

Contextual Search via Intrinsic Volumes

Renato Paes Leme
 Google Research
 New York, USA
 Email: renatoppl@google.com

Jon Schneider
 Department of Computer Science
 Princeton University
 Princeton, USA
 Email: js44@cs.princeton.com

Abstract—We study the problem of contextual search, a multidimensional generalization of binary search that captures many problems in contextual decision-making. In contextual search, a learner is trying to learn the value of a hidden vector. Every round the learner is provided an adversarially-chosen context vector, submits a guess for the inner product of the context vector with the hidden vector, learns whether their guess was too high or too low, and incurs some loss based on the quality of their guess. The learner’s goal is to minimize their total loss over the course of some number of rounds.

We present an algorithm for the contextual search problem for the symmetric loss function that achieves constant total loss. We present a new algorithm for the dynamic pricing problem (which can be realized as a special case of the contextual search problem) that achieves doubly-logarithmic total loss, improving exponentially on the previous best known upper bounds and matching the known lower bounds (up to a polynomial dependence on dimension). Both algorithms make significant use of ideas from the field of integral geometry, most notably the notion of intrinsic volumes of a convex set. To the best of our knowledge this is the first application of intrinsic volumes to algorithm design.

Keywords—online learning; integral geometry; pricing

I. INTRODUCTION

Consider the classical problem of binary search, where the goal is to find a hidden real number $x \in [0, 1]$, and where feedback is limited to guessing a number p and learning whether $p \leq x$ or whether $p > x$. One can view this as an online learning problem, where every round t a learner guesses a value $p_t \in [0, 1]$, learns whether or not $p_t < x$, and incurs some loss $\ell(x, p_t)$ (for some loss function $\ell(\cdot, \cdot)$). The goal of the learner is to minimize the total loss $\sum_{t=1}^T \ell(x, p_t)$ which can alternatively be thought of as the learner’s regret. For example, for the loss function $\ell(x, p_t) = |x - p_t|$, the learner can achieve total regret bounded by a constant via the standard binary search algorithm.

In this paper, we consider a contextual, multi-dimensional generalization of this problem which we call the *contextual search problem*. Now, the learner’s goal is to learn the value of a hidden vector $v \in [0, 1]^d$. Every round, an adversary provides a context u_t , a unit vector in \mathbb{R}^d , to the learner. The learner must now guess a value p_t , upon which they incur loss $\ell(\langle u_t, v \rangle, p_t)$ and learn whether or not $p_t \leq \langle u_t, v \rangle$. Geometrically, this corresponds to the adversary providing the learner with a hyperplane; the learner may then translate

the hyperplane however they wish, and then learn which side of the hyperplane v lies on. Again, the goal of the learner is to minimize their total loss $\sum_{t=1}^T \ell(\langle u_t, v \rangle, p_t)$.

This framework captures a variety of problems in contextual decision-making. Most notably, it captures the well-studied problem of *contextual dynamic pricing* [3], [6], [10]. In this problem, the learner takes on the role of a seller of a large number of differentiated products. Every round t the seller must sell a new product with features summarized by some vector $u_t \in [0, 1]^d$. They are selling this item to a buyer with fixed values $v \in [0, 1]^d$ for the d features (that is, this buyer is willing to pay up to $\langle u, v \rangle$ for an item with feature vector u). The seller can set a price p_t for this item, and observes whether or not the buyer buys the item at this price. If a sale is made, the seller receives revenue p_t ; otherwise the seller receives no revenue. The goal of the seller is to maximize their revenue over a time horizon of T rounds.

The dynamic pricing problem is equivalent to the contextual search problem with loss function ℓ satisfying $\ell(\theta, p) = \theta - p$ if $\theta \geq p$ and $\ell(\theta, p) = \theta$ otherwise. The one-dimensional variant of this problem was studied by Kleinberg and Leighton [13], who presented an $O(\log \log T)$ regret algorithm for this problem and showed that this was tight. Amin, Rostamizadeh and Syed [3] introduce the problem in its contextual, multi-dimensional form, but assume iid contexts. Cohen, Lobel, and Paes Leme [6] study the problem with adversarial contexts and improve the $\tilde{O}(\sqrt{T})$ -regret obtainable from general purpose contextual bandit algorithms [1] to $O(d^2 \log T)$ -regret, based on approximating the current knowledge set (possible values for v) with ellipsoids. This was later improved to $O(d \log T)$ in [14].

In this paper we present algorithms for the contextual search problem with improved regret bounds (in terms of their dependence on T). More specifically:

- 1) For the symmetric loss function $\ell(\theta, p) = |\theta - p|$, we provide an algorithm that achieves regret $O(\text{poly}(d))$. In contrast, the previous best-known algorithms for this problem (from the dynamic pricing literature) incur regret $O(\text{poly}(d) \log T)$.
- 2) For the dynamic pricing problem, we provide an algorithm that achieves regret $O(\text{poly}(d) \log \log T)$. This is tight up to a polynomial factor in d , and improves exponentially on the previous best known

bounds of $O(\text{poly}(d)\log T)$.

Both algorithms can be implemented efficiently in randomized polynomial time (and achieve the above regret bounds with high probability).

Techniques from Integral Geometry: Classical binary search involves keeping an interval of possible values (the “knowledge set”) and repeatedly bisecting it to decrease its length. In the one-dimensional case length can both be used as a potential function to measure the progress of the algorithm and as a bound for the loss. When generalizing to higher dimensions, the knowledge set becomes a higher dimensional convex set and the natural measure of progress (the volume) no longer directly bounds the loss in each step.

To address this issue we use concepts from the field of *integral geometry*, most notably the notion of *intrinsic volumes*. The field of integral geometry (also known as geometric probability) studies measures on convex subsets of Euclidean space which remain invariant under rotations/translations of the space. One of the fundamental results in integral geometry is that in d dimensions there are $d+1$ essentially distinct different measures, of which surface area and volume are two. These $d+1$ different measures are known as *intrinsic volumes*, and each corresponds to a dimension between 0 and d (for example, surface area and volume are the $(d-1)$ -dimensional and d -dimensional intrinsic volumes respectively).

A central idea in our algorithm for the symmetric loss function is to choose our guess p_t so as to divide one of the d different intrinsic volumes in half. The choice of which intrinsic volume to divide in half depends crucially on the geometry of the current knowledge set. When the knowledge set is well-rounded and ball-like, we can get away with simply dividing the knowledge set in half by volume. As the knowledge set becomes thinner and more pointy, we must use lower and lower dimensional intrinsic volumes, until finally we must divide the one-dimensional intrinsic volume in half. By performing this division carefully, we can ensure that the total sum of all the intrinsic volumes of our knowledge set (appropriately normalized) decreases by at least the loss we incur each round.

Our algorithm for the dynamic pricing problem builds on top of the ideas developed for the symmetric loss together with a new technique for charging progress based on an isoperimetric inequality for intrinsic volumes that can be obtained from the Alexandrov-Fenchel inequality. This new technique allows us to combine the doubly-exponential buckets technique of Kleinberg and Leighton with our geometric approach to the symmetric loss and obtain an $O_d(\log \log T)$ regret algorithm for the pricing loss.

One can ask whether simpler algorithms can be obtained for this setting using only the standard notions of volume and width. We analyze simpler halving algorithms and show that while they obtain $O_d(1)$ regret for the symmetric loss,

the dependency on the dimension d is exponentially worse. While the simple halving algorithms are defined purely in terms of standard geometric notions, our analysis of them still requires tools from intrinsic geometry. For the pricing loss case, we are not aware of any simpler technique just based on standard geometric notions that can achieve $O_d(\log \log T)$ regret.

Finally, we would like to mention that to the best of our knowledge this is the first application of intrinsic volumes to theoretical algorithm design.

Applications and Other Related Work: The main application of our result is to the problem of contextual dynamic pricing. The dynamic pricing problem has been extensively studied with different assumptions on contexts and valuation. Our model is the same as the one in Amin et al [3], Cohen et al [6] and Lobel et al [14] who provide regret guarantees of $O(\sqrt{T})$, $O(d^2 \log T)$ and $O(d \log T)$ respectively. The one-dimensional (non-contextual) dynamic pricing problem first appeared in [11], and was solved in [13]. The problem was also studied with stochastic valuation and additional structural assumptions on contexts in Javanmard and Nazerzadeh [10], Javanmard [9] and Qiang and Bayati [16]. This line of work relies on techniques from statistical learning, such as greedy least squares, LASSO and regularized maximum likelihood estimators. The guarantees obtained there also have $\log T$ dependency on the time horizon.

The contextual search problem was also considered with the loss function $\ell(\theta, p) = \mathbf{1}\{|\theta - p| > \epsilon\}$. For this loss function, Lobel et al [14] provide the optimal regret guarantee of $O(d \log(1/\epsilon))$. The geometric techniques developed in this line of work were later applied by Gillen et al [8] in the design of online algorithms with an unknown fairness objective. Another important application of contextual search is the problem of personalized medicine studied by Bastani and Bayati [5] in which the learner is presented with patients who are described in terms of feature vectors and needs to decide on the dosage of a certain medication. The right dosage for each patient might depend on age, gender, medical history along with various other features. After prescribing a certain dosage, the algorithm only observes if the patient was underdosed or overdosed.

Intrinsic volumes are well studied in mathematics (see [12]). Amulunxen et al apply intrinsic volumes of convex cones to study phase transitions in random instances of some problems in machine learning, including sparse recovery and dictionary learning [2].

Paper organization: The remainder of the paper is organized as follows. In Section II we define the contextual search problem and related notions. In Section III we review what is known about this problem in one dimension (where contexts are meaningless), specifically the $O(\log \log T)$ regret algorithm of Kleinberg and Leighton for the dynamic pricing problem and the corresponding $\Omega(\log \log T)$ lower bound. In Section IV, we present our algorithms for the spe-

cific case where $d = 2$, where the relevant intrinsic volumes are just the area and perimeter, and where the proofs of correctness require no more than elementary geometry (and the 2-dimensional isoperimetric inequality). In Section V, we define intrinsic volumes formally and introduce all relevant necessary facts. Finally, in Section VI, we present our two main algorithms in their general form, prove upper bounds on their regret, and argue that they can be implemented efficiently in randomized polynomial time.

II. PRELIMINARIES

A. Contextual Search

We define the *contextual search problem* as a game between a *learner* and an *adversary*. The adversary begins by choosing a point $v \in [0, 1]^d$. Then, every round for T rounds, the adversary chooses a context represented by a unit vector $u_t \in \mathbb{R}^d$ and gives it to the learner. The learner must then choose a value $p_t \in \mathbb{R}$, whereupon the learner accumulates regret $\ell(\langle u_t, v \rangle, p_t)$ (for some loss function $\ell(\cdot, \cdot)$) and learns whether $p_t \leq \langle u_t, v \rangle$ or $p_t \geq \langle u_t, v \rangle$. The goal of the learner is to minimize their total regret, which is equal to the sum of their losses over all time periods: $\text{Reg} = \sum_t \ell(\langle u_t, v \rangle, p_t)$.

