Minimum L_∞ Hausdorff Distance of Point Sets Under Translation: Generalizing Klee's Measure Problem

Timothy M. Chan
Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

Abstract

We present a (combinatorial) algorithm with running time close to $O(n^d)$ for computing the minimum directed L_∞ Hausdorff distance between two sets of n points under translations in any constant dimension d. This substantially improves the best previous time bound near $O(n^{5d/4})$ by Chew, Dor, Efrat, and Kedem from more than twenty years ago. Our solution is obtained by a new generalization of Chan's algorithm [FOCS'13] for Klee's measure problem.

To complement this algorithmic result, we also prove a nearly matching conditional lower bound close to $\Omega(n^d)$ for combinatorial algorithms, under the Combinatorial k-Clique Hypothesis.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Hausdorff distance, geometric optimization, Klee's measure problem, fine-grained complexity

Digital Object Identifier 10.4230/LIPIcs.SoCG.2023.24

Related Version *Full Version*: `https://arxiv.org/abs/2303.09122`

Funding *Timothy M. Chan*: Work supported in part by NSF Grant CCF-2224271.

1 Introduction

This paper is about the following problem:

▶ **Problem 1.** (L_∞ Translational Hausdorff) *Given a set P of n points and a set Q of m points in $\mathbb{R}^d$, compute the minimum directed L_∞ Hausdorff distance from P to Q under translation, i.e., compute $\min_{v\in\mathbb{R}^d} \vec{h}_\infty(P+v,Q)$ where $\vec{h}_\infty(P,Q) := \max_{p\in P}\min_{q\in Q}\|p-q\|_\infty$.*

The problem has been extensively studied in computational geometry in the 1990s. The analogous problem for *undirected* Hausdorff distance (defined as $h_\infty(P,Q) = \max\{\vec{h}_\infty(P,Q), \vec{h}_\infty(Q,P)\}$) is reducible [27] to the directed version if $m = \Theta(n)$. The motivation lies in measuring the resemblance between two geometric objects represented as point clouds; furthermore, a connection with an even more fundamental problem, Klee's measure problem (see next page), provides added theoretical interest (and is what attracted this author's attention in the first place). Huttenlocher and Kedem [22] introduced the problem and presented the first algorithms for $d = 2$ (a subsequent paper [23] also examined variants in L_2). Chew and Kedem [17] gave an improved algorithm with running time $O(mn\log^2(mn))$ for $d = 2$ (in L_∞), and generalized the algorithm to any constant dimension d with running time $O((mn)^{d-1}\log^2(mn))$. Chew, Dor, Efrat, and Kedem [15] described further improved algorithms in higher dimensions: in the main case $m = n$, their time bounds were $O(n^3\log^2 n)$ for $d = 3$, $O(n^{(4d-2)/3}\log^2 n)$ for $4 \le d \le 7$, and $O(n^{5d/4}\log^2 n)$ for any constant $d \ge 8$. The exponent $5d/4$ looks peculiar, and naturally raises the question of whether further improvements are still possible, but none has been found in the intervening two decades (except in the logarithmic factors [12, 14]).

Many other variants of Problem 1 have also been considered, for example, using other metrics such as L_2 (as already mentioned above), allowing rotation and/or scaling besides translation, handling other objects besides points, allowing approximations, etc. (e.g., see

39th International Symposium on Computational Geometry (SoCG 2023).
Editors: Erin W. Chambers and Joachim Gudmundsson; Article No. 24; pp. 24:1–24:13
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Table 1 Previous upper bounds [15] and new upper bounds for Problem 1 in the $m = n$ case, ignoring polylogarithmic factors.

dimension	2	3	4	5	6	7	8	9	10	11	12	$\cdots$
prev. bound	n^2	n^3	$n^{4.66\cdots}$	n^6	$n^{7.33\cdots}$	$n^{8.66\cdots}$	n^{10}	$n^{11.25}$	$n^{12.5}$	$n^{13.75}$	n^{15}	$\cdots$
new bound			n^4	n^5	n^6	n^7	n^8	n^9	n^{10}	n^{11}	n^{12}	$\cdots$

[16, 4, 21, 24, 19, 18, 3]). Several other alternatives to the Hausdorff distance have also been popularly studied in computational geometry, such as the Earth mover distance and the Fréchet distance. We will ignore all these variants in the present paper, focusing only on exact directed L_∞ Hausdorff distance for point sets under translation.

Our new result is an algorithm for Problem 1 running in $O(n^d(\log\log n)^{O(1)})$ time (using randomization) for any constant $d \geq 3$ in the main $m = n$ case (or with one extra logarithmic factor if randomization is not allowed). The exponent d is thus a *substantial* improvement over Chew et al.'s previous exponents for *every* $d \geq 4$; see Table 1. In the general case, the running time of our algorithm is $O((mn)^{d/2}(\log\log(mn))^{O(1)})$.

Connection to a generalized Klee's measure problem. It suffices to focus on the decision problem: deciding whether the minimum is at most a given value r. The original problem reduces to the decision problem, at the expense of one extra logarithmic factor in the running time by a well-known technique of Frederickson and Johnson [20] (in fact, when $d \geq 4$, a standard binary search suffices, since the optimal value lies in a universe of $O((mn)^2)$ possible values which we can explicitly enumerate). In some cases, the extra logarithmic factor can even be eliminated by a randomized optimization technique [12].

Equivalently, we want to decide whether there exists a vector $v \in \mathbb{R}^d$ with $P + v \subseteq Q + [-r, r]^d$ (where "+" denotes the Minkowski sum when it is clear from the context). Assume (by rescaling) that $r = 1/2$. Let $\mathcal{Q}$ be the set of unit hypercubes $\{q + [-1/2, 1/2]^d : q \in Q\}$, and let $S^* := \bigcup_{B \in \mathcal{Q}} B$. The condition is equivalent to $P + v \subseteq S^*$, i.e., $v \in \bigcap_{p \in P}(S^* - p)$. Thus, the decision problem is equivalent to the following:

▶ **Problem 2.** (L_∞ Translational Hausdorff Decision) *Given a set P of n points and a set $\mathcal{Q}$ of m unit hypercubes*[1] *in $\mathbb{R}^d$, decide whether $\bigcap_{p \in P}(S^* - p) = \emptyset$, where $S^* := \bigcup_{B \in \mathcal{Q}} B$.*

For each $B \in \mathcal{Q}$ and $p \in P$, create a new unit hypercube $B - p$ and give this hypercube the color p. Problem 2 then immediately reduces to the following problem on $N = mn$ colored unit hypercubes: decide whether $\bigcap_\chi S_\chi = \emptyset$, where $S_\chi := \bigcup_{B \,\in\, \mathcal{B} \text{ with color } \chi} B$. (In other words, we want to decide whether there exists a "colorful" point that lies in hypercubes of all colors.)

