

Robust Vertex Enumeration for Convex Hulls in High Dimensions*

Pranjal Awasthi
pranjal.awasthi@cs.rutgers.edu

Bahman Kalantari
kalantar@cs.rutgers.edu

Yikai Zhang
yz422@cs.rutgers.edu

Abstract

The problem of computing the vertices of the convex hull of a given finite set of points in the Euclidean space is a classic and fundamental problem, studied in the context of computational geometry, linear and convex programming, machine learning and more. In this article we present *All Vertex Triangle Algorithm* (AVTA), a robust and efficient algorithm that for a given input set $S = \{v_i \in \mathbb{R}^m : i = 1, \dots, n\}$ computes the subset \bar{S} of all K vertices of the convex hull of S . If desired AVTA computes an approximation to \bar{S} and it can also work if the input data is a perturbation of S . Let R be the diameter of S . We say $\text{conv}(S)$, the convex hull of S , is Γ_* -robust if the minimum of the distances from each vertex to the convex hull of the remaining vertices is Γ_* . Given $\gamma \leq \gamma_* = \Gamma_*/R$, the number of operations of AVTA to compute \bar{S} is $O(nK(m + \gamma^{-2}))$. Even without the knowledge of γ_* , but when K is known, using binary search, the complexity of AVTA is $O(nK(m + \gamma_*^{-2})) \log(\gamma_*^{-1})$. More generally, without the knowledge of γ_* or K , given any $t \in (0, 1)$, AVTA computes a subset \bar{S}^t of \bar{S} of cardinality $K^{(t)}$ in $O(nK^{(t)}(m + t^{-2}))$ operations so that the Euclidean distance between any point $p \in \text{conv}(S)$ to $\text{conv}(\bar{S}^t)$ is at most tR .

Next we consider AVTA under perturbation since in practice the input maybe a perturbation of S , $S_\varepsilon = \{v_i^\varepsilon : i = 1, \dots, n\}$, where $\|v_i - v_i^\varepsilon\| \leq \varepsilon R$. The set of perturbed vertices, \bar{S}_ε may differ drastically from the set of vertices of $\text{conv}(S_\varepsilon)$. Let Σ_* be the minimum of distances of vertices of $\text{conv}(S)$ to the convex hull of the remaining point of S . Under the assumption that $\sigma_* = \Sigma_*/R \geq 4\varepsilon$, given σ satisfying $4\varepsilon \leq \sigma \leq \sigma_*$, AVTA computes \bar{S}_ε in $O(nK_\varepsilon(m + \sigma^{-2}))$, where $K \leq K_\varepsilon \leq n$. When only K is known, but assuming $4\varepsilon \leq \sigma_*$, using binary search the complexity of AVTA to compute \bar{S}_ε is $O(nK(m + \sigma_*^{-2})) \log(\sigma_*^{-1})$. More generally, given any $t \in (0, 1)$, AVTA computes a subset \bar{S}_ε^t of \bar{S}_ε of cardinality $K_\varepsilon^{(t)}$ in $O(nK_\varepsilon^{(t)}(m + t^{-2}))$ so that the distance between any point $p \in \text{conv}(S)$ to $\text{conv}(\bar{S}_\varepsilon^t)$ is at most $(t + \varepsilon)R$.

We also consider the application of AVTA in the recovery of vertices through the projection of S or S_ε under a Johnson-Lindenstrauss randomized linear projection $L : \mathbb{R}^m \rightarrow \mathbb{R}^{m'}$. Denoting $U = L(S)$ and $U_\varepsilon = L(S_\varepsilon)$, by relating the robustness parameters of $\text{conv}(U)$ and $\text{conv}(U_\varepsilon)$ to those of $\text{conv}(S)$ and $\text{conv}(S_\varepsilon)$, we derive analogous complexity bounds for probabilistic computation of the vertex set of $\text{conv}(U)$ or those of $\text{conv}(U_\varepsilon)$, or an approximation to them. Finally, we apply AVTA to design new practical algorithms for two popular machine learning problems: topic modeling and non-negative matrix factorization. For topic models, our new algorithm leads to significantly better reconstruction of the topic-word matrix than state of the art approaches Arora et al. (2013); Bansal et al. (2014). Additionally, we provide a robust analysis of AVTA and empirically demonstrate that it can handle larger amounts of noise than existing methods. For non-negative matrix we show that AVTA is competitive with existing methods that are specialized for this task Arora et al. (2012a).

Keywords: Convex Hull Membership, Approximation Algorithms, Machine Learning, Linear Programming, Random Projections

*A conference version of the article will appear in the Proceedings of AISTATS 2018

1 Introduction

In this article we present *All Vertex Triangle Algorithm* (AVTA), a robust and efficient algorithm that for given input set $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$, computes the subset $\bar{S} = \{\bar{v}_1, \dots, \bar{v}_K\}$ of all vertices of $\text{conv}(S)$, the convex hull of S . More generally, given any $t \in (0, 1)$, AVTA computes a subset \bar{S}^t of \bar{S} so that the distance between any point $p \in \text{conv}(S)$ to $\text{conv}(\bar{S}^t)$ is to within a distance of tR . AVTA is also applicable if the input data is a perturbation of S .

AVTA, a *fully polynomial-time approximation scheme*, builds upon the *Triangle Algorithm* Kalantari (2015), designed to solve the *convex hull membership problem*. Specifically, given S , the Triangle Algorithm tests if a distinguished point p lies in the $\text{conv}(S)$, either by computing a point $p_\varepsilon \in \text{conv}(S)$ to within a prescribed distance to p , or a hyperplane that separates p from $\text{conv}(S)$. Before describing AVTA and its applications we wish to give an overview of the related problems and research, as well as their history, significance and connections to our work.

The convex hull membership problem is a basic problem in computational geometry and a very special case of the *convex hull problem*, see Goodman and Toth et al. (2004). Besides being a fundamental problem in computational geometry, it is a basic problem in linear programming (LP). In fact LP over the integers can be reduced to a convex hull membership problem. Furthermore, the two most famous polynomial-time LP algorithms, the ellipsoid algorithm of Khachiyan (1980) and the projective algorithm of Karmarkar (1984), are in fact explicitly or implicitly designed to solve the convex hull membership problem when $p = 0$, see Jin and Kalantari (2006). Furthermore, using an approach suggested by Chvátal, in Jin and Kalantari (2006) it can be shown that there is a direct connection between a general LP feasibility and this homogeneous case of the convex hull membership problem.

An important problem in computational geometry and machine learning is the *irredundancy problem*, the problem of computing all the vertices of $\text{conv}(S)$, see Toth et al. (2004). Clearly, any algorithm for LP feasibility can be used to solve the irredundancy problem by solving a sequence of $O(n)$ convex hull membership problems. For results that reduce the number of linear programming problems, see e.g. Clarkson (1994) and Chan (1996b). Some applications require the description of $\text{conv}(S)$ in terms of its vertices, facets and adjacencies, see Chazelle (1993). The complexity of many exact algorithms for irredundancy is exponential in terms of the dimension of the points, thus only practical in very low dimensions. On the other hand, the convex hull membership problem by itself has been studied in the context of large scale applications where simplex method or polynomial time algorithms are too expensive to run. Thus approximation schemes have been studied for the problem.

Blum et al. (2016) propose a bi-criterion algorithm based on *Nearest Neighbor Oracle*, computing a subset of vertices T satisfying two properties: i) the Hausdorff distance between $\text{conv}(T)$ and $\text{conv}(S)$ is bounded above by $(8\varepsilon^{1/3} + \varepsilon)R$ (ii) $|T| = O(K_{\text{opt}}/\varepsilon^{2/3})$. Since $T \subset S$, this implies that $\varepsilon = \Omega((K_{\text{opt}}/n)^{3/2})$. The running time of the algorithm is

$$O\left(\frac{nK_{\text{opt}}}{\varepsilon^{2/3}}\left(m + \frac{K_{\text{opt}}}{\varepsilon^{8/3}} + \frac{K_{\text{opt}}^2}{\varepsilon^{4/3}}\right)\right).$$

While there is a theoretical bound on the size of T as a polynomial in $1/\varepsilon$, it is inefficient since it uses the *Nearest Neighbor Oracle*. Indeed, in AVTA, the Triangle algorithm works as an approximate oracle which achieves great improvement in efficiency. Given that \bar{S} is γ robust and additionally γ is $\Omega(\varepsilon^{1/3})$, then we cannot use fewer than $|\bar{S}|$ vertices to give an ε approximation. This argument shows that in Blum et al. (2016) $K_{\text{opt}} = |\bar{S}|$. In a general case where γ is arbitrarily close to 0, AVTA will find all vertices in $O(nK_\varepsilon(m + \frac{1}{\varepsilon^2}))$ time. While we so far have no nontrivial bound on K_ε , it is known that $K_\varepsilon \leq n$. In this case the complexity of AVTA is $O(n^2m + n^2/\varepsilon^2)$ and Greedy clustering requires at least $O(nm/\varepsilon^2 + n/\varepsilon^{10})$ to achieve the same accuracy. It could be concluded that there exists regimes that AVTA outperforms *Greedy Clustering*. It is interesting to observe that AVTA could be used as a pre-processing algorithm for Greedy Clustering. By our analysis, AVTA only detects vertices and will not omit any of them. In case $n \gg K_\varepsilon$, we can use AVTA to delete points inside the convex hull thus reduce the size of the problem for Greedy Clustering. In summary, the two algorithms coexist.

Not only is convex hull detection a fundamental problem in computational geometry, state of the art algorithms for many machine learning problems rely on being able to solve this problem efficiently. Consider for instance the problem of non-negative matrix factorization (NMF) Lee and Seung (2001). Here, given access to a data matrix A , we want to compute non-negative, low rank matrices U and V such that $A = UV$. Although in general this problem is intractable, recent results show that under a natural *separability* assumption Donoho and Stodden (2003) such a factorization can be computed efficiently Arora et al. (2012a). The key insight in these works is that under the separability assumption, the rows of the matrix V will appear among the rows of A . Furthermore, the rows of V will be the *vertices* of the convex hull of rows of A . Hence, a fast algorithm for detecting the vertices will lead to a fast factorization algorithm as well.

A problem related to NMF is known as *topic modeling* Blei (2012). Here one is given access to a large corpus of documents, with each document represented as a long vector consisting of frequency in the document of every word in the vocabulary. This is known as the bag-of-words representation. Each document is assumed to represent a mixture of up to K hidden topics. A popular generative model for such documents is the following: For every document d , a K dimensional vector θ_d is drawn from a distribution over the simplex. Typically this distribution is the Dirichlet distribution. Then, for each word in the document, a topic is chosen according to θ_d . Finally, given a chosen topic i , a word is output according to the topic distribution vector β_i . This is known as the Latent Dirichlet Allocation (LDA) model Blei et al. (2003). The parameters of this model consist of the topic-word matrix β so that β_i defines the distribution over words for topic i . Additionally, there are hyper parameters associated with the Dirichlet distributions generating the topic distribution vector θ_d . The topic modeling problem concerns learning the topic-word matrix β and the parameters of the topic generating distribution. Similar to NMF, the problem is intractable in the worst case but can be efficiently solved under *separability* Arora et al. (2012b). In this context, the separability assumption requires that for each topic i , there exists an *anchor word* that has a non-zero probability of occurring only under topic i . Separability is an assumption that is known to hold for real world documents Arora et al. (2012b). The key component towards learning the model parameters is a fast algorithm for finding the anchor words. The algorithm of Arora et al. (2012b, 2013) uses the word-word covariance matrix and shows that under separability, the vertices of the convex hull of the rows of the matrix will correspond to the anchor words. Similarly, the work of Ding et al. (2013) shows that finding the vertices of the convex hull of the document-word matrix will also lead to detection of anchor words. Both approaches rely on the vertex detection subroutine. Furthermore, in the case of topic models, the documents are limited in size and this translates to the fact that one is given a perturbation of the set S . The goal is to use this perturbed set to approximate the original vertices \tilde{S} . Hence in this application it is crucial that the approach to finding the vertices be robust to noise.

The convex hull membership problem can be formulated as the minimization of a convex quadratic function over the unit simplex. This particular convex program finds applications in statistics, approximation theory, and machine learning, see e.g. Clarkson (2010) and Zhang (2003) who consider the analysis of a greedy algorithm for minimizing smooth convex functions over the unit simplex. The Frank-Wolfe algorithm Frank and Wolfe (1956) is a classic greedy algorithm for convex programming. When the the convex hull of a set of points does not contain the origin, the problem of computing the point in the convex hull with least norm, known as *polytope distance* is also a problem of interest. In some applications the polytope distance refers to the distance between two convex hulls, a fundamental problem in machine learning, known as SVM, see e.g. Burges (1998). Gilbert’s algorithm Gilbert (1966) for the polytope distance problem is one of the earliest known algorithms. Gärtner and Jaggi (2009) show Gilbert’s algorithm coincides with Frank-Wolfe algorithm when applied to the minimization of a convex quadratic function over a unit simplex. In this case the algorithm is known as *sparse greedy approximation*. For many results regarding the applications of the minimization of a quadratic function over a simplex, see Zhang (2003), Clarkson (2010) and Gärtner and Jaggi (2009). Clarkson (2010) analyzes the Frank-Wolfe and its variations while studying the notion of *coresets*. While the Triangle Algorithm has features that are very similar to those of Frank-Wolfe algorithm, there are other features and properties that make it an algorithm distinct from Frank-Wolfe or Gilbert’s

algorithm. To describe these differences, consider the distance between p and $\text{conv}(S)$:

$$\Delta = \min \left\{ d(p', p) \equiv \|p' - p\| : p' \in \text{conv}(S) \right\} = d(p_*, p). \quad (1)$$

Clearly, $p \notin \text{conv}(S)$, if and only if $\Delta > 0$. The goal of the convex hull membership problems (equivalently an LP feasibility) is to test feasibility, i.e. if p lies in $\text{conv}(S)$. Solving this does not require the computation of Δ when it is positive. Thus the goal of solving the convex hull membership is different from that of computing this distance Δ when positive. When $p \in \text{conv}(S)$, the analysis of complexity of the Triangle Algorithm is essentially identical with Clarkson (2010) analysis of the basic Frank-Wolfe algorithm. Gärtner and Jaggi (2009) on the other hand analyze the complexity of Gilbert's algorithm for the polytope distance problem, i.e. the approximation of Δ , however under the assumption that $\Delta > 0$. Gärtner and Jaggi (2009) do not address the case when $\Delta = 0$.

What distinguishes the Triangle Algorithm from the Frank-Wolfe and Gilbert's algorithms is the *distance dualities* which gives more flexibility to the algorithm. The algorithm we will analyze in this article, namely AVTA, is designed to generate all vertices of $\text{conv}(S)$. It makes repeated use of the distance dualities of the Triangle Algorithm, resulting in an over all efficient algorithm for computing the vertices of $\text{conv}(S)$, or very good approximation to these vertices, even under perturbation of the input set. Indeed AVTA is testimonial to the uniqueness of the Triangle Algorithm while itself is a nontrivial extension of the Triangle Algorithm. AVTA finds many applications in computational geometry and machine learning. Some of these are demonstrated here theoretically and computationally. We next describe AVTA in more detail.

To describes the complexities of AVTA we need to define some parameters. We say $\text{conv}(S)$ is Γ_* -robust, if Γ_* is the minimum of the distances from each $\bar{v}_i \in \bar{S}$ to $\text{conv}(\bar{S} \setminus \{\bar{v}_i\})$. Set $R = \max\{d(v_i, v_j), v_i, v_j \in S\}$, the diameter of S . AVTA works as follows.

