its nearest neighbor in the other set. See [1] for a similar generalization of the Hausdorff distance.

We refer to the measured distance between the sets, in both versions, as the *RMS distance*. In the former setup the measure is called the *partial-matching RMS-distance*, and in the latter we call it the *Hausdorff RMS-distance*. In both variants the sets A and B are in general not aligned, so we seek a translation of one of them that will minimize the appropriate RMS-distance, partial matching or Hausdorff.

The partial-matching RMS-distance problem. Let $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_m\}$ be two sets of points in the plane. Here we assume that $m < n$, and we seek a minimum-weight *maximum-cardinality matching* of B into A . This is a subset M of edges of the complete bipartite graph with edge set $B \times A$, so that each $b \in B$ appears in exactly one edge of M , and each $a \in A$ appears in at most one edge. The weight of an edge (b, a) is $\|b - a\|^2$, and the weight of a matching is the sum of the weights of its edges.

A maximum-cardinality matching can be identified as an injective assignment π of B into A . With a slight abuse of notation, we denote by $a_{\pi(i)}$ the point a_j that π assigns to b_i . In this notation, the minimum RMS partial-matching problem (for fixed locations of the sets) is to compute

$$M(B, A) = \min_{\pi: B \rightarrow A \text{ injective}} \sum_{i=1}^m \|b_i - a_{\pi(i)}\|^2.$$

Allowing the pattern B to be translated, we obtain the problem of the minimum partial-matching RMS-distance under translation, defined as

$$M_T(B, A) = \min_{t \in \mathbb{R}^2} M(B + t, A) = \min_{\substack{t \in \mathbb{R}^2, \pi: B \rightarrow A, \\ \pi \text{ injective}}} \sum_{i=1}^m \|b_i + t - a_{\pi(i)}\|^2.$$

The function $F(t) := M(B + t, A)$ induces a subdivision of \mathbb{R}^2 , where two points $t_1, t_2 \in \mathbb{R}^2$ are in the same region if the minimum of F at t_1 and at t_2 is attained by the same assignment $\pi : B \rightarrow A$. We refer to this subdivision, following Rote [12], as the *partial-matching subdivision* and denote it by $\mathcal{D}_{B,A}$. We say that a matching is *optimal* if it attains $F(t)$ for some $t \in \mathbb{R}^2$.

The Hausdorff RMS distance problem. Let $N_A(x)$ (resp., $N_B(x)$) denote the nearest neighbor in A (resp., in B) of a point $x \in \mathbb{R}^2$. The *unidirectional (Hausdorff) RMS distance* between B and A is defined as

$$RMS(B, A) = \sum_{b \in B} \|b - N_A(b)\|^2.$$

We also consider *bidirectional* RMS distances, in which we also measure distances from the points of A to their nearest neighbors in B . We consider two variants of this notion. The first variant is the L_1 -*bidirectional RMS distance* between A and B , which is defined as

$$RMS_1(B, A) = RMS(A, B) + RMS(B, A).$$

The second variant is the L_∞ -*bidirectional RMS distance* between A and B , and is defined as

$$RMS_\infty(B, A) = \max\{RMS(A, B), RMS(B, A)\}.$$

Allowing one of the sets (say, B) to be translated, we define the *minimum unidirectional RMS distance under translation* to be

$$RMS_T(B, A) = \min_{t \in \mathbb{R}^2} RMS(B + t, A) = \min_{t \in \mathbb{R}^2} \sum_{b \in B} \|b + t - N_A(b + t)\|^2,$$

where $B + t = \{b_1 + t, \dots, b_m + t\}$. Similarly, we define the *minimum L_1 - and L_∞ -bidirectional RMS distances under translation* to be

$$\begin{aligned} RMS_{T,1}(B, A) &= \min_{t \in \mathbb{R}^2} RMS_1(B + t, A) \quad \text{and} \\ RMS_{T,\infty}(B, A) &= \min_{t \in \mathbb{R}^2} RMS_\infty(B + t, A). \end{aligned}$$

Background. A thorough initial study of the minimum RMS partial-matching distance under translations is given by Rote [12]; see also [5, 13] for two follow-up studies, another study in [11], and an abstract of an earlier version of parts of this paper [8]. The resulting subdivision $\mathcal{D}_{B,A}$, as defined above, is shown in [12] to be a convex subdivision. Rote's main contribution for the analysis of the complexity of $\mathcal{D}_{B,A}$ was to show that a line crosses only $O(mn)$ regions of the subdivision (see Theorem 1 below). However, obtaining sharp bounds for the complexity of $\mathcal{D}_{B,A}$ is still an open issue, where the best known upper bounds are exponential.

