# Computational Complexity of the $\alpha$-Ham-Sandwich Problem


## Man-Kwun Chiu

Institut für Informatik, Freie Universität Berlin, Takustraße 9, 14195 Berlin, Germany
chiumk@inf.fu-berlin.de

## Aruni Choudhary

Institut für Informatik, Freie Universität Berlin, Takustraße 9, 14195 Berlin, Germany
arunich@inf.fu-berlin.de

## Wolfgang Mulzer

Institut für Informatik, Freie Universität Berlin, Takustraße 9, 14195 Berlin, Germany
mulzer@inf.fu-berlin.de



## Abstract

The classic Ham-Sandwich theorem states that for any $d$ measurable sets in $\mathbb{R}^d$, there is a hyperplane that bisects them simultaneously. An extension by Bárány, Hubard, and Jerónimo [DCG 2008] states that if the sets are convex and *well-separated*, then for any given $\alpha_1, \dots, \alpha_d \in [0, 1]$, there is a unique oriented hyperplane that cuts off a respective fraction $\alpha_1, \dots, \alpha_d$ from each set. Steiger and Zhao [DCG 2010] proved a discrete analogue of this theorem, which we call the *$\alpha$-Ham-Sandwich theorem*. They gave an algorithm to find the hyperplane in time $O(n(\log n)^{d-3})$, where $n$ is the total number of input points. The computational complexity of this search problem in high dimensions is open, quite unlike the complexity of the Ham-Sandwich problem, which is now known to be PPA-complete (Filos-Ratsikas and Goldberg [STOC 2019]).

Recently, Fearnley, Gordon, Mehta, and Savani [ICALP 2019] introduced a new sub-class of CLS (Continuous Local Search) called *Unique End-of-Potential Line* (UEOPL). This class captures problems in CLS that have unique solutions. We show that for the $\alpha$-Ham-Sandwich theorem, the search problem of finding the dividing hyperplane lies in UEOPL. This gives the first non-trivial containment of the problem in a complexity class and places it in the company of classic search problems such as finding the fixed point of a contraction map, the unique sink orientation problem and the $P$-matrix linear complementarity problem.



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

**Keywords and phrases** Ham-Sandwich Theorem, Computational Complexity, Continuous Local Search

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2020.31

**Category** Track A: Algorithms, Complexity and Games

**Related Version** A full version of the paper is available at https://arxiv.org/abs/2003.09266

**Funding** Supported in part by ERC StG 757609.


## 1 Introduction

The Ham-Sandwich Theorem [41] is a classic result about partitioning sets in high dimensions: for any $d$ measurable sets $S_1, \dots, S_d \subset \mathbb{R}^d$ in $d$ dimensions, there is an oriented hyperplane $H$ that simultaneously *bisects* $S_1, \dots, S_d$. More precisely, if $H^+, H^-$ are the closed half-spaces bounded by $H$, then for $i = 1, \dots, d$, the measure of $S_i \cap H^+$ equals the measure of $S_i \cap H^-$. The traditional proof goes through the Borsuk-Ulam Theorem [30]. The Ham-Sandwich

Theorem is a cornerstone of geometry and topology, and it has found applications in other areas of mathematics.

Let $[n] = \{1, \ldots, n\}$. The *discrete* Ham-Sandwich Theorem [28, 30] states that for any $d$ finite point sets $P_1, \ldots, P_d \subset \mathbb{R}^d$ in $d$ dimensions, there is an oriented hyperplane $H$ such that $H$ bisects each $P_i$, i.e., for $i \in [d]$, we have $\min\{|P_i \cap H^+|, |P_i \cap H^-|\} \geq \lceil |P_i|/2 \rceil$. We denote the associated search problem as Ham-Sandwich. Lo, Matoušek, and Steiger [28] gave an $n^{O(d)}$-time algorithm for Ham-Sandwich. They also provided a linear-time algorithm for points in $\mathbb{R}^3$, under additional constraints.

There are many alternative and more general variants of both the continuous and the discrete Ham-Sandwich Theorem. For example, Bárány and Matoušek [5] derived a version where measures in the plane can be divided into any (possibly different) ratios by *fans* instead of hyperplanes (lines). A discrete variant of this result was given by Bereg [7]. Schnider [37] and Karasev [27] studied generalizations in higher dimensions. Recently Barba, Pilz, and Schnider [6] showed that four measures in the plane can be bisected with two lines. Higher dimensional generalizations of this result were presented in [9, 25]. Zivaljević and Vrećica [44] and independently, Dol'nikov [19] proved a result called the Center Transversal Theorem that interpolates between the Ham-Sandwich Theorem and the Centerpoint Theorem [35]. There is also a no-dimensional version [14] for the Center Transversal Theorem. Schnider [38] presented a generalization based on this result among others.

Here, we focus on a version that allows for dividing the sets into arbitrary given ratios instead of simply bisecting them. The sets $S_1, \ldots, S_d \subset \mathbb{R}^d$ are *well-separated* if every selection of them can be strictly separated from the others by a hyperplane. Bárány, Hubard, and Jerónimo [4] showed that if $S_1, \ldots, S_d$ are well-separated and convex, then for any given reals $\alpha_1, \ldots, \alpha_d \in [0, 1]$, there is a unique hyperplane that divides $S_1, \ldots, S_d$ in the ratios $\alpha_1, \ldots, \alpha_d$, respectively. Their proof goes through Brouwer's Fixed Point Theorem. Steiger and Zhao [40] formulated a discrete version. In this setup, $S_1, \ldots, S_d$ are finite point sets. Again, we need that the (convex hulls of the) $S_i$ are well-separated. Additionally, we require that the $S_i$ follow a weak version of general position. Let $\alpha_1, \ldots, \alpha_d \in \mathbb{N}$ be $d$ integers with $1 \leq \alpha_i \leq |S_i|$, for $i \in [d]$. Then, there is a unique oriented hyperplane $H$ that passes through one point from each $S_i$ and has $|H^+ \cap S_i| = \alpha_i$, for $i \in [d]$ [40]. In other words, $H$ simultaneously cuts off $\alpha_i$ points from $S_i$, for $i \in [d]$. This statement does not necessarily hold if the sets are not well-separated, see Figure 1 for an example.

Steiger and Zhao called their result the *Generalized Ham-Sandwich Theorem*, yet it is not a strict generalization of the classic Ham-Sandwich Theorem. Their result requires that the point sets obey well-separation and weak general position, while the classic theorem always holds without these assumptions. Therefore, we call this result the *$\alpha$-Ham-Sandwich theorem*, for a clearer distinction. Set $n = \sum_{i \in [d]} |S_i|$. Steiger and Zhao gave an algorithm that computes the dividing hyperplane in $O\left(n(\log n)^{d-3}\right)$ time, which is exponential in $d$. Later, Bereg [8] improved this algorithm to achieve a running time of $n2^{O(d)}$, which is linear in $n$ but still exponential in $d$. We denote the associated computational search problem of finding the dividing hyperplane as Alpha-HS.

No polynomial algorithms are known for Ham-Sandwich and for Alpha-HS if the dimension is not fixed, and the notion of approximation is also not well-explored. Despite their superficial similarity, it is not immediately apparent whether the two problems are comparable in terms of their complexity. Due to the additional requirements on an input for Alpha-HS, an instance of Ham-Sandwich may not be reducible to Alpha-HS in general.

A dividing hyperplane for Alpha-HS is guaranteed to exist if the sets satisfy the conditions of well-separation and (weak) general position. Therefore, the search problem Alpha-HS

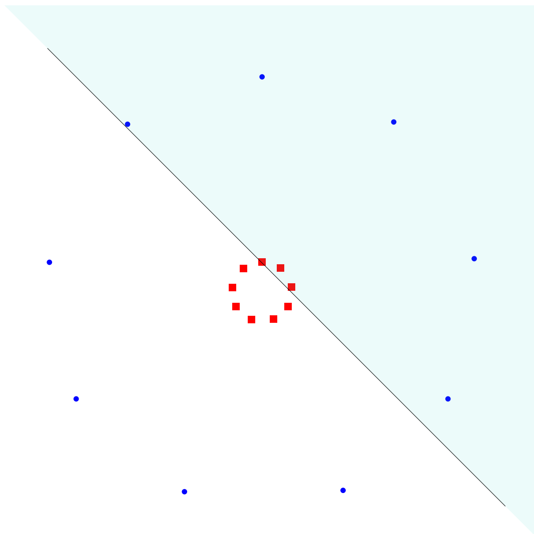

**Figure 1** The red (square) and the blue (round) point sets are not well-separated. Every halfplane that contains three red points must contain at least five blue points. Thus, there is no halfplane that contains exactly three red and three blue points.

is total, that is, there is a solution for every valid instance. In general, such problems are modelled by the complexity class TFNP (Total Function Nondeterministic Polynomial) of NP-search problems that always admit a solution. Two popular subclasses of TFNP, originally defined by Papadimitriou [34], are PPA (Polynomial Parity Argument) and its sub-class PPAD (Polynomial Parity Arguments on Directed graphs). These classes contain total search problems where the existence of a solution is based on a parity argument in an undirected or in a directed graph, respectively. Another sub-class of TFNP is PLS (polynomial local search). It models total search problems where the solutions can be obtained as minima in a local search process, while the number of steps in the local search may be exponential in the input size. The class PLS was introduced by Johnson, Papadimitriou, and Yannakakis [26]. A noteworthy sub-class of PPAD $\cap$ PLS is CLS (continuous local search) [18]. It models similar local search problems over a continuous domain using a continuous potential function.

