

Optimization of Smooth Functions With Noisy Observations: Local Minimax Rates

Yining Wang¹, Sivaraman Balakrishnan, and Aarti Singh

Abstract—We consider the problem of *global optimization* of an unknown non-convex smooth function with noisy zeroth-order feedback. We propose a *local minimax* framework to study the fundamental difficulty of optimizing smooth functions with adaptive function evaluations. We show that for functions with fast growth around their global minima, carefully designed optimization algorithms can identify a near global minimizer with many fewer queries than worst-case global minimax theory predicts. For the special case of strongly convex and smooth functions, our implied convergence rates match the ones developed for zeroth-order *convex* optimization problems. On the other hand, we show that in the worst case no algorithm can converge faster than the minimax rate of estimating an unknown function in the ℓ_∞ -norm. Finally, we show that non-adaptive algorithms, though optimal in a global minimax sense, do not attain the optimal local minimax rate.

Index Terms—Optimization of smooth functions, nonparametric statistics, local minimax analysis.

I. INTRODUCTION

GLOBAL function optimization with stochastic (zeroth-order) query oracles is an important problem in optimization, machine learning and statistics. To optimize an unknown bounded function $f : \mathcal{X} \mapsto \mathbb{R}$ defined on a known compact d -dimensional domain $\mathcal{X} \subseteq \mathbb{R}^d$, the data analyst makes n *active* queries $x_1, \dots, x_n \in \mathcal{X}$ and observes

$$y_t = f(x_t) + w_t, \quad w_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad t = 1, \dots, n. \quad (1)$$

The queries x_1, \dots, x_t are *active* in the sense that the selection of x_t can depend on the previous queries and their responses $x_1, y_1, \dots, x_{t-1}, y_{t-1}$. After n queries, an estimate $\hat{x}_n \in \mathcal{X}$ is produced that approximately minimizes the unknown function f . Such “active query” models are relevant in a broad range of (noisy) global optimization applications, for instance in hyper-parameter tuning of machine learning algorithms [1] and

Manuscript received August 10, 2018; revised April 21, 2019; accepted May 5, 2019. S. Balakrishnan was supported in part by the NSF under Grant DMS-17130003. Y. Wang and A. Singh were supported in part by the NSF under Grant CCF-1563918 and in part by the AFRL under Grant FA8750-17-2-0212. This paper was presented in part at the 2018 NeurIPS Conference.

Y. Wang and A. Singh are with the Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: yiningwa,aarti@cs.cmu.edu).

S. Balakrishnan is with the Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: siva@stat.cmu.edu).

Communicated by K. Chaudhuri, Associate Editor for Statistical Learning. Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2019.2921985

¹The exact Gaussianity of the independent noise variables ε_t is not crucial and our results can be easily generalized to sub-Gaussian noise.

sequential design in material synthesis experiments where the goal is to maximize the strength of the synthesized material as a function of experimental settings [2], [3]. We refer the readers to Section II-A for a rigorous formulation of the active query model and contrast it with the classical passive query model.

The error of the estimate \hat{x}_n is measured by the difference of $f(\hat{x}_n)$ and the *global minimum* of f :

$$\mathcal{L}(\hat{x}_n; f) := f(\hat{x}_n) - f^* \quad \text{where } f^* := \inf_{x \in \mathcal{X}} f(x). \quad (2)$$

To simplify our presentation, throughout the paper we take the domain \mathcal{X} to be the d -dimensional unit cube $[0, 1]^d$, while our results can be easily generalized to other compact domains satisfying minimal regularity conditions.

When f belongs to a smoothness class, say the Hölder class with exponent α , a straightforward global optimization method is to first sample n points uniformly at random from \mathcal{X} and then construct nonparametric estimates \hat{f}_n of f using nonparametric regression methods such as kernel smoothing or local polynomial regression [4], [5]. Classical analysis shows that the sup-norm reconstruction error $\|\hat{f}_n - f\|_\infty = \sup_{x \in \mathcal{X}} |\hat{f}_n(x) - f(x)|$ can be upper bounded by $\tilde{O}_{\mathbb{P}}(n^{-\alpha/(2\alpha+d)})^2$. This global reconstruction guarantee then implies an $\tilde{O}_{\mathbb{P}}(n^{-\alpha/(2\alpha+d)})$ upper bound on $\mathcal{L}(\hat{x}_n; f)$ by considering an estimate $\hat{x}_n \in \mathcal{X}$ for which $\hat{f}_n(\hat{x}_n) = \inf_{x \in \mathcal{X}} \hat{f}_n(x)$ (such an \hat{x}_n exists because \mathcal{X} is closed and bounded). Formally, we have the following proposition (proved in the Appendix) that converts a global reconstruction guarantee into an upper bound on the optimization error:

Proposition 1. *Suppose $\hat{f}_n(\hat{x}_n) = \inf_{x \in \mathcal{X}} \hat{f}_n(x)$. Then $\mathcal{L}(\hat{x}_n; f) \leq 2\|\hat{f}_n - f\|_\infty$.*

Typically, fundamental limits on the optimal optimization error are understood through the lens of *minimax analysis* where the object of study is the (global) minimax risk:

$$\inf_{\hat{x}_n} \sup_{f \in \mathcal{F}} \mathbb{E}_f \mathcal{L}(\hat{x}_n, f), \quad (3)$$

where \mathcal{F} is a certain class of smooth functions such as the Hölder class. Although optimization appears to be easier than global reconstruction, we show in this paper that the $n^{-\alpha/(2\alpha+d)}$ rate is *not* improvable in the global minimax sense in over Hölder classes. Such a surprising phenomenon was also noted in previous works [6]–[8] for related problems. On the

²In the $\tilde{O}(\cdot)$ or $\tilde{O}_{\mathbb{P}}(\cdot)$ notation we suppress constant factors and terms that depend poly-logarithmically on n .

other hand, extensive empirical evidence suggests that non-uniform/active allocations of query points can significantly reduce optimization error in practical global optimization of smooth, non-convex functions [1]. This raises the interesting question of understanding, from a theoretical perspective, the conditions under which the global optimization of smooth functions is *easier* than their reconstruction, and the power of *active/feedback-driven* queries that play important roles in global optimization.

In this paper, we propose a theoretical framework that partially answers the above questions. In contrast to classical *global* minimax analysis of nonparametric estimation problems, we adopt a *local analysis* which characterizes the optimal convergence rate of optimization error when the underlying function f is within a neighborhood of a “reference” function f_0 . (See Section II-B for the rigorous local minimax formulation considered in this paper.) Our main results are to characterize the local convergence rates $R_n(f_0)$ for a wide range of reference functions $f_0 \in \mathcal{F}$. Concretely, our contributions can be summarized as follows:

- 1) We design an iterative (active) algorithm whose optimization error $\mathcal{L}(\hat{x}_n; f)$ converges at a rate of $R_n(f_0)$ depending on the reference function f_0 . When the level-sets of f_0 satisfy certain regularity and polynomial growth conditions, the local rate $R_n(f_0)$ can be upper bounded by $R_n(f_0) = \tilde{O}(n^{-\alpha/(2\alpha+d-\alpha\beta)})$, where $\beta \in [0, d/\alpha]$ is a parameter depending on f_0 that characterizes the volume growth of the *level-sets* of the reference function f_0 . (See assumption (A2), Proposition 2 and Theorem 1 for details). The rate matches the global minimax convergence rate $n^{-\alpha/(2\alpha+d)}$ for worst-case f_0 where $\beta = 0$, but can be much faster when $\beta > 0$. We emphasize that our algorithm has no knowledge of the reference function f_0 and achieves this rate adaptively.
- 2) We prove *local* minimax lower bounds that match the $n^{-\alpha/(2\alpha+d-\alpha\beta)}$ upper bound, up to logarithmic factors in n . More specifically, we show that *even if* f_0 is known, no (active) algorithm can estimate f in close neighborhoods of f_0 at a rate faster than $n^{-\alpha/(2\alpha+d-\alpha\beta)}$. We further show that, if active queries are not available and queries x_1, \dots, x_n are i.i.d. uniformly sampled from \mathcal{X} , then the $n^{-\alpha/(2\alpha+d)}$ global minimax rate also applies locally regardless of how large β is. Thus, there is an explicit gap between local minimax rates in the active and uniform query models when β is large.
- 3) In the special case when f is *convex*, the global optimization problem is usually referred to as *zeroth-order convex optimization* and this problem has been widely studied [9]–[14]. Our results imply that, when f_0 is *strongly convex* and *smooth*, the local minimax rate $R_n(f_0)$ is on the order of $\tilde{O}(n^{-1/2})$, which matches the convergence rates in [11]. Additionally, our negative results (Theorem 2) indicate that the $n^{-1/2}$ rate cannot be achieved if f_0 is merely convex, which seems to contradict $n^{-1/2}$ results in [13], [14] that do not require strong convexity of f . However, it should be noted that mere convexity of f_0 does not imply convexity of f in

a neighborhood of f_0 (e.g., $\|f - f_0\|_\infty \leq \varepsilon$). Our results show significant differences in the intrinsic difficulty of zeroth-order optimization of convex and near-convex functions.

A. Related Work

Global optimization, known variously as *black-box optimization*, *Bayesian optimization* and the *continuum-armed bandit*, has a long history in the optimization research community [15], [16] and has also received a significant amount of recent interest in statistics and machine learning [1], [6], [8], [17]–[19]. Many previous works [17], [20] have derived rates for non-convex smooth payoffs in “continuum-armed” bandit problems.

The papers [21], [22] are closely related to our work. They studied the related problem of estimating the set of all optima of a smooth function in the Hausdorff distance. For Hölder smooth functions with polynomial growth, the paper [21] derives an $n^{-1/(2\alpha+d-\alpha\beta)}$ minimax rate for $\alpha < 1$ (subsequently improved to include $\alpha \geq 1$ in [23]). This result is similar to our Propositions 2 and 3. The papers [21], [22] also discussed adaptivity to unknown smoothness parameters. We however remark on several differences between our work and the papers [21], [22]. First, in [21], [22] only functions with polynomial growth are considered, while in our Theorems 1 and 2 functionals $\varepsilon_n^U(f_0)$ and $\varepsilon_n^L(f_0)$ are proposed for general reference functions f_0 satisfying mild regularity conditions, which include functions with polynomial growth as special cases. In addition, [21] considers the harder problem of estimating maxima sets in Hausdorff distance, as opposed to the problem of producing a single approximately optimal solution \hat{x}_T . As a result, the minimax lower bounds in [21] do not apply to this latter setting. An algorithm, without distinguishing between two functions with different optima sets, can nevertheless produce a good approximate optimizer as long as the two functions under consideration have *overlapping* optima sets. New constructions and information-theoretic techniques are therefore required to prove lower bounds under the weaker (one-point) approximate optimization framework. Finally, we prove minimax lower bounds when only *uniform* query points are available and demonstrate a significant gap between algorithms having access to uniformly sampled or adaptively chosen data points.

The papers [18], [19] imposed additional assumptions on the level-sets of the underlying function to obtain an improved convergence rate. The level-set assumptions considered in the mentioned references are rather restrictive and essentially require the underlying function to be uni-modal, while our assumptions are much more flexible and apply to multi-modal functions as well. In addition, [18], [19] considered a *noiseless* setting in which exact function evaluations $f(x_t)$ can be obtained, while our paper studies the noise corrupted model in (1) for which vastly different convergence rates are derived. Finally, no matching lower bounds were proved in the papers [18], [19].

The (stochastic) global optimization problem is similar to *mode estimation* of either densities or regression functions,

which has a rich literature [24]–[26]. An important difference between statistical mode estimation and global optimization is the way sample/query points $x_1, \dots, x_n \in \mathcal{X}$ are distributed: in mode estimation it is customary to assume the samples are independently and identically distributed, while in global optimization sequential designs of samples/queries are typical. Furthermore, to estimate/locate the mode of an unknown density or regression function, such a mode has to be well-defined; on the other hand, producing an estimate \hat{x}_n with small $\mathcal{L}(\hat{x}_n, f)$ is easier and results in weaker conditions imposed on the underlying function.

Methodology-wise, our proposed algorithm is conceptually similar to the abstract *Pure Adaptive Search (PAS)* framework proposed and analyzed in [27]. The iterative procedure also resembles disagreement-based active learning methods [28]–[30] and the “successive rejection” algorithm in bandit problems [31]. The intermediate steps of candidate point elimination can also be viewed as level-set estimation problems [32]–[34] or cluster-tree estimation problems [35], [36] with active queries.

Another line of research has focused on *first-order* optimization of quasi-convex or non-convex functions [37]–[42], in which exact or unbiased evaluations of function *gradients* are available at query points $x \in \mathcal{X}$. The paper [42] considered a Cheeger’s constant restriction on level-sets which is similar to our level-set regularity assumptions (A2 and A2’). The papers [43], [44] studied local minimax rates for the first-order optimization of convex functions. First-order optimization differs significantly from our setting because unbiased gradient estimation is generally impossible in the model of (1). Furthermore, most works on (first-order) non-convex optimization focus on obtaining stationary points or local minima, while we consider the problem of finding a (near) global minima.

221 B. Comparison with the HOO Algorithm

222 The HOO algorithm [17], as well as similar algorithms
223 such as Algorithm 2 in [45] and the POO algorithm in [22],
224 are theoretically well-studied methods for global optimization.
225 Below we summarize the differences of our results and the
226 ones from these works.

227 (a) Weaker Smoothness Conditions I: In Algorithm 1,
228 we use local polynomial estimation as a sub-routine
229 to obtain local estimates of the objective function
230 f . Compared to the sample average approach in
231 HOO (e.g., Algorithm 2 in [45]), local polynomial
232 estimates have the advantage of being unbiased for
233 the estimation of low-degree polynomials. This trans-
234 lates to the improved (A1) Hölder-continuity condition
235 that *only* restricts the $\lfloor \alpha \rfloor$ -th order derivatives
236 of objective functions. More specifically, the actual
237 function values of $f(x)$ and $f(x')$ for x, x' close
238 to each other can be very different, as long as such
239 differences can be perfectly modeled by low-degree
240 polynomials. This is in contrast to the smoothness
241 conditions imposed in [17], [45] which essentially
242 require $f(x)$ to be close to $f(x^*)$ for x close to x^* the
243 optima of f .

244 (b) Weaker Smoothness Conditions II: Our results in
245 Section IV-C hold on functions that are only assumed
246 to be smooth in regions close to its global minimum, in
247 contrast to Definition 1 in [45] and many other existing
248 works that place smoothness assumptions on the entire
249 domain of the objective function f .

250 (c) Spatially Restricted Queries: Our proposed algorithm is
251 “grid” based, and can be run on any sufficiently dense
252 finite grid G_n in \mathcal{X} and does not need to have the
253 capacity to query arbitrary points in \mathcal{X} . As a result,
254 our algorithm can be run in experimental settings where
255 queries are restricted to belong to a large pool of a-priori
256 chosen points.

257 (d) Results for any Smooth Function: Our algorithm and
258 lower bounds yield essentially tight results for the
259 complexity of optimization of arbitrary smooth func-
260 tions. While these rates are most interpretable under
261 the level-set growth conditions (also studied in [45]) our
262 results also yield nearly matching guarantees for other
263 (arbitrary, smooth) functions f_0 .

264 II. BACKGROUND AND NOTATION

265 We first review standard asymptotic notation that will
266 be used throughout this paper. For two sequences $\{a_n\}_{n=1}^\infty$
267 and $\{b_n\}_{n=1}^\infty$, we write $a_n = O(b_n)$ or $a_n \lesssim b_n$ if
268 $\limsup_{n \rightarrow \infty} |a_n|/|b_n| < \infty$, or equivalently $b_n = \Omega(a_n)$ or
269 $b_n \gtrsim a_n$. Denote $a_n = \Theta(b_n)$ or $a_n \asymp b_n$ if both $a_n \lesssim b_n$
270 and $a_n \gtrsim b_n$ hold. We also write $a_n = o(b_n)$ or equivalently
271 $b_n = \omega(a_n)$ if $\lim_{n \rightarrow \infty} |a_n|/|b_n| = 0$. For two sequences
272 of random variables $\{A_n\}_{n=1}^\infty$ and $\{B_n\}_{n=1}^\infty$, denote $A_n =$
273 $O_{\mathbb{P}}(B_n)$ if for every $\epsilon > 0$, there exists $C > 0$ such that
274 $\limsup_{n \rightarrow \infty} \Pr[|A_n| > C|B_n|] \leq \epsilon$. For $r > 0$, $1 \leq p \leq \infty$
275 and $x \in \mathbb{R}^d$, we denote by $B_r^p(x) := \{z \in \mathbb{R}^d : \|z - x\|_p \leq r\}$
276 the d -dimensional ℓ_p -ball of radius r centered at x , where
277 the vector ℓ_p norm is defined as $\|x\|_p := (\sum_{j=1}^d |x_j|^p)^{1/p}$
278 for $1 \leq p < \infty$ and $\|x\|_\infty := \max_{1 \leq j \leq d} |x_j|$. For any subset
279 $S \subseteq \mathbb{R}^d$ we denote by $B_r^p(x; S)$ the set $B_r^p(x) \cap S$.

280 A. Passive and Active Query Models

281 Let U be a known random quantity defined on a probability
282 space \mathcal{U} . The following definitions characterize all passive and
283 active optimization algorithms:

284 **Definition 1** (The passive query model). *Let x_1, \dots, x_n be
285 i.i.d. points uniformly sampled on \mathcal{X} and y_1, \dots, y_n be obser-
286 vations from the model (1). A passive optimization algorithm
287 \mathcal{A} with n queries is parameterized by a mapping $\phi_n :
288 (x_1, y_1, \dots, x_n, y_n, U) \mapsto \hat{x}_n$ that maps the i.i.d. obser-
289 vations $\{(x_i, y_i)\}_{i=1}^n$ to an estimated optimum $\hat{x}_n \in \mathcal{X}$, potentially
290 randomized by U .*

291 **Definition 2** (The active query model). *An active opti-
292 mization algorithm can be parameterized by mappings
293 $(\chi_1, \dots, \chi_n, \phi_n)$, where for $t = 1, \dots, n$,*

$$294 \chi_t : (x_1, y_1, \dots, x_{t-1}, y_{t-1}, U) \mapsto x_t$$

295 produces a query point $x_t \in \mathcal{X}$ based on previous observations
 296 $\{(x_i, t_i)\}_{i=1}^{t-1}$, and

$$297 \quad \phi_n : (x_1, y_1, \dots, x_n, y_n, U) \mapsto \hat{x}_n$$

298 produces the final estimate. All mappings $(\chi_1, \dots, \chi_n, \phi_n)$ can
 299 be randomized by U .

300 B. Local Minimax Rates

301 We use a classical *local minimax analysis* [46] to understand
 302 the fundamental information-theoretic limits of noisy global
 303 optimization of smooth functions. On the upper bound side,
 304 we seek (active) estimators \hat{x}_n such that

$$305 \quad \sup_{f_0 \in \Theta} \sup_{f \in \Theta', \|f - f_0\|_\infty \leq \varepsilon_n(f_0)} \Pr_f [\mathcal{L}(\hat{x}_n; f) \geq C_1 \cdot R_n(f_0)] \leq 1/4, \quad (4)$$

307 where $C_1 > 0$ is a positive constant. Here $f_0 \in \Theta$ is
 308 referred to as the *reference function*, and $f \in \Theta'$ is the true
 309 underlying function to be optimized, which is assumed to be
 310 “near” f_0 (in the ℓ_∞ norm). The minimax convergence rate
 311 of $\mathcal{L}(\hat{x}_n; f)$ is then characterized *locally* by $R_n(f_0)$ which
 312 depends on the reference function f_0 . The constant of $1/4$
 313 is chosen arbitrarily and any small constant leads to similar
 314 conclusions. To establish negative results (i.e., local minimax
 315 lower bounds), in contrast to the upper bound formulation,
 316 we assume the potential active optimization estimator \hat{x}_n has
 317 *perfect knowledge* about the reference function $f_0 \in \Theta$.
 318 We then prove local minimax lower bounds of the form

$$319 \quad \inf_{\hat{x}_n} \sup_{f \in \Theta', \|f - f_0\|_\infty \leq \varepsilon_n(f_0)} \Pr_f [\mathcal{L}(\hat{x}_n; f) \geq C_2 \cdot R_n(f_0)] \geq 1/3, \quad (5)$$

321 where $C_2 > 0$ is another positive constant and $\varepsilon_n(f_0)$, $R_n(f_0)$
 322 are desired local convergence rates for functions near the
 323 reference f_0 .

324 Although in some sense classical, the local minimax definition
 325 we propose warrants further discussion:

326 1) **Roles of Θ and Θ' :** The reference function f_0 and the
 327 true functions f are assumed to belong to different but
 328 closely related function classes Θ and Θ' . In particular,
 329 in our paper $\Theta \subseteq \Theta'$, meaning that less restrictive
 330 assumptions are imposed on the true underlying function
 331 f compared to those imposed on the reference function
 332 f_0 on which R_n and ε_n are based.

333 2) **Upper Bounds:** It is worth emphasizing that the
 334 estimator \hat{x}_n has no knowledge of the reference function
 335 f_0 . From the perspective of upper bounds, we can
 336 consider the simpler task of producing f_0 -dependent
 337 bounds (eliminating the second supremum) to instead
 338 study the (already interesting) quantity:

$$339 \quad \sup_{f_0 \in \Theta} \Pr_{f \in \Theta} [\mathcal{L}(\hat{x}_n; f_0) \geq C_1 R_n(f_0)] \leq 1/4.$$

340 As indicated above we maintain the double-supremum
 341 in the definition because fewer assumptions are imposed
 342 directly on the true underlying function f , and further
 343 because it allows to more directly compare our upper
 344 and lower bounds.