The problem of Hausdorff RMS minimization under translation has been considered in the literature (see, e.g., [1] and references therein), although only scarcely so. If A and B are sets of points on the line, the complexity of the Hausdorff RMS function, as a function of t , is $O(mn)$ (and this bound is tight in the worst case). Moreover, the function can have many local minima (up to $\Theta(mn)$ in the worst case). Hence, finding a translation that minimizes the Hausdorff RMS distance can be done in brute force, in $O(mn \log(mn))$ time,

but a worst-case near linear algorithm is not known. In practice, though, there exists a popular heuristic technique, called the ICP (Iterated Closest Pairs) algorithm, proposed by Besl and McKay [3] and analyzed in Ezra et al. [7]. Although the algorithm is reported to be efficient in practice, it might perform $\Theta(mn)$ iterations in the worst case. Moreover, each iteration takes close to linear time (to find the nearest neighbors in the present location).

The situation is worse in the plane, where the complexity of the RMS function is $O(m^2n^2)$, a bound which is worst-case tight, and the bounds for the performance of the ICP algorithm, are similarly worse. Similar degradation shows up in higher dimensions too; see, e.g., [7].

Our results. In this paper we study these two fairly different variants of the problem of minimizing the RMS distance under translation, and improve the state of the art in both of them.

In the partial-matching variant, we first analyze the complexity of $\mathcal{D}_{B,A}$. We significantly improve the bound from the naive $O(n^m)$ to $O(n^2m^{3.5}(e \ln m + e)^m)$. A preliminary informal exposition of this analysis by a subset of the authors is given in the (non-archival) note [8]. This paper expands the previous note, derives additional interesting structural properties of the subdivision, and significantly improves the complexity bound. The arguments that establish the bound can be generalized to bound the number of regions of the analogous subdivision in \mathbb{R}^d by $O((n^2m)^d(e \ln m + e)^m/\sqrt{m})$. The derivation of the upper bound proceeds by a reduction that connects partial matchings to a combinatorial question based on a game theoretical problem, which we believe to be of independent interest.

Next we present a polynomial-time algorithm for finding a local minimum of the partial-matching RMS-distance. This is significant, given that we do not have a polynomial bound on the size of the subdivision. We also fill in the details of explicitly computing the intersections of a line with $\mathcal{D}_{B,A}$. Although Rote hinted at such an algorithm in [12], by exploiting some new properties of $\mathcal{D}_{B,A}$ derived here, we manage to compute the intersections in a simple, more efficient manner.

We also note that by combining the combinatorial bound for the complexity of $\mathcal{D}_{B,A}$, along with the procedures in the algorithm for finding a local minimum of the partial-matching RMS-distance, it is possible to traverse all of $\mathcal{D}_{B,A}$, and compute a *global* minimum of the partial-matching RMS-distance in time $O(n^3m^{7.5}(e \ln m + e)^m)$. This is the best known bound for this problem.

For the Hausdorff variant, we provide improved algorithms for computing a local minimum of the RMS function, in one and two dimensions. Assuming $|A| = |B| = n$, in the one-dimensional case the algorithms run in time $O(n \log^2 n)$, and in the two-dimensional case they run in time $O(n^2 \log n)$. Our approach thus beats the worst-case running time of the ICP algorithm (used for about two decades to solve this problem). The approach is an efficient search through the (large number of) critical values of the RMS function. The techniques are reasonably standard, although their assembly is somewhat involved.

2 Properties of $\mathcal{D}_{B,A}$

We begin by reconstructing several basic properties of $\mathcal{D}_{B,A}$ that have been noted in [12]. First, if we fix the translation $t \in \mathbb{R}^2$ and the assignment π , the cost of the matching, denoted by $f(\pi, t)$, is

$$f(\pi, t) = \sum_{i=1}^m \|b_i + t - a_{\pi(i)}\|^2 = c_\pi + \langle t, d_\pi \rangle + m \|t\|^2, \quad (1)$$

where $c_\pi = \sum_{i=1}^m \|b_i - a_{\pi(i)}\|^2$ and $d_\pi = 2 \sum_{i=1}^m (b_i - a_{\pi(i)})$. For t fixed, the assignment π that minimizes $f(\pi, t)$ is the same assignment that minimizes $g(\pi, t) := c_\pi + \langle t, d_\pi \rangle$. It follows that $\mathcal{D}_{B,A}$ is the minimization diagram (the xy -projection) of the graph of the function

$$\mathcal{E}_{B,A}(t) = \min_{\pi: B \rightarrow A \text{ injective}} (c_\pi + \langle t, d_\pi \rangle), \quad t \in \mathbb{R}^2.$$