We primarily consider two loss functions:

Symmetric loss. The symmetric loss measures the absolute value between the guess and the actual dot product, i.e.,

$$\ell(\theta, p) = |\theta - p|.$$

Alternatively, $\ell(\theta, p)$ can be thought of as the distance between the learner's hyperplane $H_t := \{x \in \mathbb{R}^d; \langle u_t, x \rangle = p_t\}$ and the adversary's point v .

Pricing loss. The pricing loss corresponds to the revenue loss by pricing an item at p when the buyer's value is θ . If a price $p \leq \theta$ the product is sold with revenue p , so the loss with respect to the optimal revenue θ is $\theta - p$. If the price is $p > \theta$, the product is not sold and the revenue is zero, generating loss θ . In other words,

$$\ell(\theta, p) = \theta - p \mathbf{1}\{p \leq \theta\}.$$

The pricing loss function is highly asymmetric: underpricing by ϵ can only cause the revenue to decrease by ϵ while overpricing by ϵ can cause the item not to be sold generating a large loss.

B. Notation and framework

The algorithms we consider will keep track of a *knowledge set* $S_t \subseteq S_1 := [0, 1]^d$, which will be the set of vectors v consistent with all observations so far. In step t if the context is u_t and the guess is p_t , the algorithm will update S_{t+1} to $S_t^+(p_t; u_t)$ or $S_t^-(p_t; u_t)$ depending on the feedback obtained, where:

$$S_t^+(p_t; u_t) := \{x \in S_t; \langle u_t, x \rangle \geq p_t\}$$

and

$$S_t^-(p_t; u_t) := \{x \in S_t; \langle u_t, x \rangle \leq p_t\}.$$

Since S_1 is originally a convex set and since S_{t+1} is always obtained from S_t by intersecting it with a halfspace, our knowledge set S_t will remain convex for all t .

Given context u_t in round t , we let \underline{p}_t and \bar{p}_t be the minimum and maximum (respectively) of the dot product $\langle u_t, x \rangle$ that is consistent with S_t :

$$\underline{p}_t = \min_{x \in S_t} \langle u_t, x \rangle \quad \text{and} \quad \bar{p}_t = \max_{x \in S_t} \langle u_t, x \rangle$$

Finally, given a set S and a unit vector u we will define the width in the direction u as

$$\text{width}(S; u) = \max_{x \in S} \langle u, x \rangle - \min_{x \in S} \langle u, x \rangle.$$

We will consider strategies for the learner that map the current knowledge set S_t and context u_t to guesses p_t . In Figure 1 we summarize our general setup.

Figure 1: Contextual search framework

- 1: Adversary selects $v \in S_1 = [0, 1]^d$
- 2: **for** $t = 1$ to T **do**
- 3: Learner receives a unit vector $u_t \in \mathbb{R}^d$, $\|u_t\| = 1$.
- 4: Learner selects $p_t \in \mathbb{R}$ and incurs loss $\ell(\langle u_t, v \rangle, p_t)$.
- 5: Learner receives feedback and learns the sign of $\langle u_t, x \rangle - p_t$.
- 6: Learner updates S_{t+1} to $S_t^+(p_t; u_t)$ or $S_t^-(p_t; u_t)$ accordingly.
- 7: **end for**

Oftentimes, we will want to think of d as fixed, and consider only the asymptotic dependence on T of some quantity (e.g. the regret of some algorithm). We will use the notation $O_d(\cdot)$ and $\Omega_d(\cdot)$ to hide the dependency on d .

III. ONE DIMENSIONAL CASE AND LOWER BOUNDS

In the one dimensional case, contexts are meaningless and the adversary only gets to choose the unknown parameter $v \in [0, 1]$. Here algorithms which achieve optimal regret (up to constant factors) are known for both the symmetric loss and the pricing loss. We review them here both as a warmup for the multi-dimensional version and as a way to obtain lower bounds for the multi-dimensional problem.

For the symmetric loss function, binary search gives constant regret. If the learner keeps an interval S_t of all the values of v that are consistent with the feedback received and in each step guesses the midpoint, then the loss $\ell_t \leq |S_t| = 2^{-t}|S_1|$. Therefore the total regret

$$\text{Reg} = \sum_t \ell_t = O(1).$$

For the pricing loss, one reasonable algorithm is to perform $\log T$ steps of binary search, obtain an interval containing v of length $1/T$ and price at the lower end of this interval. This algorithm gives the learner regret $O(\log T)$. Kleinberg and Leighton [13] provide a surprising algorithm that exponentially improves upon this regret. Their policy biases the search towards lower prices to guarantee that if at some point the price p_t is above v , then the length of the interval S_t decreases by a large factor.

Kleinberg and Leighton’s algorithm works as follows. At all rounds, they maintain a knowledge set $S_t = [a_t, a_t + \Delta_t]$. If $\Delta_t > 1/T$, they choose the price $p_t = a_t + 1/2^{2^{k_t}}$ where $k_t = \lfloor 1 + \log_2 \log_2 \Delta_t^{-1} \rfloor$ (this is approximately equivalent to choosing $p_t = a_t + \Delta_t^2$). Otherwise (if $\Delta_t \leq 1/T$), they set their price equal to a_t . Moreover, they show that this bound is tight up to constant factors:

Theorem 1 (Kleinberg and Leighton [13]). *The optimal regret for the contextual search problem with pricing loss in one dimension is $\Theta(\log \log T)$.*

Their result implies a lower bound for the d -dimensional problem. If the adversary only uses coordinate vectors $e_i = (0 \dots 010 \dots 0)$ as contexts, then the problem reduces to d independent instances of the one dimensional pricing problem.

Corollary 2. *Any algorithm for the d -dimensional contextual search problem with pricing loss must incur $\Omega(d \log \log T)$ regret.*

IV. TWO DIMENSIONAL CASE

We start by showing how to obtain optimal regret for both loss functions in the two dimensional case. We highlight this special case since it is simple to visualize and conveys the geometric intuition for the general case. Moreover, it can be explained using only elementary plane geometry.

A. Symmetric loss

Our general approach will be to maintain a potential function of the current knowledge set which decreases each round by an amount proportional to the loss. Since at each time t , the loss is bounded by the width $\text{width}(S_t, u_t)$ of the knowledge S_t in direction u_t , it suffices to show that our potential function decreases each round by some amount proportional to the width of the current knowledge set.

What should we pick as our potential function? Inspired by the one-dimensional case, where one can take the potential function to be the length of the current interval, a natural candidate for the potential function is the area of the current knowledge set. Unfortunately, this does not work; if the knowledge set is long in the direction of u_t and skinny in the perpendicular direction (e.g. Figure 2a), then it can have large width but arbitrarily small area.

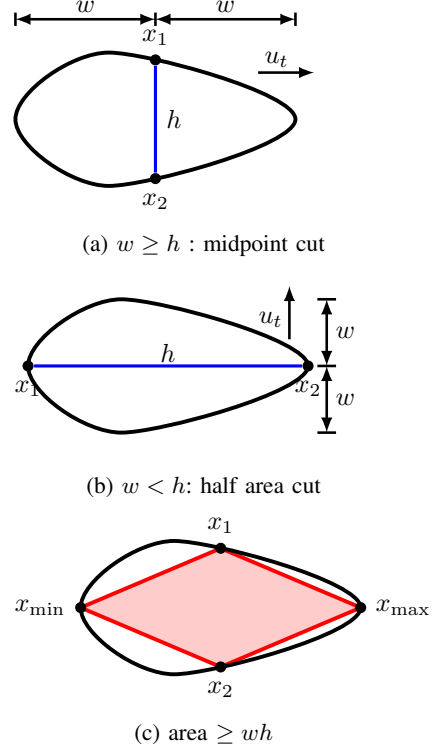


Figure 2

Ultimately we want to make the width of the knowledge set as small as possible in any given direction. This motivates a second choice of potential function: the average width of the knowledge set, i.e., $\frac{1}{2\pi} \int_0^{2\pi} \text{width}(S_t, u_\theta) d\theta$ where $u_\theta = (\cos \theta, \sin \theta)$. A result of Cauchy (see Section 5.5 in [12]) shows that the average width of a convex 2-dimensional shape is proportional to the perimeter, so this potential function can alternately be thought of as the perimeter of the knowledge set.

Unfortunately, this too does not quite work. Now, if the set S_t is thin in the direction u_t and long in the perpendicular direction (e.g. Figure 2b), any cut will result in a negligible decrease in perimeter (in particular, the perimeter decreases by $O(w^2)$ instead of $\Theta(w)$).

This motivates us to consider an algorithm that keeps track of two potential functions: the perimeter P_t , and the square root of the area $\sqrt{A_t}$. Each iteration, the algorithm will (depending on the shape of the knowledge set) choose one of these two potentials to make progress in. If S_t is long in the u_t direction, cutting it through the midpoint will allow us to decrease the perimeter by an amount proportional to the loss incurred (Figure 2a). If S_t is thin in the u_t direction, then we can charge the loss to the square root of the area (Figure 2b).

In Figure 3 we describe how to compute the guess p_t from the knowledge set S_t and u_t . Recall that the full setup

together with how knowledge sets are updated is defined in Algorithm 1.

Figure 3: 2D-SymmetricSearch

- 1: $w = \frac{1}{2}(\bar{p}_t - p_t)$ and $p_t^{\text{mid}} = \frac{1}{2}(\bar{p}_t + p_t)$
- 2: $h = \text{length of the segment } S_t \cap \{x; \langle u_t, x \rangle = p_t^{\text{mid}}\}$
- 3: **if** $w \geq h$ **then**
- 4: set $p_t = p_t^{\text{mid}}$
- 5: **else**
- 6: set p_t such that $\text{Area}(S_t^+) = \text{Area}(S_t^-)$
- 7: **end if**

Theorem 3. *The 2D-SymmetricSearch algorithm (Figure 3) has regret bounded by $8 + 2\sqrt{2}$ for the symmetric loss.*

Proof: We will keep track of the perimeter P_t and the area A_t of the knowledge set S_t and consider the potential function $\Phi_t = P_t + \sqrt{A_t}/C$, where the constant $C = (1 - \sqrt{1/2})/2$. We will show that every round this potential decreases by at least the regret we incur that round:

$$\Phi_t - \Phi_{t+1} \geq |\langle u_t, v \rangle - p_t|.$$

This implies that the total regret is bounded by $\text{Reg} \leq \Phi_1 = 4 + 2/(1 - \sqrt{1/2}) = 8 + 2\sqrt{2}$. We will write ℓ_t as shorthand for the loss $\ell(\langle u_t, v \rangle, p_t) = |\langle u_t, v \rangle - p_t|$ at time t .

We first note that both P_t and A_t are decreasing in t . This follows from the fact that S_{t+1} is a convex subset of S_t . We will show that when $w \geq h$, P_t decreases by at least ℓ_t , whereas when $w < h$, $\sqrt{A_t}$ decreases by at least ℓ_t .