- (1) If a number $0 < \gamma \leq \Gamma_*/R$ is known, the number of operations of AVTA to computes \bar{S} is.

$$O(nK(m + \gamma^{-2})). \quad (2)$$

- (2) If only K is known, the number of operations of AVTA to compute \bar{S} is

$$O(nK(m + \gamma_*^{-2})) \log \gamma_*^{-1}. \quad (3)$$

- (3) More generally, given any $t \in (0, 1)$, AVTA can compute a subset \bar{S}^t of \bar{S} so that the distance of each point in $\text{conv}(S)$ to $\text{conv}(\bar{S}^t)$ is at most tR . The corresponding number of operations is

$$O(nK^{(t)}(m + t^{-2})), \quad K^{(t)} = |\bar{S}^t|. \quad (4)$$

In practice the input set may be not S but a perturbation of it, $S_\varepsilon = \{v_1^\varepsilon, \dots, v_n^\varepsilon\}$, where $\|v_i - v_i^\varepsilon\| \leq \varepsilon R_S$. The set of perturbed vertices, $\bar{S}_\varepsilon = \{\bar{v}_1^\varepsilon, \dots, \bar{v}_K^\varepsilon\}$ may differ considerably from the set of actual vertices of $\text{conv}(S_\varepsilon)$. Under mild assumption on ε , AVTA computes $\bar{S}_\varepsilon = \{\bar{v}_1^\varepsilon, \dots, \bar{v}_K^\varepsilon\}$. More generally, given any $t \in (0, 1)$, AVTA computes a subset \bar{S}_ε^t of \bar{S}_ε so that the distance from any $p \in \text{conv}(S)$ to $\text{conv}(\bar{S}_\varepsilon^t)$ is at most $(t + \varepsilon)R$. The complexity of AVTA for this variation of the problem is analogous to the unperturbed case, however it makes use a weaker parameter. We say $\text{conv}(S)$ is Σ_* -weakly robust, if Σ_* is the minimum of the distances of each vertex in S to the convex hull of all the remaining points in S . In Figure 1 we show a simple example where Γ_* and Σ_* are shown for set of eight points.

We first prove when $\sigma_* = \Sigma_*/R \geq 4\varepsilon$, \bar{S}_ε is a subset of vertices of $\text{conv}(S_\varepsilon)$ and $\text{conv}(S_\varepsilon)$ is at least $\Sigma_*/2$ -weakly robust. Using this, we prove

- (i) If $\sigma \leq \sigma_* = \Sigma_*/R$ is known to satisfying $4\varepsilon \leq \sigma$, the number of operations of AVTA to computes \bar{S}_ε , is.

$$O(nK_\varepsilon(m + \sigma^{-2})), \quad (5)$$

where K_ε is at most the cardinality of the set of vertices of S_ε .

Clearly $\Gamma_* \geq \Sigma_*$, however we prove

$$\Sigma_* \geq .5\Gamma_*\rho_*, \quad (6)$$




Figure 1: $\Gamma_* = \Gamma_1$ and $\Sigma_* = \Sigma_2$.

where ρ_* is the minimum distance between distinct pair of points in S . This allows deriving lower bound to Σ_* from a known lower bound on Γ_* . Thus we can alternatively write

(ii) If $\gamma \leq \gamma_* = \Gamma_*/R$ is known satisfying $4\varepsilon \leq \gamma\rho_*/R$, the number of operations of *AVTA* to compute \bar{S}_ε is.

$$O(nK_\varepsilon(m + (\gamma\rho_*)^{-2})). \quad (7)$$

(iii) If only K is known, where $4\varepsilon \leq \sigma_* = \Sigma_*/R$, the number of operations of *AVTA* to compute \bar{S}_ε is.

$$O(nK_\varepsilon(m + \sigma_*^{-2})) \log(\sigma_*^{-1}). \quad (8)$$

(iv) More generally, given any $t \in (0, 1)$, *AVTA* can compute a subset \bar{S}_ε^t of \bar{S}_ε so that the distance from each p in $\text{conv}(S)$ to $\text{conv}(\bar{S}_\varepsilon^t)$ is at most $(t + \varepsilon)R$. The corresponding number of operations is

$$O(nK_\varepsilon^t(m + t^{-2})), \quad K_\varepsilon^t = |\bar{S}_\varepsilon^t|. \quad (9)$$

We also consider the application of *AVTA* in the recovery of vertices through the projection of S or S_ε under a Johnson-Lindenstrauss randomized linear projection $L : \mathbb{R}^m \rightarrow \mathbb{R}^{m'}$. By relating the robustness parameters of $\text{conv}(U)$ and $\text{conv}(U_\varepsilon)$, where $U = L(S)$ and $U_\varepsilon = L(S_\varepsilon)$, to those of $\text{conv}(S)$ and $\text{conv}(S_\varepsilon)$, we derive analogous complexity bounds for probabilistic computation of the vertex set of $\text{conv}(U)$ or those of $\text{conv}(U_\varepsilon)$, or an approximation to these subsets for a given $t \in (0, 1)$. Table 1 summarizes the complexities of computing desired sets under various cases.

The organization of the the article is as follows. In Section 2, we review the Triangle Algorithm for solving the convex hull membership problem. In Section 3, we describe an efficient implementation of the Triangle Algorithm. This will be used throughout the the article. In Section 4, we describe *All Vertex Triangle Algorithm* (*AVTA*), a modification of the Triangle Algorithm, for computing all vertices of the convex hull of a given finite set of points, S . We discuss several applications of this, in particular in solving the convex hull membership problem itself. Other applications will be described in subsequent sections. In Section 5, we consider the performance of *AVTA* under perturbation of data. In Section 6, we consider *AVTA* with Johnson-Lindenstrauss projections. Furthermore, we consider the performance of *AVTA* under perturbation of data with Johnson-Lindenstrauss projections.

2 Review of The Triangle Algorithm

The *Triangle Algorithm* described in Kalantari (2015) is a simple iterative algorithm for solving the *convex hull membership problem*, a fundamental problem in linear programming and computational geometry.

Input and Description	Computed via AVTA	Conditions	Complexity
$S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$ $R = \max\{\ v_i - v_j\ : v_i, v_j \in S\}$	\bar{S} , vertices of $\text{conv}(S)$, $ \bar{S} = K$ $\bar{S} = \{\bar{v}_1, \dots, \bar{v}_K\}$ Given $t \in (0, 1)$, $\bar{S}^t \subset \bar{S}$, $ \bar{S}^t = K^{(t)}$ $d(p, \text{conv}(\bar{S}^t)) \leq tR, \forall p \in \text{conv}(S)$	$\gamma \leq \gamma_* \equiv \Gamma_*/R$ is known Only K is known General Case	$O(nK(m + \gamma^{-2}))$ $O(nK(m + \gamma_*^{-2})) \times \log(\gamma_*^{-1})$ $O(nK^{(t)}(m + t^{-2}))$
$S_\varepsilon = \{v_1^\varepsilon, \dots, v_n^\varepsilon\}$, a perturbation of S $\ v_i^\varepsilon - v_i\ \leq \varepsilon R$	\hat{S}_ε , vertices in $\text{conv}(S_\varepsilon)$, $ \hat{S}_\varepsilon = K_\varepsilon$, $\bar{S}_\varepsilon = \{\bar{v}_1^\varepsilon, \dots, \bar{v}_{K_\varepsilon}^\varepsilon\} \subset \hat{S}_\varepsilon$ Given $t \in (0, 1)$, $\bar{S}_\varepsilon^t \subset \bar{S}_\varepsilon$, $ \bar{S}_\varepsilon^t = K_\varepsilon^{(t)}$ $d(p, \text{conv}(\bar{S}_\varepsilon^t)) \leq (t + \varepsilon)R, \forall p \in \text{conv}(S)$	$\sigma \leq \sigma_* \equiv \Sigma_*/R$ is known, $\varepsilon \leq \sigma/4$ $\gamma \leq \gamma_*$ is known, $\varepsilon \leq \gamma\rho_*/4R$ Only K is known, $\varepsilon \leq \sigma_*/4$ General Case	$O(nK_\varepsilon(m + \sigma^{-2}))$ $O(nK_\varepsilon(m + R^2/(\gamma\rho_*^2)))$ $O(nK_\varepsilon(m + \sigma_*^{-2})) \times \log(\sigma_*^{-1})$ $O(nK_\varepsilon^{(t)}(m + t^{-2}))$
J-L Projection of S $U = L(S) = \{u_1, \dots, u_n\}$ $u_i = L(v_i)$, $L : \mathbb{R}^m \rightarrow \mathbb{R}^{m'}$ R' diameter of U	\bar{U} , vertices of $\text{conv}(U)$, $ \bar{U} = K_{\varepsilon'}$ $\bar{U} = \{\bar{u}_1, \dots, \bar{u}_{K_{\varepsilon'}}\}$, $\bar{U} \subset L(\bar{S})$ Given $t \in (0, 1)$, $\bar{U}^t \subset \bar{U}$, $ \bar{U}^t = K_{\varepsilon'}^t$ $d(q, \text{conv}(\bar{U}^t)) \leq tR', \forall q \in \text{conv}(U)$	$\gamma \leq \gamma_*$ is known $m' = \varepsilon'^2/c \log n < m$, c a constant General Case	$O(nK_{\varepsilon'}(m' + (\gamma(1 - \varepsilon'))^{-2}))$ $O(nK_{\varepsilon'}^t(m' + t^{-2}))$
J-L Projection of S_ε $U_\varepsilon = L(S_\varepsilon) = \{u_1^\varepsilon, \dots, u_n^\varepsilon\}$ $u_i^\varepsilon = L(v_i^\varepsilon)$	\hat{U}_ε , vertices in $\text{conv}(U_\varepsilon)$, $ \hat{U}_\varepsilon = K_{\varepsilon\varepsilon'}$, $\bar{U}_\varepsilon = \{\bar{u}_1^\varepsilon, \dots, \bar{u}_{K_{\varepsilon\varepsilon'}}^\varepsilon\}$, $\bar{U}_\varepsilon \subset \hat{U}_\varepsilon$ Given $t \in (0, 1)$, $\bar{U}_\varepsilon^t \subset \bar{U}_\varepsilon$, $ \bar{U}_\varepsilon^t = K_{\varepsilon\varepsilon'}^{(t)}$ $d(q, \text{conv}(\bar{U}_\varepsilon^t)) \leq (t + \varepsilon)R, \forall q \in \text{conv}(U)$	$\sigma \leq \sigma_*$ is known, $\varepsilon \leq \sigma(1 - \varepsilon')/4$ $\gamma \leq \gamma_*$ is known, $\varepsilon \leq \gamma\rho_*(1 - \varepsilon')/4R$ General Case	$O(nK_{\varepsilon\varepsilon'}(m' + (\sigma(1 - \varepsilon'))^{-2}))$ $O(nK_{\varepsilon\varepsilon'}(m' + R^2/(\gamma\rho_*(1 - \varepsilon')^2)))$ $O(nK_{\varepsilon\varepsilon'}^{(t)}(m' + t^{-2}))$

Table 1: $\Gamma_* = \min\{d(\bar{v}_i, \text{conv}(\bar{S} \setminus \{\bar{v}_i\}))\}$, $\Sigma_* = \min\{d(\bar{v}_i, \text{conv}(S \setminus \{\bar{v}_i\}))\}$, $\rho_* = \min\{d(v_i, v_j), i \neq j\}$.

Formally, the convex hull membership problem is as follows: Given a set of point $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$ and a distinguished point $p \in \mathbb{R}^m$, test if $p \in \text{conv}(S)$. If $p \notin \text{conv}(S)$, find a hyperplane that separates p from $\text{conv}(S)$. If $p \in \text{conv}(S)$, the Triangle Algorithm solves the problem to within prescribed precision by generating a sequence of points inside of $\text{conv}(S)$ that get sufficiently close to p .

Given two point $u, v \in \mathbb{R}^m$ we interchangeably use $d(u, v) = \|u - v\|$. Given a point in $\text{conv}(S)$, the Triangle Algorithm searches for a *pivot* to get closer to p :

Definition 1. Given $p' \in \text{conv}(S)$, called *iterate*, we call $v \in S$ a *p-pivot* (or simply *pivot*) if

$$d(p', v) \geq d(p, v). \quad (10)$$

Equivalently, v is a pivot if and only if

$$v^T p - v^T p' \geq \frac{1}{2}(\|p\|^2 - \|p'\|^2). \quad (11)$$

Definition 2. A point $p' \in \text{conv}(S)$ is a *p-witness* (or simply *witness*) if the orthogonal bisecting hyperplane to the line segment pp' separates p from $\text{conv}(S)$.

Equivalently, $p' \in \text{conv}(S)$ is a *p-witness* if and only if

$$d(p', v_i) < d(p, v_i), \quad \forall i = 1, \dots, n. \quad (12)$$

Definition 3. Given $\varepsilon \in (0, 1)$, $p' \in \text{conv}(S)$ is an ε -approximate solution if for some $v \in S$,

$$d(p', p) \leq \varepsilon R, \quad (13)$$

where R is the diameter of S :

$$R = \max\{d(v_i, v_j) : v_i, v_j \in S\}. \quad (14)$$

Given a point $p' \in \text{conv}(S)$ that is neither an ε -approximate solution nor a witness, the Triangle Algorithm finds a *p-pivot* $v \in S$. Then on the line segment $p'v$ it compute the closest point to p , denoted by $\text{Nearest}(p; p'v)$. It then replaces p' with $\text{Nearest}(p; p'v)$ and repeats the process. It is easy to show,

Proposition 1. If an iterate $p' \in \text{conv}(S)$ satisfies $d(p', p) \leq \min\{d(p, v_i) : i = 1, \dots, n\}$, and v_j is a p -pivot, then the new iterate

$$p'' = \text{Nearest}(p; p'v_j) = (1 - \alpha)p' + \alpha v_j, \quad (15)$$

where the step-size α is

$$\alpha = \frac{(p - p')^T (v_j - p')}{d^2(v_j, p')} = \frac{p^T v_j - p'^T v_j - p^T p' + \|p'\|^2}{\|v_j\|^2 - 2p'^T v_j + \|p'\|^2}. \quad (16)$$

In particular if

$$p' = \sum_{i=1}^n \alpha_i v_i, \quad \sum_{i=1}^n \alpha_i = 1, \quad \alpha_i \geq 0, \quad \forall i, \quad (17)$$

then

$$p'' = \sum_{i=1}^n \alpha'_i v_i \quad (18)$$

where

$$\alpha'_j = (1 - \alpha)\alpha_j + \alpha, \quad \alpha'_i = (1 - \alpha)\alpha_i, \quad \forall i \neq j. \quad \square \quad (19)$$

The following duality ensures the correctness of the iterative step of the Triangle Algorithm.

Theorem 1. (Distance Duality) Precisely one of the two conditions is satisfied:

- (i) For each $p' \in \text{conv}(S)$ there exists $v_j \in S$ that is p -pivot, i.e. $d(p', v_j) \geq d(p, v_j)$.
- (ii) There exists $p' \in \text{conv}(S)$ that is p -witness, i.e. $d(p', v_i) < d(p, v_i)$ for all $v_i \in S$. \square

The following relates the gap in two consecutive iterates of the Triangle Algorithm:

Theorem 2. Let p, p', v be distinct points in \mathbb{R}^m . Suppose $d(p, p') \leq d(p, v) \leq d(p', v)$. Let $p'' = \text{Nearest}(p, p'v)$. Let $\delta = d(p', p)$, $\delta' = d(p'', p)$, and $r = d(p, v)$. Then,

$$\delta' \leq \delta \sqrt{1 - \frac{\delta^2}{4r^2}}. \quad \square \quad (20)$$




Figure 2: Reduction of gaps $\delta = \|p' - p\|$ by using a p -pivot v , $\delta' = \|p'' - p\|$.

The following gives the aggregate complexity bound.

Theorem 3. The Triangle Algorithm correctly solves the convex hull membership problem as follows:

- (i) Suppose $p \in \text{conv}(S)$. Given $\varepsilon > 0$, the number of iterations K_ε to compute $p_\varepsilon \in \text{conv}(S)$ so that $d(p, p_\varepsilon) \leq \varepsilon d(p, v_i)$, for some $v_i \in S$ is

$$K_\varepsilon \leq \frac{48}{\varepsilon^2} = O\left(\frac{1}{\varepsilon^2}\right). \quad (21)$$

- (ii) Suppose $p \notin \text{conv}(S)$. Let $R = \max\{d(v_i, v_j) : v_i, v_j \in S\}$, $\Delta = \min\{d(x, p) : x \in \text{conv}(S)\}$. The number of iterations K_Δ to compute $p_\Delta \in \text{conv}(S)$ so that $d(p_\Delta, v_i) < d(p, v_i)$ for all $v_i \in S$, satisfies

$$K_\Delta \leq \frac{48R^2}{\Delta^2} = O\left(\frac{R^2}{\Delta^2}\right). \quad \square \quad (22)$$

The straightforward implementation of each iteration of the Triangle Algorithm is easily seen to take $O(mn)$ arithmetic operations. The algorithm can be described as follows:

Triangle Algorithm ($S, p, \varepsilon \in (0, 1)$)

- **Step 0.** Set $p' = \operatorname{argmin}\{d(v_i, p)\}$.
- **Step 1.** If $d(p', p) \leq \varepsilon R$, then output p' as an ε -approximate solution, stop.
- **Step 2.** If a p -pivot $v \in S$ exists, set $p' \leftarrow \operatorname{Nearest}(p; p'v)$. Goto Step 1.
- **Step 3.** Output p' as p -witness. Stop.

