Translating Hausdorff Is Hard: Fine-Grained Lower Bounds for Hausdorff Distance Under Translation

Karl Bringmann
Universität des Saarlandes, Saarbrücken, Germany
Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

André Nusser
Saarbrücken Graduate School of Computer Science, Universität des Saarlandes, Germany
Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

Abstract

Computing the similarity of two point sets is a ubiquitous task in medical imaging, geometric shape comparison, trajectory analysis, and many more settings. Arguably the most basic distance measure for this task is the Hausdorff distance, which assigns to each point from one set the closest point in the other set and then evaluates the maximum distance of any assigned pair. A drawback is that this distance measure is not translational invariant, that is, comparing two objects just according to their shape while disregarding their position in space is impossible.

Fortunately, there is a canonical translational invariant version, the Hausdorff distance under translation, which minimizes the Hausdorff distance over all translations of one of the point sets. For point sets of size n and m, the Hausdorff distance under translation can be computed in time $\tilde{\mathcal{O}}(nm)$ for the L_1 and L_∞ norm [Chew, Kedem SWAT'92] and $\tilde{\mathcal{O}}(nm(n+m))$ for the L_2 norm [Huttenlocher, Kedem, Sharir DCG'93].

As these bounds have not been improved for over 25 years, in this paper we approach the Hausdorff distance under translation from the perspective of fine-grained complexity theory. We show (i) a matching lower bound of $(nm)^{1-o(1)}$ for L_1 and L_∞ assuming the Orthogonal Vectors Hypothesis and (ii) a matching lower bound of $n^{2-o(1)}$ for L_2 in the imbalanced case of $m = \mathcal{O}(1)$ assuming the 3SUM Hypothesis.

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness

Keywords and phrases Hausdorff Distance Under Translation, Fine-Grained Complexity Theory, Lower Bounds

Digital Object Identifier 10.4230/LIPIcs.SoCG.2021.18

Funding *Karl Bringmann*: This work is part of the project TIPEA that has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 850979).

1 Introduction

As data sets become larger and larger, the requirement for faster algorithms to handle such amounts of data becomes increasingly necessary. One very common type of data that is created during measurements is point sets in the plane, for example when recording GPS trajectories or describing shapes of objects, in medical image analysis, and in various data science applications.

A fundamental algorithmic tool for analyzing point sets is to compute the similarity of two given sets of points. There are several different measures of similarity in this setting, for example Hausdorff distance [21], geometric bottleneck matching [18], Fréchet distance [3], and Dynamic Time Warping [25]. Among these measures, the Hausdorff distance is arguably

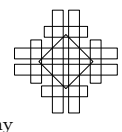

the most basic and intuitive: It assigns to each point from one set the closest point in the other set and then evaluates the maximum distance of all assigned pairs of points.[1] For a discussion of the other previously mentioned distance measures, see Section 1.1.

While these similarity measures are of great practical relevance, for some applications it is a drawback that they are not translational invariant, i.e., when translating a point set the distance can – and in most cases will – change. This is unfavorable in applications that ask for comparing the shape of two objects, meaning that the absolute position of an object is irrelevant. Examples of this task arise for example in 2D object shape similarity, medical image analysis [19], classification of handwritten characters [10], movement patterns of animals [12], and sports analysis [17].

Fortunately, any point set similarity measure has a canonical translational invariant version, by minimizing the similarity measure over all translations of the two given point sets. For the Hausdorff distance this variant is known as the *Hausdorff distance under translation*, see Section 2 for a formal definition. Given two point sets in the plane of size n and m, the Hausdorff distance under translation can be computed in time $\mathcal{O}(nm \log^2 nm)$ for the L_1 and L_∞ norm [16], and in time $\mathcal{O}(nm(n+m) \log nm)$ for the L_2 norm [22]. We are not aware of any lower bounds for this problem, not even conditional on a plausible hypothesis. The only results in this direction are $\Omega(n^3)$ lower bounds on the arrangement size [16] and on the number of connected components of the feasible translations [28] (for the decision problem on points in the plane with $n = m$). However, these bounds also hold for L_1 and L_∞, where they are "broken" by the $\mathcal{O}(nm \log^2 nm)$-time algorithm [16], so apparently these bounds are irrelevant for the running time complexity.

In this paper, we approach the Hausdorff distance under translation from the viewpoint of fine-grained complexity theory [29]. For two problem settings, we show that the known algorithms are optimal up to lower order factors assuming standard hypotheses:

1. We show an $(nm)^{1-o(1)}$ lower bound for L_1 and L_∞, matching the $\mathcal{O}(nm \log^2 nm)$-time algorithm from [16] up to lower order factors.

   This result holds conditional on the Orthogonal Vectors Hypothesis, which states that finding two orthogonal vectors among two given sets of n binary vectors in d dimensions cannot be done in time $\mathcal{O}(n^{2-\varepsilon}\text{poly}(d))$ for any $\varepsilon > 0$. It is well-known that the Orthogonal Vectors Hypothesis is implied by the Strong Exponential Time Hypothesis [30], and thus our lower bound also holds assuming the latter [23]. These two hypotheses are the most standard assumptions used in fine-grained complexity theory in the last decade [29].
2. We show an $n^{2-o(1)}$ lower bound for L_2 in the imbalanced case $m = \mathcal{O}(1)$, matching the $\mathcal{O}(nm(n+m) \log nm)$-time algorithm from [16] up to lower order factors. Previously, an $n^{2-o(1)}$ lower bound was only known for the more general problem of computing the Hausdorff distance under translation of sets of *segments* in the case that both sets have size n (a problem for which the best known algorithm runs in time[2] $\tilde{\mathcal{O}}(n^4)$) [6].

   Our result holds conditional on the 3SUM Hypothesis, which states that deciding whether among n given integers there are three that sum up to 0 requires time $n^{2-o(1)}$. This hypothesis was introduced by Gajentaan and Overmars [20], is a standard assumption in computational geometry [24], and has also found a wealth of applications beyond geometry (see, e.g., [26, 4, 2, 1]).

[1] There is a directed and an undirected variant of the Hausdorff distance, see Section 2. In this introduction, we do not differentiate between these two, since all our statements hold for both variants.

[2] By $\tilde{\mathcal{O}}$-notation we ignore logarithmic factors in n and m.

Our lower bounds close gaps that have not seen any progress over 25 years. Furthermore, note that our second lower bound shows a separation between the L_2 norm and the L_1 and L_∞ norms, as in the imbalanced case $m = \mathcal{O}(1)$ the former admits a $\tilde{\mathcal{O}}(n)$-time algorithm [16] while the latter requires time $n^{2-o(1)}$ assuming the 3SUM Hypothesis. We leave it as an open problem whether for L_2 the balanced case $n = m$ requires time $n^{3-o(1)}$.

1.1 Related work

Our work continues a line of research on fine-grained lower bounds in computational geometry, which had early success with the 3SUM Hypothesis [20] and recently got a new impulse with the Orthogonal Vectors Hypothesis (or Strong Exponential Time Hypothesis) and resulting lower bounds for the Fréchet distance [7], see also [13, 11]. Continuing this line of research is getting increasingly difficult, although there are still many classic problems from computational geometry without matching lower bounds. In this paper we obtain such bounds for two settings of the classic Hausdorff distance under translation.

