

IMPROVED REGRET FOR ZEROth-ORDER ADVERSARIAL BANDIT CONVEX OPTIMISATION

Tor Lattimore

DeepMind, London
lattimore@google.com

Abstract We prove that the information-theoretic upper bound on the minimax regret for zeroth-order adversarial bandit convex optimisation is at most $O(d^{2.5}\sqrt{n}\log(n))$, where d is the dimension and n is the number of interactions. This improves on the bound of $O(d^{9.5}\sqrt{n}\log(n)^{7.5})$ by Bubeck et al. (2017). The proof is based on identifying an improved exploratory distribution for convex functions.

Mathematics Subject Classification (2010). 52A20; 62C05.

Keywords. Bandit convex optimisation, online learning.

1 Introduction

Let $\mathcal{K} \subset \mathbb{R}^d$ be a convex body (convex, compact with non-empty interior) and \mathcal{G} be a set of convex functions from \mathcal{K} to $[0, 1]$. At the start of the game, an adversary secretly chooses a sequence $(f_t)_{t=1}^n$ with $f_t \in \mathcal{G}$. Then, in each round t , the learner chooses an action $x_t \in \mathcal{K}$, possibly with randomisation, and observes only the loss $f_t(x_t)$. The minimax regret over n rounds is

$$\mathfrak{R}_n^*(\mathcal{G}) = \inf_{\text{policies}} \sup_{(f_t)_{t=1}^n \in \mathcal{G}^n} \max_{x \in \mathcal{K}} \mathbb{E} \left[\sum_{t=1}^n f_t(x_t) - f_t(x) \right], \quad (1)$$

where the inf is over all policies of the learner that determine the actions $(x_t)_{t=1}^n$. The expectation integrates over the randomness of the actions $(x_t)_{t=1}^n$. Our contribution is a proof of the following theorem.

Theorem 1. *Suppose that \mathcal{K} contains a unit-radius Euclidean ball and \mathcal{G} is the set of all convex functions from \mathcal{K} to $[0, 1]$. Then*

$$\mathfrak{R}_n^*(\mathcal{G}) \leq \text{const} \cdot d^{2.5} \sqrt{n} \log(n \text{diam}(\mathcal{K})),$$

where const is a universal constant and $\text{diam}(\mathcal{K}) = \max_{x,y \in \mathcal{K}} |x-y|$ is the diameter of \mathcal{K} with $|\cdot|$ the standard Euclidean norm.

The assumption that \mathcal{K} contains a unit-radius Euclidean ball is relatively mild. Any convex body can be rescaled to contain a unit ball, while the regret depends only logarithmically on the diameter. As in previous work, we make use of a simple reduction that allows us to restrict slightly the class of functions available to the adversary [2, 4]. Define a constant $m = 1/((n+1)\text{diam}(\mathcal{K})^2)$ and let \mathcal{F} be the space of all convex functions \mathcal{K} to $[0, 1]$ that are

- (a) n -Lipschitz: $f(x) - f(y) \leq n|x - y|$ for all $x, y \in \mathcal{K}$; and
- (b) m -strongly convex: for all $x, y \in \mathcal{K}$ and $\lambda \in [0, 1]$,

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y) - \frac{1}{2}m\lambda(1 - \lambda)|x - y|^2.$$

Proposition 16 in Section 5 shows that it suffices to prove Theorem 1 with \mathcal{F} rather than \mathcal{G} . Briefly, the reduction works by showing that bounded convex functions from \mathcal{K} to $[0, 1]$ cannot have large directional derivatives except close to the boundary $\partial\mathcal{K}$ and must have near-minimisers that are not too close to the boundary. This means the learner can play on a subset of \mathcal{K} on which the loss functions are n -Lipschitz without sacrificing much in terms of the regret. Strong convexity is introduced by adding a small quadratic to all loss functions, which has only a negligible impact because m is small.

The next (more or less known) theorem serves as our starting point. It follows from the machinery developed by [Bubeck et al. \[3\]](#), [Bubeck and Eldan \[2\]](#) and [Russo and Van Roy \[14\]](#). For completeness, an outline of the proof is given in Appendix A.

Theorem 2. *Let $\alpha, \beta \in \mathbb{R}$ be non-negative and for $f \in \mathcal{F}$, let $f_\star = \min_{x \in \mathcal{K}} f(x)$. Suppose that for any $\bar{f} \in \mathcal{F}$ and finitely supported distribution μ on \mathcal{F} with the discrete σ -algebra there exists a probability measure ρ on \mathcal{K} such that*

$$\int_{\mathcal{K}} \bar{f}(x) d\rho(x) - \int_{\mathcal{F}} f_\star d\mu(f) \leq \alpha + \sqrt{\beta \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f}(x) - f(x))^2 d\rho(x) d\mu(f)}. \quad (2)$$

Then the minimax regret is bounded by

$$\mathfrak{R}_n^\star(\mathcal{F}) \leq 3 + n\alpha + \sqrt{\beta dn \log(3n^2 \max(1, (2\beta)^{1/2}) \text{diam}(\mathcal{K}))}.$$

The distribution ρ in Theorem 2 is called an exploratory distribution. [Bubeck and Eldan \[2\]](#) established the conditions of Theorem 2 with $\alpha = 1/\sqrt{n}$ and $\beta = d^{21} \text{polylog}(n)$. The next theorem improves on this result.

Theorem 3. *For any $\bar{f} \in \mathcal{F}$ and finitely supported distribution μ over \mathcal{F} , there exists a probability measure ρ on \mathcal{K} such that Eq. (2) holds with $\alpha = 1/n$ and $\beta = \text{const} \cdot d^4 \log(nd/m)$, where const is a universal constant.*

The measure realising Theorem 3 is a mixture over probability measures on the boundaries of level sets of \bar{f} . Combining Theorems 2 and 3 with the reduction in Proposition 16 proves Theorem 1.

Related work Online convex optimisation is usually studied under the assumption that the learner has access to the (sub-)gradient $\nabla f_t(X_t)$, or even the whole function f_t . A number of perspectives on this vast literature can be found in recent books and notes by [Cesa-Bianchi and Lugosi](#) [5], [Hazan](#) [8] and [Orabona](#) [13]. There is far less work in the zeroth-order bandit setting in which the learner only has access to point evaluations. A natural idea is to use importance-weighting to estimate the gradients. At least with current tools, however, the resulting bias/variance tradeoff leads to suboptimal regret [11].

The function class \mathcal{F} is omitted from the following regret bounds, to emphasise that the assumptions vary in minor ways. Constant factors are also omitted, while Big-O is used to hide dependence on other parameters such as smoothness/dimension/strong convexity. Information-theoretic means were used by [Bubeck et al.](#) [3] to show that the minimax regret is $\mathfrak{R}_n^* \leq \log(n)\sqrt{n}$ when $d = 1$. The multi-dimensional problem was considered by [Hazan and Li](#) [10], who showed that the minimax regret is $O(n^{1/2} \text{polylog}(n))$, but with an exponential dependence on the dimension. Shortly after, [Bubeck and Eldan](#) [2] generalised the information-theoretic machinery to prove that $\mathfrak{R}_n^* \leq d^{11} \log(n)^4 \sqrt{n}$, breaking the exponential dependence on the dimension while retaining square root dependence on the horizon. None of these works provide an efficient algorithm. More recently, [Bubeck et al.](#) [4] used kernel-based estimators and tools from online convex optimisation to show that $\mathfrak{R}_n^* \leq d^{9.5} \log(n)^{7.5} \sqrt{n}$. Furthermore, their algorithm can be implemented in polynomial time (with reasonable assumptions) with the price that the dimension-dependence increases to $d^{10.5}$. They conjectured that the optimal regret is $\mathfrak{R}_n^* \leq d^{1.5} \sqrt{n} \text{polylog}(n)$. The best known lower bound is $\mathfrak{R}_n^* \geq d\sqrt{n}$, which holds even when the function class is restricted to linear functions [6]. There is also a line of work that exploits strong convexity to obtain better bounds. In particular, if one is prepared to make an additional assumption on smoothness, then [Hazan and Levy](#) [9] proved that $\mathfrak{R}_n^* = O(d^{1.5} \sqrt{(n/m) \log(n)})$. The polynomial dependence on m means that enforcing strong convexity by adding a small quadratic to the losses blows up the bound and leads to suboptimal rates.