345 3) **Lower Bounds and the choice of the “localization
 346 radius” $\varepsilon_n(f_0)$:** Our lower bounds allow the estimator
 347 knowledge of the reference function (this makes
 348 establishing the lower bound more challenging). The
 349 lower bound in (5) implies that no estimator \hat{x}_n can
 350 effectively optimize a function f close to f_0 beyond the
 351 convergence rate of $R_n(f_0)$, even if perfect knowledge
 352 of the reference function f_0 is available a priori. The
 353 $\varepsilon_n(f_0)$ parameter that decides the “range” in which
 354 local minimax rates apply is taken to be on the same
 355 order as the actual local rate $R_n(f_0)$ in this paper.
 356 This is (up to constants) the smallest radius for which
 357 we can hope to obtain non-trivial lower-bounds: if we
 358 consider a much smaller radius than $R_n(f_0)$ then the
 359 trivial estimator which outputs the minimizer of the ref-
 360 erence function would achieve a faster rate than $R_n(f_0)$.
 361 On the other hand selecting the smallest possible radius
 362 makes establishing the lower bound most challenging
 363 but provides a refined picture of the complexity of
 364 zeroth-order optimization.

365 We remark that our primary motivation for the
 366 local-minimax analysis stems from the fact that for natural
 367 function classes the global-minimax rate for the optimization
 368 complexity is excessively pessimistic, while the local minimax
 369 analysis provides a more refined picture. In machine learning
 370 applications, there are several cases where the population risk
 371 is well-behaved (smooth, potentially non-convex) but we are
 372 only able to access/query the empirical risk which we want to
 373 minimize. Using standard concentration bounds the empirical
 374 risk and population risk are close, and the resulting problem
 375 is then to minimize the approximate-smooth empirical risk
 376 (see for instance [42], [47] for a more detailed discussion).

377 III. MAIN RESULTS

378 With this background in place we now turn our attention
 379 to our main results. We begin by collecting our assumptions
 380 about the true underlying function and the reference function
 381 in Section III-A. We state and discuss the consequences of
 382 our upper and lower bounds in Sections III-B and III-C
 383 respectively. We defer most technical proofs to Section V and
 384 turn our attention to our optimization algorithm in Section IV.

385 A. Assumptions

386 We first state and motivate assumptions that will be used.
 387 The first assumption states that f is locally Hölder smooth on
 388 its level-sets.

389 (A1) There exist constants $\kappa, \alpha, M, \zeta > 0$ such
 390 that f restricted to $\mathcal{X}_{f, \kappa, \zeta} := \{x \in \mathcal{X} : \inf_{z \in \mathcal{X}, \|z - x\|_\infty \leq \zeta} f(z) \leq f^* + \kappa\}$ belongs to the
 391 Hölder class $\Sigma^\alpha(M)$, meaning that f is k -times
 392 differentiable on $\mathcal{X}_{f, \kappa, \zeta}$ and furthermore for any
 393 $x, x' \in \mathcal{X}_{f, \kappa, \zeta}$,
 394

$$395 \quad \sum_{a_1 + \dots + a_d = k} \frac{|f^{(\alpha, k)}(x) - f^{(\alpha, k)}(x')|}{\|x - x'\|_\infty^{\alpha - k}} \leq M. \quad (6)$$

396 ³We use the ℓ_∞ -norm for convenience and it can be replaced by any
 397 equivalent vector norm.

396 Here $k = \lfloor \alpha \rfloor$ is the largest integer lower bounding α
 397 and $f^{(\alpha, j)}(x) := \partial^j f(x)/\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}$.

398 We use $\Sigma_\kappa^\alpha(M)$ to denote the class of all functions satisfying (A1). We remark that (A1) is weaker than the usual Hölder
 399 assumption in two ways. First, (6) only imposes stability
 400 conditions on the $\lfloor \alpha \rfloor$ -th order derivatives of the function f , in
 401 contrast to conditions involving all orders of derivatives in previous
 402 works [17], [45]. Second, (A1) only imposes the Hölder
 403 smoothness assumption on certain regions of \mathcal{X} , because
 404 regions with function values larger than $f^* + \kappa$ can be easily
 405 detected and removed by a pre-processing step, highlighting
 406 an important difference between optimization and ℓ_∞ -norm
 407 estimation. We give further details of the pre-processing step
 408 in Section IV-C.

410 Our next assumption concerns the “regularity” of the *level-sets* of the “reference” function f_0 . Define $L_{f_0}(\epsilon) := \{x \in$
 411 $\mathcal{X} : f_0(x) \leq f_0^* + \epsilon\}$ as the ϵ -level-set of f_0 , and
 412 $\mu_{f_0}(\epsilon) := \lambda(L_{f_0}(\epsilon))$ as the Lebesgue measure of $L_{f_0}(\epsilon)$,
 413 which we refer to as the *distribution function*. Define,
 414 $N(L_{f_0}(\epsilon), \delta)$ as the smallest number of ℓ_2 -balls of radius δ
 415 that cover $L_{f_0}(\epsilon)$. Then we make the following assumption:
 416 (A2) There exist constants $c_0 > 0$ and $C_0 > 0$ such that
 417 $N(L_{f_0}(\epsilon), \delta) \leq C_0[1 + \mu_{f_0}(\epsilon)\delta^{-d}]$ for all $\epsilon, \delta \in (0, c_0]$.
 418 We use Θ_C to denote all functions that satisfy (A2) with
 419 respect to parameters $\mathbf{C} = (c_0, C_0)$.

420 At a high-level, the regularity condition (A2) assumes that
 421 the level-sets are sufficiently “regular” such that covering them
 422 with small-radius balls does not require significantly larger
 423 total volume. For example, consider the perfectly regular case
 424 when $L_{f_0}(\epsilon)$ is the d -dimensional ℓ_2 ball of radius r : $L_{f_0}(\epsilon) =$
 425 $\{x \in \mathcal{X} : \|x - x^*\|_2 \leq r\}$. Clearly, $\mu_{f_0}(\epsilon) \asymp r^d$. In addition,
 426 the δ -covering number in ℓ_2 of $L_{f_0}(\epsilon)$ is on the order of $1 +$
 427 $(r/\delta)^d \asymp 1 + \mu_{f_0}(\epsilon)\delta^{-d}$, which satisfies the scaling in (A2).

428 When (A2) holds, uniform confidence intervals for f on
 429 its level-sets are easier to construct because little statistical
 430 efficiency is lost by slightly enlarging the level-sets so that
 431 complete (sufficiently small) d -dimensional cubes are con-
 432 tained in the enlarged level-sets. On the other hand, when
 433 regularity of level-sets fails to hold such nonparametric esti-
 434 mation can be very difficult or even impossible. As an extreme
 435 example, suppose the level-set $L_{f_0}(\epsilon)$ consists of n stand-alone
 436 and well-spaced points in \mathcal{X} : the Lebesgue measure of $L_{f_0}(\epsilon)$
 437 would be zero, but at least $\Omega(n)$ queries are necessary to
 438 construct uniform confidence intervals on $L_{f_0}(\epsilon)$. It is clear
 439 that such $L_{f_0}(\epsilon)$ violates (A2), because $N(L_{f_0}(\epsilon), \delta) \geq n$ as
 440 $\delta \rightarrow 0^+$ but $\mu_{f_0}(\epsilon) = 0$.

442 B. Upper Bound

443 The following theorem is our main result that provides
 444 an upper bound on the local minimax rate of noisy global
 445 optimization with active queries.

446 **Theorem 1.** For any $\alpha, M, \kappa, c_0, C_0 > 0$ and $f_0 \in \Sigma_\kappa^\alpha(M) \cap$
 447 Θ_C , where $\mathbf{C} = (c_0, C_0)$, define

$$448 \varepsilon_n^U(f_0) := \sup \left\{ \varepsilon > 0 : \varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \geq n/\log^\omega n \right\}, \quad (7)$$

449 where $\omega > 5 + d/\alpha$ is a large constant. Suppose also that
 450 $\varepsilon_n^U(f_0) \rightarrow 0$ as $n \rightarrow \infty$. Then for sufficiently large n ,

451 there exists an estimator \hat{x}_n with access to n active queries
 452 $x_1, \dots, x_n \in \mathcal{X}$, a constant $C_R > 0$ depending only
 453 on $\alpha, M, \kappa, c, c_0, C_0$ and a constant $\gamma > 0$ depending only
 454 on α and d such that

$$455 \sup_{f_0 \in \Sigma_\kappa^\alpha(M) \cap \Theta_C} \sup_{\substack{f \in \Sigma_\kappa^\alpha(M), \\ \|f - f_0\|_\infty \leq \varepsilon_n^U(f_0)}} \Pr_f [\mathcal{L}(\hat{x}_n, f) > \\ 456 C_R \log^\gamma n \cdot (\varepsilon_n^U(f_0) + n^{-1/2})] \leq 1/4. \quad (8)$$

457
 458 **Remark 1.** Unlike the (local) smoothness class $\Sigma_\kappa^\alpha(M)$,
 459 the additional function class Θ_C that encapsulates (A2) is
 460 imposed only on the “reference” function f_0 but not the
 461 true function f to be estimated. This makes the assumptions
 462 considerably weaker because the true function f may violate
 463 (A2) while our results remain valid.

464 **Remark 2.** The estimator \hat{x}_n does not require knowledge of
 465 parameters κ, c_0, C_0 or $\varepsilon_n^U(f_0)$, and automatically adapts to
 466 them, as shown in the next section. While the knowledge of
 467 smoothness parameters α and M is in general unavoidable
 468 in non-parametric regression (see [48]), in the zeroth-order
 469 optimization problem it is possible to adapt to α and M
 470 by running $O(\log^2 n)$ parallel sessions of \hat{x}_n on $O(\log n)$
 471 grids of α and M values, and then using $\Omega(n/\log^2 n)$
 472 single-point queries to decide on the location with the smallest
 473 function value. This adaptive strategy was suggested in [22]
 474 to remove an additional condition in [21], and also applies to
 475 our setting.

476 **Remark 3.** When the distribution function $\mu_{f_0}(\epsilon)$ does not
 477 change abruptly with ϵ the expression of $\varepsilon_n^U(f_0)$ can be
 478 significantly simplified. In particular, if for all $\epsilon \in (0, c_0]$ it
 479 holds that

$$480 \mu_{f_0}(\epsilon/\log n) \geq \mu_{f_0}(\epsilon)/[\log n]^{O(1)}, \quad (9)$$

481 then $\varepsilon_n^U(f_0)$ can be upper bounded as

$$482 \varepsilon_n^U(f_0) \leq [\log n]^{O(1)} \cdot \sup \left\{ \varepsilon > 0 : \varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \geq n \right\}. \quad (10)$$

483 If $\mu_{f_0}(\epsilon)$ scales polynomially with ϵ , i.e. $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$ for
 484 some constant $\beta \geq 0$, then (9) and (10) are both satisfied.

485 The quantity $\varepsilon_n^U(f_0) = \sup \{ \varepsilon > 0 : \varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \geq n/\log^\omega n \}$ is crucial in determining the convergence rate of
 486 optimization error of \hat{x}_n locally around the reference function
 487 f_0 . While the definition of $\varepsilon_n^U(f_0)$ is mostly implicit and
 488 involves solving an inequality involving the distribution function
 489 $\mu_{f_0}(\cdot)$, we remark that it admits a simple form when μ_{f_0}
 490 has a polynomial growth rate similar to a local Tsybakov noise
 491 condition [4], [49], as shown in the following proposition:

492 **Proposition 2.** Suppose $\mu_{f_0}(\epsilon) \lesssim \epsilon^\beta$ for some constant
 493 $\beta \in [0, 2 + d/\alpha]$. Then $\varepsilon_n^U(f_0) = \tilde{O}(n^{-\alpha/(2\alpha+d-\alpha\beta)})$.
 494 In addition, if $\beta \in [0, d/\alpha]$ then $\varepsilon_n^U(f_0) + n^{-1/2} \lesssim \varepsilon_n^U(f_0) =$
 495 $\tilde{O}(n^{-\alpha/(2\alpha+d-\alpha\beta)})$.

496 We remark that, following Proposition 1 of [45], α, β and d
 497 must satisfy the relationship that $\beta \leq d/\alpha$. Proposition 2 can

be easily verified by solving the system $\varepsilon^{-(2+d/\alpha)}\mu_{f_0}(\varepsilon) \geq n/\log^\omega n$ with the condition $\mu_{f_0}(\varepsilon) \lesssim \varepsilon^\beta$. We therefore omit its proof. The following two examples give some simple reference functions f_0 that satisfy the $\mu_{f_0}(\varepsilon) \lesssim \varepsilon^\beta$ condition in Proposition 2 with particular values of β .

Example 1. The constant function $f_0 \equiv 0$ satisfies (A1) through (A3) with $\beta = 0$.

Example 2. $f_0 \in \Sigma_\kappa^2(M)$ that is strongly convex⁴ satisfies (A1) through (A3) with $\beta = d/2$.

Example 1 is simple to verify, as the volume of level-sets of the constant function $f_0 \equiv 0$ exhibit a phase transition at $\varepsilon = 0$ and $\varepsilon > 0$. Consequently, $\beta = 0$ is the only parameter for which $\mu_{f_0}(\varepsilon) \lesssim \varepsilon^\beta$. Example 2 is more involved, and holds because the strong convexity of f_0 lower bounds the growth rate of f_0 when moving away from its minimum. We give a rigorous proof for Example 2 in the appendix. We also remark that f_0 does not need to be exactly strongly convex for $\beta = d/2$ to hold, and the example is valid for, e.g., piecewise strongly convex functions with a constant number of pieces too.

To best interpret the results in Theorem 1 and Proposition 2, it is instructive to compare the “local” rate $n^{-\alpha/(2\alpha+d-\alpha\beta)}$ with the baseline rate $n^{-\alpha/(2\alpha+d)}$, which can be attained by reconstructing f in sup-norm and applying Proposition 1. Since $\beta \geq 0$, the local convergence rate established in Theorem 1 is never slower, and the improvement compared to the baseline rate $n^{-\alpha/(2\alpha+d)}$ is dictated by β , which governs the growth rate of volume of level-sets of the reference function f_0 . In particular, for functions that grows fast when moving away from its minimum, the parameter β is large and therefore the local convergence rate around f_0 could be much faster than $n^{-\alpha/(2\alpha+d)}$.

Theorem 1 also implies concrete convergence rates for special functions considered in Examples 1 and 2. For the constant reference function $f_0 \equiv 0$, Example 1 and Theorem 1 yield that $R_n(f_0) \asymp n^{-\alpha/(2\alpha+d)}$, which matches the baseline rate $n^{-\alpha/(2\alpha+d)}$ and suggests that $f_0 \equiv 0$ is the worst-case reference function. This is intuitive, because $f_0 \equiv 0$ has a drastic level-set change at $\varepsilon \rightarrow 0^+$ and therefore small perturbations of f_0 result in changes to the optimal location. On the other hand, if f_0 is strongly smooth and convex as in Example 2, Theorem 1 leads to the bound of $R_n(f_0) \asymp n^{-1/2}$, which is significantly better than the $n^{-2/(4+d)}$ baseline rate⁵ and also matches existing works on zeroth-order optimization of convex functions [11]. The faster rate holds intuitively because strongly convex functions grow quickly when moving away from the minimum. An active query algorithm can focus most of its queries on the small level-sets of the underlying function, resulting in more accurate local function reconstruction and faster optimization error rate.

Our proof of Theorem 1 is constructive, by upper bounding the local minimax optimization error of an explicit algorithm.

⁴A twice differentiable function f_0 is strongly convex if there exists $\sigma > 0$ such that $\nabla^2 f_0(x) \succeq \sigma I, \forall x \in \mathcal{X}$.

⁵Note that f_0 being strongly smooth corresponds to $\alpha = 2$ in the local smoothness assumption.

Roughly, our algorithm partitions the n active queries evenly into $\log n$ epochs, and level-sets of f are estimated at the end of each epoch by comparing (uniform) confidence intervals on a dense grid on \mathcal{X} . It is then proved that the volume of the estimated level-sets contracts *geometrically*, until the target convergence rate $R_n(f_0)$ is attained. The algorithm is described in more detail in Section IV and the complete proof of Theorem 1 is in Section V-B.

C. Lower Bounds

We prove local minimax lower bounds that match the upper bounds in Theorem 1 up to logarithmic terms. As we remarked in Section II-B, in the local minimax lower bound formulation we assume the data analyst has full knowledge of the reference function f_0 , which makes the lower bounds stronger as more information is available a priori.

To facilitate such local minimax lower bounds, the following additional condition is imposed on the reference function f_0 of which the data analyst has perfect information.

(A2') There exist constants $c'_0, C'_0 > 0$ such that $M(L_{f_0}(\varepsilon), \delta) \geq C'_0 \mu_{f_0}(\varepsilon) \delta^{-d}$ for all $\varepsilon, \delta \in (0, c'_0]$, where $M(L_{f_0}(\varepsilon), \delta)$ is the maximum number of disjoint ℓ_2 balls of radius δ that can be packed into $L_{f_0}(\varepsilon)$.

We denote $\Theta'_{\mathbf{C}'}$ as the class of functions that satisfy (A2') with respect to parameters $\mathbf{C}' = (c'_0, C'_0) > 0$. Intuitively, (A2') can be regarded as a converse of (A2).

We are now ready to state our main negative result, which shows, from an information-theoretic perspective, that the upper bound in Theorem 1 is not improvable.

Theorem 2. Suppose $\alpha, c_0, C_0, c'_0, C'_0 > 0$ and $\kappa = \infty$. Denote $\mathbf{C} = (c_0, C_0)$ and $\mathbf{C}' = (c'_0, C'_0)$. For any $f_0 \in \Theta_{\mathbf{C}} \cap \Theta'_{\mathbf{C}'}$, define

$$\varepsilon_n^L(f_0) := \sup \left\{ \varepsilon > 0 : \varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \geq n \right\}. \quad (11)$$

Then there exists a constant $M > 0$ depending on α, d, \mathbf{C} and \mathbf{C}' such that, for any $f_0 \in \Sigma_\kappa^\alpha(M/2) \cap \Theta_{\mathbf{C}} \cap \Theta_{\mathbf{C}'}$,

$$\inf_{\hat{x}_n} \sup_{\substack{f \in \Sigma_\kappa^\alpha(M), \\ \|f - f_0\|_\infty \leq 2\varepsilon_n^L(f_0)}} \Pr_f \left[\mathcal{L}(\hat{x}_n; f) \geq \varepsilon_n^L(f_0) \right] \geq \frac{1}{3}. \quad (12)$$

Remark 4. We note in passing that for any f_0 and n it always holds that $\varepsilon_n^L(f_0) \leq \varepsilon_n^U(f_0)$.

Remark 5. If the distribution function $\mu_{f_0}(\varepsilon)$ satisfies (9) (i.e. it does not change too abruptly) in Remark 3, then $\varepsilon_n^L(f_0) \geq \varepsilon_n^U(f_0)/[\log n]^{O(1)}$. Consequently, the upper and lower bounds for these functions match up to logarithmic factors.

The following proposition derives an explicit expression for $\varepsilon_n^L(f_0)$ for reference functions whose distribution functions have a polynomial growth, which matches the upper bound in Proposition 2 up to $\log n$ factors. The proof of this Proposition is straightforward and is omitted.

Proposition 3. Suppose $\mu_{f_0}(\varepsilon) \gtrsim \varepsilon^\beta$ for some $\beta \in [0, 2 + d/\alpha]$. Then $\varepsilon_n^L(f_0) = \Omega(n^{-\alpha/(2\alpha+d-\alpha\beta)})$.

602 The following proposition additionally shows the existence
 603 of $f_0 \in \Sigma_\infty^\alpha(M) \cap \Theta_C \cap \Theta_{C'}$ that satisfies $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$ for
 604 any values of $\alpha > 0$ and $\beta \in [0, d/\alpha]$. Its proof is given in
 605 the Appendix.

606 **Proposition 4.** *Fix arbitrary $\alpha, M > 0$ and $\beta \in [0, d/\alpha]$.
 607 There exists $f_0 \in \Sigma_\kappa^\alpha(M) \cap \Theta_C \cap \Theta_{C'}$ for $\kappa = \infty$ and constants
 608 $C = (c_0, C_0)$, $C' = (c'_0, C'_0)$ that depend only on α, β, M and
 609 d such that $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$.*

610 Theorem 2 and Proposition 3 show that the $n^{-\alpha/(2\alpha+d-\alpha\beta)}$
 611 upper bound on local minimax convergence rate established in
 612 Theorem 1 is not improvable up to logarithmic factors of n .
 613 Such information-theoretic lower bounds on the convergence
 614 rates hold *even if the data analyst has perfect information of*
 615 f_0 , the reference function on which the $n^{-\alpha/(2\alpha+d-\alpha\beta)}$ local
 616 rate is based. Our results also imply an $n^{-\alpha/(2\alpha+d)}$ minimax
 617 lower bound over all α -Hölder smooth functions, showing that
 618 without additional assumptions, noisy optimization of smooth
 619 functions is as difficult as reconstructing the unknown function
 620 in sup-norm.

621 Our proof of Theorem 2 also differs from those of existing
 622 minimax lower bounds for active nonparametric models [50].
 623 The classical approach is to invoke Fano's inequality and to
 624 upper bound the KL divergence between different underlying
 625 functions f and g using $\|f - g\|_\infty$, corresponding to the
 626 point $x \in \mathcal{X}$ that leads to the largest KL divergence. Such
 627 an approach, however, does not produce tight lower bounds
 628 for our problem. To overcome such difficulties, we borrow
 629 the lower bound analysis for bandit pure exploration problems
 630 in [51]. In particular, our analysis considers the query distribution
 631 of any active query algorithm $\mathcal{A} = (\varphi_1, \dots, \varphi_n, \phi_n)$
 632 under the reference function f_0 and bounds the perturbation in
 633 query distributions between f_0 and f using Le Cam's lemma.
 634 Afterwards, an adversarial function choice f can be made
 635 based on the query distributions of the considered algorithm \mathcal{A} .
 636 We defer the complete proof of Theorem 2 to Section V-C.

637 Theorem 2 applies to any global optimization method that
 638 makes *active* queries, corresponding to the query model in
 639 Definition 2. The following theorem, on the other hand, shows
 640 that for passive algorithms (Definition 1) the $n^{-\alpha/(2\alpha+d)}$
 641 optimization rate is not improvable even with additional level-set
 642 assumptions imposed on f_0 . This demonstrates an explicit
 643 gap between passive and adaptive query models in global
 644 optimization problems.

