
Sparse Approximate Conic Hulls

Gregory Van Buskirk, Benjamin Raichel, and Nicholas Ruoizzi

Department of Computer Science

University of Texas at Dallas

Richardson, TX 75080

{greg.vanbuskirk, benjamin.raichel, nicholas.ruozzi}@utdallas.edu

Abstract

We consider the problem of computing a restricted nonnegative matrix factorization (NMF) of an $m \times n$ matrix X . Specifically, we seek a factorization $X \approx BC$, where the k columns of B are a subset of those from X and $C \in \mathbb{R}_{\geq 0}^{k \times n}$. Equivalently, given the matrix X , consider the problem of finding a small subset, S , of the columns of X such that the conic hull of S ε -approximates the conic hull of the columns of X , i.e., the distance of every column of X to the conic hull of the columns of S should be at most an ε -fraction of the angular diameter of X . If k is the size of the smallest ε -approximation, then we produce an $O(k/\varepsilon^{2/3})$ sized $O(\varepsilon^{1/3})$ -approximation, yielding the first provable, polynomial time ε -approximation for this class of NMF problems, where also desirably the approximation is independent of n and m . Furthermore, we prove an approximate conic Carathéodory theorem, a general sparsity result, that shows that any column of X can be ε -approximated with an $O(1/\varepsilon^2)$ sparse combination from S . Our results are facilitated by a reduction to the problem of approximating convex hulls, and we prove that both the convex and conic hull variants are \mathbf{d} -SUM-hard, resolving an open problem. Finally, we provide experimental results for the convex and conic algorithms on a variety of feature selection tasks.

1 Introduction

Matrix factorizations of all sorts (SVD, NMF, CU, etc.) are ubiquitous in machine learning and computer science. In general, given an $m \times n$ matrix X , the goal is to find a decomposition into a product of two matrices $B \in \mathbb{R}^{m \times k}$ and $C \in \mathbb{R}^{k \times n}$ such that the Frobenius norm between X and BC is minimized. If no further restrictions are placed on the matrices B and C , this problem can be solved optimally by computing the singular value decomposition. However, imposing restrictions on B and C can lead to factorizations which are more desirable for reasons such as interpretability and sparsity. One of the most common restrictions is non-negative matrix factorization (NMF), requiring B and C to consist only of non-negative entries (see [Berry et al., 2007] for a survey). Practically, NMF has seen widespread usage as it often produces nice factorizations that are frequently sparse. Typically NMF is accomplished by applying local search heuristics, and while NMF can be solved exactly in certain cases (see [Arora et al., 2016]), in general NMF is not only NP-hard [Vavasis, 2009] but also \mathbf{d} -SUM-hard [Arora et al., 2016].

One drawback of factorizations such as SVD or NMF is that they can represent the data using a basis that may have no clear relation to the data. CU decompositions [Mahoney and Drineas, 2009] address this by requiring the basis to consist of input points. While it appears that the hardness of this problem has not been resolved, approximate solutions are known. Most notable is the additive approximation of Frieze et al. [2004], though more recently there have been advances on the multiplicative front [Drineas et al., 2008, Çivril and Magdon-Ismail, 2012, Guruswami and Sinop, 2012]. Similar restrictions have also been considered for NMF. Donoho and Stodden [2003] introduced a separability

assumption for NMF, and Arora et al. [2016] showed that a NMF can be computed in polynomial time under this assumption. Various other methods have since been proposed for NMF under the separability (or near separability) assumption [Recht et al., 2012, Kumar et al., 2013, Benson et al., 2014, Gillis and Vavasis, 2014, Zhou et al., 2014, Kumar and Sindhvani, 2015]. The separability assumption requires that there exists a subset S of the columns of X such that $X = X_S C$ for some nonnegative matrix C . This assumption can be restrictive in practice, e.g., when an exact subset does not exist but a close approximate subset does, i.e., $X \approx X_S C$. To our knowledge, no exact or approximate polynomial time algorithms have been proposed for the general problem of computing a NMF under only the restriction that the columns must be selected from those of X .

In this work, we fill this gap by arguing that a simple greedy algorithm can be used to provide a polynomial time ε -approximation algorithm for NMF under the column subset restriction. Note that the separability assumption is not required here: our theoretical analysis bounds the error of our selected columns versus the best possible columns that could have been chosen. The algorithm is based off of recent work on fast algorithms for approximately computing the convex hull of a set of points [Blum et al., 2016]. As in previous approaches [Donoho and Stodden, 2003, Kumar et al., 2013], we formulate restricted NMF geometrically as finding a subset, S , of the columns of the matrix X whose conic hull, the set of all nonnegative combinations of columns of S , well-approximates the conic hull of X . Using gnomonic projection, we reduce the conic hull problem to a convex hull problem and then apply the greedy strategy of Blum et al. [2016] to compute the convex hull of the projected points. Given a set of points P in \mathbb{R}^m , the convex hull of $S \subseteq P$, denoted $\text{Convex}(S)$, is said to ε -approximate $\text{Convex}(P)$ if the Hausdorff distance between $\text{Convex}(S)$ and $\text{Convex}(P)$ is at most $\varepsilon \cdot \text{diameter}(P)$. For a fixed $\varepsilon > 0$, suppose the minimum sized subset of P whose convex hull ε -approximates the convex hull of P has size k , then Blum et al. [2016] show that a simple greedy algorithm gives an $\varepsilon' = O(\varepsilon^{1/3})$ approximation using at most $k' = O(k/\varepsilon^{2/3})$ points of P , with an efficient $O(nc(m + c/\varepsilon^2 + c^2))$ running time, where $c = O(k_{opt}/\varepsilon^{2/3})$. By careful analysis, we show that our reduction achieves the same guarantees for the conic problem. (Note Blum et al. [2016] present other trade-offs between k' and ε' , which we argue carry to the conic case as well). Significantly, k' and ε' are independent of n and m , making this algorithm desirable for large high dimensional point sets. Note that our bounds on the approximation quality and the number of points do not explicitly depend on the dimension as they are relative to the size of the optimal solution, which itself may or may not depend on dimension. Like the X-RAY algorithm [Kumar et al., 2013], our algorithm is easy to parallelize, allowing it to be applied to large-scale problems.

In addition to the above ε -approximation algorithm, we also present two additional theoretical results of independent interest. The first theoretical contribution provides justification for empirical observations about the sparsity of NMF [Lee and Seung, 1999, Ding et al., 2010]. Due to the high dimensional nature of many data sets, there is significant interest in sparse representations requiring far fewer points than the dimension. Our theoretical justification for sparsity is based on Carathéodory’s theorem: any point q in the convex hull of P can be expressed as a convex combination of at most $m + 1$ points from P . This is tight in the worst case for exact representation, however the approximate Carathéodory theorem [Clarkson, 2010, Barman, 2015] states there is a point q' which is a convex combination of $O(1/\varepsilon^2)$ points of P (i.e., independent of n and m) such that $\|q - q'\| \leq \varepsilon \cdot \text{diameter}(P)$. This result has a long history with significant implications in machine learning, e.g., relating to the analysis of the perceptron algorithm [Novikoff, 1962], though the clean geometric statement of this theorem appears to be not well known outside the geometry community. Moreover, this approximation is easily computable with a greedy algorithm (e.g., [Blum et al., 2016]) similar to the Frank-Wolfe algorithm. The analogous statement for the linear case does not hold, so it is not immediately obvious whether such an approximate Carathéodory theorem should hold for the conic case, a question which we answer in the affirmative. As a second theoretical contribution, we address the question of whether or not the convex/conic hull problems are actually hard, i.e., whether approximations are actually necessary. We answer this question for both problems in the affirmative, resolving an open question of Blum et al. [2016], by showing both that the conic and convex problems are d-SUM-hard.

Finally, we evaluate the performance of the greedy algorithms for computing the convex and conic hulls on a variety of feature selection tasks against existing methods. We observe that, both the conic and convex algorithms perform well for a variety of feature selection tasks, though, somewhat surprisingly, the convex hull algorithm, for which previously no experimental results had been

produced, yields consistently superior results on text datasets. We use our theoretical results to provide intuition for these empirical observations.

2 Preliminaries

Let P be a point set in \mathbb{R}^m . For any $p \in P$, we interchangeably use the terms vector and point, depending on whether or not we wish to emphasize the direction from the origin. Let $\text{ray}(p)$ denote the unbounded ray passing through p , whose base lies at the origin. Let $\text{unit}(p)$ denote the unit vector in the direction of p , or equivalently $\text{unit}(p)$ is the intersection of $\text{ray}(p)$ with the unit hypersphere $\mathbb{S}^{(m-1)}$. For any subset $X = \{x_1, \dots, x_k\} \subseteq P$, $\text{ray}(X) = \{\text{ray}(x_1), \dots, \text{ray}(x_k)\}$ and $\text{unit}(X) = \{\text{unit}(x_1), \dots, \text{unit}(x_k)\}$.

Given points $p, q \in P$, let $d(p, q) = \|p - q\|$ denote their Euclidean distance, and let $\langle p, q \rangle$ denote their dot product. Let $\text{angle}(\text{ray}(p), \text{ray}(q)) = \text{angle}(p, q) = \cos^{-1}(\langle \text{unit}(p), \text{unit}(q) \rangle)$ denote the angle between the rays $\text{ray}(p)$ and $\text{ray}(q)$, or equivalently between vectors p and q . For two sets, $P, Q \subseteq \mathbb{R}^m$, we write $d(P, Q) = \min_{p \in P, q \in Q} d(p, q)$ and for a single point q we write $d(q, P) = d(\{q\}, P)$, and the same definitions apply to $\text{angle}(\cdot)$.

For any subset $X = \{x_1, \dots, x_k\} \subseteq P$, let $\text{Convex}(X) = \{\sum_i \alpha_i x_i \mid \alpha_i \geq 0, \sum_i \alpha_i = 1\}$ denote the convex hull of X . Similarly, let $\text{Conic}(X) = \{\sum_i \alpha_i x_i \mid \alpha_i \geq 0\}$ denote the conic hull of X and $\text{DualCone}(X) = \{z \in X \mid \langle x, z \rangle \geq 0 \forall x \in X\}$ the dual cone. For any point $q \in \mathbb{R}^m$, the *projection* of q onto $\text{Convex}(X)$ is the closest point to q in $\text{Convex}(X)$, $\text{proj}(q) = \text{proj}(q, \text{Convex}(X)) = \arg \min_{x \in \text{Convex}(X)} d(q, x)$. Similarly the *angular projection* of q onto $\text{Conic}(X)$ is the angularly closest point to q in $\text{Conic}(X)$, $\text{aproj}(q) = \text{aproj}(q, \text{Conic}(X)) = \arg \min_{x \in \text{Conic}(X)} \text{angle}(q, x)$. Note that angular projection defines an entire ray of $\text{Conic}(X)$, rather than a single point, which without loss of generality we choose the point on the ray minimizing the Euclidean distance to q . In fact, abusing notation, we sometimes equivalently view $\text{Conic}(X)$ as a set of rays rather than points, in which case $\text{aproj}(\text{ray}(q)) = \text{aproj}(q)$ is the entire ray.

For $X \subset \mathbb{R}^m$, let $\Delta = \Delta_X = \max_{p, q \in X} d(p, q)$ denote the *diameter* of X . The *angular diameter* of X is $\phi = \phi_X = \max_{p, q \in X} \text{angle}(p, q)$. Similarly $\phi_X(q) = \max_{p \in X} \text{angle}(p, q)$ denotes the angular radius of the minimum radius cone centered around the ray through q and containing all of P .

