Computing L_∞ Hausdorff Distances Under Translations: The Interplay of Dimensionality, Symmetry and Discreteness

Sebastian Angrick
Karlsruhe Institute of Technology, Germany

Kevin Buchin
TU Dortmund, Germany

Geri Gokaj
Karlsruhe Institute of Technology, Germany

Marvin Künnemann
Karlsruhe Institute of Technology, Germany

Abstract

To measure the similarity of the *shape* of point sets, rather than their mere closeness in space, various notions of a *Hausdorff distance under translation* have been investigated. Specifically, let P and Q denote point sets of n and m points, respectively, in $\mathbb{R}^d$. We consider the task of computing the minimum distance $d(P, Q + \tau)$ over an admissible set of translations $\tau \in T$, where $d(\cdot, \cdot)$ denotes the Hausdorff distance under the L_∞-norm. As variants, we distinguish between *continuous* ($T = \mathbb{R}^d$) or *discrete* (T is a given finite set of t translations) as well as *directed* or *undirected* (choosing the directed or undirected Hausdorff distance for $d(\cdot, \cdot)$).

We seek to apply the paradigm of fine-grained complexity to understand the complexity of these variants, and in particular: How is the running time influenced by the dimension d, the relationship between n and m, and the specific choice of variant? As our main results, we obtain:

- The asymmetric definition of the most studied variant, the continuous directed Hausdorff distance, results in an *intrinsically asymmetric* time complexity: While (Chan, SoCG'23) established a symmetric $\tilde{O}((nm)^{d/2})$ upper bound for all $d \geq 3$ and proved it to be conditionally optimal for *combinatorial* algorithms whenever $m \leq n$, we show that this lower bound does not hold for the case $n \ll m$, by providing a combinatorial, almost-linear-time algorithm for $d = 3$ and $n = m^{o(1)}$. We further prove *general*, i.e., non-combinatorial, conditional lower bounds for $d \geq 3$, in particular: (1) $m^{\lfloor d/2 \rfloor - o(1)}$ for small n and (2) $n^{d/2 - o(1)}$ for $d = 3$ and small m.
- We observe that the directed and undirected case is closely related, in particular, all our lower bounds for $d \geq 3$ hold for both the directed and undirected variant. A remarkable exception is the case of $d = 1$ for which we provide a conditional separation. Specifically, in contrast to the undirected variants being solvable in near-linear time (Rote, IPL'91), we show that the directed variants are at least as hard as the additive problem MaxConv LowerBound introduced in (Cygan, Mucha, Wegrzycki and Wlodarczyk, TALG'19).
- We show that the discrete variants reduce to a variant of 3SUM for $d \leq 3$. This gives a barrier in proving a tight lower bound of these variants under the Orthogonal Vectors Hypothesis (OVH); in contrast, the continuous variants admit a tight conditional lower bound under OVH in $d = 2$ (Bringmann, Nusser, JoCG'21).

These results reveal an intricate interplay of dimensionality, symmetry and discreteness in determining the fine-grained complexity of computing Hausdorff distances under translation.

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

Keywords and phrases Hausdorff Distance, Fine-Grained Complexity, Computational Geometry, Translation-Invariant Similarity Measures

Digital Object Identifier 10.4230/LIPIcs.SoCG.2026.7

42nd International Symposium on Computational Geometry (SoCG 2026).
Editors: Hee-Kap Ahn, Michael Hoffmann, and Amir Nayyeri; Article No. 7; pp. 7:1–7:15
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Related Version *Full Version*: https://arxiv.org/pdf/2603.08890

Funding *Geri Gokaj*: Partially supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 462679611.
Marvin Künnemann: Partially supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 462679611.

Acknowledgements The authors would like to thank Jean Cardinal, Anita Dürr, Nick Fischer, Virginia Vassilevska Williams, and Karol Węgrzycki for stimulating discussions and contributions during the early stages of this work, as well as the Lorentz center and the organizers of its workshop on *Fine-Grained & Parameterized Computational Geometry* for facilitating these discussions.

1 Introduction

Consider the classic Hausdorff distance of given point sets P and Q, which comes in two flavors: the *directed* Hausdorff distance $\delta_{\vec{H}}(P, Q) := \max_{p\in P} \min_{q\in Q} \|p - q\|$, as well as the *undirected* Hausdorff distance $\delta_H(P, Q) := \max\{\delta_{\vec{H}}(P, Q), \delta_{\vec{H}}(Q, P)\}$. Here and throughout the paper, we fix our metric space to be $(\mathbb{R}^d, L_\infty)$. The two variants are sometimes also referred to as one-sided/bidirectional or one-way/two-way. Furthermore, δ_H yields a metric, in contrast to $\delta_{\vec{H}}$. For any constant $d \in \mathbb{N}$ and point sets P and Q of n and m points in $\mathbb{R}^d$, we can compute these distance in near-linear time[1] $\tilde{O}(n + m)$ using L_∞-nearest neighbor search, e.g., in a range tree [18].

A computationally more challenging task is to compute the Hausdorff distance under some transformation, i.e., to minimize the Hausdorff distance of P' and Q among all sets P' obtained from P using a set of admissible transformations. Such notions are well-studied from the perspective of *shape matching* or *geometric pattern matching*, see, e.g. [3, 4], where P and Q model geometric objects and we are interested in comparing their general shapes rather than their mere closeness in space. In this paper, we focus on arguably one of the most basic settings, i.e., computing L_∞ Hausdorff distances under translation. Specifically, when T denotes a set of admissible translations, the Hausdorff distance under translation is defined as $\min_{\tau\in T} \delta(P + \tau, Q)$ where δ denotes either the directed or undirected Hausdorff distance, i.e., $\delta \in \{\delta_{\vec{H}}, \delta_H\}$.

The most popular setting is $T = \mathbb{R}^d$, which yields the natural translation-invariant version of the Hausdorff distance; we refer to this as the **(continuous) Hausdorff distance under translation (HuT)**. This measure has received significant interest from the computational geometry community both for L_∞ and other norms [20, 16, 17, 21, 11, 10, 5, 7, 9], yet only for $d = 2$ tight upper and conditional lower bounds are known (see below for details).

Another natural version is to let T be a finite set given as part of the input. We refer to this as the **discrete Hausdorff distance under translation (DiscHuT)**. This version occurs naturally in the context of approximating the (continuous) Hausdorff distance under translation. More precisely, computing an α-approximation for the continuous version can be achieved by choosing $f(\alpha, d)$ many translations T such that $\min_{\tau\in T} \delta_{\vec{H}}(P + \tau, Q) \leq \alpha \min_{\tau\in\mathbb{R}^d} \delta_{\vec{H}}(P + \tau, Q)$. Since we can compute the discrete Hausdorff distance under translation in time $\tilde{O}(|T|(|P| + |Q|))$, this yields baseline α-approximation algorithms for the continuous Hausdorff distance under translation with running time $\tilde{O}(f(\alpha, d)(n + m))$; see also [15, 22]. Consequently, obtaining faster algorithms for the discrete Hausdorff distance

[1] Throughout the paper, we use $\tilde{O}(\cdot)$, to hide polylogarithmic terms in n and m. Since we assume d to be constant, this includes factors of the form $\log^{O(d)} nm$.