Besides Hausdorff distance, there are several other distance measures on point sets, including geometric bottleneck matching [18], Fréchet distance [3], and Dynamic Time Warping [25]. The geometric bottleneck matching minimizes the maximal distance in a perfect matching between the two given point sets. Fréchet distance and Dynamic Time Warping additionally take the order of the input points into account. They both consider the same class of *traversals* of the input points, and the Fréchet distance minimizes the *maximal* distance that occurs during the traversal, while Dynamic Time Warping minimizes the *sum* of distances.

Let us discuss the canonical translational invariant versions of these distance measures. For geometric bottleneck matching under translation, Efrat et al. designed an $\tilde{\mathcal{O}}(n^5)$ algorithm [18]. The discrete Fréchet distance under translation has an $\tilde{\mathcal{O}}(n^{4.66\cdots})$-time algorithm and a conditional lower bound of $n^{4-o(1)}$ [9], see also [10] for algorithm engineering work on this topic. While Dynamic Time Warping is a very popular measure (in particular for video and speech processing), no exact algorithm for its canonical translational invariant version is known in L_2 since it contains the geometric median problem as a special case [5].

Further work on the Hausdorff distance under translation includes an $\mathcal{O}((n+m)\log nm)$-time algorithm for point sets in one dimension [27]. For generalizations to dimensions $d > 2$ see [16, 15].

2 Preliminaries

In this paper, we consider finite point sets which lie in $\mathbb{R}^2$. For any $p \in \mathbb{R}^2$, we use p_x and p_y to refer to its first and second component, respectively. For a point set $A \subset \mathbb{R}^2$ and a translation $\tau \in \mathbb{R}^2$, we define $A + \tau \coloneqq \{a + \tau \mid a \in A\}$. To denote index sets, we often use $[n] \coloneqq \{1, \dots, n\}$. Given a point $x \in \mathbb{R}^2$, its p-norm is defined as

$$\|x\|_p \coloneqq \left(\sum_{i \in [d]} |x_i|^p \right)^{\frac{1}{p}}.$$

We now introduce several distance measures, which are all versions of the famous Hausdorff distance. First, let us define the most basic version. Let $A, B \subset \mathbb{R}^2$ be two point sets. The *directed Hausdorff distance* is defined as

$$\delta_{\vec{H}}(A, B) \coloneqq \max_{a \in A} \min_{b \in B} \|a - b\|_p.$$

Note that, intuitively, the directed Hausdorff distance measures the distance from A to B but not from B to A, and it is not symmetric. A symmetric variant of the Hausdorff distance, the *undirected Hausdorff distance*, is defined as

$$\delta_H(A,B) \coloneqq \max\{\delta_{\vec{H}}(A,B), \delta_{\vec{H}}(B,A)\}.$$

Note that, by definition, $\delta_{\vec{H}}(A,B) \leq \delta_H(A,B)$. Both of the above distance measures can be modified to a version which is invariant under translation. The *directed Hausdorff distance under translation* is defined as

$$\delta^T_{\vec{H}}(A,B) \coloneqq \min_{\tau\in\mathbb{R}^2} \delta_{\vec{H}}(A, B+\tau),$$

and the *undirected Hausdorff distance under translation* is defined as

$$\delta^T_H(A,B) \coloneqq \min_{\tau\in\mathbb{R}^2} \delta_H(A, B+\tau).$$

Again, it holds that $\delta^T_{\vec{H}}(A,B) \leq \delta^T_H(A,B)$. Naturally, for all of the above distance measures, the decision problem is defined such that we are given two point sets A, B and a threshold distance δ, and ask if the distance of A, B is at most δ.

For the Hausdorff distance (without translation) the undirected distance is at most as hard as the directed distance, because the undirected distance can be calculated using two calls to an algorithm computing the directed distance.[3] However, note that for the Hausdorff distance under translation, we cannot just compute the directed distance twice and then obtain the undirected distance as we have to take the maximum for the same translation.

3 OV based $(mn)^{1-o(1)}$ lower bound for L_1 and L_∞

We now present a conditional lower bound of $(mn)^{1-o(1)}$ for the Hausdorff distance under translation for L_1 and L_∞. For simplicity, we present the lower bound for the L_1 case. This construction is equivalent to the L_∞ case, via a rotation by $\frac{\pi}{4}$. Our lower bound is based on the hypothesized hardness of the Orthogonal Vectors problem.

▶ **Definition 1** (Orthogonal Vectors Problem (OV)). *Given two sets $X, Y \subset \{0,1\}^d$ with $|X| = m, |Y| = n$, decide whether there exist $x \in X$ and $y \in Y$ with $\langle x, y\rangle = 0$.*

A popular hypothesis from fine-grained complexity theory is as follows.

▶ **Definition 2** (Orthogonal Vectors Hypothesis (OVH)). *The Orthogonal Vectors problem cannot be solved in time $\mathcal{O}((nm)^{1-\epsilon}\text{poly}(d))$ for any $\epsilon > 0$.*

This hypothesis is typically stated and used for the balanced case $n = m$. However, it is known that the hypothesis for the balanced case is equivalent to the hypothesis for any unbalanced case $n = m^\alpha$ for any fixed constant $\alpha > 0$, see, e.g, [8, Lemma 5.1 in Arxiv version].

We now describe a reduction from Orthogonal Vectors to Hausdorff distance under translation. To this end, we are given two sets of d-dimensional binary vectors $X = \{x_1, \ldots, x_m\}$ and $Y = \{y_1, \ldots, y_n\}$ with $|X| = m$ and $|Y| = n$, and we construct an instance of the undirected Hausdorff distance under translation defined by point sets A and B and a

[3] Actually, the directed Hausdorff distance is also at most as hard as the undirected Hausdorff distance (thus, they are equally hard), as $\delta_{\vec{H}}(A,B) = \delta_H(A \cup B, B)$.

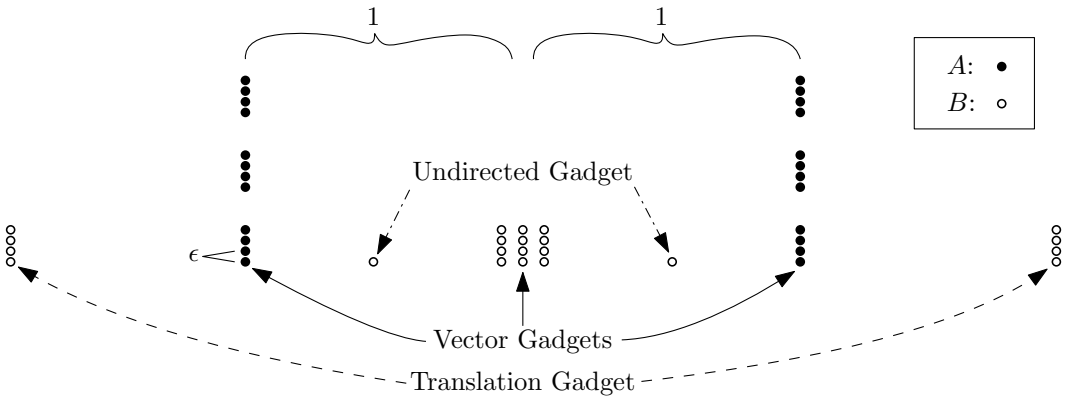

Figure 1 Sketch of the reduction from OV to the undirected Hausdorff distance under translation. The microtranslations in the order of ϵ^2 are not shown in this sketch.

decision distance $\delta = 1$. First, we describe the high-level structure of our reduction. The point set A consists only of Vector Gadgets, which encode the vectors of X using $2md$ points. The point set B consists of three types of gadgets:

- *Vector Gadgets:* They encode the vectors from Y, very similar to the Vector Gadgets of A.
- *Translation Gadget:* It restricts the possible translations of the point set B.
- *Undirected Gadget:* It makes our reduction work for the undirected Hausdorff distance under translation by ensuring that the maximum over the directed Hausdorff distances is always attained by $\delta_{\vec{H}}(B + \tau, A)$.