Case $w \geq h$. In this case, $p_t = p_t^{\text{mid}}$. We claim here that $P_t - P_{t+1} \geq w$. To see this, let x_1 and x_2 be the two endpoints of the line segment $S_t \cap \{x; \langle u_t, x \rangle = p_t^{\text{mid}}\}$ (so that $h = \|x_1 - x_2\|$). Without loss of generality, assume $S_{t+1} = S_t^-$ (the other case is analogous).

Note that the boundary of S_{t+1} is the same as the boundary of S_t , with the exception that the segment of the boundary of S_t in the half-space $\{x; \langle u_t, x \rangle \geq p_t^{\text{mid}}\}$ has been replaced with the line segment $\overline{x_1 x_2}$. The part of boundary of S_t in the halfspace $\{x; \langle u_t, x \rangle \geq p_{\text{mid}}\}$, reaches some point on the line $\langle u_t, x \rangle = \bar{p}_t$, and returns to x_2 . Since $\bar{p}_t - p_t^{\text{mid}} = w$, any such path must have length at least $2w$. From the fact that $w \geq h$, it follows that:

$$P_t - P_{t+1} \geq 2w - h \geq 2w - w \geq w \geq \ell_t.$$

Case $w < h$. We set p_t such that $A_{t+1} = A_t/2$, so

$$\begin{aligned} \Phi_t - \Phi_{t+1} &\geq (\sqrt{A_t} - \sqrt{A_{t+1}})/C \\ &\geq \sqrt{A_t}(1 - \sqrt{1/2})/C \\ &= 2\sqrt{A_t}. \end{aligned}$$

If we argue that $2\sqrt{A_t} \geq 2w \geq \ell_t$ we are done. To show this, define x_1 and x_2 as before, and let x_{\min} be a point in S_t that satisfies $\langle u_t, x_{\min} \rangle = \underline{p}_t$, and likewise let x_{\max} be a point in S_t that satisfies $\langle u_t, x_{\max} \rangle = \bar{p}_t$. Since S_t is convex, it contains the two triangles with endpoints (x_1, x_2, x_{\max}) and (x_1, x_2, x_{\min}) , see Figure 2c. But these two triangles are disjoint and each have area at least $\frac{1}{2}wh$, so $A_t \geq wh \geq w^2$, since $h > w$. It follows that $\sqrt{A_t} \geq w$. ■

B. Pricing loss

To minimize regret for pricing loss, we want to somehow combine our above insight of looking at both the area and the perimeter with Kleinberg and Leighton's bucketing procedure for the 1D case. At first we might try to do this bucketing procedure just with the area.

That is, if the area of the current knowledge set belongs to the interval $[\Delta^2, \Delta]$, choose a price that carves off a subset of total area Δ^2 . Now, if you overprice, you incur $O(1)$ regret, but the area of your knowledge set shrinks to Δ^2 (and belongs to a new "bucket"). On the other hand, if you underprice, your area decreases by at most Δ , so you can underprice at most Δ^{-1} times. If the regret you incurred this round was at most $O(\Delta)$, then this means that you would incur at most $O(1)$ total regret underpricing while your area belongs to this interval.

This would be true if the regret per round was at most the area of the knowledge set. Unfortunately, as noted earlier, this is not true; the regret per round scales with the width of the knowledge set, not the area, and you can have knowledge sets with large width and small area. The trick, as before, is to look at both area and perimeter, and argue that at each step the bucketization argument for at least one of these quantities holds.

More specifically, let P_t again be the perimeter of the knowledge set at time t , A_t be the area of the knowledge set at time t , and let $A'_t = 2\sqrt{\pi A_t}$ be a normalization of A_t . Note that by the isoperimetric inequality for dimensions (which says that of all shapes with a given area, the circle has the least perimeter), we have that $P_t \geq A'_t$.

Let $\ell_k = \exp(-1.5^k)$. Our buckets will be the intervals $(\ell_{k+1}, \ell_k]$. We will define the function $\text{bkt}(x) = k$ if $x \in (\ell_{k+1}, \ell_k]$. Our algorithm is described in Figure 4. The analysis follows.

Theorem 4. *The 2D-PricingSearch algorithm (Figure 4) has regret bounded by $O(\log \log T)$ for the pricing loss.*

Proof: We will divide the behavior of the algorithm into four cases (depending on which branch of the if statement in Figure 4 is taken), and argue the total regret sustained in each case is at most $O(\log \log T)$.

- *Case 1:* $w \leq 1/T$. Whenever this happens, we pick the minimum possible price in our convex set, so we definitely underprice and sustain regret at most $1/T$.

Figure 4: 2D-PricingSearch

- 1: define $\ell_k = \exp(-1.5^k)$ and $\text{bkt}(x) = k$ if $x \in (\ell_{k+1}, \ell_k]$.
- 2: $w = (\bar{p}_t - p_t)$
- 3: h_{max} = maximum length of a segment of the form $S_t \cap \{x; \langle u_t, x \rangle = p\}$
- 4: **if** $w < 1/T$ **then**
- 5: choose $p_t = \underline{p}_t$.
- 6: **else if** $\text{bkt}(A'_t) = \text{bkt}(P_t)$ **then**
- 7: set p_t such that $\text{Area}(S_t^-) = \ell_{\text{bkt}(A'_t)+1}^2/4\pi$.
- 8: **else if** $\text{bkt}(A'_t) > \text{bkt}(P_t)$ and $w < h_{max}$ **then**
- 9: set p_t such that $\text{Area}(S_t^-) = \ell_{\text{bkt}(A'_t)+1}^2/4\pi$.
- 10: **else if** $\text{bkt}(A'_t) > \text{bkt}(P_t)$ and $w \geq h_{max}$ **then**
- 11: set p_t such that $\text{Perimeter}(S_t) - \text{Perimeter}(S_t^+) = \frac{1}{2}\ell_{\text{bkt}(P_t)+1}$.
- 12: **end if**

The total regret sustained in this case over T rounds is therefore at most 1.

- *Case 2:* $\text{bkt}(A'_t) = \text{bkt}(P_t)$. Let $\text{bkt}(P_t) = \text{bkt}(A'_t) = k$. Note that $k \leq O(\log \log T)$, or else we would be in Case 1 (if $P_t < 1/T$, then $w < 1/T$). Therefore, fix k ; we will show that the total regret we incur for this value of k is at most $O(1)$. Summing over all $O(\log \log T)$ values of k gives us our upper bound.

If we overprice and do not make a sale, this contributes regret at most 1, but then $A'_{t+1} = \ell_{k+1}$ and we leave this bucket. If we underprice, our regret is at most $w \leq P_t \leq \ell_k$, and the decrease in area $A_{t+1} - A_t = \ell_{k+1}^2/4\pi$. It follows that we can underprice at most $4\pi\ell_k^2/\ell_{k+1}^2$ before leaving this bucket, and therefore we incur total regret at most

$$4\pi \frac{\ell_k^2}{\ell_{k+1}^2} \ell_k = 4\pi \exp(-3 \cdot 1.5^k + 2 \cdot 1.5^{k+1}) = 4\pi = O(1).$$

- *Case 3:* $\text{bkt}(A'_t) > \text{bkt}(P_t)$ and $w < h_{max}$. Let $\text{bkt}(A'_t) = r$ and $\text{bkt}(P_t) = k$. As in case 2, we will fix r and argue that total regret we incur for this r is at most $O(1)$. As before, if we overprice, A'_{t+1} becomes ℓ_{r+1} , so we incur total regret at most $O(1)$ from overpricing. Now, note that since $w < h_{max}$, S_t contains two disjoint triangles with base h_{max} and combined height w (see Figure 2c), so $A_t \geq wh_{max} > w^2$, and therefore $w < \sqrt{A_t} = A'_t/2\sqrt{\pi}$. Therefore, if we underprice, we incur regret at most $A'_t/2\sqrt{\pi} \leq \ell_r$. As before, the area decreases by at least $A_{t+1} - A_t = \ell_{r+1}^2/4\pi$ if we underprice, so we underprice at most $4\pi\ell_k^2/\ell_{k+1}^2$ before leaving this bucket, and therefore we incur total regret at most

$$4\pi \frac{\ell_r^2}{\ell_{r+1}^2} \ell_r = 4\pi \exp(-3 \cdot 1.5^r + 2 \cdot 1.5^{r+1}) = 4\pi = O(1).$$

- *Case 4:* $\text{bkt}(A'_t) > \text{bkt}(P_t)$ and $w \geq h_{max}$. Let $\text{bkt}(A'_t) = r$ and $\text{bkt}(P_t) = k$, and fix k . Now $A_t \geq wh_{max} \geq h_{max}^2$, so $h_{max} < A'_t/2\sqrt{\pi} \leq \ell_r/2\sqrt{\pi}$. First, note that it is in fact possible to set p_t so that $\text{Perimeter}(S_t) - \text{Perimeter}(S_t^+) = \frac{1}{2}\ell_{k+1}$, since the total perimeter is at least ℓ_{k+1} , so this corresponds to just cutting off a chunk (S_t^+) of S_t of perimeter $\text{Perimeter}(S_t) - \frac{1}{2}\ell_{k+1}$ (which is possible since this is nonnegative and less than $\text{Perimeter}(S_t)$).

If we overprice, then the perimeter of the new region (S_t^-) is equal to $\text{Perimeter}(S_t) - \text{Perimeter}(S_t^+) = \ell_{k+1}/2$, plus the length of the segment formed by the intersection of S_t with $\{x; \langle u_t, x \rangle = p_t\}$. The length of this segment is at most h_{max} , so the perimeter is at most $\ell_{k+1}/2 + \ell_r/2\sqrt{\pi} \leq \ell_{k+1}/2 + \ell_{k+1}/2\sqrt{\pi} \leq \ell_{k+1}$. This means we can overprice at most once before the perimeter changes buckets, and we thus incur at most $O(1)$ regret due to overpricing.

If we underprice, then we incur regret at most $w \leq P_t \leq \ell_k$, and the perimeter of the new region decreases by at least $\ell_{k+1}/2 - \ell_r/2\sqrt{\pi} \geq \ell_{k+1}/5$. This means we can underprice at most $5\ell_k/\ell_{k+1}$ times before we switch buckets, and we incur total regret at most

$$5 \frac{\ell_k}{\ell_{k+1}} \ell_k = 4 \exp(-2 \cdot 1.5^k + 1.5^{k+1}) \leq 4 = O(1).$$

V. INTERLUDE: INTRINSIC VOLUMES

The main idea in the two dimensional case was to balance between making progress in a two-dimensional measure of the knowledge set (the area) and in a one-dimensional measure (the perimeter). To generalize this idea to higher dimensions we will keep track of d potential functions, each corresponding to a j -dimensional measure of the knowledge set for each j in $\{1, 2, \dots, d\}$.