This is a lower envelope of a finite number of planes, so its graph is a convex polyhedron, and its projection $\mathcal{D}_{B,A}$ is a convex subdivision of the plane, whose faces are convex polygons. The great open question regarding minimum partial-matching RMS-distance under translation, is whether the number of regions of $\mathcal{D}_{B,A}$ is polynomial in m and n . A significant, albeit small step towards settling this question is the following result of Rote [12].

Theorem 1 (Rote [12]). *A line intersects the interior of at most $m(n-m)+1$ different regions of the partial-matching subdivision $\mathcal{D}_{B,A}$.*

The following property, observed by Rote [12], seems to be well known [16].

Lemma 1. *For any $A' \subset A$, with $|A'| = m$, the optimal assignment that realizes the minimum $M(B + t, A')$ is independent of the translation $t \in \mathbb{R}^2$.*

Next, we derive several additional properties of $\mathcal{D}_{B,A}$ which show that the diagram has, locally, low-order polynomial complexity.

Lemma 2. *Every edge of $\mathcal{D}_{B,A}$ has a normal vector of the form $a_j - a_i$ for suitable $i, j \in \{1, \dots, n\}$.*

Proof. Let E be an edge of $\mathcal{D}_{B,A}$ common to the regions associated with the injections $\pi, \sigma : B \rightarrow A$. By definition, $g(\pi, t) = g(\sigma, t) \leq g(\delta, t)$ for every injection $\delta : B \rightarrow A$ and for any $t \in E$. By Equation (1), E is contained in the line $\ell(\pi, \sigma) = \{t \in \mathbb{R}^2 : \langle t, d_\pi - d_\sigma \rangle = c_\sigma - c_\pi\}$. Let $H = (\pi \setminus \sigma) \cup (\sigma \setminus \pi)$. It is easy to see that H consists of a vertex-disjoint union of cycles and alternating paths. Let $\gamma_1, \dots, \gamma_p$ be these cycles and paths. It is not hard to see that every cycle and every path can be “flipped” independently while preserving the validity of the matching; that is, we can choose, within any of the γ_j ’s, either all the edges corresponding to π or all the ones corresponding to σ , without interfering with other cycles or paths, so that the resulting collection of edges still represents an

injection from B into A . Observe now that $\ell(\pi, \sigma) = \{t \in \mathbb{R}^2 : \langle t, \sum_{j=1}^p d_{\gamma_j} \rangle = -\sum_{j=1}^p c_{\gamma_j}\}$, where d_{γ_j} is the sum of the terms in $d_\pi - d_\sigma$ that involve only the $a_i \in A$ contained in γ_j and c_{γ_j} is analogously defined for $c_\pi - c_\sigma$. Note that d_{γ_j} is 0 for every cycle γ_j and, therefore, at least one of the γ_j 's is a path. Then, we must have $\langle t, d_{\gamma_j} \rangle = -c_{\gamma_j}$ for all $j = 1, \dots, p$ and every $t \in \ell(\pi, \sigma)$. Otherwise, a flip in a path or cycle violating the equation would contradict the optimality of π or of σ along $\ell(\pi, \sigma)$. Therefore, all the vectors d_{γ_j} must be orthogonal to $\ell(\pi, \sigma)$. Hence, the direction of $d_\pi - d_\sigma$ is the same as the one of d_{γ_j} for every path γ_j . If a path, say γ_1 , starts at some a_j and ends at some a_i , then $d_{\gamma_1} = a_j - a_i$, which concludes the proof. \square

Remark. It follows that if A is in general position then H has exactly one alternating path, and the pair a_i, a_j is unique.

Lemma 3. *i) $\mathcal{D}_{B,A}$ has at most $4m(n-m)$ unbounded regions.*

ii) Every region in $\mathcal{D}_{B,A}$ has at most $m(n-m)$ edges.

iii) Every vertex in $\mathcal{D}_{B,A}$ has degree at most $2m(n-m)$.

iv) Any convex path can intersect at most $m(n-m) + n(n-1)$ regions of $\mathcal{D}_{B,A}$, i.e., while translating B along any convex path, the optimal partial matching can change at most $m(n-m) + n(n-1)$ times.