The unit hypercube case in turn reduces to the case of *orthants* (i.e., d-sided boxes which are unbounded in one direction along each axis): we can build a uniform grid of unit-side length and solve the subproblem inside each grid cell, but inside a grid cell, a unit hypercube is identical to an orthant. (We can ignore grid cells that do not intersect hypercubes of all colors.) Since a unit hypercube intersects only $O(1)$ grid cells, these subproblems have total input size $O(N)$. All this motivates the definition of the following problem(s) on colored orthants:

[1] Throughout this paper, all hypercubes and boxes are axis-aligned.

▶ **Problem 3** (Generalized Klee's Measure Problem). *Given a set $\mathcal{B}$ of N colored orthants in $\mathbb{R}^d$,*

(a) *decide whether $\bigcap_\chi S_\chi = \emptyset$,*

(b) *or more generally, compute a point of maximum or minimum depth among the S_χ's (i.e., a point in the most or least number of regions S_χ),*

(c) *or alternatively, compute the volume of $\bigcap_\chi S_\chi$,*

where $S_\chi := \bigcup_{B \in \mathcal{B} \text{ with color } \chi} B$.

To recap, if Problem 3(a) can be solved in $T(N)$ time, then Problem 2 can automatically be solved in $O(T(N)) = O(T(mn))$ time (assuming superadditivity of $T(\cdot)$).

Problem 3(a) is a generalization of the *box coverage problem*: determine whether the union of a set of N boxes in $\mathbb{R}^d$ covers the entire space. This is because $\bigcap_\chi S_\chi = \emptyset$ iff $\bigcup_\chi \overline{S_\chi} = \mathbb{R}^d$, and a box can be expressed as the complement of a union of at most $2d$ orthants (we use $\overline{S}$ to denote the complement of a set S). Similarly, Problem 3(b) is a generalization of the *box depth problem*: determine the minimum or maximum depth among N boxes in $\mathbb{R}^d$. Problem 3(c) is a generalization of *Klee's measure problem*: compute the volume of the union of a set of N boxes in $\mathbb{R}^d$. This generalization allows us to compute the volume of the union of more general shapes, so long as each shape can be expressed as the complement of a union of orthants. (Note that we can clip to a bounding box to ensure that the volume is finite.)

The original Klee's measure problem has been extensively studied in computational geometry [2, 8, 13, 14, 28, 30]. The best known algorithm by Chan [14] for Klee's measure problem runs in $O(N^{d/2})$ time, based on a clever but simple divide-and-conquer. The box coverage and box depth problem can be solved by similar algorithms, and in fact with slightly lower time bounds by polylogarithmic factors using table lookup and bit packing tricks [14].

For $d \leq 3$, a union of orthants has linear combinatorial complexity [7] and can be constructed in near linear time. Thus, a straightforward way to solve Problem 3 is to first construct all the regions S_χ explicitly, decompose each $\overline{S_\chi}$ as a union of disjoint boxes, and then run a known algorithm for Klee's measure problem on the resulting $O(N)$ boxes. With this approach, Problem 3 can be solved in $O(N \log N)$ time for $d = 2$, and $O(N^{3/2})$ time for $d = 3$; consequently, Problem 2 can be solved in $O((mn)\log(mn))$ time for $d = 2$, and $O((mn)^{3/2})$ time for $d = 3$. This was essentially how the previous known 2D and 3D algorithms by Chew and Kedem [17] and Chew et al. [15] were designed.

However, for $d \geq 4$, a union of N orthants may have $\Theta(N^{\lfloor d/2 \rfloor})$ combinatorial complexity in the worst case [7]. So, a two-stage approach that explicitly constructs all the regions S_χ and then invokes an algorithm for Klee's measure problem would be too slow! Chew et al. [15] adapted Overmars and Yap's algorithm for Klee's measure problem [28] in a nontrivial way to obtain an $O(N^{5d/8} \log N)$-time algorithm for Problem 3, and consequently an $O(n^{5d/4} \log n)$-time algorithm for Problem 2 when $m = n$.

We present a new algorithm that solves Problem 3(c) in $O(N^{d/2} \log^{d/2} N)$ time, matching the known time bound for the original Klee's measure problem up to logarithmic factors. In fact, for Problem 3(a,b), the polylogarithmic factor can be lowered to poly-$\log\log N$ factors using table lookup and bit packing tricks. Consequently, we obtain an $O(n^d (\log\log n)^{O(1)})$ time bound for Problem 2 when $m = n$, or $O((mn)^{d/2}(\log\log(mn))^{O(1)})$ in general. Our result is obtained by directly modifying Chan's divide-and-conquer algorithm for Klee's measure problem [14]. The adaptation is not straightforward and uses interesting new ideas. As mentioned, we cannot afford to separate into two stages. Instead, within a single recursive process, we will handle two types of objects simultaneously, (i) the input orthants, and (ii) features of the regions S_χ that have been found during the process. The analysis of the recurrence is more delicate (though the overall algorithm remains simple).

Conditional lower bounds. In the other direction, recently in SoCG'21, Bringmann and Nusser [10] proved an $\Omega((mn)^{1-\delta})$ conditional lower bound for Problem 1–2 for $d = 2$ for an arbitrarily small constant $\delta > 0$, under the Orthogonal Vectors (OV) Hypothesis [29] (in particular, it holds under the Strong Exponential-Time Hypothesis (SETH) [29]). This showed that Chew and Kedem's upper bound for $d = 2$ is likely near optimal [17]. However, Bringmann and Nusser did not obtain any lower bound in higher dimensions.

As observed by Chan [13], Klee's measure problem and the box coverage problem have an $\Omega(N^{d/2-\delta})$ lower bound for combinatorial algorithms under the *Combinatorial k-Clique Hypothesis*, which states that there is no $O(n_0^{k-\delta})$-time combinatorial algorithm for detecting a k-clique in a graph with n_0 vertices, for any constant $k \geq 3$. The notion of "combinatorial" algorithms is not mathematically well-defined, but intuitively it refers to algorithms that avoid the use of fast matrix multiplication (such as Strassen's algorithm); all algorithms in this paper and in Chew et al.'s previous paper fulfill this criterion. (Recently, Künnemann [26] obtained new lower bounds for arbitrary, noncombinatorial algorithms for Klee's measure problem under the "k-Hyperclique Hypothesis", but his bounds are not tight for $d \geq 4$. See also [11, 6, 5] for conditional lower bounds for other related geometric problems. See [25] for a recent example of the usage of the Combinatorial k-Clique Hypothesis in computational geometry, and [1, 9] for other examples involving the Combinatorial k-Clique Hypothesis outside of geometry.)

Since our algorithms for Problem 3 have near $N^{d/2}$ running time, they are near optimal for combinatorial algorithms under the Combinatorial k-Clique Hypothesis. However, this does not necessarily imply optimality of our algorithms for Problem 1 or 2.

In the second part of this paper, we prove that Problems 1–2 have a conditional lower bound of $\Omega(n^{d-\delta})$ for $m = n$, or $\Omega((mn)^{d/2-\delta})$ for $m = n^\gamma$ for any fixed $\gamma \leq 1$, for combinatorial algorithms under the Combinatorial k-Clique Hypothesis. This shows that our combinatorial algorithm for Problem 1 is also conditionally near optimal. While the previous conditional lower bound for Klee's measure problem by Chan [13] was obtained by reduction from d-clique, we will reduce from clique of arbitrarily large constant size. Our new reduction is more challenging and more interesting, but still simple.