Preliminaries The d -dimensional sphere is $S^d = \{x \in \mathbb{R}^{d+1} : |x| = 1\}$. The Minkowski sum of sets A and B is denoted $A + B$. When x is a vector, $A + x = A + \{x\}$. Let vol_p denote the p -dimensional Hausdorff measure on \mathbb{R}^d , normalised to coincide with the Lebesgue measure. Given $x, y \in \mathbb{R}^d$, $[x, y] = \{tx + (1-t)y : t \in [0, 1]\}$ is the chord connecting x and y . The sets (x, y) and $(x, y]$ and (x, y) are defined similarly but with appropriate end-points removed. When $x \neq \mathbf{0}$, the hyperplane with normal x is $x^\perp = \{y \in \mathbb{R}^d : \langle x, y \rangle = 0\}$ and $P_x(y) = \arg \min_{z \in x^\perp} |y - z|$ is the Euclidean projection of y onto x^\perp . The shadow of a convex body $K \subset \mathbb{R}^d$ in direction x is $P_x(K) = \{P_x(y) : y \in K\}$. The infimum of a convex function $f : K \rightarrow \mathbb{R}$ is $f_\star = \inf_{x \in K} f(x)$. The boundary of K is ∂K .

2 Combining exploratory distributions

The plan is to establish the conditions of Theorem 2 for suitable values of α and β , which means that for any $\bar{f} \in \mathcal{F}$ and distribution μ on \mathcal{F} we need to find a probability measure ρ on \mathcal{K} satisfying Eq. (2). To make the problem more manageable, we first prove that exploratory distributions can be combined.

Lemma 4. *Let $\bar{f} \in \mathcal{F}$ and $\mathcal{F} = \cup_{i=1}^k \mathcal{F}_i$. Assume there exist probability measures $(\rho_i)_{i=1}^k$ on \mathcal{K} such that for all $i \in \{1, \dots, k\}$,*

$$\int_{\mathcal{K}} \bar{f}(x) d\rho_i(x) - f_{\star} \leq \alpha + \sqrt{\beta \int_{\mathcal{K}} (\bar{f}(x) - f(x))^2 d\rho_i(x)} \quad \text{for all } f \in \mathcal{F}_i. \quad (3)$$

Then there exists a probability measure ρ on \mathcal{K} such that

$$\int_{\mathcal{K}} \bar{f}(x) d\rho(x) - \int_{\mathcal{F}} f_{\star} d\mu(f) \leq \alpha + \sqrt{\beta k \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f}(x) - f(x))^2 d\rho(x) d\mu(f)}.$$

Proof. The argument is algebraically identical to that used by [Russo and Van Roy \[15\]](#) to bound the information ratio for Thompson sampling. Assume without loss of generality that $(\mathcal{F}_i)_{i=1}^k$ are disjoint and $\mu(\mathcal{F}_i) > 0$ for all $i \in \{1, \dots, k\}$. Let μ_i be the probability measure obtained by conditioning μ on \mathcal{F}_i : $\mu_i(A) = \mu(A \cap \mathcal{F}_i) / \mu(\mathcal{F}_i)$. Then, letting $q_i = \mu(\mathcal{F}_i)$ and $\rho = \sum_{i=1}^k q_i \rho_i$,

$$\begin{aligned} \int_{\mathcal{K}} \bar{f}(x) d\rho(x) - \int_{\mathcal{F}} f_{\star} d\mu(f) &= \sum_{i=1}^k q_i \left(\int_{\mathcal{K}} \bar{f}(x) d\rho_i(x) - \int_{\mathcal{F}} f_{\star} d\mu_i(f) \right) \\ &\leq \alpha + \sum_{i=1}^k q_i \sqrt{\beta \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f}(x) - f(x))^2 d\rho_i(x) d\mu_i(f)} \\ &\leq \alpha + \sqrt{\beta k \sum_{i=1}^k q_i^2 \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f}(x) - f(x))^2 d\rho_i(x) d\mu_i(f)} \\ &\leq \alpha + \sqrt{\beta k \sum_{i,j=1}^k q_i q_j \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f}(x) - f(x))^2 d\rho_i(x) d\mu_j(f)} \\ &= \alpha + \sqrt{\beta k \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f}(x) - f(x))^2 d\rho(x) d\mu(f)}, \end{aligned}$$

where the first inequality follows from the assumption in Eq. (3) and Jensen's inequality and the second follows from Cauchy-Schwarz. \square

Before moving on, notice that for functions $f \in \mathcal{F}$ with $f_{\star} \geq \bar{f}_{\star}$, the Dirac distribution at the minimiser of \bar{f} is a good exploratory distribution in the sense that Eq. (3) holds with $\alpha = \beta = 0$.

3 Proof of Theorem 3

The first lemma shows that convex functions with different minimisers must differ along suitably chosen rays. The situation is illustrated in Fig. 1.

Lemma 5. *Let $\mathcal{D} \subset \mathbb{R}^d$ be a convex body and $f, g : \mathcal{D} \rightarrow \mathbb{R}$ be convex functions with f minimised at $x_\star \in \mathcal{D}$ and $f_\star \leq g_\star$. Suppose that $\epsilon > 0$ and $x_\star \notin K = \{x \in \mathcal{D} : g(x) \leq g_\star + \epsilon\}$. Assume that $x, y \in \partial K$ and $x = \psi x_\star + (1 - \psi)y$ with $\psi \in (0, 1)$. Then,*

$$(f(x) - g(x))^2 + (f(y) - g(y))^2 \geq \frac{1}{2}\psi^2 (g_\star + \epsilon - f_\star)^2.$$

Proof. The definition of the level set and continuity of g means that $g(x) = g_\star + \epsilon$. Furthermore, $g(y) \leq g_\star + \epsilon$, with a strict inequality only possible if y is on the boundary of \mathcal{D} . Suppose that $\psi f_\star + (1 - \psi)f(y) \geq g_\star + \epsilon$. Then

$$(f(y) - g(y))^2 \geq \left(\frac{\psi}{1 - \psi}\right)^2 (g_\star + \epsilon - f_\star)^2 \geq \frac{1}{2}\psi^2 (g_\star + \epsilon - f_\star)^2.$$

Otherwise, using convexity of f ,

$$\begin{aligned} (g(y) - f(y))^2 + (g(x) - f(x))^2 &= (g(y) - f(y))^2 + (g_\star + \epsilon - f(x))^2 \\ &\geq (g(y) - f(y))^2 + (g_\star + \epsilon - \psi f_\star - (1 - \psi)f(y))^2 \\ &\geq \min_{a \in \mathbb{R}} [(g(y) - a)^2 + (g_\star + \epsilon - \psi f_\star - (1 - \psi)a)^2] \\ &= \frac{(\psi(g_\star + \epsilon - f_\star) + (1 - \psi)(g_\star + \epsilon - g(y)))^2}{2 - 2\psi + \psi^2} \\ &\geq \frac{1}{2}\psi^2 (g_\star + \epsilon - f_\star)^2, \end{aligned}$$

where the first inequality follows from the convexity of f and the last is naive bounding. \square




Figure 1: In the figure on the left, the horizontal grey line marks $g_\star + \epsilon = g(x)$. The proof of Lemma 5 shows that at least one of the red vertical lines is roughly ψ times the length of the black line. The figure on the right shows the top down view.

By letting $g = \bar{f}$ and otherwise using the notation in Lemma 5, the lemma shows that when ρ is the uniform distribution on x, y , then

$$\int_{\mathcal{K}} \bar{f} \, d\rho - f_{\star} \leq \bar{f}_{\star} + \epsilon - f_{\star} \leq \frac{2}{\psi} \sqrt{\int_{\mathcal{K}} (f - \bar{f})^2 \, d\rho}.$$

In other words, sampling uniformly from the entry/exit points of a ray emanating from the minimiser of f through a level set of \bar{f} is a good exploratory distribution provided that ψ , which is the depth of cut relative to the length of the chord, is reasonably large. Of course the ray depends on f , so the exploratory distribution constructed in this fashion is likely only good for a few functions $f \in \mathcal{F}$.

The next lemma generalises this argument by showing that sampling uniformly from the boundary of a level set is a good exploratory distribution for functions $f \in \mathcal{F}$ minimised at x_{\star} provided two conditions hold. First, the level set should have large shadows relative to its surface area. Second, the ‘average’ depth/length ratio of rays emanating from x_{\star} and passing through the level set should be approximately $1/d$. Lemma 4 will then be used to combine exploratory distributions produced from different level sets. Level sets with small shadows are handled by transforming coordinates so that the level set is in minimal surface area position (definition in Remark 7), and then pulling back the uniform distribution on the transformed level set.