Intuitively, the two dimensions of the translation choose the vectors $x \in X$ and $y \in Y$ by aligning a Vector Gadget from A with a Vector Gadget from B in a certain way. An alignment of distance at most 1 is only possible if x and y are orthogonal. See Figure 1 for an overview of the reduction.

3.1 Gadgets

We now describe the gadgets in detail. Let $\epsilon > 0$ be a sufficiently small constant, e.g., think of $\epsilon = \frac{1}{20mnd}$. Recall that the distance for which we want to solve the decision problem is $\delta = 1$. Furthermore, we denote the ith component of a vector v by $v[i]$.

Vector Gadget

We define a general Vector Gadget, which we then use at several places by translating it. Given a vector $v \in \{0,1\}^d$, the Vector Gadget consists of the points $p_1, \ldots, p_d \in \mathbb{R}^2$:

$$p_i = \begin{cases} (\epsilon^2, i\epsilon), & \text{if } v[i] = 0 \\ (0, i\epsilon), & \text{if } v[i] = 1 \end{cases}$$

We denote the Vector Gadget created from vector v by $V(v)$. Additionally, we define a mirrored version of the gadget $V(v)$, defined as

$$\overline{V} := V(\bar{v}),$$

where $\bar{v}$ is the inversion of v, i.e., each bit is flipped.

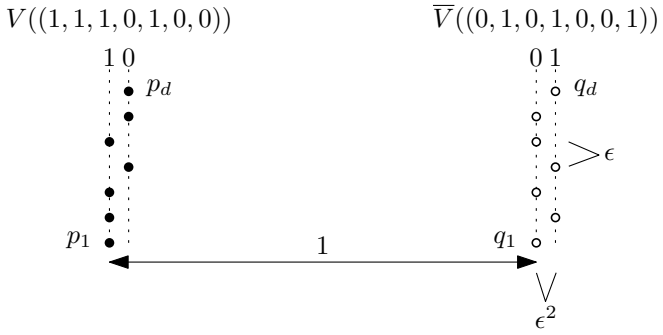

Figure 2 A depiction of the two types of Vector Gadgets and how they are placed to check for orthogonality.

▶ **Lemma 3.** *Given two vectors $v_1, v_2 \in \{0,1\}^d$ and corresponding Vector Gadgets $V_1 = V(v_1)$ and $V_2 = \overline{V}(v_2) + (1,0)$, $\delta_H(V_1, V_2) \leq 1$ if and only if $v_1 \cdot v_2 = 0$.*

Proof. Let the points of V_1 (resp. V_2) be denoted as $p_1, \ldots, p_d$ (resp. $q_1, \ldots, q_d$). First, note that $\|p_i - q_j\|_1 = 1 + |i-j|\epsilon + (v_1[i] + v_2[j] - 1)\epsilon^2 > 1$ for $i \neq j$. Thus, for the Hausdorff distance to be at most 1, we have to match p_i to q_i for all $i \in [d]$. This is possible if and only if $v_1[i] = 0$ or $v_2[i] = 0$, as p_i and q_i are only far for $v_1[i] = 1$ and $v_2[i] = 1$. ◀

See Figure 2 for an example. Note that if we swap both gadgets and invert both vectors (i.e., flip all their bits), the Hausdorff distance does not change and thus an analogous version of Lemma 3 holds in this case, as we are just performing a double inversion.

▶ **Lemma 4.** *Given two vectors $v_1, v_2 \in \{0,1\}^d$ and corresponding Vector Gadgets $V_1 = \overline{V}(v_1)$ and $V_2 = V(v_2)+(1,0)$, $\delta_H(V_1, V_2) \leq 1$ if and only if $\bar{v}_1 \cdot \bar{v}_2 = 0$, where $\bar{v}_1, \bar{v}_2$ are the inversions of v_1, v_2.*

For two Vector Gadgets $V_1 = V(v_1) + (x, y)$ and $V_2 = \overline{V}(v_2) + (x + D, y)$, we say that V_1 and V_2 are *vertically aligned*, or more precisely *vertically aligned in distance D*.

Translation Gadget

To ensure that B cannot be translated arbitrarily, we introduce a gadget to restrict the translations to the regime we require. The Translation Gadget T consists of two translated Vector Gadgets of the zero vector:

$$T \coloneqq (\overline{V}(1^d) - (2 + n\epsilon, 0)) \cup (\overline{V}(0^d) + (2 + 2\epsilon, 0))$$

We show that a simple property on the other set involved in the Hausdorff distance under translation instance already restricts the feasible translations significantly.

▶ **Lemma 5.** *Let $P \subset [-1 - \frac{1}{2}\epsilon, 1 + \frac{1}{2}\epsilon] \times \mathbb{R}$ be a point set. If $\delta_{\vec{H}}^T(T, P) \leq 1$, then $\tau_x^* \in [-(n + \frac{1}{2})\epsilon - \epsilon^2, -\frac{3}{2}\epsilon]$, where τ^* is any translation satisfying $\delta_{\vec{H}}(T, P + \tau^*) \leq 1$.*

Proof. We show the contrapositive. Therefore, assume the converse, i.e., that $\tau_x^* \notin [-(n + \frac{1}{2})\epsilon - \epsilon^2, -\frac{3}{2}\epsilon]$. If $\tau_x^* < -(n + \frac{1}{2})\epsilon - \epsilon^2$, then $-1 - \frac{1}{2}\epsilon - (-2 + n\epsilon + \epsilon^2 + \tau_x^*) > 1$ and thus the left part of T cannot contain any point of P in distance 1. If $\tau_x^* > -\frac{3}{2}\epsilon$, then $2 + 2\epsilon + \tau_x^* - (1 + \frac{1}{2}\epsilon) > 1$ and thus the right part of T cannot contain any point of P in distance 1. Thus, $\delta_{\vec{H}}^T(T, P) > 1$. ◀

Undirected Gadget

To ensure that each point in A can be matched to a point in B with distance at most 1, we add auxiliary points to B. The Undirected Gadget is defined by the point set

$$U := \{(-\frac{1}{2}, 0), (\frac{1}{2}, 0)\}.$$

▶ **Lemma 6.** *Given a set of points $P \subset [-1-\frac{1}{2}\epsilon, 1+\frac{1}{2}\epsilon]\times[-\frac{1}{8}, \frac{1}{8}]$, it holds that $\delta_{\vec{H}}(P, U+\tau) \leq 1$ for any $\tau \in [-(n+\frac{1}{2})\epsilon - \epsilon^2, (n+\frac{1}{2})\epsilon + \epsilon^2] \times [-\frac{1}{8}, \frac{1}{8}]$.*