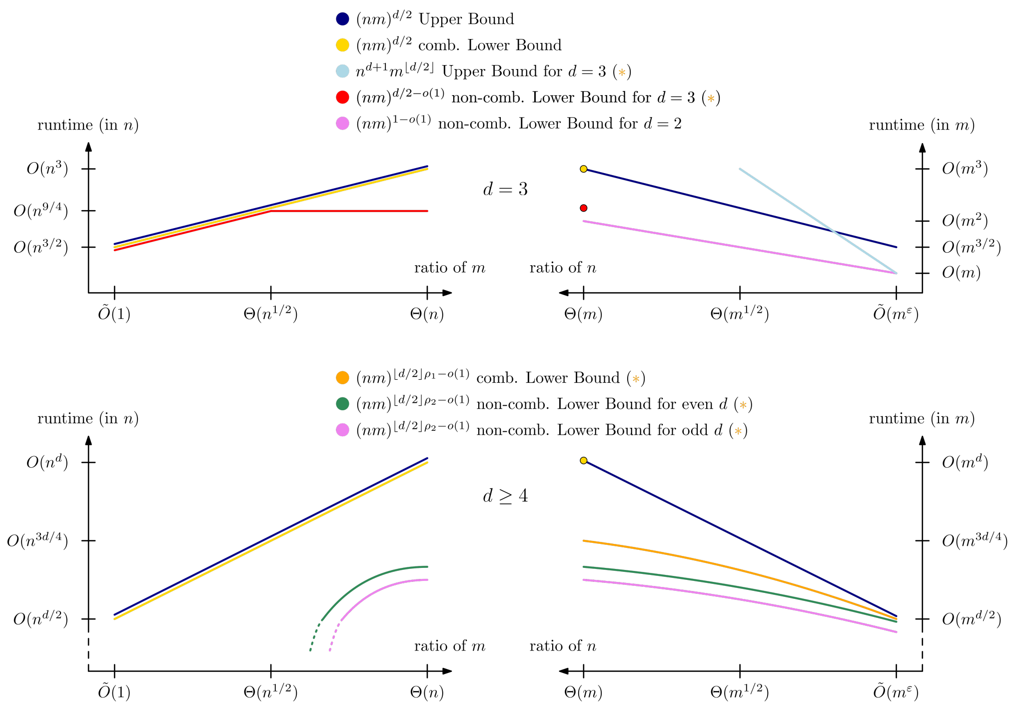

Figure 1 Schematic figure of upper and lower bounds for directed Hausdorff under Translation in 3-D (upper figure) and d-D for $d \geq 4$ (lower figure), novel results are marked with (*) . ρ_1, ρ_2 are functions on the ratio of n, m which correspond to values described in Theorem 3.7.

under translation transfers to faster approximation algorithms for the continuous version. Furthermore, the discrete version recently received interest from the perspective of a structural fine-grained complexity theory: the directed discrete Hausdorff distance under translation is part of a problem pair that is fine-grained complete for the class $\mathsf{FOP}_{\mathbb{Z}}$, a class of polynomial-time problems formed from first-order properties on finite additive structures [14] (further discussion can be found in the full version).

Previous works on Hausdorff distances under translation mostly focus –with notable exceptions such as [7, 9]– on a single particular variant, often considering particular dimensions d and the *balanced* case, in which all inputs sets (i.e., P, Q, and in the discrete case also T) have roughly equal size n. In this work, we aim to give a more comprehensive analysis, which reveals an intricate interplay between the dimensionality, symmetry (directed/undirected and the relationship between n and m), and discreteness.

1.1 Our Results

Our first question concerns the fine-grained complexity of the algorithmically most studied variant, the (continuous) directed Hausdorff distance under translation. The currently best upper bounds are $O(nm \log nm)$ for $d \leq 2$ [11], $O((nm)^{d/2} \log nm)$ for $d = 3$ [10][2], and $O((nm)^{d/2}(\log \log nm)^{O(1)})$ for general $d \geq 4$ [9]. This suggests the following question:

Question 1: *Is $(nm)^{d/2 \pm o(1)}$ the time complexity of directed Hausdorff L_∞-distance under Translation?*

[2] While [10] only considers the case $m = n$, their analysis also holds for distinct n, m. A short explanation of their approach can also be found in the full version.

Table 1 An overview of known results of the L_∞-Hausdorff distance under Translation in 1-D and 2-D for the balanced case, i.e., $n = m$ ($= t$ in the discrete setting). Upper Bounds are marked in blue, lower bounds in green, while results under non-standard hypothesis are additionally marked with (†). Novel results are denoted with (∗) . Results referring to the full version are marked with §.

$d = 1$	Directed	Undirected
Continuous	$\tilde{O}(n^2)$ ([11]) (∗) $n^{2-o(1)}$ (§)†	$\tilde{O}(n)$ ([20]) $\Omega(n)$
Discrete	$\tilde{O}(n^2)$ (§) (∗) $n^{2-o(1)}$ (3.10)†	$\tilde{O}(n)$ (§) $\Omega(n)$

$d = 2$	Directed	Undirected
Continuous	$\tilde{O}(n^2)$ ([11]) $n^{2-o(1)}$ ([7])	$\tilde{O}(n^2)$ ([11], 2.3) $n^{2-o(1)}$ ([7])
Discrete	$\tilde{O}(n^2)$ (§) (∗) $n^{2-o(1)}$ (3.14)†	$\tilde{O}(n^2)$ (§) (∗) $n^{2-o(1)}$ (3.14)†

A conditional lower bound of Bringmann and Nusser [7] (based on the Orthogonal Vectors Hypothesis) answers this question positively for $d = 2$. Subsequently, Chan [9] gave evidence of optimality even for $d \geq 3$, by proving a conditional lower bound of $(nm)^{d/2-o(1)}$ for *combinatorial* algorithms[3] in the case that $m \leq n$. However, proving such a conditional lower bound against *general* algorithms for $d \geq 3$ (or better conditional lower bounds than $(nm)^{1-o(1)}$ for the case $n \ll m$) remained open until this work. Furthermore, for $d = 1$, no $\tilde{O}(n + m + (nm)^{d/2}) = \tilde{O}(n + m)$ algorithms are known, raising the question whether any superlinear lower bound can be proven for $d = 1$.

As our first main contribution, we give a perhaps surprising *negative* answer to Question 1: For $d = 3$ and $n = m^{o(1)}$, there exists an almost-linear time $m^{1+o(1)}$ algorithm, showing that in an unbalanced setting, the upper bound of $\tilde{O}((nm)^{d/2})$ can indeed be beaten (Section 3.4). Interestingly, this stands in contrast to the case of moderately small m: For $d = 3$ and any $m \leq \sqrt{n}$ we prove a tight conditional lower bound of $(nm)^{d/2-o(1)}$, by combining the proof techniques of Chan's combinatorial k-clique lower bound [9] and the tight 3-uniform hyperclique lower bound for Klee's measure problem in 3D [19] (Section 3.2). This reveals that the time complexity of the directed Hausdorff distance under translation is *inherently asymmetric* in n and m and breaks with tradition for algorithmic results on the Hausdorff distance under translation.

We investigate the case of small n further and observe that our algorithmic approach could plausibly generalize to an $\tilde{O}(n^{d+1}m^{\lfloor d/2\rfloor})$ algorithm (we elaborate on the missing ingredient in Section 3.4) – such an algorithm would beat the baseline algorithm for all odd d and small values of n. We also show that such an algorithm would indeed be conditionally tight: we also prove a corresponding lower bound of $m^{\lfloor d/2\rfloor-o(1)}$ for all $d \geq 4$, assuming the 3-uniform hyperclique hypothesis (Section 3.1).