Before the statement, we define a measure of the average depth/length ratio for rays emanating from a point and passing through a convex body. The concepts are illustrated in Fig. 2. Let $K \subset \mathbb{R}^d$ be a convex body with $\mathbf{0} \in K$ and $x \notin K$. Let $\pi_K(x, \cdot) : P_x(K) \rightarrow \partial K$ be an ‘inverse’ of the projection P_x defined by

$$\pi_K(x, z) = \arg \min_{y \in K \cap \{z + tx : t \in \mathbb{R}\}} \langle y, x \rangle.$$

For $x \in \mathbb{R}^d$ and $z \in P_x(K)$, define the depth/distance ratio by

$$\Psi_K(x, z) = \frac{\text{vol}_1(K \cap [x, \pi_K(x, z)])}{\text{vol}_1([x, \pi_K(x, z)])}.$$

The average depth/distance ratio with respect to the uniform probability measure on $P_x(K)$ is

$$\Psi_K(x) = \frac{1}{\text{vol}_{d-1}(P_x(K))} \int_{P_x(K)} \Psi_K(x, z) \, d\text{vol}_{d-1}(z). \quad (4)$$

The maximum depth/distance ratio is $\Psi_K^{\infty}(x) = \max_{z \in P_x(K)} \Psi_K(x, z)$. A number of properties of Ψ are collected in Section 4. The most important are that $z \mapsto \Psi_K(x, z)$ is concave on its domain $P_x(K)$ and $\Psi_{TK}(Tx) = \Psi_K(x)$ for any linear bijection $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$.




Figure 2: For the configuration in the figure, $\pi_K(x, z) = y$ and $\Psi_K(x, z) = |y - w|/|y - x|$, which is the ratio of the depth of K along the chord $[x, y]$ and the length of the chord. The quantity $\Psi_K(x)$ averages $\Psi_K(x, z)$ with respect to the uniform probability measure over all z in the shadow $P_x(K)$.

Lemma 6. Let $\mathcal{D} \subset \mathbb{R}^d$ be a convex body and $f, g : \mathcal{D} \rightarrow \mathbb{R}$ be convex with $f_* \leq g_*$, $\epsilon > 0$ and $K = \{x \in \mathcal{D} : g(x) \leq g_* + \epsilon\}$. Let $x_* \in \mathcal{D}$ be the minimiser of f and assume that $\mathbf{0} \in K$ and $\Psi_K(x_*) \in [1/(128d), 1/(32d)]$. Then,

$$\int_{\partial K} g \, d\rho - f_* \leq 2^{13} d \sqrt{\max_{\theta \in S^{d-1}} \left(\frac{\text{vol}_{d-1}(\partial K)}{\text{vol}_{d-1}(P_\theta(K))} \right) \int_{\partial K} (f - g)^2 \, d\rho}.$$

where $\rho = \text{vol}_{d-1} / \text{vol}_{d-1}(\partial K)$ is the surface area probability measure on ∂K .

Proof of Lemma 6. Abbreviate $\Psi(z) = \Psi_K(x_*, z)$ and $\pi(z) = \pi_K(x_*, z)$ and $P = P_{x_*}$. The rest of the proof is divided into three steps. The first uses Lemma 15 to show that a large fraction of rays through K from x_* are reasonably deep. The second step connects surface integrals over ∂K to rays cutting K from x_* and the third puts together the pieces using Lemma 5.

Step 1: Preponderance of deep cuts Let $B \subset P(K)$ be the set given by

$$B = \left\{ z \in P(K) : \frac{\Psi(z)}{\frac{1}{\text{vol}_{d-1}(P(K))} \int_{P(K)} \Psi \, d\text{vol}_{d-1}} \in [1/4, 16] \right\},$$

which is the set of points in $z \in P(K)$ where $\Psi(z)$ is close to its mean $\Psi_K(x_*)$. By Lemma 15 and concavity of Ψ (Lemma 8),

$$\frac{\text{vol}_{d-1}(B)}{\text{vol}_{d-1}(P(K))} \geq \frac{1}{32}. \quad (5)$$

Furthermore, by the assumptions in the lemma,

$$\Psi_K(x_*) = \frac{1}{\text{vol}_{d-1}(P(K))} \int_{P(K)} \Psi \, d\text{vol}_{d-1} \in \left[\frac{1}{128d}, \frac{1}{32d} \right],$$

which implies that for $z \in B$,

$$\Psi(z) \in \left[\frac{2^{-9}}{d}, \frac{1}{2d} \right]. \quad (6)$$

Step 2: Surface area to rays Let $C = \pi(B) \subset \partial K$. Define a function $\kappa : C \rightarrow \partial K$ so that $[x_*, y] \cap K = [\kappa(y), y]$, which is the point w in Fig. 2 and is chosen so that for any $z \in P(K)$,

$$\kappa(\pi(z)) = \Psi(z)x_* + (1 - \Psi(z))\pi(z). \quad (7)$$

Let $D = \kappa(C)$. The goal in this step is to establish Eq. (8) below, which makes the connection between rays and surface area. Let $\varphi : \mathcal{D} \rightarrow [0, \infty)$ be measurable. Then the following holds:

$$\begin{aligned} \int_C \varphi \, d\text{vol}_{d-1} &\geq \int_B \varphi \circ \pi \, d\text{vol}_{d-1} \\ \int_D \varphi \, d\text{vol}_{d-1} &\geq \frac{1}{4} \int_B \varphi \circ \kappa \circ \pi \, d\text{vol}_{d-1}. \end{aligned} \quad (8)$$

The first inequality is true because the projection of ∂K onto $P(K)$ only decreases surface area. Define $\lambda : P(K) \rightarrow \mathbb{R}$ and $\Lambda : P(K) \rightarrow P(K)$ by $\lambda(z) = 1 - \Psi(z)$ and $\Lambda(z) = \lambda(z)z$. Note, $\lambda(z) \in (0, 1)$ ensures that $\Lambda(z) \in P(K)$ for all $z \in P(K)$. Also, concavity of Ψ implies convexity of λ . For the remainder of this step, choose coordinates on $P(K) \subset x_*^\perp \cong \mathbb{R}^{d-1}$ via the obvious isometry to \mathbb{R}^{d-1} , which means that gradients and similar are written as $(d-1)$ -dimensional vectors. For the second inequality in Eq. (8), differentiating Λ at $z \in B$ yields $D\Lambda(z) = \lambda(z)\text{Id} + z\nabla\lambda(z)^\top$, which exists vol_{d-1} -a.e. by convexity of λ . Recalling the matrix determinant lemma: $\det(A + uv^\top) = \det(A)(1 + v^\top A^{-1}u)$, we have

$$\det(D\Lambda(z)) = \lambda(z)^{d-1} \left(1 + \frac{\langle \nabla\lambda(z), z \rangle}{\lambda(z)} \right) \geq \lambda(z)^{d-1} \left(2 - \frac{\lambda(\mathbf{0})}{\lambda(z)} \right) \geq \frac{1}{4}. \quad (9)$$

The first inequality follows from convexity of λ , so that $\lambda(\mathbf{0}) \geq \lambda(z) - \langle \nabla\lambda(z), z \rangle$. The last inequality follows from naive bounding because for $z \in B$, $\lambda(z) \geq 1 - 1/(2d)$ and $\lambda(\mathbf{0}) \leq 1$. The claim follows since, by the definition of Λ and Eq. (7), $P \circ \kappa \circ \pi = \Lambda$, which implies that

$$\int_D \varphi \, d\text{vol}_{d-1} \geq \int_{\Lambda(B)} \varphi \circ \kappa \circ \pi \circ \Lambda^{-1} \, d\text{vol}_{d-1} \geq \frac{1}{4} \int_B \varphi \circ \kappa \circ \pi \, d\text{vol}_{d-1},$$

where in the first inequality we used the facts that $\Lambda(B) = P(D)$ and projections decrease surface area. The second inequality follows from Eq. (9).