Luckily for us, one of the central objects of study in *integral geometry* (also known as *geometric probability*) corresponds exactly to a j -dimensional measure of a d -dimensional object. Many readers are undoubtedly familiar with two of these measures, namely volume (the d -dimensional measure) and surface area (the $(d-1)$ -dimensional area) but it is less clear how to define the 1-dimensional measure of a three-dimensional convex set (indeed, Shanel [17] jokingly calls his lecture notes on the topic ‘‘What is the length of a potato?’’). These measures are known as *intrinsic volumes* and they match our intuition for how a j -dimensional measure of a d -dimensional set should behave (in particular, reducing to the regular j -dimensional volume as the set approaches a j -dimensional object).




Figure 5: Steiner's formula for 2D: $\text{Area}(K + \epsilon B) = \text{Area}(K) + \text{Perimeter}(K) \cdot \epsilon + \pi \epsilon^2$

We now present a formal definition of intrinsic volumes and summarize their most important properties. We refer to the excellent book by Klain and Rota [12] for a comprehensive introduction to integral geometry.

Intrinsic volumes can be defined as the coefficients that arise in Steiner's formula for the the volume of the (Minkowski) sum of a convex set $K \subseteq \mathbb{R}^d$ and an unit ball B . Steiner [17] shows that the $\text{Vol}(K + \epsilon B)$ is a polynomial in ϵ and the intrinsic volumes $V_j(K)$ are the (normalized) coefficients of this polynomial:

$$\text{Vol}(K + \epsilon B) = \sum_{j=0}^d \kappa_{d-j} V_j(K) \epsilon^{d-j} \quad (1)$$

where κ_{d-j} is the volume of the $(d-j)$ -dimensional unit ball. An useful exercise to get intuition about intrinsic volumes is to directly compute the intrinsic volumes of the parallelotope $[0, a_1] \times [0, a_2] \times [0, a_d]$. It is easy to check for $d = 2$ and 3 (see Figure 5) that $V_d = a_1 a_2 \dots a_d$, $V_1 = a_1 + a_2 + \dots + a_d$ and $V_0 = 1$. More generally V_j corresponds to the symmetric polynomial of degree j : $V_j = \sum_{S \subseteq [d], |S|=j} a_S$ where $a_S = \prod_{s \in S} a_s$. In particular for $[0, 1]^d$ the j -th intrinsic volume is $V_j = \binom{d}{j}$.

Definition 5 (Valuations). Let Conv_d be the class of compact convex bodies in \mathbb{R}^d . A valuation is a map $\nu : \text{Conv}_d \rightarrow \mathbb{R}$ such that $\nu(\emptyset) = 0$ and for every $S_1, S_2 \in \text{Conv}_d$ satisfying $S_1 \cup S_2 \in \text{Conv}_d$ it holds that

$$\nu(S_1 \cup S_2) + \nu(S_1 \cap S_2) = \nu(S_1) + \nu(S_2).$$

A valuation is said to be monotone if $\nu(S) \leq \nu(S')$ whenever $S \subseteq S'$. A valuation is said to be non-negative if $\nu(S) \geq 0$. Finally, a valuation is rigid if $\nu(S) = \nu(T(S))$ for every rigid motion (i.e. rotations and translations) T of \mathbb{R}^d .

To define what it means for a valuation to be continuous, we need a notion of distance between convex sets. We define the Hausdorff distance $\delta(K, L)$ between two sets $K, L \in \text{Conv}_d$ to be the minimum ϵ such that $K + \epsilon B \subseteq L$ and $L + \epsilon B \subseteq K$ where B is the unit ball. This notion of distance allows us to define limits: a sequence $K_t \in \text{Conv}_d$ converges to $K \in \text{Conv}_d$ (we write this as

$K_t \rightarrow K$) if $\delta(K_t, K) \rightarrow 0$. Continuity can now be defined in the natural way.

Definition 6 (Continuity). A valuation function ν is continuous if whenever $K_t \rightarrow K$ then $\nu(K_t) \rightarrow \nu(K)$.

Theorem 7. The intrinsic volumes are non-negative monotone continuous rigid valuations.

In fact the intrinsic volumes are quite special since they form a basis for the set of all valuations with this property. This constitutes the fundamental result of the field of integral geometry:

Theorem 8 (Hadwiger). If ν is a continuous rigid valuation of Conv_d , then there are constants c_0, \dots, c_d such that $\nu = \sum_{i=0}^d c_i V_i$.

Next we describe a few important properties of intrinsic valuations that will be useful in the analysis of our algorithms:

Theorem 9 (Homogeneity). The map V_j is j -homogenous, i.e., $V_j(\alpha K) = \alpha^j V_j(K)$ for any $\alpha \in \mathbb{R}_{\geq 0}$.

Theorem 10 (Ambient independence). Intrinsic volumes are independent of the ambient space, i.e. if $K \in \text{Conv}_d$ and K' is a copy of K embedded in a larger dimensional space

$$K' = T(\{(x, 0_k); x \in K, 0_k \in \mathbb{R}^k\}) \in \text{Conv}_{d+k}$$

for a rigid transformation T , then for any $j \leq d$, we have $V_j(K) = V_j(K')$.

We now provide an inequality between intrinsic volumes which we will use later to derive an isoperimetric inequality for intrinsic volumes. The following inequality is a consequence of the Alexandrov-Fenchel inequality due to McMullen [15].

Theorem 11 (Inequality on intrinsic volumes). If $S \in \text{Conv}_d$ and any $i \geq 1$ then

$$V_i(S)^2 \geq \frac{i+1}{i} V_{i-1}(S) V_{i+1}(S).$$

One beautiful consequence of Hadwiger's theorem is a probabilistic interpretation of intrinsic volumes as the expected volume of the projection of a set onto a random subspace. To make this precise, define the Grassmannian $\text{Gr}(d, k)$ as the collection of all k -dimensional linear subspaces of \mathbb{R}^d . The Haar measure on the Grassmannian is the unique probability measure on $\text{Gr}(d, k)$ that is invariant under rotations in \mathbb{R}^d (i.e., $SO(\mathbb{R}^d)$).

Theorem 12 (Random Projections). For any $K \in \text{Conv}_d$, the j -th intrinsic volume

$$V_j(K) = \mathbb{E}_{H \sim \text{Gr}(d, k)} [\text{Vol}(\pi_H(K))]$$

where $H \sim \text{Gr}(d, k)$ is a k -dimensional subspace H sampled according to the Haar measure, π_H is the projection

on H and $\text{Vol}(\pi_H(K))$ is the usual (k -dimensional) volume on H .

A remark on notation: we use Vol to denote the standard notion of volume and V_j to denote intrinsic volumes. When analyzing an object in a d -dimensional (sub)space, then $\text{Vol} = V_d$.

VI. HIGHER DIMENSIONS

In this section, we generalize our algorithms from Section IV from the two-dimensional case to the general multi-dimensional case.

Both results require as a central component lower bounds on the intrinsic volumes of high dimensional cones. These bounds relate the intrinsic volume of a cone to the product of the cone’s height and the intrinsic volume of the cone’s base (a sort of “Fubini’s theorem” for intrinsic volumes).

More formally, a *cone* S in \mathbb{R}^{d+1} is the convex hull of a d -dimensional convex set K and a point $p \in \mathbb{R}^{d+1}$. If the distance from p to the affine subspace containing K is h , we say the cone has *height* h and *base* K . The lemma we require is the following.

Lemma 13 (Cone Lemma). *Let K be a convex set in \mathbb{R}^d , and let S be a cone in \mathbb{R}^{d+1} with base K and height h . Then, for all $0 \leq j \leq d$,*

$$V_{j+1}(S) \geq \frac{1}{j+1} h V_j(K).$$

In the two-dimensional case, this lemma manifests itself when we use the fact that the perimeter of a convex set with height h is at least h . We note that when $j = d$, Lemma 13 holds with equality and is a simple exercise in elementary calculus. On the other hand, when $0 \leq j < d$, there is no straightforward formula for the $(j+1)$ -th intrinsic volume of a set in terms of the j -th intrinsic volume of its cross sections.

We begin by taking the Cone Lemma as true, and discuss how to use it to generalize our contextual search algorithms to higher dimensions in Sections VI-A (for symmetric loss) and VI-B (for pricing loss). We then prove the Cone Lemma in Section VI-C. Finally, in Section VI-D, we argue that both algorithms can be implemented efficiently.

A. Symmetric loss

In this section we present a $O_d(1)$ regret algorithm for the contextual search problem with symmetric loss in d dimensions. The algorithm, which we call `SymmetricSearch`, is presented in Figure 6.

Recall that in two dimensions, we always managed to choose p_t so that the loss from that round is bounded by the decrease in either the perimeter or the square root of the area. The main idea of Figure 6 is to similarly choose p_t such that the loss is bounded by the decrease in one of the intrinsic volumes, appropriately normalized. As before, if the

width is large enough, we bound the loss by the decrease in the average width (i.e. the one-dimensional intrinsic volume $V_1(S)$). As the width gets smaller, we charge the loss to progressively higher-dimensional intrinsic volumes.

Constants c_0 through c_{d-1} in Figure 6 are defined so that $c_0 = 1$ and $c_i/c_{i-1} = \frac{1}{2^i}$. In other words, $c_i = \frac{1}{2^{i-1}i!}$. Constant c_0 is only used in the analysis.

Figure 6: `SymmetricSearch`

- 1: $w = \frac{1}{2} \text{width}(S_t; u_t)$
- 2: **for** $i = 1$ to d **do**
- 3: define $p_i \in \mathbb{R}$ such that $V_i(S_t^+(p_i; u_t)) = V_i(S_t^-(p_i; u_t))$.
- 4: define $K_i = \{x \in S_t; \langle u_t, x \rangle = p_i\}$.
- 5: define $L_i = (V_i(K_i)/c_i)^{1/i}$ (set $L_0 = \infty$).
- 6: **end for**
- 7: find j such that $L_{j-1} \geq w \geq L_j$
- 8: set $p_t = p_j$.

We first argue that this algorithm is well-defined:

Lemma 14. *`SymmetricSearch` (Figure 6) is well-defined, i.e., there is always a choice of p_i and j that satisfies the required properties.*

Proof: We begin by arguing that there exists a p_i such that $V_i(S_t^+(p_i; u_t)) = V_i(S_t^-(p_i; u_t))$. To see this, note that the functions $\phi^+(x) = V_i(S_t^+(x; u_t))$ and $\phi^-(x) = V_i(S_t^-(x; u_t))$ are continuous on $[p_t, \bar{p}_t]$ since the intrinsic volumes are continuous with respect to Hausdorff distance (Definition 6 and Theorem 7). Moreover, since intrinsic volumes are monotone (Theorem 7), $\phi^+(x)$ is decreasing and ϕ^- is increasing on this interval. Finally, since $\phi^+(p_t) = \phi^-(\bar{p}_t)$, it follows from the Intermediate Value Theorem that there exists a p_i where $\phi^+(p_i) = \phi^-(p_i)$, as desired.