Remark 1. In each iteration of the Triangle Algorithm it suffices to have a representation of the iterate p' in terms of v_i 's, i.e. $p' = \sum_{i=1}^n \alpha_i v_i$, where $\sum_{i=1}^n \alpha_i = 1$, $\alpha_i \geq 0$ for all $i = 1, \dots, n$. It is not necessary to know the coordinates of p' . Rather it is enough to have an array of size n to store the vector of α_i 's. Then assuming that we have stored $p^T v_i$, $i = 1, \dots, n$, we can compute the step size α (see (16)) and p'' (the new iterate) in $O(n)$ time.

An alternate complexity bound can be stated for the Triangle Algorithm, especially when $p \in \operatorname{conv}(S)$ is well situated.

Definition 4. Given $p' \in \operatorname{conv}(S)$, $v \in S$ is a *strict* p -pivot (or simply *strict* pivot) if $\angle p'pv \geq \pi/2$.

Theorem 4. (Strict Distance Duality) Assume $p \notin S$. Then $p \in \operatorname{conv}(S)$ if and only if for each $p' \in \operatorname{conv}(S)$ there exists strict p -pivot, $v \in S$. \square

The following theorem shows that under the assumption that p is an interior point of $\operatorname{conv}(S)$ we can give an alternate complexity for the Triangle Algorithm whose number of iterations are logarithmic in $1/\varepsilon$.

Theorem 5. Suppose p is contained in a ball of radius ρ , $B_\rho(p) = \{x : d(x, p) < \rho\}$, contained in $\operatorname{conv}^\circ(S)$, the relative interior of $\operatorname{conv}(S)$. Suppose the Triangle Algorithm uses a strict pivot in each iteration. Given $\varepsilon \in (0, 1)$, the number of iterations to compute $p_\varepsilon \in \operatorname{conv}(S)$ such that $d(p, p_\varepsilon) \leq \varepsilon R$, $R = \max\{d(v_i, v_j), v_i, v_j \in S\}$ is

$$O\left(\left(\frac{R}{\rho}\right)^2 \log \frac{1}{\varepsilon}\right). \quad \square \tag{23}$$

3 Efficient Implementation of Triangle Algorithm

The worst-case complexity of each iteration in the Triangle Algorithm is $O(mn)$. Assuming that all the inner products $v_i^T v_j$ are computed the iteration complexity of Triangle Algorithm can be shown to reduce to $O(n)$. The cost of pre-computing the inner products is $O(mn^2)$. The complexity can be made more efficient. To do so it suffices to compute the inner products $v_i^T v_j$ progressively rather than pre-computing them all. Ignoring this complexity, the iteration complexity of the Triangle Algorithm reduces to $O(N)$ where $N \leq n$ is the number of points of S considered in the Triangle Algorithm. The following shows how to achieve this.

Proposition 2. Let $\hat{S} = \{\hat{v}_1, \dots, \hat{v}_N\}$ be a subset of \mathbb{R}^m . Consider testing if a given $p \in \mathbb{R}^m$ lies in $\operatorname{conv}(\hat{S})$. Suppose we have computed $\|p\|^2$, as well as $p^T \hat{v}_i$, $i = 1, \dots, N$. Suppose we have available $p' = \sum_{i=1}^N \alpha_i \hat{v}_i \in \operatorname{conv}(\hat{S})$ satisfying $d(p', p) \leq \min\{d(p, \hat{v}_i) : i = 1, \dots, N\}$. Suppose $\|p'\|^2$ is also computed. Also, suppose $p'^T \hat{v}_i$ is computed for each $i = 1, \dots, N$. Then excluding the cost of computing the entries of the $N \times N$ matrix $\widehat{M} = (\hat{v}_i^T \hat{v}_j)$, each iteration of the Triangle Algorithm can be implemented in $O(N)$ operations. More precisely,

- (i) Computation of a p -pivot \hat{v}_j at p' , if one exists, takes $O(N)$ operations.
- (ii) Given a pivot \hat{v}_j , the computation of step size α takes $O(1)$ operations.

- (iii) Computation of $\text{Nearest}(p; p'v) = p'' = (1 - \alpha)p' + \alpha\hat{v}_j = \sum_{i=1}^N \alpha'_i \hat{v}_i$ takes $O(N)$ operations.
- (iv) Computation of $\|p''\|^2$ takes $O(1)$ operations.
- (v) Computation of $p''^T \hat{v}_i$, $i = 1, \dots, N$ takes $O(N)$ operations.

Proof. The Triangle Algorithm needs to use the entries of the $N \times N$ matrix $\widehat{M} = (\hat{v}_i^T \hat{v}_j)$. However, not all entries may be needed, nor do all entries of \widehat{M} need to be computed in advance. Putting aside the complexity of computing \widehat{M} , in the following we justify the claimed complexities.

(i): From (11) and the given assumptions, to check if a particular \hat{v}_i is a pivot takes $O(1)$ operations. Thus to check if there exists a pivot takes $O(N)$ time.

(ii): From (16) and the assumptions, to compute α takes $O(1)$ operations.

(iii): From equations (18) and (19) the computation of p'' and its representation as $p'' = \sum_{i=1}^N \alpha'_i \hat{v}_i$ takes $O(N)$ operations.

(iv): Since $p'' = (1 - \alpha)p' + \alpha\hat{v}_j$, we have

$$\|p''\|^2 = p''^T p'' = (1 - \alpha)^2 \|p'\|^2 + 2\alpha(1 - \alpha)p'^T \hat{v}_j + \alpha^2 \|\hat{v}_j\|^2. \quad (24)$$

It follows that computing $\|p''\|^2$ takes $O(1)$ operations.

(v): Using that $p'' = (1 - \alpha)p' + \alpha\hat{v}_j$, the computation of $p''^T \hat{v}_i$ takes $O(1)$ computations. Hence to compute all inner products $p''^T \hat{v}_i$, $i = 1, \dots, N$ takes $O(N)$ computations. \square

The following theorem combines Theorem 3 and Proposition 2 giving an improved complexity for the Triangle Algorithm.

Theorem 6. Let $\widehat{S} = \{\hat{v}_1, \dots, \hat{v}_N\}$ be a subset of $S = \{v_1, \dots, v_n\}$. Given $p \in \mathbb{R}^m$, consider testing if $p \in \text{conv}(\widehat{S})$. Suppose $\|p\|^2$ as well as $p^T \hat{v}_i$, $i = 1, \dots, N$ are computed. Given $\varepsilon \in (0, 1)$, assume the Triangle Algorithm starts with $p' = \arg\min\{d(\hat{v}_i, p) : i = 1, \dots, N\}$. Then the complexity of testing if there exists an ε -approximate solution is

$$O\left(mN^2 + \frac{N}{\varepsilon^2}\right). \quad (25)$$

In particular, suppose in testing if $p \in \text{conv}(S)$, $S = \{v_1, \dots, v_n\}$, the Triangle Algorithm computes an ε -approximate solution p_ε by examining only the elements of a subset $\widehat{S} = \{\hat{v}_1, \dots, \hat{v}_N\}$ of S . Then the number of operations to determine if there exists an ε -approximate solution $p_\varepsilon \in \text{conv}(S)$, is as stated in (25). \square

4 All Vertex Triangle Algorithm (AVTA)

Given $S = \{v_i \in \mathbb{R}^m : i = 1, \dots, n\}$, let R be its diameter, i.e. $R = \max\{d(v_i, v_j), v_i, v_j \in S\}$. Denote the set of vertices of $\text{conv}(S)$ by

$$\overline{S} = \{\bar{v}_1, \dots, \bar{v}_K\}. \quad (26)$$

A straightforward but naive way to compute \overline{S} is to test for each v_i if it lies in $\text{conv}(S \setminus \{v_i\})$, to within an ε precision. Thus the overall this would take n times the complexity of Triangle Algorithm. This is inefficient. In what follows we describe a modification of the Triangle Algorithm with more efficient complexity than the straightforward algorithm. First we give a definition.

Definition 5. We say $\text{conv}(S)$ is Γ_* -robust if

$$\Gamma_* = \min\{d(\bar{v}_i, \text{conv}(\overline{S} \setminus \{\bar{v}_i\})) : i = 1, \dots, K\}. \quad (27)$$

As an example, given a triangle with vertices v_1, v_2, v_3 , Γ_* is the minimum of the distances from each vertex to the line segment determined by the other vertices. Thus if other points are placed inside the triangle Γ_* will not be affected.

The following is immediate from Definition 5.

Proposition 3. Let $\hat{S} = \{\hat{v}_1, \dots, \hat{v}_N\}$ be a subset of \bar{S} . Suppose $\text{conv}(S)$ is Γ_* -robust. Given $v \in S \setminus \hat{S}$, if for some $\gamma \leq \gamma_* \equiv \Gamma_*/R$ we have $d(v, \text{conv}(\hat{S})) < \gamma R$, then $v \notin \bar{S}$. \square

Theorem 7. Let $\hat{S} = \{\hat{v}_1, \dots, \hat{v}_N\}$ be a subset of \bar{S} . Given $\gamma \in (0, 1)$, consider testing if a given $v \in S \setminus \hat{S}$ satisfies $d(v, \text{conv}(\hat{S})) \leq \gamma R/2$. Suppose we are given $p' \in \text{conv}(\hat{S})$ for which $\|p'\|^2$ as well as $p'^T \hat{v}_i$, $i = 1, \dots, N$ are computed. Then the number of operations to check if $d(v, \text{conv}(\hat{S})) \leq \gamma R/2$ satisfies

$$O\left(mK^2 + \frac{K}{\gamma^2}\right). \quad (28)$$

Proof. Proof is immediate from Theorem 6 and the fact that $N \leq K$. \square

We now describe a modification of the Triangle Algorithm for computing all vertices of $\text{conv}(S)$. We call this *All Vertex Triangle Algorithm* or simply AVTA. Assume $\text{conv}(S)$ is Γ_* -robust, where Γ_* may or may not be available. However, assume we have a constant $\gamma \in (0, 1)$ known to satisfy $\gamma \leq \gamma_* = \Gamma_*/R$. AVTA works as follows. Given a working subset \hat{S} of \bar{S} , initially of cardinality $N = 1$ (see Proposition 4), a single vertex of S , it randomly selects $v \in S \setminus \hat{S}$. It then tests via the Triangle Algorithm if $d(v, \text{conv}(\hat{S})) \leq \gamma R/2$. If so, it discards v since by definition of γ it cannot belong to \bar{S} (see Proposition 3). Otherwise, it computes a v -witness $p' \in \text{conv}(\hat{S})$. It then sets $c' = v - p'$ and maximizes $c'^T x$ where x ranges in $\text{conv}(S \setminus \hat{S})$. The maximum value coincides with the maximum of $c'^T v_i$ where v_i ranges in $S \setminus \hat{S}$. If the set of optimal solution $S \setminus \hat{S}$ is denoted by S' , then $\text{conv}(S')$ is a face of $\text{conv}(S)$. A vertex v' of $\text{conv}(S')$ is a point in S' and is necessarily a vertex of $\text{conv}(S)$. Such a vertex can be computed efficiently. Having computed a new vertex v' of $\text{conv}(S)$, AVTA replaces \hat{S} with $\hat{S} \cup \{v'\}$ and the process is repeated. However, if v coincides with v' AVTA selects a new point in $S \setminus \hat{S}$. Otherwise, AVTA continues to test if the same v (for which a witness was found) is within a distance of $\gamma R/2$ of the convex hull of the augmented set \hat{S} . Also, as an iterate AVTA uses the same witness p' . In doing so each selected $v \in S$ is either determined to be a vertex itself, or it will continue to be tested if it lies to within a distance of $\gamma R/2$ of the growing set \hat{S} . If within $\gamma R/2$ distance, it will be discarded before AVTA tests another point. We will describe AVTA more precisely. However, we first prove the necessary results.

Lemma 1. Let $\hat{S} = \{\hat{v}_1, \dots, \hat{v}_N\}$ be a subset of \bar{S} . For a given $v \in S \setminus \hat{S}$ suppose $p' \in \text{conv}(\hat{S})$ is a v -witness. Let $c' = v - p'$. Then

$$\max\{c'^T x : x \in \text{conv}(S \setminus \hat{S})\} = \max\{c'^T v_i : v_i \in S \setminus \hat{S}\}. \quad (29)$$

Proof. Each $x \in \text{conv}(S \setminus \hat{S})$ can be written as a convex combination

$$x = \sum_{i: v_i \in S \setminus \hat{S}} \alpha_i v_i, \quad \sum_{i: v_i \in S \setminus \hat{S}} \alpha_i = 1, \quad \alpha_i \geq 0. \quad (30)$$

Then

$$c'^T x = \sum_{i: v_i \in S \setminus \hat{S}} \alpha_i c'^T v_i, \quad \sum_{i: v_i \in S \setminus \hat{S}} \alpha_i = 1, \quad \alpha_i \geq 0. \quad (31)$$

It follows that the maximum of $c'^T x$ over $S \setminus \hat{S}$ can be computed trivially. \square

Corollary 1. Let $c' = v - p'$ be as in Lemma 1, $p' \in \text{conv}(\hat{S})$ for which $\|p'\|^2$ as well as $p'^T \hat{v}_i$, $i = 1, \dots, N$ are computed. Then, $\max\{c'^T x : x \in \text{conv}(S \setminus \hat{S})\}$ can be computed in $O(nK)$ operations.

Proof. Since $N \leq K$, for each i , $c'^T v_i$ can be computed in $O(K)$ operations. \square

Theorem 8. Let S' be the set of optimal solutions of $\max\{c'^T x : x \in S \setminus \hat{S}\}$. Let $v' \in S'$ be a vertex of $\text{conv}(S')$. Then v' is a vertex of $\text{conv}(S)$, i.e. $v \in \bar{S} = \{\bar{v}_1, \dots, \bar{v}_K\}$.

Proof. We can write v' as a convex combination of \bar{v}_i , $i = 1, \dots, K$:

$$v' = \sum_{i=1}^K \alpha_i \bar{v}_i, \quad \sum_{i=1}^K \alpha_i = 1, \quad \alpha_i \geq 0, \quad \forall i. \quad (32)$$

The above can be rewritten as

$$v' = \sum_{i: \bar{v}_i \in S'} \alpha_i \bar{v}_i + \sum_{i: \bar{v}_i \notin S'} \alpha_i \bar{v}_i. \quad (33)$$

Since for $v_i \in S'$, $c'^T v' = c'^T v_i$ and for $v_i \notin S'$, $c'^T v' > c'^T v_i$, it follows that v' is a convex combination of \bar{v}_i for which $c'^T v' = c'^T \bar{v}_i$. But since v' is a vertex of $\text{conv}(S')$ it follows that $v' \in \bar{S}$. \square

The following shows computing a single vertex of $\text{conv}(S)$ is trivial.

Proposition 4. *Given any v in S , let $\text{Farthest}(v, S)$ return a point in S that is farthest from v . Then $\text{Farthest}(v, S)$ is a vertex of $\text{conv}(S)$, hence a member of \bar{S} .*

Proof. If $\text{Farthest}(v, S)$ is not a vertex of $\text{conv}(S)$ it can be written as a convex combination of two other points $v_1, v_2 \in \text{conv}(S)$. But then this gives a contradiction by considering the triangle with vertices $v_1, v_2, \text{Farthest}(v, S)$. \square

When $n \geq 2$, $\text{Farthest}(v, S)$ for all $v \in S$ results in a set with at least two points but it may also contain exactly two points. It can be computed in $O(n^2)$ time. Next we describe AVTA for computing all vertices of $\text{conv}(S)$.