Step 3: Combining Let $z \in B$ and $y = \pi(z) \in \partial K$ and $w = \kappa(y) = \Psi(z)x_* + (1 - \Psi(z))y \in \partial K$. Hence, by Eq. (6) and Lemma 5, the following holds on B ,

$$(f - g)^2 \circ \pi + (f - g)^2 \circ \kappa \circ \pi \geq \frac{1}{2}(g_* + \epsilon - f_*)^2 \Psi^2 \geq \frac{(g_* + \epsilon - f_*)^2}{2^{19}d^2}. \quad (10)$$

Therefore, by Eq. (5), Eq. (8) and Eq. (10), and recalling the definition of ρ in the lemma statement,

$$\begin{aligned}
\int_{\partial K} (f - g)^2 d\rho &\geq \frac{1}{\text{vol}_{d-1}(\partial K)} \int_{C \cup D} (f - g)^2 d\text{vol}_{d-1} \\
&\geq \frac{1}{4 \text{vol}_{d-1}(\partial K)} \int_B ((f - g)^2 \circ \pi + (f - g)^2 \circ \kappa \circ \pi) d\text{vol}_{d-1} \\
&\geq \frac{(g_\star + \epsilon - f_\star)^2}{2^{21} d^2} \left(\frac{\text{vol}_{d-1}(B)}{\text{vol}_{d-1}(\partial K)} \right) \\
&\geq \frac{(g_\star + \epsilon - f_\star)^2}{2^{26} d^2} \min_{\theta \in S^{d-1}} \left(\frac{\text{vol}_{d-1}(P_\theta(K))}{\text{vol}_{d-1}(\partial K)} \right) \\
&\geq \frac{1}{2^{26} d^2} \left(\int_{\mathcal{K}} g d\rho - f_\star \right)^2 \min_{\theta \in S^{d-1}} \left(\frac{\text{vol}_{d-1}(P_\theta(K))}{\text{vol}_{d-1}(\partial K)} \right),
\end{aligned}$$

where the final inequality follows since $f_\star \leq g_\star$ by assumption and $g \leq g_\star + \epsilon$ on ∂K . Rearranging yields the result. \square

Remark 7. A convex body $K \subset \mathbb{R}^d$ is in minimal surface area position if its surface area measure (as a measure on the sphere) is isotropic [1, §2.3]. Giannopoulos and Papadimitrakis [7] show that for convex bodies K in minimal surface area position,

$$\max_{\theta \in S^{d-1}} \left(\frac{\text{vol}_{d-1}(\partial K)}{\text{vol}_{d-1}(P_\theta(K))} \right) \leq 2d,$$

which is sharp when K is a cube. Furthermore, for any convex body K , there exists a linear bijection $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that TK is in minimal surface area position.

Proof of Theorem 3. The proof is broken into three parts. First we construct the basic exploratory distribution using Lemma 6. In the second step we argue that only $O(d)$ exploratory distributions are needed. The final step puts together the pieces using Lemma 4. By means of a translation, and without loss of generality, choose coordinates on \mathcal{K} such that $\mathbf{0}$ is the minimiser of \bar{f} .

Step 1: Constructing exploratory distributions Let $\epsilon > 0$ and define the level set $K_\epsilon = \{x : \bar{f}(x) \leq \bar{f}_\star + \epsilon\}$. Let \mathcal{F}_ϵ be the set of all $f \in \mathcal{F}$ for which $f_\star \leq \bar{f}_\star$ and

$$\Psi_{K_\epsilon}(\arg \min_{x \in \mathcal{K}} f(x)) \in \left[\frac{1}{128d}, \frac{1}{32d} \right].$$

Let $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a linear bijection such that TK_ϵ is in minimal surface area position and define a probability measure ρ_ϵ on ∂K by

$$\rho_\epsilon = \frac{\text{vol}_{d-1} \circ T}{\text{vol}_{d-1}(\partial(TK_\epsilon))},$$

which is the pullback of the normalised surface area measure on $\partial(TK_\epsilon)$. That is, for any measurable $\varphi : \mathcal{K} \rightarrow \mathbb{R}$,

$$\int_{\partial K_\epsilon} \varphi \, d\rho_\epsilon = \frac{1}{\text{vol}_{d-1}(\partial(TK_\epsilon))} \int_{\partial(TK_\epsilon)} \varphi \circ T^{-1} \, d\text{vol}_{d-1}.$$

Let $f \in \mathcal{F}_\epsilon$ with minimiser $x_\star \in \mathcal{K}$. Since T is a linear bijection, both $f \circ T^{-1}$ and $\bar{f} \circ T^{-1}$ are convex functions from $\mathcal{D} = TK \rightarrow [0, 1]$. By Lemma 9 in Section 4,

$$\Psi_{TK_\epsilon}(Tx_\star) = \Psi_{K_\epsilon}(x_\star) \in \left[\frac{1}{128d}, \frac{1}{32d} \right].$$

Hence, by Lemma 6 and Remark 7, for any $f \in \mathcal{F}_\epsilon$,

$$\begin{aligned} \int_{\partial K_\epsilon} \bar{f} \, d\rho_\epsilon - f_\star &= \frac{1}{\text{vol}_{d-1}(\partial(TK_\epsilon))} \int_{\partial(TK_\epsilon)} \bar{f} \circ T^{-1} \, d\text{vol}_{d-1} - f_\star \\ &\leq 2^{13} \sqrt{\frac{2d^3}{\text{vol}_{d-1}(\partial(TK_\epsilon))} \int_{\partial(TK_\epsilon)} (\bar{f} - f)^2 \circ T^{-1} \, d\text{vol}_{d-1}} \\ &= 2^{13} \sqrt{2d^3 \int_{\partial K_\epsilon} (\bar{f} - f)^2 \, d\rho_\epsilon}, \end{aligned}$$

which shows that ρ_ϵ satisfies Eq. (3) for all $f \in \mathcal{F}_\epsilon$.

Step 2: Constructing a partition We start by defining a subset $\mathcal{F}_0 \subset \mathcal{F}$ on which the Dirac at $\mathbf{0}$ is a good exploratory distribution. Let

$$\mathcal{F}_0 = \{f \in \mathcal{F} : \bar{f}_\star - f_\star \leq 2(\bar{f}_\star - f(\mathbf{0}))\} \cup \{f \in \mathcal{F} : f_\star \geq \bar{f}_\star - 1/n\}.$$

Let ρ_0 be a Dirac at $\mathbf{0}$. Then,

$$\int_{\mathcal{K}} \bar{f} \, d\rho_0 - f_\star = \bar{f}_\star - f_\star \leq \frac{1}{n} + 2|\bar{f}_\star - f(\mathbf{0})| = \frac{1}{n} + \sqrt{4 \int_{\mathcal{K}} (\bar{f} - f)^2 \, d\rho_0},$$

which confirms the claim that ρ_0 is a good exploratory distribution for \mathcal{F}_0 . On the other hand, if $f \notin \mathcal{F}_0$ with minimiser $x_\star \in \mathcal{K}$, then by the definition of \mathcal{F}_0 and the assumption that f is n -Lipschitz,

$$\frac{1}{2n} \leq f(\mathbf{0}) - f(x_\star) \leq n|x_\star|. \quad (11)$$

The above shows that $x_\star \neq \mathbf{0}$. Hence, by a continuity argument (Corollary 13), there exists an $\epsilon > 0$ such that $\Psi_{K_\epsilon}(x_\star) = 1/(64d)$. The assumption that \bar{f} is m -strongly convex means that $K_\epsilon \subset \{x \in \mathbb{R}^d : |x| \leq \sqrt{2\epsilon/m}\}$ and hence

$$\frac{1}{64d} = \Psi_{K_\epsilon}(x_\star) \leq \frac{2\sqrt{2\epsilon/m}}{|x_\star|} \stackrel{\text{Eq. (11)}}{\leq} 4n^2 \sqrt{2\epsilon/m}.$$

Rearranging shows that $\epsilon \geq m/(2^{17}d^2n^4) \triangleq \epsilon_0$. Let $\gamma = 1 + 1/(9d)$ and $\mathcal{E} = \epsilon_0\{1, \gamma, \gamma^2, \dots\} \cap [0, 1]$. Since $\epsilon \geq \epsilon_0$, there exists a $\delta \in \mathcal{E}$ such that $\delta \in [\epsilon/\gamma, \epsilon]$. By the convexity of \bar{f} and the monotonicity of level sets, $\frac{1}{\gamma}K_\epsilon \subset K_{\epsilon/\gamma} \subset K_\delta \subset K_\epsilon$. Hence, by Corollary 12,

$$\Psi_{K_\delta}(x_\star) \in \left[\frac{1}{128d}, \frac{1}{32d} \right],$$

shows that $f \in \mathcal{F}_\delta$ and hence $\mathcal{F} = \mathcal{F}_0 \cup \bigcup_{\epsilon \in \mathcal{E}} \mathcal{F}_\epsilon$. The previous step demonstrated the existence of an exploratory distribution for each \mathcal{F}_ϵ with $\epsilon \in \mathcal{E}$.