645 **Theorem 3.** *Suppose $\alpha, c_0, C_0, c'_0, C'_0 > 0$ and $\kappa = \infty$.
 646 Denote $C = (c_0, C_0)$ and $C' = (c'_0, C'_0)$. Then there exist
 647 constants $M > 0$ depending on α, d, C, C' and N depending
 648 on M such that, for any $f_0 \in \Sigma_\kappa^\alpha(M/2) \cap \Theta_C \cap \Theta_{C'}$ satisfying
 649 $\varepsilon_n^L(f_0) \leq \tilde{\varepsilon}_n^L =: [\log n/n]^{\alpha/(2\alpha+d)}$,*

$$650 \inf_{\hat{x}_n} \sup_{\substack{f \in \Sigma_\kappa^\alpha(M), \\ \|f - f_0\|_\infty \leq 2\tilde{\varepsilon}_n^L}} \Pr_f \left[\mathcal{L}(\hat{x}_n; f) \geq \tilde{\varepsilon}_n^L \right] \geq \frac{1}{3} \quad \text{for all } n \geq N. \quad (13)$$

652 Intuitively, the apparent gap demonstrated by Theorems 2
 653 and 3 between the active and passive query models stems from

654 the observation that, a passive algorithm \mathcal{A} only has access
 655 to uniformly sampled query points x_1, \dots, x_n and therefore
 656 cannot focus on a small level-set of f in order to improve
 657 query efficiency. In addition, for functions that grow faster
 658 when moving away from their minima (implying a larger
 659 value of β), the gap between passive and active query models
 660 becomes bigger as active queries can more effectively exploit
 661 the restricted level-sets of such functions.

IV. OUR ALGORITHM

662 In this section we describe a concrete algorithm that attains
 663 the upper bound in Theorem 1. We start with a cleaner
 664 algorithm that operates under the slightly stronger condition
 665 that $\kappa = \infty$ in (A1), meaning that f is α -Hölder smooth on the
 666 entire domain \mathcal{X} . The generalization to $\kappa > 0$ being a constant
 667 is given in Section IV-C with an additional pre-processing step.

668 Let $G_n \in \mathcal{X}$ be a *finite* grid of points in \mathcal{X} . We assume the
 669 finite grid G_n satisfies the following two mild conditions:

670 (B1) Points in G_n are sampled i.i.d. from an unknown distribution P_X on \mathcal{X} ; furthermore, the density p_X associated
 671 with P_X satisfies $\underline{p}_0 \leq p_X(x) \leq \bar{p}_0$ for all $x \in \mathcal{X}$, where
 672 $0 < \underline{p}_0 \leq \bar{p}_0 < \infty$ are universal constants;
 673 (B2) $|G_n| \gtrsim n^3$ and $\log |G_n| = O(\log n)$.

674 **Remark 6.** *Although typically the choices of the grid points
 675 G_n belong to the data analyst, in some applications the choices
 676 of design points are not completely unconstrained. For exam-
 677 ple, in material synthesis experiments described previously
 678 some environmental parameter settings (e.g., temperature and
 679 pressure) might not be allowed due to budget or physical con-
 680 straints. Thus, we choose to consider less restrictive conditions
 681 imposed on the design grid G_n , allowing it to be more flexible
 682 in real-world applications.*

683 **Remark 7.** *Condition (B2) ensures that the grid G_n is
 684 sufficiently dense, such that even with the smallest bandwidth
 685 our algorithm possibly uses $(h_t(x) = 1/n^2$, see (18)), each
 686 $x \in G_n$ has abundant neighboring points in G_n , so that the
 687 local polynomial estimates in (15) are well-defined.*

688 For any subset $S \subseteq G_n$ and a “weight” function
 689 $\varrho : G_n \rightarrow \mathbb{R}^+$, define the extension $S^\circ(\varrho)$ of S with respect
 690 to ϱ as

$$691 S^\circ(\varrho) := \bigcup_{x \in S} B_{\varrho(x)}^\infty(x; G_n) \quad \text{where} \\ 692 B_{\varrho(x)}^\infty(x; G_n) = \{z \in G_n : \|z - x\|_\infty \leq \varrho(x)\}. \quad (14)$$

693 The algorithm can then be formulated as two levels of iter-
 694 ations, with the outer loop shrinking the “active set” S_τ and
 695 the inner loop collecting data in order to reduce the lengths
 696 of the confidence intervals on the points in the active set.
 697 A pseudocode description of our proposed algorithm is given
 698 in Figure 1.

A. Local Polynomial Regression

701 We use local polynomial regression [5] to obtain the esti-
 702 mate \hat{f} . In particular, for any $x \in G_n$ and a bandwidth

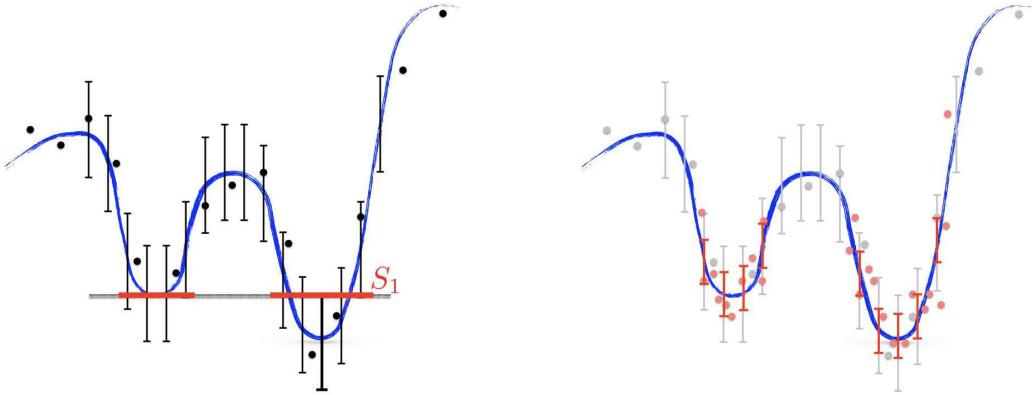


Fig. 1. An informal illustration of Algorithm 1. Solid blue curves depict the underlying function f to be optimized, black and red solid dots denote the query points and their responses $\{(x_t, y_t)\}$, and black and red vertical line segments correspond to uniform confidence intervals on function evaluations constructed using the current batch of data observed. The left figure illustrates the first epoch of our algorithm, where query points are uniformly sampled from the entire domain \mathcal{X} . Afterwards, sub-optimal locations based on the constructed confidence intervals are removed, and a shrunken ‘candidate set’ S_1 is obtained. The algorithm then proceeds to the second epoch, illustrated in the right figure, where query points (in red) are sampled only from the restricted candidate set and shorter confidence intervals (also in red) are constructed and updated. The procedure is repeated until $O(\log n)$ epochs are completed.

Parameters: α, M, δ, n

Output: \hat{x}_n , the final prediction

Initialization: $S_0 = G_n, \varrho_0(x) \equiv \infty, T = \lfloor \log_2 n \rfloor, n_0 = \lfloor n/T \rfloor$;

for $\tau = 1, 2, \dots, T$ **do**

 Compute ‘‘extended’’ sample set $S_{\tau-1}^o(\varrho_{\tau-1})$ defined in (14);

for $t = (\tau - 1)n_0 + 1$ to τn_0 **do**

 Sample x_t uniformly at random from $S_{\tau-1}^o(\varrho_{\tau-1})$ and observe $y_t = f(x_t) + w_t$;

end

 For every $x \in S_{\tau-1}$, compute bandwidth $h_\tau(x)$ using (18) and build the confidence interval

$[\ell_\tau(x), u_\tau(x)]$ in (19);

$S_\tau := \{x \in S_{\tau-1} : \ell_\tau(x) \leq \min_{x' \in S_{\tau-1}} u_\tau(x')\},$

$\varrho_\tau(x) := \min\{\varrho_{\tau-1}(x), h_\tau(x)\},$

end

Final processing: for every $x \in S_T$ let $\hat{f}_{h_T, x}(\cdot)$ be the local polynomial estimates constructed in (15) at x .

Output $\hat{x}_n = \arg \min_{x \in S_T} \min_{x' \in B_{h_T}^\infty(x; \mathcal{X})} \hat{f}_{h_T, x}(x')$.

Algorithm 1 The Main Algorithm

704 parameter $h > 0$, consider the least squares polynomial
705 estimate

$$706 \hat{f}_h \in \underset{g \in \mathcal{P}_k}{\operatorname{argmin}} \sum_{t'=1}^t \mathbb{I}[x_{t'} \in B_h^\infty(x)] \cdot (y_{t'} - g(x_{t'}))^2, \quad (15)$$

707 where $B_h^\infty(x) := \{x' \in \mathcal{X} : \|x' - x\|_\infty \leq h\}$ and \mathcal{P}_k denotes
708 all polynomials of degree k on \mathcal{X} .

709 To analyze the performance of \hat{f}_h evaluated at a certain
710 point $x \in \mathcal{X}$, define the mapping

$$711 \psi_{x, h} : z \mapsto (1, \psi_{x, h}^1(z), \dots, \psi_{x, h}^k(z))$$

712 where $\psi_{x, h}^j : z \mapsto [\prod_{\ell=1}^j h^{-1}(z_{i_\ell} - x_{i_\ell})]_{i_1, \dots, i_j=1}^d$ is the
713 degree- j polynomial mapping from \mathbb{R}^d to \mathbb{R}^{d^j} . Also define
714 $\Psi_{t, h} := (\psi_{x, h}(x_{t'}))_{1 \leq t' \leq t, x_{t'} \in B_h(x)}$ as the $m \times D$ aggregated

715 design matrix, where $m = \sum_{t'=1}^t \mathbb{I}[x_{t'} \in B_h^\infty(x)]$ and $D =$
716 $1 + d + \dots + d^k, k = \lfloor \alpha \rfloor$. The estimate \hat{f}_h defined in (15)
717 then admits the following closed-form expression:

$$718 \hat{f}_h(z) \equiv \psi_{x, h}(z)^\top (\Psi_{t, h}^\top \Psi_{t, h})^\dagger \Psi_{t, h}^\top Y_{t, h}, \quad (16)$$

719 where $Y_{t, h} = (y_{t'})_{1 \leq t' \leq t, x_{t'} \in B_h^\infty(x)}$ and A^\dagger is the
720 Moore-Penrose pseudo-inverse of A .

721 The following lemma gives a finite-sample analysis of the
722 error of $\hat{f}_h(x)$:

723 **Lemma 1.** Suppose that f satisfies (6) on $B_h^\infty(x; \mathcal{X})$,
724 $\max_{z \in B_h^\infty(x; \mathcal{X})} \|\psi_{x, h}(z)\|_2 \leq b$ and $\frac{1}{m} \Psi_{t, h}^\top \Psi_{t, h} \geq \sigma I_{D \times D}$ for
725 some $\sigma > 0$. Then for any $\delta \in (0, 1/2)$, with probability $1 - \delta$

$$726 |\hat{f}_h(x') - f(x')| \leq \underbrace{\frac{b^2}{\sigma} M d^k h^\alpha}_{b_{h, \delta}(x)} + \underbrace{b \sqrt{\frac{5D \ln(1/\delta)}{\sigma m}}}_{s_{h, \delta}(x)} =: \eta_{h, \delta}(x), \quad \forall x' \in B_h^\infty(x; \mathcal{X}). \quad (17)$$

727 **Remark 8.** $b_{h, \delta}(x)$, $s_{h, \delta}(x)$ and $\eta_{h, \delta}(x)$ depend on x because
728 σ depends on $\Psi_{t, h}$, which further depends on the sample points
729 in the neighborhood $B_h^\infty(x; \mathcal{X})$ of x .

730 In the rest of the paper we define $b_{h, \delta}(x) := (b^2/\sigma) M d^k h^\alpha$
731 and $s_{h, \delta}(x) := b \sqrt{5D \ln(1/\delta)/\sigma m}$ as the bias and standard
732 deviation terms in the error of $\hat{f}_h(x)$, respectively. We also
733 denote $\eta_{h, \delta}(x) := b_{h, \delta}(x) + s_{h, \delta}(x)$ as the overall error
734 in $\hat{f}_h(x)$.

735 Notice that when bandwidth h increases, the bias term
736 $b_{h, \delta}(x)$ increases because of the h^α term; on the other hand,
737 with h increasing the local neighborhood $B_h^\infty(x; \mathcal{X})$ grows
738 and would potentially contain more samples, implying a larger
739 m and smaller standard deviation term $s_{h, \delta}(x)$. A careful
740 selection of the bandwidth h balances $b_{h, \delta}(x)$ and $s_{h, \delta}(x)$ and
741 yields appropriate confidence intervals on $f(x)$, and we turn
742 our attention to this in the next section.

745 **B. Bandwidth Selection and Confidence Intervals**

746 Given the expressions of bias $\mathfrak{b}_{h,\delta}(x)$ and standard deviation
 747 $\mathfrak{s}_{h,\delta}(x)$ in (17), the bandwidth $h_\tau(x) > 0$ at epoch τ and point
 748 x is selected as

$$749 \quad h_\tau(x) := \frac{j_\tau(x)}{n^2} \text{ where } j_\tau(x) := \arg \max \left\{ j \in \mathbb{N}^+ : j \leq n^2 : \mathfrak{b}_{j/n^2,\delta}(x) \leq \mathfrak{s}_{j/n^2,\delta}(x) \right\}. \quad (18)$$

751 More specifically, $h_\tau(x)$ is the largest positive value in an
 752 evenly spaced grid $\{j/n^2\}$ such that the bias of $\hat{f}_{h_\tau}(x)$ is
 753 smaller than its standard deviation. This bandwidth selection
 754 is in principle similar to the Lepski's method [52], with the
 755 exception that an upper bound on the bias for any bandwidth
 756 parameter is known and does not need to be estimated from
 757 data.

758 With the selection of bandwidth $h_\tau(x)$ at epoch τ and
 759 query point x , a confidence interval on $f(x')$ for all $x' \in$
 760 $B_{h_\tau(x)}^\infty(x; \mathcal{X})$ is constructed as

$$761 \quad \ell_\tau(x) := \max_{1 \leq t' \leq \tau} \sup_{x' \in B_{h_\tau(x)}^\infty(x; \mathcal{X})} \left\{ \hat{f}_{h_{t'}(x)}(x') - \eta_{h_{t'}(x), \delta}(x) \right\};$$

$$762 \quad u_\tau(x) := \min_{1 \leq t' \leq \tau} \inf_{x' \in B_{h_\tau(x)}^\infty(x; \mathcal{X})} \left\{ \hat{f}_{h_{t'}(x)}(x') + \eta_{h_{t'}(x), \delta}(x) \right\}. \quad (19)$$

763 Note that for any $x \in \mathcal{X}$, the lower confidence edge $\ell_\tau(x)$ is
 764 a non-decreasing function in τ and the upper confidence edge
 765 $u_\tau(x)$ is a non-increasing function in τ .

767 **C. Pre-processing**

768 We describe a pre-processing step that relaxes the smoothness
 769 condition from $\kappa = \infty$ to $\kappa = \Omega(1)$, meaning that only
 770 local smoothness of f around its minimum values is required.
 771 Let $n_0 = \lfloor n/\log n \rfloor$, x_1, \dots, x_{n_0} be points i.i.d. uniformly sam-
 772 pled from \mathcal{X} and y_1, \dots, y_{n_0} be their corresponding responses.
 773 For every grid point $x \in G_n$, perform the following:

- 774 1) Compute $\hat{f}_x(\cdot)$ as the local polynomial fits of all y_i
 775 corresponding to $\|x_i - x\|_\infty \leq n_0^{-1/2d} \log^3 n =: h_0$;
- 776 2) Compute $\bar{f}(x)$ as the sample average of all y_i corre-
 777 sponding to $\|x_i - x\|_\infty \leq h_0$;
- 778 3) Remove all $x \in G_n$ from S_0 if $\bar{f}(x) \geq$
 779 $\min_{z \in G_n} \inf_{z' \in B_{h_0}^\infty(z; \mathcal{X})} \hat{f}_z(z') + 1/\log n$.

780 **Remark 9.** The $1/\log n$ term in the removal condition $\bar{f}(x) \geq$
 781 $\min_{z \in G_n} \bar{f}(z) + 1/\log n$ is not important, and can be replaced
 782 with any sequence $\{\omega_n\}$ such that $\lim_{n \rightarrow \infty} \omega_n = 0$ and
 783 $\lim_{n \rightarrow \infty} \omega_n t = \infty$ for any $t > 0$. The readers are referred to
 784 the proof of Proposition 5 in the appendix for the motivation
 785 of this term as well as the selection of the pre-processing
 786 bandwidth h_0 .

787 To analyze the pre-processing step, we state the following
 788 proposition:

789 **Proposition 5.** Assume $f \in \Sigma_\kappa^\alpha(M)$ and let S'_0 be the screened
 790 grid after step 2 of the pre-processing procedure. Then for
 791 sufficiently large n , with probability $1 - O(n^{-1})$ we have

$$792 \quad B_{h_0}^\infty(x; \mathcal{X}) \cap L_f(\kappa/2) \neq \emptyset, \quad \forall x \in S'_0, \quad (20)$$

793 where $L_f(\kappa/2) = \{x \in \mathcal{X} : f(x) \leq f^* + \kappa/2\}$.

794 To interpret Proposition 5, note that for sufficiently large n ,
 795 $f \in \Sigma_\kappa^\alpha(M)$ implies f being α -Hölder smooth (i.e., f
 796 satisfies (6)) on $\bigcup_{x \in L_f(\kappa/2)} B_{h_0}^\infty(x; \mathcal{X})$, because $\kappa > 0$ is a
 797 constant and $h_0 \rightarrow 0$ as $n \rightarrow \infty$. Subsequently, the proposition
 798 shows that with high probability, the pre-processing step will
 799 remove all grid points in G_n in non-smooth regions of f ,
 800 while maintaining the global optimal solution. This justifies
 801 the pre-processing step for $f \in \Sigma_\kappa^\alpha(M)$, because f is smooth
 802 on the grid and its close neighborhood after pre-processing.

803 The proof of Proposition 5 uses the fact that the local
 804 mean estimation is large provided that all data points in the
 805 local mean estimator are large, regardless of their underlying
 806 smoothness. The complete proof of Proposition 5 is deferred
 807 to the Appendix.

808 **V. PROOFS OF MAIN THEOREMS**809 **A. Proof of Lemma 1**

810 Our proof closely follows the analysis of asymptotic con-
 811 vergence rates for series estimators in [53]. We further work
 812 out all constants in the error bounds to arrive at a com-
 813 pletely finite-sample result, which is then used to construct
 814 finite-sample confidence intervals.

815 We start with polynomial interpolation results for all
 816 Hölder smooth functions in $B_{h_t}^\infty(x; \mathcal{X})$.

817 **Lemma 2.** Suppose f satisfies (6) on $B_h^\infty(x; \mathcal{X})$. Then there
 818 exists $\tilde{f}_x \in \mathcal{P}_k$ such that

$$819 \quad \sup_{z \in B_h^\infty(x; \mathcal{X})} |f(z) - \tilde{f}_x(z)| \leq M d^k h^\alpha. \quad (21)$$

821 *Proof.* Consider

$$822 \quad \tilde{f}_x(z) := f(x) + \sum_{j=1}^k \sum_{\alpha_1+\dots+\alpha_d=j} \frac{\partial^j f(x)}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} \prod_{\ell=1}^d (z_\ell - x_\ell)^{\alpha_\ell}. \quad (22)$$

823 By Taylor expansion with Lagrangian remainders, there exists
 824 $\xi \in (0, 1)$ such that

$$825 \quad |\tilde{f}_x(z) - f(z)| \leq \sum_{\alpha_1+\dots+\alpha_d=k} |f^{(\alpha)}(x + \xi(z-x)) - f^{(\alpha)}(x)| \cdot \prod_{\ell=1}^d |z_\ell - x_\ell|^{\alpha_\ell}.$$

827 Because f satisfies (6) on $B_h^\infty(x; \mathcal{X})$, we have that $|f^{(\alpha)}(x +$
 828 $\xi(z-x)) - f^{(\alpha)}(x)| \leq M \cdot \|z-x\|_\infty^{\alpha-k}$. Also note that $|z_\ell - x_\ell| \leq$
 829 $\|z - x\|_\infty \leq h$ for all $z \in B_h^\infty(x; \mathcal{X})$. The lemma is thus
 830 proved. \square

831 Using (16), the local polynomial estimate \hat{f}_h can be written
 832 as $\hat{f}_h(z) \equiv \psi_{x,h}(z)^\top \hat{\theta}_h$, where

$$833 \quad \hat{\theta}_h = (\Psi_{t,h}^\top \Psi_{t,h})^{-1} \Psi_{t,h}^\top Y_{t,h}. \quad (23)$$

834 In addition, because $\tilde{f}_x \in \mathcal{P}_k$, there exists $\tilde{\theta} \in \mathbb{R}^D$
 835 such that $\tilde{f}_x(z) \equiv \psi_{x,h}(z)^\top \tilde{\theta}$. Denote also that $F_{t,h} :=$
 836 $(f(x_{t'}))_{1 \leq t' \leq t, x_{t'} \in B_h^\infty(x)}$, $\Delta_{t,h} := (f(x_{t'}) -$
 $f_x(x_{t'}))$

837 $1 \leq t' \leq t, x_{t'} \in B_h^\infty(x)$ and $W_{t,h} := (w_{t'})_{1 \leq t' \leq t, x_{t'} \in B_h^\infty(x)}$. (23) can
838 then be re-formulated as

$$839 \quad \hat{\theta}_h = (\Psi_{t,h}^\top \Psi_{t,h})^{-1} \Psi_{t,h}^\top \left[\Psi_{t,h} \tilde{\theta} + \Delta_{t,h} + W_{t,h} \right] \quad (24)$$

$$840 \quad = \tilde{\theta} + \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \left[\frac{1}{m} \Psi_{t,h}^\top (\Delta_{t,h} + W_{t,h}) \right]. \quad (25)$$

841 Because $\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \geq \sigma I_{D \times D}$ and $\sup_{z \in B_h^\infty(x)} \|\psi_{x,h}(z)\|_2 \leq b$, we have that
842

$$843 \quad \|\hat{\theta}_h - \tilde{\theta}\|_2 \leq \frac{b}{\sigma} \|\Delta_{t,h}\|_\infty + \left\| \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \frac{1}{m} \Psi_{t,h}^\top W_t \right\|_2. \quad (26)$$

844 Invoking Lemma 2 we have $\|\Delta_{t,h}\|_\infty \leq M d^k h^\alpha$. In addition,
845 because $W_t \sim \mathcal{N}_m(0, I_{m \times n})$, we have that

$$846 \quad \left\| \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \frac{1}{m} \Psi_{t,h}^\top W_t \right\| \sim \mathcal{N}_D \left(0, \frac{1}{m} \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \right). \quad (27)$$

848 Applying concentration inequalities for quadratic forms of
849 Gaussian random vectors (Lemma 10), with probability $1 - \delta$
850 it holds that

$$851 \quad \left\| \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \frac{1}{m} \Psi_{t,h}^\top W_t \right\|_2 \leq \sqrt{\frac{5D \log(1/\delta)}{\sigma m}}. \quad (28)$$

852 We then have that with probability $1 - \delta$ that

$$853 \quad \|\hat{\theta}_h - \tilde{\theta}\|_2 \leq \frac{b}{\sigma_h} M d^k h^\alpha + \sqrt{\frac{5D \log(1/\delta)}{\sigma m}}. \quad (29)$$

854 Finally, noting that for all $x' \in B_h^\infty(x; \mathcal{X})$, $\|\psi_{x,h}(x')\|_2 \leq b$
855 by definition, we have that

$$856 \quad |\hat{f}_h(x') - f(x')| = |\hat{f}_h(x) - \tilde{f}_x(x')| \\ 857 \quad = |\psi_{x,h}(x')^\top (\hat{\theta}_h - \tilde{\theta})| \leq b \|\hat{\theta}_h - \tilde{\theta}\|_2,$$

858 which completes the proof of Lemma 1.