2 Algorithm

In this section, we present our new algorithm for the generalized Klee's measure problem (Problem 3) for any constant dimension $d \geq 4$. From this result, new algorithms for Problems 1–2 will immediately follow.

To solve Problem 3(c), we solve a generalization, where we are given a box "cell" γ and an extra set $\mathcal{E}$ of boxes, and we want to compute the volume of $\bigcap_\chi S_\chi \cap \bigcap_{E \in \mathcal{E}} \overline{E} \cap \gamma$. Initially, $\gamma = \mathbb{R}^d$ and $\mathcal{E} = \emptyset$. We assume that the coordinates of the input have been pre-sorted (this requires only an initial $O(N \log N)$ cost).

Call an orthant or a box *short* if some of its $(d-2)$-faces intersect γ's interior, *long* if it intersects γ's interior but is not short, and *trivial* if it does not intersect γ's interior or it completely contains γ.

Our algorithm is inspired by Chan's divide-and-conquer algorithm [14] for the original Klee's measure problem, with many similarities (for example, in how we use weighted medians to divide a cell) but also major new innovation (in how we reduce the number of long objects, and how we "convert" some objects of $\mathcal{B}$ into new objects in $\mathcal{E}$ during recursion). The analysis of our algorithm requires a new charging argument and recurrence, causing some extra logarithmic factors.

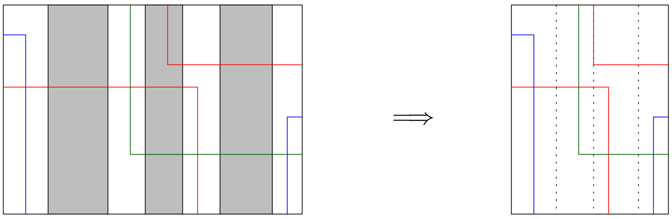

Figure 1 Shrinking the width of the shaded slabs to 0, to eliminate long boxes in $\mathcal{E}$.

Defining weights. Consider a $(d-2)$-face f of a short orthant of $\mathcal{B}$ such that f intersects γ's interior. If f is orthogonal to the i-th and the j-th axes, assign f a weight of $2^{(i+j)/d}$. Note that this weight is $\Theta(1)$, and so each short orthant of $\mathcal{B}$ contributes a total weight of $\Theta(1)$.

Similarly, consider a $(d-2)$-face f of a box of $\mathcal{E}$ such that f intersects γ's interior. If f is orthogonal to the i-th and the j-th axes, assign f a weight of $2^{(i+j)/d}/t$, where $t \geq 1$ is a parameter to be set later. Note that this weight is $\Theta(1/t)$, and so each short box of $\mathcal{E}$ contributes a total weight of $\Theta(1/t)$.

Let $T(N_{\text{long}}, W_{\text{short}})$ denote the time complexity of the problem, where N_{long} denotes the total number of long and trivial orthants in $\mathcal{B}$ and long and trivial boxes in $\mathcal{E}$, and W_{short} denotes the total weight of all short orthants in $\mathcal{B}$ and short boxes in $\mathcal{E}$. Note that the total number of orthants in $\mathcal{B}$ and boxes in $\mathcal{E}$ is upper-bounded by $O(N_{\text{long}} + tW_{\text{short}})$.

Reducing the number of long objects. First, the trivial orthants and boxes can be easily eliminated: We can remove all orthants of $\mathcal{B}$ and boxes of $\mathcal{E}$ that do not intersect γ's interior. If an entire color class of $\mathcal{B}$ does not intersect γ's interior, or if some box of $\mathcal{E}$ completely contains γ, we can return 0 as the answer. If there is an orthant of $\mathcal{B}$ completely containing γ, we can remove its color class from $\mathcal{B}$.

For each color χ, consider the long orthants of $\mathcal{B}$ with color χ; the union of these long orthants are defined by at most $2d$ orthants (since it is the complement of a box with at most $2d$ sides). Keep these $O(1)$ long orthants per color, and remove the rest. If there is a long orthant with color χ but no short orthant with that color, then $\overline{S_\chi} \cap \gamma$ is a box – add this box to $\mathcal{E}$ and remove the color class from $\mathcal{B}$ (in other words, we have "converted" an entire color class in $\mathcal{B}$ into a single box in $\mathcal{E}$). This step increases W_{short} by at most N_{long}/t. After this step, each remaining long orthant can be "charged" to a short orthant of the same color, and so the number of remaining long orthants of $\mathcal{B}$ is bounded by $O(1)$ times the number of short orthants, which is $O(W_{\text{short}})$.

Next, for each $i \in \{1, \ldots, d\}$, consider the long boxes of $\mathcal{E}$ having $(d-1)$-faces intersecting γ that are orthogonal to the i-th axis. Compute the union of these boxes by a linear scan after sorting, since this corresponds to computing the union of 1D intervals when projected to the i-th axis. The union forms a disjoint collection of slabs. Readjust all the i-th coordinates to shrink the width of these slabs to 0, without altering the volume of $\bigcap_\chi S_\chi \cap \bigcap_{E\in\mathcal{E}} \overline{E} \cap \gamma$, as illustrated in Figure 1. After doing this successively for every $i \in \{1, \ldots, d\}$, all long boxes of $\mathcal{E}$ are eliminated.

After this process, there are $O(W_{\text{short}})$ remaining long orthants of $\mathcal{B}$ and zero long boxes of $\mathcal{E}$. Thus, N_{long} is reduced to $O(W_{\text{short}})$. The weight of the short orthants of $\mathcal{B}$ may increase to at most $W_{\text{short}} + N_{\text{long}}/t$. We then have the following, for some constant c:

$$T(N_{\text{long}}, W_{\text{short}}) \leq T(cW_{\text{short}}, W_{\text{short}} + N_{\text{long}}/t) + O(N_{\text{long}} + tW_{\text{short}}). \tag{1}$$

Divide-and-conquer. Next, compute the weighted median m of the first coordinates of the $(d-2)$-faces of $\mathcal{B}$ and $\mathcal{E}$ intersecting γ's interior that are orthogonal to the first axis. Divide γ into two subcells γ_L and γ_R by the hyperplane $\{(x_1,\ldots,x_d) : x_1 = m\}$. Renumber the coordinate axes $1,2,\ldots,d$ to $2,\ldots,d,1$, and recursively solve the problem for γ_L and for γ_R.

To analyze the algorithm, consider a $(d-2)$-face f of $\mathcal{B}$ (resp. $\mathcal{E}$) orthogonal to the i-th and j-th axes with $i,j \neq 1$. After the axis renumbering, its weight changes from $2^{(i+j)/d}$ to $2^{(i-1+j-1)/d}$ (resp. from $2^{(i+j)/d}/t$ to $2^{(i-1+j-1)/d}/t$), i.e., the weight decreases by a factor of $2^{2/d}$.