Step 3: Combining By definition, $|\mathcal{E}| \leq 1 + \lceil \log_\gamma(1/\epsilon_0) \rceil$ with $\gamma = 1 + 1/(9d)$. Combining Lemma 4 with the exploratory distributions and partitions in steps 1 and 2 completes the proof. \square

4 Technical lemmas

Here we collect the necessary lemmas. The first few concern concavity, invariance, approximate monotonicity and continuity of Ψ .

Lemma 8. *Let $K \subset \mathbb{R}^d$ a convex body with $\mathbf{0} \in K$. Then, for any $x \notin K$, the function $z \mapsto \Psi_K(x, z)$ is concave on $P_x(K)$.*

Proof. With a slight abuse of notation, abbreviate $\Psi(z) = \Psi_K(x, z)$. Parameterise K by

$$K = \{z - \alpha x : f(z) \leq \alpha \leq g(z), z \in P_x(K), \alpha \in \mathbb{R}\},$$

where $f : P_x(K) \rightarrow \mathbb{R}$ is convex and $g : P_x(K) \rightarrow \mathbb{R}$ is concave (Fig. 3). Parameterise the chord connecting x and $z - g(z)x \in \partial K$ by $y(t) = (1-t)x + tz - tg(z)x$. Then,

$$\begin{aligned} 1 - \Psi(z) &= \min\{t \in (0, 1] : y(t) \in K\} \\ &= \min\{t \in (0, 1] : f(tz) \leq t + tg(z) - 1\} \\ &= \min\{1/s : sf(z/s) - g(z) + s - 1 \leq 0, s \in [1, \infty)\}. \end{aligned}$$

Let $h(z, s) = sf(z/s) - g(z) + s - 1$, which is the perspective of f minus a concave function and hence is convex on $P_x(K) \times [1, \infty)$. Note that $h(z, 1) = f(z) - g(z) \leq 0$, while the fact that $x \notin K$ implies that $\lim_{s \rightarrow \infty} h(z, s) = \infty$. Hence, for any $z, w \in P_x(K)$, there exist largest reals s and t such that $h(z, s) = 0$ and $h(w, t) = 0$ so that $1 - \Psi(z) = 1/s$ and $1 - \Psi(w) = 1/t$. By convexity, $h(\alpha z + (1 - \alpha)w, \alpha s + (1 - \alpha)t) \leq 0$ and therefore the largest value r for which

$h(\alpha z + (1 - \alpha)w, r) = 0$ is at least $\alpha s + (1 - \alpha)t$. Therefore,

$$\begin{aligned} 1 - \Psi(\alpha z + (1 - \alpha)w) &= \frac{1}{r} \leq \frac{1}{\alpha s + (1 - \alpha)t} \\ &\leq \frac{\alpha}{s} + \frac{1 - \alpha}{t} \\ &= \alpha(1 - \Psi(z)) + (1 - \alpha)(1 - \Psi(w)). \quad \square \end{aligned}$$




Figure 3: The construction used in the proof of Lemma 8.

Lemma 9. Let $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a linear bijection, $K \subset \mathbb{R}^d$ be a convex body with $\mathbf{0} \in K$ and let $x \notin K$. Then $\Psi_K(x) = \Psi_{TK}(Tx)$.

Proof. Let $z \in P_x(K)$ and $\alpha \in \mathbb{R}$ be the smallest value such that $z + \alpha x \in K$. By definition, $y \triangleq z + \alpha x = \pi_K(x, z)$ and since $z + \alpha x \in K$ if and only if $Tz + \alpha Tx \in TK$, it follows that

$$T\pi_K(x, z) = Ty = Tz + \alpha Tx = \pi_{TK}(Tx, P_{Tx}(Tz)) = \pi_{TK}(Tx, \Lambda(z)),$$

where we introduced $\Lambda = P_{Tx} \circ T$. Thus,

$$\Psi_{TK}(Tx, \Lambda(z)) = \frac{\text{vol}_{d-1}([Tx, Ty] \cap TK)}{\text{vol}_{d-1}([Tx, Ty])} = \frac{\text{vol}_{d-1}([x, y] \cap K)}{\text{vol}_{d-1}([x, y])} = \Psi_K(x, z).$$

Then, using the fact that Λ is a linear bijection between $P_x(K)$ and $P_{Tx}(TK)$,

$$\begin{aligned} \Psi_K(x) &= \frac{1}{\text{vol}_{d-1}(P_x(K))} \int_{P_x(K)} \Psi_K(x, z) \, \text{dvol}_{d-1}(z) \\ &= \frac{1}{\text{vol}_{d-1}(P_x(K))} \int_{P_x(K)} \Psi_{TK}(Tx, \Lambda(z)) \, \text{dvol}_{d-1}(z) \\ &= \frac{1}{\text{vol}_{d-1}(P_{Tx}(TK))} \int_{P_{Tx}(TK)} \Psi_{TK}(Tx, w) \, \text{dvol}_{d-1}(w) \\ &= \Psi_{TK}(Tx). \quad \square \end{aligned}$$

Lemma 10. *Let $A \subset \mathbb{R}^d$ be a convex body with $\mathbf{0} \in A$ and $\gamma > 1$. Then $\Psi_{\gamma A}(x) \geq \Psi_A(x)$ for all $x \notin \gamma A$.*

Proof. The first step is to argue that $\Psi_{\gamma A}(x, \gamma z) \geq \Psi_A(x, z)$, which follows using the notation in Fig. 4 because

$$\Psi_{\gamma A}(x, z) = \frac{|y - w|}{|y - x|} = \frac{|y - u|}{|y - \gamma x|} \geq \frac{|y - v|}{|y - \gamma x|} = \Psi_{\gamma A}(\gamma x, z).$$

The result follows from the fact that $P_x(\gamma A) = P_{\gamma x}(\gamma A)$ and Lemma 9:

$$\begin{aligned} \Psi_{\gamma A}(x) &= \int_{P_x(\gamma A)} \Psi_{\gamma A}(x, z) \, d\text{vol}_{d-1}(z) \\ &\geq \int_{P_{\gamma x}(\gamma A)} \Psi_{\gamma A}(\gamma x, z) \, d\text{vol}_{d-1}(z) = \Psi_{\gamma A}(\gamma x) \stackrel{\text{Lem. 9}}{=} \Psi_A(x). \quad \square \end{aligned}$$




Figure 4: The construction used in the proof of Lemma 10. The point y is $\pi_{\gamma A}(z)$ for some $z \in P_x(\gamma A)$. The points v and w are chosen so that $[y, \gamma x] \cap \gamma A = [v, y]$ and $[y, x] \cap A = [y, w]$. Finally, u is chosen on the chord $[\gamma x, y]$ so that $u - w$ is parallel to x .

Recall the definition of $\Psi_K^\infty(x)$ just after Eq. (4).

Lemma 11. *Let A and B be convex bodies such that $\mathbf{0} \in A \subset B \subset \gamma A$ for $\gamma > 1$. Assume that $x \notin \gamma A$ and $\Psi_{\gamma A}^\infty(x) \leq 1/2$. Then $\Psi_{\gamma A}(x) \leq (2\gamma - 1)^{d+1} \Psi_B(x)$.*

Proof. Let $\Lambda(z) = z/(2\gamma - 1)$. We claim that for all $z \in P_x(\gamma A)$,

$$\Psi_{\gamma A}(x, z) \leq (2\gamma - 1)^2 \Psi_B(x, \Lambda(z)). \quad (12)$$

Setting the proof of this aside for a moment, the consequence is that

$$\begin{aligned} \int_{P_x(\gamma A)} \Psi_{\gamma A}(x, z) \, d\text{vol}_{d-1}(z) &\leq (2\gamma - 1)^2 \int_{P_x(\gamma A)} \Psi_B(x, \Lambda(z)) \, d\text{vol}_{d-1}(z) \\ &= (2\gamma - 1)^{d+1} \int_{\Lambda(P_x(\gamma A))} \Psi_B(x, y) \, d\text{vol}_{d-1}(y) \\ &\leq (2\gamma - 1)^{d+1} \int_{P_x(B)} \Psi_B(x, y) \, d\text{vol}_{d-1}(y), \end{aligned}$$

where the last inequality is true because $\Lambda(P_x(\gamma A)) \subset P_x(A) \subset P_x(B)$, since $\gamma/(2\gamma - 1) \leq 1$. The lemma follows since $\text{vol}_{d-1}(P_x(B)) \leq \text{vol}_{d-1}(P_x(\gamma A))$. It remains to establish Eq. (12), which is a high-school exercise in length chasing. The quantities that follow are defined in the caption of Fig. 5.