Definition 2.1. Consider a subset X of a point set $P \subset \mathbb{R}^m$. X is an ε -**approximation** to $\text{Convex}(P)$ if $d_{\text{convex}}(X, P) = \max_{p \in \text{Convex}(P)} d(p, \text{Convex}(X)) \leq \varepsilon \Delta$. Note $d_{\text{convex}}(X, P)$ is the Hausdorff distance between $\text{Convex}(X)$ and $\text{Convex}(P)$. Similarly X is an ε -*approximation* to $\text{Conic}(P)$ if $d_{\text{conic}}(X, P) = \max_{p \in \text{Conic}(P)} \text{angle}(p, \text{Conic}(X)) \leq \varepsilon \phi_P$.

Note that the definition of ε -*approximation* for $\text{Conic}(P)$ uses angular rather than Euclidean distance in order to be defined for rays, i.e., scaling a point outside the conic hull changes its Euclidean distance but its angular distance is unchanged since its ray stays the same. Thus we find considering angles better captures what it means to approximate the conic hull than the distance based Frobenius norm which is often used to evaluate the quality of approximation for NMF.

As we are concerned only with angles, without loss of generality we often will assume that all points in the input set P have been scaled to have unit length, i.e., $P = \text{unit}(P)$. In our theoretical results, we will always assume that $\phi_P < \pi/2$. Note that if P lies in the non-negative orthant, then for any strictly positive q , $\phi_P(q) < \pi/2$. In the case that the P is not strictly inside the positive orthant, the points can be uniformly translated a small amount to ensure that $\phi_P < \pi/2$.

3 A Simple Greedy Algorithm

Let P be a finite point set in \mathbb{R}^m (with unit lengths). Call a point $p \in P$ *extreme* if it lies on the boundary of the conic hull (resp. convex hull). Observe that for any $X \subseteq P$, containing all the extreme points, it holds that $\text{Conic}(X) = \text{Conic}(P)$ (resp. $\text{Convex}(X) = \text{Convex}(P)$). Consider the simple greedy algorithm which builds a subset of points S , by iteratively adding to S the point angularly furthest from the conic hull of the current point set S (for the convex hull take the furthest point in distance). One can argue in each round this algorithm selects an extreme point, and thus can be used to find a subset of points whose hull captures that of P . Note if the hull is not degenerate, i.e.,

no point on the boundary is expressible as a combination of other points on the boundary, then this produces the minimum sized subset capturing P . Otherwise, one can solve a recursive subproblem as discussed by Kumar et al. [2013] to exactly recover S .

Here instead we consider finding a small subset of points (potentially much smaller than the number of extreme points) to approximate the hull. The question is then whether this greedy approach still yields a reasonable solution, which is not clear as there are simple examples showing the best approximate subset includes non-extreme points. Moreover, arguing about the conic approximation directly is challenging as it involves angles and hence spherical (rather than planar) geometry. For the convex case, Blum et al. [2016] argued that this greedy strategy does yield a good approximation. Thus we seek a way to reduce our conic problem to an instance of the convex problem, without introducing too much error in the process, which brings us to the *gnomonic projection*. Let $\text{hplane}(q)$ be the hyperplane defined by the equation $\langle (q - x), q \rangle = 0$ where $q \in \mathbb{R}^m$ is a unit length normal vector. The *gnomonic projection* of P onto $\text{hplane}(q)$, is defined as $\text{gp}^q(P) = \{\text{ray}(P) \cap \text{hplane}(q)\}$ (see Figure 3.1). Note that $\text{gp}^q(q) = q$. For any point x in $\text{hplane}(q)$, the inverse gnomonic projection is $\text{pg}^q(x) = \text{ray}(x) \cap \mathbb{S}^{(m-1)}$. Similar to other work [Kumar et al., 2013], we allow projections onto any hyperplane tangent to the unit hypersphere with normal q in the strictly positive orthant.

A key property of the gnomonic projection, is that the problem of finding the extreme points of the convex hull of the projected points is equivalent to finding the extreme points of the conic hull of P . (Additional properties of the gnomonic projection are discussed in the full version.) Thus the strategy to approximate the conic hull should now be clear. Let $P' = \text{gp}^q(P)$. We apply the greedy strategy of Blum et al. [2016] to P' to build a set of extreme points S , by iteratively adding to S the point furthest from the convex hull of the current point set S . This procedure is shown in Algorithm 1.

We show that Algorithm 1 can be used to produce an ε -approximation to the restricted NMF problem. Formally, for $\varepsilon > 0$, let $\text{opt}(P, \varepsilon)$ denote any minimum cardinality subset $X \subseteq P$ which ε -approximates $\text{Conic}(P)$, and let $k_{\text{opt}} = |\text{opt}(P, \varepsilon)|$. We consider the following problem.

Problem 3.1. *Given a set P of n points in \mathbb{R}^m such that $\phi_P \leq \pi/2 - \gamma$, for a constant $\gamma > 0$, and a value $\varepsilon > 0$, compute $\text{opt}(P, \varepsilon)$.*

Alternatively one can fix k rather than ε , defining $\text{opt}(P, k) = \arg \min_{X \subseteq P, |X|=k} d_{\text{conic}}(X, P)$ and $\varepsilon_{\text{opt}} = d_{\text{conic}}(\text{opt}(P, k), P)$. Our approach works for either variant, though here we focus on the version in Problem 3.1. Note the bounded angle assumption applies to any collection of points in the strictly positive orthant (a small translation can be used to ensure this for any nonnegative data set).

In this section we argue Algorithm 1 produces an (α, β) -approximation to an instance (P, ε) of Problem 3.1, that is a subset $X \subseteq P$ such that $d_{\text{conic}}(X, P) \leq \alpha$ and $|X| \leq \beta \cdot k_{\text{opt}} = \beta \cdot |\text{opt}(P, \varepsilon)|$.

For $\varepsilon > 0$, similarly define $\text{opt}^{\text{convex}}(P, \varepsilon)$ to be any minimum cardinality subset $X \subseteq P$ which ε -approximates $\text{Convex}(P)$. Blum et al. [2016] gave (α, β) -approximation for the following.

Problem 3.2. *Given a set P of n points in \mathbb{R}^m , and a value $\varepsilon > 0$, compute $\text{opt}^{\text{convex}}(P, \varepsilon)$.*

Note the proofs of correctness and approximation quality from Blum et al. [2016] for Problem 3.2 do not immediately imply the same results for using Algorithm 1 for Problem 3.1. To see this, consider any points u, v on $\mathbb{S}^{(m-1)}$. Note the angle between u and v is the same as their geodesic distance on $\mathbb{S}^{(m-1)}$. Intuitively, we want to claim the geodesic distance between u and v is roughly the same as the Euclidean distance between $\text{gp}^q(u)$ and $\text{gp}^q(v)$. While this is true for points near q , as we move away from q the correspondence breaks down (and is unbounded as you approach $\pi/2$). This non-uniform distortion requires care, and thus the proofs had to be moved to the full version.

Finally, observe that Algorithm 1, requires being able to compute the point furthest from the convex hull. To do so we use the (convex) approximate Carathéodory, which is both theoretically and practically very efficient, and produces provably sparse solutions. As a stand alone result, we first prove the conic analog of the approximate Carathéodory theorem. This result is of independent interest since it can be used to sparsify the returned solution from Algorithm 1, or any other algorithm.

3.1 Sparsity and the Approximate Conic Carathéodory Theorem

Our first result is a conic approximate Carathéodory theorem. That is, given a point set $P \subseteq \mathbb{R}^m$ and a query point q , then the angularly closest point to q in $\text{Conic}(P)$ can be approximately expressed as

Algorithm 1: Greedy Conic Hull

Data: A set of n points, P , in \mathbb{R}^m such that $\phi_P < \pi/2$, a positive integer k , and a normal vector q in $\text{DualCone}(P)$.

Result: $S \subseteq P$ such that $|S| = k$

$Y \leftarrow \text{gp}^q(P)$;

Select an arbitrary starting point $p^0 \in Y$;

$S \leftarrow \{p^0\}$;

for $i = 2$ **to** k **do**

Select

$p^* \in \arg \max_{p \in Y} d_{\text{convex}}(p, S)$;

$S \leftarrow S \cup \{p^*\}$;

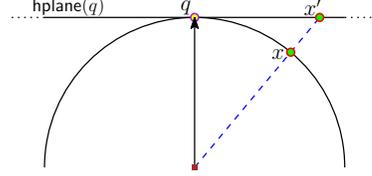


Figure 3.1: Side view of gnomonic projection.

a sparse combination of point from P . More precisely, one can compute a point t which is a conic combination of $O(1/\varepsilon^2)$ points from P such that $\text{angle}(q, t) \leq \text{angle}(q, \text{Conic}(P)) + \varepsilon\phi_P$.

The significance of this result is as follows. Recall that we seek a factorization $X \approx BC$, where the k columns of B are a subset of those from X and the entries of C are non-negative. Ideally each point in X is expressed as a sparse combination from the basis B , that is each column of C has very few non-zero entries. So suppose we are given any factorization BC , but C is dense. Then no problem, just throw out C , and use our Carathéodory theorem to compute a new matrix C' with sparse columns. Namely treat each column of X as the query q and run the theorem for the point set $P = B$, and then the non-zero entries of corresponding column of C' are just the selected combination from B . Not only does this mean we can sparsify any solution to our NMF problem (including those obtained by other methods), but it also means conceptually that rather than finding a good pair BC , one only needs to focus on finding the subset B , as is done in [Algorithm 1](#). Note that [Algorithm 1](#) allows non-negative inputs in P because $\phi_P < \pi/2$ ensures P can be rotated into the positive orthant.

While it appears the conic approximate Carathéodory theorem had not previously been stated, the convex version has a long history (e.g., implied by [[Novikoff, 1962](#)]). The algorithm to compute this sparse convex approximation is again a simple and fast greedy algorithm, which roughly speaking is a simplification of the Frank-Wolfe algorithm for this particular problem. Specifically, to find the projection of q onto $\text{Convex}(P)$, start with any point $t_0 \in \text{Convex}(P)$. In the i th round, find the point $p_i \in P$ most extreme in the direction of q from t_{i-1} (i.e., maximizing $\langle q - t_{i-1}, p_i \rangle$) and set t_i to be the closest point to q on the segment $t_{i-1}p_i$ (thus simplifying Frank Wolfe, as we ignore step size issues). The standard analysis of this algorithm (e.g., [[Blum et al., 2016](#)]) gives the following.

Theorem 3.3 (Convex Carathéodory). *For a point set $P \subseteq \mathbb{R}^m$, $\varepsilon > 0$, and $q \in \mathbb{R}^m$, one can compute, in $O(|P|m/\varepsilon^2)$ time, a point $t \in \text{Convex}(P)$, such that $d(q, t) \leq d(q, \text{Convex}(P)) + \varepsilon\Delta$, where $\Delta = \Delta_P$. Furthermore, t is a convex combination of $O(1/\varepsilon^2)$ points of P .*

Again by exploiting properties of the gnomonic projection we are able to prove a conic analog of the above theorem. Note for $P \subset \mathbb{R}^m$, P is contained in the linear span of at most m points from P , and similarly the exact Carathéodory theorem states any point $q \in \text{Convex}(P)$ is expressible as a convex combination of at most $m + 1$ points from P . As the conic hull lies between the linear case (with all combinations) and the convex case (with non-negative combinations summing to one), it is not surprising an exact conic Carathéodory theorem holds. However, the linear analog of the approximate convex Carathéodory theorem does not hold, and so the following conic result is not a priori obvious.

