Advanced Algorithms (6.854) Final Project - Chasing Convex Bodies

Angelos Assos, Binwei Yan, Andi Liu
assos@mit.edu bineva@mit.edu andiliu@mit.edu

June 21, 2023

---

1 Introduction

1.1 Introduction to the problem

The *convex body chasing* problem, first introduced by Friedman and Linial [FL93], is a rather important competitive analysis problem with a long line of work. The problem itself is easy to state and understand: We control a point in R^d initially at $x_0 = 0$, and for each time step $t \in N$, the adversary gives a convex body $K_t \subseteq R^d$ as a request, and the player picks a point inside the convex body, i.e. pick an x_t such that $x_t \in K_t$. (In essence we are moving our point to lie within K_t.) The goal is to minimize the total distance moved over all time steps T, which is:

$$\sum_{t=1}^{T} ||x_t - x_{t-1}||$$

In this paper, we work in the framework of competitive analysis. The Friedman-Linial Conjecture states it is possible to achieve constant competitiveness in any d dimensional space R^d. In fact, this constant is at least $\sqrt{d}$, as shown later in section 5.1. This conjecture has remained open for over two decades, but it recently gained a lot of attention because of its connections to machine learning.

1.2 Past Results and Roadmap for this paper

As we mentioned before, the idea of the Convex Bodies Chasing problem was initiated by the questions asked by Linial and Friedman [FL93] back in the 90's, when they were trying to understand the interplay between geometry and competitive ratios in online problems on metric spaces. They provided us with a lower bound of $\sqrt{d}$ for the competitive ratio where d is the number of dimensions. After many years of little to no substantial progress (especially in the cases where $d \geqslant 3$), people start to shift their attention to the restricted variant of the problem: Nested Convex Bodies Chasing, in which the sequence of bodies (K_t) must be nested (i.e. inside the last request):

$$K_1 \supseteq K_2 \supseteq \cdots \supseteq K_T$$

Bansal et al. [BBE+18] gave an exponential competitive algorithm for the nested case of the problem. Shortly after, Argue et al. [ABC+19] improved the algorithm by moving to the centroid of each requested convex body and managed to get $O(d \log d)$ competitiveness on the nested case (covered in section 5.3). The trick is that every time we do that, the volume of the body shrinks by at least a constant factor, and so after a small amount of moves, we are going to 'eliminate' a dimension (imagine if we had a 3d body, after a few moves it will look like a pancake).

A year later Bubeck et al. [BLLS18] managed to chase an even better point than the centroid, namely the *Steiner Point* (defined in Section 2). They manage to improve the bound to $O(\min(d, \sqrt{d \log T}))$ where T is the number of bodies to be followed, which we would examine in section 6.

Notice that all the aforementioned bounds are for the nested case. Argue et al. [AGTG21] managed to generalize the $O(\min(d, \sqrt{d \log T}))$ competitive algorithm of the nested case to the general case. We will give a brief overview of the algorithm in Section 7. At the end of this paper we give current results for all versions of the problems and all norms.

2 Preliminaries

Before we dive in the specifics of convex body chasing, this section is introducing some of the basic concepts of computational geometry and some objects that we are going to examine through this paper, including l_p norm, centroid, and most importantly, the Steiner Point.

Definition 2.1. For a real number $l \geqslant 1$, the l^p norm of a vector x is defined by:

$$||x||_p = (|x_1|^p + |x_2|^p + \cdots + |x_n|^p)^{1/p}$$

In the special case of $p = 2$, we have the Euclidean norm:

$$||x|| = ||x||_2 = \sqrt{(x_1^2 + x_2^2 + \cdots + x_n^2)}$$

In the special case of $p = \infty$, we have the maximum norm:

$$||x|| = ||x||_\infty = \max\{|x_1|, |x_2|, \cdots, |x_n|\}$$

For most of the paper, we will be talking about Euclidean norm. However, there are results for the maximum norm at the end of the paper that will be briefly talked about.

Definition 2.2. First, a ball in dimension d with radius r is denoted by B_r, i.e.

$$B_r = \{x \in \mathbb{R}^d : \|x\| \leqslant r\}$$

and we use B to represent the unit circle, i.e. $B = B_1$.

Definition 2.3. We denote the *centroid* or center of mass of a convex body K as $cg(K)$. It is the average of every points, formally defined as:

$$cg(K) = \frac{\int_{x \in K} x \ \ dx}{\int_{x \in K} dx}$$

Now, we have the following 3 definitions for the Steiner Point of a convex body K, denoted by $st(K)$:

Definition 2.4. For any direction $\theta \in \mathbb{S}^{d-1}$ (where $\mathbb{S}^{d-1}$ denotes the boundary of a hyper-sphere in d-dimension), let $f_K(\theta) = \arg\max_{x \in K} (\theta \cdot x)$ which is the extreme point in K at direction θ. Then we have:

$$st(K) = \int_{\theta \in \mathbb{S}^{d-1}} f_K(\theta) d\theta$$

Note that f_K is the *gradient* of the body K. Also note that here we integrate over the uniform probability measure on $\mathbb{S}^{d-1}$