859 *B. Proof of Theorem 1*

860 In this section we prove Theorem 1. We prove the theorem
861 by considering every reference function $f_0 \in \Sigma_\kappa^\alpha(M) \cap \Theta_C$
862 separately. For simplicity, we assume $\kappa = \infty$ throughout
863 the proof. The $0 < \kappa < \infty$ can be handled by replacing
864 \mathcal{X} with S_0 which is the grid after the pre-processing step
865 described in Section IV-C. We also suppress dependency on
866 $d, \alpha, M, \mathbf{C}, \underline{p}_0, \bar{p}_0$ in $O(\cdot), \Omega(\cdot), \Theta(\cdot), \gtrsim, \lesssim$ and \asymp notations.
867 We further suppress logarithmic terms of n in $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$
868 notations.

869 The following lemma is our main lemma, which shows
870 that the active set S_τ in our proposed algorithm shrinks
871 geometrically before it reaches a certain level. To simplify
872 notations, denote $\tilde{c}_0 := 10c_0$ and (A2) then hold for all
873 $\epsilon, \delta \in [0, \tilde{c}_0]$ for all $f_0 \in \Theta_C$.

874 **Lemma 3.** For $\tau = 1, \dots, T$ define $\varepsilon_\tau := \max\{\tilde{c}_0 \cdot 2^{-\tau}, C_3[\varepsilon_n^U(f_0) + n^{-1/2} \log^2 n]\}$, where $C_3 > 0$ is a constant depending only on $d, \alpha, M, \underline{p}_0, \bar{p}_0$ and \mathbf{C} . Denote also $\rho_\tau^* := \max_{x \in S_\tau} \varrho_\tau(x)$. Then for sufficiently large n , with

875 probability $1 - O(n^{-1})$ the following holds uniformly for all
876 outer iterations $\tau = 1, \dots, T$:

$$877 \quad B_{\rho_\tau^*}^\infty(x; \mathcal{X}) \cap L_f(\varepsilon_\tau) \neq \emptyset. \quad (30)$$

881 Lemma 3 shows that the level ε_τ in $L_f(\varepsilon_\tau)$ that
882 contains $S_{\tau-1}$ shrinks geometrically, until the condition $\varepsilon_\tau \geq C_3[\varepsilon_n^U(f_0) + n^{-1/2} \log^2 n]$ is violated. If the condition is
883 never violated, then at the end of the last epoch τ^* we
884 have $\varepsilon_{\tau^*} = O(n^{-1})$ because $\tau^* = \log n$. On the other
885 hand, because $S_\tau \subseteq S_{\tau-1}$ always holds, we have $\varepsilon_{\tau^*} \lesssim [C_3[\varepsilon_n^U(f_0) + n^{-1/2} \log^2 n]]^{1/2}$. Combining both cases we have that
886 $\varepsilon_{\tau^*} \lesssim [C_3[\varepsilon_n^U(f_0) + n^{-1/2} \log^2 n]]^{1/2} n + n^{-1}$. Theorem 1 is thus
887 proved.

888 In the rest of this section we prove Lemma 3. We need
889 several technical lemmas and propositions. Except for Proposition
890 6 that is straightforward, the proofs of the other technical
891 lemmas are deferred to the end of this section.

892 Denote $x_n^* := \arg\min_{x \in G_n} f(x)$ as the point on the grid G_n
893 with the smallest objective value. The following proposition
894 shows that with high probability, the confidence intervals
895 constructed in the algorithm are truthful and the successive
896 rejection procedure will never exclude the true optimizer of f
897 on G_n .

898 **Proposition 6.** Suppose $\delta = 1/n^4 |G_n|$. Then with probability
899 $1 - O(n^{-1})$ the following hold:

- 900 1) $f(x') \in [\ell_t(x), u_t(x)]$ for all $1 \leq t \leq n$ and $x \in G_n$,
901 $x' \in B_{h_t(x)}^\infty(x; \mathcal{X})$;
- 902 2) $x_n^* \in S_\tau$ for all $0 \leq \tau \leq n$.

903 *Proof.* The first property is true by applying the union bound
904 over all $t = 1, \dots, n$ and $x \in G_n$. The second property
905 then follows, because $\ell_t(x_n^*) \leq f(x_n^*)$ and $\min_{x \in S_{\tau-1}} u_t(x) \geq$
906 $f(x_n^*)$ for all τ . \square

907 The following lemma shows that every small box centered
908 around a certain sample point $x \in G_n$ contains a sufficient
909 number of sample points whose least eigenvalue can be
910 bounded with high probability under the polynomial mapping
911 $\psi_{x,h}$ defined in Section III-B.

912 **Lemma 4.** For any $x \in G_n$, $1 \leq m \leq n$ and $h > 0$, let $K_{h,m}^1(x), \dots, K_{h,m}^n(x)$ be n independent point sets, where
913 each point set consists of m points sampled i.i.d. uniformly at
914 random from $B_h^\infty(x; G_n) = G_n \cap B_h^\infty(x; \mathcal{X})$. With probability
915 $1 - O(n^{-1})$ the following holds true uniformly for all $x \in G_n$,
916 $h \in \{j/n^2 : j \in \mathbb{N}, j \leq n^2\}$ and $K_{h,m}^\ell(x)$, $\ell \in [n]$ as $n \rightarrow \infty$:

- 917 1) $\sup_{h>0} \sup_{z \in B_h^\infty(x)} \|\psi_{x,h}(z)\|_2 \asymp \Theta(1)$;
- 918 2) $|B_h^\infty(x; G_n)| \asymp h^d |G_n|$;
- 919 3) $\sigma_{\min}(K_{h,m}^\ell(x)) \asymp \Theta(1)$ for all $m \geq \Omega(\log^2 n)$ and
920 $m \leq |G_n|$, where $\sigma_{\min}(K_{h,m}^\ell(x))$ is the least eigenvalue
921 of $\frac{1}{m} \sum_{z \in K_{h,m}^\ell(x)} \psi_{x,h}(z) \psi_{x,h}(z)^\top$.

922 **Remark 10.** It is possible to improve the concentration result
923 in (48) using the strategies adopted in [35] based on sharper
924 Bernstein type concentration inequalities. Such improvements
925 are, however, not important in establishing the main results of
926 this paper.

The next lemma shows that, the bandwidth h_t selected at the end of each outer iteration τ is near-optimal, being sandwiched between two quantities determined by the size of the active sample grid $\tilde{S}_{\tau-1} := S_{\tau-1}^\circ(\rho_{\tau-1})$.

Lemma 5. *There exist constants $C_1, C_2 > 0$ depending only on $d, \alpha, M, \underline{p}_0, \bar{p}_0$ and \mathbf{C} such that with probability $1 - O(n^{-1})$, the following holds for every outer iteration $\tau \in \{1, \dots, T\}$ and all $x \in S_{\tau-1}$:*

$$C_1[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} - \tau/n \leq \varrho_\tau(x) \leq C_2[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} \log n + \tau/n, \quad (31)$$

where $\tilde{v}_{\tau-1} := |G_n|/|\tilde{S}_{\tau-1}|$.

We are now ready to state the proof of Lemma 3, which is based on an inductive argument over the epochs $\tau = 1, \dots, T$.

Proof. We use induction to prove this lemma. For the base case $\tau = 1$, because $\|f - f_0\|_\infty \leq \varepsilon_n^U(f_0)$ and $\varepsilon_n^U(f_0) \rightarrow 0$ as $n \rightarrow \infty$, it suffices to prove that $B_{\rho_1^*}^\infty(x; \mathcal{X}) \cap L_{f_0}(\tilde{c}_0/4) \neq \emptyset$ for all $x \in S_1$ and sufficiently large n . Because $\tilde{S}_0 = S_0 = G_n$, invoking Lemmas 5 and 1 we have that $|\eta_{h_t(x), \delta}(x)| = \tilde{O}(n^{-\alpha/(2\alpha+d)})$ for all $x \in G_n$ with high probability at the end of the first outer iteration $\tau = 1$. Therefore, for sufficiently large n we conclude that $\sup_{x \in G_n} |\eta_{h_t(x), \delta}(x)| \leq c_0/16$ and hence $B_{\rho_1^*}^\infty(x; \mathcal{X}) \cap L_{f_0}(\tilde{c}_0/4) \neq \emptyset$ for all $x \in S_1$.

We now prove the lemma for $\tau \geq 2$, assuming it holds for $\tau - 1$. We also assume that n (and hence n_0) is sufficiently large, such that the maximum CI length $\max_{x \in G} |\eta_{h_t(x), \delta}(x)|$ after the first outer iteration $\tau = 1$ is smaller than $c_0/2$.

Because $\|f - f_0\|_\infty \leq \varepsilon_n^U(f_0)$ and $\varepsilon_{\tau-1} \geq C_3 \varepsilon_n^U(f_0) \log^2 n$, for appropriately chosen constant C_3 that is not too small, we have that $\|f - f_0\|_\infty \leq \varepsilon_{\tau-1}$. By the inductive hypothesis we have

$$\forall x \in S_{\tau-1}, \quad B_{\rho_{\tau-1}^*}^\infty(x; \mathcal{X}) \cap L_f(\varepsilon_{\tau-1}) \neq \emptyset;$$

Equivalently,

$$\begin{aligned} S_{\tau-1} &\subseteq L_f^\circ(\varepsilon_{\tau-1}, \rho_{\tau-1}^*) \subseteq L_{f_0}^\circ(\varepsilon_{\tau-1} + \|f - f_0\|_\infty, \rho_{\tau-1}^*) \\ &\subseteq L_{f_0}^\circ(2\varepsilon_{\tau-1}, \rho_{\tau-1}^*). \end{aligned} \quad (32)$$

Subsequently,

$$\tilde{S}_{\tau-1} = S_{\tau-1}^\circ \subseteq L_{f_0}^\circ(2\varepsilon_{\tau-1}, 2\rho_{\tau-1}^*). \quad (33)$$

Let $\bigcup_{x \in H_n} B_{2\rho_{\tau-1}^*}^2(x)$ be the smallest covering set of $L_{f_0}(2\varepsilon_{\tau-1})$, meaning that $L_{f_0}(2\varepsilon_{\tau-1}) \subseteq \bigcup_{x \in H_n} B_{2\rho_{\tau-1}^*}^2(x)$, where $B_{2\rho_{\tau-1}^*}^2(x) = \{z \in \mathcal{X} : \|z - x\|_2 \leq 2\rho_{\tau-1}^*\}$ is the ℓ_2 ball of radius $2\rho_{\tau-1}^*$ centered at x . By (A2), we know that $|H_n| \lesssim 1 + [\rho_{\tau-1}^*]^{-d} \mu_{f_0}(2\varepsilon_{\tau-1})$. In addition, the enlarged level-set satisfies $L_{f_0}^\circ(2\varepsilon_{\tau-1}, 2\rho_{\tau-1}^*) \subseteq \bigcup_{x \in H_n} B_{4\rho_{\tau-1}^*}^\infty(x)$. Subsequently,

$$\mu_{f_0}^\circ(2\varepsilon_{\tau-1}, \rho_{\tau-1}^*) \lesssim |H_n| \cdot [\rho_{\tau-1}^*]^d \lesssim \mu_{f_0}(2\varepsilon_{\tau-1}) + [\rho_{\tau-1}^*]^d. \quad (34)$$

By Lemma 5, the monotonicity of $|\tilde{S}_{\tau-1}|$ and the fact that $\underline{p}_0 \leq p_X(z) \leq \bar{p}_0$ for all $z \in \mathcal{X}$, we have

$$\rho_{\tau-1}^* \lesssim [\mu_f^\circ(\varepsilon_{\tau-1}, \rho_{\tau-1}^*)]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log n \quad (35)$$

$$\leq [\mu_{f_0}^\circ(2\varepsilon_{\tau-1}, \rho_{\tau-1}^*)]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log n \quad (36)$$

$$\lesssim \left(\mu_{f_0}(2\varepsilon_{\tau-1}) + [\rho_{\tau-1}^*]^d \right)^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log n. \quad (37)$$

Re-arranging terms on both sides of (37) we have

$$\rho_{\tau-1}^* \lesssim \max \left\{ [\mu_{f_0}(2\varepsilon_{\tau-1})]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log n, n_0^{-1/(2\alpha+d)} \log n \right\}. \quad (38)$$

On the other hand, according to the selection procedure of the bandwidth $h_t(x)$, we have that $\eta_{h_t(x), \delta}(x) \lesssim b_{h_t(x), \delta}(x)$. Invoking Lemma 5 we have for all $x \in S_{\tau-1}$ that

$$\eta_{h_t(x), \delta}(x) \lesssim b_{h_t(x), \delta}(x) \lesssim [h_t(x)]^\alpha \quad (39)$$

$$\lesssim [\tilde{v}_{\tau-1}n_0]^{-\alpha/(2\alpha+d)} \log n \quad (40)$$

$$\lesssim [\tilde{v}_{\tau-2}n_0]^{-\alpha/(2\alpha+d)} \log n \quad (41)$$

$$\lesssim [\rho_{\tau-1}^*]^\alpha \log n. \quad (42)$$

Here (40) holds by invoking the upper bound on $h_t(x)$ in Lemma 5, (41) holds because $\tilde{v}_{\tau-1} \geq \tilde{v}_{\tau-2}$, and (42) holds by again invoking the lower bound on $\varrho_{\tau-1}(x)$ in Lemma 5. Combining Eqs. (38,42) we have

$$\max_{x \in S_{\tau-1}} \eta_{h_t(x), \delta}(x) \quad (43)$$

$$\lesssim \max \left\{ [\mu_{f_0}(2\varepsilon_{\tau-1})]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log^2 n, n_0^{-1/(2\alpha+d)} \log n \right\}. \quad (44)$$

Recall that $n_0 = n/\log n$ and $\varepsilon_n^U(f_0) \leq \varepsilon_{\tau-1}$, provided that C_3 is not too small. By definition, every $\varepsilon \geq \varepsilon_n^U(f_0)$ satisfies $\varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \leq n/\log^\omega n$ for some large constant $\omega > 5 + d/\alpha$. Subsequently,

$$[\mu_{f_0}(2\varepsilon_{\tau-1})]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log^2 n \quad (45)$$

$$\lesssim 2\varepsilon_{\tau-1} n^{1/(2\alpha+d)} \log^{-\frac{\omega\alpha}{2\alpha+d}} n \cdot n_0^{-1/(2\alpha+d)} \log^2 n \quad (46)$$

$$\lesssim \varepsilon_{\tau-1} / [\log n]^{\frac{(5+d/\alpha)\omega}{2\alpha+d}}. \quad (47)$$

Because $\omega > 5 + d/\alpha$, the right-hand side of (46) is asymptotically dominated⁶ by $\varepsilon_{\tau-1}$. In addition, $n_0^{-1/(2\alpha+d)} \log n$ is also asymptotically dominated by $\varepsilon_{\tau-1}$ because $\varepsilon_{\tau-1} \geq C_3 n^{-1/2} \log^\omega n$. Therefore, for sufficiently large n we have

$$\max_{x \in S_{\tau-1}} \eta_{h_t(x), \delta}(x) \leq \varepsilon_{\tau-1}/4. \quad (48)$$

Lemma 3 is thus proved. \square

⁶We say $\{a_n\}$ is asymptotically dominated by $\{b_n\}$ if $\lim_{n \rightarrow \infty} |a_n|/|b_n| = 0$.

1011 *1) Proof of Lemma 4:*

1012 *Proof.* We first show that the first property holds almost surely.
1013 Recall the definition of $\psi_{x,h}$, we have that $1 \leq \|\psi_{x,h}(z)\|_2 \leq$
1014 $D \cdot [\max_{1 \leq j \leq d} h^{-1} |z_j - x_j|]^k$. Because $\|z - x\|_\infty \leq h$ for
1015 all $z \in B_h^\infty(x)$, $\sup_{z \in B_h^\infty(x)} \|\psi_{x,h}(z)\|_2 \lesssim O(1)$ for all $h > 0$.
1016 Thus, $\sup_{h>0} \sup_{z \in B_h^\infty(x)} \|\psi_{x,h}(z)\|_2 \asymp \Theta(1)$ for all $x \in G_n$.

1017 For the second property, by Hoeffding's inequality
1018 (Lemma 9) and the union bound, with probability $1 - O(n^{-1})$
1019 we have that

$$1020 \max_{x,h} \left| \frac{|B_h^\infty(x; G_n)|}{|G_n|} - P_X(z \in B_h^\infty(x)) \right| \lesssim \sqrt{\frac{\log n}{|G_n|}}. \quad (48)$$

1021 In addition, note that $P_X(z \in B_h^\infty(x; \mathcal{X})) \geq$
1022 $\underline{p}_0 \lambda(B_h^\infty(x; \mathcal{X})) \gtrsim h^d$ and $P_X(z \in B_h^\infty(x; \mathcal{X})) \leq$
1023 $\overline{p}_0 \lambda(B_h^\infty(x; \mathcal{X})) \lesssim h^d$, where $\lambda(\cdot)$ denotes the Lebesgue
1024 measure on \mathcal{X} . Subsequently, $|B_h^\infty(x; G_n)|$ is lower bounded by
1025 $\Omega(h^d |G_n| - \sqrt{|G_n| \log n})$ and upper bounded by
1026 $O(h^d |G_n| + \sqrt{|G_n| \log n})$. The second property is then
1027 proved by noting that $h_d \gtrsim n^{-d}$ and $|G_n| \gtrsim n^{3d/\min(\alpha, 1)}$.

1028 We next prove the third property. Because $\underline{p}_0 \leq p_X(z) \in \overline{p}_0$
1029 for all $z \in \mathcal{X}$, we have that

$$1030 \underline{p}_0 \int_{B_h^\infty(x; \mathcal{X})} \psi_{x,h}(z) \psi_{x,h}(z)^\top dU_{x,h}(z) \\ 1031 \leq \mathbb{E} \left[\frac{1}{m} \sum_{z \in K_{h,m}^\ell} \psi_{x,h}(z) \psi_{x,h}(z)^\top \right] \quad (49)$$

$$1032 \leq \overline{p}_0 \int_{B_h^\infty(x; \mathcal{X})} \psi_{x,h}(z) \psi_{x,h}(z)^\top dU_{x,h}(z), \quad (50)$$

1033 where $U_{x,h}$ is the uniform distribution on $B_h^\infty(x; \mathcal{X})$. Note
1034 also that

$$1035 \int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z) \\ 1036 \leq \int_{B_h^\infty(x; \mathcal{X})} \psi_{x,h}(z) \psi_{x,h}(z)^\top dU_{x,h}(z) \quad (51)$$

$$1037 \leq 2^d \int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z) \quad (52)$$

1038 where U is the uniform distribution on $\mathcal{X} = [0, 1]^d$. The
1039 following proposition upper and lower bounds the eigenvalues
1040 of $\int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z)$, which is proved in the appendix.