Next consider a $(d-2)$-face f of $\mathcal{B}$ (resp. $\mathcal{E}$) orthogonal to the first and the j-th axes with $j \neq 1$. After the axis renumbering, its weight changes from $2^{(1+j)/d}$ to $2^{(d+j-1)/d}$ (resp. from $2^{(1+j)/d}/t$ to $2^{(d+j-1)/d}/t$), i.e., the weight increases by a factor of $2^{(d-2)/d}$. But when γ is divided into subcells γ_L and γ_R, the weight within each subcell decreases by a factor of 2; the net decrease in weight is thus a factor of $2^{2/d}$.

Hence, $W_{\rm short}$ decreases by a factor of $2^{2/d}$ in either subcell. (On the other hand, $N_{\rm long}$ may not necessarily decrease.) We then have

$$T(N_{\rm long}, W_{\rm short}) \;\le\; 2\,T(N_{\rm long}, W_{\rm short}/2^{2/d}) + O(N_{\rm long} + tW_{\rm short}). \tag{2}$$

Putting it together. By combining (2) and (1) and letting $T(N) := T(cN, N)$, we obtain

$$\begin{aligned} T(N) \;\le\; 2\,T(cN, N/2^{2/d}) + O(tN) \;&\le\; 2\,T(cN/2^{2/d}, N/2^{2/d} + cN/t) + O(tN) \\ &\le\; 2\,T\left(\tfrac{1+2c/t}{2^{2/d}}N\right) + O(tN). \end{aligned}$$

For the base case, we have $T(O(1)) = O(t^{d/2})$: when $W_{\rm short} = O(1)$, there are $O(1)$ orthants of $\mathcal{B}$ and $O(t)$ boxes of $\mathcal{E}$, and so the problem can be solved by running a known algorithm for Klee's measure problem on $O(t)$ boxes [14].

By the master theorem, the solution to the recurrence is

$$T(N) \;=\; O(t^{d/2} N^{1/\log_2(2^{2/d}/(1+2c/t))}) \;=\; O(t^{d/2} N^{d/2+O(1/t)}).$$

Choosing $t = \log N$ yields $T(N) = O(N^{d/2}\log^{d/2} N)$. This completes the description and analysis of the main algorithm.

Shaving logs by bit packing. For Problem 3(a), we can obtain a minor (but not-very-practical) improvement in the polylogarithmic factors by using more technical but standard bit-packing tricks, as we now briefly explain (see [14] for more details on these kinds of tricks): The main observation is that actual coordinate values do not matter here, only their relative order, so we can replace them with their ranks in the sorted list. Thus, the $O(N_{\rm long} + tW_{\rm short})$ input objects can be represented by $O((N_{\rm long} + tW_{\rm short})\log(N_{\rm long} + tW_{\rm short}))$ bits and packed in $O(((N_{\rm long} + tW_{\rm short})\log(N_{\rm long} + tW_{\rm short}))/w)$ words, assuming a w-bit word RAM model of computation. The $O(N_{\rm long} + tW_{\rm short})$ cost for various linear scans during recursion can be reduced to $O(((N_{\rm long} + tW_{\rm short})\log^2(N_{\rm long} + tW_{\rm short}))/w)$, since sorting k b-bit numbers can be done in $O((kb\log k)/w)$ time by a packed version of mergesort. Thus, the recurrence changes to $T(N) \;\le\; 2\,T(\frac{1+2c/t}{2^{2/d}}N) + O(1 + (tN\log^2(tN))/w)$.

For the base case, we can use a bit-packed version of Chan's algorithm for the box coverage problem on $O(t)$ boxes, which runs in $T(O(1)) = O(1 + (t/w)^{d/2}\log^{O(1)} w)$ time [14, Section 3.1]. The recurrence solves to $T(N) = O(N^{d/2+O(1/t)} \cdot (1 + (t/w)^{d/2}\log^{O(1)} w))$. Choosing $t = \log N$ yields $T(N) = O(N^{d/2} \cdot (1 + ((\log N)/w)^{d/2}\log^{O(1)} w))$.

The above may require nonstandard operations on w-bit words. By choosing $w = \delta \log N$ for a sufficiently small constant $\delta > 0$, such operations can be simulated in constant time via table lookup after preprocessing in $2^{O(w)} = N^{O(\delta)}$ time. Hence, we obtain the final time bound of $O(N^{d/2}(\log\log N)^{O(1)})$.

Problem 3(b) can be solved similarly, with some modification to the steps to reduce the number of long boxes in $\mathcal{E}$, but this is identical to the modification of Chan's algorithm for the original box depth problem [14, Section 3.1]. To summarize, we have obtained the following theorem:

▶ **Theorem 1.** *For any constant $d \geq 4$, Problem 3(c) can be solved in $O(N^{d/2}\log^{d/2} N)$ time, and Problems 3(a,b) can be solved in $O(N^{d/2}(\log\log N)^{O(1)})$ time.*

▶ **Corollary 2.** *For any constant $d \geq 4$, Problem 2 can be solved in $O((mn)^{d/2}(\log\log(mn))^{O(1)})$ time. Problem 1 can be solved in $O((mn)^{d/2}\log(mn)(\log\log(mn))^{O(1)})$ time deterministically, or in $O((mn)^{d/2}(\log\log(mn))^{O(1)})$ expected time with randomization.*

Proof. As mentioned in Section 1, Problem 1 reduces to Problem 2 by Frederickson and Johnson's technique [20] or by ordinary binary search, with an extra logarithmic factor.

Chan [12, Section 4.2] has described how to apply his randomized optimization technique to reduce the following problem to its decision problem without losing a logarithmic factor:

> Given N colored points in $\mathbb{R}^d$, find the smallest hypercube that contains points of all colors.

As noted in [12], Problem 1 reduces to this problem. On the other hand, as noted in Section 1, the decision version of this problem (equivalent to finding a point that is inside unit hypercubes of all colors) reduces to Problem 3(a), which we have just solved. ◀

3 Conditional Lower Bound

In this section, we prove a nearly matching conditional lower bound for Problems 1–2 for combinatorial algorithms under the Combinatorial k-Clique Hypothesis. We first introduce an intermediate problem which is more convenient to work with. Roughly speaking, Problem 2 considers the intersection of translates of a *single* shape (the shape being a union of unit hypercubes), whereas the problem below considers the intersection of translates of *multiple* shapes (each shape being a union of orthants).

▶ **Problem 4.** *Let $\mathcal{Z}$ be a set of shapes, where each shape is a union of orthants in $\mathbb{R}^d$. Let m be the total number of orthants over all shapes of $\mathcal{Z}$. Given a set $\mathcal{S}$ of n objects where each object is a translate of some shape in $\mathcal{Z}$, decide whether $\bigcap_{S\in\mathcal{S}} S = \emptyset$.*

▶ **Lemma 3.** *Problem 4 reduces to Problem 2 on $O(n)$ points and $O(m)$ unit hypercubes.*

Proof. Assume (by rescaling) that the coordinates of all orthants of $\mathcal{Z}$ and all translation vectors used in $\mathcal{S}$ are in $[0,1/2]$. In particular, if $\bigcap_{S\in\mathcal{S}} S$ is nonempty, it must contain a point in $[0,1]^d$. Inside $[0,1]^d$, each orthant of $\mathcal{Z}$ may be replaced by an equivalent unit hypercube.