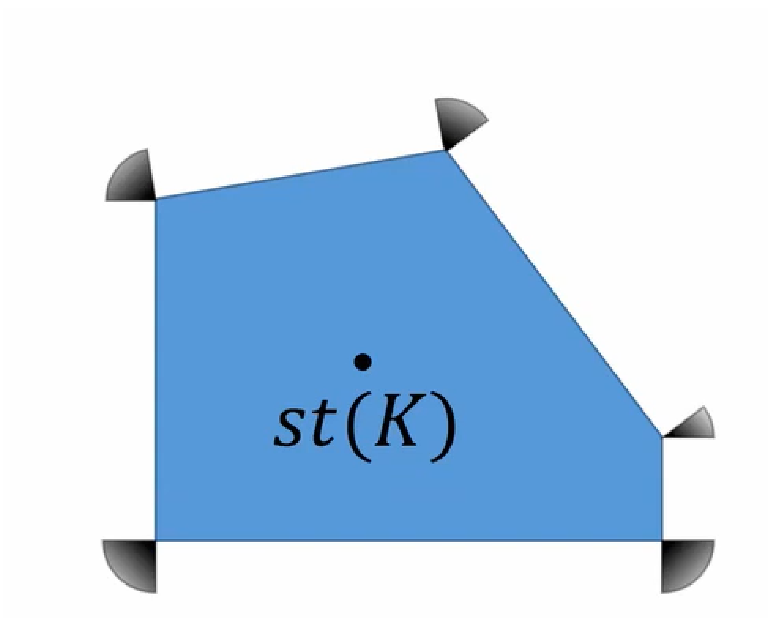

Figure 1: the Steiner Point is the weighted average of extreme points, where the weight is represented by the area of these circular sectors. Note they are perpendicular to the edges.

Definition 2.5. For any direction $\theta \in \mathbb{S}^{d-1}$, let $h_K(\theta) = \max_{x \in K} (\theta \cdot x)$, so that h_K is the *support function* of K. Then we have:

$$st(K) = d \int_{\theta \in \mathbb{S}^{d-1}} \theta \cdot h_K(\theta) d\theta$$

Definition 2.6. [Prz96]: the Steiner Point of a convex body is defined as:

$$st(K) = \lim_{s \to \infty} cg(K + sB)$$

The above three definitions are all useful in their own way; the first definition tells us that the Steiner Point will always be inside the body K. The second definition is used to bound the movement. The third definition will also be helpful in the general case. Actually, the integrals of the first two definitions of the Steiner Point can be proved to be equivalent due to the divergence theorem, once we noteme that $f_K(\theta) = \nabla h_K(\theta)$.

Let us finally define the work function which we will use in section 7. The work function up to query t of point x, is the optimal trajectory that satisfies the first t queries and ends up at point x. Therefore:

$$w_t(x) = \min_{x_1, x_2, \ldots, x_t} \sum_{i=1..t} \|x_i - x_{i-1}\| + \|x - x_t\|$$
$$\text{s.t. } x_i \in K_i$$

Definition 2.7. A 1-dimensional projection on the k'th dimension is the set

$$S_k = \{x_k : (x_1, x_1, ..., x_k, x_{k+1}, ..., x_d) \in K\|$$

In other words, it is all the points x_k such that there is a point in K that contains x_k in their k'th coordinate. The length of the projection is defined as $l_k = \max_{x \in S_k} x - \min_{x \in S_k} x$.

Definition 2.8. The *width* of a convex body K $w(K)$ is the average length random 1-dimensional projection on the body K.

The last two definitions will only be used in the later sections of the paper. Now, let us look at some special cases of the algorithm to provide some intuition.

3 Function chasing

Before we dive into the specifics of the algorithms, we would first like to present the general problem of *chasing convex functions* and how it can be reduced to the easier subproblem of chasing convex bodies using the technique from [BLLS19]. When chasing convex functions, each request, instead of corresponding to a convex set as in convex bodies chasing, corresponds to a convex function $f_t : R^d \to R_+ \cup \{+\infty\}$ (that is, each point in the space has an extra cost called service cost, which is a positive number that can be infinity). In other words, in addition to the movement cost between two points x_{t-1} and x_t, there is now also a service cost of $f_t(x_t)$ that varies on the position the point ends up at each step and varies between request, and we are minimizing

$$\sum_{t=1}^{T}(||x_t - x_{t-1}|| + f_t(x_t))$$

It is easy to see that convex body chasing is a special case of convex function chasing by essentially making the service cost infinity anywhere outside of the convex set and zero inside the convex set. We can even furthermore argue that these problems are roughly equivalent in