Proof. By symmetry, we can restrict to proving that the distance of the point set

$$P' = P \cap [0, (n+\frac{1}{2})\epsilon + \epsilon^2] \times [-\frac{1}{8}, \frac{1}{8}]$$

to $(\frac{1}{2}, 0)+\tau$ is at most 1. For any $p' \in P'$, we have $|p'_x - (\frac{1}{2}+\tau_x)| \leq \frac{1}{2} + \mathcal{O}(n\epsilon)$ and $|p'_y - \tau_y| \leq \frac{1}{4}$. Thus, $\|p' - ((\frac{1}{2}, 0) + \tau)\|_1 = \frac{3}{4} + \mathcal{O}(n\epsilon)$, which is less than 1 for small enough ϵ. ◀

3.2 Reduction and correctness

We now describe the reduction and prove its correctness. We construct the point sets of our Hausdorff distance under translation instance as follows. The first set, i.e., set A, consists only of Vector Gadgets:

$$A := \left(\bigcup_{i\in[m]} V(x_i) + (-1 - \frac{1}{2}\epsilon, i \cdot 2d\epsilon)\right) \cup \left(\bigcup_{i\in[m]} V(1^d) + (1 + \frac{1}{2}\epsilon, i \cdot 2d\epsilon)\right)$$

The second set, i.e., set B, consists of Vector Gadgets, the Translation Gadget, and the Undirected Gadget:

$$B := \left(\bigcup_{j\in[n]} \overline{V}(y_j) + (j\epsilon, 0)\right) \cup T \cup U$$

See Figure 1 for a sketch of the above construction. To reference the vector gadgets as they are used in the reduction, we use the notation

$$V_r(x_i) := V(x_i) + (-1 - \frac{1}{2}\epsilon, i \cdot 2d\epsilon) \quad \text{and} \quad \overline{V}_r(y_j) := \overline{V}(y_j) + (j\epsilon, 0).$$

We can now prove correctness of our reduction. In the reduction, we return some canonical positive instance, if the 0^d vector is contained in any of the two OV sets. This allows us to drop all 1^d vectors from the input, as they cannot be orthogonal to any other vector. Thus, we can assume that all vectors in our input contain at least one 0-entry and at least one 1-entry.

▶ **Theorem 7.** *Computing the directed or undirected Hausdorff distance under translation in L_1 or L_∞ for two sets of size n and m cannot be solved in time $\mathcal{O}((mn)^{1-\gamma})$ for any $\gamma > 0$, unless the Orthogonal Vectors Hypothesis fails.*

Proof. Recall that we only have to consider the L_1 case. We first prove that there is a pair of orthogonal vectors $x \in X$ and $y \in Y$ if and only if $\delta_H^T(A, B) \leq 1$.

$\Rightarrow$: Assume that there exist $x_i \in X$, $y_j \in Y$ and $\langle x_i, y_j \rangle = 0$. Then consider the translation $\tau = (-(j+\frac{1}{2})\epsilon, i \cdot 2d\epsilon)$ which vertically aligns the Vector Gadgets $V_r(x_i)$ and $\overline{V}_r(y_j)+\tau$ in distance 1. As x_i and y_j are orthogonal, it follows from Lemma 3 that $\delta_{\vec{H}}(\overline{V}_r(y_j)+\tau, A) \leq 1$. It remains to show that all remaining points of $B + \tau$ have a point in distance at most 1. The Vector Gadgets in $B + \tau$ which correspond to $y_{j'}$ with $j' < j$ are strictly to the left of $\overline{V}_r(y_j) + \tau$ and are thus also in Hausdorff distance at most 1 from $V_r(x_i)$. If $j = n$, then we are done with the Vector Gadgets. Otherwise, consider the Vector Gadget $\overline{V}_r(y_{j+1}) + \tau$. We claim that each point of it is in distance at most 1 from $V(1^d) + (1 + \frac{1}{2}\epsilon, i \cdot 2d\epsilon)$. As the two gadgets are vertically aligned, we just have to check their horizontal distance, which is

$$1 + \frac{1}{2}\epsilon - ((j+1)\epsilon - (j + \frac{1}{2})\epsilon) = 1.$$

Thus, by Lemma 3, we have $\delta_{\vec{H}}(\overline{V}_r(y_{j+1}) + \tau, A) \leq 1$. Now, by the same argument as above, all gadgets $\overline{V}_r(y_{j'}) + \tau$ with $j' > j + 1$ are in directed Hausdorff distance at most 1 from A.
As the points of the Undirected Gadget $U + \tau$ are closer by a distance of almost $\frac{1}{2}$ to A than the Vector Gadgets in $B + \tau$, also $\delta_{\vec{H}}(U + \tau, A) \leq 1$ holds. Finally, we have to show that the Translation Gadget $T + \tau$ is in distance at most 1 from A. As the left part of T and $V_r(x_i)$ are aligned vertically, we only have to check the horizontal distance. The horizontal distance is

$$-1 - \frac{1}{2}\epsilon - (-2 + n\epsilon - (j + \frac{1}{2})\epsilon) = 1 - (n - j)\epsilon \leq 1$$

for any $j \in [n]$. Similarly, the distance of the right part of the Translation Gadget from the vertically aligned $V(1^d)$ in A is

$$2 + 2\epsilon - (j + \frac{1}{2})\epsilon - (1 + \frac{1}{2}\epsilon) = 1 - (j - 1)\epsilon \leq 1$$

for any $j \in [n]$. Thus, by Lemma 3 and Lemma 4, it holds that $\delta_{\vec{H}}(T + \tau, A) \leq 1$. As $\tau \in [-(n + \frac{1}{2})\epsilon - \epsilon^2, -\frac{3}{2}\epsilon] \times [-\frac{1}{8}, \frac{1}{8}]$, we know by Lemma 6 that $\delta_{\vec{H}}(A, B + \tau) \leq 1$ and thus also $\delta_H^T(A, B) \leq 1$.

$\Leftarrow$: Now, assume that $\delta_H^T(A, B) \leq 1$ and let τ be any translation for which $\delta_{\vec{H}}(B + \tau, A) \leq 1$. We used the directed Hausdorff distance in the previous statement on purpose, as we prove hardness for both versions. Lemma 5 implies that $\tau_x \in [-(n + \frac{1}{2})\epsilon - \epsilon^2, -\frac{3}{2}\epsilon]$.
Let $\overline{V}_r(y_j) + \tau, \overline{V}_r(y_{j+1}) + \tau$ be the Vector Gadgets such that $\overline{V}_r(y_j) + \tau$ has directed Hausdorff distance at most 1 to the left Vector Gadgets of A and $\overline{V}_r(y_{j+1})+\tau$ has directed Hausdorff distance at most 1 to the right Vector Gadgets of A. This is well-defined as the left Vector Gadgets of A and the right Vector Gadgets of A are in distance at least $2+\epsilon-\epsilon^2$ from each other, and thus no Vector Gadget of $B+\tau$ can be in distance at most 1 from both sides. Furthermore, as $\tau_x \leq -\frac{3}{2}\epsilon$, there has to be a Vector Gadget $\overline{V}_r(y_j)+\tau$ that has directed Hausdorff distance at most 1 to the left Vector Gadgets of A, as

$$j\epsilon - \frac{3}{2}\epsilon - (-1 - \frac{1}{2}\epsilon) = 1 + (j - 1)\epsilon \leq 1$$

for $j = 1$. If $j = n$, then $\overline{V}_r(y_{j+1}) + \tau$ is undefined.