AVTA ($S, \gamma \in (0, 1)$)

- **Step 0.** Set $\hat{S} = \{\text{Farthest}(v, S)\}$ for some $v \in S$.
- **Step 1.** Randomly select $v \in S \setminus \hat{S}$.
- **Step 2.** Call **Triangle Algorithm** ($\hat{S}, v, \gamma/2$).
- **Step 3.** If the output p' of Step 2 is a v -witness then Goto Step 4. Otherwise, p' is a $\gamma/2$ -approximate solution to v . Set $S \leftarrow S \setminus \{v\}$. If $S = \emptyset$, stop. Otherwise, Goto Step 1.
- **Step 4.** Let $c' = v - p'$.
Compute S' , the set of optimal solutions of $\max\{c'^T x : x \in S \setminus \hat{S}\}$. Randomly select $v' \in S'$.
 $v' \leftarrow \text{Farthest}(v', S')$, $\hat{S} \leftarrow \hat{S} \cup \{v'\}$.
- **Step 5.** If $v = v'$, Goto Step 1. Otherwise, Goto Step 2.

Remark 2. Here we make remarks about the steps of AVTA. In Step 0 AVTA selects the first vertex. In Step 1 it randomly select a v in $S \setminus \hat{S}$. In Step 2 AVTA checks if the point v selected in Step 1 is sufficiently close to the convex hull of the current set of vertices, \hat{S} . If so, in Step 3 v is discarded from further considerations. Otherwise, a v -witness p' is at hand. Step 4 then uses this witness to compute a direction, $c' = v - p'$, where the maximization of $c'^T x$ gives a subset S' of S consisting of the optimal solutions. Then a vertex $\text{conv}(S')$ will necessarily be a vertex of $\text{conv}(S)$. A vertex of $\text{conv}(S')$ is selected by choosing an arbitrary $v' \in S'$ and computing its farthest point in S' . It maybe the case that the vertex v' found in Step 4 coincides with v . Step 5 checks if $v' = v$ in which case it select a new v in the updated $S \setminus \hat{S}$ in Step 1 for consideration. Otherwise, when this new vertex v' is not v itself, in Step 5 in AVTA v is sent back to Step 2 to be reexamined if v is within $\gamma R/2$ distance of the convex hull of augmented \hat{S} .

Example 1. We consider an example of AVTA, see Figure 3. In this example $S = \{v_1, \dots, v_{11}\}$. Note that the set of vertices is $\bar{S} = \{v_4, v_{10}, v_6, v_1, v_9, v_2, v_5, v_8\}$. Suppose the current working subset of vertices

\bar{S} consists of $\hat{S} = \{v_1, v_9, v_2, v_5\}$ and $v = v_3$ is randomly selected to be tested if it lies in $\text{conv}(\hat{S})$. A witness $p' \in \text{conv}(\hat{S})$ is computed and with $c' = p' - v$ maximum of $c'^T x$ over $\text{conv}(\bar{S} \setminus \hat{S})$ is attained at $S' = \{v_4, v_7, v_{10}\}$. Subsequently one of the two points v_4 or v_{10} will become the next vertex to be placed in \hat{S} .




Figure 3: An example where $\hat{S} = \{v_1, v_9, v_2, v_5\}$, then $v = v_3$ is randomly selected from $S \setminus \hat{S}$ and is tested if it lies in $\text{conv}(\hat{S})$. A witness p' is found. Then using $c' = v - p'$ the set $S' = \{v_4, v_7, v_{10}\}$ is computed and one of vertices of $\text{conv}(S')$, i.e. v_4 or v_{10} is selected for inclusion in \hat{S} .

The following theorem is one of the main results:

Theorem 9. Let $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$. Let R be the diameter of S . Let $\bar{S} = \{\bar{v}_1, \dots, \bar{v}_K\}$ be the set of vertices of $\text{conv}(S)$. Suppose that $\text{conv}(S)$ is Γ_* -robust. Let $\gamma_* = \Gamma_*/R$.

(1) If a number γ is known satisfying $0 < \gamma \leq \gamma_*$, the number of arithmetic operations to compute \bar{S} is

$$O\left(nK\left(m + \frac{1}{\gamma^2}\right)\right). \quad (34)$$

(2) If only K is known, the complexity of computing \bar{S} is

$$O\left(\left(nK\left(m + \frac{1}{\gamma_*^2}\right)\right) \log \frac{1}{\gamma_*}\right). \quad (35)$$

(3) More generally, given any prescribed $t \in (0, 1)$ in

$$O\left(nK^{(t)}\left(m + \frac{1}{t^2}\right)\right) \quad (36)$$

operations AVTA computes a subset \bar{S}^t of \bar{S} of size $K^{(t)}$ so that the distance from each p in $\text{conv}(S)$ to $\text{conv}(\bar{S}^t)$ is at most tR .

Proof. (1): Initially in AVTA the subset \hat{S} consists of a single element of \bar{S} . It continues to grow until it reaches \bar{S} . By Theorem 6 for each $v \in S \setminus \hat{S}$ the cost of Step 2 in AVTA is $O(mK^2 + K/\gamma^2)$. The needed inner

products in Step 2 are $\widehat{v}_i^T \widehat{v}_j$. However, these inner products need to be computed only once and since there are at most K of \widehat{v}_i 's, these inner products can be computed at the cost of $O(mK^2)$ operations. We can store the values of the inner products in an array. Then we use them again as they arise in subsequent iterations. This kind of storing can be done for other inner products that may need to be computed in the course of the algorithm. When a selected v is within the distance of $\gamma R/2$ to $\text{conv}(\widehat{S})$, Step 3 eliminates it from further considerations. If v is not eliminated, it either gives rise to a new vertex $v' \in \overline{S}$, or v is a vertex itself. In either case, in order to identify a new vertex of \overline{S} , after a witness has become available, it requires the minimization of $c'^T v_i$ as v_i ranges over current set of vertices, $S \setminus \widehat{S}$. Since $c' = v - p'$, $p' = \sum_{j=1}^N \alpha_j \widehat{v}_j$, where $N = |\widehat{S}|$, the evaluation of $c'^T v_i$ requires the computation of $v^T v_i$, and $v_i^T \widehat{v}_j$, $j = 1, \dots, N$. This requires $O(Nm)$ operations. Since such computation is only required of each vertex in \overline{S} , over all the computation of all $c'^T v_i$ requires $O((n - K)mK) = O(nmK)$ operations. These together with Theorem 6 imply that the over all complexity is $O(mK^2 + nmK + nK/\gamma^2)$ which is the claimed complexities in (1).

(2): When only K is known, we execute AVTA, first selecting $\gamma = .5$. If we compute K vertices with this estimate of $\gamma_* = \Gamma_*/R$, we stop. Otherwise, we halve γ and repeat the process. Eventually in $O(\log(\gamma_*^{-1}))$ calls to AVTA we accumulate all K vertices in \overline{S} .

(3): For each input $t \in (0, 1)$, AVTA computes a subset \overline{S}^t of \overline{S} with $K^{(t)}$ elements. The proof of complexity is analogous to the previous cases. Next we prove for each $p \in \text{conv}(S)$, the distance from p to $\text{conv}(\overline{S})$ is at most tR . We have

$$p = \sum_{i=1}^K \alpha_i \overline{v}_i, \quad \sum_{i=1}^K \alpha_i = 1, \quad \alpha_i \geq 0. \quad (37)$$

For each i let $p_i \in \text{conv}(S^t)$ be the closest point to \overline{v}_i . Now consider

$$p^t = \sum_{i=1}^K \alpha_i p_i. \quad (38)$$

Then $p^t \in \text{conv}(S^t)$. On the other hand, by the triangle inequality

$$\|p - p^t\| \leq \sum_{i=1}^K \alpha_i \|\overline{v}_i - p_i\| \leq tR \sum_{i=1}^K \alpha_i = tR. \quad (39)$$

□

Remark 3. If neither K nor an estimate γ to $\gamma_* = \Gamma_*/R$ are known, initially we select $t = 0.5$ and with this value of t compute a subset of vertices with $K^{(t)}$ elements. We can then halve t and repeat the process. Intuitively, if for two consecutive values of t no more vertices are generated we can terminate the process, or decrease t by a factor of four. If Γ_* is not too small we will produce a reasonably good subset of \overline{S} within a reasonable number of calls to AVTA. In either case we are assured of an approximation of $\text{conv}(S)$ according to (3) in Theorem 9.

4.1 Application of AVTA in Solving the Convex Hull Membership

Suppose we wish to solve the convex hull membership problem: Test if a particular point p lies in $\text{conv}(S)$, $S = \{v_1, \dots, v_n\}$. This is equivalent to linear programming and thus can be solved with variety of algorithms, including polynomial-time algorithms, the simplex method, Frank-Wolfe, or triangle Algorithm. Whichever algorithm we use, the number n plays a role in the complexity. Thus if we compute the set of vertices of $\text{conv}(S)$, \overline{S} , we can then test if p lies in $\text{conv}(\overline{S})$ with K instead of n . This approach may seem to be inefficient, however depending upon the accuracy to which we wish to solve the problem and the size of γ_* it may result in a more efficient algorithm. The next theorem considers the application of Theorem 9 in solving the convex hull membership problem.

Theorem 10. Let $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$. Let R be the diameter of S . Let $\bar{S} = \{\bar{v}_1, \dots, \bar{v}_K\}$ be the set of vertices of $\text{conv}(S)$. Suppose $\text{conv}(S)$ is Γ_* -robust. Given any $0 < \gamma \leq \Gamma_*/R$, the number of operations to test if for a given $p \in \mathbb{R}^m$ admits an ε -approximate solution is

$$O\left(nmK + \frac{nK}{\gamma^2} + \frac{K}{\varepsilon^2}\right). \quad (40)$$

Proof. To test if p admits an ε -approximate solution can be achieved by first computing the vertices in S , followed by testing if p admits an ε -approximate solution in $\text{conv}(S)$. From Theorem 8 and Theorem 3 it follows that the total complexity is as claimed. \square

Remark 4. It is easy to check that for some values of $\varepsilon < \gamma$ the computations of \bar{S} followed by testing if p lies in $\text{conv}(\bar{S})$ could be more efficient than solving the convex hull membership without computing \bar{S} . This is especially true when $K = o(n)$.

5 AVTA Under Input Perturbation

As in the previous section, we assume $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$, R the diameter of S , and $\bar{S} = \{\bar{v}_1, \dots, \bar{v}_K\}$ the set of vertices of $\text{conv}(S)$. Assume $\text{conv}(S)$ is Γ_* -robust.

As before we wish to compute \bar{S} or a reasonable subset of it. However, in practice the input set S may be not S but a perturbation of S . This changes the set of vertices, robustness parameter and more. We wish to study perturbations under which we can recover the corresponding perturbation of \bar{S} and extend AVTA to computing this perturbation.

Definition 6. For a given $\varepsilon \in (0, 1)$ the ε -perturbations of S is the set S_ε defined as

$$S_\varepsilon = \{v_1^\varepsilon, \dots, v_n^\varepsilon\}, \quad \|v_i - v_i^\varepsilon\| \leq \varepsilon R. \quad (41)$$

The ε -perturbations of \bar{S} is the set \bar{S}_ε , denoted by

$$\bar{S}_\varepsilon = \{\bar{v}_1^\varepsilon, \dots, \bar{v}_K^\varepsilon\}, \quad (42)$$

where \bar{v}_i^ε is the perturbation of \bar{v}_i .

In practice we may be given S_ε as opposed to S . The first question that arises is: What is the relationship between the vertices of S and those of S_ε ? Without any assumptions, the vertices of $\text{conv}(S_\varepsilon)$ could change drastically, even under small perturbations.

Example 2. Consider a triangle with three additional interior points, very close to its vertices. It may be the case that even under small perturbation all six points become vertices, or that the interior points become the new vertices while the vertices become the new interior points. Thus there is a need to make some assumptions before we can say anything about the nature of perturbed points.

We would hope that for appropriate range of values of ε , \bar{S}_ε would at least be a subset of the set of vertices of S_ε . First we need a definition.

Definition 7. We say $\text{conv}(S)$ is Σ_* -weakly robust if

$$\Sigma_* = \min\{d(v, \text{conv}(S \setminus \{v\})) : v \in \bar{S}\}. \quad (43)$$

Example 3. Suppose that S consists of the vertices of a non-degenerate triangle with vertices v_1, v_2, v_3 . Suppose one additional point is placed inside the triangle. Then clearly $\Sigma_* < \Gamma_*$.

More generally we have

Proposition 5. Given $S = \{v_1, \dots, v_n\}$, we have

$$\Sigma_* \leq \Gamma_*. \quad \square \quad (44)$$

Other than the inequality in Proposition 5, Σ_* and Γ_* corresponding to the set S may seem unrelated, however in the following theorem we establish a relationship between the two that is useful in the analysis of AVTA for computing \bar{S}_ε .

Theorem 11. Let S and \bar{S} be as before. Suppose $\text{conv}(S)$ is Γ_* -robust, also Σ_* -weakly robust. Let $\rho_* = \min\{d(v_i, v_j) : v_i, v_j \in S, i \neq j\}$. We have

$$\Sigma_* \geq \frac{\rho_*}{R} \Gamma_* = \rho_* \gamma_*. \quad (45)$$

Proof. For each vertex $v \in \text{conv}(S)$, let Γ_v be the distance from v to the convex hull of the remaining vertices in S . Specifically,

$$\Gamma_v = d(v, \text{conv}(\bar{S} \setminus \{v\})). \quad (46)$$

Also let Σ_v be the distance from v to the convex hull of all other points in S . Specifically,

$$\Sigma_v = d(v, \text{conv}(S \setminus \{v\})). \quad (47)$$

Clearly we have,

$$\Sigma_v \leq \Gamma_v. \quad (48)$$

Assume v is a vertex for which $\Sigma_v < \Gamma_v$. If no such a vertex exists then $\Sigma_* = \Gamma_*$ (see Figure 4). Let u be the closest point to v lying in the convex hull of the the other vertices of S . Thus

$$\Gamma_v = d(v, u), \quad u \in \text{conv}(\bar{S}). \quad (49)$$

Let H_u be the hyperplane orthogonal to the line segment vu , passing through u . By definition of u and Carathéodory's theorem u is a convex combination of vertices of $\text{conv}(S)$ lying on H_u . Thus for some subset T of \bar{S}

$$u = \sum_{\bar{v}_i \in T \subset \bar{S}} \alpha_i \bar{v}_i, \quad \sum_{\bar{v}_i \in T \subset \bar{S}} \alpha_i = 1, \quad \alpha_i \geq 0. \quad (50)$$

Figure 4 gives a depiction of this property for a simple example. In the example u is a convex combination of \bar{v} and \bar{v}' , vertices of $\text{conv}(S)$ lying in the intersection of H_u and $\text{conv}(S)$. Consider one of these vertices, say \bar{v} . Moving the hyperplane H_u parallel to itself toward v , it intersects the line segment uv at a unique point w that lies on a facet of $\text{conv}(S \setminus \{v\})$. Such w exists because $\Sigma_* < \Gamma_*$. In other words, if H_w is a hyperplane parallel to H_u passing through w , then the region of $\text{conv}(S)$ enclosed between the halfspace defined by H_w and v contains no point of S in its interior (see shaded area in Figure 4. This implies

$$\Sigma_v \geq d(v, w). \quad (51)$$

Now consider the intersection of H_w and each ray connecting v to $\bar{v}_i \in T$. Denote this intersection by y_i . In the figure the intersection of H_w and the ray connecting $v\bar{v}$ is denoted by y . By definition of w and Carathéodory's theorem there must exist a point $v_j \in S$ lying on H_w . Furthermore, v_j can be written as a convex combination of all the y_i 's. Thus may write

$$v_j = \sum_{\bar{v}_i \in T \subset \bar{S}} \beta_i y_i, \quad \sum_{\bar{v}_i \in T \subset \bar{S}} \beta_i = 1, \quad \beta_i \geq 0. \quad (52)$$

Since by definition of ρ_* , $d(v, v_j) \geq \rho_*$, at least for one y_i we must have $d(v, y_i) \geq \rho_*$. This implies we could assume \bar{v} was chosen so that the corresponding y satisfies

$$d(v, y) \geq \rho_*. \quad (53)$$

On the one hand we have

$$u - v^\varepsilon = \sum_{i=2}^n \alpha_i (v_i - v_i^\varepsilon). \quad (59)$$

Then by the triangle inequality

$$\|u - v^\varepsilon\| \leq \sum_{i=2}^n \alpha_i \|v_i - v_i^\varepsilon\| \leq \sum_{i=2}^n \alpha_i \varepsilon R = \varepsilon R. \quad (60)$$

On the other hand, v is in \bar{S} . Without loss of generality assume $v = \bar{v}_1$. From this assumption and since by (58) $u \in \text{conv}(\bar{S} \setminus \{\bar{v}_1\})$ we have

$$u = \sum_{i=2}^K \gamma_i \bar{v}_i, \quad \sum_{i=2}^K \gamma_i = 1, \quad \gamma_i \geq 0. \quad (61)$$

Since $\text{conv}(S)$ is Σ_* -weakly robust on \bar{S} and $\sigma \leq \sigma_* = \Sigma_*/R$ we have,

$$\|u - v\| \geq \sigma R. \quad (62)$$

However, from (60), the fact that $\|v - v^\varepsilon\| \leq \varepsilon R$ and the triangle inequality we may write.

$$\|u - v\| = \|u - v^\varepsilon + v^\varepsilon - v\| \leq \|u - v^\varepsilon\| + \|v^\varepsilon - v\| \leq \varepsilon R + \varepsilon R = 2\varepsilon R. \quad (63)$$

This contradicts the assumption that $2\varepsilon < \sigma$. Hence v^ε is a vertex of $\text{conv}(S_\varepsilon)$. \square

Remark 5. The theorem implies that if the input to AVTA is S_ε instead of S , AVTA will still return at least K vertices. However, the set of vertices of $\text{conv}(S_\varepsilon)$ may have more elements than K , possibly all of S_ε . Moreover, the weakly robustness parameter Σ_* will change. We thus need to revise AVTA if we wish to extract the subset $\bar{S}_\varepsilon = \{\bar{v}_1^\varepsilon, \dots, \bar{v}_K^\varepsilon\}$ from the set of vertices of $\text{conv}(S_\varepsilon)$.