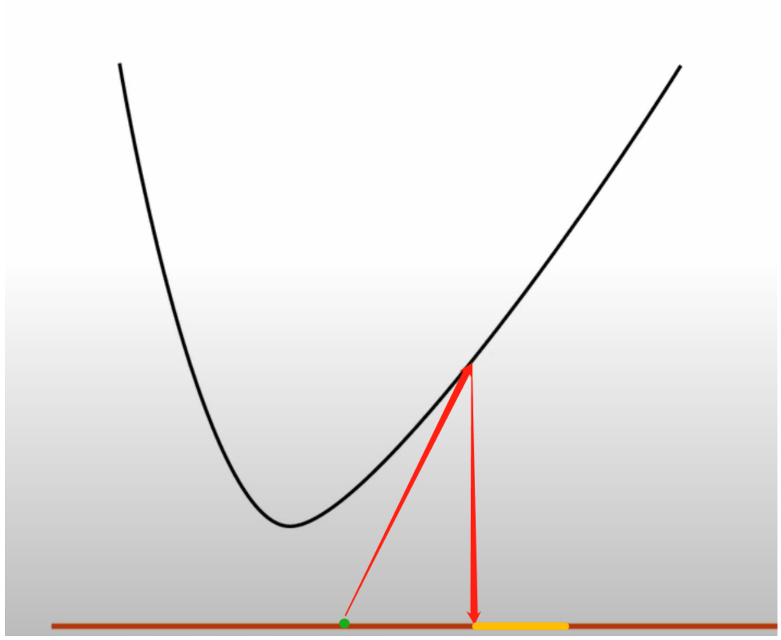

Figure 2: The reduction from 1d Convex Function Chasing to 2d Convex Body Chasing

competitive analysis, as argued by Bubeck et al, in that a convex *function* chasing problem in dimension d can be reduced to a convex *body* chasing problem in dimension $d+1$ (with only a loss of constant factor). This is achieved via replacing the function request with two requests: one requesting the point to be on the area above the graph of the function, (which is feasible since the function is convex itself), another requesting the point to be back on the axis (formally $R^d \times \{0\}$.), and thus the moving cost is covered by the horizontal distance in first part and the service cost covered by the vertical distance in the second part. (see figure 2) Because they are roughly equivalent, we would be dealing with the Convex *Body* Chasing case with no service cost inside the convex set.

4 Line Chasing

Line Chasing is a special case of convex set chasing, where all the convex set are lines. In its introduction by Friedman and Linial [FL93], a 28.53 competitive ratio in 2d plane is achieved by doing greedy algorithm on the lines after some transformations (this ratio increases as the dimension gets higher). After decades, Bienkowski et al [BBC$^+$18]. proposed a extremely elegant algorithm named DRIFT that achieves a competitive ratio of 3 for all dimensions.

The DRIFT algorithm (see figure 3) has the following intuition: suppose the last requested line is L and we need to move from $P \in L$ to L', the greedy algorithm (making a perpendicular line to move the shortest distance) will move to $\overline{P} \in L'$, which could potentially be bad if we simply alternate between L and L', where DRIFT drifts a distance $x = \frac{1}{\sqrt{2}}(h+s-r)$ (as in figure 3) toward the intersection. (If they do not intersect, we do the same as greedy.) Its 3-competitiveness can be proved by analyzing the potential function $\phi(P, A)$ depending on the

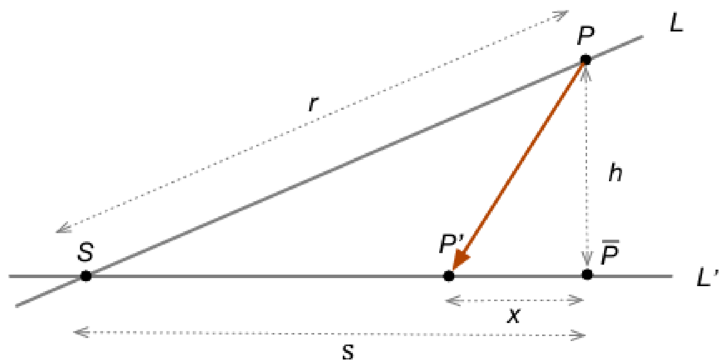

Figure 3: Instead of greedily going to $\overline{P}$, DRIFT drift toward the intersection s

locations $P, A \in L$ of the algorithm's and the adversary's point. The higher-dimension cases can solved by projecting the lines to a carefully chosen plane.

This idea of drifting toward the center will be used later.

5 Lower bound and Baseline Ideas

5.1 Lower bound

An easily illustrated lower bound for the competitive ratio is $\sqrt{d}$. In a d-dimensional cube, we can cut the dimensions one by one by requesting a $(d-1)$-dimensional cube, and each time we are forced to use the center of the $(d-1)$-dimensional cube because otherwise our adversary could take advantage of that by requesting the opposite of wherever we drifted.

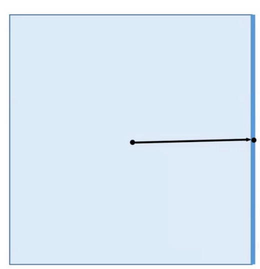

Figure 4: in a square (2d), we move to the center of requested line (1d)

Eventually our request will be just a point, the end of the path. While our algorithm moved the Manhattan distance to get to the point, the clairvoyant algorithm knows that it could just move to the final point, costing the Euclidean distance.

According to Pythagoras's theorem for Euclidean space (see figure 5), we know that the ratio between the two would be $\sqrt{d}$.