Up to very recently, these complexity classes had mostly been studied in the context of algorithmic game theory. These classes have also found relevance in the study of fairness [33] and markets [10, 12]. However, there have been increasing efforts towards mapping the complexity landscape of existence theorems in high-dimensional discrete geometry. Computing an approximate solution for the search problem associated with the Borsuk-Ulam Theorem is in PPA. In fact, this problem is complete for this class. The discrete analogue of the Borsuk-Ulam Theorem, Tucker's Lemma [42], is also PPA-complete [1, 34]. Therefore, since the traditional proof of the Ham-Sandwich Theorem goes through the Borsuk-Ulam Theorem, it follows that Ham-Sandwich lies in PPA. In fact, Filos-Ratsikas and Goldberg [21] recently showed that Ham-Sandwich is complete for PPA. The (presumably smaller) class PPAD is associated with fixed-point type problems: computing an approximate Brouwer fixed point is a prototypical complete problem for PPAD. The discrete analogue of Brouwer's Fixed Point Theorem, Sperner's Lemma, is also complete for PPAD [34]. The computational version of the Hairy Ball Theorem has recently been shown to be PPAD-complete [24]. In a celebrated result, the relevance of PPAD for algorithmic game theory was made clear when it turned out that computing a Nash-equilibrium in a three player game is PPAD-complete [17].

Subsequently, this was also shown for the two player game [11]. In discrete geometry, finding a solution to the Colorful Carathéodory problem [3] was shown to lie in the intersection PPAD ∩ PLS [31, 32]. This further implies that finding a *Tverberg* partition (and computing a centerpoint) also lies in the intersection [29, 36, 43]. The problem of computing the (unique) fixed point of a contraction map is known to lie in CLS [18].

Recently, at ICALP 2019, Fearnley, Gordon, Mehta, and Savani defined a sub-class of CLS that represents a family of total search problems with unique solutions [20]. They named the class *Unique End of Potential Line* (UEOPL) and defined it through the canonical complete problem UniqueEOPL. This problem is modelled as a directed graph. There are polynomially-sized Boolean circuits that compute the successor and predecessor of each node, and a potential value that always increases on a directed path. There is supposed to be only a single vertex with no predecessor (*start of line*). Under these conditions, there is a unique path in the graph that ends on a vertex (called *end of line*) with the highest potential along the path. This vertex is the solution to UniqueEOPL. Since the uniqueness of the solution is guaranteed only under certain assumptions, such a formulation is called a *promise* problem. Since there seems to be no efficient way to verify the assumptions, the authors allow two possible outcomes of the search algorithm: either report a correct solution, or provide any solution that was found to be in violation of the assumptions. This formulation turns UniqueEOPL into a *non-promise* problem and places it in TFNP, since a correct solution is bound to exist when there are no violations, and otherwise a violation can be reported as a solution. Fearnley et al. [20] also introduced the concept of a *promise-preserving* reduction between two problems $A$ and $B$, such that if an instance of $A$ has no violations, then the reduced instance of $B$ is also free of violations. This notion is particularly meaningful for non-promise problems.

**Contributions.** We provide the first non-trivial containment in a complexity class for the $\alpha$-Ham-Sandwich problem by locating it in UEOPL. More precisely, we formulate Alpha-HS as a non-promise problem in which we allow for both valid solutions representing the correct dividing hyperplane, as well as violations accounting for the lack of well-separation and/or (weak) general position of the input point sets. A precise formulation of the problem is given in Definition 4 in Section 2. We then show a promise-preserving reduction from Alpha-HS to UniqueEOPL. This implies that Alpha-HS lies in UEOPL, and hence in CLS ⊆ PPAD ∩ PLS. See Figure 2 for a pictorial description.

It is not surprising to discover that Alpha-HS lies in PPAD, since the proof of the continuous version in [4] was based on Brouwer's Fixed Point Theorem. The observation that it also lies in PLS is new and noteworthy, putting Alpha-HS into the reach of local search algorithms. In contrast, given our current understanding of total search problems, it is unlikely that the problem Ham-Sandwich would be in PLS.

Since Alpha-HS lies in PPAD ⊆ PPA, it is computationally easier than Ham-Sandwich, which is PPA-complete. This implies the existence of a polynomial-time reduction from Alpha-HS to Ham-Sandwich. A reduction in the other direction is unlikely. It thus turns out that well-separation brings down the complexity of the problem significantly.

Often, problems in TFNP come in the guise of a polynomial-size Boolean circuit with some property. In contrast, Alpha-HS is a purely geometric problem that has no circuit in its problem definition. Apart from the $P$-Matrix Linear complementarity problem, this is one of the few problems in UEOPL and hence in CLS that do not have a description in terms of circuits.

Our local-search formulation is based on the intuition of rotating a hyperplane until we

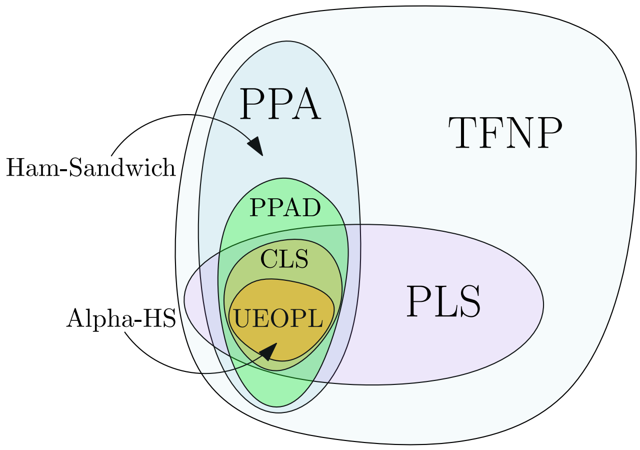


**Figure 2** The hierarchy of complexity classes.

reach the desired solution. We essentially start with a hyperplane that is tangent to the convex hull of each input set, and we deterministically rotate the hyperplane until it hits a new point. This rotation can be continued whenever the hyperplane hits a new point, until we reach the correct dividing hyperplane. In other words, we can follow a local-search argument to find the solution. We show that this sequence of rotations can be modelled as a canonical path in a grid graph, and we give a potential function that guides the rotation and always increases along this path. Every violation of well-separation and (weak) general position can destroy this path. Furthermore, no efficient methods to verify these two assumptions are known. This poses a major challenge in handling the violations. One of our main technical contributions is to handle the violation solutions concisely.

An alternative approach would have been to look at the dual space of points where we get an arrangement of hyperplanes. The dividing hyperplane could then be found by looking at the correct level sets of the arrangement. However, this approach has the problem that the orientations of the hyperplanes in the original space and the dual space are not consistent. This complicates the arguments on the level sets, so we found it more convenient to use our notion of rotating hyperplanes. We show that we can maintain a consistent orientation throughout the rotation, and an inconsistent rotation is detected as a violation of the promise.

**Outline of the paper.** We discuss the background about the $\alpha$-Ham-sandwich Theorem and UniqueEOPL in Section 2. In Section 3, we describe our instance of Alpha-HS and give an overview of the reduction and violation-handling. We conclude in Section 4. The technical details of the reduction and some proofs can be found in the full version of the paper in [13].

# 2 Preliminaries

## 2.1 The $\alpha$-Ham-Sandwich problem

For conciseness, we describe the discrete version of $\alpha$-Ham-Sandwich Theorem [40] here. The continuous version [4] follows a similar formulation.

Let $P_1, \dots, P_d \subset \mathbb{R}^d$ be a collection of $d$ finite point sets. Let $n_1, \dots, n_d$ denote the sizes of $P_1, \dots, P_d$, respectively. For each $i \in [d]$ we say that the point set $P_i$ represents a unique color and let $P := P_1 \cup \dots \cup P_d$ denote the union of all the points. A set of points

$\{p_1, \ldots, p_m\}$ is said to be *colorful* if there are no two points $p_i, p_j$ both from the same color. Indeed a colorful point set can have size at most $d$.