1041 **Proposition 7.** *There exist constants $0 < \psi_0 \leq \Psi_0 < \infty$
1042 depending only on d, D such that*

$$1043 \psi_0 I_{D \times D} \leq \int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z) \leq \Psi_0 I_{D \times D}. \quad (53)$$

1044 Using Proposition 7 and Eqs. (51,52), we conclude that

$$1045 \Omega(1) \cdot I_{D \times D} \leq \mathbb{E} \left[\frac{1}{m} \sum_{z \in K_{h,m}^\ell} \psi_{x,h}(z) \psi_{x,h}(z)^\top \right] \leq O(1) \cdot I_{D \times D}. \quad (54)$$

1046 Applying matrix Chernoff bound (Lemma 11) and the union
1047 bound, we have that with probability $1 - O(n^{-1})$,
1048

$$1049 \max_{x,h,m,\ell} \left\| \frac{1}{m} \sum_{z \in K_{h,m}^\ell(x)} \psi_{x,h}(z) \psi_{x,h}(z)^\top \right\|_{\text{op}} \\ 1050 - \mathbb{E} \left[\psi_{x,h}(z) \psi_{x,h}(z)^\top \mid z \in B_h^\infty(x) \right] \lesssim \sqrt{\frac{\log n}{m}}. \quad 1051$$

1052 Combining Eqs. (54,55) and applying Weyl's inequality
1053 (Lemma 12) we have

$$1054 \Omega(1) - O(\sqrt{\log n/m}) \lesssim \sigma_{\min}(K_{h,m}^\ell(x)) \\ 1055 \lesssim O(1) - O(\sqrt{\log n/m}). \quad 1055$$

1056 The third property is therefore proved. \square

1057 *2) Proof of Lemma 5: Proof.* We use induction to prove this
1058 lemma. For the base case of $\tau = 1$, we have $\tilde{S}_0 = S_0 = G_n$
1059 and therefore $\tilde{v}_{\tau-1} = 1$. Furthermore, applying Lemma 4 we
1060 have that for all $h = j/n^2$,

$$1061 \mathfrak{b}_{h,\delta}(x) \asymp h^\alpha, \quad \mathfrak{s}_{h,\delta}(x) \asymp \sqrt{\frac{\log n}{h^d n_0}}. \quad 1061$$

1062 Thus, for h selected according to (18) as the largest bandwidth
1063 of the form j/n^2 , $j \in \mathbb{N}$ such that $\mathfrak{b}_{h,\delta}(x) \leq \mathfrak{s}_{h,\delta}(x)$,
1064 both $\mathfrak{b}_{h,\delta}(x), \mathfrak{s}_{h,\delta}(x)$ are on the order of $n_0^{-1/(2\alpha+d)}$ up to
1065 logarithmic terms of n , and therefore one can pick appropriate
1066 constants $C_1, C_2 > 0$ such that $C_1 n_0^{-1/(2\alpha+d)} \leq \varrho_1(x) \leq$
1067 $C_2 n_0^{-1/(2\alpha+d)} \log n$ holds for all $x \in G_n$.

1068 We next prove the lemma for $\tau > 1$, assuming it holds
1069 for $\tau - 1$. We first establish the lower bound part. Define
1070 $\rho_{\tau-1}^* := \min_{z \in S_{\tau-1}} \varrho_{\tau-1}(z)$. By inductive hypothesis, $\rho_{\tau-1}^* \geq$
1071 $C_1 [\tilde{v}_{\tau-2} n_0]^{-1/(2\alpha+d)} - (\tau - 1)/n$. Note also that $\tilde{v}_{\tau-1} \geq \tilde{v}_{\tau-2}$
1072 because $\tilde{S}_{\tau-1} \subseteq \tilde{S}_{\tau-2}$, which holds because $S_{\tau-1} \subseteq S_{\tau-2}$
1073 and $\varrho_{\tau-1}(z) \leq \varrho_{\tau-2}(z)$ for all z . Let h_t^* be the smallest
1074 number of the form j_t^*/n^2 , $j_t^* \in [n^2]$ such that $h_t^* \geq$
1075 $C_1 [\tilde{v}_{\tau-1} n_0]^{-1/(2\alpha+d)} - \tau/n$. We then have $h_t^* \leq \rho_{\tau-1}^*$ and
1076 therefore query points in epoch τ are uniformly distributed in
1077 $B_{h_t^*}^\infty(x; G_n)$. Subsequently, applying Lemma 4 we have with
1078 probability $1 - O(n^{-1})$ that

$$1079 \mathfrak{b}_{h_t^*,\delta}(x) \leq C' [h_t^*]^\alpha, \quad \mathfrak{s}_{h_t^*,\delta}(x) \geq C'' \sqrt{\frac{\log n}{[h_t^*]^d \tilde{v}_{\tau-1} n}}, \quad 1079$$

1080 where $C', C'' > 0$ are constants that depend on
1081 $d, \alpha, M, \underline{p}_0, \overline{p}_0$ and \mathbf{C} , but not C_1, C_2, τ or h_t^* . By choosing
1082 C_1 appropriately (depending on C' and C'') we can make
1083 $\mathfrak{b}_{h_t^*,\delta}(x) \leq \mathfrak{s}_{h_t^*,\delta}(x)$ holds for all $x \in S_{\tau-1}$, thus establishing
1084 $\varrho_\tau(x) \geq \min\{\varrho_{\tau-1}(x), h_t^*\} \geq C_1 [\tilde{v}_{\tau-1} n_0]^{-1/(2\alpha+d)} - \tau/n$.

1085 We next prove the upper bound part. For any $h_t = j_t/n^2$
1086 where $j_t \in [n^2]$, invoking Lemma 4 we have that

$$1087 \mathfrak{b}_{h,\delta}(x) \geq \tilde{C}' h^\alpha, \quad \mathfrak{s}_{h,\delta}(x) \leq \tilde{C}'' \sqrt{\frac{\log n}{\min\{h, \rho_{\tau-1}^*\}^d \cdot \tilde{v}_{\tau-1} n}}, \quad 1087$$

1088

where \tilde{C}' and \tilde{C}'' are again constants depending on $d, \alpha, M, p_0, \bar{p}_0$ and \mathbf{C} , but not C_1, C_2 . Note also that $\rho_{\tau-1}^* \geq C_1[\tilde{v}_{\tau-2}n_0]^{-1/(2\alpha+d)} - (\tau-1)/n \geq C_1[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} - \tau/n$, because $\tilde{v}_{\tau-1} \geq \tilde{v}_{\tau-2}$. By selecting constant $C_2 > 0$ carefully (depending on \tilde{C}', \tilde{C}'' and C_1), we can ensure $b_{h,\delta}(x) > s_{h,\delta}(x)$ for all $h \geq C_2[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} + \tau/n$. Therefore, $\varrho_\tau(x) \leq h_\tau(x) \leq C_2[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} + \tau/n$. \square

1096 C. Proof of Theorem 2

1097 In this section we prove the main negative result in Theorem 2. To simplify presentation, we suppress dependency on 1098 α, d, c_0 and C_0 in $\lesssim, \gtrsim, \asymp, O(\cdot)$ and $\Omega(\cdot)$ notations. However, 1099 we do not suppress dependency on C_R or M in any of the 1100 above notations.

1101 Let $\varphi_0 : [-2, 2]^d \rightarrow \mathbb{R}^*$ be a non-negative function 1102 defined on \mathcal{X} such that $\varphi_0 \in \Sigma_{\kappa}^{[\alpha]}(1)$ with $\kappa = \infty$, 1103 $\sup_{x \in \mathcal{X}} \varphi_0(x) = \Omega(1)$ and $\varphi_0(z) = 0$ for all $\|z\|_2 \geq 1$. Here 1104 $[\alpha]$ denotes the smallest integer that upper bounds α . Such 1105 functions exist and are the cornerstones of the construction of 1106 information-theoretic lower bounds in nonparametric estimation 1107 problems [50]. One typical example is the ‘‘smoothstep’’ 1108 function (see for example [54])

$$1110 S_N(x) := \frac{1}{Z} x^{N+1} \sum_{n=0}^N \binom{N+n}{n} \binom{2N+1}{N-n} (-x)^n, \\ 1111 N = 0, 1, 2, \dots,$$

1112 where $Z > 0$ is a scaling parameter. The smoothstep function 1113 S_N is defined on $[0, 1]$ and satisfies the Hölder condition in (6) 1114 of order $\alpha = N$ on $[0, 1]$. It can be easily extended to $\tilde{S}_{N,d} : 1115 [-2, 2]^d \rightarrow \mathbb{R}$ by considering $\tilde{S}_{N,d}(x) := 1/Z - S_N(a\|x\|_1)$ 1116 where $\|x\|_1 = |x_1| + \dots + |x_d|$ and $a = 1/(2d)$. It is easy 1117 to verify that, with Z chosen appropriately, $\tilde{S}_{N,d} \in \Sigma_{\infty}^N(1)$, 1118 $\sup_{x \in \mathcal{X}} \tilde{S}_{N,d}(x) = 1/Z = \Omega(1)$ and $\tilde{S}_{N,d}(z) = 0$ for all 1119 $\|z\|_2 \geq 1$, where $M > 0$ is a constant.

1120 For any $x \in \mathcal{X}$ and $h > 0$, define $\varphi_{x,h} : \mathcal{X} \rightarrow \mathbb{R}^*$ as

$$1121 \varphi_{x,h}(z) := \mathbb{I}[z \in B_h^{\infty}(x)] \cdot \frac{Mh^{\alpha}}{2} \varphi_0\left(\frac{z-x}{h}\right). \quad (59)$$

1122 It is easy to verify that $\varphi_{x,h} \in \Sigma_{\infty}^{\alpha}(M/2)$, and furthermore 1123 $\sup_{z \in \mathcal{X}} \varphi_{x,h}(z) \asymp Mh^{\alpha}$ and $\varphi_{x,h}(z) = 0$ for all $z \notin B_h^{\infty}(x)$.

1124 Let $L_{f_0}(\varepsilon_n^L(f_0))$ be the level-set of f_0 at $\varepsilon_n^L(f_0)$. Let $H_n \subseteq 1125 L_{f_0}(\varepsilon_n^L(f_0))$ be the largest packing set such that $B_h^{\infty}(x)$ are 1126 disjoint for all $x \in H_n$, and $\bigcup_{x \in H_n} B_h^{\infty}(x) \subseteq L_{f_0}(\varepsilon_n^L(f_0))$. 1127 By (A2') and the definition of $\varepsilon_n^L(f_0)$, we have that

$$1128 |H_n| \geq M(L_{f_0}(\varepsilon_n^L(f_0)), 2\sqrt{dh}) \\ 1129 \gtrsim \mu_{f_0}(\varepsilon_n^L(f_0)) \cdot h^{-d} \geq [\varepsilon_n^L(f_0)]^{2+d/\alpha} \cdot nh^{-d}. \quad (60)$$

1130 For any $x \in H_n$, construct $f_x : \mathcal{X} \rightarrow \mathbb{R}$ as

$$1131 f_x(z) := f_0(z) - \varphi_{x,h}(z). \quad (61)$$

1132 Let $\mathcal{F}_n := \{f_x : x \in H_n\}$ be the class of functions indexed 1133 by $x \in H_n$. Let also $h \asymp (\varepsilon_n^L(f_0)/M)^{1/\alpha}$ such that $\|\varphi_{x,h}\|_{\infty} = 1134 2\varepsilon_n^L(f_0)$. We then have that $\|f_x - f_0\|_{\infty} \leq 2\varepsilon_n^L(f_0)$ and $f_x \in 1135 \Sigma_{\infty}^{\alpha}(M)$, because $f_0, \varphi_{x,h} \in \Sigma_{\infty}^{\alpha}(M/2)$.

1136 The next lemma shows that, with n adaptive queries to the 1137 noisy zeroth-order oracle $y_t = f(x_t) + w_t$, it is information 1138 theoretically not possible to identify a certain f_x in \mathcal{F}_n with 1139 high probability.

1140 **Lemma 6.** Suppose $|\mathcal{F}_n| \geq 2$. Let $\mathcal{A}_n = (\chi_1, \dots, \chi_n, \phi_n)$ 1141 be an active optimization algorithm operating with a sample 1142 budget n , which consists of samplers $\chi_{\ell} : \{(x_i, y_i)\}_{i=1}^{\ell-1} \mapsto x_{\ell}$ 1143 and an estimator $\phi_n : \{(x_i, y_i)\}_{i=1}^n \mapsto \hat{f}_n \in \mathcal{F}_n$, both can be 1144 deterministic or randomized functions. Then

$$1145 \inf_{\mathcal{A}_n} \sup_{f_x \in \mathcal{F}_n} \Pr_{f_x} \left[\hat{f}_n \neq f_x \right] \geq \frac{1}{2} - \sqrt{\frac{n \cdot \sup_{f_x \in \mathcal{F}_n} \|f_x - f_0\|_{\infty}^2}{2|\mathcal{F}_n|}}. \quad (62)$$

1146 **Lemma 7.** There exists constant $M > 0$ depending on 1147 α, d, c_0, C_0 such that the right-hand side of (62) is lower 1148 bounded by $1/3$.

1149 Lemmas 6 and 7 are proved at the end of this section. 1150 Combining both lemmas and noting that for any distinct 1151 $f_x, f_{x'} \in \mathcal{F}_n$ and $z \in \mathcal{X}$, $\max\{\mathcal{L}(z; f_x), \mathcal{L}(z; f_{x'})\} \geq \varepsilon_n^L(f_0)$, 1152 we proved the minimax lower bound formulated in Theorem 2. 1153

1154 1) *Proof of Lemma 6:* Our proof is inspired by the negative 1155 result of multi-arm bandit pure exploration problems estab- 1156 lished in [51].

1157 *Proof.* For any $x \in H_n$, define

$$1158 n_x := \mathbb{E}_{f_0} \left[\sum_{i=1}^n \mathbb{I}[x \in B_h^{\infty}(x)] \right]. \quad (63)$$

1159 Because $B_h^{\infty}(x)$ are disjoint for $x \in H_n$, we have $\sum_{x \in H_n} n_x \leq 1160 n$. Also define, for every $x \in H_n$,

$$1161 \varphi_x := \Pr_{f_0} \left[\hat{f}_n = f_x \right]. \quad (64)$$

1162 Because $\sum_{x \in H_n} \varphi_x = 1$, by pigeonhole principle there is at 1163 most one $x \in H_n$ such that $\varphi_x > 1/2$. Let $x_1, x_2 \in H_n$ 1164 be the points that have the smallest and second smallest n_x . 1165 Then there exists $x \in \{x_1, x_2\}$ such that $\varphi_x \leq 1/2$ and 1166 $n_x \leq 2n/|\mathcal{F}_n|$. By Le Cam’s and Pinsker’s inequality (see, 1167 for example, [4]) we have that

$$1168 \Pr_{f_x} \left[\hat{f}_n = f_x \right] \leq \Pr_{f_0} \left[\hat{f}_n = f_x \right] + d_{\text{TV}}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n}) \quad (65)$$

$$1169 \leq \Pr_{f_0} \left[\hat{f}_n = f_x \right] + \sqrt{\frac{1}{2} \text{KL}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n})} \quad (66)$$

$$1170 = \varphi_x + \sqrt{\frac{1}{2} \text{KL}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n})} \quad (67)$$

$$1171 \leq \frac{1}{2} + \sqrt{\frac{1}{2} \text{KL}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n})}. \quad (68)$$

1172 It remains to upper bound KL divergence of the active 1173 queries made by \mathcal{A}_n . Using the standard lower bound analysis 1174 for active learning algorithms [50], [55] and the fact that

1175 $f_x \equiv f_0$ on $\mathcal{X} \setminus B_h^\infty(x)$, we have

$$1176 \quad \text{KL}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n}) = \mathbb{E}_{f_0, \mathcal{A}_n} \left[\log \frac{P_{f_0, \mathcal{A}_n}(x_{1:n}, y_{1:n})}{P_{f_x, \mathcal{A}_n}(x_{1:n}, y_{1:n})} \right] \quad (69)$$

$$1177 \quad = \mathbb{E}_{f_0, \mathcal{A}_n} \left[\log \frac{\prod_{i=1}^n P_{f_0}(y_i | x_i) P_{\mathcal{A}_n}(x_i | x_{1:(i-1)}, y_{1:(i-1)})}{\prod_{i=1}^n P_{f_x}(y_i | x_i) P_{\mathcal{A}_n}(x_i | x_{1:(i-1)}, y_{1:(i-1)})} \right] \quad (70)$$

$$1179 \quad = \mathbb{E}_{f_0, \mathcal{A}_n} \left[\log \frac{\prod_{i=1}^n P_{f_0}(y_i | x_i)}{\prod_{i=1}^n P_{f_x}(y_i | x_i)} \right] \quad (71)$$

$$1180 \quad = \mathbb{E}_{f_0, \mathcal{A}_n} \left[\sum_{x_i \in B_h^\infty(x)} \log \frac{P_{f_0}(y_i | x_i)}{P_{f_x}(y_i | x_i)} \right] \quad (72)$$

$$1181 \quad \leq n_x \cdot \sup_{z \in B_h^\infty(x; \mathcal{X})} \text{KL}(P_{f_0}(\cdot | z) \| P_{f_x}(\cdot | z)) \quad (73)$$

$$1182 \quad \leq n_x \cdot \|f_0 - f_x\|_\infty^2. \quad (74)$$

1183 Therefore,

$$1184 \quad \Pr_{f_x} \left[\hat{f}_x = f_x \right] \leq \frac{1}{2} + \sqrt{\frac{1}{4} n_x \varepsilon_n^2} \leq \frac{1}{2} + \sqrt{\frac{n \|f_x - f_0\|_\infty^2}{2 |\mathcal{F}_n|}}. \quad (75)$$

□

1186 2) Proof of Lemma 7:

1187 *Proof.* By construction, $n \sup_{f_x \in \mathcal{F}_x} \|f_x - f_0\|_\infty^2 \lesssim M^2 nh^{2\alpha}$ and $|\mathcal{F}_n| = |H_n| \gtrsim [\underline{C}_\varepsilon \varepsilon_n^\perp(f_0)]^{2+d/\alpha} nh^{-d}$. Note also that $h \asymp (\varepsilon/M)^{1/\alpha} \asymp (\underline{C}_\varepsilon \varepsilon_n^\perp(f_0)/M)^{1/\alpha}$ because $\|f_x - f_0\|_\infty = \varepsilon = \underline{C}_\varepsilon \varepsilon_n^\perp(f_0)$. Subsequently,

$$1191 \quad \frac{n \sup_{f_x \in \mathcal{F}_x} \|f_x - f_0\|_\infty^2}{2 |\mathcal{F}_n|} \lesssim \frac{n [\underline{C}_\varepsilon \varepsilon_n^\perp(f_0)]^2}{n [\underline{C}_\varepsilon \varepsilon_n^\perp(f_0)]^2 \cdot M^{d/\alpha}} = M^{-d/\alpha}. \quad (76)$$

1192 By choosing the constant $M > 0$ to be sufficiently large, the right-hand side of the above inequality is upper bounded by $1/36$. The lemma is thus proved. □

1193 D. Proof of Theorem 3

1194 The proof of Theorem 3 is similar to the proof of Theorem 2, but is much more standard by invoking the *Fano's inequality* [4]. In particular, adapting the Fano's inequality on any finite function class \mathcal{F}_n constructed we have the following lemma:

1202 **Lemma 8** (Fano's inequality). *Suppose $|\mathcal{F}_n| \geq 2$, and $\{(x_i, y_i)\}_{i=1}^n$ are i.i.d. random variables. Then*

$$1204 \quad \inf_{\hat{f}_x} \sup_{f_x \in \mathcal{F}_n} \Pr_{f_x} \left[\hat{f}_x \neq f_x \right] \\ 1205 \quad \geq 1 - \frac{\log 2 + n \cdot \sup_{f_x, f_{x'} \in \mathcal{F}_n} \text{KL}(P_{f_x} \| P_{f_{x'}})}{\log |\mathcal{F}_n|}, \quad (77)$$

1206 where P_{f_x} denotes the distribution of (x, y) under the law of f_x .

1208 Let \mathcal{F}_n be the function class constructed in the previous proof of Theorem 2, corresponding to the largest packing set H_n of $L_{f_0}(\varepsilon_n^\perp)$ such that $B_h^\infty(x)$ for all $x \in H_n$ are disjoint, where $h \asymp (\varepsilon_n^\perp/M)^{1/\alpha}$ such that $\|\varphi_{x,h}\|_\infty = 2\varepsilon_n^\perp$ for

1212 all $x \in H_n$. Because f_0 satisfies (A2'), we have that $|\mathcal{F}_n| = 1213 |H_n| \gtrsim \mu_{f_0}(\varepsilon_n^\perp) h^{-d}$. Under the condition that $\varepsilon_n^\perp(f_0) \leq \varepsilon_n^\perp$, it 1214 holds that $\mu_{f_0}(\varepsilon_n^\perp) \geq [\varepsilon_n^\perp]^{2+d/\alpha} n$. Therefore,

$$1215 \quad |\mathcal{F}_n| \gtrsim [\varepsilon_n^\perp]^{2+d/\alpha} \cdot nh^{-d} \gtrsim [\varepsilon_n^\perp]^2 \cdot n M^{d/\alpha}. \quad (78)$$

1216 Because $\log(n/\varepsilon_n^\perp) \gtrsim \log n$ and $M > 0$ is a constant, we have 1217 that $\log |\mathcal{F}_n| \geq c \log n$ for all $n \geq N$, where $c > 0$ is a constant 1218 depending only on α, d and $N \in \mathbb{N}$ is a constant depending on M . 1219

1220 Let U be the uniform distribution on \mathcal{X} . Because $x \sim U$ 1221 and $f_x \equiv f_{x'}$ on $\mathcal{X} \setminus B_h^\infty(x)$, we have that 1222

$$1222 \quad \text{KL}(P_{f_x} \| P_{f_{x'}}) = \frac{1}{2} \int_{\mathcal{X}} |f_x(z) - f_{x'}(z)|^2 dU(z) \quad (79)$$

$$1223 \quad \leq \frac{1}{2} \Pr_U [z \in B_h^\infty(x)] \cdot \|f_x - f_{x'}\|_\infty^2 \quad (80)$$

$$1224 \quad \leq \frac{1}{2} \lambda(B_h^\infty(x)) \cdot [\varepsilon_n^\perp]^2 \quad (81)$$

$$1225 \quad \lesssim h^d [\varepsilon_n^\perp]^2 \lesssim [\varepsilon_n^\perp]^{2+d/\alpha} / M^{d/\alpha}. \quad (82)$$

1226 By choosing M to be sufficiently large, the right-hand side of 1227 (77) can be lower bounded by an absolute constant. The 1228 theorem is then proved following the same argument as in the 1229 proof of Theorem 2.

1230 APPENDIX A 1231 SOME CONCENTRATION INEQUALITIES

1232 In this section, to ease readability of our paper, we provide 1233 some concentration inequalities and other standard results that 1234 we use extensively.