In what follows we will first show how under a mild assumptions on the relationship between Σ_*/R and ε , AVTA can compute a subset \hat{S}_ε of the vertices of $\text{conv}(S_\varepsilon)$ containing \bar{S}_ε (Theorem 13). We then show how AVTA can efficiently extract from \hat{S}_ε the desired set, namely \bar{S}_ε . The next lemma establishes a lower bound on the weak-robustness of $\text{conv}(S_\varepsilon)$. It also shows how spurious vertices of $\text{conv}(S_\varepsilon)$ are situated with respect to the convex hull of the remaining vertices. This will be used in Theorem 13 in pruning such vertices.

Lemma 2. *Suppose $\text{conv}(S)$ is Σ_* -weakly robust. Suppose $\varepsilon < \Sigma_*/2R$. Let v^ε be any point in \bar{S}_ε . Let \hat{S}_ε be any subset of vertices of $\text{conv}(S_\varepsilon)$ containing \bar{S}_ε . Then,*

$$d(v^\varepsilon, \text{conv}(\hat{S}_\varepsilon)) \geq (\Sigma_* - 2\varepsilon R). \quad (64)$$

Moreover let \hat{v}^ε be any (spurious) point in $\hat{S}_\varepsilon \setminus \bar{S}_\varepsilon$. Then

$$d(\hat{v}^\varepsilon, \text{conv}(\hat{S}_\varepsilon \setminus \{\hat{v}^\varepsilon\})) \leq \varepsilon R. \quad (65)$$

Proof. By Theorem 12, \bar{S}_ε is a subset of vertices of $\text{conv}(S_\varepsilon)$. Given $v^\varepsilon \in \bar{S}_\varepsilon$, let v be the corresponding vertex in \bar{S} . Given w^ε in $\text{conv}(S_\varepsilon \setminus \{v^\varepsilon\})$, let w in $\text{conv}(S \setminus \{v\})$ be the corresponding point, i.e. defined with respect to the same convex combination of corresponding vertices. Then

$$\|v - v^\varepsilon\| \leq \varepsilon R, \quad \|w - w^\varepsilon\| \leq \varepsilon R. \quad (66)$$

From the above it is easy to show

$$|d(v, w) - d(v^\varepsilon, w^\varepsilon)| \leq 2\varepsilon R. \quad (67)$$

But this implies

$$d(v, w) - d(v^\varepsilon, w^\varepsilon) \leq 2\varepsilon R. \quad (68)$$

Equivalently,

$$d(v, w) - 2\varepsilon R \leq d(v^\varepsilon, w^\varepsilon). \quad (69)$$

But $d(v, w) \geq \sigma_* R = \Sigma_*$. This proves (64).

To prove (65), let \hat{v} be the point in S corresponding to \hat{v}^ε . We have

$$\hat{v} = \sum_{i=1}^K \alpha_i \bar{v}_i, \quad \sum_{i=1}^K \alpha_i = 1, \quad \alpha_i \geq 0. \quad (70)$$

Define

$$\hat{w} = \sum_{i=1}^K \alpha_i \bar{v}_i^\varepsilon, \quad \sum_{i=1}^K \alpha_i = 1, \quad \alpha_i \geq 0. \quad (71)$$

It is now easy to show

$$\|\hat{v}^\varepsilon - \hat{w}\| \leq \|\hat{v}^\varepsilon - \hat{v}\| + \|\hat{v} - \hat{w}\| \leq 2\varepsilon R. \quad (72)$$

This proves (65). \square

Theorem 13. Let $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$. Assume $\text{conv}(S)$ is Σ_* -weakly robust. Suppose $\varepsilon \leq \Sigma_*/4R$.

(i) Given σ satisfying, $4\varepsilon \leq \sigma \leq \sigma_* = \Sigma_*/R$, AVTA can be modified to compute a subset \hat{S}_ε of the set of vertices of S_ε containing \bar{S}_ε , then compute from this subset \bar{S}_ε itself. If K_ε is the cardinality of \hat{S}_ε , the total number of operations satisfies

$$O\left(nK_\varepsilon\left(m + \frac{1}{\sigma^2}\right)\right). \quad (73)$$

(ii) Given γ , satisfying $4\varepsilon \leq \gamma\rho_* \leq \Gamma_*\rho_*/R = \gamma_*\rho_*$, AVTA can be modified to compute a subset \hat{S}_ε of the set of vertices of S_ε containing \bar{S}_ε , then compute from this subset \bar{S}_ε itself. If K_ε is the cardinality of \hat{S}_ε the total number of operations satisfies

$$O\left(nK_\varepsilon\left(m + \frac{1}{(\rho_*\gamma)^2}\right)\right). \quad (74)$$

(iii) Given only K , where $4\varepsilon \leq \Sigma_*/R$, the number of operations of AVTA to compute \bar{S}_ε is.

$$O\left(nK_\varepsilon\left(m + \frac{1}{\sigma_*^2}\right)\log\left(\frac{1}{\sigma_*}\right)\right). \quad (75)$$

(iv) More generally, given any $t \in (0, 1)$, AVTA can be modified to compute a subset \bar{S}_ε^t of the set of vertices of $\text{conv}(S_\varepsilon)$ of cardinality $K_\varepsilon^{(t)}$ so that the distance from each point in $\text{conv}(S_\varepsilon)$ to $\text{conv}(\bar{S}_\varepsilon^t)$ is at most t . In particular, the distance from each point in $\text{conv}(S)$ to $\text{conv}(S_\varepsilon^t)$ is at most $(t + \varepsilon)R$. The complexity of the computation of \bar{S}_ε^t is

$$O\left(nK_\varepsilon^{(t)}\left(m + \frac{1}{t^2}\right)\right). \quad (76)$$

Proof. By Theorem 12, \bar{S}_ε is a subset of vertices of $\text{conv}(S_\varepsilon)$. Let $\sigma_o = (\Sigma_* - 2\varepsilon R)/R$. Then since $\varepsilon \leq \Sigma_*/4R$, $\sigma_o \geq \Sigma_*/2R$. Then by Lemma 2, for each $v^\varepsilon \in \bar{S}_\varepsilon$, we have

$$d(v^\varepsilon, \text{conv}(S^\varepsilon \setminus \{v^\varepsilon\})) \geq \Sigma_*/2R. \quad (77)$$

Now consider a modification of AVTA that replaces $\gamma/2$, by $\sigma/2$. Such modified AVTA will compute a subset \bar{S}_ε of vertices of $\text{conv}(S_\varepsilon)$ that must necessarily contain \bar{S}_ε . Analogous to Theorem 9, (1), the complexity of this part is as stated in part (i) of the present theorem.

Now consider $\text{conv}(\widehat{S}_\varepsilon)$ and assume v^ε is a vertex of it within a distance of less than $\sigma/2$, say $\sigma/4$. Then by Lemma 2, $v^\varepsilon \notin \overline{S}^\varepsilon$. We can thus apply the Triangle Algorithm to remove any vertex of $\text{conv}(\widehat{S}_\varepsilon)$ that is within a distance of less than $\sigma/2$ of the convex hull of the other vertices in $\text{conv}(\widehat{S}_\varepsilon)$. Again analogous to Theorem 9 the over all complexity of this step is bounded by

$$O\left(mK_\varepsilon^2 + \frac{K_\varepsilon^2}{\sigma_\circ^2}\right) = O\left(mK_\varepsilon^2 + \frac{K_\varepsilon^2}{\sigma^2}\right). \quad (78)$$

This is dominated by the complexity of the first part. This proves (i). Proof of (ii) follows from Theorem 11, (45), that $\gamma\rho_* \leq \sigma_*$.

To prove (iii), we start by $\sigma = 1/2$ and run AVTA. Then as previous case prune unwanted vertices. If we end up with \overline{S}_ε , we are done. If not, we repeat the process with $\sigma = 1/4$ and so on. Eventually we will recover \overline{S}_ε .

The proof of (iv) is analogous to the proof of Theorem 12, part (3). □

Remark 6. Ideally, K_ε is within a constant multiple of K , in which case the complexities are analogous to those of Theorem 9. In the worst-case $\widehat{S}_\varepsilon = S_\varepsilon$, i.e. $K_\varepsilon = n$. On the other hand, ignoring the size of K_ε , suppose $\sigma_\circ \geq (\sqrt{n/K})\varepsilon$, then the complexity of generating the vertices of $\text{conv}(S_\varepsilon)$ is

$$O\left(nmK_\varepsilon + \frac{nK}{\varepsilon^2}\right). \quad (79)$$

6 Triangle Algorithm with Johnson-Lindenstrauss Projections

Consider again $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$. We wish to compute the subset $\overline{S} = \{\bar{v}_1, \dots, \bar{v}_K\}$ of all vertices of $\text{conv}(S)$. Johnson-Lindenstrauss lemma allows embedding the n points of S in an m' -dimensional Euclidean space, where $\mathbb{R}^{m'}$, $m' < m$, via a randomized linear map so that the distances between every pair of points in S and those of their images in $\mathbb{R}^{m'}$ remain approximately the same, with high probability. More specifically, there is a universal constant c such if ε' satisfies,

$$\frac{c \log n}{m} \leq \varepsilon'^2 < 1, \quad (80)$$

and $m' < m$ is an integer satisfying

$$m' \approx \frac{c \log n}{\varepsilon'^2}, \quad (81)$$

then there exists a randomized linear map $L : \mathbb{R}^m \rightarrow \mathbb{R}^{m'}$ so that if $u_i = L(v_i)$, and

$$U = L(S) = \{u_1, \dots, u_n\} \subset \mathbb{R}^{m'}, \quad (82)$$

then for for each $i, j \in \{1, \dots, n\}$ we have

$$\Pr\left(d(v_i, v_j)(1 - \varepsilon') \leq d(u_i, u_j) \leq d(v_i, v_j)(1 + \varepsilon')\right) > 1 - \frac{2}{n}. \quad (83)$$

The projection of each point takes $O(m \log n)$ operations so that the overall number of operations to project all the n points is

$$O(nm \log n). \quad (84)$$

In this section we consider computing \overline{S} , the set of vertices of $\text{conv}(S)$ by using the Johnson-Lindenstrauss projections and then computing the set of vertices of $\text{conv}(U)$ via AVTA. Let \overline{U} denote the set of vertices of $\text{conv}(U)$ and let its cardinality be K' . First we state some properties of $\text{conv}(U)$.

Lemma 3. Given $v \in S$, $L : \mathbb{R}^m \rightarrow \mathbb{R}^{m'}$, a randomized linear map, suppose $u = L(v)$ is a vertex of $\text{conv}(U)$. Then v is a vertex of $\text{conv}(S)$.

Proof. Suppose v is not a vertex of $\text{conv}(S)$. Then $v = \sum_{i=1}^n \alpha_i v_i$, $\sum_{i=1}^n \alpha_i = 1$, $\alpha_i \geq 0$, $i = 1, \dots, n$, with some $0 < \alpha_j < 1$. By linearity of L we have

$$u = L(v) = \sum_{i=1}^n \alpha_i L(v_i) = \sum_{i=1}^n \alpha_i u_i. \quad (85)$$

This implies u is not a vertex of $\text{conv}(U)$, a contradiction. \square

The next theorem gives an estimate of the robustness parameters of $\text{conv}(U)$ in terms of those $\text{conv}(S)$.

Theorem 14. Suppose $\text{conv}(S)$ is Γ_* -robust and Σ_* -weakly robust. Let $U = L(S)$, L a randomized linear map, $L : \mathbb{R}^m \rightarrow \mathbb{R}^{m'}$. Let m' and ε' be related as in (81). If $\text{conv}(U)$ is Γ'_* -robust, Σ'_* -weakly robust, then with probability at least $(1 - 2/n)$, we have

$$\Gamma'_* \geq \Gamma_*(1 - \varepsilon'), \quad \Sigma'_* \geq \Sigma_*(1 - \varepsilon'). \quad (86)$$

Proof. Suppose u is a vertex of $\text{conv}(U)$ and \widehat{U} a subset of its vertices not containing u . Let v and \widehat{S} be the preimages of u and \widehat{U} under the linear map L . By Lemma 3 v and the elements of \widehat{S} are all vertices of $\text{conv}(S)$. From (83) it is easy to argue that with probability at least $(1 - 2/n)$ we have

$$d(u, \text{conv}(\widehat{U})) \geq d(v, \text{conv}(\widehat{S}))(1 - \varepsilon'). \quad (87)$$

The claimed inequalities follow. \square

From Theorem 14 and Theorem 9 we can state the following:

Theorem 15. Given $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$ let $U = L(S) = \{u_1, \dots, u_n\} \subset \mathbb{R}^{m'}$, L a randomized linear map, m' , ε' as before. Let $\overline{U} = \{\overline{u}_1, \dots, \overline{u}_{K_{\varepsilon'}}\}$ be the set of vertices of $\text{conv}(U)$. Suppose $\text{conv}(S)$ is Γ_* -robust and $\text{conv}(U)$ is Γ'_* -robust. Then with probability at least $(1 - 2/n)$,

(1) The number of arithmetic operations of AVTA to compute \overline{U} is

$$O\left(nK_{\varepsilon'}(m' + \frac{nK_{\varepsilon'}R^2}{\Gamma_*^2})\right) = O\left(\frac{n \log nK_{\varepsilon'}}{\varepsilon'^2} + \frac{nK_{\varepsilon'}}{\gamma_*^2(1 - \varepsilon')^2}\right). \quad (88)$$

(2) Given any prescribed positive $t \in (0, 1)$, AVTA in

$$O\left(nK_{\varepsilon'}^t(m' + \frac{1}{t^2})\right) = O\left(\frac{n \log nK_{\varepsilon'}^t}{\varepsilon'^2} + \frac{nK_{\varepsilon'}^t}{t^2}\right) \quad (89)$$

operations can compute a subset \overline{U}^t of \overline{U} of size $K_{\varepsilon'}^t$, so that the distance from each point in $\text{conv}(U)$ to $\text{conv}(\overline{U}^t)$ is at most t . \square

Remark 7. The results in this section and the above theorem suggest a heuristic approach as an alternative to using AVTA directly to compute all the vertices of $\text{conv}(S)$: Compute $U = L(S)$, the Johnson-Lindenstrauss projection of S under a randomized linear map L . Then apply AVTA to compute all the vertices of $\text{conv}(U)$, \overline{U} . This identifies $|\overline{U}| \leq K$ vertices of $\text{conv}(S)$. Next move up to the full dimension and continue with AVTA to recover the remaining vertices of $\text{conv}(S)$. Alternatively, we can repeat randomized projections and compute the corresponding vertices. We would have to delete duplications which is not difficult, given that we store the computed vertices via their vector of representation of convex combination coefficients. We would expect that when sufficient number of projections are applied all vertices of $\text{conv}(S)$ can be recovered. However, in the remaining of the section we analyze the probability that under a random projection, the projection of a vertex of $\text{conv}(S)$ is a vertex of the projection.