Our conditional lower bounds can be adapted to the case of $n = \Theta(m)$, but no longer match the upper bound. In fact, we obtain a conditional lower bound of $n^{2.25-o(1)}$ for $d = 3$ which still leaves a gap of $n^{0.75\pm o(1)}$ to the upper bound of $\tilde{O}(n^3)$. It remains an open problem to settle the full time complexity for *any* relationship between n and m. An intriguing possibility consistent with our results is the existence of an $\tilde{O}(n^{\lceil d/2\rceil}m^{\lfloor d/2\rfloor})$-time algorithm.

[3] Here, combinatorial algorithms refers to algorithms that avoid the use of algebraic techniques underlying fast matrix multiplication algorithms beating $n^{3-o(1)}$ running time.

We turn to our second main question.

Question 2: *How does the choice of dimension, directed vs. undirected and discrete vs. continuous affect the time complexity?*

We observe that the time complexity of computing Hausdorff distances under translation exhibits an intricate interplay between the different aspects of the problem, illustrated in Table 1 and Figure 1.

On directed vs. undirected

A folklore reduction from computing the undirected Hausdorff distance to computing the directed Hausdorff distance under translation (Observation 2.3) shows that the undirected case is never harder than the (balanced) directed case. Could there exist a reduction in the other direction?

We show that in full generality, such a reduction appears unlikely: While for all $d \geq 3$, we can transfer our conditional lower bounds for the directed case to analogous bounds in the undirected case (analogous to the OVH-based lower bound for $d = 2$ [7], which also transfers to the undirected case) (Section 3.5), we obtain a conditional separation for $d = 1$. Notably, Rote [20] obtains a $\tilde{O}(n+m)$-time algorithm for computing the undirected Hausdorff distance under translation in $d = 1$. For the directed setting, we give a fine-grained reduction from the quadratic-time L_∞-Necklace Alignment problem [6], which in turn is at least as hard as the MaxConv LowerBound problem studied in [12]. We remark that we are not aware of any conditional lower bound for MaxConv LowerBound under an established hypothesis from fine-grained complexity theory; however, the question of subquadratic equivalence between this problem and the well-studied $(\min, +)$-convolution problem was raised as an open problem in 2017 [12]. We believe that our reduction to the (discrete or continuous) directed Hausdorff distance adds further reason to study the apparent quadratic-time complexity of MaxConv LowerBound, in hopes to either refute or justify this hardness barrier.

Circumventing the separation in $d = 1$, we give a general reduction from the undirected setting to the directed setting that blows up the dimension (Section 3.5). This allows us to transfer the lower bounds from the directed case for $d = 1$ to the undirected case in $d = 2$. In particular, we conclude that if MaxConv LowerBound truly requires quadratic-time, the undirected setting is near-linear time solvable if and only if $d = 1$.

On discrete vs. continuous

The tight lower bound based on OVH [7] convincingly settles the fine-grained complexity of the continuous setting in $d = 2$. In contrast, the only justification of hardness of the discrete setting is given by our reductions from MaxConv LowerBound and L_∞-Necklace Alignment (see above and the full version). Is there any reason why a tight lower bound based on OVH does not exist for the discrete setting?

We find evidence for the non-existence of such a reduction by showing a fine-grained reduction from (directed or undirected) discrete Hausdorff distance under translation in $d \leq 3$ to the All-Ints 3SUM problem (in the full version). In particular, this implies a tight fine-grained reduction from the balanced setting to 3SUM. Thus, establishing quadratic-time OVH-hardness of the balanced, discrete setting implies a tight reduction from OV to 3SUM, which is a central open problem in fine-grained complexity theory, and suffers from barriers based on the Nondeterministic Strong Exponential Time Hypothesis, see [8].

We remark that we obtain further connections from the class of additive problem $\mathsf{FOP}_{\mathbb{Z}}$ to the discrete variants. For space reasons, these results are deferred to the full version.

2 Preliminaries

A brief definition of relevant problems as well as the used hardness hypotheses can be found in the full version. We there also give a brief overview of known algorithms for the continuous and discrete variant.

We recall the Hausdorff distance as a distance measure between two point sets P, Q of sizes n, m. The metric space used for this paper will be $(\mathbb{R}^d, L_\infty)$. The undirected (symmetric) and directed (asymmetric) variant are defined as:

$$\delta_{\overrightarrow{H}}(P,Q) := \max_{p\in P}\min_{q\in Q} \|p-q\|_\infty \qquad \delta_H(P,Q) := \max\{\delta_{\overrightarrow{H}}(P,Q), \delta_{\overrightarrow{H}}(Q,P)\}$$

We show almost all of our conditional lower bounds already for the *decision variant* of Hausdorff under Translation, which receives a test distance $\delta > 0$ as additional input[4]. Clearly, the corresponding optimization problems, finding a translation minimizing δ, are at least as hard. Specifically, we study:

▶ **Definition 2.1** (Hausdorff under Translation (*HuT*)). *Given sets $P, Q \subset \mathbb{R}^d$ of n and m points, respectively, and a real value $\delta > 0$. Decide whether:*

$$\exists \tau \in \mathbb{R}^d : \delta_{\overrightarrow{H}}(P+\tau, Q) \le \delta \quad \equiv \quad \exists \tau \in \mathbb{R}^d\ \forall p \in P\ \exists q \in Q : \|(p+\tau)-q\|_\infty \le \delta$$

▶ **Definition 2.2** (Discrete Hausdorff under Translation (*DiscHuT*)). *Given sets $P, Q \subset \mathbb{Z}^d$ of n and m points, respectively, a real value $\delta > 0$ as well as a set of at most t translations $T \subset \mathbb{Z}^d$. Decide whether:*

$$\exists \tau \in T : \delta_{\overrightarrow{H}}(P+\tau, Q) \le \delta \quad \equiv \quad \exists \tau \in T\ \forall p \in P\ \exists q \in Q : \|(p+\tau)-q\|_\infty \le \delta$$

Note that our definition of the discrete variant assumes all input entries to be polynomially bounded integers[5].

To obtain the *undirected* versions of the above problems, we replace $\delta_{\overrightarrow{H}}(P+\tau, Q)$ by $\delta_H(P+\tau, Q)$, and denote the problem by *(Disc)uHuT*. If not specified otherwise, we consider the directed versions. We call the case in which all input sets have roughly equal size, i.e., $Q = \Theta(|P|)$ and additionally for the discrete case $|T| = \Theta(|P|)$, the *balanced* parameter setting, otherwise we say the parameters are *lopsided*.

Lastly, we give a folklore reduction from undirected to directed Hausdorff under Translation that works for both the continuous and discrete variant. By setting $\sigma = (M, 0, \ldots, 0)$ for a sufficiently large M, we see that $\delta_H(P+\tau, Q) = \delta_{\overrightarrow{H}}((P \cup (\sigma - Q)) + \tau, Q \cup (\sigma - P))$.

▶ **Observation 2.3.** *The undirected (Discrete) Hausdorff under Translation problem with input sizes n, m reduces to the directed version with input sizes $n' = m' = n + m$.*

While this reduction introduces additional points to the input, we discuss in the full version that an input size maintaining reduction to a single instance of directed HuT is unlikely.

As our model of computation, we use the Real RAM for continuous HuT, while for the discrete HuT variant with point sets in $\mathbb{Z}^d$ we assume a Word RAM machine model with word size $\Theta(\log n)$.