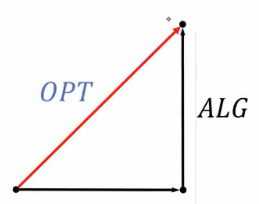

Figure 5: Manhattan distance v.s. Euclidean Distance

5.2 Greedy Algorithm

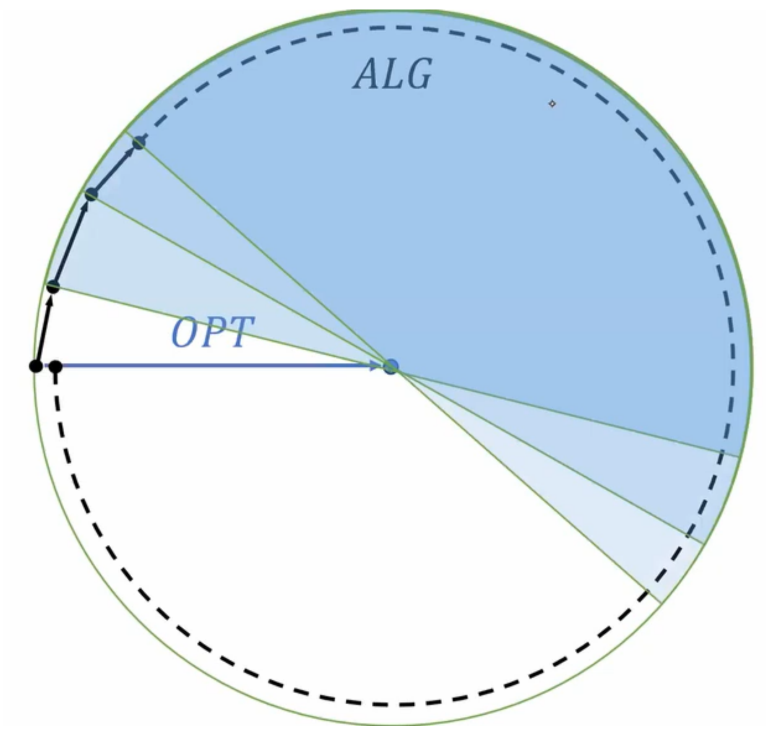

Figure 6: being greedy can make us non-competitive with OPT

Being greedy and choosing the closest point in the convex body from where we currently are does not work. For a counterexample, see figure 6. If we start near the perimeter of a circle and each request is a slightly more tilted semicircle, we are just greedily making a perpendicular line and going in spirals, where the adversary could have just gone to the center and solved the requests once and for all. This results in an arbitrarily large competitive ratio, depending on how much we tilt the semicircle each time (after all, a tangent line of a circle is always perpendicular to the diameter).

5.3 Centroid Algorithm

Notice that in greedy algorithms we just go to the edge of the convex shape. The opposite approach would be always moving toward the center of the convex shape to save time for the next request. The high level idea behind the centroid algorithm by Argue et al. [ABC+19] is that the point always moves to the centroid of the shapes (when they are not already in the request, of course), where centroid is defined by the mean of all points. The competitive ratio for this algorithm is $O(d \log d)$ (in the nested case). Here is a rough proof. The worst case for the centroid algorithm, in a 2d example, is when we start from a rectangle, then the adversary alternately requests slanted rectangles and normal rectangles, resulting in a zig-zag pattern. (see figure 7)

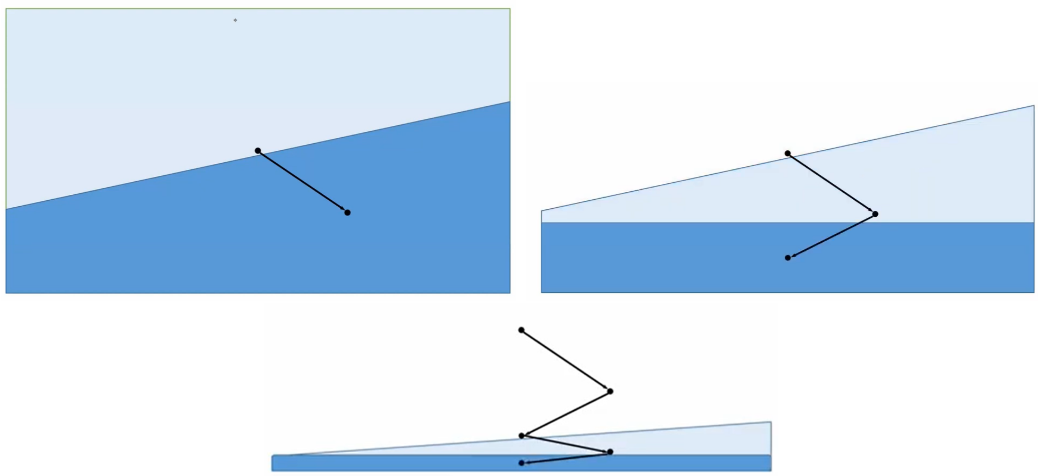

Figure 7: Worst case: we follow the centroid through the slanted shapes

Suppose the rectangle has a side length of r, while the optimum can just spend $O(r)$ to move to the side, our algorithm needs to basically "thin out" a dimension (known as the "pancake" argument) after logarithmic number of request (According to Grunbaum's Theorem, any half-plane that cuts through the centroid reduces the volume by $1/e$.). Each zig-zag costs $O(r)$, resulting in a total cost of $O(r \log d)$ for one dimension. There are a total of d dimensions, each costing $O(r \log d)$, and in the ratio, r cancel each other out, which is where the $O(d \log d)$ comes from.