Proof. i) Take a bounding box that encloses all the vertices of the diagram. By Theorem 1, every edge of the bounding box crosses at most $m(n-m) + 1$ regions of $\mathcal{D}_{B,A}$. The edges of the box traverse only unbounded regions, and cross every unbounded region exactly once, except for the coincidences of the last region traversed by an edge and the first region traversed by the next edge.

ii) By Lemma 2, the normal vector of every edge of a region corresponding to the injection π is a multiple of $a_j - a_i$ for some $a_i \in \pi(B)$ and $a_j \notin \pi(B)$. There are exactly $m(n-m)$ such possibilities.

iii) Let v be a vertex of $\mathcal{D}_{B,A}$. Draw two generic parallel lines close enough to each other to enclose v and no other vertex. Each edge adjacent to v is crossed by one of the lines, and by Theorem 1 each of these lines crosses at most $m(n-m)$ edges.

iv) We use the following property that was observed in Rote's proof of Theorem 1. Suppose that we translate B along a line in some direction v . Rank the points of A by their order in the v -direction, i.e., $a < a'$ means that $\langle a, v \rangle < \langle a', v \rangle$ (for simplicity, assume that v is generic so there are no ties). Let Φ denote the sum of the ranks of the m points of A that participate in the optimal partial match. As Rote has shown, whenever the optimal assignment changes, Φ must increase. Now follow our convex path γ , which, without loss of generality, can be assumed to be polygonal. As we traverse an edge of γ , Φ obeys the above property, increasing every time we cross into a new region of $\mathcal{D}_{B,A}$. When we turn (counterclockwise) at a vertex of γ , the ranking of A may change, but each such change consists of a sequence of swaps of consecutive elements in the present ranking. At each such swap, Φ can decrease by at most 1. Since γ is convex, each pair of points of A can be swapped at most twice, so the total

decrease in Φ is at most $2\binom{n}{2} = n(n-1)$. Hence, the accumulated increase in Φ , and thus also the total number of regions of $\mathcal{D}_{B,A}$ crossed by γ , is at most $(n + (n-1) + \dots + (n-m+1)) - (1 + 2 + \dots + m) + n(n-1)$. \square

In the remainder of this section, we focus on establishing a global bound on the complexity of the diagram $\mathcal{D}_{B,A}$. We begin by deriving the following technical auxiliary results.

Lemma 4. *Let π be an optimal assignment for a fixed translation $t \in \mathbb{R}^2$.*

- i) *There is no cyclic sequence $(i_1, i_2, \dots, i_k, i_1)$ satisfying $\|b_{i_j} + t - a_{\pi(i_j)}\| < \|b_{i_j} + t - a_{\pi(i_{j+1})}\|$ for all $j \in \{1, \dots, k\}$ (modulo k).*
- ii) *Each point of $B + t$ is matched to one of its m nearest neighbors in A .*
- iii) *At least one point in $B + t$ is matched to its nearest neighbor in A .*
- iv) *There exists an ordering $\langle b_1, \dots, b_m \rangle$ of the elements of B , such that each b_k is assigned by π to one of its k nearest neighbors in A , for $k = 1, \dots, m$.*

Proof. i) For the sake of contradiction, we assume that there exists a cyclic sequence that satisfies all the prescribed inequalities. Consider the assignment σ defined by $\sigma(i_j) = \pi(i_{(j-1) \bmod k})$ for all $j \in \{1, \dots, k\}$ and $\sigma(\ell) = \pi(\ell)$ for all other indices ℓ . Since π is a one-to-one matching, we have that $\pi(i_j) \neq \pi(i_{j'})$ for all different $j, j' \in \{1, \dots, k\}$ and, consequently, σ is one-to-one as well. It is easily checked that $f(\sigma, t) < f(\pi, t)$, contradicting the optimality of π .

ii) For contradiction, assume that for some point $b \in B$, $b + t$ is not matched by π to one of its m nearest neighbors in A . Then, at least one of these neighbors, say a , cannot be matched (because these m points can be claimed only by the remaining $m-1$ points of $B+t$). Thus, we can reduce the cost of π by matching $b+t$ to a , a contradiction that establishes the claim.