In what follows we first state a result on Johnson-Lindenstrauss random projections on the convex hull membership problem from Vu et al. (2017). Next we state an alternative result.

Proposition 6. (Vu et al. (2017), Proposition 3.3) *Given $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$, $p \in \mathbb{R}^m$ such that $p \notin \text{conv}(S)$, let $d = \min\{d(p, x) : x \in \text{conv}(S)\}$ and $D = \max\{d(p, v_i) : i = 1, \dots, n\}$. Let $T : \mathbb{R}^m \rightarrow \mathbb{R}^k$ be a random linear map. Then*

$$\text{Prob}\left(T(p) \notin T(\text{conv}(S))\right) \geq 1 - 2n^2 e^{-c(\varepsilon^2 - \varepsilon^3)k} \quad (90)$$

for some constant c (independent of m, n, k, d, D) and $\varepsilon < d^2/D^2$. \square

Remark 8. Note that $k = O(\ln n/\varepsilon^2) = O(\ln n D^4/d^4)$.

The following is an alternative to Proposition 6 based on the Distance Duality theorem (1) and generally gives a better estimate of ε , hence a smaller k than Proposition 6.

Theorem 16. *Given $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$, $p \in \mathbb{R}^m$ such that $p \notin \text{conv}(S)$, let $d = \min\{d(p, x) : x \in \text{conv}(S)\}$, $p_* = \text{argmin}\{d(p, x) : x \in \text{conv}(S)\}$ and $D = \max\{d(p, v_i) : i = 1, \dots, n\}$. Let*

$$E = \min \left\{ \frac{d(p, v_i)}{d(p_*, v_i)} : i = 1, \dots, n \right\}. \quad (91)$$

Let $T : \mathbb{R}^m \rightarrow \mathbb{R}^k$ be a random linear map. Then

$$\text{Prob}\left(T(p) \notin T(\text{conv}(S))\right) \geq 1 - 2n^2 e^{-c\varepsilon^2 k}, \quad (92)$$

for some constant c (independent of m, n, k, d, D) and $\varepsilon < (E-1)/(E+1)$. Furthermore, $(E-1)/(E+1) > d^2/4D^2$.

Proof. Since p_* is the closest point to p in $\text{conv}(S)$, it is easy to show that it is a p -witness, i.e.

$$d(p_*, v_i) < d(p, v_i), \quad \forall i = 1, \dots, n. \quad (93)$$

Let $\bar{p} = T(p)$, $\bar{p}_* = T(p_*)$, and for $i = 1, \dots, n$, $\bar{v}_i = T(v_i)$. We now consider the set of $n+1$ points $\{v_0 = p, v_1, \dots, v_n\}$ and their random projections and find condition on ε such that \bar{p}_* will be an \bar{p} -pivot with respect to $T(\text{conv}(S))$, probabilistically. By the Johnson-Lindenstrauss Lemma we have,

$$\text{Prob}\left((1-\varepsilon)d(v_i, v_j) \leq d(\bar{v}_i, \bar{v}_j) \leq (1+\varepsilon)d(v_i, v_j)\right) \geq 1 - 2(n+1)^2 e^{-c\varepsilon^2 k}, \quad (94)$$

for some constant c (independent of m, n, k). From (94) and definition of E , for each $i = 1, \dots, n$ with probability at least $1 - 2(n+1)^2 e^{-c\varepsilon^2 k}$ we have,

$$d(\bar{p}_*, \bar{v}_i) \leq (1+\varepsilon)d(p_*, v_i) \leq \frac{(1+\varepsilon)}{E}d(p, v_i) \leq \frac{(1+\varepsilon)}{(1-\varepsilon)} \frac{1}{E}d(\bar{p}, \bar{v}_i). \quad (95)$$

Note that assuming $n \geq 2$, $1 < E < \infty$. We thus restrict ε to satisfy

$$\frac{(1+\varepsilon)}{(1-\varepsilon)} \frac{1}{E} < 1. \quad (96)$$

Equivalently,

$$\varepsilon < \frac{E-1}{E+1}. \quad (97)$$

Thus with ε satisfying the above, \bar{p}_* is a witness with high probability.

Next we find a lower bound on the right-hand-side of the above. Since E is finite, $E = d(p, v_j)/d(p_*, v_j)$ for some j , i.e. $p_* \neq v_j$. Consider the triangle with vertices p , v_j and p_* . With $d(p, p_*)$ and $d(p, v_j)$ fixed, the maximum value of $d(p_*, v_j)$ is $\sqrt{d^2(p, v_i) - d^2(p, p_*)}$. Using this we may write

$$E = \frac{d(p, v_j)}{d(p_*, v_j)} \geq \frac{d(p, v_j)}{\sqrt{d^2(p, v_i) - d^2(p, p_*)}} = \frac{1}{\sqrt{1 - d^2(p, p_*)/d^2(p, v_j)}}. \quad (98)$$

But $d(p, p_*) = d$ and $d(p, v_j) \leq D$. Thus

$$E \geq \frac{1}{\sqrt{1 - d^2/D^2}} = \frac{D}{\sqrt{D^2 - d^2}}. \quad (99)$$

The function $(x-1)/(x+1)$ is monotonically increasing. Thus from (100) we have

$$\frac{E-1}{E+1} \geq \frac{D - \sqrt{D^2 - d^2}}{D + \sqrt{D^2 - d^2}} = \frac{d^2}{(D + \sqrt{D^2 - d^2})^2} \geq \frac{d^2}{4D^2}. \quad (100)$$

□

Remark 9. We would expect that $(E-1)/(E+1)$ is generally a larger number than $d^2/4D^2$. Thus Theorem 16 gives generally a better estimate of ε and k than those of Proposition 6. An additional advantage of Theorem 16 is that it shows the applicability of the Triangle Algorithm in solving the convex hull membership problem using random projections.

We now state a corollary of the theorem on computation of all vertices of $\text{conv}(S)$.

Corollary 2. *Given $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$, suppose $\text{conv}(S)$ is Γ_* -robust. Let R be the diameter of S . Suppose v_j is a vertex of $\text{conv}(S)$. Let $T : \mathbb{R}^m \rightarrow \mathbb{R}^k$ be a random linear map. Then the probability that $T(v_j)$ is a vertex of $T(\text{conv}(S))$ is at least $1 - 2n^2e^{-c\varepsilon^2k}$, for some constant c (independent of m, n, k) and $\varepsilon < \gamma_*^2/4$.*

Proof. We apply the previous theorem with v_j as b and considering the probability that under a random projection of v_j lies in projection of the convex hull of the remaining points. Note that $d(v_j, \text{conv}(S \setminus \{v_j\})) \geq \Gamma_*$ and $\max\{d(v_j, v_i) : v_i \in S \setminus \{v_j\}\} \leq R$. Thus we can replace for b/D in (100) in the previous theorem by $\gamma_* = \Gamma_*/R$. Thus we can write $(E-1)/(E+1) \geq \gamma_*^2/4$. This gives the upper bound on ε . □

6.1 AVTA Under Perturbation and Johnson-Lindenstrass Projection

Let S_ε be as before and U_ε , a subset of $\mathbb{R}^{m'}$ the perturbation of U . Let \bar{U}_ε be the perturbation of \bar{U} . Based on the results in this section and previous complexity bounds we have

Theorem 17. *Let $S = \{v_1, \dots, v_n\} \subset \mathbb{R}^m$. Assume $\text{conv}(S)$ is Σ_* -weakly robust. Suppose $\varepsilon < \Sigma_*/4R$. Let $\sigma_o = (\Sigma_* - 2\varepsilon R)/R = \sigma_* - 2\varepsilon$. Then with probability at least $(1 - 2/n)$,*

(i) *AVTA can be modified to compute a subset \hat{U}_ε of U_ε , of cardinality $K_{\varepsilon\varepsilon'}$ such that it contains \bar{U}_ε . Then AVTA can compute from this subset \bar{U}_ε itself, where the total number of operations satisfies*

$$O\left(nm'K_{\varepsilon\varepsilon'} + \frac{nK_{\varepsilon\varepsilon'}}{\sigma_o^2(1 - \varepsilon')^2}\right) = O\left(\frac{n \log nK_{\varepsilon\varepsilon'}}{\varepsilon^2} + \frac{nK_{\varepsilon\varepsilon'}}{\sigma_o^2(1 - \varepsilon')^2}\right). \quad (101)$$

(ii) *Given any prescribed positive $t \in (0, 1)$, in*

$$O\left(\frac{n \log nK_{\varepsilon\varepsilon'}^{(t)}}{\varepsilon^2} + \frac{nK_{\varepsilon\varepsilon'}^{(t)}}{\sigma_o^2(1 - \varepsilon')^2}\right) \quad (102)$$

operations the modified AVTA can compute a subset U_ε^t of \bar{U}_ε of size $K_{\varepsilon\varepsilon'}^{(t)}$ so that the distance from each point in $\text{conv}(U_\varepsilon)$ to $\text{conv}(U_\varepsilon^t)$ is at most t .

7 Applications

While the modified AVTA algorithm comes with theoretical guarantees, in certain cases the algorithm might output many more vertices, K_ε , than desired. Here we present a practical implementation that always outputs exactly K vertices, provided K is known. When K is unknown, our experiments in the next section reveal that the algorithm can automatically detect a slightly larger set that contains a good approximation to the K vertices of interest. Notice that we want a fast way to detect good approximations to the original vertices of the set S and prune out spurious points, i.e., additional vertices of the set S_ε . The key insight on top of the AVTA algorithm is the following: *If the perturbed set is randomly projected onto a lower dimensional space, it is more likely for an original vertex to still be a vertex than for a spurious vertex.* Using this insight the algorithm outlined below runs the modified AVTA algorithm over several random projections and outputs the set of points that appear as vertices in many random projections.

AVTA with multiple random projections ($S = \{v_1, \dots, v_n\}$, K , γ , M)

- **Step 0.** Set $\text{Freq} \leftarrow 0^{|S|}$.
- **Step 1.** For $i = 1$ to M :
 - $S' \leftarrow S$: Project data on to randomly chosen $\frac{4\log(n)}{\epsilon^2}$ dimensions.
 - $\hat{S} \leftarrow \text{AVTA}(S', \gamma)$
 - For each $d_j \in \hat{S}$, $\text{Freq}[j] = \text{Freq}[j] + 1$.
- **Step 2.** Output top K frequent vertices.

We now show how AVTA can be used to solve various problems in computational geometry and machine learning.

Application of AVTA in Linear Programming: Consider linear programming feasibility problem of testing if $P = \{x \in \mathbb{R}^n : Ax = b, x \geq 0\}$ is nonempty, where A is $m \times n$, $b \in \mathbb{R}^m$. Suppose n is much larger than m . If we reduce the size of A the problem would be more efficiently solvable, no matter what algorithm we use to solve it.

Proposition 7. *Given $P = \{x \in \mathbb{R}^n : Ax = b, x \geq 0\}$, let $\text{conv}(A)$ denote the convex hull of columns of A . Let A' denote the $m \times n'$ submatrix A whose columns form the set of all vertices of $\text{conv}(A)$. Let*

$$P' = \{x' \in \mathbb{R}^{n'} : A'x' = b, x' \geq 0\}. \quad (103)$$

Then P is feasible if and only if P' is feasible.

Proof. Clearly, if P' is feasible then P is feasible. Assume P is feasible. Thus for some $x \in \mathbb{R}^n$, $x \geq 0$, $Ax = b$. Denote the columns of A by $a^{(i)}$. Then each $a^{(i)}$ is a convex combination of columns of A' . That is, for each $i = 1, \dots, n$, there exists

$$\alpha^{(i)} \in S_{n'} = \{s \in \mathbb{R}^{n'} : \sum_{i=1}^{n'} s_i = 1, s \geq 0\}, \quad (104)$$

where

$$a^{(i)} = A'\alpha^{(i)}. \quad (105)$$

Thus

$$Ax = \sum_{i=1}^n x_i a^{(i)} = \sum_{i=1}^n x_i A' \alpha^{(i)} = A' \sum_{i=1}^n x_i \alpha^{(i)}. \quad (106)$$

Letting

$$x' = \sum_{i=1}^n x_i \alpha^{(i)}, \quad (107)$$

$$A'x' = b, x' \geq 0. \quad \square$$

Proposition 8. Assume $P = \{x \in \mathbb{R}^n : Ax = b, x \geq 0\}$ is nonempty. Consider the linear program $\min\{c^T x : x \in P\}$. Let B be the $(m+1) \times n$ matrix whose first row is c^T and the remaining rows are A . Let B' be the $(m+1) \times n'$ matrix whose columns form the vertices of the convex hull of the columns of B . Let c'^T be the first row of B' and A' the remaining $m \times n'$ submatrix of B' . Then

$$\min\{c^T x : Ax = b, x \geq 0\} = \min\{c'^T x' : A'x' = b, x' \geq 0\}. \quad (108)$$

Proof. Consider any feasible solution x_0 of original LP. Then by Proposition 7 the set $\{c'^T x' = c^T x_0, A'x' = b, x' \geq 0\}$ is feasible. This implies the original LP has a finite optimal value if and only if the restricted problem does. In particular, the optimal objective values of the two problems coincide. \square

The above propositions imply that AVTA has potential applications in the reduction of the LP feasibility or optimization, whether we solve the problem via simplex method or other methods.

AVTA for topic modeling in the presence of anchor words: Arora et al. (2013) provide a practical algorithm for topic modeling with provable guarantees. Their algorithm works under the assumptions that the topic-word matrix is *separable*. In particular, they assume that corresponding to each topic i , there exists an *anchor word* w_i that has a non zero probability of appearing only under topic i . Under this assumption, the algorithm of Arora et al. (2013) consists of two stages: a) find the anchor words, and b) use the anchor words to learn the topic word matrix. The problem of finding anchor words corresponds to finding the vertices of the convex hull of the word-word covariance matrix. They propose an algorithm named *fast anchor words* in order to find the vertices. Since AVTA works in general setting, we can instead use AVTA to find the anchor words. Additionally, the fast anchor words algorithm needs to know the value of the number of anchor words, as an input. On the other hand, from the statements of Theorems 9 and 13 it is easy to see that AVTA can work in a variety of settings when other properties of the data are known such as the robustness. We argue that robustness is a parameter that can be tuned in a better manner than trying different values of the number of anchor words. In fact, one can artificially add random noise to the data and make it robust up to certain value. One can then run AVTA with the lower bound on robustness as input and let the algorithm automatically discover the number of anchor words. This is much more desirable in practical settings. Our first implementation of AVTA is named AVTA+RecoverL2 that uses AVTA to detect anchor words and then uses the anchor words to learn the topic word matrix using the approach from Arora et al. (2013). AVTA is also theoretically superior than fast anchor words and achieves slightly better run times in the regime when the number of topics is $o(\log n)$, where n is the number of words in the vocabulary. This is usually the case in most practical scenarios.

AVTA for topic modeling the absence of anchor words: The presence of anchor words is a strong assumption that often does not hold in practice. Recently, the work of Bansal et al. (2014) designed a new practical algorithm for topic models under the presence of catch words. Catch words for topic i correspond to set S_i such that it's total probability of appearing under topic i is significantly higher than in any other topic. Their algorithm called TSVD recovers much better reconstruction of the topic-word matrix in terms of the ℓ_1 error. They also assume that for each topic i , there are a few dominant documents that mostly contain words from topic i . The TSVD algorithm works in two stages. In stage 1, the (thresholded) word-document data matrix is projected onto a K -SVD space to compute a different embedding of the documents. Then, the documents are clustered into K clusters. Under the assumptions mentioned above, one can show that the dominant documents for each topic will be clustered correctly. In stage 2, a simple post processing algorithm can approximate the topic-word matrix from the clustering.