This is one step closer to our main topic, the Steiner Point algorithm, where the center is not the mean of all points, but the weighted average of the extreme points.

6 Follow the Steiner Point (Nested Case)

In this section we are going to analyze the competitiveness of the Follow the Steiner Point algorithm, for the nested case. We will first prove an $O(d)$ ratio that can be achieved by using the second definition of the Steiner Point as seen in Section 2, and then we will get a bound that depends on the number of queries.

Before we jump to out first proof, let us prove some useful lemmas:

Lemma 6.1. *For $\theta \in \mathbb{S}^{d-1}$, and a vector $v \in \mathbb{R}^d$, we have:*

$$\theta \cdot v \leqslant v$$

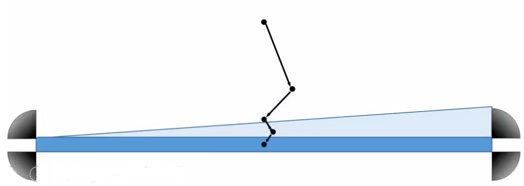

Figure 8: Following the Steiner Point is much smoother

Proof. θ has magnitude of 1. $\square$

Lemma 6.2. *For convex bodies K and K' where $K \supseteq K'$ we have that:*

$$h_K(\theta) \geqslant h_{K'}(\theta)$$

Proof. Recall that $h_K(\theta) = \max_{x \in K} (\theta \cdot x)$. Then we have that if x is the vector that maximizes $h_{K'}(\theta)$ for K', then since $K \supseteq K'$, $x \in K$ as well, Therefore

$$\max_{x \in K} (\theta \cdot x) \geqslant \max_{x \in K'} (\theta \cdot x)$$

$\square$

Lemma 6.3. *For a convex body K we have that:*

$$h_K(\theta) + h_K(-\theta) \geqslant 0$$

Proof. Recall that $h_K(\theta) = \max_{x \in K} (\theta \cdot x)$. Assume that $h_K(\theta) \geqslant 0$. Then $h_K(-\theta) = \max_{x \in K} (-\theta \cdot x) \geqslant (-\theta) \cdot x$ and thus:

$$h_K(\theta) + h_K(-\theta) \geqslant \theta \cdot x + (-\theta) \cdot x = 0$$

$\square$

We are now ready to get our $O(d)$ bound for the nested case. Note that in the following theorem we assume that we start from the unit ball for simplicity:

Theorem 6.4. *Let $B = K_1 \supseteq K_2 \supseteq ... \supseteq K_T$ be a sequence of nested convex bodies. Then:*

$$\sum_{i=1}^{T-1} \left\| st(K_{i+1}) - st(K_i) \right\| \leqslant d$$

Proof. The idea is that simply instead of summing the distances, we can sum over all θ since for each fixed θ, the intergrand decreases by a total of at most 2 (since we are in the 1 unit ball). So the total budget for movement is 2d. To save the factor 2 we combine $\pm\theta$, noting that they can change by at most 2 in total. We have:

$$\sum_{i=1}^{T-1} \left\| st(K_{i+1}) - st(K_i) \right\| = \sum_{i=1}^{T-1} \left\| d \int_{\theta \in \mathbb{S}^{d-1}} \theta \cdot h_{K_i}(\theta) d\theta - d \int_{\theta \in \mathbb{S}^{d-1}} \theta \cdot h_{K_{i+1}}(\theta) d\theta \right\|$$

from the second definition of the Steiner Point as defined in section 2. Now we use Lemma 6.1 to simplify the above expression and we also rearrange the sum with the integral:

$$d\sum_{i=1}^{T-1}\left\|\int_{\theta\in\mathbb{S}^{d-1}}\theta\cdot h_{K_i}(\theta)d\theta-\int_{\theta\in\mathbb{S}^{d-1}}\theta\cdot h_{K_{i+1}}(\theta)d\theta\right\|\leqslant d\sum_{i=1}^{T-1}\left\|\int_{\theta\in\mathbb{S}^{d-1}}h_{K_i}(\theta)d\theta-\int_{\theta\in\mathbb{S}^{d-1}}h_{K_{i+1}}(\theta)d\theta\right\|$$

$$\leqslant d\int_{\theta\in\mathbb{S}^{d-1}}\sum_{i=1}^{T-1}\left|h_{K_i}(\theta)-h_{K_{i+1}}(\theta)\right|$$