To see that there exists a j such that $L_{j-1} \geq w \geq L_j$, note that $L_d = 0$ since K_d is in a $d-1$ -dimensional hyperplane, so the segments $[L_j, L_{j-1}]$ for $L_j < L_{j-1}$ cover the entire $[0, \infty)$. It follows that one such interval must contain w . ■

Before we proceed to the regret bound, we will show the following two lemmas. The first lemma shows that if we pick j in this manner, then $V_j(S_t)^{1/j}$ will be at least $\Omega(w)$.

Lemma 15. $V_j(S_t) \geq \frac{1}{j} c_{j-1} w^j$.

Proof: For $j = 1$, note that S_t contains a segment of length $2w$, so $V_1(S_t)$ is at most the 1-dimensional intrinsic volume of that segment, which is exactly $2w$ (Theorem 10).

For $j > 1$, we know that S_t contains a cone with base K_{j-1} and height w (since the width of S_t is $2w$, there is a point at least distance w from the plane $H_{j-1} = \{x; \langle u_t, x \rangle = p_{j-1}\}$). By Theorem 13 and the fact that V_j is monotone, this implies that:

$$V_j(S_t) \geq \frac{1}{j} w V_{j-1}(K_{j-1}).$$

Since $w < L_{j-1} = (V_{j-1}(K_{j-1})/c_{j-1})^{1/j-1}$, we have that $V_{j-1}(K_{j-1}) \geq c_{j-1}w^{j-1}$. Substituting this into the previous expression, we obtain the desired result. ■

The second lemma shows that if we pick j in this manner, then $V_j(S_{t+1})/V_j(S_t)$ is bounded above by a constant strictly less than 1.

Lemma 16. $V_j(S_{t+1}) \leq \frac{3}{4}V_j(S_t)$.

Proof: The set S_{t+1} is equal to either $S^+ = S_t^+(p_j; u_t)$ or $S^- = S_t^-(p_j; u_t)$. Our choice of p_j is such that $V_j(S^-) = V_j(S^+)$. Therefore:

$$\begin{aligned} 2V_j(S_{t+1}) &= V_j(S^+) + V_j(S^-) \\ &= V_j(S^- \cap S^+) + V_j(S^- \cup S^+) \\ &= V_j(K_j) + V_j(S_t) \end{aligned}$$

To bound $V_j(K_j)$ in terms of $V_j(S_t)$ we observe that $w \geq L_j = (V_j(K_j)/c_j)^{1/j}$ so $V_j(K_j) \leq c_jw^j$. Plugging the previous lemma we get $V_j(K_j) \leq j \frac{c_j}{c_{j-1}} V_j(S_t) = \frac{1}{2}V_j(S_t)$ by the choice of constants. Substituting this inequality into the above equation gives us the desired result. ■

Together, these lemmas let us argue that each round, the sum of the normalized intrinsic volumes $V_i(S_t)^{1/i}$ decreases by at least $\Omega(w)$ (and hence the total regret is constant).

Theorem 17. *The SymmetricSearch algorithm (Figure 6) has regret bounded by $O(d^4)$ for the symmetric loss.*

Proof: We will show that for the potential function $\Phi_t = \sum_{i=1}^d i^2 V_i(S_t)^{1/i}$ we can always charge the loss to the decrease in potential, i.e., $\Phi_t - \Phi_{t+1} \geq \Omega(w) \geq \Omega(\ell_t)$ and therefore, $\text{Reg} \leq \sum_{t=1}^{\infty} \ell_t \leq O(\Phi_1)$. The initial potential is

$$\begin{aligned} \Phi_1 &= \sum_{i=1}^d i^2 V_i([0, 1]^d)^{1/i} \\ &= \sum_{i=1}^d i^2 \binom{d}{i}^{1/i} \\ &\leq \sum_{i=1}^d i^2 O(d) \\ &= O(d^4) \end{aligned}$$

Since $V_j(S_t) \geq V_j(S_{t+1})$ by monotonicity, we can bound the potential change by $\Phi_t - \Phi_{t+1} \geq j^2[V_j(S_t)^{1/j} - V_j(S_{t+1})^{1/j}]$. We now show that this last term is $\Omega(w)$:

$$\begin{aligned} &j^2[V_j(S_t)^{1/j} - V_j(S_{t+1})^{1/j}] \\ &\geq j^2 \left(1 - \left(\frac{3}{4}\right)^{1/j}\right) V_j(S_t)^{1/j} \\ &\geq j^2 \left(1 - \left(\frac{3}{4}\right)^{1/j}\right) \left(\frac{c_{j-1}}{j}\right)^{1/j} w \\ &\geq j^2 (1 - (1 - \log(4/3)/j)) \left(\frac{1}{2^{j-2}j!}\right)^{1/j} w \\ &\geq j^2 \left(\frac{\log(4/3)}{j}\right) \left(\frac{e}{2j}\right) w \\ &\geq \Omega(w). \end{aligned}$$

Here the first inequality follows from Lemma 16 and the second from Lemma 15. ■

B. Pricing loss

In the 2-dimensional version of the dynamic pricing problem, we decomposed the range of each potential into $O(\log \log T)$ buckets and used the isoperimetric inequality $\sqrt{4\pi A_t} \leq P_t$ to argue that (when suitably normalized), the area always belonged to a higher bucket than the perimeter. To apply the same idea here, we will apply our inequality on intrinsic volumes (Theorem 11) to obtain an isoperimetric inequality for intrinsic volumes:

Lemma 18 (Isoperimetric inequality). *For any $S \in \text{Conv}_d$ and any $i \geq 1$ it holds that*

$$(i!V_i(S))^{1/i} \geq ((i+1)!V_{i+1}(S))^{1/(i+1)}.$$

Proof: We proceed by induction. For $i = 1$, note that Theorem 11 gives us that $V_1(S)^2 \geq 2V_0(S)V_2(S)$. Since $V_0(S)$ equals 1 for any convex set S , this reduces to $V_1(S) \geq \sqrt{2!V_2(S)}$.

Now assume via the inductive hypothesis that we have proven the claim for all $j \leq i$. From Theorem 11 we have that

$$\begin{aligned} V_i(S)^2 &\geq \frac{i+1}{i} V_{i-1}(S)V_{i+1}(S) \\ &= \frac{i+1}{i!} ((i-1)!V_{i-1}(S))V_{i+1}(S) \\ &\geq \frac{i+1}{i!} (i!V_i(S))^{(i-1)/i} V_{i+1}(S) \\ &= \frac{1}{i!^{(i+1)/i}} V_i(S)^{(i-1)/i} (i+1)!V_{i+1}(S). \end{aligned}$$

Rearranging, this reduces to $(i!V_i(S))^{(i+1)/i} \geq (i+1)!V_{i+1}(S)$, and therefore $(i!V_i(S))^{1/i} \geq ((i+1)!V_{i+1}(S))^{1/(i+1)}$. ■

Inspired by the isoperimetric inequality we will keep track of the following ‘‘potentials’’ (these vary with t , but we will omit the subscript for notational convenience):

$$\varphi_i = (i!V_i(S_t))^{1/i}$$

Since $S_1 = [0, 1]^d$, their initial values will be given by $\varphi_i = (i! \cdot \binom{d}{i})^{1/i} < di \leq d^2$. Since those quantities are monotone non-increasing, they will be in the interval $[0, d^2]$. We will divide this interval in ranges of doubly-exponentially decreasing length (as in one and two dimensions). The ranges will be $(\ell_{k+1}, \ell_k]$ where

$$\ell_k = d^2 \exp(-\alpha^k) \quad \text{for } \alpha = 1 + 1/d$$

To keep track of which range each of our potentials ϕ_i belongs to, define k_i so that $\varphi_i \in (\ell_{k_i+1}, \ell_{k_i}]$. By the isoperimetric inequality we know that:

$$\varphi_1 \geq \varphi_2 \geq \dots \geq \varphi_d \quad k_1 \leq k_2 \leq \dots \leq k_d$$

Recall that in the 2-dimensional case, whenever the perimeter and the area were in the same range, we chose to make progress in the area. To extend this idea to higher dimensions, whenever many ϕ_i belong to the same range and we decide to make progress on that range, we will always choose the largest such φ_i :

$$M(i) = \max\{j; k_i = k_j\}$$

The complete method is summarized in Figure 7. As before, constants c_0 through c_{d-1} in Figure 7 are defined so that $c_0 = 1$ and $c_i/c_{i-1} = \frac{1}{2^i}$. In other words, $c_i = \frac{1}{2^{i-1}i!}$.

We begin by arguing that our algorithm is well-defined. We ask the reader to recall the notation $\underline{p}_t = \min_{x \in S_t} \langle u_t, x \rangle$ and $\bar{p}_t = \max_{x \in S_t} \langle u_t, x \rangle$.

Figure 7: PricingSearch

```

1:  $w = \frac{1}{2} \text{width}(S_t; u_t)$ 
2: for  $i = 1$  to  $d$  do
3:   let  $\varphi_i = (i! \cdot V_i(S_t))^{1/i}$  and  $k_i$  such that  $\varphi_i \in (\ell_{k_i+1}, \ell_{k_i}]$ 
4:   if  $V_i(S_t) - V_i(S_t^+(\bar{p}_t; u_t)) > \ell_{k_i+1}^i / (2 \cdot i!)$  then
5:     choose  $p_i$  such that  $V_i(S_t) - V_i(S_t^+(p_i; u_t)) = \ell_{k_i+1}^i / (2 \cdot i!)$ 
6:   else
7:     choose  $p_i = \bar{p}_t$ 
8:   end if
9:   define  $K_i = \{x \in S_t; \langle u_t, x \rangle = p_i\}$ 
10:  define  $L_i = (V_i(K_i)/c_i)^{1/i}$  (define  $L_0 = \infty$ )
11: end for
12: if  $w < 1/T$  then
13:   set  $p_t = \underline{p}_t$ .
14: else
15:   let  $M(i) = \max\{j; k_i = k_j\}$ 
16:   find a  $j$  such that  $L_{j-1} \geq w \geq L_{M(j)}$ 
17:   let  $J = M(j)$  and set  $p_t = p_J$ .
18: end if

```

Lemma 19. *PricingSearch (Figure 7) is well defined, i.e., it is always possible to choose p_i and j with the desired properties.*

Proof: For the choice of p_i , if $V_i(S_t) - V_i(S_t^+(\bar{p}_t; u_t)) > \ell_{k_i+1}^i / (2 \cdot i!)$, then the function $\phi_i : [p_t, \bar{p}_t] \rightarrow \mathbb{R}$, $\phi_i(p) = V_i(S_t) - V_i(S_t^+(p; u_t))$ is continuous and monotone with $\phi_i(p_t) = 0$ and $\phi_i(\bar{p}_t) > \frac{1}{2}i! \cdot \ell_{k_i+1}^i$ so this guarantees the existence of such p_i .