iii) Again we assume for contradiction that π does not match any of the points of $B+t$ to its nearest neighbor in A . We construct the following cyclic sequence in the matching π . We start at some arbitrary point $b_1 \in B$, and denote by a_1 its nearest neighbor in A (to simplify the presentation, we do not explicitly mention the translation t in what follows). By assumption, b_1 is not matched to a_1 . If a_1 is also not claimed in π by any of the points of B , then b_1 could have claimed it, thereby reducing the cost of π , which is impossible. Let then b_2 denote the point that claims a_1 in π . Again, by assumption, a_1 is not the nearest neighbor a_2 of b_2 , and the preceding argument then implies that a_2 must be claimed by some other point b_3 of B . We continue this process, and obtain an alternating path $(b_1, a_1, b_2, a_2, b_3, \dots)$ such that the edges (b_i, a_i) are not in π , and the edges (b_{i+1}, a_i) belong to π , for $i = 1, 2, \dots$. The process must terminate when we reach a point b_k that either coincides with b_1 , or is such that its nearest neighbor is among the already encountered points a_i , $i < k$. We thus obtain a cyclic sequence as in part i), reaching a contradiction.

iv) Start with some point $b_1 \in B$ such that $b_1 + t$ goes to its nearest neighbor a_1 in A in the optimal partial-matching π ; such a point exists by part iii). Delete b_1 from B , and a_1 from A . The optimal matching of $B \setminus \{b_1\}$ into $A \setminus \{a_1\}$ (relative to t) is equal to the restriction of π to the points in $B \setminus \{b_1\}$, because

otherwise we could have improved π itself. We apply part iii) to the reduced sets, and obtain a second point $b_2 \in B \setminus \{b_1\}$ whose translation $b_2 + t$ is matched to its nearest neighbor a_2 in $A \setminus \{a_1\}$, which is either its first or second nearest neighbor in the original set A . We keep iterating this process until the entire set B is exhausted. At the k -th step we obtain a point $b_k \in B \setminus \{b_1, \dots, b_{k-1}\}$, such that the nearest neighbor a_k in $A \setminus \{a_1, \dots, a_{k-1}\}$ is matched to b_k by π , so a_k is among the k nearest neighbors in A of $b_k + t$. \square

Observe, that the geometric properties in Lemma 4 can be interpreted in purely combinatorial terms. Indeed, for t fixed, associate with each $b_i \in B$ an ordered list $L_t(b_i)$, called its *preference list*, which consists of the points of A sorted by their distances from $b_i + t$. In general, given m such ordered lists on n elements, an injective assignment from $\{1, \dots, m\}$ to $\{1, \dots, n\}$ such that there is no cycle as in part i) is called *stable* or *Pareto efficient*. The problem of finding a stable matching was studied, for the case $m = n$, in the game theory literature under the name of the *House Allocation Problem* [14]. Note also that the proofs of parts ii)–iv) can be carried out in this abstract setting, and hold for any stable matching. Note that part iv) immediately yields an upper bound of $m!$ on the number of stable matchings and, in addition, implies that only the first m elements of each $L_t(b_i)$ are relevant. This bound is tight for the combinatorial problem, since if the ordered lists all coincide there are $m!$ different stable matchings. A recent article, motivated by the extended abstract [8] prior to this work, studied this combinatorial problem and derived the following.

Lemma 5 (Asinowski et al. [2]). *The number of elements that belong to some stable matching on m ordered preference lists is at most $m(\ln m + 1)$.*

The properties derived so far imply the following significantly improved upper bound on the complexity of $\mathcal{D}_{B,A}$.

Theorem 2. *The combinatorial complexity of $\mathcal{D}_{B,A}$ is $O(n^2 m^{3.5} (e \ln m + e)^m)$.*

Proof. The proof has two parts. First, we identify a convex subdivision K such that in each of its regions the first m elements of the ordered preference lists $L_t(b)$ of neighbors of each $b+t$, according to their distance from $b+t$, are fixed for all $b \in B$. We show that the complexity of K is only polynomial; specifically, it is $O(n^2 m^4)$. Second, we give an upper bound on how many regions of $\mathcal{D}_{B,A}$ can intersect a given region of K , using Lemma 5. Together, these imply an upper bound on the complexity of $\mathcal{D}_{B,A}$. The proof of the first part, which is based on a somewhat non-standard application of the Clarkson-Shor technique, is omitted in this version. We now consider all possible translations t in the interior of some fixed region τ of K and their corresponding optimal matchings. Lemma 4(i) ensures that all of them must be stable with respect to the fixed preference lists $L_t(b)$, for $b \in B$, over $t \in \tau$. In addition, Lemma 1 ensures that we only need to bound the number of different image sets of such stable matchings. Using the bound in Lemma 5, we can derive that the number of optimal matchings for translations in τ is then $O\left(\binom{m(\ln m + 1)}{m}\right) = O\left(\frac{m^m (\ln m + 1)^m}{m!}\right) = O\left(\frac{(e \ln m + e)^m}{\sqrt{m}}\right)$,