Now from lemma 6.2 we get that $h_{K_i}(\theta)\geqslant h_{K_{i+1}}(\theta)$. Then we can simplify the above telescoping sums to:

$$d\int_{\theta\in\mathbb{S}^{d-1}}\sum_{i=1}^{T-1}\left|h_{K_i}(\theta)-h_{K_{i+1}}(\theta)\right|=d\int_{\theta\in\mathbb{S}^{d-1}}\sum_{i=1}^{T-1}h_{K_i}(\theta)-h_{K_{i+1}}(\theta)=\int_{\theta\in\mathbb{S}^{d-1}}\sum_{i=1}^{T-1}h_{K_1}(\theta)-\int_{\theta\in\mathbb{S}^{d-1}}\sum_{i=1}^{T-1}h_{K_T}(\theta)$$

To end the proof note that the first intergrand cannot be bigger than 1 since we are in the unit ball and the second intergrand is nonnegative because $h_{K_T}(\theta)+h_{K_T}(-\theta)\geqslant 0$ (from lemma 6.3). Thus:

$$\int_{\theta\in\mathbb{S}^{d-1}}\sum_{i=1}^{T-1}h_{K_1}(\theta)-\int_{\theta\in\mathbb{S}^{d-1}}\sum_{i=1}^{T-1}h_{K_T}(\theta)\leqslant 1$$

and the problem is proved. □

The above algorithm can be generalized for bodies outside of the unit ball using the same algorithm. Now that we found our $O(d)$ competitive algorithm, we can move on to calculate the other bound we want.

Notice that the idea in the $O(d)$ algorithm was to bound the total distance we could move in each direction by a constant. Since we could bound that, we can get an $O(d)$ distance at most. Now, the idea is to bound the distance we can move in each step. Lemma 6.5 below shows that we can bound the distance we travel in each step by $O(\sqrt{d}\cdot f(\lambda))$. The distance depends on λ which is defined as:

$$\lambda=\frac{w(K_i)-w(K_{i+1})}{2}$$

The following lemma from [BLLS18] bounds the distance travelled in one move and will help us achieve the bound we need for every time step:

Lemma 6.5. *[BLLS18]: For any convex bodies K' and K such that $K'+B_1\supseteq K\supseteq K'$ and $\lambda=\frac{w(K)-w(K')}{2}$ we have:*

$$\left\|s(K)-s(K')\right\|\leqslant\lambda\sqrt{d\log\left(\lambda^{-1}\right)}$$

Recall that $w(K)$ is the width of the convex body, i.e. the maximum distance we can travel in one dimension on that body. We are not going to give a proof of Lemma 6.2 as it is too technical, but we will try to give the intuition as to why it is true.
The idea is that initially we have a 'budget' we want to spend, which is

$$d\int_{\theta\in\mathbb{S}^{d-1}} h_{K_1}(\theta) - d\int_{\theta\in\mathbb{S}^{d-1}} h_{K_T}(\theta)$$

Note that every time we move to the Steiner Point we spend

$$d\int_{\theta\in\mathbb{S}^{d-1}} h_{K_i}(\theta) - d\int_{\theta\in\mathbb{S}^{d-1}} h_{K_{i+1}}(\theta)$$

from our budget, but we actually go a distance

$$d\int_{\theta\in\mathbb{S}^{d-1}} h_{K_i}(\theta)\theta - d\int_{\theta\in\mathbb{S}^{d-1}} h_{K_{i+1}}(\theta)$$

The intuition is if we make a move that costs a lot of budget, we most likely made a move that does not correspond to a big distance. This is just because a lot of the θ's end up cancelling up in the calculation and therefore not adding up in the right way. In other words, the moves that take a lot of our budget, are not moves that hurt our distance sum, and once we lose budget we can't claim it back since we are in the nested case. Therefore the best thing for an adversary is to make us use our budget we have slowly.

Now we have everything ready to proceed to the proof of the main theorem.

Theorem 6.6. *Following the Steiner Point in nested case, starting from B_1, and following requests $B_1 = K_1 \supseteq K_2 \supseteq \cdots \supseteq K_T$ gives total movement*

$$O(\min(d, \sqrt{dlogT}))$$

after T requests

Proof. Set $\lambda_i = \frac{w(K_i)-w(K_{i+1})}{2}$ for $i = 0, ..., T-1$. Then for each step we have from Lemma 3.2:

$$\sum_{i=1}^{T-1} \left\| st(K_i) - st(K_{i+1}) \right\| \leqslant d^{1/2} \cdot \sum_{i=1}^{T-1} \lambda_i \sqrt{\log \lambda_i^{-1}}$$

Now note that we have that $\sum_i \lambda_i \leqslant 1$ since it is a telescoping sum and note that $\sqrt{\log x}$ is a concave function. From Jensen inequality we have that for concave functions f such that $a_1 + a_2 + ... + a_n \leqslant 1$:

$$f(a_1 x_1 + ... + a_n x_n) \geqslant a_1 f(x_1) + ... + a_n f(x_n)$$

Now using $f(x) = \sqrt{\log x}$ we have that:

$$\sum_{i=1}^{T-1} \lambda_i \sqrt{\log \lambda_i^{-1}} = \sqrt{\log \sum_{i=1}^{T-1} \lambda_i \lambda_i^{-1}} = \sqrt{\log T}$$

In the end we have after combining everything:

$$\sum_{i=1}^{T-1} \left\| st(K_i) - st(K_{i+1}) \right\| \leqslant d^{1/2} \sum_{i=1}^{T-1} \lambda_i \sqrt{\log \lambda_i^{-1}} \leqslant d^{1/2} \sqrt{\log T}$$

as desired. Adding the bound from theorem 6.1, we get the final bound.