We improve on TSVD by asking the following question: *is K -SVD the right representation of the data?* Our key insight is that if dominant documents are present in the topic, it is easy to show that most other

documents will be approximated by a convex combination of the dominant topics. Furthermore, the coefficients in the convex combinations will provide a much more faithful low dimensional embedding of the data. Using this insight, we propose a new algorithm that runs AVTA on the data matrix to detect vertices and to approximate each point using a convex combination of the vertices. We then use the coefficient matrix as the new representation of the data that needs to be clustered. Once the clustering is obtained, the same post processing step from Bansal et al. (2014) can be used to recover the topic-word matrix. Our results show that the embedding produced by AVTA leads to much better reconstruction error than of that produced by TSVD. Furthermore, K -SVD is an expensive procedure and very sensitive to the presence of outliers in the data. In contrast, our new algorithm called AVTA+CatchWord is much more stable to noise in the data.

AVTA+CatchWord ($S = \{v_1, \dots, v_n\}, \gamma, K, \epsilon$)

- **Step 0.** Randomly project S onto $2K$ dimensions to get \hat{S} .
- **Step 1.** Compute a super set of vertices \bar{V} by $AVTA(\hat{S}, \gamma)$.
- **Step 2.** Prune \bar{V} into \hat{V} (of size K) by iteratively picking $\bar{v} \in \{\bar{V}\} \setminus \{\hat{V}\}$ which has the maximum distance to $conv(\hat{V})$.
- **Step 3.** For each projected point $\hat{v}_i \in \hat{S} \setminus \hat{V}$, compute a vector α_i such that $\|\hat{V}\alpha_i - \hat{v}_i\| \leq \epsilon$.
- **Step 4.** Initialize cluster assignment for each point by majority weight: $\operatorname{argmax}_{j \in [K]} \alpha_j$.
- **Step 5.** Clustering using Lloyd's algorithm on the embedding provided by the α vectors.
- **Step 6.** Use the post processing as described in Bansal et al. (2014) to recover the topic-word matrix from the clustering.

AVTA for NMF: The work of Arora et al. (2012a) showed that convex hull detection can be used to solve the non-negative matrix factorization problem under the separability assumption. We show that by using the more general AVTA algorithm for solving the convex hull problem results in comparable performance guarantee.

8 Applications and Experiments ¹

8.1 Feasibility problem

In this section, we present experimental results which empirically show when the problem is 'over complete', AVTA can be a 'shortcut' solution. In another word, given an $m \times n$ matrix A as data, where the convex hull of the columns of A , denoted by $conv(A)$, has K vertices, $K \ll n$. We apply the AVTA to solve 2 classical problems which appear in many applications.

Convex hull membership problem:

In the experiments, vertices of the convex hull are generated by the Gaussian distribution, i.e. $v_i \sim \mathcal{N}(0, \mathcal{I}_m), i \in [K]$. Having generated the vertices, the 'redundant' points d_j where $d_j \in conv(S), j \in [n - K]$ are produced using random convex combination $d_j = \sum_{i=1}^K \alpha_i v_i$. Here α_i are scaled standard uniform random variable where α_i are scaled so that $\sum_{i=1}^K \alpha_i = 1$. Specifically, comparison is by fixing $K = 100, m = 50$ and n varying from $5,000 \sim 500,000$. We compare the efficiency of 4 algorithms on solving this problem: the Simplex method Chvatal (1983), the Frank Wolfe Algorithm (FW) Jaggi (2013), the Triangle Algorithm (TA) Kalantari (2015), and our algorithm on solving the convex hull membership query problem.

Results on Convex hull membership query: Table 2 shows when $n \gg K$, AVTA is more efficient than other algorithms solving the convex hull membership problem. This result supports the output sensitivity

¹Resources: <https://github.com/yikaizhang/AVTA>

Table 2: Running time of convex hull membership (secs)

# of redundant pts	AVTA	TA	FW	Simplex
5,000	1.75	0.21	0.52	0.9
20,000	1.49	0.66	1.94	2.76
45,000	2.94	1.84	5.51	6.16
80,000	2.71	3.22	10.87	10.63
125,000	3.83	4.28	17.67	15.95
180,000	4.15	5.38	23.14	24.13
245,000	6.95	9.56	33.42	36.96
320,000	8.09	13.24	44.99	44.26
405,000	10.01	14.75	56.35	59.5
500,000	14.12	15.69	70.7	90.41

property of AVTA.

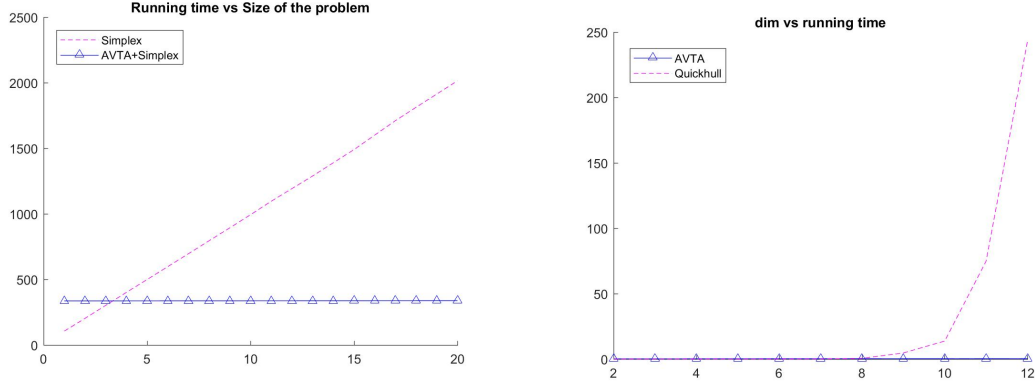
Non-negative linear system: The non-negative linear system problem is to find a feasible solution of :

$$\begin{aligned} A\alpha &= p \\ \alpha &\geq 0 \end{aligned} \tag{109}$$

In another word, to test if $p \in \text{cone}(A_j)$ where A_j are columns of A . In case when A is over complete, any feasible p can be represented using only the generators of $\text{cone}(A)$ the set $\bar{A} \subset A$. By scaling A so that columns of AD ($D_{ii} = \frac{b}{a \cdot A_i}$) are in a $m - 1$ dimensional hyperlane $\langle a, \alpha \rangle = b$, one can find the generators of $\text{cone}(A)$ by finding the vertices of the convex hull of the projected points. This could be done efficiently by AVTA. Suppose we have a linear system A and series of query points p , it is sufficient to run AVTA once for dimension reduction and solve the subproblem $\bar{A}\alpha' = p, \alpha' \geq 0$ using simplex method. We compare the running time of Simplex Method with AVTA+Simplex Method. The generator \bar{A} is entrywise independent $\text{uniform}(0, 1)$ random matrix and the 'overcomplete' part of the matrix $\bar{A}^c = A/\bar{A}$ are generated by $\bar{A}^c = \bar{A}B$ where $B \in \mathbb{R}^{K \times (n-K)}$ is entrywise independent $\text{uniform}(0, 10)$ random matrix. We set the number of generators $K = 100$, the dimension $m = 50$, and the number of 'redundant' columns $n = 50,000$. We simply set half of the query points feasible and rest infeasible. The feasible points p are generated as $p = Ax$ where $x \in \mathbb{R}^n$ is entrywise independent $\text{uniform}(0, 1)$ random vector and the infeasible points are generated in the same way as generators.

Table 3: Running time of linear programming feasibility (secs)

# of query	AVTA+Simplex	Simplex	# of query	AVTA+Simplex	Simplex
1.00	241.24	152.09	11.00	241.94	1810.72
2.00	241.36	303.86	12.00	242.01	1967.93
3.00	241.41	477.89	13.00	242.07	2125.62
4.00	241.45	660.95	14.00	242.16	2289.91
5.00	241.54	853.91	15.00	242.23	2490.52
6.00	241.61	1016.77	16.00	242.29	2680.61
7.00	241.69	1177.30	17.00	242.32	2866.23
8.00	241.72	1336.38	18.00	242.41	3065.50
9.00	241.83	1495.70	19.00	242.44	3245.78
10.00	241.91	1652.84	20.00	242.47	3412.39



(a) Running time for algorithms to find a feasible solution. (b) Running time for algorithms to find all vertices.

It can be observed from Figure 5a and Table 3 that the running time of AVTA+Simplex doesn't have obvious increase while Simplex increases drastically. This suggests the potential applications of AVTA in linear programming feasibility problem.

8.2 Computing all vertices

Compute vertices of convex hull: In this section, we compare the efficiency of AVTA with another popular algorithm for finding vertices Quickhull Barber et al. (1996). We generate vertices according to a Gaussian distribution $\mathcal{N}(0, 10)^m$. Having generated K such points, n interior points are generated as convex combination of the vertices, where the weights are generated scaled i.i.d uniform distribution.

Experiment and results: In the experiment, we set $K = 100$, $n = 500$ and m varying from 2 \sim 12.² The computational results is shown in Table 4. In high dimension $m \geq 9$, when $\text{conv}(S)$ is γ robust for some $\gamma > 0$, the AVTA algorithm successfully find all vertices of the convex hull efficiently while the Quick hull algorithm is stuck by its explosion of complexity in dimension m .

Table 4: Running time (secs)

dim	Qhull	AVTA	dim	Qhull	AVTA
2	0.13	14.82	7	2.92	41.51
3	0.02	16.62	8	16.48	39.63
4	0.04	24.49	9	82.09	44.21
5	0.12	32.76	10	391.36	45.79
6	0.59	37.66	11	1479.51	51.19

Compute vertices of simplex in high dimension: The Fast Anchor Word can be used to detect the vertices of a simplex. In this section, we compare the efficiency of AVTA with Fast Anchor Word when convex hull is a simplex with $K = 50$ and $m = 100$. The number of points in the convex hull n varies from 100 \sim 100,000.

Results of running time in simplex case: The running of efficiency comparison between AVTA and Fast Anchor Word in simplex case is presented in Figure 6c. In regime $n \geq 30,000$, AVTA has less running time.

Compute vertices with perturbation: In this section, we compare the robustness of AVTA with multiple random projections presented in section 7 with Fast Anchor Word Arora et al. (2013). Instead of actual set

² The maximum of dimension is 12 in the experiment because of the explosion of running time of the Quick hull algorithm

of points S as input, the algorithm is given a perturbed set S_\circ , i.e. S is corrupted by some noise. Having fixed $K = 100, n = 500, m = 100$, we choose a Gaussian perturbation from $\mathcal{N}(0, \tau)^m$ where τ varies from 0.3 to 3. In case of general convex hull, a failure of Fast Anchor Word on computing vertices of general convex hull is presented. The data is generated by setting $\tau = 0.3, m = 50, n = 500$ and let K varies from 10 \sim 100. We do an error analysis and evaluate the output of the algorithms by measuring the l_2 distance between true vertices and the convex hull of output vertices of the two algorithms. More precisely, given a true vertex $v_i \in S$ and \hat{S} , the output of an algorithm, the error in recovering v_i is defined to be $\min_{u \in \text{conv}(\hat{S})} \|u - v_i\|_2$. We

add up all the errors to get the total accumulated error.

Results on computing perturbed vertices:

The recovery error in robustness comparison is shown in Table 5. The AVTA with multiple random projection has a better recovery error in the simplex case.

It can also be observed from Figure 6b that in general case, as number of vertices exceeds the number of dimensions, Fast Anchor Word fails to recover more vertices and its error explodes.

Table 5: Recovery error (Simplex)

Var	AVTA+Multiple Rp	Fast Anhor	variance	AVTA+Multiple Rp	Fast Anhor
0.3	2.96	2.96	1.8	16.60	17.98
0.6	5.79	5.79	2.1	19.40	20.58
0.9	8.61	9.36	2.4	21.93	23.77
1.2	11.34	12.00	2.7	23.69	24.90
1.5	14.16	15.44	3	26.72	28.78

8.3 Topic modeling

We compare our algorithms with the Fast Anchor + Recoverl2 algorithm of Arora et al. (2013) and the TSVD algorithm of Bansal et al. (2014) on two types of data sets: semi-synthetic data and real world data. We next describe our methodology and empirical results in detail.

Semi Synthetic Data: For Semi-Synthetic data set, we use similar methodology as in Arora et al. (2013). We first train the model on real data set using Gibbs sampling with 1,000 iterations. We choose 50 as the number of topics which follows Bansal et al. (2014). Given the parameters learned from dataset, we generate documents with α set to be 0.01. The average document length is 1,000. Then the reconstruction error is measured by the l_1 distance of bipartite matched pairs between the true word-topic distribution and the word-topic distribution Arora et al. (2013). We then average the errors to compute the final mean error.

Real Data: We use the NIPS data set with 1500 documents, and a pruned vocabulary of 2K words, and the NYTimes Corpus with sub sampled 30,000 documents, and a pruned vocabulary of 5k words.³ For the real world data set, as in prior works Arora et al. (2013); Bansal et al. (2014), we evaluate the coherence to measure topic quality Yao et al. (2009). Given a set of words \mathcal{W} associated with a learned topic, the coherence is computed as: $Coherence(\mathcal{W}) = \sum_{w_1, w_2 \in \mathcal{W}} \log \frac{D(w_1, w_2) + \epsilon}{D(w_2)}$, where $D(w_1)$ and $D(w_1, w_2)$ are the number of documents where w_1 appears and (w_1, w_2) appear together respectively Arora et al. (2013), and ϵ is set to 0.01 to avoid w_1, w_2 that never co-occur Stevens et al. (2012). The total coherence is the sum of the coherence of each topic. In the NIPS dataset, 1,000 out of the 1,500 documents were selected as the training set to learn the word-topic distributions. The rest of the documents were used as the testing set.

Implementation Details: We compare 4 algorithms, AVTA+CatchWord, TSVD, the Fast Anchor + Recoverl2 and the AVTA+Recoverl2. We implement our own version of Fast Anchor + Recoverl2 as described in Arora et al. (2013). TSVD is implemented using the code provided by the authors in Bansal et al. (2014). AVTA+Recoverl2 corresponds to using AVTA to detect anchor words from the word-word covariance matrix