**Weak general position.** We say that $P$ has *very weak general position* [40], if for every choice of points $x_1 \in P_1, \ldots, x_d \in P_d$, the affine hull of the set $\{x_1, \ldots, x_d\}$ is a $(d-1)$-flat and does not contain any other point of $P$. This definition is sufficient for the result of Steiger and Zhao, where they simply call it as weak general position. Of course, this definition of weak general position has no restriction on sets $\{x_1, \ldots, x_d\}$ that contain multiple points from the same color. To simplify our proofs we need a slightly stronger form of general position. We discuss how to deal with very weak general position at the end of Section 3. We say that $P$ has *weak general position* if the above restriction also applies to sets having exactly $d-1$ colors. That means, each color may contribute at most one point to the set, except perhaps one color which is allowed to contribute two points. A certificate for checking violations of weak general position is a set of $d+1$ points whose affine hull has dimension at most $d-1$, with at least $d-1$ colors in the set. Testing whether a point set is in general position can be shown to be NP-Hard, using the result in [23]. It is easy to see that when $d = 2$, weak general position is equivalent to general position.

**Well-separation.** The point set $P$ is said to be *well-separated* [4, 40], if for every choice of points $y_1 \in \text{conv}(P_{i_1}), \ldots, y_k \in \text{conv}(P_{i_k})$, where $i_1, \ldots, i_k$ are distinct indices and $1 \leq k \leq d$, the affine hull of $\{y_1, \ldots, y_k\}$ is a $(k-1)$-flat. An equivalent definition is as follows: $P$ is well-separated if and only if for every disjoint pair of index sets $I, J \subset [d]$, there is a hyperplane that separates the set $\{\cup_{i \in I} P_i\}$ from the set $\{\cup_{j \in J} P_j\}$ strictly. Formally:

▶ **Lemma 1.** *Let $y_1, \ldots, y_d$ be a colorful set of points in the corresponding $\text{conv}(P_i)$. The affine hull of $y_1, \ldots, y_d$ has dimension $d-2$ or less if and only if there is a partition of $[d]$ into index sets $I, J$ such that $\text{conv}(\{\cup_{i \in I} P_i\}) \cap \text{conv}(\{\cup_{j \in J} P_j\}) \neq \emptyset$.*

*Given such a colorful set, the partition of $[d]$ can be computed in $\text{poly}(n, d)$ time. Vice-versa, given such a partition, the colorful set can be computed in $\text{poly}(n, d)$ time.*

A certificate for checking violations of well-separation is a colorful set $\{x_1, \ldots, x_d\}$ whose affine hull has dimension at most $d-2$. Another certificate is a partition $I, J \subset [d]$ such that the convex hulls of the indexed sets are not separable. Due to Lemma 1, both certificates are equivalent and either can be converted into the other in polynomial time. To the best of our knowledge, the complexity of testing well-separation is unknown.

Given any set of positive integers $\{\alpha_1, \ldots, \alpha_d\}$ satisfying $1 \leq \alpha_i \leq n_i$, $i \in [d]$, an $(\alpha_1, \ldots, \alpha_d)$-*cut* is an oriented hyperplane $H$ that contains one point from each color and satisfies $|H^+ \cap P_i| = \alpha_i$ for $i \in [d]$, where $H^+$ is the closed positive half-space defined by $H$.

▶ **Theorem 2** ($\alpha$-Ham-Sandwich Theorem [40])**.** *Let $P_1, \ldots, P_d$ be finite, well-separated point sets in $\mathbb{R}^d$. Let $\alpha = (\alpha_1, \ldots, \alpha_d)$ be a vector, where $\alpha_i \in [n_i]$ for $i \in [d]$.*

1. *If an $\alpha$-cut exists, then it is unique.*
2. *If $P$ has weak general position, then an $\alpha$-cut exists for each choice of $\alpha$.*

That means, every colorful $d$-tuple of $P$ represents an oriented hyperplane that corresponds to exactly one $\alpha$-vector. Steiger and Zhao [40] also presented an algorithm to compute the cut in $O(n(\log n)^{d-3})$ time, where $n = \sum_{i=1}^{d} n_i$. The algorithm proceeds inductively in dimension and employs a prune-and-search technique. Bereg [8] improved the pruning step to improve the runtime to $n2^{O(d)}$.

## 2.2 Unique End of Potential Line

We briefly explain the *Unique end of potential line* problem that was introduced in [20]. More details about the problem and the associated class can be found in the above reference.

▶ **Definition 3** (from [20]). *Let $n, m$ be positive integers. The input consists of*

- *a pair of Boolean circuits $\mathbb{S}, \mathbb{P} : \{0,1\}^n \to \{0,1\}^n$ such that $\mathbb{P}(0^n) = 0^n \neq \mathbb{S}(0^n)$, and*
- *a Boolean circuit $\mathbb{V} : \{0,1\}^n \to \{0,1,\ldots,2^m-1\}$ such that $\mathbb{V}(0^n) = 0$,*

*each circuit having $\text{poly}(n,m)$ size. The* UniqueEOPL *problem is to report one of the following:*

**(U1).** *A point $v \in \{0,1\}^n$ such that $\mathbb{P}(\mathbb{S}(v)) \neq v$.*

**(UV1).** *A point $v \in \{0,1\}^n$ such that $\mathbb{S}(v) \neq v$, $\mathbb{P}(\mathbb{S}(v)) = v$, and $\mathbb{V}(\mathbb{S}(v)) - \mathbb{V}(v) \leq 0$.*

**(UV2).** *A point $v \in \{0,1\}^n$ such that $\mathbb{S}(\mathbb{P}(v)) \neq v \neq 0^n$.*

**(UV3).** *Two points $v, u \in \{0,1\}^n$ such that $v \neq u$, $\mathbb{S}(v) \neq v$, $\mathbb{S}(u) \neq u$, and either $\mathbb{V}(v) = \mathbb{V}(u)$ or $\mathbb{V}(v) < \mathbb{V}(u) < \mathbb{V}(\mathbb{S}(v))$.*

The problem defines a graph $G$ with up to $2^n$ vertices. Informally, $\mathbb{S}(\cdot), \mathbb{P}(\cdot), \mathbb{V}(\cdot)$ represent the *successor*, *predecessor* and *potential* functions that act on each vertex in $G$. The in-degree and out-degree of each vertex is at most one. There is an edge from vertex $u$ to vertex $v$ if and only if $\mathbb{S}(u) = v$, $\mathbb{P}(v) = u$ and $\mathbb{V}(u) < \mathbb{V}(v)$. Thus, $G$ is a directed acyclic path graph (line) along which the potential strictly increases. The condition $\mathbb{S}(\mathbb{P}(x)) \neq x$ means that $x$ is the start of the line, $\mathbb{P}(\mathbb{S}(x)) \neq x$ means that $x$ is the end of the line, and $\mathbb{P}(\mathbb{S}(x)) = x$ occurs when $x$ is neither. The vertex $0^n$ is a given start of the line in $G$.

**(U1)** is a solution representing the end of a line. **(UV1)**, **(UV2)** and **(UV3)** are violations. **(UV1)** gives a vertex $v$ that is not the end of line, and the potential of $\mathbb{S}(v)$ is not strictly larger than that of $v$, which is a violation of our assumption that the potential increases strictly along the line. **(UV2)** gives a vertex that is the start of a line, but is not $0^n$. **(UV3)** shows that $G$ has more than one line, which is witnessed by the fact that $v$ and $u$ cannot lie on the same line if they have the same potential, or if the potential of $u$ is sandwiched between that of $v$ and the successor of $v$. Under the promise that there are no violations, $G$ is a single line starting at $0^n$ and ending at a vertex that is the unique solution. UniqueEOPL is formulated in the non-promise setting, placing it in the class TFNP.

The complexity class UEOPL represents the class of problems that can be reduced in polynomial time to UniqueEOPL. This has been shown to lie in CLS and contains three classical problems in [20]: finding the fixed point of a piecewise-linear contraction map, solving the P-Matrix Linear complementarity problem, and finding the unique sink of a directed graph (with arbitrary edge orientations such that each face has a sink) on the 1-skeleton of a hypercube. Note that finding the fixed point of a contraction map is in CLS [18], but is not known to lie in UEOPL.

A notion of *promise-preserving* reductions is also defined in [20]. Let $X$ and $Y$ be two problems both having a formulation that allows for valid and violation solutions. A reduction from $X$ to $Y$ is said to be promise-preserving, if whenever it is promised that $X$ has no violations, then the reduced instance of $Y$ also has no violations. Thus a promise-preserving reduction to UniqueEOPL would mean that whenever the original problem is free of violations, then the reduced instance always has a single line that ends at a valid solution.

## 2.3 Formulating the search problem

We formalize the search problem for $\alpha$-Ham-Sandwich in a non-promise setting:

▶ **Definition 4** (Alpha-HS). *Given $d$ finite sets of points $P = P_1 \cup \ldots \cup P_d$ in $\mathbb{R}^d$ and a vector $(\alpha_1, \ldots, \alpha_d)$ of positive integers such that $\alpha_i \leq |P_i|$ for all $i \in [d]$, the* Alpha-HS *problem is to find one of the following:*

**(G1).** *An $(\alpha_1, \ldots, \alpha_d)$-cut.*

**(GV1).** *A subset of $P$ of size $d+1$ and at least $d-1$ colors that lies on a hyperplane.*

**(GV2).** *A disjoint pair of sets $I, J \subset [d]$ such that $\operatorname{conv}(\{\cup_{i\in I} P_i\}) \cap \operatorname{conv}(\{\cup_{j\in J} P_j\}) \neq \emptyset$.*

Here a solution of type **(G1)** corresponds to a solution representing a valid cut, while solutions of type **(GV1)** and **(GV2)** refer to violations of weak general position and well-separation, respectively. From Theorem 2 we see that a valid solution is guaranteed if no violations are presented, which shows that Alpha-HS is a total search problem.