Theorem 3.4. *Let $P \subset \mathbb{R}^m$ be a point set, let q be such that $\phi_P(q) < \pi/2 - \gamma$ for some constant $\gamma > 0$, and let $\varepsilon > 0$ be a parameter. Then one can find, in $O(|P|m/\varepsilon^2)$ time, a point $t \in \text{Conic}(P)$ such that $\text{angle}(q, t) \leq \text{angle}(q, \text{Conic}(P)) + \varepsilon\phi_P(q)$. Moreover, t is a conic combination of $O(1/\varepsilon^2)$ points from P .*

Due to space constraints, the detailed proof of [Theorem 3.4](#) appears in the full version. In the proof, the dependence on γ is made clear but we make a remark about it here. If ε is kept fixed, γ shows up

in the running time roughly by a factor of $\tan^2(\pi/2 - \gamma)$. Alternatively, if the running time is fixed, the approximation error will roughly depend on the factor $1/\tan(\pi/2 - \gamma)$.

We now give a simple example of a high dimensional point set which shows our bounded angle assumption is required for the conic Carathéodory theorem to hold. Let P consist of the standard basis vectors in \mathbb{R}^m , let q be the all ones vector, and let ε be a parameter. Let X be a subset of P of size k , and consider $\text{aproj}(q) = \text{aproj}(q, X)$. As P consists of basis vectors, each of which have all but one entry set to zero, $\text{aproj}(q)$ will have at most k non-zero entries. By the symmetry of q it is also clear that all non-zero entries in $\text{aproj}(q)$ should have the same value. Without loss of generality assume that this value is 1, and hence the magnitude of $\text{aproj}(q)$ is \sqrt{k} . Thus for $\text{aproj}(q)$ to be an ε -approximation to q , $\text{angle}(\text{aproj}(q), q) = \cos^{-1}(\frac{k}{\sqrt{k}\sqrt{m}}) = \cos^{-1}(\sqrt{k/m}) < \varepsilon$. Hence for a fixed ε , the number of points required to ε -approximate q depends on m , while the conic Carathéodory theorem should be independent of m .

3.2 Approximating the Conic Hull

We now prove that [Algorithm 1](#) yields an approximation to the conic hull of a given point set and hence an approximation to the nonnegative matrix factorization problem. As discussed above, previously [Blum et al. \[2016\]](#) provided the following (α, β) -approximation for [Problem 3.2](#).

Theorem 3.5 ([Blum et al., 2016]). *For a set P of n points in \mathbb{R}^m , and $\varepsilon > 0$, the greedy strategy, which iteratively adds the point furthest from the current convex hull, gives a $((8\varepsilon^{1/3} + \varepsilon)\Delta, O(1/\varepsilon^{2/3}))$ -approximation to [Problem 3.2](#), and has running time $O(nc(m + c/\varepsilon^2 + c^2))$ time, where $c = O(k_{\text{opt}}/\varepsilon^{2/3})$.*

Our second result, is a conic analog of the above theorem.

Theorem 3.6. *Given a set P of n points in \mathbb{R}^m such that $\phi_P \leq \frac{\pi}{2} - \gamma$ for a constant $\gamma > 0$, and a value $\varepsilon > 0$, [Algorithm 1](#) gives an $((8\varepsilon^{1/3} + \varepsilon)\phi_P, O(1/\varepsilon^{2/3}))$ -approximation to [Problem 3.1](#), and has running time $O(nc(m + c/\varepsilon^2 + c^2))$, where $c = O(k_{\text{opt}}/\varepsilon^{2/3})$.*

Bounding the approximation error requires carefully handling the distortion due to the gnomonic project, and the details are presented in the full version. Additionally, [Blum et al. \[2016\]](#) provide other (α, β) -approximations, for different values of α and β , and in the full version these other results are also shown to hold for the conic case.

4 Hardness of the Convex and Conic Problems

This section gives a reduction from **d-SUM** to the convex approximation of [Problem 3.2](#), implying it is **d-SUM-hard**. In the full version a similar setup is used to argue the conic approximation of [Problem 3.1](#) is **d-SUM-hard**. Actually if [Problem 3.1](#) allowed instances where $\phi_P = \pi/2$ the reduction would be virtually the same. However, arguing that the problem remains hard under our requirement that $\phi_P \leq \pi/2 - \gamma$, is non-trivial and some of the calculations become challenging and lengthy. The reductions to both problems are partly inspired by [Arora et al. \[2016\]](#). However, here, we use the somewhat non-standard version of **d-SUM** where repetitions are allowed as described below.

Problem 4.1 (d-SUM). *In the d-SUM problem we are given a set $S = \{s_1, s_2, \dots, s_N\}$ of N values, each in the interval $[0, 1]$, and the goal is to determine if there is a set of d numbers (not necessarily distinct) whose sum is exactly $d/2$.*

It was shown by [Patrascu and Williams \[2010\]](#) that if **d-SUM** can be solved in $N^{o(d)}$ time then **3-SAT** has a sub-exponential time algorithm, i.e., that the Exponential Time Hypothesis is false.

Theorem 4.2 (d-SUM-hard). *Let $d < N^{0.99}$, $\delta < 1$. If **d-SUM** on N numbers of $O(d \log(N))$ bits can be solved in $O(N^{\delta d})$ time, then **3-SAT** on n variables can be solved in $2^{o(n)}$ time.*

We will prove the following decision version of **Problem 3.2** is **d-SUM-hard**. Note in this section the dimension will be denoted by d rather than m , as this is standard for **d-SUM** reductions.

Problem 4.3. *Given a set P of n points in \mathbb{R}^d , a value $\varepsilon > 0$, and an integer k , is there a subset $X \subseteq P$ of k points such that $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$, where Δ is the diameter of P .*

Given an instance of **d-SUM** with N values $S = \{s_1, s_2, \dots, s_N\}$ we construct an instance of **Problem 4.3** where $P \subset \mathbb{R}^{d+2}$, $k = d$, and $\varepsilon = 1/3$ (or any sufficiently small value). The idea is to create d clusters each containing N points corresponding to a choice of one of the s_i values. The clusters are positioned such that exactly one point from each cluster must be chosen. The $d + 2$ coordinates are labeled a_i for $i \in [d]$, w , and v . Together, a_1, \dots, a_d determine the cluster. The w dimension is used to compute the sum of the chosen s_i values. The v dimension is used as a threshold to determine whether **d-SUM** is a yes or no instance to **Problem 4.3**. Let $w(p_j)$ denote the w value of an arbitrary point p_j .

We assume $d \geq 2$ as **d-SUM** is trivial for $d = 1$. Let $e_1, e_2, \dots, e_d \in \mathbb{R}^d$ be the standard basis in \mathbb{R}^d , $e_1 = (1, \dots, 0)$, $e_2 = (0, 1, \dots, 0)$, \dots , and $e_d = (0, \dots, 1)$. Together they form the unit d -simplex, and they define the d clusters in the construction. Finally, let $\Delta^* = \sqrt{2 + (\varepsilon s_{\max} - \varepsilon s_{\min})^2}$ be a constant where s_{\max} and s_{\min} are, respectively, the maximum and minimum values in S .

Definition 4.4. The set of points $P \subset \mathbb{R}^{d+2}$ are the following

- p_j^i points: For each $i \in [d]$, $j \in [N]$, set $(a_1, \dots, a_d) = e_i$, $w = \varepsilon s_j$ and $v = 0$
- q point: For each $i \in [d]$, $a_i = 1/d$, $w = \varepsilon/2$, $v = 0$
- q' point: For each $i \in [d]$, $a_i = 1/d$ and $w = \varepsilon/2$, $v = \varepsilon\Delta^*$

Lemma 4.5 (Proof in full version). *The diameter of P , Δ_P , is equal to Δ^* .*

We prove completeness and soundness of the reduction. Below $P^i = \cup_j p_j^i$ denotes the i th cluster.

Lemma 4.6 (Completeness). *If there is a subset $\{s_{k_1}, s_{k_2}, \dots, s_{k_d}\}$ of d values (not necessarily distinct) such that $\sum_{i \in [d]} s_{k_i} = d/2$, then the above described instance of **Problem 4.3** is a true instance, i.e. there is a d sized subset $X \subseteq P$ with $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$.*

Proof: For each value s_{k_i} consider the point $x_i = (e_i, \varepsilon \cdot s_{k_i}, 0)$, which by **Definition 4.4** is a point in P . Let $X = \{x_1, \dots, x_d\}$. We now prove $\max_{p \in P} d(p, \text{Convex}(X)) \leq \varepsilon\Delta$, which by **Observation C.3** implies that $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$.

First observe that for any p_j^i in P , $d(p_j^i, x_i) = \sqrt{(w(p_j^i) - w(x_i))^2} \leq |\varepsilon s_j - \varepsilon s_{k_i}| \leq \varepsilon\Delta$. The only other points in P are q and q' . Note that $d(q, q') = \varepsilon\Delta^* = \varepsilon\Delta$ from **Lemma 4.5**. Thus if we can prove that $q \in \text{Convex}(X)$ then we will have $\max_{p \in P} d(p, \text{Convex}(X)) \leq \varepsilon\Delta$. Specifically, we prove that the convex combination $x = \frac{1}{d} \sum_i x_i$ is the point q . As X contains exactly one point from each set P^i , and in each such set all points have $a_i = 1$ and all other $a_j = 0$, it holds that x has $1/d$ for all the a coordinates. All points in X have $v = 0$ and so this holds for x as well. Thus we only need to verify that $w(x) = w(q) = \varepsilon/2$, for which we have $w(x) = \frac{1}{d} \sum_i w(x_i) = \frac{1}{d} \sum_i \varepsilon s_{k_i} = \frac{1}{d} (\varepsilon d/2) = \varepsilon/2$. ■

Proving soundness requires some helper lemmas. Note that in the above proof we constructed a solution to **Problem 4.3** that selected exactly one point from each cluster P^i . We now prove that this is a required property.

Lemma 4.7 (Proof in full version). *Let $P \subset \mathbb{R}^{d+2}$ be as defined above, and let $X \subseteq P$ be a subset of size d . If $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$, then for all i , X contains exactly one point from P^i .*

Lemma 4.8 (Proof in full version). *If $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$, then $q \in \text{Convex}(X)$ and moreover $q = \frac{1}{d} \sum_{x_i \in X} x_i$.*

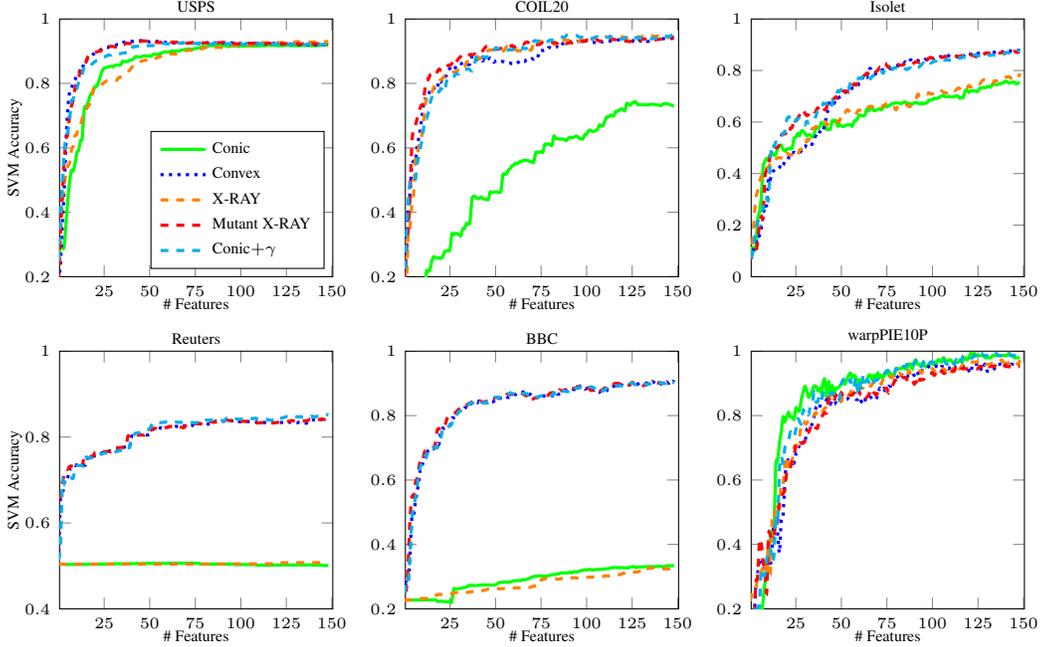


Figure 4.1: Experimental results for feature selection on six different data sets. Best viewed in color.