Let $Z_1,\ldots,Z_\ell$ be the shapes of $\mathcal{Z}$. Let $\mathcal{B}_i$ be the unit hypercubes corresponding to the orthants defining Z_i (so that $(Z_i+t)\cap[0,1]^d = \bigcup_{B\in\mathcal{B}_i}(B+t)\cap[0,1]^d$ for any $t\in[0,1/2]^d$).

We construct an instance of Problem 2 as follows: For each $B\in\mathcal{B}_i$, add the shifted unit hypercube $B+u_i$ to $\mathcal{Q}$ where $u_i := (4i,0,\ldots,0)\in\mathbb{R}^d$. (This operation distributes objects in different classes $\mathcal{B}_i$ to different parts of space, since the vectors u_i are at least 4 units apart from each other along the first axis.) For each object $S \in \mathcal{S}$, if S is the translate $Z_i + t$, add the point $u_i - t$ to P. Lastly, add two auxiliary unit hypercubes $[0,1]^d$ and $u_{\ell+1} + [0,1]^d$ to $\mathcal{Q}$, and two auxiliary points u_0 and $u_{\ell+1}$ to P.

We solve Problem 2 on these points of P and these unit hypercubes of $\mathcal{Q}$, to determine whether $\bigcap_{p\in P}(S^* - p) = \emptyset$, where $S^* := \bigcup_{B\in\mathcal{Q}} B$. For correctness, we just observe that $\bigcap_{p\in P}(S^* - p)$ is identical to $\bigcap_{S\in\mathcal{S}} S$ inside $[0,1]^d$. This is because for each $B \in \mathcal{B}_i$, $(B + u_i) - (u_{i'} - t)$ may intersect $[0,1]^d$ only if $i = i'$ (since u_i and $u_{i'}$ are far apart if $i \neq i'$), assuming $t \in [0,1/2]^d$. ◀

We now prove hardness of Problem 4 by reduction from the clique problem for graphs. We first warm up with two simpler reductions yielding weaker lower bounds, before presenting the final reduction in Lemma 6. (Readers who do not need intuition building may go straight to Lemma 6's proof.)

First attempt. First observe that the box coverage problem (deciding whether n boxes in $\mathbb{R}^d$ cover the entire space, i.e., deciding whether the intersection of the complements of n boxes is empty) easily reduces to Problem 4 with n orthants and n objects in $\mathbb{R}^d$, since the complement of a box is the union of $O(1)$ orthants. By Lemma 3, we immediately obtain an $\Omega(n^{d/2-\delta})$ conditional lower bound for Problem 4, since the box coverage problem has an $\Omega(n^{d/2-\delta})$ lower bound under the Combinatorial k-Clique Hypothesis [13].

In the following lemma, we directly modify the (very simple) known reduction from clique to the box coverage problem [13], to show the same lower bound even when the number of orthants m is $O(1)$:

▶ **Lemma 4.** *Detecting a d-clique in a graph with n_0 vertices reduces to Problem 4 with $m = O(1)$ orthants and $n = O(n_0^2)$ objects in $\mathbb{R}^d$.*

Proof. Let $G = (V, E)$ be the given graph, with $V = [n_0] = \{0, \ldots, n_0 - 1\}$. We will construct a set $\mathcal{S}$ of objects whose intersection is

$$\{(x_1, \ldots, x_d) \in [0, n_0)^d : \ \{\lfloor x_1 \rfloor, \ldots, \lfloor x_d \rfloor\} \text{ is a } d\text{-clique of } G\}. \tag{3}$$

It would then follow that the intersection is nonempty iff a d-clique exists.

The construction is very simple: for each $\alpha, \beta \in \{1, \ldots, d\}$ with $\alpha \neq \beta$ and for each $u, v \in [n_0]$ with $uv \notin E$, add the complement of the box

$$B_{\alpha,\beta,u,v} := \{(x_1, \ldots, x_d) : \ \lfloor x_\alpha \rfloor = u, \ \lfloor x_\beta \rfloor = v\}$$

to $\mathcal{S}$. (Note that if $u = v$, we consider $uv \notin E$.) These boxes are unit squares when projected to the α-th and β-th axes, and are thus translates of $O(1)$ fixed boxes, and the complement of each such fixed box can obviously be expressed as a union of $O(1)$ orthants and can be added to $\mathcal{Z}$. Lastly, add the $O(1)$ halfspaces bounding $[0, n_0)^d$ to $\mathcal{S}$. Then $\mathcal{S}$ has a total of $O(n_0^2)$ translates and clearly satisfies the desired property (3). ◀

In combination with Lemma 3, the above lemma indeed implies an $\Omega(n^{d/2-\delta})$ conditional lower bound for Problem 2: if Problem 2 for $m = O(1)$ has an $O(n^{d/2-\delta})$-time combinatorial algorithm, then the d-clique detection problem for a graph with n_0 vertices has a combinatorial algorithm with running time $O((n_0^2)^{d/2-\delta}) = O(n_0^{d-2\delta})$, contradicting the Combinatorial k-Clique Hypothesis.

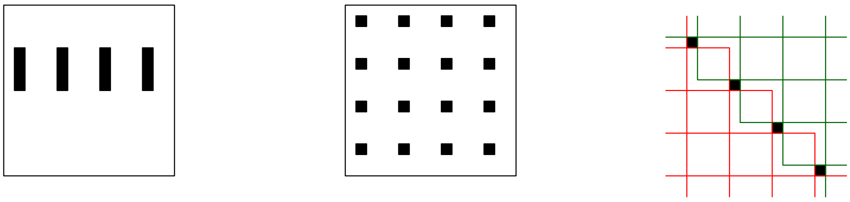

Figure 2 (left) A region $Y_{\alpha,\beta,a,b,u,v}$ for $(a,b)=(0,1)$. (middle) A region $Y_{\alpha,\beta,a,b,u,v}$ for $(a,b)=(0,0)$. (right) A diagonal $D_{\alpha,\beta}$ (whose complement can be expressed as a union of the red and the green orthants).

Second attempt. We now improve the lower bound by reducing from clique of a large size $2d$. We use the following key idea: encode a pair of vertices in a single coordinate value.

▶ **Lemma 5.** *Detecting a $(2d)$-clique in a graph G with n_0 vertices reduces to Problem 4 with $m=O(n_0)$ orthants and $n=O(n_0^3)$ objects in $\mathbb{R}^d$.*

Proof. Let $G=(V,E)$ be the given graph, with $V=[n_0]$. For any $x\in[0,n_0^2)$, let $\phi_0(x)=\lfloor x\rfloor \bmod n_0$ and $\phi_1(x)=\lfloor x/n_0\rfloor$. We will construct a set $\mathcal{S}$ of objects whose intersection is

$$\{(x_1,\ldots,x_d)\in[0,n_0^2)^d:\ \{\phi_0(x_1),\phi_1(x_1),\ldots,\phi_0(x_d),\phi_1(x_d)\}\text{ is a }(2d)\text{-clique in }G\}. \quad (4)$$

It would then follow that the intersection is nonempty iff a $(2d)$-clique exists.