where in the second step we used Stirling’s approximation. Hence, by multiplying this bound by the number of regions in K , we conclude that the number of assignments corresponding to optimal matchings, and thus also the complexity of $\mathcal{D}_{B,A}$, is at most $O(n^2 m^{3.5} (e \ln m + e)^m)$. \square

The following proposition (proof omitted in this version) sets an obstruction for the combinatorial approach alone to yield a polynomial bound for $\mathcal{D}_{B,A}$.

Proposition 1. *For every $n \geq \lfloor \frac{m}{2} \rfloor + m$, there exists m preference lists of $\{1, \dots, n\}$ with $\Omega\left(\frac{2^m}{\sqrt{m}}\right)$ different images of stable matchings.*

3 Finding a local minimum of the partial-matching RMS-distance under translation

The high-level algorithm. We now concentrate on the algorithmic problem of computing, in polynomial time, a local minimum of the partial-matching RMS-distance under translation.

We “home in” on a local minimum of $F(t)$ by maintaining a vertical slab I in the plane that is known to contain such a local minimum in its interior, and by repeatedly shrinking it until we obtain a slab I^* that does not contain any vertex of $\mathcal{D}_{B,A}$. That is, any (vertical) line contained in I^* intersects the same sequence of regions, and, by Theorem 1, the number of these regions is $O(mn)$. We compute these regions, find the optimal partial matching assignment in each region, and the corresponding explicit (quadratic) expression of $F(t)$, and search for a local minimum within each region.

A major component of the algorithm is a procedure, that we call $\Pi_1(\ell)$, which, for a given input line ℓ , constructs the intersection of $\mathcal{D}_{B,A}$ with ℓ , computes the global minimum t^* of F on ℓ , and determines a side of ℓ , in which F attains strictly smaller values than $F(t^*)$. If no such decrease is found in the neighborhood of t^* then it is a local minimum of F , and we stop. Using Lemma 2 and the Hungarian algorithm [6, 10], $\Pi_1(\ell)$ runs in $O(m^5 n^2)$ time.

We use this “decision procedure” as follows. Suppose we have a current vertical slab I , bounded on the left by a line ℓ^- and on the right by a line ℓ^+ . We assume that Π_1 has been executed on ℓ^- and on ℓ^+ , and that we have determined that F assumes smaller values than its global minimum on ℓ^- to the right of ℓ^- , and that it assumes smaller values than its global minimum on ℓ^+ to the left of ℓ^+ . This is easily seen to imply that F must contain a local minimum in the interior of I . (We note that just finding a *local* minimum of F along ℓ^+ or ℓ^- is not sufficient; see the full version for details.) Let ℓ be some vertical line passing through I . We run Π_1 on ℓ . If it determines that F attains smaller values to its left (resp., to its right), we shrink I to the slab bounded by ℓ^- and ℓ (resp., the slab bounded by ℓ and ℓ^+). By what has just been argued, this ensures that the new slab also contains a local minimum of F in its interior.

To initialize the slab I , we choose an arbitrary *horizontal* line λ , and run Π_1 on λ , to find the sequence S of its intersection points with the edges of $\mathcal{D}_{B,A}$. We

run a binary search through S , where at each step we execute Π_1 on the vertical line through the current point. When the search terminates, we have a vertical slab I_0 whose intersection with λ is contained in a single region σ_0 of $\mathcal{D}_{B,A}$.

After this initialization, we find the region σ_1 that lies directly above σ_0 and that the final slab I^* should cross. In general, there are possibly many such regions, but fortunately, by Lemma 3(ii), their number is only at most $m(n-m)$.