Lemma 4.9 (Soundness). *Let P be an instance of Problem 4.3 generated from a d -SUM instance S , as described in Definition 4.4. If there is a subset $X \subseteq P$ of size d such that $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$, then there is a choice of d values from S that sum to exactly $d/2$.*

Proof: From Lemma 4.7 we know that X consist of exactly one point from each cluster P^i . Thus for each $x_i \in X$, $w(x_i) = \varepsilon s_{k_i}$ for some $s_{k_i} \in S$. By Lemma 4.8, $q = \frac{1}{d} \sum_i x_i$, which implies $w(q) = \frac{1}{d} \sum_i w(x_i)$. By Definition 4.4 $w(q) = \varepsilon/2$, which implies $\varepsilon/2 = \frac{1}{d} \sum_i w(x_i) = \frac{1}{d} \sum_i \varepsilon s_{k_i}$. Thus we have a set $\{s_{k_1}, \dots, s_{k_d}\}$ of d values from S such that $\sum_i s_{k_i} = d/2$. ■

Lemma 4.6 and Lemma 4.9 immediately imply the following.

Theorem 4.10. *For point sets in \mathbb{R}^{d+2} , Problem 4.3 is d -SUM-hard.*

5 Experimental Results

We report an experimental comparison of the proposed greedy algorithm for conic hulls, the greedy algorithm for convex hulls (the conic hull algorithm without the projection step) [Blum et al., 2016], the X-RAY (max) algorithm [Kumar et al., 2013], a modified version of X-RAY, dubbed mutant X-RAY, which simply selects the point furthest away from the current cone (i.e., with the largest residual), and a γ -shifted version of the conic hull algorithm described below. Other methods such as Hottopixx [Recht et al., 2012, Gillis and Luce, 2014] and SPA [Gillis and Vavasis, 2014] were not included due to their similar performance to the above methods. For our experiments, we considered the performance of each of the methods when used to select features for a variety of SVM classification tasks on various image, text, and speech data sets including several from the Arizona State University feature selection repository [Li et al., 2016] as well as the UCI Reuters dataset and the BBC News dataset [Greene and Cunningham, 2006]. The Reuters and BBC text datasets are represented using the TF-IDF representation. For the Reuters dataset, only the ten most frequent topics were used for classification. In all datasets, columns (corresponding to features) that were identically equal to zero were removed from the data matrix.

For each problem, the data is divided using a 30/70 train/test split, the features are selected by the indicated method, and then an SVM classifier is trained using only the selected features. For the conic

and convex hull methods, ϵ is set to 0.1. The accuracy (percent of correctly classified instances) is plotted versus the number of selected features for each method in Figure 4.1. Additional experimental results can be found in the full version. Generally speaking, the convex, mutant X-RAY, and shifted conic algorithms seem to consistently perform the best on the tasks. The difference in performance between convex and conic is most striking on the two text data sets Reuters and BBC. In the case of BBC and Reuters, this is likely due to the fact that many of the columns of the TF-IDF matrix are orthogonal. We note that the quality of both X-RAY and conic is improved if thresholding is used when constructing the feature matrix, but they still seem to under perform the convex method for text datasets.

The text datasets are also interesting as not only do they violate the explicit assumption in our theorems that the angular diameter of the conic hull be strictly less than $\pi/2$, but that there are many such mutually orthogonal columns of the document-feature matrix. This observation motivates the γ -shifted version of the conic hull algorithm that simply takes the input matrix X and adds γ to all of the entries (essentially translating the data along the all ones vector) and then applies the conic hull algorithm. Let $1^{a,b}$ denote the $a \times b$ matrix of ones. After a nonnegative shift, the angular assumption is satisfied, and the restricted NMF problem is that of approximating $(X + \gamma 1^{m,n})$ as $(B + \gamma 1^{m,k})C$, where the columns of B are again chosen from those of X . Under the Frobenus norm $\|(X + \gamma 1^{m,n}) - (B + \gamma 1^{m,k})C\|_2^2 = \sum_{i,j} (X_{ij} - B_{i,:}C_{:,j} + \gamma(1 - \|C_{:,j}\|_1))^2$. As C must be a nonnegative matrix, the shifted conic case acts like the original conic case plus a penalty that encourages the columns of C to sum to one (i.e., it is a hybrid between the conic case and the convex case). The plots illustrate the performance of the γ -shifted conic hull algorithm for $\gamma = 10$. After the shift, the performance more closely matches that of the convex and mutant X-RAY methods on TF-IDF features.

Given these experimental results and the simplicity of the proposed convex and conic methods, we suggest that both methods should be added to practitioners’ toolboxes. In particular, the superior performance of the convex algorithm on text datasets, compared to X-RAY and the conic algorithm, seems to suggest that these types of “convex” factorizations may be more desirable for TF-IDF features.

Acknowledgments

Greg Van Buskirk and Ben Raichel were partially supported by NSF CRII Award-1566137. Nicholas Ruozi was partially supported by DARPA Explainable Artificial Intelligence Program under contract number N66001-17-2-4032 and NSF grant III-1527312

References

- M. Berry, M. Browne, A. Langville, V. Pauca, and R. Plemmons. Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics & Data Analysis*, 52(1): 155–173, 2007.
- S. Arora, R. Ge, R. Kannan, and A. Moitra. Computing a nonnegative matrix factorization - provably. *SIAM J. Comput.*, 45(4):1582–1611, 2016.
- S. Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2009.
- M. Mahoney and P. Drineas. Cur matrix decompositions for improved data analysis. *Proceedings of the National Academy of Sciences*, 106(3):697–702, 2009.
- A. Frieze, R. Kannan, and S. Vempala. Fast monte-carlo algorithms for finding low-rank approximations. *J. ACM*, 51(6):1025–1041, 2004.
- P. Drineas, M. Mahoney, and S. Muthukrishnan. Relative-error CUR matrix decompositions. *SIAM J. Matrix Analysis Applications*, 30(2):844–881, 2008.
- A. Çivril and M. Magdon-Ismail. Column subset selection via sparse approximation of SVD. *Theor. Comput. Sci.*, 421:1–14, 2012.

- V. Guruswami and A. Sinop. Optimal column-based low-rank matrix reconstruction. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1207–1214, 2012.
- D. L. Donoho and V. Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in Neural Information Processing Systems (NIPS)*, 2003.
- B. Recht, C. Re, J. Tropp, and V. Bittorf. Factoring nonnegative matrices with linear programs. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1214–1222, 2012.
- A. Kumar, V. Sindhwani, and P. Kambadur. Fast conical hull algorithms for near-separable non-negative matrix factorization. In *International Conference on Machine Learning (ICML)*, pages 231–239, 2013.
- A. R. Benson, J. D. Lee, B. Rajwa, and D. F. Gleich. Scalable methods for nonnegative matrix factorizations of near-separable tall-and-skinny matrices. In *Advances in Neural Information Processing Systems (NIPS)*, pages 945–953, 2014.
- N. Gillis and S. A. Vavasis. Fast and robust recursive algorithms for separable nonnegative matrix factorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(4):698–714, 2014.
- T. Zhou, J. A. Bilmes, and C. Guestrin. Divide-and-conquer learning by anchoring a conical hull. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1242–1250, 2014.
- A. Kumar and V. Sindhwani. *Near-separable Non-negative Matrix Factorization with l_1 and Bregman Loss Functions*, pages 343–351. 2015.
- A. Blum, S. Har-Peled, and B. Raichel. Sparse approximation via generating point sets. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 548–557, 2016.
- D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.
- C. H. Q. Ding, T. Li, and M. I. Jordan. Convex and semi-nonnegative matrix factorizations. *IEEE transactions on pattern analysis and machine intelligence*, 32(1):45–55, 2010.
- K. L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. 6(4), 2010.
- S. Barman. Approximating nash equilibria and dense bipartite subgraphs via an approximate version of caratheodory’s theorem. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC)*, pages 361–369, 2015.
- A.B.J. Novikoff. On convergence proofs on perceptrons. In *Proc. Symp. Math. Theo. Automata*, volume 12, pages 615–622, 1962.
- M. Patrascu and R. Williams. On the possibility of faster SAT algorithms. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1065–1075, 2010.
- Nicolas Gillis and Robert Luce. Robust near-separable nonnegative matrix factorization using linear optimization. *Journal of Machine Learning Research*, 15(1):1249–1280, 2014.
- J. Li, K. Cheng, S. Wang, F. Morstatter, T. Robert, J. Tang, and H. Liu. Feature selection: A data perspective. *arXiv:1601.07996*, 2016.
- Derek Greene and Pádraig Cunningham. Practical solutions to the problem of diagonal dominance in kernel document clustering. In *Proceedings of the 23rd international conference on Machine learning*, pages 377–384. ACM, 2006.

A Gnomonic Projection

Here we provide some basic facts about the gnomonic projection. The lemma below follows by elementary trigonometry (see [Figure 3.1](#)).

Observation A.1. For any $q \in \mathbb{S}^{(m-1)}$, the function $\text{gp}^q(\cdot)$ defines a one to one correspondence between $\text{hplane}(q)$ and the part of the hypersphere at angle $< \pi/2$ from q , i.e., the gnomonic projection projects the hemisphere defined by positive dot product with q onto $\text{hplane}(q)$.

Lemma A.2. Let q, x be any two points on $\mathbb{S}^{(m-1)}$ such that $\text{angle}(q, x) < \pi/2$. By the cosine formula for the dot product, $\text{gp}^q(x) = \frac{x}{\langle x, q \rangle}$. Moreover, for any point x' in $\text{hplane}(q)$, by considering the triangle defined by the origin, q , and x' , the definition of the tangent function gives $\tan(\text{angle}(q, x')) = d(q, x')$.

In particular, for any point set $P \subseteq \mathbb{S}^{(m-1)}$ such that $\phi_P(q) < \pi/2$, $\text{diameter}(\text{gp}^q(P)) \leq 2 \tan(\phi_P(q))$, and notably this holds for $q = p$ for any $p \in P$ when $\phi_P < \pi/2$.

As $\tan(\theta)$ is monotonically increasing for $0 \leq \theta < \pi/2$, we have the following corollary.

Corollary A.3. Let q, x, y be points on $\mathbb{S}^{(m-1)}$ such that $\text{angle}(q, x), \text{angle}(q, y) < \pi/2$. Then $\text{angle}(q, x) \leq \text{angle}(q, y)$ if and only if $d(q, \text{gp}^q(x)) \leq d(q, \text{gp}^q(y))$. Conversely, for any two points $x', y' \in \text{hplane}(q)$, $d(q, x') \leq d(q, y')$ if and only if $\text{angle}(q, \text{pg}^q(x')) \leq \text{angle}(q, \text{pg}^q(y'))$.

The proofs of the following lemmas are straightforward and are provided for completeness.

Lemma A.4. Given a point set $P \subset \mathbb{R}^m$, $\text{Conic}(P) = \text{ray}(\text{Convex}(P))$.