For the choice of j , let $0 = i_0 < i_1 < \dots < i_a = d$ be the indices i such that $M(i) = i$. Notice that the intervals $[L_{i_{s+1}}, L_{i_s})$ are of the form $[L_{M(i)}, L_{i-1})$ for $i = i_s + 1$. Finally notice that the intervals $[L_{i_{s+1}}, L_{i_s})$ cover the entire interval $[L_d, L_0) = [0, \infty)$ so one of them must contain w . ■

Before we proceed to the main analysis, we begin by proving a couple of lemmas regarding the ranges $(\ell_{k+1}, \ell_k]$ of the intrinsic volumes before and after each iteration. The first lemma says that if we overprice (i.e. S_t^- is chosen) the quantity φ_J jumps from the range $[\ell_{k_J+1}, \ell_{k_J})$ to the next range $(\ell_{k_J+2}, \ell_{k_J+1}]$.

Lemma 20. $[J! \cdot V_J(S_t^-(p_J; u_t))]^{1/J} \leq \ell_{k_J+1}$

Proof: We abbreviate $S_t^-(p_J; u_t)$ and $S_t^+(p_J; u_t)$ by S^- and S^+ respectively. Using the fact that V_J is a valuation and that $S^- \cap S^+ = K_J$ we have that:

$$\begin{aligned} V_J(S^-) &= V_J(S_t) - V_J(S^+) + V_J(K_J) \\ &\leq \ell_{k_J+1}^J / (2 \cdot J!) + V_J(K_J) \end{aligned}$$

It remains to show that $V_J(K_J) \leq \ell_{k_J+1}^J / (2 \cdot J!)$. To do this, we will again use the Cone Lemma to obtain the following inequalities:

$$\begin{aligned} \frac{1}{J+1} V_J(K_J) w &\leq V_{J+1}(S_t) \leq \frac{1}{(J+1)!} \ell_{k_{J+1}}^{J+1} \\ &\leq \frac{1}{(J+1)!} \ell_{(k_J)+1}^{J+1} \end{aligned}$$

The first inequality is the Cone Lemma (Lemma 13) applied to the fact that S_t contains a cone of base K_J and height at least w . The second inequality comes from the definition of k_J and the third inequality comes from the fact that $J = M(J)$ so $k_{J+1} \geq k_J + 1$.

Finally, observe that because of our choice of J , $w \geq L_J = (V_J(K_J)/c_J)^{1/J}$. Substituting in the previous equation we obtain:

$$\frac{1}{J+1} V_J(K_J)^{(J+1)/J} (c_J)^{-1/J} \leq \frac{1}{(J+1)!} \ell_{(k_J)+1}^{J+1}$$

Substituting the definition of c_J and simplifying, we get the desired bound of $V_J(K_J) \leq \ell_{k_J+1}^J / (2 \cdot J!)$. ■

We next show that, for our chosen J , if we underprice, then the J th intrinsic volume of our knowledge set decreases by at least $\ell_{k_J+1}^J$. This will allow us to bound the number of times we can potentially underprice before k_J changes (in particular, it is at most $2\ell_{k_J}^J / \ell_{k_J+1}^J$).

Lemma 21. $V_J(S_t) - V_J(S_t^+(p_J; u_t)) = \ell_{k_J+1}^J / (2 \cdot J!)$

Proof: Note that this equality is guaranteed by the algorithm's choice of p_J , except when $V_J(S_t) - V_J(S_t^+(\bar{p}_t; u_t)) < \ell_{k_J+1}^J / (2 \cdot J!)$ and $p_J = \bar{p}_t$. However, in this case, $S_t^-(\bar{p}_t; u_t) = S_t$ by the definition of \bar{p}_t . Lemma 20 then implies that $[J! \cdot V_J(S_t)]^{1/J} \leq \ell_{k_J+1}$, but this contradicts the definition of k_J . ■

We now show that in each round, the width of the knowledge set (and thus our loss) is at most $2\ell_{k_J}$.

Lemma 22. $w \leq 2\ell_{k_J}$

Proof: We will derive both an upper and lower bound on $V_j(S_t)$. For the upper bound we again apply the Cone Lemma (Lemma 13).

$$V_j(S_t) \geq \frac{1}{j} V_{j-1}(K_{j-1})w \geq \frac{1}{j} (c_{j-1}w^{j-1})w$$

If $j > 1$, then the first inequality holds since S_t contains a cone of base K_{j-1} and height w , and the second inequality follows from the fact that $w \leq L_{j-1}$. If $j = 1$, then we observe that S_t contains a segment of length w , so $V_1(S_t) \leq w$.

To get a lower bound on $V_j(S_t)$, simply note that

$$V_j(S_t) \leq \ell_{k_j}^j / (j!) = \ell_{k_J}^j / (j!)$$

where the first inequality follows from the definition of k_j and the second from the fact that $k_j = k_J$ since $J = M(j)$.

Together the bounds imply that $c_{j-1}w^j / j \leq \ell_{k_J}^j / (j!)$. Substituting in the value of c_{j-1} and simplifying we obtain that $w \leq 2\ell_{k_J}$. ■

Finally, we argue that if w is large enough (at least $1/T$), then k_J is at most $O_d(\log \log T)$. Once w is at most $1/T$, we can always price at \underline{p}_t and incur at most $O(1)$ additional regret, so this provides a bound for the number of times we can e.g. overprice.

Lemma 23. *In iterations where $w \geq 1/T$, then $k_J \leq O(d \log \log(dT))$.*

Proof: It follows directly from Lemma 22: $1/T \leq w \leq 2\ell_{k_J} = 2d^2 \exp(-\alpha^{k_J})$. Simplifying the expression we get $k_J \leq O(d \log \log(dT))$. ■

We are now ready to prove our main result:

Theorem 24. *The total loss of PricingSearch (Figure 7) is bounded by $O(d^4 \log \log(dT))$.*

Proof: We sum the loss in different cases. The first is when $w < 1/T$ and the algorithm prices at \underline{p}_t . In those occasions the algorithm always sells and the loss is at most $2w \leq 2/T$, so the total loss is at most 2.

The second case is when the algorithm overprices and doesn't sell. If the algorithm doesn't sell, then by Lemma 20, then ϕ_J goes from range $(\ell_{k_J+1}, \ell_{k_J}]$ to the next range $(\ell_{k_J+2}, \ell_{k_J+1}]$. Since $k_J \leq O(d \log \log(dT))$ by Lemma 23

this can happen at most this many times for each index J . Since there are d such indices and the loss of each event is at most 1, the total loss is bounded by $O(d^2 \log \log(dT))$.

The final case is when the algorithm underprices. The loss in this case is bounded by the width $2w$. We sum the total loss of events in which the algorithm overprices. We fix the selected index J and k_J . The loss in such a case is at most $2w \leq 4\ell_{k_J}$ by Lemma 22. Whenever this happens $S_{t+1} = S_t^+(p_J; u_t)$ so the J -th intrinsic volume decreases by $\ell_{k_J+1}^J / (2J!)$ since $V_J(S_t) - V_J(S_{t+1}) = V_J(S_t) - V_J(S_t^+) = \ell_{k_J+1}^J / (2J!)$ by Lemma 21. Since $V_J(S_t) \leq \ell_{k_J}^J / (J!)$. Therefore the total number of times it can happen is: $2\ell_{k_J}^J / \ell_{k_J+1}^J$. The total loss is at most the number of times the event can happen multiplied by the maximum loss for an event, which is:

$$\begin{aligned} \frac{2\ell_{k_J}^J}{\ell_{k_J+1}^J} \cdot (4\ell_{k_J}) &= 8 \frac{\ell_{k_J}^{J+1}}{\ell_{k_J+1}^J} \\ &= 8d^2 \exp(J\alpha^{k_J+1} - (J+1)\alpha^{k_J}) \\ &\leq 8d^2 \exp(\alpha^{k_J}(d\alpha - (d+1))) \\ &= 8d^2 \end{aligned}$$

since $\alpha = 1 + 1/d$. By summing over all d possible values of J and all $O(d \log \log(dT))$ values of k_J we obtain a total loss of $O(d^4 \log \log(dT))$. ■

C. Proof of the Cone Lemma

We will prove the Cone Lemma in three steps. We start by proving some geometric lemmas about how linear transformations affect intrinsic volumes. We then use these lemmas to bound the intrinsic volumes of cylinders. Finally, by approximating a cone as a stack of thin cylinders, we apply these bounds to prove the Cone Lemma.

1) *Geometric lemmas:* Define an α -stretch of \mathbb{R}^d as a linear transformation which contracts \mathbb{R}^d along some axis by a factor of α , leaving the remaining axes untouched (in other words, there is some coordinate system in which an α -stretch T_α sends (x_1, x_2, \dots, x_d) to $(\alpha x_1, x_2, \dots, x_d)$).

A contraction is a linear transformation $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $\|Tx\| \leq \|x\|$ for all $x \in \mathbb{R}^d$. An α -stretch is a contraction whenever $\alpha \in [0, 1]$.

Lemma 25. *Let H and H' be two (d -dimensional) hyperplanes in \mathbb{R}^{d+1} , whose normals are separated by angle θ . Let K be a convex body contained in H , and let K' be the projection of K onto H' . Then K' is (congruent to) a $(\cos \theta)$ -stretch of K .*

Proof: Without loss of generality, let H' be the hyperplane with orthonormal basis e_1, e_2, \dots, e_d , and let H be the hyperplane with orthonormal basis $e'_1 = (\cos \theta)e_1 + (\sin \theta)e_{d+1}, e'_2 = e_2, \dots, e'_d = e_d$. Note that a point $a_1 e'_1 + a_2 e'_2 + \dots + a_n e'_n$ in H , projects to the point $(\cos \theta)a_1 e_1 +$

$a_2e_2 + \dots + a_n e_n$ in H' . This is the definition of a $(\cos \theta)$ -stretch. ■

The next lemma bounds the change in the d -th volume of a $(d+1)$ -dimensional object when it is transformed by a contraction. The analysis will be based on the fact that for a $(d+1)$ -dimensional convex set S , $V_d(S)$ corresponds to half of the surface area. This fact can be derived either from Hadwiger's theorem (Theorem 8) or from Cauchy's formula for the surface area together with Theorem 12.

It is simpler to reason about the surface area of polyhedral convex sets (i.e. sets that can be described as a finite intersection of half-spaces). The boundary of a polyhedral convex set in \mathbb{R}^{d+1} can be described as a finite collection of facets, which are convex sets of dimension d . The surface area corresponds to the sum of the d -dimensional volume of the facets. For a 2-dimensional polytope the surface area correspond to the perimeter. For a 3-dimensional polytope the surface area corresponds to the sum of the area (the 2-dimensional volume) of the facets. For a general convex set K , the surface area can be computed as the limit of the surface area of K_t where K_t are polyhedral sets that converge (in the Hausdorff sense) to K . This is equivalent to the usual definition of the surface area as the surface integral of a volume element.

Given the discussion in the previous paragraph, to reason about how the surface area transforms after a linear transformation, it is enough to reason how the volume of d -dimensional convex sets (the facets) transform when the ambient \mathbb{R}^{d+1} space is transformed by a linear transformation.