1235 **Lemma 9** ([56]). *Suppose X_1, \dots, X_n are i.i.d. random 1236 variables such that $a \leq X_i \leq b$ almost surely. Then for any 1237 $t > 0$,*

$$1238 \quad \Pr \left[\left| \frac{1}{n} \sum_{i=1}^n X_i - \mathbb{E}X \right| > t \right] \leq 2 \exp \left\{ - \frac{nt^2}{2(b-a)^2} \right\}.$$

1239 **Lemma 10** ([57]). *Suppose $x \sim \mathcal{N}_d(0, I_{d \times d})$ and let A be 1240 a $d \times d$ positive semi-definite matrix. Then for all $t > 0$, 1241*

$$1242 \quad \Pr \left[x^\top Ax > \text{tr}(A) + 2\sqrt{\text{tr}(A^2)t} + 2\|A\|_{\text{op}}t \right] \leq e^{-t}. \quad (83)$$

1244 **Lemma 11** ([58], simplified). *Suppose A_1, \dots, A_n are 1245 i.i.d. positive semidefinite random matrices of dimension d and 1246 $\|A_i\|_{\text{op}} \leq R$ almost surely. Then for any $t > 0$, 1247*

$$1247 \quad \Pr \left[\left\| \frac{1}{n} \sum_{i=1}^n A_i - \mathbb{E}A \right\|_{\text{op}} > t \right] \leq 2 \exp \left\{ - \frac{nt^2}{8R^2} \right\}.$$

1249 **Lemma 12** (Weyl's inequality). *Let A and $A + E$ 1250 be $d \times d$ matrices with $\sigma_1, \dots, \sigma_d$ and $\sigma'_1, \dots, \sigma'_d$ be 1251 their singular values, sorted in descending order. Then 1252 $\max_{1 \leq i \leq d} |\sigma_i - \sigma'_i| \leq \|E\|_{\text{op}}$. 1253*

1253 APPENDIX B
1254 ADDITIONAL PROOFS

1255 *Proof of Proposition 1.* Consider arbitrary $x^* \in \mathcal{X}$ such
1256 that $f(x^*) = \inf_{x \in \mathcal{X}} f(x)$. Then we have that $\mathcal{L}(\hat{x}_n; f) =$
1257 $f(\hat{x}_n) - f(x^*) \leq [\hat{f}_n(\hat{x}_n) + \|\hat{f}_n - f\|_\infty] - [\hat{f}_n(x^*) - \|\hat{f}_n -$
1258 $f\|_\infty] \leq 2\|\hat{f}_n - f\|_\infty$, where the last inequality holds because
1259 $\hat{f}_n(\hat{x}_n) \leq f_n(x^*)$ by optimality of \hat{x}_n . \square

1260 *Proof of Example 2.* Because $f_0 \in \Sigma_\kappa^2(M)$ is strongly convex,
1261 there exists $\sigma > 0$ such that $\nabla^2 f_0(x) \succeq \sigma I$ for all
1262 $x \in \mathcal{X}_{f_0, \kappa}$, where $\mathcal{X}_{f_0, \kappa} := L_{f_0}(\kappa)$ is the κ -level-set of f_0 .
1263 Let $x^* = \arg \min_{x \in \mathcal{X}} f_0(x)$, which is unique because f_0 is
1264 strongly convex. The smoothness and strong convexity of f_0
1265 implies that

$$1266 f_0^* + \frac{\sigma}{2} \|x - x^*\|_\infty^2 \leq f_0(x) \leq f_0^* + \frac{M}{2} \|x - x^*\|_\infty^2 \quad \forall x \in \mathcal{X}_{f_0, \kappa}. \quad (83)$$

1268 Subsequently, there exist constants $c_0, C_1, C_2 > 0$ depending
1269 only on σ, M, κ and d such that for all $\epsilon \in (0, c_0]$,

$$1270 B_{C_1 \sqrt{\epsilon}}^\infty(x^*; \mathcal{X}) \subseteq L_{f_0}(\epsilon) \subseteq B_{C_2 \sqrt{\epsilon}}^\infty(x^*; \mathcal{X}). \quad (84)$$

1271 The property $\mu_{f_0}(\epsilon) \lesssim \epsilon^\beta$ holds because $\mu(L_{f_0}(\epsilon)) \leq$
1272 $\mu(B_{C_1 \sqrt{\epsilon}}^\infty(x^*; \mathcal{X})) \lesssim \epsilon^{d/2}$. To prove (A2), note that
1273 $N(L_{f_0}(\epsilon), \delta) \leq N(B_{C_2 \sqrt{\epsilon}}^\infty(x^*; \mathcal{X}), \delta) \lesssim 1 + (\sqrt{\epsilon}/\delta)^d$. Because
1274 $\epsilon^{d/2} \lesssim \mu(L_{f_0}(\epsilon)) = \mu_{f_0}(\epsilon)$, we conclude that
1275 $N(L_{f_0}(\epsilon), \delta) \lesssim 1 + \delta^{-d} \mu_{f_0}(\epsilon)$ and (A2) is thus proved. \square

1276 *Proof of Proposition 4.* Consider $f_0 \equiv 0$ if $\beta = 0$ and $f_0(z) :=$
1277 $a_0 [z_1^p + \dots + z_d^p]$ for all $z = (z_1, \dots, z_d) \in [0, 1]^d$, where
1278 $a_0 > 0$ is a constant depending on α, M , and $p = d/\beta$ for
1279 $\beta \in (0, d/\alpha]$. The $\beta = 0$ case where $f_0 \equiv 0$ trivially holds.
1280 So we shall only consider the case of $\beta \in (0, d/\alpha]$.

1281 We first show $f_0 \in \Sigma_\kappa^\alpha(M)$ with $\kappa = \infty$, provided that a_0 is
1282 sufficiently small. For any $j \leq k = \lfloor \alpha \rfloor$ and $\alpha_1 + \dots + \alpha_d = j$,
1283 we have

$$1284 \frac{\partial^j}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} f_0(z) = \begin{cases} a_0 j! \cdot z_\ell^{p-j} & \text{if } \alpha_\ell = j, \ell \in [d]; \\ 0 & \text{otherwise.} \end{cases} \quad (85)$$

1286 Because $z_1, \dots, z_d \in [0, 1]$ and $p = d/\beta \geq \alpha \geq j$, it's clear
1287 that $0 \leq \partial^j f_0(z)/\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d} \leq a_0 j!$. In addition, for any
1288 $z, z' \in [0, 1]^d$ and $\alpha_\ell = k$, $\ell \in [d]$, we have

$$1289 \left| \frac{\partial^k}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} f_0(z) - \frac{\partial^k}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} f_0(z') \right| \leq a_0 k! \cdot |[z_\ell]^{p-k} - [z'_\ell]^{p-k}| \quad (86)$$

$$1291 \leq a_0 k! \cdot |z_\ell - z'_\ell|^{\min\{p-k, 1\}}, \quad (87)$$

1292 where the last inequality holds because x^t is $\min\{t, 1\}$ -Hölder
1293 continuous on $[0, 1]$ for $t \geq 0$. The $|z_\ell - z'_\ell|^{\min\{p-k, 1\}}$
1294 term can be further upper bounded by $\|z - z'\|_\infty^{\alpha-k}$, because
1295 $p = d/\beta \geq \alpha$. By selecting $a_0 > 0$ to be sufficiently small
1296 (depending on M) we have $f_0 \in \Sigma_\infty^\alpha(M)$.

1297 We next prove f_0 satisfies $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$ with parameter β
1298 depending on a_0 and p . For any $\epsilon > 0$, the level-set $L_{f_0}(\epsilon)$ can

1299 be expressed as $L_{f_0}(\epsilon) = \{z \in [0, 1]^d : z_1^p + \dots + z_d^p \leq \epsilon/a_0\}$.
1300 Subsequently,

$$1301 \left[0, \left(\frac{\epsilon}{a_0 d} \right)^{1/p} \right]^d \subseteq L_{f_0}(\epsilon) \subseteq \left[0, \left(\frac{\epsilon}{a_0} \right)^{1/p} \right]^d. \quad (88)$$

1302 Therefore,

$$1303 [\epsilon/(a_0 d)]^{dp} \leq \mu_{f_0}(\epsilon) \leq [\epsilon/a_0]^{dp}. \quad (89)$$

1304 Because a_0, d are constants and $dp = \beta$, we established
1305 $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$ for $\beta = dp$.

1306 Finally, note that for any $\epsilon > 0$, $L_{f_0}(\epsilon)$ is sandwiched
1307 between two cubics whose volumes only differ by a constant.
1308 This proves (A2) and (A2') on the covering and packing
1309 numbers of $L_{f_0}(\epsilon)$. \square

1310 *Proof of Proposition 5.* By the Chernoff bound and the
1311 union bound, with probability $1 - O(n^{-1})$ uniformly over all
1312 $x \in G_n$, there are $\Omega(\sqrt{n} \log^2 n)$ uniform samples in
1313 $B_{h_0}^\infty(x; \mathcal{X})$. Because $h_0 \leq \zeta$ for sufficiently large n_0 (ζ is
1314 defined in condition (A1)), by Lemma 1 it holds that

$$1315 |\check{f}_x(x') - f_x(x')| \lesssim h_0^\alpha + n_0^{-1/4} \lesssim n_0^{-\alpha/2d} + n_0^{-1/4}, \quad 1316 \forall x \in G_n, x' \in B_{h_0}^\infty(x; \mathcal{X}). \quad (90)$$

1317 Also, using the standard Gaussian concentration inequality,
1318 with probability $1 - O(n^{-1})$ we have

$$1319 \inf_{x' \in B_{h_0}^\infty(x; \mathcal{X})} f(x) - O(n_0^{-1/4}) \quad 1320 \leq \bar{f}(x) \leq \sup_{x' \in B_{h_0}^\infty(x; \mathcal{X})} f(x) + O(n_0^{-1/4}) \quad \forall x \in G_n. \quad (91)$$

1321 Let x^* be the minimizer of f on \mathcal{X} and $x \in G_n$ such
1322 that $\|x - x^*\|_\infty \leq h_0$. By (90), we have with probability
1323 $1 - O(n^{-1})$ that $\inf_{x' \in B_{h_0}^\infty(x; \mathcal{X})} \check{f}_x(x') \leq f^* + O(n_0^{-\alpha/2d} +$
1324 $n_0^{-1/4}) \leq f^* + 1/2 \log n$, where $f^* = f(x^*)$. Now consider
1325 arbitrary $z \in G_n$ such that $B_{h_0}^\infty(z; \mathcal{X}) \cap L_f(\kappa/2) = \emptyset$,
1326 meaning that for all $z' \in \mathcal{X}$, $\|z' - z\|_\infty \leq h_0$, $f(z') >$
1327 $\kappa/2$. By (90), $\bar{f}(z) \geq \kappa/2 - O(n_0^{-1/4}) \geq \kappa/2 - 1/2 \log n$.
1328 Hence when n_0 is sufficiently large, $z \notin S'_0$, which is to be
1329 demonstrated. \square

1330 *Proof of Proposition 7.* The upper bound part of (53) triv-
1331 ially holds because the absolute values of every element in
1332 $\psi_{0,1}(z)\psi_{0,1}(z)^\top$ for $z \in \mathcal{X} = [0, 1]^d$ is upper bounded by
1333 $O(1)$. To prove the lower bound part, we only need to show
1334 $\int_{\mathcal{X}} \psi_{0,1}(z)\psi_{0,1}(z)^\top dU(z)$ is invertible. Assume the contrary.
1335 Then there exists $v \in \mathbb{R}^D \setminus \{0\}$ such that

$$1336 v^\top \left[\int_{\mathcal{X}} \psi_{0,1}(z)\psi_{0,1}(z)^\top dU(z) \right] v = \int_{\mathcal{X}} |\psi_{0,1}(z)^\top v|^2 dU(z) = 0. \quad (92)$$

1337 Therefore, $\langle \psi_{0,1}(z), v \rangle = 0$ almost everywhere on $z \in$
1338 $[0, 1]^d$. Because $h > 0$, by re-scaling with constants this

implies the existence of non-zero coefficient vector ξ such that

$$P(z_1, \dots, z_m) := \sum_{\alpha_1+\dots+\alpha_m \leq k} \xi_{\alpha_1, \dots, \alpha_m} z_1^{\alpha_1} \dots z_m^{\alpha_m} = 0$$

almost everywhere on $z \in [0, 1]^d$.

We next use induction to show that, for any degree- k polynomial P of s variables z_1, \dots, z_s that has at least one non-zero coefficient, the set $\{z_1, \dots, z_s \in [0, 1]^d : P(z_1, \dots, z_s) = 0\}$ must have zero measure. This would then result in the desired contradiction. For the base case of $s = 1$, the fundamental theorem of algebra asserts that $P(z_1) = 0$ can have at most k roots, which is a finite set and of measure 0.

We next consider the case where $P(z_1, \dots, z_s)$ takes on s variables. Re-organizing the terms we have

$$P(z_1, \dots, z_s) \equiv P_0(z_1, \dots, z_{s-1}) + z_s P_1(z_1, \dots, z_{s-1}) + \dots + z_s^k P_k(z_1, \dots, z_{s-1}), \quad (93)$$

where P_1, \dots, P_k are degree- k polynomials of z_1, \dots, z_{s-1} . Because P has a non-zero coefficient, at least one P_j must also have a non-zero coefficient. By the inductive hypothesis, the set $\{z_1, \dots, z_{s-1} : P_j(z_1, \dots, z_{s-1}) = 0\}$ has measure 0. On the other hand, if $P_j(z_1, \dots, z_{s-1}) \neq 0$, then invoking the fundamental theorem of algebra again on z_s we know that there are finitely many z_s such that $P(z_1, \dots, z_s) = 0$. Therefore, $\{z_1, \dots, z_s : P(z_1, \dots, z_s) = 0\}$ must also have measure zero. \square

REFERENCES

- [1] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*, vol. 1. Cambridge, MA, USA: MIT Press, 2006.
- [2] B. Reeja-Jayan, K. L. Harrison, K. Yang, C.-L. Wang, A. E. Yilmaz, and A. Manthiram, "Microwave-assisted low-temperature growth of thin films in solution," *Sci. Rep.*, vol. 2, Dec. 2012, Art. no. 1003.
- [3] N. Nakamura, J. Seepaul, J. B. Kadane, and B. Reeja-Jayan, "Design for low-temperature microwave-assisted crystallization of ceramic thin films," *Appl. Stochastic Models Bus. Ind.*, vol. 33, no. 3, pp. 314–321, 2017.
- [4] A. B. Tsybakov, *Introduction to Nonparametric Estimation* (Springer Series in Statistics). New York, NY, USA: Springer, 2009.
- [5] J. Fan and I. Gijbels, *Local Polynomial Modelling and its Applications*. Boca Raton, FL, USA: CRC Press, 1996.
- [6] A. D. Bull, "Convergence rates of efficient global optimization algorithms," *J. Mach. Learn. Res.*, vol. 12, pp. 2879–2904, Oct. 2011.
- [7] J. Scarlett, I. Bogunovic, and V. Cevher, "Lower bounds on regret for noisy Gaussian process bandit optimization," in *Proc. Annu. Conf. Learn. Theory (COLT)*, 2017, pp. 1723–1742.
- [8] E. Hazan, A. Klivans, and Y. Yuan, "Hyperparameter optimization: A spectral approach," 2017, *arXiv:1706.00764*. [Online]. Available: <https://arxiv.org/abs/1706.00764#>
- [9] A. S. Nemirovski and D. B. Yudin, *Problem Complexity and Method Efficiency in Optimization*. Hoboken, NJ, USA: Wiley, 1983.
- [10] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, "Online convex optimization in the bandit setting: Gradient descent without a gradient," in *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, 2005, pp. 385–394.
- [11] A. Agarwal, O. Dekel, and L. Xiao, "Optimal algorithms for online convex optimization with multi-point bandit feedback," in *Proc. Annu. Conf. Learn. Theory (COLT)*, 2010, pp. 28–40.
- [12] K. G. Jamieson, R. Nowak, and B. Recht, "Query complexity of derivative-free optimization," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2012, pp. 2672–2680.
- [13] A. Agarwal, D. P. Foster, D. Hsu, S. M. Kakade, and A. Rakhlin, "Stochastic convex optimization with bandit feedback," *SIAM J. Optim.*, vol. 23, no. 1, pp. 213–240, 2013.
- [14] S. Bubeck, Y. T. Lee, and R. Eldan, "Kernel-based methods for bandit convex optimization," in *Proc. 49th Annu. ACM SIGACT Symp. Theory Comput. (STOC)*, 2017, pp. 72–85.
- [15] A. H. G. R. Kan and G. T. Timmer, "Stochastic global optimization methods part I: Clustering methods," *Math. Program.*, vol. 39, no. 1, pp. 27–56, 1987.
- [16] A. H. G. R. Kan and G. T. Timmer, "Stochastic global optimization methods part II: Multi level methods," *Math. Program.*, vol. 39, no. 1, pp. 57–78, 1987.
- [17] S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári, "X-armed bandits," *J. Mach. Learn. Res.*, vol. 12, pp. 1655–1695, May 2011.
- [18] C. Malherbe, E. Contal, and N. Vayatis, "A ranking approach to global optimization," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2016.
- [19] C. Malherbe and N. Vayatis, "Global optimization of Lipschitz functions," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2017.
- [20] R. D. Kleinberg, "Nearly tight bounds for the continuum-armed bandit problem," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2005, pp. 697–704.
- [21] S. Minsker, "Estimation of extreme values and associated level sets of a regression function via selective sampling," in *Proc. Conf. Learn. Theory (COLT)*, 2013, pp. 105–121.
- [22] J.-B. Grill, M. Valko, and R. Munos, "Black-box optimization of noisy functions with unknown smoothness," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2015, pp. 667–675.
- [23] S. Minsker, "Non-asymptotic bounds for prediction problems and density estimation," Ph.D. dissertation, Georgia Inst. Technol., Atlanta, Georgia, 2012.
- [24] J. Kiefer and J. Wolfowitz, "Stochastic estimation of the maximum of a regression function," *Ann. Math. Stat.*, vol. 23, no. 3, pp. 462–466, 1952.
- [25] E. Purzen, "On estimation of a probability density and mode," *Ann. Math. Statist.*, vol. 39, no. 3, pp. 1065–1076, 1962.
- [26] H. Chen, "Lower rate of convergence for locating a maximum of a function," *Ann. Statist.*, vol. 16, no. 3, pp. 1330–1334, 1988.
- [27] Z. B. Zabinsky and R. L. Smith, "Pure adaptive search in global optimization," *Math. Program.*, vol. 53, no. 1, pp. 323–338, 1992.
- [28] M.-F. Balcan, A. Beygelzimer, and J. Langford, "Agnostic active learning," *J. Comput. Syst. Sci.*, vol. 75, no. 1, pp. 78–89, 2009.
- [29] S. Dasgupta, D. J. Hsu, and C. Monteleoni, "A general agnostic active learning algorithm," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2008, pp. 353–360.
- [30] S. Hanneke, "A bound on the label complexity of agnostic active learning," in *Proc. 24th Int. Conf. Mach. Learn. (ICML)*, 2007, pp. 353–360.
- [31] E. Even-Dar, S. Mannor, and Y. Mansour, "Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems," *J. Mach. Learn. Res.*, vol. 7, pp. 1079–1105, Jun. 2006.
- [32] W. Polonik, "Measuring mass concentrations and estimating density contour clusters—an excess mass approach," *Ann. Statist.*, vol. 23, no. 3, pp. 855–881, 1995.
- [33] P. Rigollet and R. Vert, "Optimal rates for plug-in estimators of density level sets," *Bernoulli*, vol. 15, no. 4, pp. 1154–1178, 2009.
- [34] A. Singh, C. Scott, and R. Nowak, "Adaptive Hausdorff estimation of density level sets," *Ann. Statist.*, vol. 37, no. 5B, pp. 2760–2782, 2009.
- [35] K. Chaudhuri, S. Dasgupta, S. Kpotufe, and U. V. Luxburg, "Consistent procedures for cluster tree estimation and pruning," *IEEE Trans. Inf. Theory*, vol. 60, no. 12, pp. 7900–7912, Dec. 2014.
- [36] S. Balakrishnan, S. Narayanan, A. Rinaldo, A. Singh, and L. Wasserman, "Cluster trees on manifolds," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2013, pp. 2679–2687.
- [37] Y. Nesterov and B. T. Polyak, "Cubic regularization of Newton method and its global performance," *Math. Program.*, vol. 108, no. 1, pp. 177–205, 2006.
- [38] E. Hazan, K. Levy, and S. Shalev-Shwartz, "Beyond convexity: Stochastic quasi-convex optimization," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2015, pp. 1594–1602.
- [39] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points—Online stochastic gradient for tensor decomposition," in *Proc. Annu. Conf. Learn. Theory (COLT)*, 2015, pp. 797–842.
- [40] N. Agarwal, Z. Allen-Zhu, B. Bullins, E. Hazan, and T. Ma, "Finding approximate local minima faster than gradient descent," in *Proc. 49th Annu. ACM SIGACT Symp. Theory Comput. (STOC)*, 2017, pp. 1195–1199.
- [41] Y. Carmon, O. Hinder, J. C. Duchi, and A. Sidford, "'Convex until proven guilty': Dimension-free acceleration of gradient descent on non-convex functions," 2017, *arXiv:1705.02766*. [Online]. Available: <https://arxiv.org/abs/1705.02766>

AQ:3

1477 [42] Y. Zhang, P. Liang, and M. Charikar, "A hitting time analysis of
1478 stochastic gradient Langevin dynamics," in *Proc. Annu. Conf. Learn.
1479 Theory (COLT)*, 2017, pp. 1–43.

1480 [43] Y. Zhu, S. Chatterjee, J. Duchi, and J. Lafferty, "Local minimax
1481 complexity of stochastic convex optimization," in *Proc. NIPS*, 2016,
1482 pp. 3431–3439.