As $\delta_{\vec{H}}(B+\tau, A) \leq 1$, we know that $\overline{V}_r(y_j)+\tau$ has directed Hausdorff distance at most 1 to a gadget $V_r(x)$ for some $x \in X$. We claim that this distance cannot be closer than 1 as $\overline{V}_r(y_{j+1})+\tau$ must have a directed Hausdorff distance at most 1 from the right side of A or, in case $j=n$, due to the restrictions imposed by the Translation Gadget. Let us consider the case $j \neq n$ first. Any translation τ' which places $\overline{V}_r(y_{j+1})+\tau'$ in directed Hausdorff distance at most 1 from the right side of A needs to fulfill

$$1+\frac{1}{2}\epsilon-((j+1)\epsilon+\tau_1') \leq 1$$

and thus $\tau_1' \geq -(j+\frac{1}{2})\epsilon$, using the fact that each vector in Y contains at least one 0-entry. This, on the other hand, implies that $\overline{V}_r(y_j)+\tau'$ is in Hausdorff distance at least

$$j\epsilon-(j+\frac{1}{2})\epsilon-(-1-\frac{1}{2}\epsilon)=1$$

from $V_r(x)$. Now consider the case $j=n$. As by Lemma 5 we have $\tau_x \geq -(n+\frac{1}{2})\epsilon-\epsilon^2$, it follows that $\overline{V}_r(y_n)+\tau$ is in Hausdorff distance at least

$$n\epsilon-(n+\frac{1}{2})\epsilon-(-1-\frac{1}{2}\epsilon)=1$$

from $V_r(x)$, using the fact that each vector in Y contains at least one 0-entry (this is the reason why the ϵ^2 disappears).
By the arguments above, the two gadgets $\overline{V}_r(y_j)+\tau$ and $V_r(x)$ have to be horizontally aligned as required by Lemma 3. They also have to be vertically aligned as a vertical deviation would incur a Hausdorff distance larger than 1 for the pair of points in the two gadgets that are in horizontal distance 1. Then, applying Lemma 3, it follows that x an y_j are orthogonal.

It remains to argue why the above reduction implies the lower bound stated in the theorem. Assume we have an algorithm that computes the Hausdorff distance under translation for L_1 or L_∞ in time $(mn)^{1-\gamma}$ for some $\gamma > 0$. Then, given an Orthogonal Vectors instance X, Y with $|X|=m$ and $|Y|=n$, we can use the described reduction to obtain an equivalent Hausdorff under translation instance with point sets A, B of size $|A|=\mathcal{O}(md)$ and $|B|=\mathcal{O}(nd)$ and solve it in time $\mathcal{O}((mn)^{1-\gamma}\text{poly}(d))$, contradicting the Orthogonal Vectors Hypothesis. ◀

3.3 Generalization to L_p

We believe that we can extend the above construction such that it works for all L_p norms with $p \neq \infty$ by changing the spacing between 0 and 1 points of the Vector Gadgets and also set ϵ accordingly. More precisely, it seems that we can use ϵ^{2p} as spacing (instead of ϵ^2) and set $\epsilon < \frac{1}{40pmnd}$. The proofs should then be analogous to the L_1 case.

4 3Sum based $n^{2-o(1)}$ lower bound for $m \in \mathcal{O}(1)$

We now present a hardness result for the unbalanced case of the directed and undirected Hausdorff distance under translation. We base our hardness on another popular hypothesis of fined-grained complexity theory: the 3Sum Hypothesis. Before stating the hypothesis, let us first introduce the 3Sum problem.

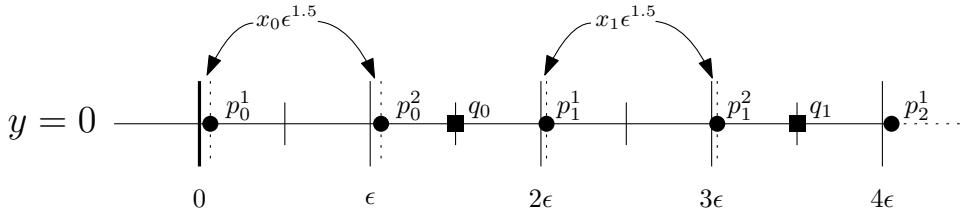

Figure 3 The A set of the low-level gadget of the 3Sum reduction, which is used to build the high-level gadgets. We just show the leftmost part of the gadget, but the remainder is similar.

▶ **Definition 8** (3Sum). *Given three sets of positive integers X, Y, Z all of size n, do there exist $x \in X, y \in Y, z \in Z$ such that $x + y = z$?*

The corresponding hardness assumption is the 3Sum Hypothesis.

▶ **Definition 9** (3Sum Hypothesis). *There is no $\mathcal{O}(n^{2-\epsilon})$ algorithm for 3Sum for any $\epsilon > 0$.*

There are several equivalent variants of the 3Sum problem. Most important for us is the convolution 3Sum problem, abbreviated as Conv3Sum [26, 14].

▶ **Definition 10** (Conv3SUM). *Given a sequence of positive integers $X = (x_1, \ldots, x_n)$ of size n, do there exist i, j such that $x_i + x_j = x_{i+j}$?*

This problem has a trivial $\mathcal{O}(n^2)$ algorithm and, assuming the 3Sum Hypothesis, this is also optimal up to lower order factors. As 3Sum and Conv3Sum are equivalent, a lower bound conditional on Conv3Sum implies a lower bound conditional on 3Sum.

Therefore, given a Conv3Sum instance defined by the set of integers X with $|X| = n$, we create an equivalent instance of the directed Hausdorff distance under translation for L_2 by constructing two sets of points A and B with $|A| = \mathcal{O}(n)$ and $|B| = \mathcal{O}(1)$ and providing a decision distance δ. Intuitively, we define a low-level gadget from which we build three high-level gadgets by rotation and scaling. Recall that in the Conv3Sum problem we have to find values i, j which fulfill the equation $x_i + x_j = x_{i+j}$. We encode the choice of these two values into the two dimensions of the translation. These three high-level gadgets then verify if the Conv3Sum equation is fulfilled. In the remainder of this section, we present the details of our reduction and prove that it implies the claimed lower bound.

4.1 Construction

Given an integer Conv3Sum instance with $X \subset [M]$ where $n = |X|$, we now describe the construction of the Hausdorff distance under translation instance with point sets A, B and threshold distance δ. We use a small enough ϵ, e.g., $\epsilon = (4Mn^2)^{-2}$, as value for microtranslations. Furthermore, we set $\delta = 1+4n^2\epsilon^2$. The additional $4n^2\epsilon^2$ term compensates for the small variations in distance that occur on microtranslations due to the curvature of the L_2-ball.

4.1.1 Low-level gadget

We use a single low-level gadget, which is then scaled and rotated to obtain high-level gadgets. This gadget consists of two point sets A_l and B_l. The point set A_l contains what we call *number points* p_i^1, p_i^2 and *filling points* q_i for $0 \leq i < n$. The set B_l just contains two points:

r_1 and r_2. The number points p_i^1, p_i^2 encode the number x_i, while the filling points make sure that no other translations than the desired ones are possible. See Figure 3 for an overview. All of the points in this gadget are of the form $(x, 0)$. The number points are

$$p_i^1 = \left(2i\epsilon + x_i\epsilon^{1.5}, 0\right), \quad p_i^2 = p_i^1 + (\epsilon, 0)$$

for $0 \leq i < n$. The filling points are

$$q_i = \left(\left(2i + \frac{3}{2}\right)\epsilon, 0\right)$$

for $0 \leq i < n$.