Lemma 26. *Let $K \in \text{Conv}_{d+1}$ and T be a contraction, then*

$$V_d(T(K)) \geq \det T \cdot V_d(K)$$

Proof: By the previous discussion, $V_d(K)$ is proportional to the surface area of K . By taking finer and finer approximations of K by polytopes, it suffices to prove the result for a polyhedral set. We only need to argue how the d -dimensional volume of the facets is transformed by T . The change in volume of a facet corresponds to the determinant of the transformation induced by T on the tangent space of that facet¹. More precisely, given linearly independent vectors $v_1, \dots, v_d \in \mathbb{R}^{d+1}$, let P be the parallelepiped generated by them and let $V_d(P)$ be its volume. Let also n be the unit vector orthogonal to affine subspace containing P and N an unit segment in that direction, i.e., the set of points of the form tn for $t \in [0, 1]$, then:

$$V_{d+1}(T(P+N)) = (\det T) \cdot V_{d+1}(P+N) = (\det T) \cdot V_d(P)$$

where the first equality follows from how the (standard) volume transforms and the second since N is orthogonal to P and has size 1.

¹The tangent space of a facet is the space of all vectors that are parallel to that facet

Now, since $T(P+N) = T(P) + T(N)$, the volume $V_{d+1}(T(P+N))$ can be written as $V_{d+1}(T(P))$ times the projection of N in the orthogonal direction of $T(P)$, which is $\langle Tn, n' \rangle \leq \|Tn\| \cdot \|n'\| \leq 1$ where n' is the orthogonal vector to $T(P)$ and $\|Tn\| \leq 1$ follows from the fact that T is a contraction. Therefore:

$$V_d(T(P)) \geq V_{d+1}(T(P+N)) = (\det T) \cdot V_d(P).$$

■

2) *Intrinsic volumes of cylinders:* Given a convex set K in \mathbb{R}^d , an *orthogonal cylinder* with base K and height w is the convex set in \mathbb{R}^{d+1} formed by taking the Minkowski sum of K (embedded into \mathbb{R}^{d+1}) and a line segment of length w orthogonal to K .

Lemma 27. *Let K be a convex set in \mathbb{R}^d , and let S be an orthogonal cylinder with base K and height h . Then, for all $0 \leq j \leq d$,*

$$V_{j+1}(S) = V_{j+1}(K) + hV_j(K).$$

Proof: Embed K into \mathbb{R}^{d+1} so that it lies in the hyperplane $x_{d+1} = 0$, and let L be the line segment from 0 to he_{d+1} , so that $S = K + L$ is an orthogonal cylinder with base K and height h . We will compute $\text{Vol}_{d+1}(S + \varepsilon B_{d+1})$. Recall that Vol refers to the standard volume. Whenever we add subscripts (e.g. Vol_d) we do so to highlight that we are talking about the standard volume of a convex set in a d -dimensional (sub)space.

We claim we can decompose $S + \varepsilon B_{d+1}$ into two parts; one with total volume $\text{Vol}_{d+1}(K + \varepsilon B_{d+1})$, and one with total volume $h\text{Vol}_d(K + \varepsilon B_d)$. To begin, consider the intersection of $S + \varepsilon B_{d+1}$ with $\{x_{d+1} \in [0, h]\}$. We claim this set has volume at least $h\text{Vol}_d(K + \varepsilon B_d)$. In particular, note that (since S is an orthogonal cylinder) every cross-section of the form $(S + \varepsilon B_{d+1}) \cap \{x_{d+1} = t\}$ for $t \in [0, h]$ is congruent to the set $K + \varepsilon B_d$. It follows that the volume of this region is $h\text{Vol}_d(K + \varepsilon B_d)$.

Next, consider the intersection of $S + \varepsilon B_{d+1}$ with the set $\{x_{d+1} \notin [0, h]\}$. This intersection has two components: a component S^+ , the intersection of $S + \varepsilon B_{d+1}$ with the set $\{x_{d+1} \geq h\}$, and a component S^- , the intersection of $S + \varepsilon B_{d+1}$ with the set $\{x_{d+1} \leq 0\}$ (see Figure 8a). Now, define K^+ to be the intersection of $K + \varepsilon B_{d+1}$ with $\{x_{d+1} \geq 0\}$, and let K^- be the intersection of $K + \varepsilon B_{d+1}$ with $\{x_{d+1} \leq 0\}$. It is straightforward to verify that K^+ is congruent to S^+ and that K^- is congruent to S^- , and therefore the volume of this region is equal to $\text{Vol}(K^+) + \text{Vol}(K^-) = \text{Vol}_{d+1}(K + \varepsilon B_{d+1})$.

We therefore have that $\text{Vol}_{d+1}(S + \varepsilon B_{d+1}) = \text{Vol}_{d+1}(K + \varepsilon B_{d+1}) + h\text{Vol}_d(K + \varepsilon B_d)$. Expanding out all parts via Steiner's formula (1), we have that:

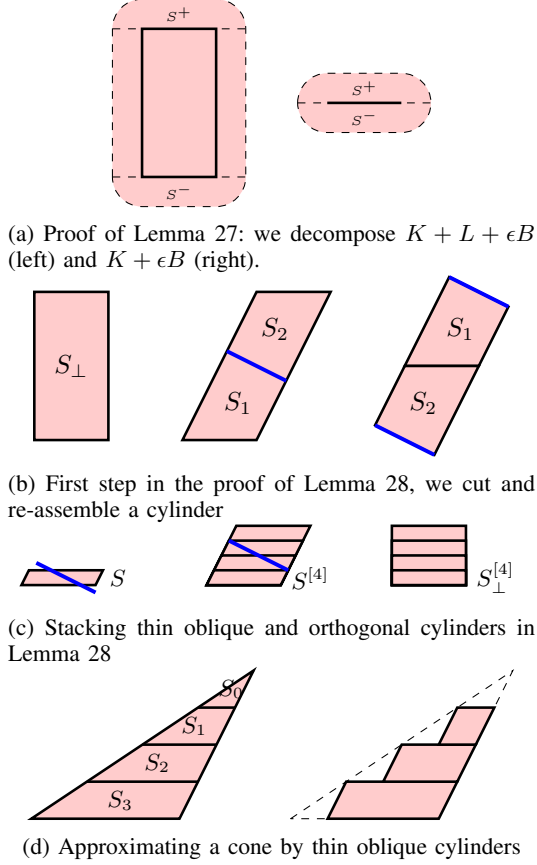


Figure 8: Illustration of the cylinder and cone proofs. In all cases, the x -axis is a d -dimensional space and the y -axis a 1-dimensional space

$$\begin{aligned}
& \sum_{j=0}^{d+1} \kappa_{d+1-j} V_j(S) \varepsilon^{d+1-j} \\
= & \sum_{j=0}^d \kappa_{d+1-j} V_j(K) \varepsilon^{d+1-j} + h \sum_{j=0}^d \kappa_{d-j} V_j(K) \varepsilon^{d-j}.
\end{aligned}$$

Equating coefficients of ε^{d-j} , we find that

$$V_{j+1}(S) = V_{j+1}(K) + hV_j(K).$$

An *oblique cylinder* in \mathbb{R}^{d+1} is formed by taking the Minkowski sum of a convex set $K \subset \mathbb{R}^d$ and a line segment L not necessarily perpendicular to K . The *height* of an oblique cylinder is equal to the length of the component of L orthogonal to the affine subspace containing K .

Lemma 28. *Let K be a convex set in \mathbb{R}^d . If S is an oblique cylinder with base K and height h , and S_\perp is an orthogonal cylinder with base K and height h , then (for all $1 \leq j \leq d+1$)*

$$V_j(S) \geq V_j(S_\perp).$$

Proof: Note that when $j = d+1$, $V_j(S) = V_j(S_\perp)$ as they are related by a linear transformation with determinant 1. For the remaining cases, we will first prove for $j = d$ and then reduce all other cases to $j = d$.

Case $\mathbf{j} = \mathbf{d}$ (tall cylinder): Write $S = K + L$, where L is a line segment of length ℓ (with orthogonal component h with respect to K). We will begin by choosing a hyperplane H perpendicular to L that intersects S along its lateral surface, dividing it into two sections S_1 and S_2 (see Figure 8b). Note that this is only possible if the height h of this cylinder is large enough with respect to the diameter of K and angle L makes with K . We address the case of the short cylinder in the next case. Let $K' = H \cap S$. By Theorem 7, we know that $V_d(S) = V_d(S_1) + V_d(S_2) - V_d(K')$.

Note that it is possible to reassemble S_1 and S_2 by gluing them along their copies of K to form an orthogonal cylinder with base K' and height ℓ . Call this cylinder S' . Again by Theorem 7, we have that $V_d(S') = V_d(S_1) + V_d(S_2) - V_d(K')$, and therefore $V_d(S) = V_d(S') + V_d(K) - V_d(K')$. But by Lemma 27, $V_d(S') = V_d(K') + \ell V_{d-1}(K')$, so $V_d(S) = V_d(K) + \ell V_{d-1}(K')$. On the other hand (also by Lemma 27), $V_d(S_\perp) = V_d(K) + hV_{d-1}(K)$. Therefore, to show that $V_d(S) \geq V_d(S_\perp)$, it suffices to show that $\ell V_{d-1}(K') \geq hV_{d-1}(K)$.

Now, note that K' is the projection of K onto the hyperplane H . The normal to H is parallel to L . Since L has length ℓ and orthogonal component h with respect to K , the angle between L and the normal to K equals $\arccos(h/\ell)$, from which it follows from Lemma 25 that K' is an (h/ℓ) stretch of K . By Lemma 26, it follows that $V_{j-1}(K') \geq (h/\ell)V_{j-1}(K)$, from which the desired inequality follows.

Case $\mathbf{j} = \mathbf{d}$ (short cylinder): Finally, what if the original cylinder was not tall enough to divide into two components in the desired manner? To deal with this, let $S^{[n]}$ denote n copies of S stacked on top of each other (i.e. $S^{[n]} = K + nL$), and let $S_\perp^{[n]}$ denote n copies of S_\perp stacked on top of each other (i.e. an orthogonal cylinder with base K and height nh). Repeatedly applying Theorem 7, we have that $V_d(S^{[n]}) = nV_d(S) + (n-1)V_d(K)$, and that $V_d(S_\perp^{[n]}) = nV_d(S_\perp) + (n-1)V_d(K)$. Therefore, to show that $V_d(S) \geq V_d(S_\perp)$, it suffices to show that $V_d(S^{[n]}) \geq V_d(S_\perp^{[n]})$. For some n , $S^{[n]}$ will be tall enough to divide as desired, which completes the proof.

Reducing $\mathbf{j} < \mathbf{d}$ to $\mathbf{j} = \mathbf{d}$: We can without loss of generality assume that K (the base of the cylinder) is in the plane spanned by the first d coordinate vectors. Also, let $\pi_d : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^d$ be the projection in the first d coordinates.

Recall that the j th intrinsic volume $V_j(S)$ is equal to the expected volume of the projection of S onto a randomly

chosen j -dimensional subspace of \mathbb{R}^{d+1} , where the distribution over subspaces is given by the Haar measure over $\text{Gr}(d+1, k)$ (see Theorem 12).