³<https://archive.ics.uci.edu/ml/datasets/bag+of+words>

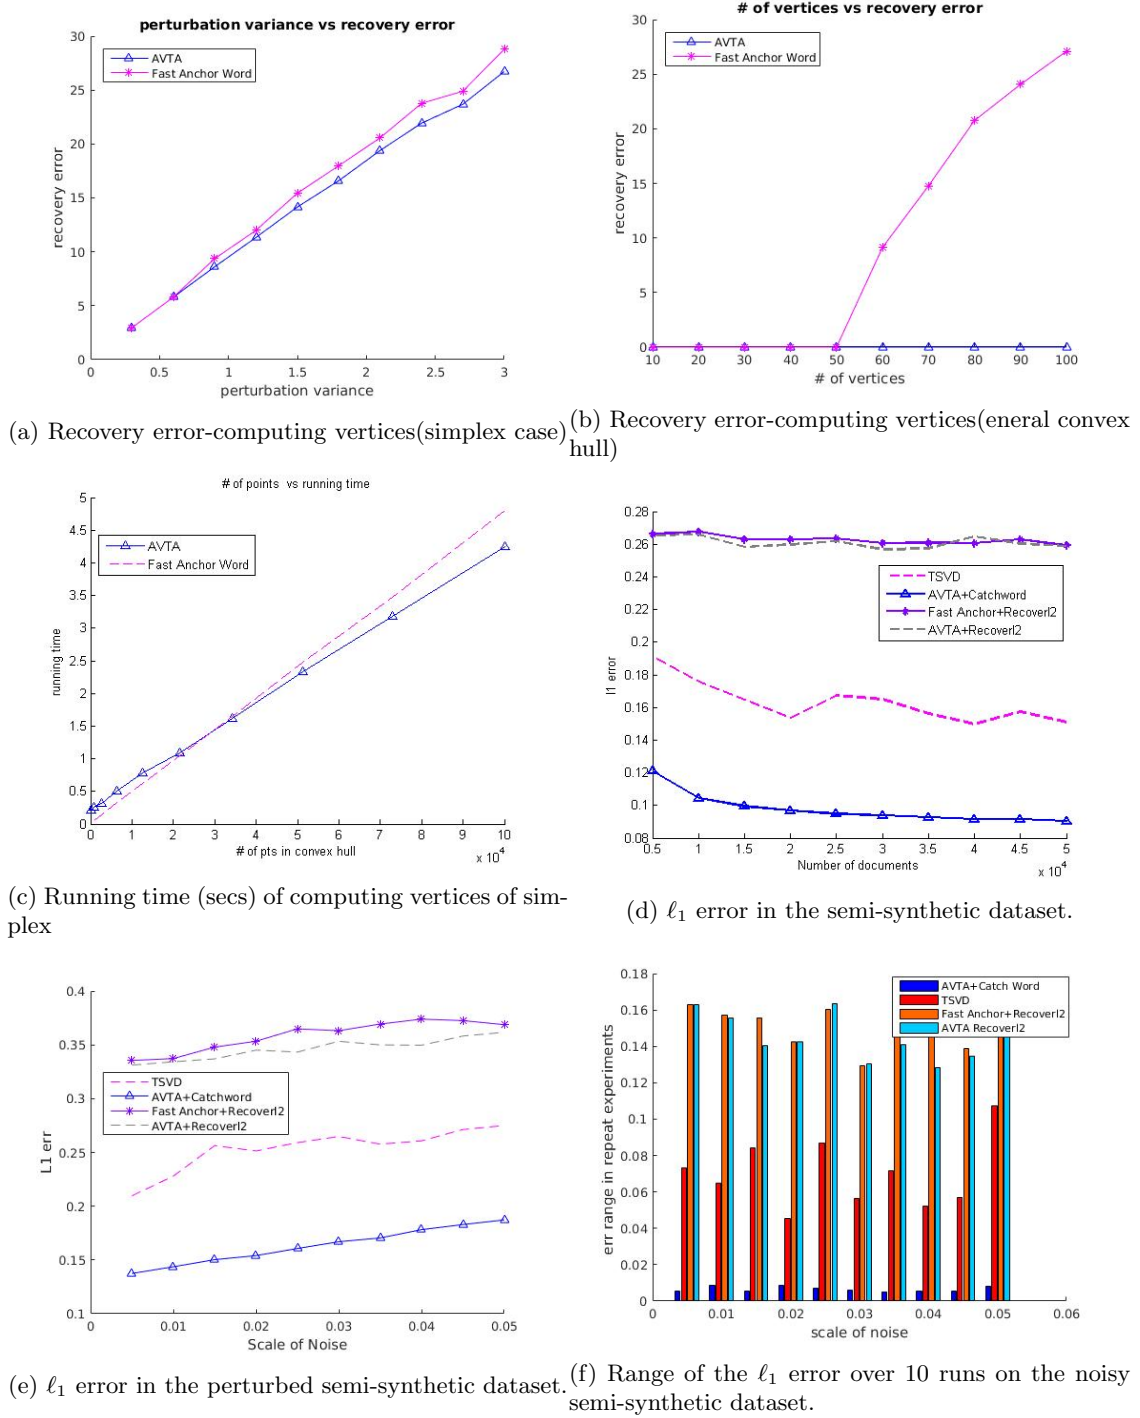


Figure 6

and then using the Recoverl2 procedure from Arora et al. (2013) to get the topic-word matrix. AVTA + CatchWord corresponds to finding the low dimensional embedding of each document in terms of the coefficient vector of its representation in the convex hull of the vertices. The next step is to cluster these

points. In practice, one could use the Lloyd’s algorithm for this step which could be sensitive to initialization. To remedy this, we use similar heuristic as Bansal et al. (2014) of the initialization step. We repeat AVTA for 3 times and pick the set of vertices with highest quality where the quality is measured by sum of distances of each vertex to convex hull of other vertices. We set the number of output vertices $K = 50$ which is the same as the number of topics. i.e. each vertex corresponds to a topic. We found that initializing by simply assigning clusters using neighborhoods of highest degree vertices works effectively. As a final step, we use the post processing step from Bansal et al. (2014) to recover the topic-word matrix from the clustering.

Robustness: We also generate perturbed version of the semi synthetic data. We generate a random matrix with i.i.d. entries uniformly distributed with different scales varying from 0.005 – 0.05. We test all the algorithms with the document-word matrix added with the noise matrix.

Results on Semi Synthetic Data: Figures 6d and 6e show the ℓ_1 reconstruction of all the four algorithms under both clean and noisy versions of the semi synthetic data set. For topic i , let A_i be the ground truth topic vector and \hat{A}_i be the topic vector recovered by the algorithm. Then the ℓ_1 error is defined as $\frac{1}{K} \sum_{i=1}^K \|A_i - \hat{A}_i\|_1$. The plots show that AVTA+CatchWord is consistently better than both TSVD and Fast Anchor + Recoverl2 and produces significantly more accurate topic vectors. In order to further test the robustness of our approach, we plot in Figure 6f the range of the ℓ_1 error obtained across multiple runs of the algorithms on the same data set. The range is defined to be the difference between the maximum and the minimum error recovered by the algorithm across different runs. We see that AVTA+CatchWord produces solutions that are much more stable to the effect of the noise as compared to other algorithms. Table 6 shows the running time of the experiments of 4 algorithms. As can be seen, when using AVTA to learn topic models via the anchor words approach, our algorithm has comparable run time to Fast Anchor + Recoverl2. In CatchWord based learning, computing vertices is expensive compared to K-SVD step of TSVD thus AVTA has longer running time.

Table 6: Running time of algorithms on semi synthetic data (secs)

Num of documents	5,000	15,000	30,000	50,000
Fast anchor+Recoverl2	5.49	6.00	10.30	13.60
AVTA+Recoverl2	7.82	7.68	12.84	16.40
TSVD	17.02	43.27	81.24	112.80
AVTA+Catch Word	29.89	120.04	372.17	864.30

Results on Real Data: Table 7 shows the topic coherence obtained by the algorithms. One can see that in both the approaches, either via anchor words or the clustering approach, AVTA based algorithms perform comparably to state of the art methods ⁴. The running time is presented in Table 8. The AVTA+CatchWord has less running time in the real data experiments. Per our observation, the convex hull of word-document vectors in real data set has more vertices than K , the number of topics. The AVTA catches K vertices efficiently due to its small number of iterations on line search for γ . In semi-synthetic data set, the number of ‘robust’ vertices is approximately the same as number of topics K thus AVTA needs to find almost all vertices. To catch enough vertices, AVTA needs several iterations decreasing γ which is computationally expensive.

8.4 Non-negative matrix factorization

AVTA for NMF: For our experiments on NMF we use the Swimmer data set Donoho and Stodden (2003) that consists of 256 swimmer figures with each a 32×32 binary pixel images. One can interpret each image as a document and pixels as a word in the document Ding et al. (2013). All swimmers consist of 4 limbs with each limb having 4 different possible poses. One can then consider the different poses of limbs as the true

⁴The topic coherence results for TSVD do not match the ones presented in Bansal et al. (2014) since in their experiments, the authors look at top 10 most frequent words in each topic. In our experiments we compute coherence for the top 5 most frequent words in each topic.

underlying topics Donoho and Stodden (2003). We compare the algorithm proposed in Arora et al. (2012a) with AVTA+RecoverL2 on the swimmer data set. We construct a noisy version by adding spurious poses to original swimmer data set. Let $\Omega(A)$ be a function that outputs a randomly chosen 32×8 block of an image. We generate a 'spurious pose' of size 32×8 by $\Omega(M_i)$ where M_i is a randomly chosen swimmer image. Then we take another randomly chosen image M_j and compute the corrupted image as $M'_j = M_j + c \cdot \Omega(M_i)$ where we simply set $c = 0.1$. An illustration of the noise data set is shown in Figure 7b. Since the true underlying topics are known, we will plot the output of the algorithms and compare it with the underlying truth.

Results on NMF: We compare the performance of AVTA on these data sets with the performance of the Separable NMF algorithm proposed in Arora et al. (2012a). Figures 7c and 7d show the output of the Separable NMF algorithm and that of our algorithm respectively on the noisy data set. Our approach produces competitive results as compared to the Separable NMF algorithm.

Table 7: Topic coherence on real data

	Fast Anchor+RecoverL2	AVTA+RecoverL2	TSVD	AVTA+Catch Word
NIPS	-15.8 \pm 2.24	-16.04 \pm 2.09	-16.86 \pm 1.66	-18.65 \pm 1.78
NYTimes	-32.15 \pm 2.7	-32.13 \pm 2.43	-29.39 \pm 1.43	-30.13 \pm 1.98

Table 8: Running time on real data experiments (secs)

	Fast Anchor+RecoverL2	AVTA+RecoverL2	TSVD	AVTA+Catch Word
NIPS	3.22	4.41	56.58	22.78
NYTimes	26.05	27.79	237.6	101.07

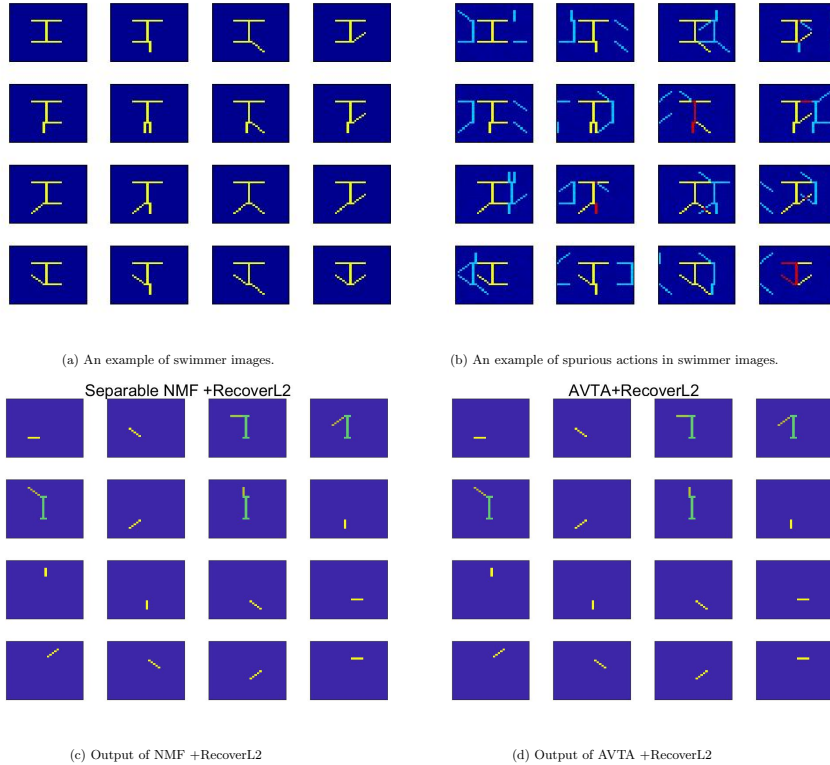


Figure 7

9 Conclusion

In this work we have presented a fast and robust algorithm for computing the vertices of the convex hull of a set of points. Our algorithm efficiently computes the vertices of convex hulls in high dimensions and even in the special case of the simplex is competitive with the state of the art approaches in terms of running time Arora et al. (2013). Furthermore, our algorithm leads to an improved algorithm for topic modeling that is more robust and produces better approximations to the topic-word matrix. It will be interesting to provide theoretical claims supporting this observation in the context of specific applications. Furthermore, we believe that our algorithm will have more applications in machine learning problems beyond the ones investigated here as well as applications in computational geometry and in linear programming.

References

- Anima Anandkumar, Dean P Foster, Daniel J Hsu, Sham M Kakade, and Yi-Kai Liu. A spectral algorithm for latent dirichlet allocation. In *Advances in Neural Information Processing Systems*, pages 917–925, 2012.
- Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra. Computing a nonnegative matrix factorization—provably. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 145–162. ACM, 2012a.
- Sanjeev Arora, Rong Ge, and Ankur Moitra. Learning topic models—going beyond svd. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 1–10. IEEE, 2012b.
- Sanjeev Arora, Rong Ge, Yonatan Halpern, David Mimno, Ankur Moitra, David Sontag, Yichen Wu, and Michael Zhu. A practical algorithm for topic modeling with provable guarantees. In *International Conference on Machine Learning*, pages 280–288, 2013.
- Sanjeev Arora, Rong Ge, and Ankur Moitra. New algorithms for learning incoherent and overcomplete dictionaries. In *COLT*, pages 779–806, 2014.
- Pranjal Awasthi and Andrej Risteski. On some provably correct cases of variational inference for topic models. In *Advances in Neural Information Processing Systems*, pages 2098–2106, 2015.
- Trapit Bansal, Chiranjib Bhattacharyya, and Ravindran Kannan. A provable svd-based algorithm for learning topics in dominant admixture corpus. In *Advances in Neural Information Processing Systems*, pages 1997–2005, 2014.
- C Bradford Barber, David P Dobkin, and Hannu Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software (TOMS)*, 22(4):469–483, 1996.
- David M Blei. Probabilistic topic models. *Communications of the ACM*, 55(4):77–84, 2012.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- Avrim Blum, Sariel Har-Peled, and Benjamin Raichel. Sparse approximation via generating point sets. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 548–557. Society for Industrial and Applied Mathematics, 2016.
- Christopher JC Burges. A tutorial on support vector machines for pattern recognition. *Data mining and knowledge discovery*, 2(2):121–167, 1998.
- Timothy M Chan. Optimal output-sensitive convex hull algorithms in two and three dimensions. *Discrete & Computational Geometry*, 16(4):361–368, 1996a.

- Timothy M Chan. Output-sensitive results on convex hulls, extreme points, and related problems. *Discrete & Computational Geometry*, 16(4):369–387, 1996b.
- Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 10(1):377–409, 1993.
- Vasek Chvatal. *Linear programming*. Macmillan, 1983.
- Kenneth L Clarkson. More output-sensitive geometric algorithms. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 695–702. IEEE, 1994.
- Kenneth L Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Transactions on Algorithms (TALG)*, 6(4):63, 2010.
- Weicong Ding, Mohammad Hossein Rohban, Prakash Ishwar, and Venkatesh Saligrama. Topic discovery through data dependent and random projections. In *ICML (3)*, pages 1202–1210, 2013.
- David Donoho and Victoria Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in Neural Information Processing Systems*, 2003.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics (NRL)*, 3(1-2):95–110, 1956.
- Bernd Gärtner and Martin Jaggi. Coresets for polytope distance. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 33–42. ACM, 2009.
- Elmer G Gilbert. An iterative procedure for computing the minimum of a quadratic form on a convex set. *SIAM Journal on Control*, 4(1):61–80, 1966.
- Sariel Har-Peled, Dan Roth, and Dav Zimak. Maximum margin coreset for active and noise tolerant learning. In *IJCAI*, pages 836–841, 2007.
- Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. 2013.
- Ray A Jarvis. On the identification of the convex hull of a finite set of points in the plane. *Information Processing Letters*, 2(1):18–21, 1973.
- Yi Jin and Bahman Kalantari. A procedure of chvátal for testing feasibility in linear programming and matrix scaling. *Linear algebra and its applications*, 416(2-3):795–798, 2006.
- William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- Bahman Kalantari. A characterization theorem and an algorithm for a convex hull problem. *Annals of Operations Research*, 226(1):301–349, 2015.
- Narendra Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 302–311. ACM, 1984.
- Leonid G Khachiyan. Polynomial algorithms in linear programming. *USSR Computational Mathematics and Mathematical Physics*, 20(1):53–72, 1980.
- Daniel D Lee and H Sebastian Seung. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, pages 556–562, 2001.
- Jiří Matoušek and Otfried Schwarzkopf. Linear optimization queries. In *Proceedings of the eighth annual symposium on Computational geometry*, pages 16–25. ACM, 1992.

- Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607, 1996.
- Christos H Papadimitriou, Hisao Tamaki, Prabhakar Raghavan, and Santosh Vempala. Latent semantic indexing: A probabilistic analysis. In *Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 159–168. ACM, 1998.
- Daniel A Spielman, Huan Wang, and John Wright. Exact recovery of sparsely-used dictionaries. In *COLT*, pages 37–1, 2012.
- Keith Stevens, Philip Kegelmeyer, David Andrzejewski, and David Buttler. Exploring topic coherence over many models and many topics. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 952–961. Association for Computational Linguistics, 2012.
- Csaba D Toth, Joseph O’Rourke, and Jacob E Goodman. *Handbook of discrete and computational geometry*. CRC press, 2004.
- Godfried T Toussaint. Solving geometric problems with the rotating calipers. In *Proc. IEEE Melecon*, volume 83, page A10, 1983.
- Ky Vu, Pierre-Louis Poirion, and Leo Liberti. Random projections for linear programming. *arXiv preprint arXiv:1706.02768*, 2017.
- Limin Yao, David Mimno, and Andrew McCallum. Efficient methods for topic model inference on streaming document collections. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 937–946. ACM, 2009.
- Tong Zhang. Sequential greedy approximation for certain convex optimization problems. *IEEE Transactions on Information Theory*, 49(3):682–691, 2003.