$$\frac{|p-u|}{|p-x|} = \frac{|y-w|}{|y-x|} = \frac{|y-v||u-q|}{|y-x||v-q|} = \frac{|u-q|}{|v-q|} \Psi_{\gamma A}(x, z) = \frac{1}{\gamma} \Psi_{\gamma A}(x, z).$$

Hence, provided that $p \in B$, there exists a $t \in [r, z]$ such that $\Psi_B(x, t) \geq \Psi_{\gamma A}(x, z)/\gamma$. Furthermore,

$$\frac{|u-w|}{|y-p|} = 1 - \frac{|p-u|}{|p-x|} = 1 - \frac{1}{\gamma} \Psi_{\gamma A}(x, z) \quad \frac{|u-w|}{|y-q|} = \frac{|v-u|}{|v-q|} = \frac{\gamma-1}{\gamma}.$$

Dividing one by the other in the third equality below yields

$$\frac{|r-q|}{|z-q|} = \frac{|p-q|}{|y-q|} = 1 - \frac{|y-p|}{|y-q|} = 1 - \left(\frac{\gamma-1}{\gamma} \right) \frac{1}{1 - \Psi_{\gamma A}(x, z)/\gamma} \in \left[\frac{1}{2\gamma-1}, \frac{1}{\gamma} \right],$$

where the final relation holds because $\Psi_{\gamma A}^\infty(x) \leq 1/2$ by assumption. Extracting from the above that $|p-q|/|y-q| \leq 1/\gamma$ shows that $p \in A \subset B$, which confirms there exists a $t \in [r, z]$ such that $\Psi_B(x, t) \geq \Psi_{\gamma A}(x, z)/\gamma$. Furthermore, the above display also shows that $r \in [\Lambda(z), z]$ and hence, by the concavity of Ψ_B ,

$$\Psi_B(x, \Lambda(z)) \geq \frac{1}{2\gamma-1} \max_{t \in [\Lambda(z), z]} \Psi_B(x, t) \geq \frac{\Psi_{\gamma A}(x, z)}{\gamma(2\gamma-1)} \geq \frac{\Psi_{\gamma A}(x, z)}{(2\gamma-1)^2}. \quad \square$$




Figure 5: The construction used in the proof of Lemma 11, which has $q = \mathbf{0}$, $y = \pi_{\gamma A}(x, z)$. The point v is chosen so that $[x, y] \cap \gamma A = [v, y]$. The point u is such that $[v, q] \cap A = [u, q]$ and p is the intersection of $[q, y]$ and the affine hull $\text{aff}(\{x, u\})$. Lastly, w is the point in $[v, y]$ such that $[u, w]$ is parallel to $[q, y]$ and $r = P_x(p)$.

Corollary 12. *Suppose that A and B are convex bodies and $\mathbf{0} \in \frac{1}{\gamma}A \subset B \subset A$ for $\gamma = 1 + 1/(9d)$. Then, for any $x \notin A$ with $\Psi_A(x) \leq 1/(2d)$,*

$$\frac{\Psi_B(x)}{\Psi_A(x)} \in [1/2, 2].$$

Proof. The second part of Lemma 14 (given next) and concavity of $z \mapsto \Psi_A(x, z)$ shows that $\Psi_A^\infty(x) \leq 1/2$. By assumption, $\frac{1}{\gamma}A \subset B \subset A$, which implies that $\frac{1}{\gamma}B \subset \frac{1}{\gamma}A \subset B$ and so by Lemmas 10 and 11,

$$\begin{aligned} (2\gamma - 1)^{2d+2} \Psi_B(x) &\geq (2\gamma - 1)^{d+1} \Psi_A(x) && \text{(Lemma 11)} \\ &\geq (2\gamma - 1)^{d+1} \Psi_{A/\gamma}(x) && \text{(Lemma 10)} \\ &\geq \Psi_B(x), && \text{(Lemma 11)} \end{aligned}$$

where the last inequality also uses the fact that $\Psi_B^\infty(x) \leq \Psi_A^\infty(x) \leq 1/2$. The result follows from the choice of γ and naive bounding. \square

Corollary 13. *Suppose that $f \in \mathcal{F}$ is minimised at $\mathbf{0}$ and let $K_\epsilon = \{y \in \mathcal{K} : f(y) \leq f_* + \epsilon\}$. Then, for any $x \neq \mathbf{0}$, there exists an $\epsilon > 0$ such that*

$$\Psi_{K_\epsilon}(x) = \frac{1}{64d}.$$

Proof. That $\epsilon \mapsto \Psi_{K_\epsilon}^\infty(x)$ is continuous and non-decreasing is straightforward. Using the intermediate value theorem and the facts that $\lim_{\epsilon \rightarrow 0} \Psi_{K_\epsilon}(x) = 0$ and that $\Psi_{K_\epsilon}(x) = 1$ when ϵ is such that $x \in \partial K_\epsilon$, there exists a smallest value ϵ_{\max} such that

$$\Psi_{K_{\epsilon_{\max}}}^\infty(x) = 1/64.$$

By the second part of Lemma 14 means that $\Psi_{K_{\epsilon_{\max}}}(x) \geq 1/(64d)$. Strong convexity of f ensures level sets contract to a point as ϵ tends to zero and hence, $\lim_{\epsilon \rightarrow 0} \Psi_{K_\epsilon}(x) = 0$. Hence, by the intermediate value theorem it suffices to show that $\epsilon \mapsto \Psi_{K_\epsilon}(x)$ is continuous for $\epsilon \in (0, \epsilon_{\max}]$. Let $\epsilon \in (0, \epsilon_{\max}]$ and $\gamma > 1$. By convexity of f , $\frac{1}{\gamma}K_{\gamma\epsilon} \subset K_\epsilon \subset K_{\gamma\epsilon}$. Repeating the argument in the proof of Corollary 12 shows that $\Psi_{K_\epsilon}(x)$ tends to $\Psi_{K_{\gamma\epsilon}}(x)$ as γ tends to 1. \square

The next two lemmas are probably known. They concern the law of a concave random variable under the uniform probability measure on the domain, which is shown to have constant mass about its expectation.

Lemma 14. *Let $A \subset \mathbb{R}^{d-1}$ be convex and $\varphi : A \rightarrow [0, \infty)$ be concave. Then,*

$$\frac{1}{\text{vol}_{d-1}(A)} \int_A \varphi^2 \, \text{dvol}_{d-1} \leq 2^{5/2} \left(\frac{\int_A \varphi \, \text{dvol}_{d-1}}{\text{vol}_{d-1}(A)} \right)^2.$$

Furthermore, $\max_{x \in A} \varphi(x) \leq \frac{d}{\text{vol}_{d-1}(A)} \int_A \varphi \, \text{dvol}_{d-1}$.

Proof. Let $B = \{(x, y) : x \in A, |y| \leq \varphi(x)\} \subset \mathbb{R}^d$, which is convex. Define $\theta = (0, \dots, 0, 1)$ and $h(t) = \text{vol}_{d-1}(B \cap (\theta^\perp + t\theta))$. By Brunn's concavity principle [1, theorem 1.2.1], $t \mapsto h(t)^{1/(d-1)}$ is concave on its support. By the arithmetic-geometric inequality, for $\lambda \in [0, 1]$ and s, t in the support of h ,

$$h(\lambda s + (1 - \lambda)t)^{\frac{1}{d-1}} \geq \lambda h(s)^{\frac{1}{d-1}} + (1 - \lambda)h(t)^{\frac{1}{d-1}} \geq h(s)^{\frac{\lambda}{d-1}} h(t)^{\frac{1-\lambda}{d-1}},$$