□

7 Follow the Steiner Point (General Case)

In this section, we will try to generalize the ideas shown in the nested case, and try to prove some bounds for the general problem of chasing convex bodies, without the restriction of them being nested.

We take the algorithm that follows the Steiner Point and use it in the general case as well. The high level idea is that instead of just blindly moving to the Steiner Point of the newly asked shape, we move to the Steiner Point of the body we are asked that is extended to a certain amount.

7.1 The Algorithm

The algorithm is very similar to the algorithm presented for the nested case. Instead of following the Steiner Point of the next body K_t, we follow the Steiner Point of a somehow extended K_t body, extended by a distance $2r$, where r is an estimate of the distance covered so far by the optimal route, $r \in [OPT/2, OPT]$, and r is calculated with the help of the work function, as defined in section 2.

Algorithm 1 ExtendedSteinerPointChasing

$x_0 \leftarrow 0$
$r \leftarrow dist(x_0, K_1)$ ▷ Initializing our variables
for $t = 1$ to T **do** ▷ For every query
    $\Omega_t \leftarrow \{x | w_t(x) \leqslant 2r\}$ ▷ Get the 2r hitting set
    while $\Omega_t = \varnothing$ **do** ▷ Go on until you find a point that is almost optimal
        $r \leftarrow 2r$
        $\Omega_t \leftarrow \{x | w_t(x) \leqslant 2r\}$
    end while
    $x_t \leftarrow st(\Omega_t)$ ▷ Chase the Steiner Point
end for

The natural question that comes after looking at the above algorithm is why is the Steiner Point of Ω_t falling inside the body K_t. This is the main Lemma we need to prove in order to guarantee correctness for our algorithm. Once we prove that, we can use Theorem 6.3 to get the bound we want, and the proof will be over. In order to prove that Ω_t lies within the body K_t we will take a useful lemma:

Lemma 7.1. *The work function is a convex function and the Ω_t's for $t = 1, 2, ..., T$ are bounded convex sets.*

Now we are now ready to prove our important Lemma.

Lemma 7.2. *For x_i as chosen in the algorithm above, we get that x_i lies within the body K_t.*

Proof. By the third definition of the Steiner Point, we only want to prove that:

$$cg(\Omega_t + sB) \in K_t$$

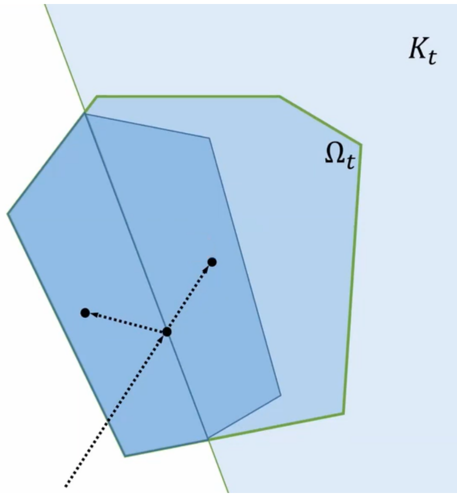

Figure 9: Since there are more inside, the Steiner Point is going to be inside

To prove that, we prove that $cg(\Omega_t + sB)$ is in every halfspace that contains K_t.

Let $H := \{x \in \mathbb{R}^d : a \cdot x \geqslant b\}$ be a halfspace containing K_t and let $H^=$ be the plane cutting the border of this halfspace. Define $\rho(x)$ to be the reflection of x across $H^=$:

$$\rho(x) = x + 2(b - (a \cdot x))a$$

Recall that the work function is:

$$w_i(x) = \min_{y \in K_t} \|x - y\| + w_{i-1}(y)$$

and let y be the argmin of the expression on the right. Since $y \in K_t \subseteq H$, if we choose a point x that is in Ω_t but not in H, we have:

$$w_t(x) = \|x - y\| + w_{t-1}(y) \geqslant \big\|\rho(x) - y\big\| + w_{t-1}(y) \geqslant w_t(\rho(x))$$

just because $\rho(x) \in H$ but $x \notin H$. Therefore we have that if a point $x \in \Omega_t/H$, then its reflection is going to be in Ω_t as well. Similarly we can prove this for $\Omega_t + sB$ as well.