To find σ_1 , we compute the boundary of σ_0 ; this is done similarly to the execution of Π_1 ; see the full version for details. Once we have explored the boundary of σ_0 , we take the sequence of all vertices of σ_0 , and run a Π_1 -guided binary search on the vertical lines passing through them, exactly as we did with the vertices of S , to shrink I_0 into a slab I_1 , so that σ_0 intersects I_1 in a trapezoid (or a triangle), with a single (portion of an) edge at the top and a single edge at the bottom. This allows us to determine σ_1 , which is the region lying on the other (higher) side of the top edge, in $O(m^5 n^2 \log(mn))$ time. A symmetric variant of this procedure will find the region lying directly below σ_0 in the final slab.

We repeat the previous step to find the entire stack of $O(nm)$ regions that I^* crosses, where each step shrinks the current slab and then crosses to the next region in the stack. Once this is completed, we find a local minimum within I^* as explained above. Again, details are omitted in this version.

In summary, we have the following main result of this section.

Theorem 3. *Given two finite point sets A, B in \mathbb{R}^2 , with $n = |A| > |B| = m$, and such that for every two pairs $(a_1, a_2), (a_3, a_4) \in A \times A$ the vectors $a_1 - a_2$ and $a_3 - a_4$ are non-parallel, a local minimum of the partial-matching RMS-distance under translation can be computed in $O(m^6 n^3 \log n)$ time.*

4 Finding a local minimum of the Hausdorff RMS-distance under translation

In this section, we turn to the simpler problem involving the Hausdorff RMS-distance, and present efficient algorithms for computing a local minimum of the RMS function in one and two dimensions. Due to lack of space, most of the material in this section is omitted, and we only provide here a high-level review of our algorithms.

The one-dimensional unidirectional case. Let $N_A(b+t)$ be the nearest neighbor in A of $b+t$, for $b \in B$, and $t \in \mathbb{R}$. The function $r(t) := RMS(B+t, A) = \sum_{b \in B} (b+t - N_A(b+t))^2$ is continuous and piecewise parabolic, with $O(mn)$ non-smooth breakpoints, which are the breakpoints of the step functions $N_A(b+t)$. For any given t_0 , it is easy to compute, in $O(m \log n)$ time, the derivative $r'(t_0)$, or its left and right one-sided versions $r'(t_0)^-, r'(t_0)^+$ (when t_0 is a breakpoint). A simple observation is that if $I = [t_1, t_2]$ is an interval satisfying $r'(t_1)^+ < 0$ and $r'(t_2)^- > 0$ then I contains a local minimum of r . We thus start with a large interval I that contains all breakpoints of r , and keep shrinking it, halving

the number of breakpoints in I in each step, until it contains only linearly many breakpoints, in which case r can be constructed explicitly over I , and searched for a local minimum, in near-linear time. Specifically, we obtain:

Theorem 4. *Given two finite point sets A, B on the real line, with $|A| = n$ and $|B| = m$, a local minimum of the unidirectional RMS distance under translation from B to A can be obtained in time $O(m \log^2 n + n \log n)$.*

The one-dimensional bidirectional case. Simple extensions of the procedure given above apply to the two variants of the minimum bidirectional Hausdorff RMS-distance, as defined in the introduction. Omitting the fairly routine details of these extensions, we obtain:

Theorem 5. *Given two finite point sets A, B on the real line, with $|A| = n$ and $|B| = m$, a local minimum under translation of the L_1 -bidirectional or L_∞ -bidirectional RMS distance between A and B , can be computed in time $O((n \log m + m \log n) \log \min \{m, n\})$.*

Minimum Hausdorff RMS-distance under translation in two dimensions. Here the function $r(t) := \text{RMS}(B+t, A) = \sum_{b \in B} \|b+t - N_A(b+t)\|^2$ induces a convex subdivision of the plane, where in each of its regions σ , all the m values $N_A(b+t)$, for $b \in B$, are fixed for $t \in \sigma$. This subdivision is simply the overlay M of the m shifted copies $\mathcal{V}(A - b)$, for $b \in B$, of the Voronoi diagram of A . These copies have a total of $O(mn)$ edges, and their overlay has thus complexity $O(m^2 n^2)$ (which is tight in the worst case). Over each region of M , $r(t)$ is a quadratic function (a paraboloid), and the explicit expression for $r(t)$ can be updated in $O(1)$ time as we cross from one region to an adjacent one.