Proof: First note that any $x \in \text{Conic}(P)$ is a scaling of some point $z \in \text{Convex}(P)$, i.e. $x = \lambda z$ for some value λ . To see this, let $x = \sum \alpha_i p_i$ such that $\alpha_i \geq 0$. Now let $\lambda = \sum \alpha_i$ and $z = \frac{1}{\lambda} \sum \alpha_i p_i$. The other direction can be shown in a similar way. ■

Lemma A.5. Given a point set $P \subset \mathbb{R}^m$ and a point $q \in \mathbb{S}^{(m-1)}$ such that $\phi_P(q) < \pi/2$, it holds that $\text{Conic}(P) \cap \text{hplane}(q) = \text{Convex}(\text{gp}^q(P))$

Proof: Let $P = \{p_1, \dots, p_n\}$. As $\phi_P(q) < \pi/2$, by [Observation A.1](#), each p_i defines a unique $p'_i = \text{gp}^q(p_i)$ in $\text{hplane}(q)$, and by [Lemma A.2](#), $p'_i = p_i / \langle p_i, q \rangle$. Let $\text{gp}^q(P) = P' = \{p'_1, \dots, p'_n\}$. First we show $\text{Conic}(P) \cap \text{hplane}(q) \supseteq \text{Convex}(P')$. Consider any point in $x' \in \text{Convex}(P')$. Note that $x' = \sum_i \alpha_i p'_i = \sum_i \frac{\alpha_i}{\langle p_i, q \rangle} p_i$, where $\sum_i \alpha_i = 1$ and $\alpha_i \geq 0$ for all i . Thus $x' \in \text{Conic}(P)$ as the coefficients $\frac{\alpha_i}{\langle p_i, q \rangle}$ are non-negative. Moreover, $x' \in \text{hplane}(q)$ as it is a convex combination of points in $\text{hplane}(q)$.

Now we argue $\text{Conic}(P) \cap \text{hplane}(q) \subseteq \text{Convex}(P')$. Consider any point x in $\text{Convex}(P)$, that is $x = \sum_i \alpha_i p_i$, where $\sum_i \alpha_i = 1$ and $\alpha_i \geq 0$ for all i . First observe since x is a convex combination of the p_i , each of which has positive dot product with q , by linearity of the dot product, so does x , and hence $\text{gp}^q(x)$ is a well defined point on $\text{hplane}(q)$. We now argue $\text{gp}^q(x) \in \text{Convex}(P')$, which by [Lemma A.4](#) implies the claim as $\text{gp}^q(x)$ is on $\text{ray}(x)$ for $x \in \text{Convex}(P)$.

$$\text{gp}^q(x) = \frac{x}{\langle x, q \rangle} = \frac{\sum_i \alpha_i p_i}{\langle (\sum_j \alpha_j p_j), q \rangle} = \frac{\sum_i \alpha_i p_i}{\sum_j \alpha_j \langle p_j, q \rangle} = \frac{\sum_i \alpha_i \langle p_i, q \rangle p'_i}{\sum_j \alpha_j \langle p_j, q \rangle} = \sum_i \frac{\alpha_i \langle p_i, q \rangle}{\sum_j \alpha_j \langle p_j, q \rangle} p'_i.$$

Thus $\text{gp}^q(x) \in \text{Convex}(P')$ as the coefficients are non-negative and $\sum_i \frac{\alpha_i \langle p_i, q \rangle}{\sum_j \alpha_j \langle p_j, q \rangle} = 1$. ■

B Proof of [Theorem 3.4](#)

Theorem B.1 (Restatement of [Theorem 3.4](#)). Let $P \subset \mathbb{R}^m$ be a point set, let $q \in \mathbb{S}^{(m-1)}$ be such that $\phi_P(q) < \pi/2 - \gamma$ for some constant $\gamma > 0$, and let $\varepsilon > 0$ be a parameter. Then one can find, in $O(|P|m/\varepsilon^2)$ time, a point $t \in \text{Conic}(P)$ such that $\text{angle}(q, t) \leq \text{angle}(q, \text{Conic}(P)) + \varepsilon \phi_P(q)$. Moreover, t is a conic combination of $O(1/\varepsilon^2)$ points from P .

Proof: First compute the set $\text{gp}^q(P)$. (Note that since $\phi_P(q) < \pi/2$, by **Observation A.1**, every point $p \in P$ defines a unique point $\text{gp}^q(p) \in \text{hplane}(q)$.) Let $\varepsilon' = \varepsilon\phi_P(q)/(2 \tan(\phi_P(q)))$. Note that $\phi_P(q) < \pi/2 - \gamma$ implies $c \leq \phi_P(q)/(2 \tan(\phi_P(q))) \leq 1/2$, for some positive constant c depending only on the constant γ , and thus $\varepsilon' = \Theta(\varepsilon)$.

Let $\text{opt}' \in \text{Convex}(\text{gp}^q(P))$ be a point in the convex hull of $\text{gp}^q(P)$ that is closest to q . By **Theorem 3.3**, one can compute a point $t \in \text{Convex}(\text{gp}^q(P))$ such that $d(q, t) \leq d(q, \text{opt}') + \varepsilon'\Delta$, where $\Delta = \text{diameter}(\text{gp}^q(P))$, in $O(|P|m/\varepsilon'^2) = O(|P|m/\varepsilon^2)$ time, which is a convex combination of a subset Z of $O(1/\varepsilon'^2) = O(1/\varepsilon^2)$ points of $\text{gp}^q(P)$. We return $\text{pg}^q(Z)$ as our solution. (Note $\text{pg}^q(Z)$ is returned instead of Z , only because the theorem requires returning a subset of P .) Thus the running time and size of $\text{pg}^q(Z)$ in the theorem statement hold, and what remains is proving t yields the desired approximation.

By **Lemma A.2**, $\Delta \leq 2 \tan(\phi_P(q))$. By the same lemma, for any $x \in \text{Convex}(\text{gp}^q(P))$, $\text{angle}(x, q) = \tan^{-1}(d(x, q))$. Since $\tan^{-1}(\theta)$ is monotonically increasing for $0 \leq \theta \leq \pi/2$,

$$\begin{aligned} \text{angle}(q, t) &= \tan^{-1}(d(q, t)) \leq \tan^{-1}(d(q, \text{opt}') + \varepsilon'\Delta) \\ &\leq \tan^{-1}(d(q, \text{opt}') + 2\varepsilon' \cdot \tan(\phi_P(q))) = \tan^{-1}(d(q, \text{opt}') + \varepsilon\phi_P(q)). \end{aligned}$$

Furthermore, because $\tan^{-1}(\theta)$ is concave between $0 \leq \theta \leq \pi/2$ and $\frac{d}{d\theta} \tan^{-1}(\theta)|_0 = 1$,

$$\tan^{-1}(d(q, \text{opt}') + \varepsilon\phi_P(q)) \leq \tan^{-1}(d(q, \text{opt}')) + \varepsilon\phi_P(q) = \text{angle}(q, \text{opt}') + \varepsilon\phi_P(q)$$

Let $\text{opt} \in \text{Conic}(P)$ be a point in the conic hull of P that has smallest angle to q . To complete the proof, observe that **Lemma A.5** implies both that $\text{gp}^q(\text{opt}) \in \text{Convex}(\text{gp}^q(P))$ and $t \in \text{Conic}(P)$. By definition $d(\text{opt}', q) \leq d(\text{gp}^q(\text{opt}), q)$, and thus by **Corollary A.3**, $\text{angle}(\text{opt}', q) \leq \text{angle}(\text{opt}, q)$ (which actually implies $\text{angle}(\text{opt}', q) = \text{angle}(\text{opt}, q)$). Combining with the above inequality gives $\text{angle}(q, t) \leq \text{angle}(q, \text{opt}') + \varepsilon\phi_P(q) \leq \text{angle}(q, \text{opt}) + \varepsilon\phi_P(q)$ \blacksquare

C Proof of Theorem 3.6

The following helper lemma is crucial in bounding the approximation quality of our algorithm.

Lemma C.1. *Let $P \subset \mathbb{S}^{(m-1)}$ be a point set on the hypersphere such that $\phi_P \leq \frac{\pi}{2} - \gamma$ for some constant $\gamma > 0$, and let $X \subseteq P$ be a non-empty subset. For any $q \in X$ and $p \in P$, let $y = \text{pg}^q(\text{proj}(\text{gp}^q(p), \text{Convex}(\text{gp}^q(X))))$ be the projection of $\text{gp}^q(p)$ onto $\text{Convex}(\text{gp}^q(X))$ pulled back to $\mathbb{S}^{(m-1)}$. Then it holds that $\text{angle}(p, y) \leq 3\ell$, where $\ell = d(\text{gp}^q(p), \text{gp}^q(y))$.*

Proof: First observe that if $y = q$, the angle between p and y is simply $\tan^{-1}(\ell) \leq \ell < 3\ell$. Also, if $y = p$, i.e. $\text{gp}^q(p)$ is in $\text{Convex}(\text{gp}^q(X))$, then $\text{angle}(p, y) = 0$. So for the remainder of the proof we assume $y \neq q$ and $\text{gp}^q(p) \notin \text{Convex}(\text{gp}^q(X))$.

We need to show some basic relations about these points and angles that are needed in the proof. Consider the spherical triangle defined by points p, y, q . Let $\beta = \text{angle}(q, p)$, and let θ be the interior angle opposite the side py , see **Figure C.1**. First observe that $\ell \leq \tan(\beta)$ as otherwise, q would be closer to $\text{gp}^q(p)$ than $\text{gp}^q(y)$. Next, observe that $\text{angle}(\text{gp}^q(y), q) \leq \text{angle}(\text{gp}^q(p), q)$. To see this, assume the contrary and note that we could then draw a perpendicular line from $\text{gp}^q(p)$ to the line $q \text{gp}^q(y)$. The point of intersection would then be closer to $\text{gp}^q(p)$ than $\text{gp}^q(y)$ which is a contradiction.

The last relation we need involves a new point, z , which we now define. Consider two circles on the sphere. The first is the great circle through q and y . The second is the latitude circle at p with respect to q (i.e. q acts as the north pole), which is the set of all points on the sphere with fixed angle $\beta = \text{angle}(q, p)$ from q . Let z be the intersection point of these two circles (technically there are two intersections, so take the one closer to p). See **Figure C.1**. Let $\theta_1 = \angle \text{gp}^q(y) \text{gp}^q(p) \text{gp}^q(z)$ and $\theta_2 = \angle \text{gp}^q(z) \text{gp}^q(p)$ (see **Figure C.2**) and note that $\theta_1 \leq \theta_2$. Hence by the law of sines, we obtain $d(\text{gp}^q(y), \text{gp}^q(z)) \leq \ell$.

To prove the lemma, we invoke the spherical triangle inequality, that is $\text{angle}(p, y) \leq \text{angle}(p, z) + \text{angle}(z, y)$. We now upper bound $\text{angle}(p, z)$ and $\text{angle}(z, y)$ separately.

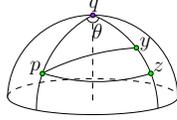


Figure C.1

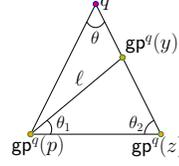


Figure C.2

To bound $\text{angle}(y, z)$, first observe that q, y , and z are all on the same great circle, and so $\text{angle}(q, z) = \text{angle}(q, y) + \text{angle}(y, z)$. By [Lemma A.2](#),

$$\begin{aligned} \text{angle}(y, z) &= \text{angle}(q, z) - \text{angle}(q, y) = \tan^{-1}(d(\text{gp}^q(z), q)) - \tan^{-1}(d(\text{gp}^q(y), q)) \\ &= \tan^{-1}(d(\text{gp}^q(y), q) + d(\text{gp}^q(y), \text{gp}^q(z))) - \tan^{-1}(d(\text{gp}^q(y), q)) \\ &\leq \tan^{-1}(d(\text{gp}^q(y), q)) + d(\text{gp}^q(y), \text{gp}^q(z)) - \tan^{-1}(d(\text{gp}^q(y), q)) = d(\text{gp}^q(y), \text{gp}^q(z)) \leq \ell, \end{aligned}$$

where the first inequality follows as $\tan^{-1}(\zeta)$ is concave for $0 \leq \zeta \leq \pi/2$ and $\frac{d}{d\zeta} \tan^{-1}(\zeta)|_0 = 1$ (i.e. the same argument from [Theorem B.1](#)).