## 3 Alpha-HS is in UEOPL

In this section we describe our instance of Alpha-HS in more detail and briefly outline a reduction to UniqueEOPL.

**Setup.** The input consists of $d$ finite point sets $P_1, \ldots, P_d \subset \mathbb{R}^d$ each representing a unique color, of sizes $n_1, \ldots, n_d$, respectively, and a vector of integers $\alpha = (\alpha_1, \ldots, \alpha_d)$ such that $\alpha_i \in [n_i]$ for each $i \in [d]$. Let $k$ denote the number of coordinates of $\alpha$ that are not equal to 1. Without loss of generality, we assume that $\{\alpha_1, \ldots, \alpha_k\}$ are the non-unit entries in $\alpha$. Let $P$ denote the union $P_1 \cup \cdots \cup P_d$. For each $i \in [d]$ we define an arbitrary order $\prec_i$ on $P_i$. Concatenating the orders $\prec_1, \prec_2, \ldots, \prec_d$ in sequence gives a global order $\prec$ on $P$. That means, $p \prec q$ if $p \in P_i, q \in P_j$ and $i < j$ or $p, q \in P_j$ and $p \prec_j q$.

We follow the notation of [40] to define the orientation of a hyperplane in $\mathbb{R}^d$ that has a non-empty intersection with the convex hull of each $P_i$. For any hyperplane $H$ passing via $\{x_1 \in \operatorname{conv}(P_1), \ldots, x_d \in \operatorname{conv}(P_d)\}$, the normal is the unit vector $\hat{n} \in \mathbb{R}^d$ that satisfies $\langle x_i, \hat{n}\rangle = t$ for some fixed $t \in \mathbb{R}$ and each $i \in [d]$, and $\det \begin{vmatrix} x_1 & x_2 & \ldots & x_d & \hat{n} \\ 1 & 1 & \ldots & 1 & 0 \end{vmatrix} > 0$, where the columns of the matrix are determined using the order $\prec$. The positive and negative half-spaces of $H$ are defined accordingly. In [4, Proposition 2], the authors show that the choice of $\hat{n}$ does not depend on the choice of $x_i \in \operatorname{conv}(P_i)$ for any $i$, if the colors are well-separated. Notice that if the colors are not well-separated, then the dimension of the affine hull of $\{x_1, \ldots, x_d\}$ may be less than $d-1$. This makes the value of the determinant above to be zero, so the orientation is not well-defined.

We call a hyperplane *colorful* if it passes through a colorful set $\{p_1, \ldots, p_d\} \subset P$. Otherwise, we call the hyperplane *non-colorful*. There is a natural orientation for colorful hyperplanes using the definition above. In order to define an orientation for non-colorful hyperplanes, one needs additional points from the convex hulls of unused colors on the hyperplane. Let $H'$ denote a hyperplane that passes through points of $(d-1)$ colors. Let $P_j$ denote the missing color in $H'$. To define an orientation for $H'$, we choose a point from $\operatorname{conv}(P_j)$ that lies on $H'$ as follows. We collect the points of $P_j$ on each side of $H'$, and choose the highest ranked points under the order $\prec_j$. Let these points on opposite sides of $H'$ be denoted by $x$ and $y$. Let $z$ denote the intersection of the line segment $xy$ with $H'$. By convexity, $z$ is a point in $\operatorname{conv}(P_j)$, so we choose $z$ to define the orientation of $H'$. The intersection point $z$ does not change if $x$ and $y$ are interchanged, giving a valid definition of orientation for $H'$. We can also extend this construction to define orientations for hyperplanes containing points from fewer than $d-1$ colors, but for our purpose this definition suffices. The *$\alpha$-vector* of any oriented hyperplane $H$ is a $d$-tuple $(\alpha_1, \ldots, \alpha_d)$ of integers where $\alpha_i$ is the number of points of $P_i$ in the closed halfspace $H^+$ for $i \in [d]$.

## 3.1 An overview of the reduction

We give a short overview of the ideas used in the reduction from Alpha-HS to UniqueEOPL. The details are technical and we encourage the interested reader to go through the details of our reduction in [13].

Our intuition is based on rotating a colorful hyperplane $H$ to another colorful hyperplane $H'$ through a sequence of local changes of the points on the hyperplanes such that the $\alpha$-vector of $H'$ increases in some coordinate by one from that of $H$. We next define the rotation operation in a little more detail. An *anchor* is a colorful $(d-1)$-tuple of $P$ which spans a $(d-2)$-flat. The following procedure takes as input an anchor $R$ and some point $p \in P \setminus R$ and determines the next hyperplane obtained by a rotation. The output is $(R', p')$, where $R'$ is an anchor and $p' \in P \setminus R'$ is some point.

**Procedure** $(R', p') = \mathit{NextRotate}(R, p)$

1. Let $H$ denote the hyperplane defined by $R \cup \{p\}$ and $t_1$ be the missing color in $R$.
2. If the orientation of $H$ is not well-defined, report a violation of weak general position and well-separation.
3. Let $P^+_{t_1}$ be the subset of $P_{t_1}$ that lies in the closed halfspace $H^+$ and $P^-_{t_1}$ be the subset of $P_{t_1}$ that lies in the open halfspace $H^-$. Let $x \in P^+_{t_1}$ be the highest ranked point according to the order $\prec_{t_1}$ and $y \in P^-_{t_1}$ be the highest ranked point according to $\prec_{t_1}$.
4. If $p$ has color $t_1$ and $|P^+_{t_1}| = n_{t_1}$, report out of range.
5. We rotate $H$ around the anchor $R$ in a direction such that the hyperplane is moving away from $x$ along the segment $xy$ until it hits some point $q \in P$.
6. If the hyperplane hits multiple points at the same time, report a violation of weak general position.
7. If $q$ is not color $t_1$, set $R' := R \cup \{q\} \setminus \{r\}$ and $p' = r$, where $r$ is a point in $R$ with the same color as $q$. Otherwise, set $R' = R$ and $p' = q$.
8. Return $(R', p')$.

Figure 3 shows an application of this procedure, rotating $H_0$ to $H_4$ through $H_1, H_2, H_3$.

This rotation function can be interpreted as a function that assigns each hyperplane to the next hyperplane. The set of colorful hyperplanes can be interpreted as vertices in a graph with the rotation function determining the connectivity of the graph.

**Canonical path.** Each colorful hyperplane $H$ is incident to a colorful set of $d$ points. This set of points defines $d$ possible anchors, and each anchor can be used to rotate $H$ in a different fashion. To define a unique sequence of rotations, we pick a specific order as follows: first, we assume that the colorful hyperplane $H$ whose $\alpha$-vector is $(1, \dots, 1)$ is given (we show later how this assumption can be removed). We start at $H$ and pick the anchor that excludes the first color, then apply a sequence of rotations until we hit another colorful hyperplane with $\alpha$-vector $(2, 1, \dots, 1)$. Similarly, we move to a colorful hyperplane with $\alpha$-vector $(3, 1, \dots, 1)$ and so on until we reach $(\alpha_1, 1, \dots, 1)$. Then, we repeat this for the other colors in order to reach $(\alpha_1, \alpha_2, 1, \dots, 1)$ and so on until we reach the target $\alpha$-vector. This pattern of $\alpha$-vectors helps in defining a potential function that strictly increases along the path. We can encode this sequence of rotations as a unique path in the UniqueEOPL instance, and we call it *canonical path.*

A natural way to define the UniqueEOPL graph would be to consider hyperplanes as the vertices in the graph. However, this leads to complications. Figure 3 shows a rotation from $H_0$ to $H_4$, with $\alpha$-vectors $(3, 2)$ and $(3, 3)$ respectively. During the rotation, we encounter a hyperplane $H_2$ for which its $\alpha$-vector is $(4, 2)$, which differs from our desired sequence