which implies that h is log-concave. Then,

$$\begin{aligned} \frac{1}{\text{vol}_{d-1}(A)} \int_A \varphi^2 \, \text{dvol}_{d-1} &= \frac{1}{\text{vol}_{d-1}(A)} \int_B |\langle x, \theta \rangle| \, \text{dvol}_d(x) \\ &\leq \frac{1}{\text{vol}_{d-1}(A)} \left(\text{vol}_d(B) \int_B \langle x, \theta \rangle^2 \, \text{dvol}_d(x) \right)^{1/2} \\ &= \frac{1}{\text{vol}_{d-1}(A)} \left(\text{vol}_d(B) \int_{-\infty}^{\infty} t^2 h(t) \, dt \right)^{1/2} \\ &\leq \frac{1}{\text{vol}_{d-1}(A)} \left(\frac{2 \text{vol}_d(B)}{h(0)^2} \left(\int_{-\infty}^{\infty} h(t) \, dt \right)^3 \right)^{1/2} \\ &= 2^{5/2} \left(\frac{\int_A \varphi \, \text{dvol}_{d-1}}{\text{vol}_{d-1}(A)} \right)^2, \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz and the second from corollary 2.24 in the notes by Tkocz [16] and log-concavity of h and $\int_{-\infty}^{\infty} th(t) \, dt = 0$, which holds because h is symmetric. The last equality follows since $h(0) = \text{vol}_{d-1}(A)$ and $\int_{-\infty}^{\infty} h(t) \, dt = \text{vol}_d(B) = 2 \int_A \varphi \, \text{dvol}_{d-1}$. For the second part, let $x \in \arg \max_{y \in A} \varphi(y)$ and $C \subset B$ be the convex hull of $(x, \varphi(x))$, $(x, -\varphi(x))$ and $A \times \{0\}$, which is the union of two cones with the same base. Then,

$$\int_A \varphi \, \text{dvol}_{d-1} = \frac{1}{2} \text{vol}_d(B) \geq \frac{1}{2} \text{vol}_d(C) = \frac{\varphi(x) \text{vol}_{d-1}(A)}{d},$$

where the second equality follows from the formula for the volume of a cone. \square

Lemma 15. *Let $A \subset \mathbb{R}^{d-1}$ be convex and $\varphi : A \rightarrow [0, \infty)$ be concave and $\nu = \text{dvol}_{d-1} / \text{vol}_{d-1}(A)$ be the uniform probability measure on A . Then,*

$$\nu \left(\left\{ x \in A : \frac{\varphi(x)}{\int_A \varphi \, \text{d}\nu} \in [1/4, 16] \right\} \right) \geq 1/32.$$

Proof. By Markov's inequality,

$$\nu \left(\left\{ x \in A : \varphi(x) \leq 16 \int_A \varphi \, \text{d}\nu \right\} \right) \geq 1 - \frac{1}{16}.$$

For non-negative random variable X and any $\theta \in (0, 1)$, the Payley-Zygmund inequality says that $\mathbb{P}(X \geq \theta \mathbb{E}[X]) \geq (1 - \theta)^2 \mathbb{E}[X]^2 / \mathbb{E}[X^2]$. Combining this with

Lemma 14 shows that

$$\nu \left(\left\{ x \in A : \varphi(x) \geq \frac{1}{4} \int_A \varphi \, d\nu \right\} \right) \geq \left(1 - \frac{1}{4} \right)^2 \frac{(\int_A \varphi \, d\nu)^2}{\int_A \varphi^2 \, d\nu} \geq \frac{9}{2^{13/2}} \geq \frac{3}{32}.$$

Combining the previous displays yields the result. \square

5 Lipschitz and strong convexity relaxation

The last ingredient of the proof is to show that the Lipschitz and strong convexity assumptions are indeed mild. More or less the same argument has been used elsewhere [4, 2].

Proposition 16. *Theorem 1 follows from Theorems 2 and 3.*

Proof. Let \mathcal{F} be the set of n -Lipschitz and m -strongly convex functions from convex body $\mathcal{J} \subset \mathbb{R}^d$ to $[0, 1]$. Theorems 2 and 3 show that

$$\mathfrak{R}_n^*(\mathcal{F}) \leq \text{const} \cdot d^{2.5} \sqrt{n} \log(nd \text{diam}(\mathcal{J})/m). \quad (13)$$

Theorem 1 assumes that \mathcal{K} contains a unit-radius Euclidean ball. By translation we may assume that $\{x \in \mathbb{R}^d : |x| \leq 1\} \subset \mathcal{K}$. Let $(f_t)_{t=1}^n$ be an arbitrary sequence of convex functions from \mathcal{K} to $[0, 1]$, possibly non-Lipschitz and non-strongly convex. Define $\text{dist}(x, A) = \min_{y \in A} |x - y|$ and $\mathcal{J} = \{x \in \mathcal{K} : \text{dist}(x, \partial\mathcal{K}) \geq 1/n\}$, which is a convex subset of \mathcal{K} . Next, let $(\hat{f}_t)_{t=1}^n$ be the sequence of convex functions from $\mathcal{K} \rightarrow [0, 1]$ given by

$$\hat{f}_t(x) = \left(\frac{n}{n+1} \right) \left(f_t(x) + \frac{1}{n} \left(\frac{|x|}{\text{diam}(\mathcal{K})} \right)^2 \right) \in [0, 1].$$

That \hat{f}_t is m -strongly convex with $m = 1/((n+1) \text{diam}(\mathcal{K})^2)$ and $\hat{f}_t(x) \in [0, 1]$ is immediate. To see that \hat{f}_t is n -Lipschitz on \mathcal{J} , let $g \in \partial \hat{f}_t(x)$ be a subgradient of \hat{f}_t for $x \in \mathcal{J}$ and let $\alpha > 0$ be such that $x + \alpha g/|g| \in \partial\mathcal{K}$. Then, by convexity and boundedness of \hat{f}_t , $1 \geq \hat{f}_t(x + \alpha g/|g|) - \hat{f}_t(x) \geq \langle g, \alpha g/|g| \rangle = \alpha |g|$. By the definition of \mathcal{J} , $\alpha \geq 1/n$ and hence $|g| \leq n$. Running a policy witnessing Eq. (13),

$$\begin{aligned} \text{const} \cdot d^{2.5} \sqrt{n} \log(nd \text{diam}(\mathcal{J})/m) &\geq \max_{y \in \mathcal{J}} \mathbb{E} \left[\sum_{t=1}^n \hat{f}_t(x_t) - \hat{f}_t(y) \right] \\ &\geq \frac{n}{n+1} \left[\max_{y \in \mathcal{J}} \mathbb{E} \left[\sum_{t=1}^n f_t(x_t) - f_t(y) \right] - 1 \right] \\ &\geq \frac{n}{n+1} \left[\max_{x \in \mathcal{K}} \mathbb{E} \left[\sum_{t=1}^n f_t(x_t) - f_t(x) \right] - 2 \right]. \end{aligned} \quad (14)$$

Only the last inequality presents any challenge. To see why this is true, let $x \in \mathcal{K}$ be the minimiser of $\sum_{t=1}^n f_t$ on \mathcal{K} and let $y = (1 - 1/n)x$. Since $\{x : |x| \leq 1\} \subset \mathcal{K}$, it follows that $y \in \mathcal{K}$. Furthermore, since f_t is convex and bounded in $[0, 1]$,

$$f_t(y) \leq (1 - 1/n)f_t(x) + \frac{1}{n}f_t(\mathbf{0}) \leq f_t(x) + \frac{1}{n}.$$

The result follows by rearranging Eq. (14), noting that $\text{diam}(\mathcal{J}) \leq \text{diam}(\mathcal{K})$ and by substituting the value of m . Note, the dependence on the dimension in the logarithm is dropped, since when $d \geq n$ the regret is better bounded by n . \square

A Proof of Theorem 2

For completeness, we outline the key steps in the proof of Theorem 2, giving the main arguments and referring to the specific parts of the literature where necessary. The best reference is the article by [Bubeck and Eldan \[2\]](#).

Discretisation Let $\mathcal{C} \subset \mathcal{K}$ be a finite collection of points such that for all $x \in \mathcal{K}$ there exists a $y \in \mathcal{C}$ with $|x - y| \leq \epsilon \triangleq 1/(n^2 \max(1, (2\beta)^{1/2}))$. Standard covering number results [[1](#), corollary 4.1.14] show that \mathcal{C} may be chosen so that

$$\log |\mathcal{C}| \leq d \log \left(\frac{3 \text{diam}(\mathcal{K})}{\epsilon} \right). \quad (15)$$

Let $\Pi(x) = \arg \min_{y \in \mathcal{C}} |x - y|$ and $y^* = \arg \min_{y \in \mathcal{C}} \sum_{t=1}^n f_t(y)$. Since functions $f \in \mathcal{F}$ are n -Lipschitz, for any policy,

$$\mathfrak{R}_n = \max_{x^* \in \mathcal{K}} \mathbb{E} \left[\sum_{t=1}^n f_t(x_t) - f_t(x^*) \right] \leq 1 + \mathbb{E} \left[\sum_{t=1}^n f_t(x_t) - f_t(y^*) \right]. \quad (16)$$

Exploratory distributions can also be discretised. Given a distribution μ on \mathcal{F} and an exploratory distribution ρ , let $\xi = \rho \circ \Pi^{-1}$, which is supported on \mathcal{C} . Then, repeatedly using the assumption that $f \in \mathcal{F}$ are n -Lipschitz and naive bounding,

$$\begin{aligned} \int_{\mathcal{K}} \bar{f} d\xi - \int_{\mathcal{F}} f_* d\mu(f) &\leq \frac{1}{n} + \int_{\mathcal{K}} \bar{f} d\rho - \int_{\mathcal{F}} f_* d\mu(f) \\ &\leq \frac{1}{n} + \alpha + \sqrt{\beta \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f} - f)^2 d\rho d\mu(f)} \\ &\leq \frac{2}{n} + \alpha + \sqrt{2\beta \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f} - f)^2 d\xi d\mu(f)}. \end{aligned} \quad (17)$$

Hence, ξ is a good exploratory distribution with slightly larger constants than ρ .