---

[4] Note that the case $\delta = 0$, i.e., finding an exact match, allows an almost-linear-time algorithm [13] and is thus uninteresting for our purposes.

[5] One may also assume rational coordinates, which is equivalent by a scaling argument.

2.1 Translation Problems

For clarity of presentation, it is often useful to view the Hausdorff under Translation decision problem as a translation problem involving certain *shapes*.

▶ **Definition 2.4** (Translation Problem). *Given a set of n points P and m shapes $\mathcal{Q}$ in $\mathbb{R}^d$. Compute a translation $\tau \in T \subseteq \mathbb{R}^d$ such that*

$$\forall p \in P \; \exists Q \in \mathcal{Q} : p + \tau \in Q.$$

For an instance of HuT with point sets P, Q and $\delta > 0$, if we construct hypercubes $\mathcal{Q}$ with side length 2δ and centers $q \in Q$, the region $T^* = \bigcap_{p \in P} \bigcup_{\hat{Q} \in \mathcal{Q}} \hat{Q} - p$ is exactly the region of feasible translations, i.e., we have $\tau \in T^* \iff \forall p \in P \; \exists q \in Q : \delta_{\overrightarrow{H}}(P + \tau, Q) \leq \delta$ [1]. We call this case, where all shapes $\mathcal{Q}$ are hypercubes of the same size, the *Translation Problem with Hypercubes (TPwC)*. In case all shapes of Q are axis-aligned boxes (of arbitrary size), we call it the *Translation Problem with Boxes (TPwB)*. While the TPwC problem is equivalent to HuT, there is no reduction to HuT known within the same ambient space from the Translation Problem with Boxes or even from the translation problem with hypercubes of different sizes. We further use a third variant.

▶ **Definition 2.5** (Translation Problem with Orthants (TPwO)). *Given a target box B with a maximum side length δ_0. Given a set of k Translation Problem with Hypercube instances, each on a point set P_i and hypercubes $\mathcal{Q}_i$ with side length $\delta > \delta_0$, such that $\sum_{i \in [k]} |P_i| = n$, $\sum_{i \in [k]} |\mathcal{Q}_i| = m$. Compute a translation $\tau \in \mathbb{R}^d$ such that $\forall i \in [k] : P_i + \tau \subset B$ and*

$$\forall i \in [k] \; \forall p_i \in P_i \; \exists Q_i \in \mathcal{Q}_i : p_i + \tau \in Q_i.$$

We call the shapes of the TPwO instance, i.e., the shapes of its sub-instances, orthants as within the target box all hypercubes are indistinguishable from orthants. An important property of this problem, which we will use later, is that we may fix the side-length of all orthants to a value $\delta > \delta_0$. Similar to [9, Lemma 3], the TPwO problem reduces to the TPwC and thus has a reduction to HuT, its proof is deferred to the full version. A converse reduction, from HuT to TPwO, is generally not known.

▶ **Lemma 2.6.** *The Translation Problem with Orthants on n points and m orthants reduces in linear time to the Translation Problem with Hypercubes on $\Theta(n)$ points and $\Theta(m)$ hypercubes.*

3 Technical Overview for Selected Results

In this section, we discuss the main ideas of our lower bounds and algorithm. We present the general structure of the proofs and give sketches on the most important technical details. For space reasons, we defer further details as well as a complete discussion of our results and further extensions to the full version.

3.1 (Non-)Combinatorial Lower Bounds for Arbitrary n, m

For our lower bounds in higher dimensions ($d \geq 4$), we give a reduction in three steps from the k-hyperclique problem to HuT, see the full version for a formal definition and its hardness hypothesis.

In the first step, we encode any k-tuple of vertices of a hypergraph $H = (V, E)$ as cells inside a $\lfloor d/2 \rfloor$-dimensional hypercube, see Figure 2. The used encoding was introduced by Chan [9]. For any non-edge $\bar{e} \in \binom{V}{3} \setminus E$, the *feasible region* of $\bar{e}$ are the cells inside the

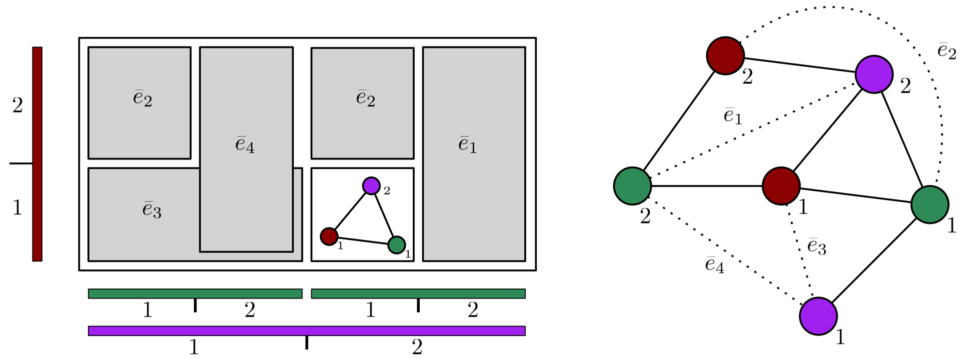

Figure 2 An example of an encoding of a 3-partite graph to a 2-D box (skewed cube). Normal lines represent edges, dotted lines represent a subset of non-edges. The gray boxes correspond to the forbidden regions of the respective non-edge. Note that in our reductions, we aim to cover the complement of these regions, the feasible regions.

hypercube for which the corresponding k-tuple does not include $\bar{e}$. In contrast, the cells for which the corresponding k-tuple includes all vertices of $\bar{e}$ are called the *forbidden region*. For each non-edge, we construct a covering of its feasible region with axis-aligned boxes[6], see Lemma 3.5 for the exact covering. As a k-clique is a k-tuple which does not contain any non-edge, a cell that lies within the feasible regions of all non-edges encodes a k-clique. In Lemma 3.6 we reduce finding a cell inside all feasible regions to the intermediate Translated Box-Shape Problem (Definition 3.2) with our desired values of n, m.

In the second step, the resulting Translated Shape Problem is reduced to the Translation Problem with Boxes, see Lemma 3.3.

Lastly, in our third step we reduce the Translation Problem with Boxes in $\lfloor d/2 \rfloor$-D to the Translation Problem with Orthants in d dimensions with n points and m orthants, see Lemma 3.1. As given in Lemma 2.6, the Translation Problem with Orthants has a linear time reduction to directed Hausdorff under Translation, which yields our lower bound.

We go through the steps in backwards order, starting with the last, third step.

Step 3: From Boxes to Orthants

We reduce the Translation Problem with Boxes to the Translation Problem with Orthants. We implement the natural idea to cover boxes of arbitrary aspect ratios by orthants (large hypercubes), at the cost of doubling the underlying dimension. As this basic idea allows additional, unintended solutions, we need a technical translation gadget that ensures soundness. Figure 3 gives a visual intuition of the proof in 1-D.

▶ **Lemma 3.1.** *The Translation Problem with Boxes in d-D with input parameters n, m reduces in linear time to the Translation Problem with Orthants in $2d$-D with input parameters $n' = \Theta(n), m' = \Theta(m)$.*

Proof Sketch. We define the coordinate-wise transformations

$$c : p \mapsto (p, p), \qquad c^+ : [a, b] \mapsto [a, \infty) \times (-\infty, b], \qquad c^- : [a, b] \mapsto (-\infty, b] \times [a, \infty).$$

[6] As we only consider axis-aligned boxes in this paper, we may only call these boxes.