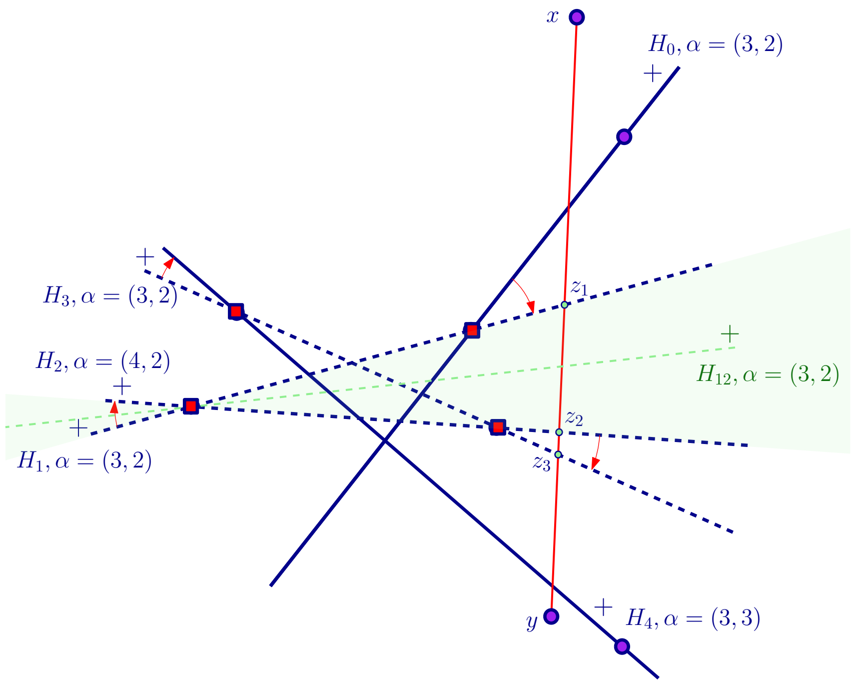


**Figure 3** An example showing a sequence of rotations from $H_0$ to $H_4$ through $H_1, H_2, H_3$. Red (square) is the first color and purple (disk) is the second color. This sequence represents a path between two vertices in the UniqueEOPL graph that is generated in the reduction. The double-wedge is shaded and its angular bisector $H_{12}$ has the desired $\alpha$-vector.

of $(3, 2), \ldots, (3, 2), (3, 3)$. This makes it difficult to define a potential function in the graph that strictly increases along the path $v_{H_0}, \ldots, v_{H_4}$ where $v_{H_i}$ is the vertex representing hyperplane $H_i$. One way to alleviate this problem is to not use $H_i$ as a vertex directly, but the *double-wedge* that is traced out by the rotation from $H_i$ to $H_{i+1}$. If the $\alpha$-vector is now measured using the hyperplane that bisects the double-wedge, then we get the desired sequence of $(3, 2), \ldots, (3, 2), (3, 3)$. See Figure 3 for an example.

With additional overhead, the rotation function can be extended to double-wedges. This in turn also leads to a neighborhood graph where the vertices are the double-wedges and the rotations can be used to define the edges. The graph is connected and has a grid-like structure that may be of independent interest. Due to lack of space, the description of double-wedges and the associated graph can be found in [13].

**Distance parameter and potential function.** The $\alpha$-vector is not sufficient to define the potential function, since the sequence of rotations between two colorful hyperplanes may have the same $\alpha$-vector. For instance, the bisectors of the rotations in $H_0, \ldots, H_3$ in Figure 3 all have the same $\alpha$-vector. Hence, we need an additional measurement in order to determine the direction of rotation that increases the $\alpha$-vector.

Similar to how we define the orientation for a non-colorful hyperplane, let $H$ denote a hyperplane that passes through points of $(d-1)$ colors. Let $P_j$ denote the missing color in $H$. Let $x, y \in P_j$ be the highest ranked points under $\prec_j$ in $H^+$ and $H^-$ respectively. Let $z$ denote the intersection of $xy$ and $H$. We define a distance parameter called *dist-value* of $H$ to be the distance $\|x - z\|$. In Figure 3, we can see that rotating from $H_0$ to $H_4$ sweeps the segment $xy$ in one direction, with the dist-value of the hyperplanes increasing strictly. This is sufficient to break ties and hence determine the correct direction of rotation. The precise statement is given in Lemma 6. We can extend this definition to the domain of double-wedges. We define a potential value for each vertex on the canonical path in UniqueEOPL using the sum of weighed components of $\alpha$-vector and dist-value for the tie-breaker.

**Correctness.** We show that if there are no violations, we can always apply **Procedure** $NextRotate$ to increment the $\alpha$-vector until we find the desired solution, which implies that the canonical path exists. If the input satisfies weak general position, we can see that the rotating hyperplane always hits a unique point in Step 5, which may be swapped to form a new anchor in Step 7.

The well-separation condition guarantees that the potential function always increases along the rotation. Let $H_1, H_2$ denote a pair of hyperplanes that are the input and output of **Procedure** $NextRotate$ respectively. Let $H$ denote any intermediate hyperplane during the rotation from $H_1$ to $H_2$ through the common anchor. Let $P_j$ be the color missing from the anchor and $x$ be the highest ranked point under $\prec_j$ in $H_1^+$. We say that the orientation of $H_2$ (resp. $H$) is *consistent* with that of $H_1$ if $x \in H_2^+$ (resp. $x \in H^+$). Lemma 5 shows that the orientations are always consistent when $H_1$ and $H_2$ are non-colorful hyperplanes even without the assumption of well-separation.

▶ **Lemma 5** (consistency of orientation)**.** *Assume that weak general position holds. Let $H_1, H_2$ be the input and output of* **Procedure** $NextRotate$ *respectively. Let $H$ denote any intermediate hyperplane within the rotation. The orientations of $H_1$ (resp. $H_2$) and $H$ are consistent when $H_1$ (resp. $H_2$) is a non-colorful hyperplane.*

**Proof.** Since $H_1$ is a non-colorful hyperplane, let $P_j$ denote the color missing from $H_1$. $H_1$ and $H$ give the same partition of $P_j$ into two sets because the continuous rotation from $H_1$ to $H$ does not hit any point in $P_j$. Let $x$ and $y$ be the highest ranked points under $\prec_j$ in each set. Since we have weak general position, the segment $xy$ cannot pass through the anchor of the rotation so that the orientations of $H_1$ and $H$ are well-defined by the $(d-1)$ colored points in the anchor and the intersections of the hyperplanes with the segment $xy$. Thus, the determinant defining the normal of the rotating hyperplane from $H_1$ to $H$ for the orientation is always non-zero. Since the intersection of the rotating hyperplane from $H_1$ to $H$ and the segment $xy$ moves continuously along $xy$, by a continuity argument, the normal of the hyperplane does not flip during the rotation. Without loss of generality, assume that $x \in H_1^+$. This implies that $x$ is always in the positive half-space of $H$ and hence $H$ has a consistent orientation as $H_1$. The same proof holds for $H_2$. ◀

Next, we show that the dist-value is strictly increasing for all the intermediate hyperplanes in the sequence of rotations from one colorful hyperplane to another colorful hyperplane.

▶ **Lemma 6.** *Assume that weak general position holds. Let $H_0$ be a colorful hyperplane and $H_k$ be the first colorful hyperplane obtained by a sequence of rotations by* **Procedure** $NextRotate$*. We denote by $H_1, \dots, H_{k-1}$ the non-colorful hyperplanes obtained from the above sequence of rotations. The dist-values of $H_1, \dots, H_{k-1}$ are strictly increasing.*

**Proof.** Let $P_j$ denote the color missing from $H_1$. Then, $H_2, \dots, H_{k-1}$ all miss the color $P_j$, otherwise $H_k$ is not the first colorful hyperplane obtained by the rotations. Therefore, each $H_i$ gives the same partition of $P_j$ into two sets for $i = 1, \dots, k-1$ because the continuous rotations from $H_1$ to $H_{k-1}$ does not hit any point in $P_j$. Let $x$ and $y$ be the highest ranked points under $\prec_j$ in each set. Without loss of generality, assume that $x \in H_1^+$. Since $H_1, \dots, H_{k-1}$ are non-colorful hyperplanes, by Lemma 5, the consistent of the orientation can carry from

$H_1$ to $H_2$ and so on. Then we have $x \in H_1^+, \ldots, x \in H_{k-1}^+$ and $y \in H_1^-, \ldots, y \in H_{k-1}^-$. Let $z_1 = xy \cap H_1, \ldots, z_{k-1} = xy \cap H_{k-1}$. According to Step 5 of **Procedure** *NextRotate*, each rotation is performed by moving away from $x$ along the segment $xy$. Hence we have $\|x - z_1\| < \|x - z_2\| < \cdots < \|x - z_{k-1}\|$. ◀

The last step for proving that the potential function always increases along the canonical path is to show that the $\alpha$-vector increases in some coordinate from one colorful hyperplane to another colorful hyperplane through **Procedure** *NextRotate*. This requires the assumption of well-separation. Lemma 7 shows that if the orientations of $H_1, H_2$ and $H$ are inconsistent, then well-separation is violated. By the contrapositive, if well-separation is satisfied, then all hyperplanes in the rotation always give consistent orientations. Then, it implies that rotating from a colorful hyperplane $H_0$ to another colorful hyperplane $H_k$ through a sequence of non-colorful hyperplanes that miss color $P_j$, we have $H_0^+ \cap P_j \subset H_k^+ \cap P_j$ and $H_k$ contains one additional point in $P_j$ that is hit by the last rotation. Therefore, $\alpha_j$ is increased by 1 and other $\alpha_i$s keep the same value because of the way we swap the point of repeated color with the one in the anchor and the direction of rotation.