To bound $\text{angle}(p, z)$, first observe that $\beta = \text{angle}(q, p) = \text{angle}(q, z)$, that is we have an isosceles spherical triangle. Consider the corresponding isosceles triangle in $\text{hplane}(q)$ defined by the points $q, \text{gp}^q(p)$, and $\text{gp}^q(z)$, see [Figure C.2](#). Note that the interior spherical angle θ defined above is also the angle at q in this planar triangle, that is $\theta = \angle \text{gp}^q(p) q \text{gp}^q(z)$. Also, $\theta < \pi/2$ as otherwise $d(q, \text{gp}^q(p)) \leq d(\text{gp}^q(p), \text{gp}^q(y))$ contradicting the assumption $\text{gp}^q(y) = \text{proj}(p)$. By [Lemma A.2](#), $d(\text{gp}^q(p), q) = d(\text{gp}^q(z), q) = \tan(\beta)$, so by the law cosines,

$$\cos(\theta) = \frac{\tan^2(\beta) + \tan^2(\beta) - d(\text{gp}^q(p), \text{gp}^q(z))^2}{2 \cdot \tan(\beta) \cdot \tan(\beta)} = 1 - \frac{d(\text{gp}^q(p), \text{gp}^q(z))^2}{2 \cdot \tan^2(\beta)}.$$

Going back to the sphere, by the spherical cosine law,

$$\begin{aligned} \cos(\text{angle}(p, z)) &= \cos(\beta) \cos(\beta) + \sin(\beta) \sin(\beta) \cos(\theta) = \\ &= \cos^2(\beta) + \sin^2(\beta) \cdot \left(1 - \frac{d(\text{gp}^q(p), \text{gp}^q(z))^2}{2 \cdot \tan^2(\beta)}\right) = 1 - \frac{\sin^2(\beta) \cdot d(\text{gp}^q(p), \text{gp}^q(z))^2}{2 \cdot \sin^2(\beta) / \cos^2 \beta} \\ &= 1 - \cos^2(\beta) \cdot \frac{d(\text{gp}^q(p), \text{gp}^q(z))^2}{2} \geq 1 - 2\ell^2 \cdot \cos^2(\beta), \end{aligned}$$

where the last inequality follows from $d(\text{gp}^q(p), \text{gp}^q(z)) \leq d(\text{gp}^q(p), \text{gp}^q(y)) + d(\text{gp}^q(y), \text{gp}^q(z)) \leq 2\ell$, by the above relations. Remember that $\ell \leq \tan(\beta)$ and $\tan^{-1}(\cdot)$ is monotonically increasing. Therefore, as $\cos^{-1}(\cdot)$ is a monotonically decreasing function,

$$\text{angle}(p, z) \leq \cos^{-1}(1 - 2\ell^2 \cdot \cos^2(\beta)) \leq \cos^{-1}(1 - 2\ell^2 \cdot \cos^2(\tan^{-1}(\ell))) = \cos^{-1}\left(1 - \frac{2\ell^2}{1 + \ell^2}\right).$$

We now show $\cos^{-1}(1 - 2\ell^2/(1 + \ell^2)) \leq 2\ell$. Consider the ratio $\frac{\cos^{-1}(1 - 2\ell^2/(1 + \ell^2))}{\ell}$ and observe that the derivative is negative for $\ell > 0$ (ℓ is always positive). Thus, $\frac{\cos^{-1}(1 - 2\ell^2/(1 + \ell^2))}{\ell}$ is increasing as ℓ approaches 0^+ . By L'Hopital's rule,

$$\lim_{\ell \rightarrow 0^+} \frac{\cos^{-1}(1 - 2\ell^2/(1 + \ell^2))}{\ell} = \lim_{\ell \rightarrow 0^+} \frac{2}{\ell + 1} = 2.$$

Therefore, the ratio is upper bounded by 2, and $\text{angle}(p, z) \leq \cos^{-1}(1 - 2\ell^2/(1 + \ell^2)) \leq 2\ell$ holds. Note that because $2\ell^2/(1 + \ell^2) < 2$, the upper bound we obtained for $\text{angle}(p, z)$, i.e. $\cos^{-1}(1 - 2\ell^2/(1 + \ell^2))$, is well defined for any value $\ell > 0$.

Finally, the spherical triangle inequality tell us our bounds $\text{angle}(y, z) \leq \ell$ and $\text{angle}(p, z) \leq 2\ell$, imply $\text{angle}(p, y) \leq 3\ell$. \blacksquare

We wish to prove the conic analog of the following theorem from [\[Blum et al., 2016\]](#), which is a more general form of [Theorem 3.5](#) from [Section 3.2](#).

Theorem C.2 (Blum et al., 2016). Given a set P of n points in \mathbb{R}^m , and a value $\varepsilon > 0$, in polynomial one can compute: (i) an $(\varepsilon\Delta, O(m \log k_{opt}))$ -approximation, (ii) a $((1+\delta)\varepsilon\Delta, O(\log(n)/(\varepsilon\delta)))$ -approximation, and (iii) an $((8\varepsilon^{1/3} + \varepsilon)\Delta, O(1/\varepsilon^{2/3}))$ -approximation to **Problem 3.2**. Moreover, for (iii), the run time is $O(nc(m + c/\varepsilon^2 + c^2))$, where $c = O(k_{opt}/\varepsilon^{2/3})$.

We are now ready to prove our main theorem, which is a more general form of **Theorem 3.6**. The proof will make use of the following simple observation about $d_{convex}(X, P)$ and $d_{conic}(X, P)$, which we formally state as it is also used in our d-SUM-harness proofs.

Observation C.3. If $\max_{p \in P} d(p, \text{Convex}(X)) \leq \varepsilon\Delta$, then $d_{convex}(X, P) \leq \varepsilon\Delta$: For point sets A and $B = \{b_1, \dots, b_m\}$, if we fix $a \in \text{Convex}(A)$, then for any $b \in \text{Convex}(B)$ we have $\|a - b\| = \|a - \sum_i \alpha_i b_i\| = \|\sum_i \alpha_i (a - b_i)\| \leq \sum_i \alpha_i \|a - b_i\| \leq \max_i \|a - b_i\|$.

Similarly, if $\max_{p \in P} d(p, \text{Conic}(X)) \leq \varepsilon\phi_P$, then $d_{conic}(X, P) \leq \varepsilon\phi_P$: Again for sets A and B with $\phi_{A \cup B} < \pi/2$, fix $a \in \text{Conic}(A)$ and $b \in \text{Conic}(B)$. **Lemma A.5** implies that $\text{gp}^a(b)$ is a convex combination of $\text{gp}^a(b_i)$, and thus the argument for the convex case implies $\|a - \text{gp}^a(b)\| \leq \max_i \|a - \text{gp}^a(b_i)\|$. **Corollary A.3** then implies $\text{angle}(a, b) \leq \max_i \text{angle}(a, b_i)$.

Theorem C.4. Given a set P of n points in \mathbb{R}^m such that $\phi_P \leq \frac{\pi}{2} - \gamma$ for a constant $\gamma > 0$, and a value $\varepsilon > 0$, in polynomial time one can compute: (i) an $(\varepsilon\phi_P, O(m \log k_{opt}))$ -approximation, (ii) a $((1 + \delta)\varepsilon\phi_P, O(\log(n)/(\varepsilon\delta)))$ -approximation, and (iii) an $((8\varepsilon^{1/3} + \varepsilon)\phi_P, O(1/\varepsilon^{2/3}))$ -approximation to **Problem 3.1**. Moreover, for case (iii), the run time is $O(nc(m + c/\varepsilon^2 + c^2))$, where $c = O(k_{opt}/\varepsilon^{2/3})$.

Proof: First note that as we are only concerned with angles, without loss of generality we assume $P \subset \mathbb{S}^{(m-1)}$. Set $\varepsilon' = c \cdot \varepsilon$ for a constant $c > 0$ to be determined shortly, and let $q = \text{unit}(p)$ for an arbitrary point p in P . Compute the set $P' = \text{gp}^q(P)$, and consider the instance (P', ε') of **Problem 3.2**. By **Theorem C.2**, in polynomial time we can compute a subset $X' \subseteq P'$ that is an (α, β) -approximation to **Problem 3.2** where (α, β) is either (i) $(\varepsilon'\Delta, O(m \log k_{opt}))$, (ii) $((1 + \delta)\varepsilon'\Delta, O(\log(n)/(\varepsilon\delta)))$, or (iii) $((8\varepsilon')^{1/3} + \varepsilon')\Delta, O(1/\varepsilon'^{2/3})$, where we used the fact that $\varepsilon' = \Theta(\varepsilon)$.

Note that $X = \text{unit}(X')$ is a subset of P . To prove the theorem, we apply **Lemma C.1** to the sets X and P , which we now argue implies X is a $(3\alpha, \beta)$ -approximation to the instance (P, ε) of **Problem 3.1**. First observe that in order to apply the lemma, q must be in X , so if it is not already, simply add q to X as this does not asymptotically change β . By **Definition 2.1**, if X' is an (α, β) -approximation to an instance (P', ε') of **Problem 3.2**, then for any point $\text{gp}^q(p) \in P'$ the distance to its projection onto $\text{Convex}(X')$, $y' = \text{proj}(\text{gp}^q(p), \text{Convex}(X'))$, is at most α . Thus setting $\ell = d(\text{gp}^q(p), y')$, **Lemma C.1** implies $\text{angle}(p, \text{pg}^q(y')) \leq 3\ell \leq 3\alpha$. Note by **Lemma A.5**, $\text{pg}^q(y') \in \text{Conic}(X)$, and so every point $p \in P$ has a point in $\text{Conic}(X)$ at angle $\leq 3\alpha$, which by **Observation C.3** implies X is a $(3\alpha, \beta)$ -approximation to **Problem 3.1**.

To complete the proof, we provide constants c such that either (i) $3(c\varepsilon\Delta) \leq \varepsilon\phi_P$, or (ii) $3(1 + \delta)c\varepsilon\Delta \leq (1 + \delta)\varepsilon\phi_P$, or (iii) $3(8(c\varepsilon)^{1/3} + c\varepsilon)\Delta \leq (8\varepsilon^{1/3} + \varepsilon)\phi_P$. For case (i) and (ii), set $c = \frac{\phi_P}{6 \tan(\phi_P)}$. Note that c can indeed be treated as a positive constant because $\frac{\phi_P}{6 \tan(\phi_P)}$ is always less than 1, and is bounded away from 0 by our assumption that $\phi_P \leq \pi/2 - \gamma$ for some constant γ (and thus we still have $\varepsilon' = \Theta(\varepsilon)$). From **Lemma A.2**, $\Delta \leq 2 \tan(\phi_P)$, and so $3(c\varepsilon\Delta) = \frac{3\varepsilon\Delta\phi_P}{6 \tan(\phi_P)} \leq \frac{6\varepsilon\phi_P \tan(\phi_P)}{6 \tan(\phi_P)} = \varepsilon\phi_P$, thus satisfying case (i). Similarly, $3(1 + \delta)c\varepsilon\Delta \leq (1 + \delta)\varepsilon\phi_P$, satisfying case (ii).