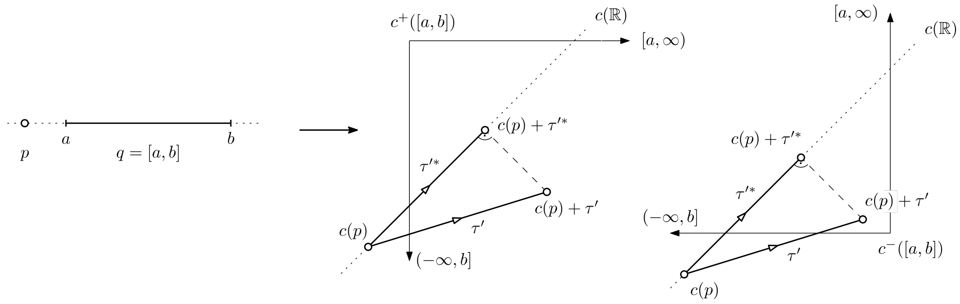

Figure 3 The reduction of a box in 1-D, i.e., an interval $[a, b]$, to orthants in 2-D. We prove that for each feasible translation τ' there exists a feasible translation τ'^* which lies on the diagonal $c(\mathbb{R})$. For that, we always assume that the translation τ' lies between the diagonal $c(\mathbb{R})$ and the vertex of the orthant, i.e., for this example the orthant of c^- is relevant.

For points in $\mathbb{R}^d$, we define the transformation f as the coordinate-wise application of c. For our analysis, we further apply f to sets or intervals in which case we apply f element-wise. For boxes in $\mathbb{R}^d$, we define the transformations $f_1, ..., f_{2^d}$ as all possible combinations of the coordinate-wise application of c^+, c^-. Note that the images of f_i are orthants.

Given an instance of the Translation Problem with Boxes with P being the point set and $\mathcal{B}$ being the set of boxes. We construct the sub-instances $(f(P), f_1(\mathcal{B})), ..., (f(P), f_{2^d}(\mathcal{B}))$ of the resulting Translation Problem with Orthants instance, the target box is simply the bounding box of $\mathcal{B}$. We claim that

$$\begin{aligned} &\exists \tau \in \mathbb{R}^d \ \forall p \in P \ \exists B \in \mathcal{B} : p + \tau \in B \\ \iff &\exists \tau' \in \mathbb{R}^{2d} \ \forall i \in [2^d] \ \forall p' \in f(P) \ \exists C \in f_i(\mathcal{B}) : p' + \tau' \in C. \end{aligned}$$

For the proof, we see that a feasible translation $f(\tau) \in f(\mathbb{R}^d)$ for the right side for *any* $i \in [2^d]$ is a feasible translation for the left side, so we prove that a feasible translation $\tau' \notin f(\mathbb{R}^d)$ can be transformed to a feasible translation $\tau'^* \in f(\mathbb{R})$. Choosing the $i \in [2^d]$ such that τ' always lies on the correct side of $f(\mathbb{R}^d)$, where all vertices of our constructed orthants $f_i(\mathcal{B})$ lie, enables us to move τ' to its closest point τ'^* on $f(\mathbb{R}^d)$, giving $p' + \tau' \in C \implies p' + \tau'^* \in C$ for all $p' \in f(P), C \in f_I(\mathcal{B})$. The full proof details can be found in the full version. ◀

Step 2: The intermediate Translated Box-Shape Problem

We introduce an intermediate problem, the Translated Box-Shape Problem, similar to the one of [9, Problem 4] using hypercubes. In contrast to the Translation Problem with Boxes, it allows composing boxes into shapes for which translated copies can be instantiated with a single translation object. Thus, a set of n^2 boxes, structured appropriately, can be described by a shape of n boxes and n translation objects.

▶ **Definition 3.2.** (Translated (Box-)Shape Problem). *Let $\mathcal{Z}$ be a set of shapes in $\mathbb{R}^d$, composed of a union of boxes. Let m be the total number of boxes over all shapes of $\mathcal{Z}$. Given a set $\mathcal{S}$ of n objects where each object is the translate of some shape in $\mathcal{Z}$, decide whether $\bigcap_{S \in \mathcal{S}} S \neq \emptyset$.*

We further give a reduction from the Translated Shape Problem to the Translation Problem with Boxes, analog to [9, Lemma 3]. Its proof again is deferred to the full version.

▶ **Lemma 3.3.** *The Translated (Box-)Shape Problem with n translation objects and a total of m boxes reduces to the Translation Problem with Boxes with $|P| = \Theta(n)$ and $|\mathcal{Q}| = \Theta(m)$.*

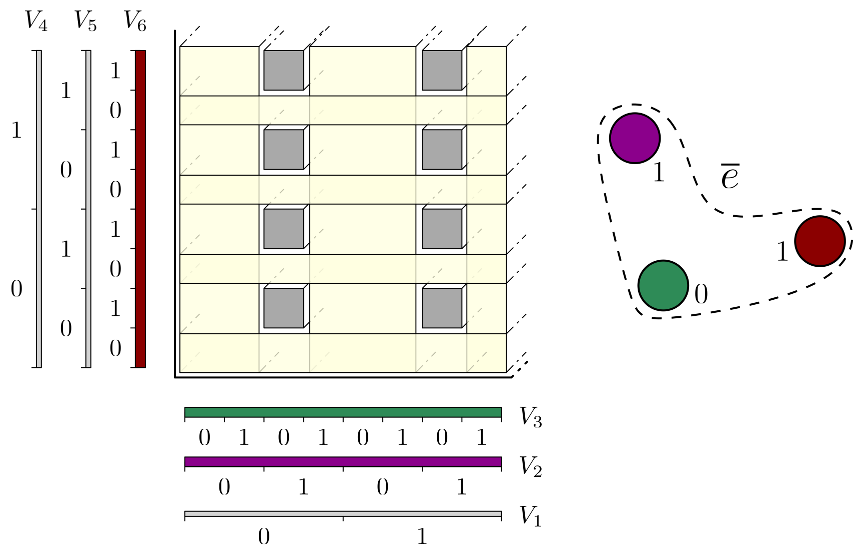

Figure 4 An example of the feasible region of a non-edge $\bar{e} \in V_2 \times V_3 \times V_6 \setminus E$ covered by boxes, depicted in yellow, within a d-dimensional hypercube projected to 2 dimensions.

Lastly, we give a simple observation allowing, by repeated application, to compose arbitrary many individual Translated (Box-)Shape Problems without additional overhead.

▶ **Observation 3.4.** *Given two instances of the Translated Shape Problem with inputs $\mathcal{S}_1, \mathcal{Z}_1$ and $\mathcal{S}_2, \mathcal{Z}_2$ of size n_1, m_1 and n_2, m_2 in a shared universe. Deciding whether $\bigcap_{S \in \mathcal{S}_1 \cup \mathcal{S}_2} S \neq \emptyset$ reduces to a single instance of the Translated Shape Problem with $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$, $\mathcal{Z} = \mathcal{Z}_1 \cup \mathcal{Z}_2$ of size $n \leq n_1 + n_2$ and $m \leq m_1 + m_2$.*

Step 1: From k-Hyperclique to the Translated Box-Shape Problem

In the remaining step, along the basic idea presented in the beginning of Section 3.1, we encode the k-hyperclique problem to a d-dimensional hypercube and reduce finding a cell that lies in the feasible regions of all non-edges to the Translated (Box-)Shape Problem with the desired distribution of parameters n, m. As our approach will also provide novel combinatorial lower bounds, we more generally consider a u-uniform hypergraph with $u \in \{2, 3\}$.