▶ **Lemma 7.** *Assume that weak general position holds. Let $H_1, H_2$ be the input and output of* **Procedure** *NextRotate respectively. Let $R$ denote the anchor of the rotation from $H_1$ to $H_2$, and $P_j$ denote the color missing from $R$. Let $H$ denote any intermediate hyperplane within the rotation. If the orientations of $H_1$ (resp. $H_2$) and $H$ are inconsistent, then $H_1$ (resp. $H_2$) is a colorful hyperplane and we can find a colorful set $R \cup \{x'\}$ lying in a $(d-2)$-flat where $x' \in \mathrm{conv}(P_j)$, in $O(d^3)$ arithmetic operations. The set $R \cup \{x'\}$ witnesses the violation of well-separation.*

**Proof.** Since the orientations of $H_1$ and $H$ are inconsistent, $H_1$ must be a colorful hyperplane by Lemma 5. Therefore, the point in $H_1$ that is not in the anchor is in $P_j$, denoted by $p$.

Let $x$ and $y$ be the points defined in Lemma 5 such that $x, y \in P_j$, and $x$ and $y$ are on different sides of $H_1$ and $H$. The $(d-2)$-flat containing $R$ separates $H_1$ and $H$ into two $(d-1)$-dimensional half-subspaces each. Let $H_{1,R}^+$ and $H_R^+$ be the half-subspaces intersecting with $xy$ on $H_1$ and $H$ respectively, and let us denote the intersection points by $z_p$ and $z$, respectively. The opposite half-subspaces are denoted by $H_{1,R}^-$ and $H_R^-$, respectively. By definition of the orientation for non-colorful hyperplanes, the orientation of $H$ is defined by $R \cup \{z\}$. Although the orientation of $H_1$ is defined by $R \cup \{p\}$, if we consider the determinant defining the orientation using $R \cup \{z_p\}$, it gives an orientation consistent with that of $H$. Therefore, it must be that $p \in H_{1,R}^-$. Then, we can see that the line segment $pz_p$ intersects the $(d-2)$-flat of $R$. We can compute $z_p$ and also the intersection point $x'$ of $pz_p$ and the $(d-2)$-flat of $R$ by solving systems of linear equations with $d$ equations and $d$ variables in $O(d^3)$ arithmetic operations. Since $x' \in \mathrm{conv}(P_j)$, $R \cup \{x'\}$ is a colorful set contained in the $(d-2)$-flat of $R$. ◀

In order to guarantee that there is no other path in UNIQUEEOPL apart from the canonical path, we introduce self-loops for vertices that are not on the canonical path. The detailed proof in [13] shows that if there are no violations, then the reduced instance of UNIQUEEOPL only gives a **(U1)** solution, which readily translates to a **(G1)** solution, so our reduction is promise-preserving, and this can be done in polynomial time.

Since we do not know the hyperplane with $\alpha$-vector $(1, \ldots, 1)$ in advance, we split the problem into two sub-problems: in the first we start with any colorful hyperplane. We reverse the direction of the canonical path determined by the potential and construct an ALPHA-HS instance for which the vertex with $\alpha$-vector $(1, \ldots, 1)$ is the solution. In the second, we use this vertex as the input to the main Alpha-HS instance. If the input is free of violations, then both sub-problems give valid solutions and together they answer the original question. To merge the two sub-problems into one UniqueEOPL instance, we can make two layer copies of the vertices with an additional flag variable to indicate which copy is in the first layer. In the first layer, we build the canonical path from any colorful vertex to the colorful vertex with $\alpha$-vector $(1, \ldots, 1)$, which connects to the colorful vertex with $\alpha$-vector $(1, \ldots, 1)$ in the second layer. Similarly, in the second layer, we build the canonical path from the colorful vertex with $\alpha$-vector $(1, \ldots, 1)$ to the vertex with the target $\alpha$-vector. Then, we can also easily modify the potential function accordingly.

An alternative approach is to define the canonical path directly from any colorful vertex to the target vertex. In this case, each coordinate of the current $\alpha$-vector may increase or decrease depending on the signed distance to the target $\alpha$-vector along the canonical path. However, the potential function can still be defined in a way that it is strictly increasing along the path.

**Handling violations.** The reduction maps violations of Alpha-HS to violations of the UniqueEOPL instance, and certificates for the violations can be recovered from additional processing. When a violation of weak general position is witnessed on a vertex that lies on the canonical path, a hyperplane incident to $d$ colors may contain additional points. This in turn implies that some $\alpha$-cut is missing, so that the correct solution for the target may not exist. For cuts that exist in spite of the violation, reporting either the correct solution or the violation are sufficient for Alpha-HS.

In addition, the (highest-ranked) points $x, y$ from the missing color that we choose to define the orientation of a non-colorful hyperplane may form a segment $xy$ that passes through the $(d-2)$-flat spanned by the anchor. In that case the orientation of the hyperplane is not well-defined. In the reduction, these problematic vertices are removed from the canonical path, thereby creating some additional starting points and end points in the reduced instance. These violations can be captured by **(U1)** with a wrong $\alpha$-vector or **(UV2)**. Furthermore, the hyperplanes that contain the degenerate point sets could be represented by different choices of anchors and an additional point on the plane. Each such pair represents a vertex in the reduced instance. We join these vertices in the form of a cycle in the UniqueEOPL instance with all vertices having the same potential value, so that the violations can also be captured by **(UV1)** and **(UV3)**.

When a violation of well-separation is witnessed on a vertex on the canonical path, the orientations of the two hyperplanes paired by **Procedure** *NextRotate* may be inconsistent, which may not guarantee that the $\alpha$-vector is incremented in one component by one (See Figure 4). Hence, the canonical path is split into two paths that can be captured by **(UV2)**. Furthermore, a violation of well-separation also creates multiple colorful hyperplanes with the same $\alpha$-vector (See Figure 4, left). Two vertices in the UniqueEOPL graph with the same potential value, which could correspond to some colorful or non-colorful hyperplanes, can be reported by **(UV3)**. We show that this gives a certificate of violation of well-separation in the following lemmas, where $m_0$ is the number of bits used to represent each coordinate of points of $P$.

▶ **Lemma 8.** *Given two colorful hyperplanes $H_p$, $H_q$ with the same $\alpha$-vector, we can find a colorful set $\{x_1 \in \mathrm{conv}(P_1), \ldots, x_d \in \mathrm{conv}(P_d)\}$ that lies on a $(d-2)$-flat in $\mathrm{poly}(n, d, m_0)$ time.*

▶ **Lemma 9.** *Given two non-colorful hyperplanes that both contain $d-1$ points and have*

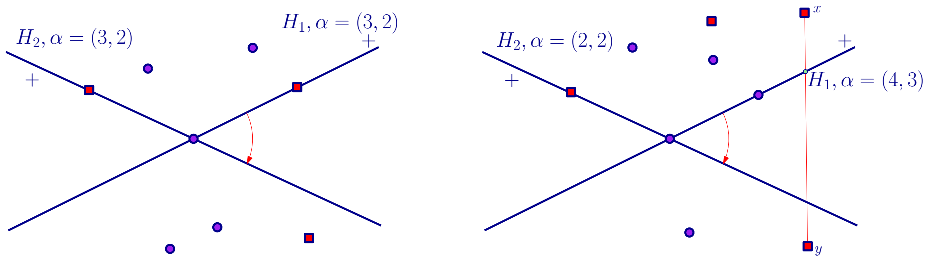


**Figure 4** The examples show two sets of points that are not well-separated. Purple (circle) represents the first color and red (square) represents the second color. In both examples the rotation procedure does not increase the $\alpha$-vector. Both examples show that the orientation of the hyperplane may be flipped after the rotation, so the resulting $\alpha$-vector can go wrong.

*the same missing color, $\alpha$-vector and dist-value, we can find a colorful set of points $\{x_1 \in \mathrm{conv}(P_1), \ldots, x_d \in \mathrm{conv}(P_d)\}$ that lies on a $(d-2)$-flat in $\mathrm{poly}(n, d, m_0)$ time.*

For the second output $(\mathbb{V}(v) < \mathbb{V}(u) < \mathbb{V}(\mathbb{S}(v)))$ of **(UV3)**, there are two cases to consider. In the first case, if both $v$ and $\mathbb{S}(v)$ correspond to the same $\alpha$-vector, then $u$ also has the same $\alpha$-vector and its dist-value is between that of $v$ and $\mathbb{S}(v)$. Since rotating the hyperplane from $v$ to $\mathbb{S}(v)$ does not pass through $u$, we can find a different hyperplane that is interpolated by $v$ and $\mathbb{S}(v)$ and has the same dist-value as $u$. Hence, we apply Lemma 9 again to find a witness of the violation. For the second case that the $\alpha$-vector of $\mathbb{S}(v)$ increases in one coordinate by one from that of $v$, since the role of dist-value is dominated by the role of $\alpha$-vector in the potential function, the dist-value of $u$ can be arbitrarily large. Therefore, we may not be able to apply the interpolation technique from the first again. We argue that we can transform $P$ to a point set $P'$ satisfying $\mathrm{conv}(P'_i) \subseteq \mathrm{conv}(P_i)$ for all $i \in [d]$, such that the hyperplanes of $v$ and $u$ become colorful. Then, we apply Lemma 8 to show that $P'$ is not well-separated, which also implies that $P$ is not well-separated. The precise statement and proof are given in [13]. We also show

- how to compute a **(GV1)** solution from a **(UV1)** solution,
- how to compute a **(GV1)** or **(GV2)** solution, given a **(UV2)** or **(UV3)** solution, and
- a **(GV1)** or **(GV2)** solution that can occur with a **(U1)** solution that has the incorrect $\alpha$-vector.