To finish the proof, split $\Omega_t + sB$ into three parts; $\Omega_t + sB = \Omega^- \cup \Omega^+ \cup \hat{\Omega}$ where $\Omega- = (\Omega_t + sB)/H$, $\Omega^+ = \rho(\Omega^-)$ and $\hat{\Omega} = ((\Omega_t + sB) \cap H)/\Omega^+$. The idea is that the combined centroid of Ω^- and Ω^+ is going to be on the border of the halfspace and $\hat{\Omega}$ will have its centroid strictly inside the the halfspace, therefore overall the centroid of $\Omega + sB$ is going to be inside the halfsapce and therefore we are done. (see 9 for a better illustration)

□

Now we are ready to prove our important theorem:

Theorem 7.3. *Following the Steiner Point in the general case we get a*

$$O(\min(d, \sqrt{d \log T}))$$

competitive algorithm after T requests

Proof. We consider the progression of the algorithm in phases; each new phase begins when r changes. Suppose the phase corresponding to some value of r consists of times $\{t_1, t_1 + 1, ..., t_2\}$. From the fact that the work function is non-decreasing over time, i.e., $w_t \leqslant w_{t+1}$, as we move more and more, and also the fact that $\Omega_t = \{x | w_t(x) \leqslant 2r\}$ it follows that the bigger our work function is, the smallest the Ω_t set is. Therefore:

$$\Omega_{t_1} \supseteq \Omega_{t_1+1} \supseteq ..\Omega_{t_2}$$

Bounding Ω_1 in a ball of radius $2r$ allows us to use Theorem 6.3 for the above Ω's. Following the Steiner Point of the Ω's leads to a

$$(2r) \cdot O(\min(d, \sqrt{d \log T}))$$

for each phase, where the optimal algorithm spends $O(r)$ distance. In total we spend $r_{final} \cdot O(\min(d, \sqrt{d \log T}))$, where r_{final} is the sum of all the r for all the phases. OPT uses only $O(r_{final})$ therefore the competitive ratio desired is achieved. □

8 Conclusion

8.1 Summary

In our paper we tried to give a short overview of the problem of Convex Body Chasing. We gave problems that are similar and belong in the same family of problems, examined a very interesting and well studied case of the problem, the nested case, and gave some of the current best known algorithms, both for the general but also in the nested case. Below we summarize the current best known competitive algorithms for the nested and general case for all the different l^p norms, where $p \geqslant 2$ and the lower bounds for each of them:

Norm	Problem Case	Lower Bound	Current best	By who
l^2	Nested	$O(\sqrt{d})$	$O(\sqrt{d \log d})$	[BLLS18]
l^2	General	$O(\sqrt{d})$	$O(\min(\sqrt{d \log T}, d))$	[AGTG21]
$l^p, p > 2$	Nested	$O(d^{1-\frac{1}{p}})$	$O(d^{1-\frac{1}{p}} \sqrt{\log d})$	[BLLS18]
$l^p, p > 2$	General	$O(d^{1-\frac{1}{p}})$	Not yet found	-
l^∞	Nested	$O(d)$	$O(d)$	[BLLS18]
l^∞	General	$O(d)$	Not yet found	-

As we can see from the table, a lot of ground has been covered in the last few years. For the nested case, we have almost found the tightest algorithms that are only a factor of $\sqrt{\log d}$ apart from the optimal ratio. [BLLS18] proved that the Steiner Point is the optimal point of a convex body to chase. Lastly, recently [Sel20] proved that following the Steiner Point in the function setting, as described in section 3, achieves competitive ratio $O(d \log T)$

In the general case, we still need progress. the Steiner Point helped us achieve a good algorithm which is at worse $O(\sqrt{d})$ away from the optimal ratio. In the cases where the norm is greater than 2 we could not find any papers that achieved any good competitive ratio, so it would be no surprise if in the near future we see some new developments in that area.s

References

[ABC+19] CJ Argue, Sébastien Bubeck, Michael B Cohen, Anupam Gupta, and Yin Tat Lee. A nearly-linear bound for chasing nested convex bodies. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 117–122. SIAM, 2019.

[AGTG21] CJ Argue, Anupam Gupta, Ziye Tang, and Guru Guruganesh. Chasing convex bodies with linear competitive ratio. *Journal of the ACM (JACM)*, 68(5):1–10, 2021.

[BBC+18] Marcin Bienkowski, Jaroslaw Byrka, Marek Chrobak, Christian Coester, Lukasz Jez, and Elias Koutsoupias. Better bounds for online line chasing. *CoRR*, abs/1811.09233, 2018.

[BBE+18] Nikhil Bansal, Martin Böhm, Marek Eliáš, Grigorios Koumoutsos, and Seeun William Umboh. Nested convex bodies are chaseable. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1253–1260. SIAM, 2018.

[BLLS18] Sébastien Bubeck, Yin Tat Lee, Yuanzhi Li, and Mark Sellke. Chasing nested convex bodies nearly optimally. *CoRR*, abs/1811.00999, 2018.

[BLLS19] Sébastien Bubeck, Yin Tat Lee, Yuanzhi Li, and Mark Sellke. Competitively chasing convex bodies. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 861–868, New York, NY, USA, 2019. Association for Computing Machinery.

[FL93] Joel Friedman and Nathan Linial. On convex body chasing. *Discrete Comput. Geom.*, 9(3):293–321, dec 1993.

[Prz96] Krzysztof Przesławski. Centres of convex sets inlpmetrics. *journal of approximation theory*, 85(3):288–296, 1996.

[Sel20] Mark Sellke. Chasing convex bodies optimally. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1509–1518. SIAM, 2020.