We first give the encoding of the hypergraph to a d-dimensional hypercube and show that the feasible regions of the non-edges can be covered efficiently. The same encoding for a graph can exemplarily be seen in Figure 2, note that for illustrative reasons we there have not covered the feasible but the forbidden regions.

▶ **Lemma 3.5.** *Let $d \in \mathbb{N}$ and $u \in \{2, 3\}$ be constant and let $k \in \mathbb{N}$ be constant with $d|k$. Given a u-uniform hypergraph $H = (V, E)$ with $|V| = n$, where V is partitioned into k color classes $V = V_1 \cup V_2 \cup ... \cup V_k$. There is an encoding of $V_1 \times ... \times V_k$ to a d-dimensional hypercube so that for each $\bar{e} \in \binom{V}{u} \setminus E$ the feasible region of $\bar{e}$ can be covered by $O(n^{k/d-1})$ boxes, which can be computed in $O(n^{k/d-1})$ time.*

Proof Sketch. By scaling, we assume that the d-dimensional hypercube has a side-length of $n^{k/d}$. Given the numbers $x_1, ..., x_{k/d} \in [n]$, we use the canonical function

$$\text{ind} : (x_1, ..., x_{k/d}) \mapsto x_1 \cdot n^{k/d-1} + x_2 \cdot n^{k/d-2} + ... + x_{k/d}$$

to address an integer point in the range $[0, n^{k/d} - 1]$ We denote each unit hypercube as a cell, its relative origin can be addressed by $(\text{ind}(X_1), ..., \text{ind}(X_d))$ where $X_1, ..., X_d \in [n]^{k/d}$.

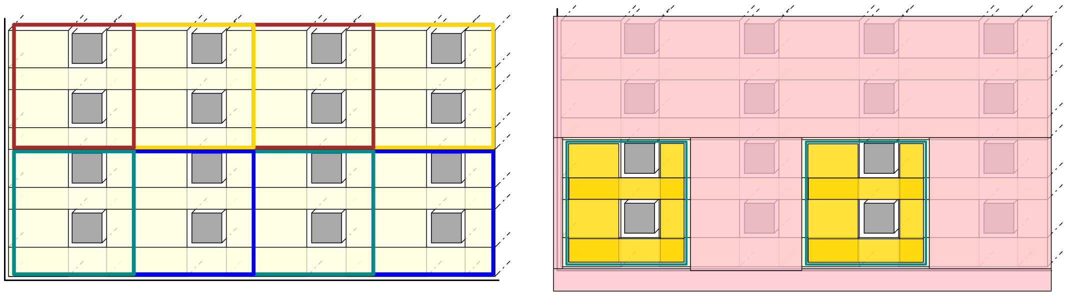

Figure 5 Left: The decomposition of a (skewed) hypercube in 2-D into four identical sub-regions, marked in red, yellow, green, and blue. Boxes in equal color denote a common sub-region, here in each row there are two sub-regions that are split into two parts each.
Right: The boxes used to cover the feasible region of the green sub-regions, in yellow, as well as the boxes covering the complement of the green sub-region, marked in pink.

The encoding of H to the hypercube is the same as used in [9]. We evenly assign the k color-classes to the d dimensions, i.e., the color classes $V_1, ..., V_{k/d}$ are assigned to the first dimension, the classes $V_{k/d+1}, ..., V_{2k/d}$ to the second dimension, etc. A k-tuple $(v_1, ..., v_k) \in V_1 \times ... \times V_k$ is bijectively assigned to the cell

$$\text{cell}((v_1, ..., v_k)) = (\text{ind}(v_1, ..., v_{k/d}), \text{ind}(v_{k/d+1}, ..., v_{2k/d}), ..., \text{ind}(v_{(d-1)k/d}, ..., v_k)).$$

For any non-edge $\bar{e} \in V^u \setminus E$, we cover the feasible region of $\bar{e}$ using at most $O(dn^{k/d-1})$ boxes $\mathcal{B}$, as can be seen in Figure 4. A full proof can be found at in the full version. ◀

Using Lemma 3.5, we can already turn a k-(hyper-)clique instance into a Translated (Box-)Shape instance with $n = O(|V|^u), m = O(|V|^{k/d-1+u})$ as we can treat each feasible region of the $O(|V^u|)$ non-edges as an independent shape with $O(|V|^{k/d-1})$ boxes each. Observation 3.4 allows to compose the resulting independent instances.

To extend the reduction to arbitrary parameter choices of n, m, independently for each non-edge we decompose the hypercube into n smaller identical parts such that in each part the feasible region of a non-edge can be covered by m boxes as shown in Lemma 3.5. This allows a reduction to the Translated Shape Problem with m boxes and n translation objects.

▶ **Lemma 3.6.** *Let $\lambda \in [0, 1]$, $u \in \{2, 3\}$ be constant, and $d, k \in \mathbb{N}$ be constants with $d|k$. Given a u-regular hypergraph $H = (V, E)$ and its encoding to a d-dimensional hypercube as constructed in Lemma 3.5. We reduce the problem of finding a k-(hyper-)clique in H in $O(|V|^{k/d})$ time to an instance of the Translated (Box-)Shape Problem with $n = O(|V|^{u\lambda k/d+u})$ translates of common shapes consisting of $m = O(|V|^{(1-\lambda)k/d+u})$ boxes.*

Proof Sketch. By Observation 3.4, it suffices to consider the feasible regions of a single non-edge $\bar{e}$ and reduce finding a cell that lies in its feasible regions to a Translated (Box-)Shape instance with $n' = O(|V|^{u\lambda k/d})$ translates of a single shape with $m' = O(|V|^{(1-\lambda)k/d})$ boxes.

We consider the encoding of the hypergraph to the d-dimensional hypercube as given in Lemma 3.5. As $\bar{e}$ is a u-tuple, its included vertices are only encoded in up to u dimensions, it thus suffices to consider the hypercube projected to these u dimensions.

We construct $((|V|^{k/d})^\lambda)^u = n'$ many sub-regions that decompose the u-dimensional hypercube: We decompose each single dimension into $(|V|^{k/d})^\lambda$ many (not necessarily continuous) slices that are identical under translation, i.e., when translated so that two slices have the same relative origin, the feasible and forbidden regions of $\bar{e}$ within these slices match. These one-dimensional slices can be combined independently in all dimensions, yielding the desired decomposition of the full u-dimensional hypercube into identical sub-regions.

The feasible region of $\bar{e}$ in each sub-region can be covered by m boxes. As the sub-regions are identical, we can construct a Translated (Box-)Shape instance with n' translation objects and a single shape with m boxes. A full proof can be found in the full version. ◀

Resulting Lower Bounds

Combining our results from Lemmas 3.1, 3.3, and 3.6, we get lower bounds for arbitrary distributions of n, m. However, as the bounds deteriorate with growing values of n, we only get tight lower bounds for $n = m^{o(1)}$. We consider other cases in the full theorem statement which can be found with its proof in the full version.