For each $\alpha,\beta\in\{1,\ldots,d\}$ and $a,b\in\{0,1\}$ with $(\alpha,a)\neq(\beta,b)$, and for each $u,v\in[n_0]$ with $uv\notin E$, define the region

$$Y_{\alpha,\beta,a,b,u,v}\ :=\ \{(x_1,\ldots,x_d)\in[0,n_0^2)^d:\ \phi_a(x_\alpha)=u,\ \phi_b(x_\beta)=v\}.$$

If $\alpha=\beta$, then $Y_{\alpha,\beta,a,b,u,v}$ is just a unit interval when projected to the α-th axis, and is thus a translate of one of $O(1)$ fixed boxes, and the complement of each such fixed box can be expressed as a union of $O(1)$ orthants and can be added to $\mathcal{Z}$. From now on, assume $\alpha\neq\beta$.

If $(a,b)=(1,1)$, then $Y_{\alpha,\beta,a,b,u,v}$ is just an $n_0\times n_0$ square when projected to the α-th and β-th axes, and is thus a translate of one of $O(1)$ fixed boxes, and the complement of each such fixed box can be expressed as a union of $O(1)$ orthants and can be added to $\mathcal{Z}$.

If $(a,b)=(0,1)$ (or $(a,b)=(1,0)$), then $Y_{\alpha,\beta,a,b,u,v}$ is a union of n_0 rectangles of dimension $1\times n_0$ (or $n_0\times 1$) when projected to the α-th and β-th axes (see Figure 2(left)), and is thus a union of n_0 translates of one of $O(1)$ fixed boxes, and the complement of each fixed box can be expressed as a union of $O(1)$ orthants and can be added to $\mathcal{Z}$.

If $(a,b)=(0,0)$, then $Y_{\alpha,\beta,a,b,u,v}$ forms a $n_0\times n_0$ grid pattern when projected to the α-th and β-th axes (see Figure 2(middle)). Although the complement of this region can't be expressed as a union of orthants, we can decompose the grid into subregions that can. The most obvious approach is to decompose into rows or columns, but this still doesn't work. Instead, we will decompose into "diagonals". More precisely, define

$$\begin{aligned} D_{\alpha,\beta}\ :=\ \{(x_1,\ldots,x_d):\quad & \lfloor x_\alpha\rfloor \bmod n_0=0,\ \lfloor x_\beta\rfloor \bmod n_0=0,\\ & \lfloor x_\alpha/n_0\rfloor+\lfloor x_\beta/n_0\rfloor=n_0,\ x_\alpha,x_\beta\geq 0\}.\end{aligned}$$

Since $D_{\alpha,\beta,a,b}$ can be viewed as the region sandwiched between two staircases when projected to 2D, its complement $\overline{D_{\alpha,\beta,a,b}}$ can be expressed as a union of $O(n_0)$ orthants (see Figure 2(right)). Add the shape $\overline{D_{\alpha,\beta,a,b}}$ to $\mathcal{Z}$. The region $Y_{\alpha,\beta,a,b,u,v}$ can be expressed as a union of $O(n_0)$ translates of $D_{\alpha,\beta}$ when clipped to $[0,n_0^2)^d$.

In any case, $\overline{Y_{\alpha,\beta,a,b,u,v}} \cap [0,n_0^2)^d$ can be expressed as the intersection of $[0,n_0^2)^d$ with $O(n_0)$ translates of shapes from $\mathcal{Z}$. Add these translates to $\mathcal{S}$, for each $\alpha,\beta \in \{1,\ldots,d\}$ and $a,b \in [g]$ with $(\alpha,a) \neq (\beta,b)$, and for each $u,v \in [n_0]$ with $uv \notin E$. Lastly, add the $O(1)$ halfspaces bounding $[0,n_0^2)^d$ to $\mathcal{S}$. Then $\mathcal{S}$ has a total of $O(n_0^3)$ translates and satisfies the desired property (4). ◀

The above lemma implies a larger $\Omega(n^{2d/3-\delta})$ conditional lower bound: if Problem 2 for $m = n^{1/3}$ has an $O(n^{2d/3-\delta})$-time combinatorial algorithm, then the $(2d)$-clique detection problem for a graph with n_0 vertices has a combinatorial algorithm with running time $O((n_0^3)^{2d/3-\delta}) = O(n_0^{2d-3\delta})$, contradicting the Combinatorial k-Clique Hypothesis.

Final reduction. We obtain our final lower bound by generalizing the idea further. We reduce from clique of still larger size and now encode g-tuples of vertices instead of pairs (incidentally, the idea of encoding tuples has also appeared recently in Künnemann's conditional lower bound proofs for Klee's measure problem [26]):

▶ **Lemma 6.** *Let g be any integer constant. Detecting a (dg)-clique in a graph G with n_0 vertices reduces to Problem 4 with $m = O(n_0^{g-1})$ orthants and $n = O(n_0^{g+1})$ objects in $\mathbb{R}^d$.*

More generally, for any given $m \le n_0^{g-1}$, detecting a (dg)-clique in a graph G with n_0 vertices reduces to Problem 4 with m orthants and $n = O(n_0^{2g}/m)$ objects in $\mathbb{R}^d$.

Proof. Let $G = (V,E)$ be the given graph, with $V = [n_0]$. For any $x \in [0,n_0^g)$ and $a \in [g]$, let $\phi_a(x)$ be the $(a+1)$-th least significant digit of $\lfloor x \rfloor$ in base n. We will construct a set $\mathcal{S}$ of objects whose intersection is

$$\begin{aligned}\{(x_1,\ldots,x_d) \in [0,n_0^g)^d : \quad & \{\phi_0(x_1),\ldots,\phi_{g-1}(x_1),\ldots,\phi_0(x_d),\ldots,\phi_{g-1}(x_d)\} \\ & \text{is a } (dg)\text{-clique in } G\}. \end{aligned} \tag{5}$$

It would then follow that the intersection is nonempty iff a (dg)-clique exists.

For each $\alpha,\beta \in \{1,\ldots,d\}$ and $a,b \in [g]$ with $(\alpha,a) \neq (\beta,b)$, and for each $u,v \in [n_0]$ with $uv \notin E$, define the region

$$\begin{aligned} Y_{\alpha,\beta,a,b,u,v} \quad &:= \quad \{(x_1,\ldots,x_d) \in [0,n_0^d)^d : \ \phi_a(x_\alpha) = u, \ \phi_b(x_\beta) = v\} \\ &= \quad \{(x_1,\ldots,x_d) : \ x_\alpha \in [in_0^{a+1} + un_0^a, in_0^{a+1} + (u+1)n_0^a), \\ &\qquad\qquad\qquad\qquad x_\beta \in [jn_0^{b+1} + vn_0^b, jn^{b+1} + (v+1)n_0^b) \\ &\qquad\qquad\qquad\qquad \text{for some } i \in [n_0^{g-a-1}],\ j \in [n_0^{g-b-1}]\}. \end{aligned}$$

If $\alpha = \beta$, then $Y_{\alpha,\beta,a,b,u,v}$ is a union of at most $O(n_0^{g-1})$ unit intervals when projected to the α-th axis, and is thus a union of $O(n_0^{g-1})$ translates of $O(1)$ fixed boxes, and the complement of each fixed box can be expressed as a union of $O(1)$ orthants and can be added to $\mathcal{Z}$.