We show that converting these solutions always takes $\mathrm{poly}(n, d)$ time. The violations may be detected in either the first sub-problem or the second sub-problem. Our constructions thus culminate in the promised result:

▶ **Theorem 10.** *Alpha-HS* $\in$ UEOPL $\subseteq$ CLS.

**Handling very weak general position.** We have described our construction for the case when weak general position holds. If we only assume that very weak general position holds, then there may exist a hyperplane that passes through more than $d$ points of at most $d-1$ colors. Therefore, in Step 5 of **Procedure** *NextRotate* the rotating hyperplane may hit more than one point so that it is not clear how to define the new anchor in Step 7. From the point of view of the reduction, there are many non-colorful vertices that represent the same hyperplane. We need a new approach to define a unique path to traverse these vertices with respect to this hyperplane. In other words, we charge the computational time of finding the new anchor to traversing these vertices on the path instead of considering it as one operation.

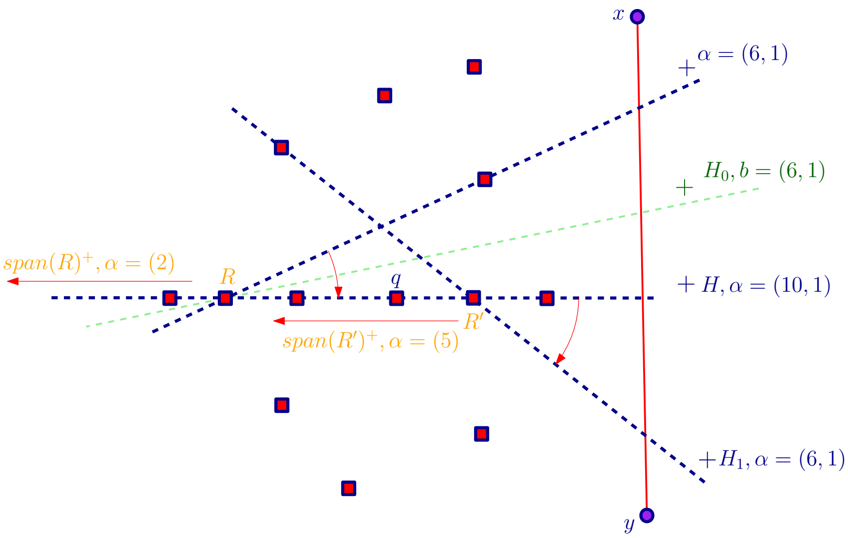


**Figure 5** An example showing the relationship between the $\alpha$-vector in a subproblem in $\mathbb{R}$ and the $\alpha$-vector in the original problem in $\mathbb{R}^2$. Red (square) is the first color and purple (disk) is the second color. The orientation of $span(R)$ in $H$ is defined such that it is consistent with $H_0$. $b = (6, 1)$ is the $\alpha$-vector of $H_0$. $k_1 = 4$ is the number of red points in $H^+ \setminus H$. The $\alpha$-vector of the starting vertex (i.e., $R$) with respect to $H$ is $(6 - 4 = 2)$. The $\alpha$-vector of the end vertex is $(6 + 1 - 2 = 5)$. We can see that $q \in span(R')^+$ and $q$ moves to the negative side of $H_1$ when rotating from $H$ to $H_1$.

If we consider the space of all the points lying on the hyperplane, we have $d-1$ sets of points each representing a unique color in an affine subspace of $d-1$ dimensions. Thus, we can consider it as a new instance of Alpha-HS in one dimension lower. Let $H$ be the rotating hyperplane that hits more than one point and contains $d-1$ colors. Without loss of generality, we assume that $d$ is the missing color. We denote by $Q = Q_1 \cup Q_2 \cup \ldots \cup Q_{d-1}$ the $d-1$ sets of points in $H$ such that $Q_i \subseteq P_i$ and denote by $\widehat{Q}_i$ the set of points represented in the new coordinate system in $\mathbb{R}^{d-1}$ for $Q_i$ in $H$. First, we claim that if $P$ is well-separated and in very weak general position, then $\widehat{Q}$ is also well-separated and in very weak general position. Since $Q \subset P$, it is clear that well-separation follows. Suppose that $\widehat{Q}$ violates very weak general position, then there exists a $(d-2)$-flat that contains more than $d-1$ points of $d-1$ colors in $Q$. In particular, any $(d-1)$-flat spanned by the $(d-2)$-flat and any point in $P_d$ contains more than $d$ points of $d$ colors, which contradicts the fact that $P$ is in very weak general position.

Suppose that $P$ is well-separated and in very weak general position. Now we define what is the unique path with respect to $\widehat{Q}$. Let $b = (b_1, \ldots, b_d)$ be the $\alpha$-vector of the rotating hyperplane $H_0$ just before rotating to $H$ at the anchor $R$. In the new instance of Alpha-HS, we would pick the orientation of $(d-2)$-flats in $\mathbb{R}^{d-1}$ such that every point $p \in Q$ lies in $H_0^+$ if and only if the corresponding point $\widehat{p} \in \widehat{Q}$ lies in $span(\widehat{R})^+$. Let $k_1, \ldots k_{d-1}$ denote the number of points of $P_1, \ldots, P_{d-1}$ in $H^+$, but not in $Q_i$. Then, we can see that the number of points in $\widehat{Q}_i$ lying in $span(\widehat{R})^+$ is equal to $b_i - k_i$. Thus, the $\alpha$-vector of $span(\widehat{R})^+$ is $(b_1 - k_1, \ldots, b_{d-1} - k_{d-1})$, which is the $\alpha$-vector of the starting vertex of the path. On the other hand, the $\alpha$-vector of the end vertex is $(|Q_1| + 1 - b_1 + k_1, \ldots, |Q_{d-1}| + 1 - b_{d-1} + k_{d-1})$. It is because the points in $H_0^+ \setminus H_0$ become in the opposite side after the rotation passes through $H$. Therefore, if we rotate at the new anchor with $\alpha$-vector $(|Q_1| + 1 - b_1 + k_1, \ldots, |Q_{d-1}| + 1 - b_{d-1} + k_{d-1})$ in $\widehat{Q}$, then the $\alpha$-vector of the new rotating hyperplane is still $(b_1, \ldots, b_d)$. The next question is that if the vertex only stores any $d$ points of $H$, we cannot recover $b$ and $H_0$ so that the orientation cannot be defined consistently and the target $\alpha$-vector for $\widehat{Q}$ is not known. To handle this problem, we need to redefine the double-wedge to be $(R_1, p_1, R_2, p_2)$ instead of $(R, p, q)$ in such a way that $R_1 = R_2$ if the double-wedge contains exactly $d+1$ points, otherwise $R_1 \subset span(R_2 \cup \{p_2\})$. For instance, if $(\widehat{R}_1, \widehat{q}_1)->\ldots->(\widehat{R}_m, \widehat{q}_m)$ is the unique path in $\widehat{Q}$, where $\widehat{R}_i$ is an anchor of size $d-2$ so that $\widehat{R}_i$ and $\widehat{q}_i$ represent a $(d-2)$-flat in $\mathbb{R}^{d-1}$, then the corresponding path in the original problem is $(R, p, R_1 \cup \{q_1\}, p_1)->(R, p, R_2 \cup \{q_2\}, p_2)->\ldots->(R, p, R_m \cup \{q_m\}, p_m)$, where $p_i$ is some point in $H$ that is picked under $\prec$ in a way that the tuple is uniquely defined in the path. Hence, $b$ can be computed from the bisector of $(R, p)$ and $(R_i \cup \{q_i\}, p_i)$, and the orientation of $(d-2)$-flats can also be defined by the bisector. There may exist some other double-wedge $(*, *, R_i \cup \{q_i\}, p_i)$ that is incident to $H$, but it will not have the same $b$.

In conclusion, the unique path in the reduction can be defined recursively as above in an ALPHA-HS instance of one dimension lower. As a result, the representation of the double-wedges gets more complicated and the size is increased by a factor of $O(d)$. The potential function becomes a weighted sum of the potential function in each recursive level, but the number of bits is still in polynomial size. For handling violations, there are not many changes. Instead of reporting the violation of weak general position, we now report the violation of very weak general position when the rotating hyperplane in $R^i$ contains more than $i$ points of $i$ colors. If any recursive subproblem violates very weak general position, it also implies that the original input $P$ violates very weak general position.