1483 [44] J. Duchi and F. Ruan, "Asymptotic optimality in stochastic opti-
1484 mization," 2016, *arXiv:1612.05612*. [Online]. Available: <https://arxiv.org/abs/1612.05612>

1485 [45] A. Locatelli and A. Carpentier, "Adaptivity to smoothness in X-armed
1486 bandits," in *Proc. Conf. Learn. Theory (COLT)*, 2018, pp. 1463–1492.

1487 [46] A. W. van der Vaart, *Asymptotic Statistics*, vol. 3. Cambridge, U.K.:
1488 Cambridge Univ. Press, 1998.

1489 [47] C. Jin, L. T. Liu, R. Ge, and M. I. Jordan, "On the local minima of
1490 the empirical risk," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*,
1491 2018, pp. 1–10.

1492 [48] T. T. Cai and M. G. Low, "An adaptation theory for nonparametric
1493 confidence intervals," *Ann. Statist.*, vol. 32, no. 5, pp. 1805–1840, 2004.

1494 [49] A. P. Korostelev and A. B. Tsybakov, *Minimax Theory of Image
1495 Reconstruction*, vol. 82. Springer, 2012.

1496 [50] R. M. Castro and R. D. Nowak, "Minimax bounds for active learning,"
1497 *IEEE Trans. Inf. Theory*, vol. 54, no. 5, pp. 2339–2353, May 2008.

1498 [51] S. Bubeck, R. Munos, and G. Stoltz, "Pure exploration in multi-armed
1499 bandits problems," in *Proc. Int. Conf. Algorithmic Learn. Theory (ALT)*,
1500 2009, pp. 23–37.

1501 [52] O. V. Lepski, E. Mammen, and V. G. Spokoiny, "Optimal spatial
1502 adaptation to inhomogeneous smoothness: An approach based on kernel
1503 estimates with variable bandwidth selectors," *Ann. Statist.*, vol. 25, no. 3,
1504 pp. 929–947, 1997.

1505 [53] W. K. Newey, "Convergence rates and asymptotic normality for series
1506 estimators," *J. Econometrics*, vol. 79, no. 1, pp. 147–168, 1997.

1507 [54] D. S. Ebert, *Texturing & Modeling: A Procedural Approach*. San Mateo,
1508 CA, USA: Morgan Kaufmann, 2003.

1509 [55] R. M. Castro, "Adaptive sensing performance lower bounds for sparse
1510 signal detection and support estimation," *Bernoulli*, vol. 20, no. 4,
1511 pp. 2217–2246, 2014.

1512 [56] W. Hoeffding, "Probability inequalities for sums of bounded random
1513 variables," *J. Amer. Stat. Assoc.*, vol. 58, no. 301, pp. 13–30, 1963.

1514 [57] D. Hsu, S. M. Kakade, and T. Zhang, "A tail inequality for quadratic
1515 forms of subgaussian random vectors," *Electron. Commun. Probab.*,
1516 vol. 17, no. 52, pp. 1–6, 2012.

1517 [58] J. A. Tropp, "An introduction to matrix concentration inequalities,"
1518 *Found. Trends Mach. Learn.*, vol. 8, nos. 1–2, pp. 1–230, 2015.

Yining Wang received the B.Eng. degree in computer science and technology 1520
in 2014 from Tsinghua University, Beijing China, the M.S. degree in machine 1521
learning in 2017 from Carnegie Mellon University, Pittsburgh, PA, USA. 1522
He is currently a Ph.D. student in machine learning in the machine learning 1523
department at Carnegie Mellon University, Pittsburgh, PA, USA. His research 1524
interests are primarily in statistical machine learning, with emphasis on 1525
interactive methods, active learning, adaptive sampling. 1526

Sivaraman Balakrishnan is an Assistant Professor in the Department of 1527
Statistics and Data Science at Carnegie Mellon University. Prior to this he 1528
received his Ph.D. from the School of Computer Science at Carnegie Mellon 1529
University and was a postdoctoral researcher in the Department of Statistics at 1530
UC Berkeley. His Ph.D. work was supported by several fellowships including 1531
the Richard King Mellon Fellowship and a grant from the Gates Foundation. 1532
He is broadly interested in problems that lie at the interface between computer 1533
science and statistics. Some particular areas that have provided motivation 1534
for his past and current research include the applications of statistical 1535
methods in ranking problems, computational biology, clustering, topological 1536
data analysis, nonparametric statistics, robust statistics and non-convex 1537
optimization. 1538

Aarti Singh received the B.E. degree in electronics and communication 1539
engineering from the University of Delhi, New Delhi, India, in 2001, and 1540
the M.S. and Ph.D. degrees in electrical and computer engineering from the 1541
University of Wisconsin–Madison, Madison, WI, USA, in 2003 and 2008, 1542
respectively. She was a Postdoctoral Research Associate at the Program in 1543
Applied and Computational Mathematics, Princeton University, from 2008 to 1544
2009, before joining the School of Computer Science, Carnegie Mellon 1545
University, Pittsburgh, PA, USA, where she has been an Associate Professor 1546
since 2009. Her research interests include the intersection of machine learning, 1547
statistics and signal processing, and focus on designing statistically and 1548
computationally efficient algorithms that can leverage inherent structure of the 1549
data in the form of clusters, graphs, subspaces, and manifold using direct, 1550
compressive, and active queries. Her work is recognized by the NSF Career Award, 1551
the United States Air Force Young Investigator Award, A. Nico Habermann 1552
Faculty Chair Award, Harold A. Peterson Best Dissertation Award, and a best 1553
student paper award at Allerton. 1554

AUTHOR QUERIES

AUTHOR PLEASE ANSWER ALL QUERIES

PLEASE NOTE: We cannot accept new source files as corrections for your paper. If possible, please annotate the PDF proof we have sent you with your corrections and upload it via the Author Gateway. Alternatively, you may send us your corrections in list format. You may also upload revised graphics via the Author Gateway.

AQ:1 = Author: Please confirm or add details for any funding or financial support for the research of this article.

AQ:2 = Please provide the page range for Refs. [18] and [19].

AQ:3 = Please provide the publisher location for Ref. [49].

Optimization of Smooth Functions With Noisy Observations: Local Minimax Rates

Yining Wang[✉], Sivaraman Balakrishnan, and Aarti Singh

Abstract—We consider the problem of *global optimization* of an unknown non-convex smooth function with noisy zeroth-order feedback. We propose a *local minimax* framework to study the fundamental difficulty of optimizing smooth functions with adaptive function evaluations. We show that for functions with fast growth around their global minima, carefully designed optimization algorithms can identify a near global minimizer with many fewer queries than worst-case global minimax theory predicts. For the special case of strongly convex and smooth functions, our implied convergence rates match the ones developed for zeroth-order *convex* optimization problems. On the other hand, we show that in the worst case no algorithm can converge faster than the minimax rate of estimating an unknown function in the ℓ_∞ -norm. Finally, we show that non-adaptive algorithms, though optimal in a global minimax sense, do not attain the optimal local minimax rate.

Index Terms—Optimization of smooth functions, nonparametric statistics, local minimax analysis.

I. INTRODUCTION

GLOBAL function optimization with stochastic (zeroth-order) query oracles is an important problem in optimization, machine learning and statistics. To optimize an unknown bounded function $f : \mathcal{X} \mapsto \mathbb{R}$ defined on a known compact d -dimensional domain $\mathcal{X} \subseteq \mathbb{R}^d$, the data analyst makes n *active* queries $x_1, \dots, x_n \in \mathcal{X}$ and observes

$$y_t = f(x_t) + w_t, \quad w_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad t = 1, \dots, n. \quad (1)$$

The queries x_1, \dots, x_t are *active* in the sense that the selection of x_t can depend on the previous queries and their responses $x_1, y_1, \dots, x_{t-1}, y_{t-1}$. After n queries, an estimate $\hat{x}_n \in \mathcal{X}$ is produced that approximately minimizes the unknown function f . Such “active query” models are relevant in a broad range of (noisy) global optimization applications, for instance in hyper-parameter tuning of machine learning algorithms [1] and

Manuscript received August 10, 2018; revised April 21, 2019; accepted May 5, 2019. S. Balakrishnan was supported in part by the NSF under Grant DMS-17130003. Y. Wang and A. Singh were supported in part by the NSF under Grant CCF-1563918 and in part by the AFRL under Grant FA8750-17-2-0212. This paper was presented in part at the 2018 NeurIPS Conference.

Y. Wang and A. Singh are with the Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: yiningwa,aarti@cs.cmu.edu).

S. Balakrishnan is with the Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: siva@stat.cmu.edu).

Communicated by K. Chaudhuri, Associate Editor for Statistical Learning. Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2019.2921985

¹The exact Gaussianity of the independent noise variables ε_t is not crucial and our results can be easily generalized to sub-Gaussian noise.

sequential design in material synthesis experiments where the goal is to maximize the strength of the synthesized material as a function of experimental settings [2], [3]. We refer the readers to Section II-A for a rigorous formulation of the active query model and contrast it with the classical passive query model.

The error of the estimate \hat{x}_n is measured by the difference of $f(\hat{x}_n)$ and the *global minimum* of f :

$$\mathcal{L}(\hat{x}_n; f) := f(\hat{x}_n) - f^* \quad \text{where } f^* := \inf_{x \in \mathcal{X}} f(x). \quad (2)$$

To simplify our presentation, throughout the paper we take the domain \mathcal{X} to be the d -dimensional unit cube $[0, 1]^d$, while our results can be easily generalized to other compact domains satisfying minimal regularity conditions.

When f belongs to a smoothness class, say the Hölder class with exponent α , a straightforward global optimization method is to first sample n points uniformly at random from \mathcal{X} and then construct nonparametric estimates \hat{f}_n of f using nonparametric regression methods such as kernel smoothing or local polynomial regression [4], [5]. Classical analysis shows that the sup-norm reconstruction error $\|\hat{f}_n - f\|_\infty = \sup_{x \in \mathcal{X}} |\hat{f}_n(x) - f(x)|$ can be upper bounded by $\tilde{O}_{\mathbb{P}}(n^{-\alpha/(2\alpha+d)})^2$. This global reconstruction guarantee then implies an $\tilde{O}_{\mathbb{P}}(n^{-\alpha/(2\alpha+d)})$ upper bound on $\mathcal{L}(\hat{x}_n; f)$ by considering an estimate $\hat{x}_n \in \mathcal{X}$ for which $\hat{f}_n(\hat{x}_n) = \inf_{x \in \mathcal{X}} \hat{f}_n(x)$ (such an \hat{x}_n exists because \mathcal{X} is closed and bounded). Formally, we have the following proposition (proved in the Appendix) that converts a global reconstruction guarantee into an upper bound on the optimization error:

Proposition 1. *Suppose $\hat{f}_n(\hat{x}_n) = \inf_{x \in \mathcal{X}} \hat{f}_n(x)$. Then $\mathcal{L}(\hat{x}_n; f) \leq 2\|\hat{f}_n - f\|_\infty$.*

Typically, fundamental limits on the optimal optimization error are understood through the lens of *minimax analysis* where the object of study is the (global) minimax risk:

$$\inf_{\hat{x}_n} \sup_{f \in \mathcal{F}} \mathbb{E}_f \mathcal{L}(\hat{x}_n, f), \quad (3)$$

where \mathcal{F} is a certain class of smooth functions such as the Hölder class. Although optimization appears to be easier than global reconstruction, we show in this paper that the $n^{-\alpha/(2\alpha+d)}$ rate is *not* improvable in the global minimax sense in over Hölder classes. Such a surprising phenomenon was also noted in previous works [6]–[8] for related problems. On the

²In the $\tilde{O}(\cdot)$ or $\tilde{O}_{\mathbb{P}}(\cdot)$ notation we suppress constant factors and terms that depend poly-logarithmically on n .

other hand, extensive empirical evidence suggests that non-uniform/active allocations of query points can significantly reduce optimization error in practical global optimization of smooth, non-convex functions [1]. This raises the interesting question of understanding, from a theoretical perspective, the conditions under which the global optimization of smooth functions is *easier* than their reconstruction, and the power of *active/feedback-driven* queries that play important roles in global optimization.

In this paper, we propose a theoretical framework that partially answers the above questions. In contrast to classical *global* minimax analysis of nonparametric estimation problems, we adopt a *local analysis* which characterizes the optimal convergence rate of optimization error when the underlying function f is within a neighborhood of a “reference” function f_0 . (See Section II-B for the rigorous local minimax formulation considered in this paper.) Our main results are to characterize the local convergence rates $R_n(f_0)$ for a wide range of reference functions $f_0 \in \mathcal{F}$. Concretely, our contributions can be summarized as follows:

- 1) We design an iterative (active) algorithm whose optimization error $\mathcal{L}(\hat{x}_n; f)$ converges at a rate of $R_n(f_0)$ depending on the reference function f_0 . When the level-sets of f_0 satisfy certain regularity and polynomial growth conditions, the local rate $R_n(f_0)$ can be upper bounded by $R_n(f_0) = \tilde{O}(n^{-\alpha/(2\alpha+d-\alpha\beta)})$, where $\beta \in [0, d/\alpha]$ is a parameter depending on f_0 that characterizes the volume growth of the *level-sets* of the reference function f_0 . (See assumption (A2), Proposition 2 and Theorem 1 for details). The rate matches the global minimax convergence rate $n^{-\alpha/(2\alpha+d)}$ for worst-case f_0 where $\beta = 0$, but can be much faster when $\beta > 0$. We emphasize that our algorithm has no knowledge of the reference function f_0 and achieves this rate adaptively.
- 2) We prove *local* minimax lower bounds that match the $n^{-\alpha/(2\alpha+d-\alpha\beta)}$ upper bound, up to logarithmic factors in n . More specifically, we show that *even if* f_0 is known, no (active) algorithm can estimate f in close neighborhoods of f_0 at a rate faster than $n^{-\alpha/(2\alpha+d-\alpha\beta)}$. We further show that, if active queries are not available and queries x_1, \dots, x_n are i.i.d. uniformly sampled from \mathcal{X} , then the $n^{-\alpha/(2\alpha+d)}$ global minimax rate also applies locally regardless of how large β is. Thus, there is an explicit gap between local minimax rates in the active and uniform query models when β is large.
- 3) In the special case when f is *convex*, the global optimization problem is usually referred to as *zeroth-order convex optimization* and this problem has been widely studied [9]–[14]. Our results imply that, when f_0 is *strongly convex* and *smooth*, the local minimax rate $R_n(f_0)$ is on the order of $\tilde{O}(n^{-1/2})$, which matches the convergence rates in [11]. Additionally, our negative results (Theorem 2) indicate that the $n^{-1/2}$ rate cannot be achieved if f_0 is merely convex, which seems to contradict $n^{-1/2}$ results in [13], [14] that do not require strong convexity of f . However, it should be noted that mere convexity of f_0 does not imply convexity of f in

a neighborhood of f_0 (e.g., $\|f - f_0\|_\infty \leq \varepsilon$). Our results show significant differences in the intrinsic difficulty of zeroth-order optimization of convex and near-convex functions.

A. Related Work

Global optimization, known variously as *black-box optimization*, *Bayesian optimization* and the *continuum-armed bandit*, has a long history in the optimization research community [15], [16] and has also received a significant amount of recent interest in statistics and machine learning [1], [6], [8], [17]–[19]. Many previous works [17], [20] have derived rates for non-convex smooth payoffs in “continuum-armed” bandit problems.

The papers [21], [22] are closely related to our work. They studied the related problem of estimating the set of all optima of a smooth function in the Hausdorff distance. For Hölder smooth functions with polynomial growth, the paper [21] derives an $n^{-1/(2\alpha+d-\alpha\beta)}$ minimax rate for $\alpha < 1$ (subsequently improved to include $\alpha \geq 1$ in [23]). This result is similar to our Propositions 2 and 3. The papers [21], [22] also discussed adaptivity to unknown smoothness parameters. We however remark on several differences between our work and the papers [21], [22]. First, in [21], [22] only functions with polynomial growth are considered, while in our Theorems 1 and 2 functionals $\varepsilon_n^U(f_0)$ and $\varepsilon_n^L(f_0)$ are proposed for general reference functions f_0 satisfying mild regularity conditions, which include functions with polynomial growth as special cases. In addition, [21] considers the harder problem of estimating maxima sets in Hausdorff distance, as opposed to the problem of producing a single approximately optimal solution \hat{x}_T . As a result, the minimax lower bounds in [21] do not apply to this latter setting. An algorithm, without distinguishing between two functions with different optima sets, can nevertheless produce a good approximate optimizer as long as the two functions under consideration have *overlapping* optima sets. New constructions and information-theoretic techniques are therefore required to prove lower bounds under the weaker (one-point) approximate optimization framework. Finally, we prove minimax lower bounds when only *uniform* query points are available and demonstrate a significant gap between algorithms having access to uniformly sampled or adaptively chosen data points.

The papers [18], [19] imposed additional assumptions on the level-sets of the underlying function to obtain an improved convergence rate. The level-set assumptions considered in the mentioned references are rather restrictive and essentially require the underlying function to be uni-modal, while our assumptions are much more flexible and apply to multi-modal functions as well. In addition, [18], [19] considered a *noiseless* setting in which exact function evaluations $f(x_t)$ can be obtained, while our paper studies the noise corrupted model in (1) for which vastly different convergence rates are derived. Finally, no matching lower bounds were proved in the papers [18], [19].

The (stochastic) global optimization problem is similar to *mode estimation* of either densities or regression functions,

which has a rich literature [24]–[26]. An important difference between statistical mode estimation and global optimization is the way sample/query points $x_1, \dots, x_n \in \mathcal{X}$ are distributed: in mode estimation it is customary to assume the samples are independently and identically distributed, while in global optimization sequential designs of samples/queries are typical. Furthermore, to estimate/locate the mode of an unknown density or regression function, such a mode has to be well-defined; on the other hand, producing an estimate \hat{x}_n with small $\mathcal{L}(\hat{x}_n, f)$ is easier and results in weaker conditions imposed on the underlying function.

Methodology-wise, our proposed algorithm is conceptually similar to the abstract *Pure Adaptive Search (PAS)* framework proposed and analyzed in [27]. The iterative procedure also resembles disagreement-based active learning methods [28]–[30] and the “successive rejection” algorithm in bandit problems [31]. The intermediate steps of candidate point elimination can also be viewed as level-set estimation problems [32]–[34] or cluster-tree estimation problems [35], [36] with active queries.

Another line of research has focused on *first-order* optimization of quasi-convex or non-convex functions [37]–[42], in which exact or unbiased evaluations of function *gradients* are available at query points $x \in \mathcal{X}$. The paper [42] considered a Cheeger’s constant restriction on level-sets which is similar to our level-set regularity assumptions (A2 and A2’). The papers [43], [44] studied local minimax rates for the first-order optimization of convex functions. First-order optimization differs significantly from our setting because unbiased gradient estimation is generally impossible in the model of (1). Furthermore, most works on (first-order) non-convex optimization focus on obtaining stationary points or local minima, while we consider the problem of finding a (near) global minima.

221 B. Comparison with the HOO Algorithm

222 The HOO algorithm [17], as well as similar algorithms
223 such as Algorithm 2 in [45] and the POO algorithm in [22],
224 are theoretically well-studied methods for global optimization.
225 Below we summarize the differences of our results and the
226 ones from these works.

227 (a) Weaker Smoothness Conditions I: In Algorithm 1,
228 we use local polynomial estimation as a sub-routine
229 to obtain local estimates of the objective function
230 f . Compared to the sample average approach in
231 HOO (e.g., Algorithm 2 in [45]), local polynomial
232 estimates have the advantage of being unbiased for
233 the estimation of low-degree polynomials. This trans-
234 lates to the improved (A1) Hölder-continuity condition
235 that *only* restricts the $\lfloor \alpha \rfloor$ -th order derivatives
236 of objective functions. More specifically, the actual
237 function values of $f(x)$ and $f(x')$ for x, x' close
238 to each other can be very different, as long as such
239 differences can be perfectly modeled by low-degree
240 polynomials. This is in contrast to the smoothness
241 conditions imposed in [17], [45] which essentially
242 require $f(x)$ to be close to $f(x^*)$ for x close to x^* the
243 optima of f .

244 (b) Weaker Smoothness Conditions II: Our results in
245 Section IV-C hold on functions that are only assumed
246 to be smooth in regions close to its global minimum, in
247 contrast to Definition 1 in [45] and many other existing
248 works that place smoothness assumptions on the entire
249 domain of the objective function f .

250 (c) Spatially Restricted Queries: Our proposed algorithm is
251 “grid” based, and can be run on any sufficiently dense
252 finite grid G_n in \mathcal{X} and does not need to have the
253 capacity to query arbitrary points in \mathcal{X} . As a result,
254 our algorithm can be run in experimental settings where
255 queries are restricted to belong to a large pool of a-priori
256 chosen points.

257 (d) Results for any Smooth Function: Our algorithm and
258 lower bounds yield essentially tight results for the
259 complexity of optimization of arbitrary smooth func-
260 tions. While these rates are most interpretable under
261 the level-set growth conditions (also studied in [45]) our
262 results also yield nearly matching guarantees for other
263 (arbitrary, smooth) functions f_0 .

264 II. BACKGROUND AND NOTATION

265 We first review standard asymptotic notation that will
266 be used throughout this paper. For two sequences $\{a_n\}_{n=1}^\infty$
267 and $\{b_n\}_{n=1}^\infty$, we write $a_n = O(b_n)$ or $a_n \lesssim b_n$ if
268 $\limsup_{n \rightarrow \infty} |a_n|/|b_n| < \infty$, or equivalently $b_n = \Omega(a_n)$ or
269 $b_n \gtrsim a_n$. Denote $a_n = \Theta(b_n)$ or $a_n \asymp b_n$ if both $a_n \lesssim b_n$
270 and $a_n \gtrsim b_n$ hold. We also write $a_n = o(b_n)$ or equivalently
271 $b_n = \omega(a_n)$ if $\lim_{n \rightarrow \infty} |a_n|/|b_n| = 0$. For two sequences
272 of random variables $\{A_n\}_{n=1}^\infty$ and $\{B_n\}_{n=1}^\infty$, denote $A_n =$
273 $O_{\mathbb{P}}(B_n)$ if for every $\epsilon > 0$, there exists $C > 0$ such that
274 $\limsup_{n \rightarrow \infty} \Pr[|A_n| > C|B_n|] \leq \epsilon$. For $r > 0$, $1 \leq p \leq \infty$
275 and $x \in \mathbb{R}^d$, we denote by $B_r^p(x) := \{z \in \mathbb{R}^d : \|z - x\|_p \leq r\}$
276 the d -dimensional ℓ_p -ball of radius r centered at x , where
277 the vector ℓ_p norm is defined as $\|x\|_p := (\sum_{j=1}^d |x_j|^p)^{1/p}$
278 for $1 \leq p < \infty$ and $\|x\|_\infty := \max_{1 \leq j \leq d} |x_j|$. For any subset
279 $S \subseteq \mathbb{R}^d$ we denote by $B_r^p(x; S)$ the set $B_r^p(x) \cap S$.