Therefore, choose H according to this measure and let $P = \pi_d(H)$. Note that (almost surely) P is an element of $\text{Gr}(d, j)$ and P is distributed according to the Haar measure of this Grassmannian. By the law of total expectation, we can write

$$\begin{aligned} V_j(S) &= \mathbb{E}_{H \sim \text{Gr}(d+1, j)} [V_j(\Pi_H S)] \\ &= \mathbb{E}_{P \sim \text{Gr}(d, j)} [\mathbb{E}_{H \sim \text{Gr}(d+1, j)} [V_j(\Pi_H S) \mid \pi_d(H) = P]] \end{aligned}$$

Let P' be the element of $\text{Gr}(d+1, j+1)$ spanned by P and e_{d+1} . Note that since $H \subset P'$, $\Pi_H S = \Pi_H \Pi_{P'} S$. We therefore claim that

$$\mathbb{E}_{H \sim \text{Gr}(d+1, j)} [V_j(\Pi_H S) \mid \pi_d(H) = P] = V_j(\Pi_{P'} S).$$

Indeed, conditioned on $\pi_d(H) = P$, H is a (Haar-)uniform subspace of dimension j of the $j+1$ -dimensional space P' , from which the above equality follows. Therefore, we have that

$$V_j(S) = \mathbb{E}_{P \sim \text{Gr}(d, j)} [V_j(\Pi_{P'} S)]$$

and similarly

$$V_j(S_\perp) = \mathbb{E}_{P \sim \text{Gr}(d, j)} [V_j(\Pi_{P'} S_\perp)].$$

Now, since e_{d+1} belongs to P' , if S_\perp is an orthogonal cylinder with base K and height h in \mathbb{R}^{d+1} , then $\Pi_{P'} S_\perp$ is an orthogonal cylinder with base $\Pi_{P'} K$ and height h in P' . Likewise, $\Pi_{P'} S$ is an oblique cylinder with base $\Pi_{P'} K$ and height h in P' . Since P' is $j+1$ dimensional, it follows the previous cases that: $V_j(\Pi_{P'} S) \geq V_j(\Pi_{P'} S_\perp)$ so

$$\begin{aligned} V_j(S) &= \mathbb{E}_{P \sim \text{Gr}(d, j)} [V_j(\Pi_{P'} S)] \\ &\geq \mathbb{E}_{P \sim \text{Gr}(d, j)} [V_j(\Pi_{P'} S_\perp)] \\ &\geq V_j(S_\perp) \end{aligned}$$

■

3) *Intrinsic volumes of cones:* A cone in \mathbb{R}^{d+1} is the convex hull of a d -dimensional convex set K and a point $p \in \mathbb{R}^{d+1}$. If the distance from p to the affine subspace containing K is h , we say the cone has *height* h and *base* K .

Lemma 29. *Let K be a convex set in \mathbb{R}^d , and let S be a cone in \mathbb{R}^{d+1} with base K and height h . Then, for all $0 \leq j \leq d$,*

$$V_{j+1}(S) \geq \frac{1}{j+1} h V_j(K).$$

Proof: Choose a positive integer n , and divide S into n parts via the hyperplanes $H_i = \{x_{d+1} = \frac{n-i}{n} h\}$ (for

$0 \leq i \leq n$). For $0 \leq i < n$, let K_i be the intersection of H_i with S , and let S_i be the region of S bounded between hyperplanes H_i and H_{i+1} (see Figure 8d). Note that each S_i is a frustum with bases K_i and K_{i+1} and height h/n , and furthermore that each K_i is congruent to $\frac{i}{n} K$.

By repeatedly applying Theorem 7, we know that

$$V_{j+1}(S) = \sum_{i=0}^{n-1} V_{j+1}(S_i) - \sum_{i=1}^{n-1} V_{j+1}(K_i).$$

Note that each set S_i contains an oblique cylinder with base K_i (since K_i is a contraction of K_{i+1} , some translate of K_i is strictly contained inside K_{i+1}) and height h/n . It follows from Lemmas 27 and 28 that $V_{j+1}(S_i) \geq V_{j+1}(K_i) + \frac{h}{n} V_j(K_i)$. It follows that

$$\begin{aligned} V_{j+1}(S) &\geq \sum_{i=1}^{n-1} \frac{h}{n} V_j(K_i) \\ &= \sum_{i=1}^{n-1} \frac{h}{n} V_j\left(\frac{i}{n} K\right) \\ &= \sum_{i=1}^{n-1} \frac{h}{n} \left(\frac{i}{n}\right)^j V_j(K) \\ &= \left(\sum_{i=1}^{n-1} \left(\frac{i}{n}\right)^j \frac{1}{n}\right) h V_j(K). \end{aligned}$$

As n goes to infinity, this sum approaches $\int_0^1 x^j dx = \frac{1}{j+1}$, and therefore we have that $V_{j+1}(S) \geq \frac{1}{j+1} h V_j(K)$. ■

D. Efficient implementation

We have thus far ignored issues of computational efficiency. In this subsection, we will show that algorithms `SymmetricSearch` (Figure 6) and `PricingSearch` (Figure 7) can be implemented in polynomial time by a randomized algorithm that succeeds with high probability.

The main primitive we require to implement both algorithms is a way to efficiently compute the intrinsic volumes of a convex set (and in particular a convex polytope, since our knowledge set starts as $[0, 1]^d$ and always remains a convex polytope). Unfortunately, even computing the ordinary volume of a convex polytope (presented as an intersection of half-spaces) is known to be $\#P$ -hard [4]. Fortunately, there exist efficient randomized algorithms to compute arbitrarily good multiplicative approximations of the volume of a convex set.

Theorem 30 (Dyer, Frieze, and Kannan [7]). *Let K be a convex subset of \mathbb{R}^d with an efficient membership oracle (which given a point, returns whether or not $x \in K$). Then there exists a randomized algorithm which, given input $\varepsilon > 0$, runs in time $\text{poly}(d, \frac{1}{\varepsilon})$ and outputs an ε -approximation to $\text{Vol}(K)$ with high probability.*

We will show how we can extend this to efficiently compute (approximately, with high probability) the intrinsic volumes of a convex polytope presented as an intersection of half-spaces.

Theorem 31. *Let K be a polytope in \mathbb{R}^d defined by the intersection of n half-spaces and contained in $[0, 1]^d$. Then there exists a randomized algorithm which, given input $\varepsilon > 0$ and $1 \leq i \leq d$, runs in time $\text{poly}(d, n, \frac{1}{\varepsilon})$, and outputs an ε -approximation to $V_i(K)$ with high probability.*

Proof: We use the fact (Theorem 12) that $V_i(K)$ is the expected volume of the projection of K onto a randomly chosen i -dimensional subspace (sampled according to the Haar measure). Since K is contained inside $[0, 1]^d$, any i -dimensional projection of K will be contained within an i -dimensional projection of $[0, 1]^d$, whose i -dimensional volume is at most $\text{poly}(d)$. By Hoeffding’s inequality, we can therefore obtain an ε -approximation to $V_i(K)$ by taking the average of $\text{poly}(d, \frac{1}{\varepsilon})$ ($\varepsilon/2$)-approximations for volumes of projections of K onto i -dimensional subspaces.

To approximately compute the volume of a projection of K onto an i -dimensional subspace S , we will apply Theorem 12. Note that we can check whether a point belongs in the projection of K into S by solving an LP (the point adds i additional linear constraints to the constraints defining K). This can be done efficiently in polynomial time, and therefore we have a polynomial-time membership oracle for this subproblem. ■

We now briefly argue that Theorem 31 allows us to implement efficient randomized variants of SymmetricSearch and PricingSearch which succeed with high probability. To do this, it suffices to note that all of the analysis of both algorithms is robust to tiny perturbations in computations of intrinsic volumes. For example, in SymmetricSearch the analysis carries through even if instead of K_i dividing S_t into two regions such that $V_i(S^+) = V_i(S^-)$, it divides them into regions satisfying $V_i(S^+) \in [(1 - \varepsilon)V_i(S^-), (1 + \varepsilon)V_i(S^-)]$ for some constant ε .

The only remaining implementation detail is how to use hyperplanes K_i that divide the i th intrinsic volume of S_i equally (or in the case of PricingSearch, divide off a fixed amount of intrinsic volume). Since intrinsic volumes are monotone (Theorem 7), this can be accomplished via binary search.

REFERENCES

- [1] Alekh Agarwal, Daniel J. Hsu, Satyen Kale, John Langford, Lihong Li, and Robert E. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 1638–1646, 2014.
- [2] Dennis Amelunxen, Martin Lotz, Michael B McCoy, and Joel A Tropp. Living on the edge: Phase transitions in convex programs with random data. *Information and Inference: A Journal of the IMA*, 3(3):224–294, 2014.
- [3] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 622–630, 2014.
- [4] Imre Bárány and Zoltán Füredi. Computing the volume is difficult. *Discrete & Computational Geometry*, 2(4):319–326, 1987.
- [5] Hamsa Bastani and Mohsen Bayati. Online decision-making with high-dimensional covariates. *Working paper, Stanford University*, 2016.
- [6] Maxime C. Cohen, Ilan Lobel, and Renato Paes Leme. Feature-based dynamic pricing. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC ’16, Maastricht, The Netherlands, July 24-28, 2016*, page 817, 2016.
- [7] Martin Dyer, Alan Frieze, and Ravi Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *Journal of the ACM (JACM)*, 38(1):1–17, 1991.
- [8] Stephen Gillen, Christopher Jung, Michael Kearns, and Aaron Roth. Online learning with an unknown fairness metric. *arXiv preprint arXiv:1802.06936*, 2018.
- [9] Adel Javanmard. Perishability of data: dynamic pricing under varying-coefficient models. *The Journal of Machine Learning Research*, 18(1):1714–1744, 2017.
- [10] Adel Javanmard and Hamid Nazerzadeh. Dynamic pricing in high-dimensions. *Working paper, University of Southern California*, 2016.
- [11] Richard Karp, Elias Koutsoupias, Christos Papadimitriou, and Scott Shenker. Optimization problems in congestion control. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 66–74. IEEE, 2000.
- [12] Daniel A Klain and Gian-Carlo Rota. *Introduction to geometric probability*. Cambridge University Press, 1997.
- [13] Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 594–605. IEEE, 2003.
- [14] Ilan Lobel, Renato Paes Leme, and Adrian Vladu. Multidimensional binary search for contextual decision-making. *Operations Research*, 2017.
- [15] Peter McMullen. Inequalities between intrinsic volumes. *Monatshefte für Mathematik*, 111(1):47–53, 1991.
- [16] Sheng Qiang and Mohsen Bayati. Dynamic pricing with demand covariates. *Available at SSRN 2765257*, 2016.
- [17] Stephen H Schanuel. What is the length of a potato? *Lecture Notes in Mathematics, Springer*, pages 118–126, 1986.