The points in B_l should introduce a gap to only allow alignment of the number gadgets such that the microtranslations (i.e., those in the order of $\epsilon^{1.5}$) correspond to the number of the gap in the number gadget. Thus B_l contains the points

$$r_1 = (-1, 0), \quad r_2 = (1 + \epsilon, 0).$$

Before we prove properties of the low-level gadget, we first prove that the error that is happening due to the curvature of the L_2-ball is small.

▶ **Lemma 11.** *Let $(p_x, p_y), (q_x, q_y) \in \mathbb{R}^2$ be two points with $|p_x - q_x| \in [\frac{1}{2}, 2]$ and $p_y = q_y$. For any $\tau \in [0, (2n-1)\epsilon]^2$, we have*

$$|p_x - (q_x + \tau_x)| \leq \|p - (q + \tau)\|_2 \leq |p_x - (q_x + \tau_x)| + 4n^2\epsilon^2.$$

Proof. As each component is a lower bound to the L_2 norm, the first inequality follows. Thus, let us prove the second inequality. We first transform

$$\|p - (q + \tau)\|_2 = \sqrt{(p_x - (q_x + \tau_x))^2 + \tau_y^2} = |p_x - (q_x - \tau_x)|\sqrt{1 + \tau_y^2/(p_x - (q_x + \tau_x))^2}.$$

Because $\sqrt{1+x} \leq 1 + \frac{x}{2}$ for any $x \geq 0$, we have

$$\|p - (q + \tau)\|_2 \leq |p_x - (q_x - \tau_x)| + \tau_y^2/(2|p_x - (q_x - \tau_x)|).$$

As $\tau_y \leq 2(n-1)\epsilon$ and $|p_x - (q_x - \tau_x)| \geq \frac{1}{2}$, we obtain the desired upper bound. ◀

An analogous statement holds when swapping the x and y coordinates. Note that the $4n^2\epsilon^2$ term also occurs in the value of δ that we chose, as this is how we compensate for these errors in our construction. While we have to consider this error in the following arguments, it already seems that it will be insignificant due to its magnitude.

We now state two lemmas which show how the Hausdorff distance under translation decision problem is related to the structure of the low-level gadget.

▶ **Lemma 12.** *Given a low-level gadget A_l, B_l as constructed above and the translation being restricted to $\tau \in [0, (2n-1)\epsilon]^2$, it holds that if $\delta_{\vec{H}}(A_l, B_l + \tau) \leq \delta$, then*

$$\exists i \in \mathbb{N} : \tau_x = 2i\epsilon + x_i\epsilon^{1.5} \pm 4n^2\epsilon^2.$$

Proof. Let $\tau \in [0, (2n-1)\epsilon]^2$ and assume $\delta_{\vec{H}}(A_l, B_l + \tau) \leq \delta$. Then all points in A_l are in distance at most δ from one of the two points in B_l. Furthermore, both points in $B_l + \tau$ also have at least one close point in A_l, as

$$\|r_1+\tau-p_0^1\|_2 \leq 1-\tau_x+4n^2\epsilon^2 \leq \delta \quad \text{and} \quad \|r_2+\tau-q_{n-1}\|_2 \leq 1+\tau_x-(2n-\frac{1}{2})\epsilon+4n^2\epsilon^2 < \delta,$$

using Lemma 11.

The gaps between neighboring points in A_l either have width close to $\frac{1}{2}\epsilon$, if the gap is between a number point and a filling point (p_i^1 and q_{i-1}, or p_i^2 and q_i), or they have a width of ϵ, if the gap is between two number points (p_i^1 and p_i^2). Furthermore, the two points in B_l have distance $2+\epsilon$, so there is an $\epsilon - 8n^2\epsilon^2$ gap between their δ-balls. Thus, there is an i such that p_i^1 has distance at most δ to r_1, and p_i^2 has distance at most δ to r_2. This alignment of the gadgets can only be realized by a translation τ for which

$$\tau_x = 2i\epsilon + x_i\epsilon^{1.5} \pm 4n^2\epsilon^2,$$

which completes the proof. ◀

▶ **Lemma 13.** *Given a low-level gadget A_l, B_l as constructed above and the translation being restricted to $\tau \in [0,(2n-1)\epsilon]^2$, it holds that if*

$$\exists i \in \mathbb{N} : \tau_x = 2i\epsilon + x_i\epsilon^{1.5},$$

then $\delta_H(A_l, B_l+\tau) \leq \delta$.

Proof. Let $i \in \mathbb{N}$ and let $\tau_x = 2i\epsilon + x_i\epsilon^{1.5}$. Consider any translations $\tau \in \{\tau_x\} \times [0, 2(n-1)\epsilon]$. Due to the restricted translation and Lemma 11, we can disregard the error terms that arise from the vertical translation τ_y as they are compensated for by δ. Then all the points in A_l before and including p_i^1 are in distance at most δ from $r_1 \in B_l + \tau$ and all the points afterwards are in distance at most δ from $r_2 \in B_l + \tau$. Clearly, both points in $B_l + \tau$ then also have points from A_l in distance δ, and thus $\delta_H(A_l, B_l+\tau) \leq \delta$. ◀

4.1.2 High-level gadgets

This construction is inspired by the hard instance that was given in [28]. We want to obtain a grid of translations of spacing ϵ with some microtranslations in the $\mathcal{O}(\epsilon^{1.5})$ range. We already defined the low-level gadget above, and we now define the high-level gadgets.

Column Gadget

The column gadget induces columns in translational space, i.e., it enforces that valid translations have to lie on one of these columns. The column gadget is actually the low-level gadget we already described above. You can see a sketch of this gadget in Figure 4a. To semantically distinguish it from the low-level gadget, we refer to the point sets of the column gadget as A_c and B_c.

Row Gadget

The row gadget induces rows in translational space, i.e., it enforces that valid translations have to lie on one of these rows. We obtain the row gadget by rotating all points in the low-level gadget around the origin by $\pi/2$ counterclockwise. You can see a sketch of this gadget in Figure 4b. We call the point sets of the row gadget A_r and B_r.

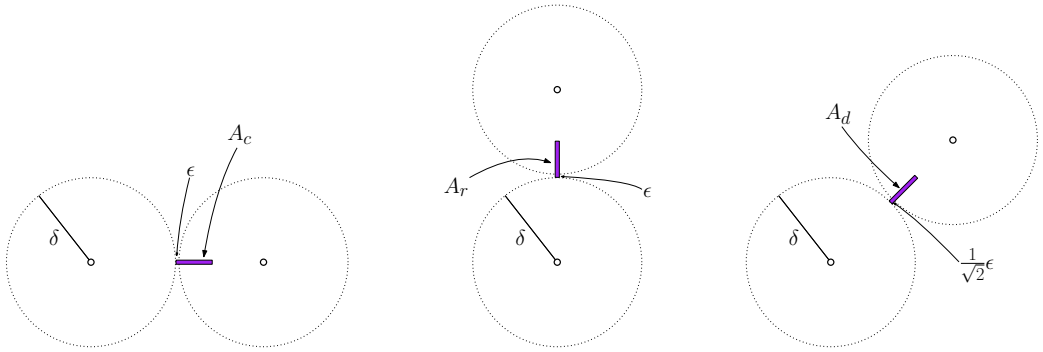

(a) Column Gadget. **(b)** Row Gadget. **(c)** Diagonal Gadget.