If $\alpha \neq \beta$, then $Y_{\alpha,\beta,a,b,u,v}$ forms a $O(n_0^{g-1}) \times O(n_0^{g-1})$ grid pattern when projected to the α-th and β-th axes. Define the "diagonal"

$$\begin{aligned} D_{\alpha,\beta,a,b} \ = \ \{(x_1,\ldots,x_d) : \quad & x_\alpha \in [in_0^{a+1}, in_0^{a+1} + n_0^a),\ x_\beta \in [jn_0^{b+1}, jn_0^{b+1} + n_0^b) \\ & \text{for some } i,j \in [m] \text{ with } i+j = m\}. \end{aligned}$$

Since $D_{\alpha,\beta,a,b}$ can be viewed as the region sandwiched between two staircases when projected to 2D, its complement $\overline{D_{\alpha,\beta,a,b}}$ can be expressed as a union of $O(m)$ orthants. Add the shape $\overline{D_{\alpha,\beta,a,b}}$ to $\mathcal{Z}$. The region $Y_{\alpha,\beta,a,b,u,v}$ can be expressed as a union of $O(n_0^{2(g-1)}/m)$ translates of $D_{\alpha,\beta,a,b}$ when clipped to $[0,n_0^g)^d$.

In any case, $\overline{Y_{\alpha,\beta,a,b,u,v}} \cap [0, n_0^g)^d$ can be expressed as an intersection of $[0, n_0^g)^d$ with $O(n_0^{2(g-1)}/m)$ translates of shapes from $\mathcal{Z}$. Add all these translates to $\mathcal{S}$, for each $\alpha, \beta \in \{1, \ldots, d\}$ and $a, b \in [g]$ with $(\alpha, a) \neq (\beta, b)$, and for each $u, v \in [n_0]$ with $uv \notin E$. Lastly, add the $O(1)$ halfspaces bounding $[0, n_0^g)^d$ to $\mathcal{S}$. Then $\mathcal{S}$ has a total of $O(n_0^2 \cdot n_0^{2(g-1)}/m) = O(n_0^{2g}/m)$ translates and satisfies the desired property (5). ◀

The above lemma implies an $\Omega(n^{gd/(g+1)-\delta})$ conditional lower bound for any integer constant g: if Problem 2 for $m = n^{(g-1)/(g+1)}$ has an $O(n^{gd/(g+1)-\delta})$-time combinatorial algorithm, then the (dg)-clique detection problem for a graph with n_0 vertices has a combinatorial algorithm with running time $O((n_0^{g+1})^{gd/(g+1)-\delta}) = O(n_0^{dg-(g+1)\delta})$, contradicting the Combinatorial k-Clique Hypothesis. The exponent $gd/(g+1) - \delta$ exceeds $d - 2\delta$, by picking a sufficiently large $g \geq d/\delta$.

More generally, for any constant $\gamma \leq (g-1)/(g+1)$, if Problem 2 for $m = n^\gamma$ has an $O((mn)^{d/2-\delta})$-time combinatorial algorithm, then the (dg)-clique detection problem for a graph with n_0 vertices has a combinatorial algorithm with running time $O(((n_0^{2g}/m)\cdot m)^{d/2-\delta}) = O(n_0^{dg-2g\delta})$, contradicting the Combinatorial k-Clique Hypothesis.

▶ **Theorem 7.** *Under the Combinatorial k-Clique Hypothesis, Problem 1 or 2 for n points and n unit hypercubes in $\mathbb{R}^d$ does not have an $O(n^{d-\delta})$-time combinatorial algorithm for any constant $\delta > 0$.*

More generally, under the same hypothesis, for any fixed constant $\gamma \leq 1$, Problem 1 or 2 for $m = n^\gamma$ points and n unit hypercubes in $\mathbb{R}^d$ does not have an $O((mn)^{d/2-\delta})$-time combinatorial algorithm for any constant $\delta > 0$.

4 Final Remarks

To summarize, we have studied the L_∞ translational Hausdorff distance problem for point sets, a fundamental problem with a long history in computational geometry. We have obtained a substantially improved upper bound for this problem, and the first conditional lower bound in dimension 3 and higher, which nearly match the upper bound. Our technique for the upper bound is interesting, in that it implies a natural colored generalization of Klee's measure problem can be solved in roughly the same time bound as the original Klee's problem. Our lower bound proof is interesting, in that it adds to a growing body of recent work on fine-grained complexity in computational geometry, and more specifically illustrates the power of the Combinatorial Clique Hypothesis.

Our near-$O((mn)^{d/2})$ upper bound also applies to the variant of the problem for *undirected* Hausdorff distance, since the undirected version of Problem 1 can also be reduced to Problem 3(a) with $N = O(mn)$. However, more effort might be needed to adapt our lower bounds to the undirected problem (although we have not tried seriously).

For noncombinatorial algorithms, our reduction implies a lower bound of $\Omega((mn)^{d\omega/6-\delta})$, under the standard hypothesis that the k-clique problem for graphs with n_0 vertices requires $\Omega(n_0^{d\omega/3-\delta'})$ time, where $\omega \in [2, 2.373)$ denotes the matrix multiplication exponent. Proving better conditional lower bounds for noncombinatorial algorithms remains open. This might require further new techniques, as we currently do not have tight conditional lower bounds for the original Klee's measure problem for noncombinatorial algorithms for $d \geq 4$ [26].

As mentioned, Bringmann and Nusser [10] proved a near-mn lower bound for $d = 2$ under the OV Hypothesis; their result is in some sense more robust (it holds for noncombinatorial algorithms) and applies also to the L_2 case (and L_p for any $1 \leq p \leq \infty$). However, the problem for L_2 probably has higher complexity than for L_∞: the best upper bounds are near

n^3 for $d = 2$ and near n^5 for $d = 3$ [23], and near $n^{\lceil 3d/2\rceil+1}$ for $d \geq 4$ [15], in the $m = n$ case. (See Bringmann and Nusser's paper for a 3SUM-based lower bound for the L_2 problem for $d = 2$ in the "unbalanced" case when m is constant.)

References

1 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the current clique algorithms are optimal, so is Valiant's parser. *SIAM J. Comput.*, 47(6):2527–2555, 2018. Preliminary version in FOCS 2015. doi:10.1137/16M1061771.

2 Pankaj K. Agarwal. An improved algorithm for computing the volume of the union of cubes. In *Proc. 26th ACM Symposium on Computational Geometry* (SoCG), pages 230–239, 2010. doi:10.1145/1810959.1811000.

3 Pankaj K. Agarwal, Sariel Har-Peled, Micha Sharir, and Yusu Wang. Hausdorff distance under translation for points and balls. *ACM Trans. Algorithms*, 6(4):71:1–71:26, 2010. Preliminary version in SoCG 2003. doi:10.1145/1824777.1824791.