For case (iii) set $c = \left(\frac{\phi_P}{6 \tan(\phi_P)}\right)^3$, which is still a constant. Since $\frac{\phi_P}{6 \tan(\phi_P)} \leq 1$ we have $3(8(c\varepsilon)^{1/3} + c\varepsilon)\Delta \leq 3\left(8\left(\frac{\phi_P}{6 \tan(\phi_P)}\right)\varepsilon^{1/3} + \left(\frac{\phi_P}{6 \tan(\phi_P)}\right)^3 \varepsilon\right)2 \tan(\phi_P) \leq (8\varepsilon^{1/3} + \left(\frac{\phi_P}{6 \tan(\phi_P)}\right)^2 \varepsilon)\phi_P \leq (8\varepsilon^{1/3} + \varepsilon)\phi_P$. ■

D Additional Hardness Results

D.1 Problem 3.1 is d-SUM-hard

This section proves the following decision version of **Problem 3.1** is d-SUM-hard. The approach is the same as in the convex reduction where we construct a point set P with d clusters each having N

points. However, we now require that $\phi_P \leq \pi/2 - \gamma$ to match the problem statement in [Problem 3.1](#). Without this requirement, the reduction would be virtually the same as before, except with angles not distances. As a consequence of the bounded angle assumption, some of the calculations become very technically challenging.

Problem D.1. Given a set P of n points in \mathbb{R}^d such that $\phi_P \leq \pi/2 - \gamma$, a value $\varepsilon > 0$, and an integer k , is there a subset $X \subseteq P$ of k points such that $d_{\text{conic}}(X, P) \leq \varepsilon \phi_P$, where ϕ_P is the angular diameter of P .

Given an instance of **d-SUM** with N values $S = \{s_1, s_2, \dots, s_N\}$ we construct an instance of [Problem D.1](#) where $P \subset \mathbb{R}^{d+2}$, $k = d$, and $\varepsilon = 1/7$. Let $u = (1, 1, \dots, 1)$ be the all ones vector in \mathbb{R}^d , e_i be the i th standard basis vector in \mathbb{R}^d , and $w^* = \tan(\varepsilon \cos^{-1}(\frac{d+2+1/5}{d+3}))$ be a constant. Note that setting $\varepsilon = 1/7$ still implies the general ε case is hard, and moreover the problem remains hard even if we restrict to cases with $\varepsilon \geq 1/7$. In this section we assume $d \geq 4$, which is not necessary but it simplifies some of our later calculations.

Definition D.2. The set of points $P \subset \mathbb{R}^{d+2}$ are the following

- p_j^i points: For each $i \in [d]$, $j \in [N]$, set $(a_1, \dots, a_d) = u + e_i$, $w = s_j \cdot w^*$ and $v = 0$
Let $R = \cup_{i,j} p_j^i$ be the set of all p_j^i points
- q point: For each $i \in [d]$, $a_i = \frac{1+d}{d}$, $w = w^*/2$, $v = 0$
- q' point: For each $i \in [d]$, $a_i = \frac{1+d}{d}$, $w = w^*/2$, $v = \|q\| \cdot \tan(\varepsilon \phi_R)$

Again, we let $P^i = \cup_j p_j^i$ denote the i th cluster. Observe that w^* is the largest w value any point can have because $s_k \in [0, 1]$. Since w^* is largest when $d = 4$ and since $\varepsilon = 1/7$ we have the upper bound $w^* < 1/\sqrt{5}$, which will be used in the following calculations.

Lemma D.3. $\phi_R = \cos^{-1}\left(\frac{d+2+s_{\min}s_{\max}(w^*)^2}{\sqrt{d+3+(s_{\min}w^*)^2}\sqrt{d+3+(s_{\max}w^*)^2}}\right)$ where s_{\min} and s_{\max} are, respectively, the minimum and maximum values in S .

Proof: We consider the angle between two points in different clusters and two points in the same cluster. Let p_j^i, p_k^i, p_l^m be three arbitrary points in R where $j \neq k$ and $m \neq i$.

The inter-cluster angle is $\text{angle}(p_j^i, p_l^m) = \cos^{-1}\left(\frac{d+2+w(p_j^i)w(p_l^m)(w^*)^2}{\sqrt{d+3+(w(p_j^i)w^*)^2}\sqrt{d+3+(w(p_l^m)w^*)^2}}\right)$.

The intra-cluster angle is $\text{angle}(p_j^i, p_k^i) = \cos^{-1}\left(\frac{d+3+w(p_j^i)w(p_k^i)(w^*)^2}{\sqrt{d+3+(w(p_j^i)w^*)^2}\sqrt{d+3+(w(p_k^i)w^*)^2}}\right)$.

Note that the inter-cluster angle is maximized when $w(p_j^i) = s_{\max}$ and $w(p_l^m) = s_{\min}$. This angle is exactly what we claim ϕ_R to be. So it is left to show that this maximized inter-cluster angle is always larger than any intra-cluster angle. Note that any inter-cluster angle is lower bounded by $\text{angle}(p_j^i, p_l^m) \geq \cos^{-1}\left(\frac{d+2+1/5}{d+3}\right)$ and thus we have

$$\text{angle}(p_j^i, p_k^i) \leq \cos^{-1}\left(\frac{d+3}{d+3+\frac{1}{5}}\right) < \cos^{-1}\left(\frac{d+2+\frac{1}{5}}{d+3}\right) \leq \text{angle}(p_j^i, p_l^m). \quad \blacksquare$$

Lemma D.4. The angular diameter of P , ϕ_P , is equal to ϕ_R .

Proof: To prove the lemma we show the lower bound for ϕ_R , seen in [Lemma D.3](#), is larger than $\text{angle}(p_j^i, q)$, $\text{angle}(p_j^i, q')$, and $\text{angle}(q, q')$ for a point $p_j^i \in R$. First note that by [Definition D.2](#), $\text{angle}(q, q') = \varepsilon \phi_R < \phi_P$, since ϕ_P is at least ϕ_R .

It is clear that $\text{angle}(p_j^i, q) \leq \text{angle}(p_j^i, q')$, so we only need to consider the latter. We first find an upper bound on $\text{angle}(p_j^i, q')$. Because $\|q\|^2 \leq (d+2+1/d+1/20)$, we obtain

$$\begin{aligned} \cos(\text{angle}(p_j^i, q')) &\geq \frac{(d+1)^2}{d\sqrt{d+3+\frac{1}{5}}\sqrt{d+2+\frac{1}{d}+\frac{(w^*)^2}{4}}+\|q\|^2 \tan^2(\varepsilon\phi_R)} \\ &\geq \frac{(d+1)^2}{d\sqrt{d+\frac{16}{5}}\sqrt{d+\frac{1}{d}+\frac{23}{10}}}. \end{aligned}$$

The last inequality holds because $\phi_R \leq \cos^{-1}\left(\frac{d+2}{d+4}\right)$ and $\|q\|^2 \tan^2(\varepsilon \cos^{-1}\left(\frac{d+2}{d+4}\right)) < 1/4$ for all $d \geq 4$ and $\varepsilon = 1/7$. To verify $\text{angle}(p_j^i, q') \leq \phi_R$, we use the this upper bound on $\text{angle}(p_j^i, q')$ and the same lower bound on ϕ_R used in the proof of [Lemma D.3](#). The limit of the ratio of these two bounds is

$$\lim_{d \rightarrow \infty} \frac{\cos^{-1}\left(\frac{d^2+2d+1}{d\sqrt{d+3+\frac{1}{5}}\sqrt{d+\frac{1}{d}+\frac{23}{10}}}\right)}{\cos^{-1}\left(\frac{d+2+\frac{1}{5}}{d+3}\right)} = \sqrt{\frac{15}{16}} < 1.$$

Since the derivative of this ratio is positive for all $d \geq 4$, this limit is an upper bound. Thus, $\text{angle}(p_j^i, q') \leq \phi_R$ and hence $\phi_P = \phi_R$. \blacksquare

Corollary D.5. *The diameter of P is less than $\pi/4$.*

Proof: From the above proofs we have the bound $\phi_P = \phi_R \leq \cos^{-1}\left(\frac{d+2}{d+3+\frac{1}{5}}\right)$. Thus for any $d \geq 4$, $\phi_P < \pi/4$, and moreover this upper bound is a decreasing function whose limit is 0 as d goes to infinity. \blacksquare

We now prove the completeness and soundness of the above reduction using the fact that $\phi_P = \phi_R$ from [Lemma D.4](#).

Lemma D.6 (Completeness). *If there is a subset $\{s_{k_1}, s_{k_2}, \dots, s_{k_d}\}$ of d values (not necessarily distinct) such that $\sum_{i \in [d]} s_{k_i} = d/2$, then the above described instance of [Problem D.1](#) is a true instance, i.e. there is a d sized subset $X \subseteq P$ such that $d_{\text{conic}}(X, P) \leq \varepsilon\phi_P$.*

Proof: Let $x_i = (u + e_i, s_{k_i} \cdot w^*, 0)$, which by [Definition D.2](#) is a point in P and let $X = \{x_1, \dots, x_d\}$. By [Observation C.3](#), $\max_{p \in P} d(p, \text{Conic}(X)) \leq \varepsilon\phi_P$ implies $d_{\text{conic}}(X, P) \leq \varepsilon\phi_P$. We now show that $\max_{p \in P} \text{angle}(p, \text{Conic}(X)) \leq \varepsilon\phi_P$ which completes the lemma.

First observe that for any point p_j^i in P , x_i will have the same values as p_j^i in all the a_i coordinates. Consider the right triangle formed by the origin, p_j^i with $w(p_j^i)$ set to 0, and p_j^i with $w(p_j^i)$ set to w^* . Both p_j^i and x_i lie somewhere on the short leg since their w values are at most $w^* = \tan(\varepsilon \cos^{-1}\left(\frac{d+2+1/5}{d+3}\right)) \leq \tan(\varepsilon\phi_P)$. Thus their angle is less than the angle opposite the short leg. Since $\tan^{-1}(\cdot)$ is monotonically increasing and the length of the longer leg is at least \sqrt{d} , we get

$$\text{angle}(p_j^i, x_i) \leq \tan^{-1}\left(\frac{w^*}{\sqrt{d}}\right) \leq \tan^{-1}(\tan(\varepsilon\phi_P)) \leq \varepsilon\phi_P.$$

The only other points in P are q and q' and by definition $\text{angle}(q, q') = \varepsilon\phi_P$. Again, we claim $x = \frac{1}{d} \sum_i x_i \in \text{Conic}(X)$ is the point q which implies $\max_{p \in P} \text{angle}(p, \text{Conic}(X)) \leq \varepsilon\phi_P$. As X contains exactly one point from each set P^i , and in each such set all points have $a_i = 2$ and all other $a_j = 1$, it holds that x has $\frac{1+d}{d}$ for all the a coordinates. All points in X have $v = 0$ and so this holds for x as well. Thus we only need to verify that $w(x) = w(q) = \frac{w^*}{2}$, for which we have $w(x) = \frac{1}{d} \sum_i w(x_i) = \frac{w^*}{d} \sum_i s_{k_i} = \frac{w^*}{2}$. \blacksquare

Lemma D.7. *Let $P \subset \mathbb{R}^{d+2}$ be as defined above, and let $X \subseteq P$ be a subset of size d . Then if $d_{\text{conic}}(X, P) \leq \varepsilon\phi_P$, then for all i , X contains exactly one point from P^i .*

Proof: Assume that there is a set P^t such that $X \cap P^t = \emptyset$. Consider an arbitrary point p_j^t in P^t . We prove that $\text{angle}(x, p_j^t) > \varepsilon\phi_P$ for any point x in $\text{Conic}(X)$, which contradicts the assumption that $d_{\text{conic}}(X, P) \leq \varepsilon\phi_P$.