The goal is to search for a local minimum of r without explicitly constructing these many features of M . Similarly to the one-dimensional case, we maintain a vertical slab I , known to contain a local minimum, and keep shrinking it until it contains no vertices of M . In this case it overlaps only $O(mn)$ regions of M , vertically stacked above one another, and it is straightforward to enumerate all of them, get the explicit expressions of r over each of them, and search for a local minimum in each part, in a total of $O(mn)$ time.

The shrinking of I is performed in two phases. We first enumerate all $O(mn)$ Voronoi vertices of the original diagrams, and run a binary search through them, as above. The resulting intermediate slab contains no original vertices, so the edges that cross it behave like lines. They might still intersect at $O(m^2 n^2)$ points within I , but we can run a binary search through them efficiently, using the (dual version of the) *slope selection* algorithm of [4], so that each step takes only $O(mn \log mn)$ time.

Concretely, we obtain:

Theorem 6. *Given two finite point sets A, B in \mathbb{R}^2 , with $|A| = n$ and $|B| = m$, a local minimum of the unidirectional Hausdorff RMS-distance from B to A under translation can be computed in time $O(mn \log^2 mn)$.*

The bidirectional variants can be handled in much the same way, and, omitting the details, we get:

Theorem 7. *Given two finite point sets A, B in \mathbb{R}^2 , with $|A| = n$ and $|B| = m$, a local minimum of the L_1 -bidirectional or the L_∞ -bidirectional Hausdorff RMS-distance between A and B under translation can be computed in $O(mn \log^2 mn)$ time.*

References

1. P. K. Agarwal, S. Har-Peled, M. Sharir and Y. Wang, Hausdorff distance under translation for points, disks, and balls, *ACM Trans. on Algorithms* 6 (2010), 1–26.
2. A. Asinowski, B. Keszegh and T. Miltzow, Counting houses of Pareto optimal matchings in the House Allocation Problem, arXiv:1401.5354v2.
3. P. J. Besl and N. D. McKay, A method for registration of 3-d shapes, *IEEE Trans. Pattern Anal. Mach. Intell.* 14 (1992), 239–256.
4. R. Cole, J. Salowe, W. Steiger, and E. Szemerédi, An optimal-time algorithm for slope selection, *SIAM J. Comput.* 18 (1989), 792–810.
5. A. Dumitrescu, G. Rote and C.D. Tóth, Monotone paths in planar convex subdivisions and polytopes, in *Discrete Geometry and Optimization*, Károly Bezdek, Antoine Deza, and Yinyu Ye, editors, Fields Institute Communications 69, Springer-Verlag, 2013, pp. 79–104.
6. J. Edmonds and R. M. Karp., Theoretical Improvements in Algorithmic Efficiency for Network Flow Problems, *J. ACM* 19, 2 (April 1972), pp. 248–264.
7. E. Ezra, M. Sharir and A. Efrat, On the ICP Algorithm, *Comput. Geom. Theory Appl.* 41 (2008), 77–93.
8. M. Henze, R. Jaume, and B. Keszegh, On the complexity of the partial least-squares matching Voronoi diagram, in *Proc. 29th European Workshop Comput. Geom. (EuroCG'13)*, pp. 193–196, 2013.
9. I. Jung and S. Lacroix, A robust interest points matching algorithm, in *Proc. ICCV'01*, volume 2, pages 538–543, 2001.
10. H. W. Kuhn, The Hungarian method for the assignment problem, *Naval Research Logistics Quarterly* 2(1–2) (1955), 83–97.
11. J. M. Phillips and P. K. Agarwal, On bipartite matching under the RMS distance, in *Proc. 18th Canadian Conf. Comput. Geom. (CCCG'06)*, pp. 143–146, 2006.
12. G. Rote, Partial least-squares point matching under translations, in *Proc. 26th European Workshop Comput. Geom. (EuroCG'10)*, pp. 249–251, 2010.
13. G. Rote, Long monotone paths in convex subdivisions, in *Proc. 27th European Workshop Comput. Geom. (EuroCG'11)*, pp. 183–184, 2011.
14. L. S. Shapley and H. Scarf, On cores and indivisibility, *J. Math. Economics*, 1(1974), 23–37.
15. S. Umeyama, Least-squares estimation of transformation parameters between two point patterns, *IEEE Trans. Pattern Anal. Mach. Intell.* 13(4) (1991), 376–380.
16. K. Zikan and T. M. Silberberg, The Frobenius metric in image registration, in L. Shapiro and A. Rosenfeld, editors, *Computer Vision and Image Processing*, pp. 385–420, Elsevier, 1992.