4 Oswin Aichholzer, Helmut Alt, and Günter Rote. Matching shapes with a reference point. *Int. J. Comput. Geom. Appl.*, 7(4):349–363, 1997. Preliminary version in SoCG 1994. doi:10.1142/S0218195997000211.

5 Boris Aronov and Jean Cardinal. Geometric pattern matching reduces to k-SUM. In *Proc. 31st International Symposium on Algorithms and Computation* (ISAAC), volume 181, pages 32:1–32:9, 2020. doi:10.4230/LIPIcs.ISAAC.2020.32.

6 Gill Barequet and Sariel Har-Peled. Polygon containment and translational min-Hausdorff-distance between segment sets are 3SUM-hard. *Int. J. Comput. Geom. Appl.*, 11(4):465–474, 2001. doi:10.1142/S0218195901000596.

7 Jean-Daniel Boissonnat, Micha Sharir, Boaz Tagansky, and Mariette Yvinec. Voronoi diagrams in higher dimensions under certain polyhedral distance functions. In *Proc. 11th Annual Symposium on Computational Geometry* (SoCG), pages 79–88, 1995. doi:10.1145/220279.220288.

8 Karl Bringmann. An improved algorithm for Klee's measure problem on fat boxes. *Comput. Geom.*, 45(5-6):225–233, 2012. Preliminary version in SoCG 2010. doi:10.1016/j.comgeo.2011.12.001.

9 Karl Bringmann, Allan Grønlund, and Kasper Green Larsen. A dichotomy for regular expression membership testing. In *Proc. 58th IEEE Annual Symposium on Foundations of Computer Science* (FOCS), pages 307–318, 2017. doi:10.1109/FOCS.2017.36.

10 Karl Bringmann and André Nusser. Translating Hausdorff is hard: Fine-grained lower bounds for Hausdorff distance under translation. In *Proc. 37th International Symposium on Computational Geometry* (SoCG), pages 18:1–18:17, 2021. doi:10.4230/LIPIcs.SoCG.2021.18.

11 Sergio Cabello, Panos Giannopoulos, and Christian Knauer. On the parameterized complexity of d-dimensional point set pattern matching. *Inf. Process. Lett.*, 105(2):73–77, 2008. doi:10.1016/j.ipl.2007.08.003.

12 Timothy M. Chan. Geometric applications of a randomized optimization technique. *Discret. Comput. Geom.*, 22(4):547–567, 1999. Preliminary version in SoCG 1998. doi:10.1007/PL00009478.

13 Timothy M. Chan. A (slightly) faster algorithm for Klee's measure problem. *Comput. Geom.*, 43(3):243–250, 2010. Preliminary version in SoCG 2008. doi:10.1016/j.comgeo.2009.01.007.

14 Timothy M. Chan. Klee's measure problem made easy. In *Proc. 54th IEEE Symposium on Foundations of Computer Science* (FOCS), pages 410–419, 2013. doi:10.1109/FOCS.2013.51.

15 L. Paul Chew, Dorit Dor, Alon Efrat, and Klara Kedem. Geometric pattern matching in d-dimensional space. *Discret. Comput. Geom.*, 21(2):257–274, 1999. Preliminary version in ESA 1995. doi:10.1007/PL00009420.

16 L. Paul Chew, Michael T. Goodrich, Daniel P. Huttenlocher, Klara Kedem, Jon M. Kleinberg, and Dina Kravets. Geometric pattern matching under Euclidean motion. *Comput. Geom.*, 7:113–124, 1997. doi:10.1016/0925-7721(95)00047-X.

17 L. Paul Chew and Klara Kedem. Getting around a lower bound for the minimum Hausdorff distance. *Comput. Geom.*, 10(3):197–202, 1998. Preliminary version in SWAT 1992. doi:10.1016/S0925-7721(97)00032-1.

18 Minkyoung Cho and David M. Mount. Improved approximation bounds for planar point pattern matching. *Algorithmica*, 50(2):175–207, 2008. doi:10.1007/s00453-007-9059-9.

19 Alon Efrat, Piotr Indyk, and Suresh Venkatasubramanian. Pattern matching for sets of segments. *Algorithmica*, 40(3):147–160, 2004. Preliminary version in SODA 2001. doi:10.1007/s00453-004-1089-y.

20 Greg N. Frederickson and Donald B. Johnson. Generalized selection and ranking: Sorted matrices. *SIAM J. Comput.*, 13(1):14–30, 1984. doi:10.1137/0213002.

21 Michael T. Goodrich, Joseph S. B. Mitchell, and Mark W. Orletsky. Approximate geometric pattern matching under rigid motions. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(4):371–379, 1999. Preliminary version in SoCG 1994. doi:10.1109/34.761267.

22 Daniel P. Huttenlocher and Klara Kedem. Computing the minimum Hausdorff distance for point sets under translation. In *Proc. 6th ACM Symposium on Computational Geometry* (SoCG), pages 340–349, 1990. doi:10.1145/98524.98599.

23 Daniel P. Huttenlocher, Klara Kedem, and Micha Sharir. The upper envelope of Voronoi surfaces and its applications. *Discret. Comput. Geom.*, 9:267–291, 1993. Preliminary version in SoCG 1991. doi:10.1007/BF02189323.

24 Piotr Indyk, Rajeev Motwani, and Suresh Venkatasubramanian. Geometric matching under noise: Combinatorial bounds and algorithms. In *Proc. 10th ACM-SIAM Symposium on Discrete Algorithms* (SODA), pages 457–465, 1999. URL: http://dl.acm.org/citation.cfm?id=314500.314601.

25 Ce Jin and Yinzhan Xu. Tight dynamic problem lower bounds from generalized BMM and OMv. In *Proc. 54th Annual ACM Symposium on Theory of Computing* (STOC), pages 1515–1528, 2022. doi:10.1145/3519935.3520036.

26 Marvin Künnemann. A tight (non-combinatorial) conditional lower bound for Klee's Measure Problem in 3D. In *Proc. 63rd IEEE Symposium on Foundations of Computer Science* (FOCS), pages 555–566, 2022.

27 André Nusser. *Fine-Grained Complexity and Algorithm Engineering of Geometric Similarity Measures.* PhD thesis, Saarland University, Saarbrücken, Germany, 2021. URL: https://publikationen.sulb.uni-saarland.de/handle/20.500.11880/33904.

28 Mark H. Overmars and Chee-Keng Yap. New upper bounds in Klee's measure problem. *SIAM J. Comput.*, 20(6):1034–1045, 1991. Preliminary version in FOCS 1988. doi:10.1137/0220065.

29 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the ICM*, volume 3, pages 3431–3472. World Scientific, 2018. URL: https://people.csail.mit.edu/virgi/eccentri.pdf.

30 Hakan Yildiz and Subhash Suri. Computing Klee's measure of grounded boxes. *Algorithmica*, 71(2):307–329, 2015. Preliminary version in SoCG 2012. doi:10.1007/s00453-013-9797-9.