Minimax duality Using the assumption that $f \in \mathcal{F}$ are n -Lipschitz, the definition of the discretisation and then a minimax theorem [12, theorem 1], we obtain

$$\begin{aligned} \mathfrak{R}_n^*(\mathcal{F}) &\leq 1 + \inf_{\mathcal{C}\text{-policies}} \sup_{(f_t)_{t=1}^n \in \mathcal{F}^n} \mathbb{E} \left[\sum_{t=1}^n f_t(x_t) - f_t(y^*) \right] \\ &= 1 + \sup_{\nu} \inf_{\mathcal{C}\text{-policies}} \underbrace{\mathbb{E} \left[\sum_{t=1}^n f_t(x_t) - f_t(y^*) \right]}_{\mathfrak{BR}_n}. \end{aligned} \quad (18)$$

where the inf over policies is restricted to policies playing in \mathcal{C} and the sup on the right-hand side is over all finitely supported distributions on \mathcal{F} . The expectation on the right-hand side is now over the randomness in both $(x_t)_{t=1}^n$ and $(f_t)_{t=1}^n$, with the latter sampled from ν . Note, the comparator y^* is now a random variable. The expectation on the right-hand side is called the Bayesian regret, denoted by \mathfrak{BR}_n , which depends on prior and policy. For the remainder of this section we fix a prior ν and design a policy for which the Bayesian regret with respect to ν is bounded in terms of n , d , $\text{diam}(\mathcal{K})$ and the constants α and β that define the conditions for an exploratory distribution in Theorem 2.

Note, [Bubeck et al. \[3\]](#) also propose a minimax theorem for convex bandits without a discretisation argument, which inspired the argument for partial monitoring quoted above [12]. Some details in the former proof are missing, however, such as the choice of topology on the space of convex functions and the continuity of the regret as a function of a deterministic policy.

Posterior, policy and information ratio The policy will sample x_t from a probability measure ξ_t on \mathcal{C} , where ξ_t is an exploratory distribution provided by combining the conditions of Theorem 2 and Eq. (17). Let \mathbb{E}_t be the conditional expectation with respect to $(x_s)_{s=1}^t$ and $(f_s(x_s))_{s=1}^t$. Then, for each $y \in \mathcal{C}$ and $x \in \mathcal{K}$, let

$$f_{t,y}(x) = \mathbb{E}_{t-1}[f_t(x)|y^* = y] \quad \bar{f}_t(x) = \mathbb{E}_{t-1}[f_t(x)].$$

Convexity, m -strong convexity and n -Lipschitzness are preserved under averaging, so $f_{t,y}$ and \bar{f}_t are in \mathcal{F} almost surely. Next, let μ_t be the finitely supported probability distribution on \mathcal{F} for which $\mu_t(\{f_{t,y}\}) = \mathbb{E}_{t-1}[\mathbf{1}_{y^*=y}]$. By the conditions in Theorem 2 and Eq. (17), there exists a probability measure ξ_t supported on \mathcal{C} such that

$$\int_{\mathcal{K}} \bar{f}_t d\xi_t - \int_{\mathcal{F}} f_{\star} d\mu_t(f) \leq \frac{2}{n} + \alpha + \sqrt{2\beta \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f}_t(x) - f(x))^2 d\xi_t(x) d\mu_t(f)}.$$

Then, considering the policy that samples x_t from ξ_t and using Eq. (16),

$$\begin{aligned} \mathfrak{B}\mathfrak{R}_n &= \mathbb{E} \left[\sum_{t=1}^n \int_{\mathcal{K}} \bar{f}_t \, d\xi_t - \int_{\mathcal{F}} f(y^*) \, d\mu_t(f) \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^n \int_{\mathcal{K}} \bar{f}_t \, d\xi_t - \int_{\mathcal{F}} f_* \, d\mu_t(f) \right] \\ &\leq 2 + n\alpha + \mathbb{E} \left[\sum_{t=1}^n \sqrt{2\beta \int_{\mathcal{F}} \int_{\mathcal{K}} (\bar{f}_t(x) - f(x))^2 \, d\xi_t(x) \, d\mu_t(f)} \right]. \end{aligned}$$

The final step is the bound

$$\mathbb{E} \left[\sum_{t=1}^n \sqrt{2\beta v_t} \right] \leq \sqrt{2n\beta \mathbb{E} \left[\sum_{t=1}^n v_t \right]} \leq \sqrt{n\beta \log(|\mathcal{C}|)},$$

where the first inequality is Jensen's and the second follows from Pinsker's inequality to bound the squared difference of expectations in v_t and the chain rule for mutual information [3, lemma 5]. Therefore,

$$\mathfrak{B}\mathfrak{R}_n \leq 2 + n\alpha + \sqrt{n\beta \log |\mathcal{C}|}.$$

Since this bound is independent of the prior ν , it follows from Eq. (18) and Eq. (15) that

$$\mathfrak{R}_n^*(\mathcal{F}) \leq 3 + n\alpha + \sqrt{\beta dn \log (3n^2 \max(1, (2\beta)^{1/2}) \text{diam}(\mathcal{K}))}.$$

Acknowledgements Many thanks to my colleagues András György and Marcus Hutter for their useful suggestions.

References

- [1] S. Artstein-Avidan, A. Giannopoulos, and V. D. Milman. *Asymptotic geometric analysis, Part I*, volume 202. American Mathematical Soc., 2015.
- [2] S. Bubeck and R. Eldan. Exploratory distributions for convex functions. *Mathematical Statistics and Learning*, 1(1):73–100, 2018.
- [3] S. Bubeck, O. Dekel, T. Koren, and Y. Peres. Bandit convex optimization: \sqrt{T} regret in one dimension. In *Proceedings of the 28th Conference on Learning Theory*, pages 266–278, Paris, France, 2015. JMLR.org.
- [4] S. Bubeck, Y-T. Lee, and R. Eldan. Kernel-based methods for bandit convex optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 72–85, 2017.

- [5] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [6] V. Dani, T. P. Hayes, and S. M. Kakade. Stochastic linear optimization under bandit feedback. In *Proceedings of the 21st Conference on Learning Theory*, pages 355–366, 2008.
- [7] A. Giannopoulos and M. Papadimitrakis. Isotropic surface area measures. *Mathematika*, 46(1):1–13, 1999.
- [8] E. Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [9] E. Hazan and K. Levy. Bandit convex optimization: Towards tight bounds. In *Advances in Neural Information Processing Systems*, pages 784–792, 2014.
- [10] E. Hazan and Y. Li. An optimal algorithm for bandit convex optimization. *arXiv preprint arXiv:1603.04350*, 2016.
- [11] X. Hu, LA. Prashanth, A. György, and Cs. Szepesvári. (bandit) convex optimization with biased noisy gradient oracles. In *Artificial Intelligence and Statistics*, pages 819–828, 2016.
- [12] T. Lattimore and Cs. Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. In *Proceedings of the 32nd Conference on Learning Theory*, pages 2111–2139, Phoenix, USA, 2019. PMLR.
- [13] F. Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- [14] D. Russo and B. Van Roy. Learning to optimize via information-directed sampling. In *Advances in Neural Information Processing Systems*, pages 1583–1591. Curran Associates, Inc., 2014.
- [15] D. Russo and B. Van Roy. An information-theoretic analysis of Thompson sampling. *Journal of Machine Learning Research*, 17(1):2442–2471, 2016. ISSN 1532-4435.
- [16] T. Tkocz. *Asymptotic Convex Geometry Lecture Notes*. 2018.