▶ **Theorem 3.7.** *Let $\varepsilon > 0, \lambda \in (0, 1], d \in \mathbb{N}_{\geq 4}$ be constant. Under the k-hyperclique hypothesis, there is no $O((nm)^{\lfloor \frac{d}{2} \rfloor \frac{\lambda+3}{3\lambda+3} - \varepsilon})$ algorithm for directed HuT in d-D with $n = \Theta(m^\lambda)$.*

3.2 Tight Non-Combinatorial Lower Bound in 3-D

Chan [9] has proven a tight lower bound that applies for distributions $m \in O(n)$ in 3-D. We give a non-combinatorial version of this $(nm)^{d/2-o(1)}$ lower bound for $m \in O(n^{1/2})$.

We reduce from the k-Hyperclique problem, combining Chan's ideas with the work of Künnemann [19] using *prefix covering designs*. As done in Section 3.1, we encode a k-tuple of vertices inside a d-dimensional hypercube. However, we use a redundant encoding based on the prefix covering designs, that is, each color class is encoded more than once to our hypercube. The full proof is presented in the full version.

▶ **Corollary 3.8.** *Let $\varepsilon > 0$. For any $\lambda \in (0, 1/2]$, there is no $O((nm)^{3/2-\varepsilon})$ (non-combinatorial) algorithm for directed Hausdorff under Translation in 3-D with $m = \Theta(n^\lambda)$ under the k-hyperclique hypothesis.*

3.3 Connections to Additive Problems

We draw several connections between the Hausdorff distance under Translation and several additive problems. Below we highlight a conditional lower bound in 1D and conditional upper bound in 3D. Omitted proofs and additional results are given in the full version.

We first show that even for one-dimensional instances, directed HuT inherits hardness from quadratic-time barriers from problems such as *MaxConv LowerBound*.

▶ **Definition 3.9** (MaxConv LowerBound, MaxConvLB)**.** *Given integer arrays A, B, C of length[7] n. Determine whether $C[k] \leq \max_{i+j=k}(A[i] + B[j])$ holds for all $k \in \{2, ..., n\}$.*

In an investigation of hardness results from $(\min, +)$-convolution, Cygan, Mucha, Wegrzycki and Wlodarczyk [12] studied the closely related MaxConv LowerBound problem and gave a subquadratic reduction to L_∞-Necklace Alignment (in turn, L_∞-Necklace Alignment admits a subquadratic reduction to $(\min, +)$-Convolution [6]). Cygan et al. raised the natural question whether MaxConv LowerBound is subquadratic equivalent to $(\min, +)$-convolution; to the best of our knowledge, neither a positive or negative resolution is known.

For the (continuous) directed HuT we get hardness by a reduction from the Linear Alignment problem, shown in the full version. Further, using a reduction from L_∞-Necklace Alignment to Linear Alignment, we get a reduction from MaxConv LowerBound to directed

[7] Other common definitions use for C an array of size $2n$. For convenience, we stick to this definition.

HuT. Interestingly, we obtain such a lower bound even for discrete directed HuT by directly reducing from MaxConv LowerBound. These reductions illuminate why no subquadratic algorithms are known for directed variants, despite near-linear algorithms for their undirected counterparts [20].

▶ **Theorem 3.10.** *If 1-dimensional discrete HuT can be solved in time $T(n)$, then the MaxConv LowerBound problem can be solved in time $O(T(n))$.*

Proof sketch. Given arrays A, B, C of positive values as inputs from MaxConv LowerBound. We construct T, P, Q as arrays for DiscHuT. By abuse of notation, we allow Q to be indexed negatively. We define $T = [-(kM + C[k])]_{k\in[n]}$, $P = [iM + A[i] + M/4]_{i\in[n]}$, $Q = [-zM + M/4]_{z\in\{-n,\dots,0\}} \cup [-(jM + B[j])]_{j\in[n]}$, $\delta = \frac{M}{4} - 1$, where M is a large number with $M > 4 \cdot \max_{i,j,k} (A[i] + B[j] + C[k])$.

Now, if A, B, C is a NO-instance to MaxConv LowerBound, then there is a k^* such that for all $i < k^*$ we have $A[i]+B[k^*-i] < C[k^*]$. From this we can derive that $|p_i+\tau_{k^*}-q_{k^*-i}| \leq \delta$, where p_i, q_{k^*-i}, and τ_{k^*} denote the points in P, Q, and T, respectively, corresponding to $A[i]$, $B[k^*-i]$ and $C[k^*]$. Thus, we have a YES-instance to DiscHuT. Analogously, a YES-instance to MaxConv results in a NO-instance to DiscHuT. See the full version for the full proof. ◀

We give an upper bound to discrete HuT, which stems from a reduction from 3-dimensional Discrete HuT to 3-SUM. On a very high level, the basic idea consists of a complement trick, which enables a quantifier switch in the discrete HuT problem.

▶ **Corollary 3.11.** *Let $d \leq 3$. If 3-SUM can be solved in time $T(n)$, then we can solve d-dimensional balanced Discrete HuT in time $\tilde{O}(T(n\log(n)) + n\log(n))$.*

3.4 A Faster Algorithm for $m \gg n$

For $n = m^{o(1)}$, Theorem 3.7 gives a $(nm)^{\lfloor d/2\rfloor - o(1)}$ lower bound, which for even d matches the best known algorithm running in time $O((nm)^{d/2})$ [9], while for odd d it seems to have some slack. Intuitively, one might expect that the known $(nm)^{d/2-o(1)}$ combinatorial lower bound for $m = O(n)$ [9] also transfers to the case of $m \gg n$. Surprisingly, in 3-D the $(nm)^{1-o(1)}$ lower bound of [7] is tight for $n = m^{o(1)}$ as we give an $O(m^{1+o(1)})$ algorithm.

The algorithm uses a structural insight from [10, Theorem 7], which we turn algorithmic. The full description and proof can be found in the full version. Our algorithm relies on an algorithm recently introduced by Agarwal and Steiger [2], which in 3-D computes the vertices of the union of m unit cubes in an output-sensitive runtime of $\tilde{O}(m)$. Assuming an algorithm that computes the vertices of the union of m unit hypercubes in d-D in time $\tilde{O}(m^{\lfloor d/2\rfloor})$, our algorithm directly generalizes to higher dimensions with a runtime of $\tilde{O}(m^{\lfloor d/2\rfloor} n^{d+1})$.

▶ **Theorem 3.12.** *Directed Hausdorff under Translation in 3-D can be solved in time $\tilde{O}(n^4 m)$.*

Its existence forbids $(nm)^{d/2-o(1)}$ lower bounds for arbitrary distributions of n, m for $d = 3$, and we give compelling intuition for a similar result in odd $d > 3$.

▶ **Corollary 3.13.** *Let $d = 3$, $\kappa > 0$ and set $\lambda = \frac{\kappa}{d+2}$. For $n = O(m^\lambda)$, directed Hausdorff under Translation in d-D can be solved in time $O(m^{\lfloor d/2\rfloor} n^{d+1}) < O(m^{\lfloor d/2\rfloor+\kappa})$.*

This result points to a discrepancy in complexity, at least for $d = 3$ but potentially also for all odd d. While we have a $\tilde{O}(m)$ algorithm for $n = m^{o(1)}$, there is a $(nm)^{d/2-o(1)}$ tight (combinatorial) lower bound for $m \in O(n)$ [9]. It remains unknown where the cutoff point is located, i.e., for which values of n, m do we start to get faster than $(nm)^{d/2\pm o(1)}$ algorithms?