We first find a point x' that lower bounds $\text{angle}(x, p_j^t)$ for all x in $\text{Conic}(X)$. Since every point in $P \setminus P^t$ has at most $\frac{1+d}{d}$ in the a_t coordinate, any point x in $\text{Conic}(X)$ will have at most this value for a_t . Thus, we let x' have the same value as p_j^t in every coordinate except for a_t which is $\frac{1+d}{d}$. Note that the angle between x' and p_j^t is minimized when they have the same w values. Thus we let $w(p_j^t) = w(x') = w$ and their angle becomes

$$\text{angle}(x, p_j^t) \geq \text{angle}(x', p_j^t) = \cos^{-1} \left(\frac{2 + \frac{2}{d} + d - 1 + w^2}{\sqrt{(d+3+w^2)\left(\left(\frac{1+d}{d}\right)^2 + d - 1 + w^2\right)}} \right)$$

The derivative with respect to w of the above inner ratio is positive for any $d \geq 4$, $w \geq 0$. Therefore, as $w^* < 1/\sqrt{5}$ is an upper bound on the w value of any point, we let $w(p_j^t) = w(x') = 1/\sqrt{5}$ which maximizes the inner ratio and consequently minimizes the angle between the two points. Thus for any $x \in \text{Conic}(X)$,

$$\text{angle}(x, p_j^t) \geq \text{angle}(x', p_j^t) \geq \cos^{-1} \left(\frac{2 + \frac{2}{d} + d - 1 + \frac{1}{5}}{\sqrt{(d+3+\frac{1}{5})\left(\left(\frac{1+d}{d}\right)^2 + d - 1 + \frac{1}{5}\right)}} \right)$$

Remember from the proof of [Lemma D.4](#) that $\cos^{-1}\left(\frac{d+2}{d+4}\right)$ is an upper bound for ϕ_P . Thus to show $\text{angle}(x, p_j^t) > \varepsilon\phi_P$ it suffices to argue the ratio

$$\frac{\cos^{-1} \left(\frac{2 + \frac{2}{d} + d - 1 + \frac{1}{5}}{\sqrt{(d+3+\frac{1}{5})\left(\left(\frac{1+d}{d}\right)^2 + d - 1 + \frac{1}{5}\right)}} \right)}{\cos^{-1} \left(\frac{d+2}{d+4} \right)},$$

is $> \varepsilon = 1/7$ for any $d \geq 4$. Thus the claim follows as the derivative of this ratio is positive for all $d \geq 4$, and thus the minimum occurs at $d = 4$, which itself is greater than $1/7$. ■

Lemma D.8. *If $d_{\text{conic}}(X, P) \leq \varepsilon\phi_P$, then $q \in \text{Conic}(X)$ and moreover $q = \frac{1}{d} \sum_{x_i \in X} x_i$.*

Proof: By [Lemma D.7](#), X cannot contain the point q' , and since all points in P other than q' have v coordinate equal to 0, $\text{Conic}(X)$ is contained in the hyperplane $v = 0$. Note that the projection of q' onto this hyperplane is the point q . As the angle between q' and q is exactly $\varepsilon\phi_P$ by definition, it then must be that $q \in \text{Conic}(X)$.

For the second part of the lemma, consider the coordinate a_i for any i . q has this coordinate set to $\frac{1+d}{d}$. Again by [Lemma D.7](#), X contains exactly one point with this coordinate set to 2 and all other points in X have this coordinate set to 1. Thus the convex combination for q must take at least a $1/d$ fraction of this one point in order to achieve a total value of $\frac{1+d}{d}$ in coordinate a_i . On the other hand, the convex combination cannot take more than a $1/d$ fraction of this point, as by symmetry it would imply for some other coordinate a_j , the total value would be strictly less than $\frac{1+d}{d}$. Thus for all i , the convex combination for q must take exactly a $1/d$ fraction of x_i . ■

Lemma D.9 (Soundness). *Let P be an instance of [Problem 3.1](#) generated from a d -SUM instance S , as described in [Definition D.2](#). If there is a subset $X \subseteq P$ of size d such that $d_{\text{conic}}(X, P) \leq \varepsilon\phi_P$, then there is a choice of d values from S that sum to exactly $d/2$.*

Proof: From [Lemma D.7](#) we know that X consists of exactly one point from each cluster P^i . Thus for each $x_i \in X$, $w(x_i) = s_{k_i} \cdot w^*$ for some $s_{k_i} \in S$. By [Lemma D.8](#), $q = \frac{1}{d} \sum_i x_i$, which implies $w(q) = \frac{1}{d} \sum_i w(x_i)$. By [Definition D.2](#) $w(q) = \frac{w^*}{2}$, which implies:

$$\frac{w^*}{2} = \frac{1}{d} \sum_i w(x_i) = \frac{w^*}{d} \sum_i s_{k_i}.$$

Thus we have a set $\{s_{k_1}, \dots, s_{k_d}\}$ of d values from S such that $\sum_i s_{k_i} = d/2$. ■

[Lemma D.6](#) and [Lemma D.9](#) immediately imply the following.

Theorem D.10. *For point sets in \mathbb{R}^{d+2} , [Problem D.1](#) is d -SUM-hard.*

E Additional Experimental Results

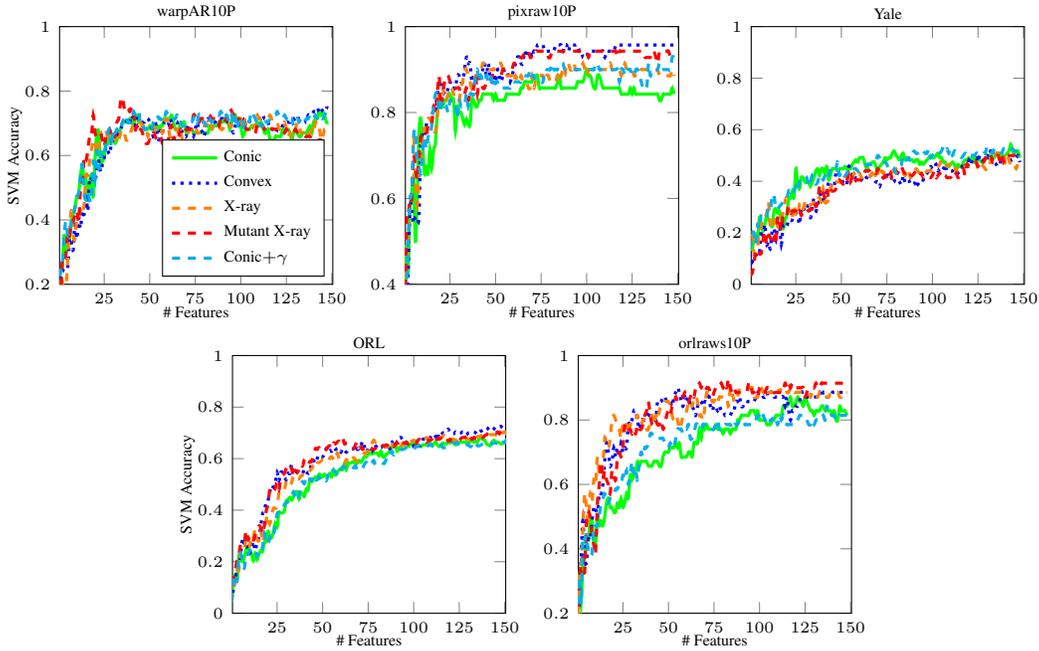


Figure E.1: Additional experimental results for feature selection.

F Proofs from [Section 4](#)

Lemma F.1 (Restatement of [Lemma 4.5](#)). *The diameter of P , Δ_P , is equal to Δ^* .*

Proof: Observe that there are only five different types of pairs of points in P that can determine Δ_P . Consider three points p_j^i, p_k^i, p_l^m in P where $j \neq k$ and $m \neq i$, i.e. p_j^i and p_k^i are in the same cluster and p_l^m is in a different cluster. The five distances are: $d(p_j^i, p_k^i)$, $d(p_j^i, p_l^m)$, $d(p_j^i, q')$, $d(p_k^i, q')$, $d(q, q')$. Consider $d(p_j^i, p_l^m)$. The distance between these two points is $\sqrt{2 + (\varepsilon w(p_j^i) - \varepsilon w(p_l^m))^2}$. This distance is maximized when $w(p_j^i)$ and $w(p_l^m)$ are the maximum and minimum values in S , which is exactly Δ^* . Thus we need to show every other distance is less than Δ^* .

First we have $d(p_j^i, p_k^i) \leq \sqrt{(\varepsilon w(p_j^i) - \varepsilon w(p_k^i))^2} \leq \varepsilon$ and $d(q, q') = \varepsilon \Delta^* < \Delta^*$. Next, note that $(\varepsilon s_j - \varepsilon/2)^2 \leq 1/36$ and $(\varepsilon \Delta^*)^2 \leq 7/27$ since $\varepsilon \leq 1/3$. Thus,

$$d(p_j^i, q) \leq d(p_j^i, q') = \sqrt{(1 - 1/d)^2 + (d - 1)/d^2 + (\varepsilon s_j - \varepsilon/2)^2 + (\varepsilon \Delta^*)^2} < \sqrt{2} \leq \Delta^*. \quad \blacksquare$$

Lemma F.2 (Restatement of Lemma 4.7). *Let $P \subset \mathbb{R}^{d+2}$ be as defined above, and let $X \subseteq P$ be a subset of size d . If $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$, then for all i , X contains exactly one point from P^i .*

Proof: Assume that there is a set P^t such that $X \cap P^t = \emptyset$. Consider an arbitrary point p_j^t in P^t . We prove that $d(x, p_j^t) > \varepsilon\Delta$ for any point x in $\text{Convex}(X)$, which is a contradiction as it implies $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$. This will imply each set P^i has non-empty intersection with X , and as the size of X is the same as the number of P^i sets, this in turn implies X contains exactly one point from each P^i set.

First observe the only points in $P \setminus P^t$ with non-zero a_t coordinate are q and q' , each of which has $a_t = 1/d$. As x is a convex combination of points from $P \setminus P^t$ this implies that its a_t coordinate is at most $1/d$. On the other hand, p_j^t has coordinate a_t set to 1. This implies that $d(x, p_j^t) \geq \sqrt{(1 - 1/d)^2} = 1 - 1/d \geq 1/2 > \varepsilon\Delta^* = \varepsilon\Delta$. (Other coordinates can be ignored as differing values in other coordinates can only increase $d(x, p_j^t)$.) ■

Lemma F.3 (Restatement of Lemma 4.8). *If $d_{\text{convex}}(X, P) \leq \varepsilon\Delta$, then $q \in \text{Convex}(X)$ and moreover $q = \frac{1}{d} \sum_{x_i \in X} x_i$.*

Proof: By Lemma 4.7, observe that X cannot contain the point q' . On the other hand, all points in P other than q' have v coordinate equal to 0, and thus $\text{Convex}(X)$ is contained in the hyperplane $v = 0$. At the same time, q' has v coordinate equal to $\varepsilon\Delta^*$, and thus has distance $\varepsilon\Delta$ to the hyperplane $v = 0$. Thus the projection of q' onto the $v = 0$ hyperplane must be contained in $\text{Convex}(X)$. However, the projection of q' onto this hyperplane is the point q , and thus $q \in \text{Convex}(X)$.

For second part of the lemma, consider the coordinate a_i for any i . q has this coordinate set to $1/d$. Again by Lemma 4.7, X contains exactly one point with this coordinate set to 1 and all other points in X have this coordinate set to zero. Thus any convex combination of X realizing q must take a $1/d$ fraction of this point. As this is true for all i , the claim follows. ■