## 4 Conclusion and future work

We gave a complexity-theoretic upper bound for ALPHA-HS. No hardness results are known for this search problem, and the next question is determining if this is hard for UEOPL. One challenge is that UNIQUEEOPL is formulated as Boolean circuits, whereas ALPHA-HS is purely geometric. Emulating circuits using purely geometric arguments is highly non-trivial. Filos-Ratsikas and Goldberg showed a reduction of this form in [21]. They reduced the PPA-complete 2D-Tucker circuit to HAM-SANDWICH, going via the *Consensus-Halving* [39], and the *Necklace-splitting problems* [2]. A simplified argument was recently presented in [22]. It could be a worthwhile exercise to investigate if their techniques can provide insights for hardness of ALPHA-HS.

Some related problems are determining the complexity of answering whether a point set is well-separated, whether it is in weak general position, or whether a given $\alpha$-cut exists for the point set. A given $\alpha$-cut may exist even when both assumptions are violated. On a related note, deciding whether the Linear Complementarity problem has a solution is NP-complete [15]. The solution is unique if the problem involves a $P$-matrix, but checking this condition is coNP-complete [16]. However, using witnesses to verify whether a matrix is P-matrix or not, a total search version is shown to be in UEOPL. Our result for ALPHA-HS would go in a similar vein, if the complexities of the above problems were better determined.

Another line to work could be to determine the computational complexities of other extensions of the Ham-Sandwich theorem. For other geometric problems that are total and admit unique solutions, it could be worthwhile to explore their place in the class UEOPL. Faster algorithms for computing the $\alpha$-cut can also be explored.

## References


1. James Aisenberg, Maria Luisa Bonet, and Sam Buss. 2-D Tucker is PPA-complete. *J. Comput. System Sci.*, 108:92–103, 2020.
2. Noga Alon and Douglas B. West. The Borsuk-Ulam theorem and bisection of necklaces. *Proc. American Mathematical Society*, 98(4):623–628, 1986.
3. Imre Bárány. A generalization of Carathéodory's theorem. *Discrete Mathematics*, 40(2-3):141–152, 1982.
4. Imre Bárány, Alfredo Hubard, and Jesús Jerónimo. Slicing convex sets and measures by a hyperplane. *Discrete Comput. Geom.*, 39(1):67–75, 2008.
5. Imre Bárány and Jiří Matoušek. Simultaneous partitions of measures by $k$-fans. *Discrete Comput. Geom.*, 25(3):317–334, 2001.
6. Luis Barba, Alexander Pilz, and Patrick Schnider. Sharing a pizza: bisecting masses with two cuts. *CoRR*, abs/1904.02502, 2019.
7. Sergey Bereg. Equipartitions of measures by 2-fans. *Discrete Comput. Geom.*, 34(1):87–96, 2005.
8. Sergey Bereg. Computing generalized ham-sandwich cuts. *Inform. Process. Lett.*, 112(13):532–534, 2012.
9. Pavle V. M. Blagojević , Aleksandra Dimitrijević Blagojević , Roman Karasev, and Jonathan Kliem. More bisections by hyperplane arrangements, 2018. `arXiv:1809.05364`.
10. Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. Settling the complexity of arrow-debreu equilibria in markets with additively separable utilities. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009*, pages 273–282, 2009.
11. Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3):14:1–14:57, 2009.
12. Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. The complexity of non-monotone markets. *J. ACM*, 64(3):20:1–20:56, 2017.
13. Man-Kwun Chiu, Aruni Choudhary, and Wolfgang Mulzer. Computational complexity of the $\alpha$-ham-sandwich problem, 2020. `arXiv:2003.09266`.
14. Aruni Choudhary and Wolfgang Mulzer. No-dimensional Tverberg theorems and algorithms. In *Proc. 36th Int. Sympos. Comput. Geom. (SoCG)*, page to appear, 2020.
15. S. J. Chung. NP-completeness of the Linear Complementarity Problem. *Journal of Optimization Theory and Applications*, 60(3):393–399, 1989.
16. Gregory E. Coxson. The P-matrix problem is co-NP-complete. *Mathematical Programming*, 64(1):173–178, 1994.
17. Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash Equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
18. Constantinos Daskalakis and Christos H. Papadimitriou. Continuous Local Search. In *Proc. 22nd Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 790–804, 2011.
19. Vladimir L. Dol'nikov. A generalization of the Ham-sandwich theorem. *Mathematical Notes*, 52(2):771–779, 1992.
20. John Fearnley, Spencer Gordon, Ruta Mehta, and Rahul Savani. Unique End of Potential Line. In *Proc. 46th Internat. Colloq. Automata Lang. Program. (ICALP)*, pages 56:1–56:15, 2019.
21. Aris Filos-Ratsikas and Paul W. Goldberg. The complexity of splitting Necklaces and bisecting Ham sandwiches. In *Proc. 51st Annu. ACM Sympos. Theory Comput. (STOC)*, pages 638–649, 2019.
22. Aris Filos-Ratsikas, Alexandros Hollender, Katerina Sotiraki, and Manolis Zampetakis. Consensus-halving: Does it ever get easier? *CoRR*, abs/2002.11437, 2020.
23. Vincent Froese, Iyad A. Kanj, André Nichterlein, and Rolf Niedermeier. Finding points in general position. *Internat. J. Comput. Geom. Appl.*, 27(4):277–296, 2017.
24. Paul W. Goldberg and Alexandros Hollender. The Hairy Ball problem is PPAD-complete. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, pages 65:1–65:14, 2019.

25 Alfredo Hubard and Roman Karasev. Bisecting measures with hyperplane arrangements. *Mathematical Proceedings of the Cambridge Philosophical Society*, pages 1–9, 2019.

26 David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *J. Comput. System Sci.*, 37(1):79–100, 1988.

27 Roman N. Karasev. Topological methods in combinatorial geometry. *Russian Mathematical Surveys*, 63(6):1031–1078, 2008.

28 Chi-Yuan Lo, Jiří Matoušek, and William L. Steiger. Algorithms for Ham-sandwich cuts. *Discrete Comput. Geom.*, 11:433–452, 1994.

29 Jesús De Loera, Xavier Goaoc, Frédéric Meunier, and Nabil Mustafa. The discrete yet ubiquitous theorems of Carathéodory, Helly, Sperner, Tucker, and Tverberg. *Bulletin of the American Mathematical Society*, 56(3):415–511, 2019.

30 Jiří Matoušek. *Using the Borsuk-Ulam theorem.* Springer-Verlag Berlin Heidelberg, 2003.

31 Frédéric Meunier, Wolfgang Mulzer, Pauline Sarrabezolles, and Yannik Stein. The rainbow at the end of the line: A PPAD formulation of the Colorful Carathéodory theorem with applications. In *Proc. 28th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 1342–1351, 2017.

32 Wolfgang Mulzer and Yannik Stein. Computational aspects of the Colorful carathéodory theorem. *Discrete Comput. Geom.*, 60(3):720–755, 2018.

33 Abraham Othman, Christos H. Papadimitriou, and Aviad Rubinstein. The complexity of fairness through equilibrium. *ACM Trans. Economics and Comput.*, 4(4):20:1–20:19, 2016.

34 Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. System Sci.*, 48(3):498–532, 1994.

35 Richard Rado. A theorem on general measure. *Journal of the London Mathematical Society*, s1-21(4):291–300, 1946.

36 Karanbir S. Sarkaria. Tverberg's theorem via number fields. *Israel Journal of Mathematics*, 79(2-3):317–320, 1992.

37 Patrick Schnider. Equipartitions with wedges and cones. *CoRR*, abs/1910.13352, 2019.

38 Patrick Schnider. Ham-sandwich cuts and center transversals in subspaces. In *Proc. 35th Int. Sympos. Comput. Geom. (SoCG)*, pages 56:1–56:15, 2019.

39 Forest W. Simmons and Francis Edward Su. Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Mathematical Social Sciences*, 45(1):15–25, 2003.

40 William Steiger and Jihui Zhao. Generalized Ham-sandwich cuts. *Discrete Comput. Geom.*, 44(3):535–545, 2010.

41 Arthur H. Stone and John W. Tukey. Generalized "Sandwich" theorems. *Duke Mathematical Journal*, 9(2):356–359, 06 1942.

42 Albert W. Tucker. Some topological properties of disk and sphere. In *Proc. First Canadian Math. Congress, Montreal, 1945*, pages 285–309. University of Toronto Press, Toronto, 1946.

43 Helge Tverberg. A generalization of Radon's theorem. *Journal of the London Mathematical Society*, s1-41(1):123–128, 1966.

44 Rade T. Zivaljević and Siniša T. Vrećica. An extension of the Ham sandwich theorem. *Bulletin of the London Mathematical Society*, 22(2):183–186, 1990.