3.5 Transferring Lower Bounds to the Undirected Setting

As of now, all here presented lower bounds only apply to the directed Hausdorff under Translation setting. Chan [9] already raised the question if these bounds transfer to the undirected setting, however there are no such lower bounds known so far for $d \geq 3$. Only for $d = 2$, lower bounds for the undirected setting are known [7], which match the bounds for the directed variant. For $d = 1$, we have an optimal $\tilde{O}(n + m)$ algorithm [20].

We partially resolve this question by providing a reduction from the Translation Problem with Orthants to undirected Hausdorff under Translation, maintaining all instance parameters.

▶ **Theorem 3.14.** *If we can decide the undirected Hausdorff under Translation problem in d-D with point sets of size n, m in time $T(n, m)$, we can decide the Translation Problem with Orthants in d-D with n points and m orthants in time $T(n, m) + O(n + m)$.*

Proof Sketch. We only give intuition on the reduction to a uHuT instance from a single TPwC sub-instance of our TPwO input. The full proof is given in the full version.

For an individual TPwC instance, given the point set P and the set of hypercubes $\mathcal{C}$ of side length δ. Note that the target box B of our TPwO instance implies that the translated point set $P + T$ and the set of feasible translations T is restricted, we may assume that $T \subseteq [-\delta/3, \delta/3]^d$ and $P + T \subseteq B \subseteq b + [-2\delta/3, 2\delta/3]^d$, where $b \in \mathbb{R}^d$ is some relative origin.

We strategically place additional points in the TPwC sub-instance to allow a reduction to uHuT. These points are located on the vertices and middle points of all faces of a $b + [-\delta, \delta]^d$ extended target box. As the translation and translated point set are restricted, any hypercube that intersects the target box $b + [-2\delta/3, 2\delta/3]^d$ also intersects the extended target box $b + [-\delta, \delta]^d$ when translated. Thus, we can place auxiliary points such that under all translations each hypercube C of our TPwC instance has an auxiliary point which lies inside C. Having this property, the reduction to uHuT is imminent. ◀

All lower bounds for HuT in $d \geq 3$ dimensions presented so far, as well as the combinatorial $(nm)^{d/2-o(1)}$ lower bound for $m \in O(n)$ of Chan [9], can be directly applied to undirected Hausdorff under Translation as they use the Translation Problem with Orthants as an intermediate problem. At the cost of doubling the dimension using Lemma 3.1, we generally reduce any instance of directed Hausdorff under Translation to the undirected setting.

▶ **Corollary 3.15** (Excerpt from full version)**.** *We get the following results:*

- *For arbitrary n, m, there is no $O((nm)^{d/2-\varepsilon})$ combinatorial algorithm for undirected Hausdorff under Translation in d-D under the k-clique hypothesis.*
- *For $\min(n, m) \in O(\max(n, m)^{1/2})$, there is no $O((nm)^{3/2-\varepsilon})$ algorithm for undirected Hausdorff under Translation in 3-D under the k-hyperclique hypothesis.*
- *If balanced (discrete) undirected Hausdorff under Translation in 2-D can be solved in time $T_2(n)$, then MaxConv LowerBound can be solved in time $O(T_2(n) + n)$.*

References

1 P. K. Agarwal, M. Sharir, and S. Toledo. Applications of parametric searching in geometric optimization. *J. Algorithms*, 17(3):292–318, 1994. doi:10.1006/jagm.1994.1038.

2 Pankaj K. Agarwal and Alex Steiger. An output-sensitive algorithm for computing the union of cubes and fat boxes in 3D. *proc48th*, 198:10:1–10:20, 2021. doi:10.4230/LIPIcs.ICALP.2021.10.

3 Helmut Alt and Leonidas J. Guibas. Discrete geometric shapes: Matching, interpolation, and approximation. In Jörg-Rüdiger Sack and Jorge Urrutia, editors, *Handbook of Computational Geometry*, pages 121–153. North Holland / Elsevier, 2000. doi:10.1016/B978-044482537-7/50004-8.

4 Helmut Alt, Kurt Mehlhorn, Hubert Wagener, and Emo Welzl. Congruence, similarity, and symmetries of geometric objects. *Discret. Comput. Geom.*, 3:237–256, 1988. doi:10.1007/BF02187910.

5 Gill Barequet and Sariel Har-Peled. Polygon Containment and Translation in Hausdorff-Distance Between Segment Sets are 3SUM-Hard. *International Journal of Computational Geometry & Applications*, 2001. doi:10.1142/S0218195901000596.

6 David Bremner, Timothy M. Chan, Erik D. Demaine, Jeff Erickson, Ferran Hurtado, John Iacono, Stefan Langerman, Mihai Pǎtraşcu, and Perouz Taslakian. Necklaces, convolutions, and X+Y. *Algorithmica*, 69(2):294–314, 2014. doi:10.1007/S00453-012-9734-3.

7 Karl Bringmann and André Nusser. Translating Hausdorff is hard: fine-grained lower bounds for Hausdorff distance under translation. *J. Comput. Geom.*, 13(2):30–50, 2021. doi:10.20382/JOCG.V13I2A3.

8 Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *7th*, pages 261–270, 2016. doi:10.1145/2840728.2840746.

9 Timothy M. Chan. Minimum L_∞ Hausdorff distance of point sets under translation: Generalizing Klee's measure problem. In *proc39th*, volume 258, pages 24:1–24:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.SOCG.2023.24.

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

11 L. Paul Chew and Klara Kedem. Getting around a lower bound for the minimum Hausdorff distance. *Computational Geometry*, 10(3):197–202, June 1998. doi:10.1016/S0925-7721(97)00032-1.

12 Marek Cygan, Marcin Mucha, Karol Wegrzycki, and Michal Wlodarczyk. On problems equivalent to (min,+)-convolution. *ACM Transactions on Algorithms*, 15(1):14:1–14:25, 2019. doi:10.1145/3293465.

13 Nick Fischer. Deterministic sparse pattern matching via the Baur-Strassen theorem. In *proc35th*, pages 3333–3353. SIAM, 2024. doi:10.1137/1.9781611977912.119.

14 Geri Gokaj and Marvin Künnemann. Completeness theorems for k-SUM and geometric friends: Deciding fragments of linear integer arithmetic. In *proc16th*, volume 325 of *LIPIcs*, pages 55:1–55:25, 2025. ISSN: 1868-8969. doi:10.4230/LIPIcs.ITCS.2025.55.

15 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. doi:10.1109/34.761267.

16 Daniel P. Huttenlocher, Klara Kedem, and Micha Sharir. The upper envelope of Voronoi surfaces and its applications. *Discrete & Computational Geometry*, 9:267–291, 1993. doi:10.1007/BF02189323.

17 Daniel P. Huttenlocher, Gregory A. Klanderman, and William Rucklidge. Comparing images using the Hausdorff distance. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15(9):850–863, 1993. doi:10.1109/34.232073.

18 Sanjiv Kapoor and Michiel Smid. New techniques for exact and approximate dynamic closest-point problems. *SIAM Journal on Computing*, 25(4):775, 1996.

19 Marvin Künnemann. A tight (non-combinatorial) conditional lower bound for Klee's measure problem in 3D. In *proc63rd*, pages 555–566, 2022. doi:10.1109/FOCS54457.2022.00059.

20 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.

21 William Rucklidge. Lower bounds for the complexity of the graph of the hausdorff distance as a function of transformation. *Discret. Comput. Geom.*, 16(2):135–153, 1996. doi:10.1007/BF02716804.

22 Carola Wenk. *Shape Matching in Higher Dimensions*. PhD thesis, Freie Universität Berlin, 2003. doi:10.17169/refubium-8310.