Figure 4 Three of the high-level gadgets. The points of A are all in the low-level gadgets, while the points in B are explicitly shown including their δ-ball.

Diagonal Gadget

The diagonal gadget induces diagonals in translational space, i.e., it enforces that valid translations have to lie on one of these diagonals. As opposed to the column and row gadget, the diagonal gadget also has to be scaled. We scale the sets A_l and B_l separately. We scale A_l such that the gap between the number point pairs p_i^1, p_i^2 becomes $\frac{1}{\sqrt{2}}\epsilon$. And we scale B_l such that the gap between the points becomes $2 + \frac{1}{\sqrt{2}}\epsilon$. After scaling, we rotate the points counterclockwise around the origin by $\pi/4$. You can see a sketch of this gadget in Figure 4c. We call the point sets of the diagonal gadget A_d and B_d.

Translation Gadget

To restrict the translations for the directed Hausdorff distance under translation, we introduce another gadget. The first set of points A_t contains

$$z_l := (-1 + (2n-1)\epsilon, 0), \quad z_r := (1,0), \quad z_b := (0, -1 + (2n-1)\epsilon), \quad z_t := (0,1).$$

The second point set B_t only contains the origin $z_c := (0,0)$. We want to make sure that this gadget behaves well in a certain range.

▶ **Lemma 14.** *Given $\tau \in [0, (2n-1)\epsilon]^2$, it holds that $\delta_H(A_t, B_t + \tau) \leq \delta$.*

Proof. As B_t has a point on all sides, clearly $\delta_{\vec{H}}(B_t + \tau, A_t) \leq \delta$. Furthermore,

$$\|z_l - (z_c + \tau)\|_2 \leq 1 + 4n^2\epsilon^2 \leq \delta \quad \text{and} \quad \|z_r - (z_c + \tau)\|_2 \leq \delta,$$

using Lemma 11. Analogous statements hold for z_b and z_t. Thus, also $\delta_{\vec{H}}(A_t, B_t + \tau) \leq \delta$. ◀

4.1.3 Complete construction

To obtain the final sets of the reduction, we now place all four described high-level gadgets (i.e., column gadget, row gadget, diagonal gadget, and translation gadget) far enough apart. More explicitly, the point sets A, B of the Hausdorff distance under translation instance are defined as

$$A := A_c \cup (A_r + (10,0)) \cup (A_d + (20,0)) \cup (A_t + (30,0))$$

and

$$B := B_c \cup (B_r + (10,0)) \cup (B_d + (20,0)) \cup (B_t + (30,0)).$$

The far placement ensures that the two point sets of the respective gadgets have to be matched to each other for a decision distance δ.

4.2 Proof of correctness

First, we want to ensure that everything relevant happens in a very small range of translations.

▶ **Lemma 15.** *Let $\tau \in \mathbb{R}^2$. If $\delta_{\vec{H}}(A, B + \tau) \leq \delta$, then $\tau \in [0, (2n-1)\epsilon]^2$.*

Proof. Note that for a Hausdorff distance at most δ, the sets A_c and B_c have to matched to each other and analogously for A_r, B_r, and A_d, B_d, and A_t, B_t. To show the contrapositive, now assume $\tau \notin [0, (2n-1)\epsilon]^2$. For simplicity, we refer to the points in the high-level gadgets with the notation of the low-level gadget. Additionally, due to the translation gadget, we have

$$\|z_l - (z_c + \tau)\|_2 > \delta \quad \text{for} \quad \tau_x > (2n-1)\epsilon + 4n^2\epsilon^2,$$

and

$$\|z_r - (z_c + \tau)\|_2 > \delta \quad \text{for} \quad \tau_x < -4n^2\epsilon^2.$$

We now show that under these restricted translations and as $\delta_{\vec{H}}(A, B + \tau) \leq \delta$, both points r_1, r_2 in B_c have at least one point of A_c in distance δ. In the column gadget for $\tau_x \in [-4n^2\epsilon^2, 0)$, we have

$$\|(r_1 + \tau) - p_0^1\|_2 \geq |-1 - (p_0^1)_x + \tau_x| > \delta \quad \text{and} \quad \|(r_2 + \tau) - p_0^1\|_2 \geq 1 + \epsilon - \mathcal{O}(\epsilon^{1.5}) > \delta,$$

for small enough ϵ. On the other hand, for $\tau_x \in ((2n-1)\epsilon, (2n-1)\epsilon + 4n^2\epsilon^2]$, we have

$$\|r_2 + \tau - p_{n-1}^2\|_2 \geq 1 + \epsilon + \tau_x - (2n-1)\epsilon > \delta \quad \text{and} \quad \|r_1 + \tau - p_{n-1}^2\|_2 \geq 1 + \epsilon - \mathcal{O}(\epsilon^{1.5}) > \delta$$

for small enough ϵ. An analogous argument holds for the row gadget and τ_y, as the row gadget is just a rotated version of the column gadget and the translation gadget is symmetric with respect to these gadgets. ◀

We can now prove the main result of this section.

▶ **Theorem 16.** *Computing the directed or undirected Hausdorff distance under translation in L_2 for two sets of size n and $\mathcal{O}(1)$ cannot be solved in time $\mathcal{O}(n^{2-\gamma})$ for any $\gamma > 0$, unless the 3Sum Hypothesis fails.*

Proof. We construct a Hausdorff under translation instance in this proof from a Conv3Sum instance as described previously in this section, and then show that they are equivalent. We first consider how to apply Lemma 12 and Lemma 13 to the diagonal gadget. More precisely, we consider which translations align the gaps of A_d and B_d as is used in these two lemmas. Due to the scaling of the gadget, these translations are of the form $\sqrt{2}\tau_x = 2k\epsilon + x_k\epsilon^{1.5}$. By the rotation, we then obtain translations of the form

$$\frac{\sqrt{2}(\tau_x + \tau_y)}{\sqrt{2}} = \|\tau\|_1 = 2(i+j)\epsilon + x_{i+j}\epsilon^{1.5}.$$

$\Leftarrow$: Assume X is a positive Conv3Sum instance. Then there exist x_i, x_j such that $x_i + x_j = x_{i+j}$. Consider $\tau = (2i\epsilon + x_i\epsilon^{1.5}, 2j\epsilon + x_j\epsilon^{1.5})$ as translation. Due to Lemma 13, we have that $\delta_H(A_c, B_c + \tau) \leq \delta$ and analogously $\delta_H(A_r, B_r + \tau) \leq \delta$. By the initial observation we can also apply Lemma 13 to the diagonal gadget, and thus $\delta_H(A_d, B_d + \tau) \leq \delta$. Finally, by Lemma 14, we also have that $\delta_H(A_t, B_t + \tau) \leq \delta$ for the given τ.

$\Rightarrow$: Assume $\delta_{\bar{H}}^{T}(A,B) \leq \delta$. From Lemma 15, it follows that $\tau \in [0,(2n-1)\epsilon]^2$. Then, due to Lemma 12 and the initial observation about the diagonal gadget, we have that there exist i,j,k that fulfill

$$\tau_x = 2i\epsilon + x_i\epsilon^{1.5} \pm 4n^2\epsilon^2 \quad \text{and} \quad \tau_y = 2j\epsilon + x_j\epsilon^{1.5} \pm 4n^2\epsilon^2 \quad \text{and} \quad \tau_x + \tau_y = 2k\epsilon + x_k\epsilon^{1.5} \pm 4n^2\epsilon^2.$$

It follows that

$$2i\epsilon + x_i\epsilon^{1.5} + 2j\epsilon + x_j\epsilon^{1.5} \pm 8n^2\epsilon^2 = 2k\epsilon + x_k\epsilon^{1.5} \pm 4n^2\epsilon^2,$$

and thus $i+j=k$ and $x_i + x_j = x_k$.

It remains to argue why the above reduction implies the lower bound stated in the theorem. Assume we have an algorithm that computes the Hausdorff distance under translation in L_2 in time $\mathcal{O}(n^{2-\gamma})$ for some $\gamma > 0$. Then, given a Conv3Sum instance X with $|X| = n$, we can use the described reduction to obtain an equivalent Hausdorff under translation instance with point sets A, B of size $|A| = \mathcal{O}(n)$ and $|B| = \mathcal{O}(1)$ and solve it in time $\mathcal{O}(n^{2-\gamma})$, contradicting the 3Sum Hypothesis. ◀

References

1. Amir Abboud, Arturs Backurs, Karl Bringmann, and Marvin Künnemann. Fine-grained complexity of analyzing compressed data: Quantifying improvements over decompress-and-solve. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 192–203. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.26.
2. Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. Consequences of faster alignment of sequences. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 39–51. Springer, 2014. doi:10.1007/978-3-662-43948-7_4.
3. Helmut Alt and Michael Godau. Computing the Fréchet Distance between Two Polygonal Curves. *Int. J. Comput. Geometry Appl.*, 5:75–91, 1995. doi:10.1142/S0218195995000064.
4. Amihood Amir, Timothy M. Chan, Moshe Lewenstein, and Noa Lewenstein. On hardness of jumbled indexing. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 114–125. Springer, 2014. doi:10.1007/978-3-662-43948-7_10.
5. Chanderjit Bajaj. The algebraic degree of geometric optimization problems. *Discrete & Computational Geometry*, 3(2):177–191, 1988.
6. Gill Barequet and Sariel Har-Peled. Polygon containment and translational in-hausdorff-distance between segment sets are 3sum-hard. *International Journal of Computational Geometry & Applications*, 11(04):465–474, 2001. doi:10.1142/S0218195901000596.
7. Karl Bringmann. Why walking the dog takes time: Frechet distance has no strongly subquadratic algorithms unless SETH fails. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 661–670. IEEE Computer Society, 2014. doi:10.1109/FOCS.2014.76.
8. Karl Bringmann and Marvin Künnemann. Multivariate fine-grained complexity of longest common subsequence. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1216–1235. SIAM, 2018. doi:10.1137/1.9781611975031.79.

9 Karl Bringmann, Marvin Künnemann, and André Nusser. Fréchet distance under translation: Conditional hardness and an algorithm via offline dynamic grid reachability. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2902–2921. SIAM, 2019. doi:10.1137/1.9781611975482.180.

10 Karl Bringmann, Marvin Künnemann, and André Nusser. When lipschitz walks your dog: Algorithm engineering of the discrete fréchet distance under translation. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, volume 173 of *LIPIcs*, pages 25:1–25:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ESA.2020.25.

11 Karl Bringmann and Wolfgang Mulzer. Approximability of the discrete fréchet distance. *J. Comput. Geom.*, 7(2):46–76, 2016. doi:10.20382/jocg.v7i2a4.

12 Kevin Buchin, Anne Driemel, Natasja van de L'Isle, and André Nusser. klcluster: Center-based clustering of trajectories. In Farnoush Banaei Kashani, Goce Trajcevski, Ralf Hartmut Güting, Lars Kulik, and Shawn D. Newsam, editors, *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2019, Chicago, IL, USA, November 5-8, 2019*, pages 496–499. ACM, 2019. doi:10.1145/3347146.3359111.

13 Kevin Buchin, Tim Ophelders, and Bettina Speckmann. SETH says: Weak fréchet distance is faster, but only if it is continuous and in one dimension. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2887–2901. SIAM, 2019. doi:10.1137/1.9781611975482.179.

14 Timothy M. Chan and Qizheng He. Reducing 3sum to convolution-3sum. In Martin Farach-Colton and Inge Li Gørtz, editors, *3rd Symposium on Simplicity in Algorithms, SOSA@SODA 2020, Salt Lake City, UT, USA, January 6-7, 2020*, pages 1–7. SIAM, 2020. doi:10.1137/1.9781611976014.1.

15 L. P. Chew, D. Dor, A. Efrat, and K. Kedem. Geometric pattern matching in d-dimensional space. *Discrete & Computational Geometry*, 21(2):257–274, February 1999. doi:10.1007/PL00009420.

16 L. Paul Chew and Klara Kedem. Improvements on geometric pattern matching problems. In Otto Nurmi and Esko Ukkonen, editors, *Algorithm Theory — SWAT '92*, Lecture Notes in Computer Science, pages 318–325. Springer Berlin Heidelberg, 1992.

17 Mark de Berg, Atlas F. Cook, and Joachim Gudmundsson. Fast fréchet queries. *Computational Geometry*, 46(6):747–755, 2013. doi:10.1016/j.comgeo.2012.11.006.

18 A. Efrat, A. Itai, and M. J. Katz. Geometry Helps in Bottleneck Matching and Related Problems. *Algorithmica*, 31(1):1–28, 2001. doi:10.1007/s00453-001-0016-8.

19 Andriy Fedorov, Eric Billet, Marcel Prastawa, Guido Gerig, Alireza Radmanesh, Simon K Warfield, Ron Kikinis, and Nikos Chrisochoides. Evaluation of brain mri alignment with the robust hausdorff distance measures. In *International Symposium on Visual Computing*, pages 594–603. Springer, 2008.

20 Anka Gajentaan and Mark H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom.*, 5:165–185, 1995. doi:10.1016/0925-7721(95)00022-2.

21 Felix Hausdorff. *Grundzüge der Mengenlehre*, volume 7. von Veit, 1914.

22 Daniel P Huttenlocher, Klara Kedem, and Micha Sharir. The upper envelope of voronoi surfaces and its applications. *Discrete & Computational Geometry*, 9(3):267–291, 1993.

23 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.

24 James King. A survey of 3sum-hard problems, 2004.

25 Meinard Müller. *Information retrieval for music and motion*. Springer, 2007. doi:10.1007/978-3-540-74048-3.

26 Mihai Patrascu. Towards polynomial lower bounds for dynamic problems. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, STOC '10, page 603–610, New York, NY, USA, 2010. Association for Computing Machinery. doi:10.1145/1806689.1806772.

27 Günter Rote. Computing the minimum Hausdorff distance between two point sets on a line under translation. *Information Processing Letters*, 38(3):123–127, 1991. doi:10.1016/0020-0190(91)90233-8.

28 W. J. Rucklidge. Lower bounds for the complexity of the graph of the Hausdorff distance as a function of transformation. *Discrete & Computational Geometry*, 16(2):135–153, 1996. doi:10.1007/BF02716804.

29 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proc. ICM*, volume 3, pages 3431–3472. World Scientific, 2018.

30 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005. doi:10.1016/j.tcs.2005.09.023.