280 A. Passive and Active Query Models

281 Let U be a known random quantity defined on a probability
282 space \mathcal{U} . The following definitions characterize all passive and
283 active optimization algorithms:

284 **Definition 1** (The passive query model). *Let x_1, \dots, x_n be
285 i.i.d. points uniformly sampled on \mathcal{X} and y_1, \dots, y_n be obser-
286 vations from the model (1). A passive optimization algorithm
287 \mathcal{A} with n queries is parameterized by a mapping $\phi_n :
288 (x_1, y_1, \dots, x_n, y_n, U) \mapsto \hat{x}_n$ that maps the i.i.d. observations
289 $\{(x_i, y_i)\}_{i=1}^n$ to an estimated optimum $\hat{x}_n \in \mathcal{X}$, potentially
290 randomized by U .*

291 **Definition 2** (The active query model). *An active opti-
292 mization algorithm can be parameterized by mappings
293 $(\chi_1, \dots, \chi_n, \phi_n)$, where for $t = 1, \dots, n$,*

$$294 \chi_t : (x_1, y_1, \dots, x_{t-1}, y_{t-1}, U) \mapsto x_t$$

295 produces a query point $x_t \in \mathcal{X}$ based on previous observations
 296 $\{(x_i, t_i)\}_{i=1}^{t-1}$, and

$$297 \quad \phi_n : (x_1, y_1, \dots, x_n, y_n, U) \mapsto \hat{x}_n$$

298 produces the final estimate. All mappings $(\chi_1, \dots, \chi_n, \phi_n)$ can
 299 be randomized by U .

300 B. Local Minimax Rates

301 We use a classical *local minimax analysis* [46] to understand
 302 the fundamental information-theoretic limits of noisy global
 303 optimization of smooth functions. On the upper bound side,
 304 we seek (active) estimators \hat{x}_n such that

$$305 \quad \sup_{f_0 \in \Theta} \sup_{f \in \Theta', \|f - f_0\|_\infty \leq \varepsilon_n(f_0)} \Pr_f [\mathcal{L}(\hat{x}_n; f) \geq C_1 \cdot R_n(f_0)] \leq 1/4, \quad (4)$$

307 where $C_1 > 0$ is a positive constant. Here $f_0 \in \Theta$ is
 308 referred to as the *reference function*, and $f \in \Theta'$ is the true
 309 underlying function to be optimized, which is assumed to be
 310 “near” f_0 (in the ℓ_∞ norm). The minimax convergence rate
 311 of $\mathcal{L}(\hat{x}_n; f)$ is then characterized *locally* by $R_n(f_0)$ which
 312 depends on the reference function f_0 . The constant of $1/4$
 313 is chosen arbitrarily and any small constant leads to similar
 314 conclusions. To establish negative results (i.e., local minimax
 315 lower bounds), in contrast to the upper bound formulation,
 316 we assume the potential active optimization estimator \hat{x}_n has
 317 *perfect knowledge* about the reference function $f_0 \in \Theta$.
 318 We then prove local minimax lower bounds of the form

$$319 \quad \inf_{\hat{x}_n} \sup_{f \in \Theta', \|f - f_0\|_\infty \leq \varepsilon_n(f_0)} \Pr_f [\mathcal{L}(\hat{x}_n; f) \geq C_2 \cdot R_n(f_0)] \geq 1/3, \quad (5)$$

321 where $C_2 > 0$ is another positive constant and $\varepsilon_n(f_0)$, $R_n(f_0)$
 322 are desired local convergence rates for functions near the
 323 reference f_0 .

324 Although in some sense classical, the local minimax definition
 325 we propose warrants further discussion:

326 1) **Roles of Θ and Θ' :** The reference function f_0 and the
 327 true functions f are assumed to belong to different but
 328 closely related function classes Θ and Θ' . In particular,
 329 in our paper $\Theta \subseteq \Theta'$, meaning that less restrictive
 330 assumptions are imposed on the true underlying function
 331 f compared to those imposed on the reference function
 332 f_0 on which R_n and ε_n are based.

333 2) **Upper Bounds:** It is worth emphasizing that the
 334 estimator \hat{x}_n has no knowledge of the reference function
 335 f_0 . From the perspective of upper bounds, we can
 336 consider the simpler task of producing f_0 -dependent
 337 bounds (eliminating the second supremum) to instead
 338 study the (already interesting) quantity:

$$339 \quad \sup_{f_0 \in \Theta} \Pr_{f_0} [\mathcal{L}(\hat{x}_n; f_0) \geq C_1 R_n(f_0)] \leq 1/4.$$

340 As indicated above we maintain the double-supremum
 341 in the definition because fewer assumptions are imposed
 342 directly on the true underlying function f , and further
 343 because it allows to more directly compare our upper
 344 and lower bounds.

345 3) **Lower Bounds and the choice of the “localization
 346 radius” $\varepsilon_n(f_0)$:** Our lower bounds allow the estimator
 347 knowledge of the reference function (this makes
 348 establishing the lower bound more challenging). The
 349 lower bound in (5) implies that no estimator \hat{x}_n can
 350 effectively optimize a function f close to f_0 beyond the
 351 convergence rate of $R_n(f_0)$, even if perfect knowledge
 352 of the reference function f_0 is available a priori. The
 353 $\varepsilon_n(f_0)$ parameter that decides the “range” in which
 354 local minimax rates apply is taken to be on the same
 355 order as the actual local rate $R_n(f_0)$ in this paper.
 356 This is (up to constants) the smallest radius for which
 357 we can hope to obtain non-trivial lower-bounds: if we
 358 consider a much smaller radius than $R_n(f_0)$ then the
 359 trivial estimator which outputs the minimizer of the ref-
 360 erence function would achieve a faster rate than $R_n(f_0)$.
 361 On the other hand selecting the smallest possible radius
 362 makes establishing the lower bound most challenging
 363 but provides a refined picture of the complexity of
 364 zeroth-order optimization.

365 We remark that our primary motivation for the
 366 local-minimax analysis stems from the fact that for natural
 367 function classes the global-minimax rate for the optimization
 368 complexity is excessively pessimistic, while the local minimax
 369 analysis provides a more refined picture. In machine learning
 370 applications, there are several cases where the population risk
 371 is well-behaved (smooth, potentially non-convex) but we are
 372 only able to access/query the empirical risk which we want to
 373 minimize. Using standard concentration bounds the empirical
 374 risk and population risk are close, and the resulting problem
 375 is then to minimize the approximate-smooth empirical risk
 376 (see for instance [42], [47] for a more detailed discussion).

377 III. MAIN RESULTS

378 With this background in place we now turn our attention
 379 to our main results. We begin by collecting our assumptions
 380 about the true underlying function and the reference function
 381 in Section III-A. We state and discuss the consequences of
 382 our upper and lower bounds in Sections III-B and III-C
 383 respectively. We defer most technical proofs to Section V and
 384 turn our attention to our optimization algorithm in Section IV.

385 A. Assumptions

386 We first state and motivate assumptions that will be used.
 387 The first assumption states that f is locally Hölder smooth on
 388 its level-sets.

389 (A1) There exist constants $\kappa, \alpha, M, \zeta > 0$ such
 390 that f restricted to $\mathcal{X}_{f, \kappa, \zeta} := \{x \in \mathcal{X} : \inf_{z \in \mathcal{X}, \|z - x\|_\infty \leq \zeta} f(z) \leq f^* + \kappa\}$ belongs to the
 391 Hölder class $\dot{\Sigma}^\alpha(M)$, meaning that f is k -times
 392 differentiable on $\mathcal{X}_{f, \kappa, \zeta}$ and furthermore for any
 393 $x, x' \in \mathcal{X}_{f, \kappa, \zeta}$,
 394

$$395 \quad \sum_{a_1 + \dots + a_d = k} \frac{|f^{(\alpha, k)}(x) - f^{(\alpha, k)}(x')|}{\|x - x'\|_\infty^{\alpha - k}} \leq M. \quad (6)$$

396 ³We use the ℓ_∞ -norm for convenience and it can be replaced by any
 397 equivalent vector norm.

396 Here $k = \lfloor \alpha \rfloor$ is the largest integer lower bounding α
 397 and $f^{(\alpha, j)}(x) := \partial^j f(x)/\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}$.

398 We use $\Sigma_\kappa^\alpha(M)$ to denote the class of all functions satisfying (A1). We remark that (A1) is weaker than the usual Hölder
 399 assumption in two ways. First, (6) only imposes stability
 400 conditions on the $\lfloor \alpha \rfloor$ -th order derivatives of the function f , in
 401 contrast to conditions involving all orders of derivatives in previous
 402 works [17], [45]. Second, (A1) only imposes the Hölder
 403 smoothness assumption on certain regions of \mathcal{X} , because
 404 regions with function values larger than $f^* + \kappa$ can be easily
 405 detected and removed by a pre-processing step, highlighting
 406 an important difference between optimization and ℓ_∞ -norm
 407 estimation. We give further details of the pre-processing step
 408 in Section IV-C.

409 Our next assumption concerns the “regularity” of the *level-sets* of the “reference” function f_0 . Define $L_{f_0}(\epsilon) := \{x \in$
 410 $\mathcal{X} : f_0(x) \leq f_0^* + \epsilon\}$ as the ϵ -level-set of f_0 , and
 411 $\mu_{f_0}(\epsilon) := \lambda(L_{f_0}(\epsilon))$ as the Lebesgue measure of $L_{f_0}(\epsilon)$,
 412 which we refer to as the *distribution function*. Define,
 413 $N(L_{f_0}(\epsilon), \delta)$ as the smallest number of ℓ_2 -balls of radius δ
 414 that cover $L_{f_0}(\epsilon)$. Then we make the following assumption:
 415 (A2) There exist constants $c_0 > 0$ and $C_0 > 0$ such that
 416 $N(L_{f_0}(\epsilon), \delta) \leq C_0[1 + \mu_{f_0}(\epsilon)\delta^{-d}]$ for all $\epsilon, \delta \in (0, c_0]$.
 417 We use Θ_C to denote all functions that satisfy (A2) with
 418 respect to parameters $\mathbf{C} = (c_0, C_0)$.

419 At a high-level, the regularity condition (A2) assumes that
 420 the level-sets are sufficiently “regular” such that covering them
 421 with small-radius balls does not require significantly larger
 422 total volume. For example, consider the perfectly regular case
 423 when $L_{f_0}(\epsilon)$ is the d -dimensional ℓ_2 ball of radius r : $L_{f_0}(\epsilon) =$
 424 $\{x \in \mathcal{X} : \|x - x^*\|_2 \leq r\}$. Clearly, $\mu_{f_0}(\epsilon) \asymp r^d$. In addition,
 425 the δ -covering number in ℓ_2 of $L_{f_0}(\epsilon)$ is on the order of $1 +$
 426 $(r/\delta)^d \asymp 1 + \mu_{f_0}(\epsilon)\delta^{-d}$, which satisfies the scaling in (A2).

427 When (A2) holds, uniform confidence intervals for f on
 428 its level-sets are easier to construct because little statistical
 429 efficiency is lost by slightly enlarging the level-sets so that
 430 complete (sufficiently small) d -dimensional cubes are con-
 431 tained in the enlarged level-sets. On the other hand, when
 432 regularity of level-sets fails to hold such nonparametric esti-
 433 mation can be very difficult or even impossible. As an extreme
 434 example, suppose the level-set $L_{f_0}(\epsilon)$ consists of n standalone
 435 and well-spaced points in \mathcal{X} : the Lebesgue measure of $L_{f_0}(\epsilon)$
 436 would be zero, but at least $\Omega(n)$ queries are necessary to
 437 construct uniform confidence intervals on $L_{f_0}(\epsilon)$. It is clear
 438 that such $L_{f_0}(\epsilon)$ violates (A2), because $N(L_{f_0}(\epsilon), \delta) \geq n$ as
 439 $\delta \rightarrow 0^+$ but $\mu_{f_0}(\epsilon) = 0$.

442 B. Upper Bound

443 The following theorem is our main result that provides
 444 an upper bound on the local minimax rate of noisy global
 445 optimization with active queries.

446 **Theorem 1.** For any $\alpha, M, \kappa, c_0, C_0 > 0$ and $f_0 \in \Sigma_\kappa^\alpha(M) \cap$
 447 Θ_C , where $\mathbf{C} = (c_0, C_0)$, define

$$448 \varepsilon_n^U(f_0) := \sup \left\{ \varepsilon > 0 : \varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \geq n/\log^\omega n \right\}, \quad (7)$$

449 where $\omega > 5 + d/\alpha$ is a large constant. Suppose also that
 450 $\varepsilon_n^U(f_0) \rightarrow 0$ as $n \rightarrow \infty$. Then for sufficiently large n ,

451 there exists an estimator \hat{x}_n with access to n active queries
 452 $x_1, \dots, x_n \in \mathcal{X}$, a constant $C_R > 0$ depending only
 453 on $\alpha, M, \kappa, c, c_0, C_0$ and a constant $\gamma > 0$ depending only
 454 on α and d such that

$$455 \sup_{f_0 \in \Sigma_\kappa^\alpha(M) \cap \Theta_C} \sup_{\substack{f \in \Sigma_\kappa^\alpha(M), \\ \|f - f_0\|_\infty \leq \varepsilon_n^U(f_0)}} \Pr_f [\mathcal{L}(\hat{x}_n, f) > \\ 456 C_R \log^\gamma n \cdot (\varepsilon_n^U(f_0) + n^{-1/2})] \leq 1/4. \quad (8)$$

457
 458 **Remark 1.** Unlike the (local) smoothness class $\Sigma_\kappa^\alpha(M)$,
 459 the additional function class Θ_C that encapsulates (A2) is
 460 imposed only on the “reference” function f_0 but not the
 461 true function f to be estimated. This makes the assumptions
 462 considerably weaker because the true function f may violate
 463 (A2) while our results remain valid.

464 **Remark 2.** The estimator \hat{x}_n does not require knowledge of
 465 parameters κ, c_0, C_0 or $\varepsilon_n^U(f_0)$, and automatically adapts to
 466 them, as shown in the next section. While the knowledge of
 467 smoothness parameters α and M is in general unavoidable
 468 in non-parametric regression (see [48]), in the zeroth-order
 469 optimization problem it is possible to adapt to α and M
 470 by running $O(\log^2 n)$ parallel sessions of \hat{x}_n on $O(\log n)$
 471 grids of α and M values, and then using $\Omega(n/\log^2 n)$
 472 single-point queries to decide on the location with the smallest
 473 function value. This adaptive strategy was suggested in [22]
 474 to remove an additional condition in [21], and also applies to
 475 our setting.

476 **Remark 3.** When the distribution function $\mu_{f_0}(\epsilon)$ does not
 477 change abruptly with ϵ the expression of $\varepsilon_n^U(f_0)$ can be
 478 significantly simplified. In particular, if for all $\epsilon \in (0, c_0]$ it
 479 holds that

$$480 \mu_{f_0}(\epsilon/\log n) \geq \mu_{f_0}(\epsilon)/[\log n]^{O(1)}, \quad (9)$$

481 then $\varepsilon_n^U(f_0)$ can be upper bounded as

$$482 \varepsilon_n^U(f_0) \leq [\log n]^{O(1)} \cdot \sup \left\{ \varepsilon > 0 : \varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \geq n \right\}. \quad (10)$$

483 If $\mu_{f_0}(\epsilon)$ scales polynomially with ϵ , i.e. $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$ for
 484 some constant $\beta \geq 0$, then (9) and (10) are both satisfied.

485 The quantity $\varepsilon_n^U(f_0) = \sup \{ \varepsilon > 0 : \varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \geq n/\log^\omega n \}$ is crucial in determining the convergence rate of
 486 optimization error of \hat{x}_n locally around the reference function
 487 f_0 . While the definition of $\varepsilon_n^U(f_0)$ is mostly implicit and
 488 involves solving an inequality involving the distribution function
 489 $\mu_{f_0}(\cdot)$, we remark that it admits a simple form when μ_{f_0}
 490 has a polynomial growth rate similar to a local Tsybakov noise
 491 condition [4], [49], as shown in the following proposition:

492 **Proposition 2.** Suppose $\mu_{f_0}(\epsilon) \lesssim \epsilon^\beta$ for some constant
 493 $\beta \in [0, 2 + d/\alpha]$. Then $\varepsilon_n^U(f_0) = \tilde{O}(n^{-\alpha/(2\alpha+d-\alpha\beta)})$.
 494 In addition, if $\beta \in [0, d/\alpha]$ then $\varepsilon_n^U(f_0) + n^{-1/2} \lesssim \varepsilon_n^U(f_0) =$
 495 $\tilde{O}(n^{-\alpha/(2\alpha+d-\alpha\beta)})$.

496 We remark that, following Proposition 1 of [45], α, β and d
 497 must satisfy the relationship that $\beta \leq d/\alpha$. Proposition 2 can

be easily verified by solving the system $\epsilon^{-(2+d/\alpha)}\mu_{f_0}(\epsilon) \geq n/\log^\omega n$ with the condition $\mu_{f_0}(\epsilon) \lesssim \epsilon^\beta$. We therefore omit its proof. The following two examples give some simple reference functions f_0 that satisfy the $\mu_{f_0}(\epsilon) \lesssim \epsilon^\beta$ condition in Proposition 2 with particular values of β .

Example 1. The constant function $f_0 \equiv 0$ satisfies (A1) through (A3) with $\beta = 0$.

Example 2. $f_0 \in \Sigma_\kappa^2(M)$ that is strongly convex⁴ satisfies (A1) through (A3) with $\beta = d/2$.

Example 1 is simple to verify, as the volume of level-sets of the constant function $f_0 \equiv 0$ exhibit a phase transition at $\epsilon = 0$ and $\epsilon > 0$. Consequently, $\beta = 0$ is the only parameter for which $\mu_{f_0}(\epsilon) \lesssim \epsilon^\beta$. Example 2 is more involved, and holds because the strong convexity of f_0 lower bounds the growth rate of f_0 when moving away from its minimum. We give a rigorous proof for Example 2 in the appendix. We also remark that f_0 does not need to be exactly strongly convex for $\beta = d/2$ to hold, and the example is valid for, e.g., piecewise strongly convex functions with a constant number of pieces too.

To best interpret the results in Theorem 1 and Proposition 2, it is instructive to compare the “local” rate $n^{-\alpha/(2\alpha+d-\alpha\beta)}$ with the baseline rate $n^{-\alpha/(2\alpha+d)}$, which can be attained by reconstructing f in sup-norm and applying Proposition 1. Since $\beta \geq 0$, the local convergence rate established in Theorem 1 is never slower, and the improvement compared to the baseline rate $n^{-\alpha/(2\alpha+d)}$ is dictated by β , which governs the growth rate of volume of level-sets of the reference function f_0 . In particular, for functions that grows fast when moving away from its minimum, the parameter β is large and therefore the local convergence rate around f_0 could be much faster than $n^{-\alpha/(2\alpha+d)}$.

Theorem 1 also implies concrete convergence rates for special functions considered in Examples 1 and 2. For the constant reference function $f_0 \equiv 0$, Example 1 and Theorem 1 yield that $R_n(f_0) \asymp n^{-\alpha/(2\alpha+d)}$, which matches the baseline rate $n^{-\alpha/(2\alpha+d)}$ and suggests that $f_0 \equiv 0$ is the worst-case reference function. This is intuitive, because $f_0 \equiv 0$ has a drastic level-set change at $\epsilon \rightarrow 0^+$ and therefore small perturbations of f_0 result in changes to the optimal location. On the other hand, if f_0 is strongly smooth and convex as in Example 2, Theorem 1 leads to the bound of $R_n(f_0) \asymp n^{-1/2}$, which is significantly better than the $n^{-2/(4+d)}$ baseline rate⁵ and also matches existing works on zeroth-order optimization of convex functions [11]. The faster rate holds intuitively because strongly convex functions grow quickly when moving away from the minimum. An active query algorithm can focus most of its queries on the small level-sets of the underlying function, resulting in more accurate local function reconstruction and faster optimization error rate.

Our proof of Theorem 1 is constructive, by upper bounding the local minimax optimization error of an explicit algorithm.

⁴A twice differentiable function f_0 is strongly convex if there exists $\sigma > 0$ such that $\nabla^2 f_0(x) \succeq \sigma I, \forall x \in \mathcal{X}$.

⁵Note that f_0 being strongly smooth corresponds to $\alpha = 2$ in the local smoothness assumption.

Roughly, our algorithm partitions the n active queries evenly into $\log n$ epochs, and level-sets of f are estimated at the end of each epoch by comparing (uniform) confidence intervals on a dense grid on \mathcal{X} . It is then proved that the volume of the estimated level-sets contracts *geometrically*, until the target convergence rate $R_n(f_0)$ is attained. The algorithm is described in more detail in Section IV and the complete proof of Theorem 1 is in Section V-B.

C. Lower Bounds

We prove local minimax lower bounds that match the upper bounds in Theorem 1 up to logarithmic terms. As we remarked in Section II-B, in the local minimax lower bound formulation we assume the data analyst has full knowledge of the reference function f_0 , which makes the lower bounds stronger as more information is available a priori.

To facilitate such local minimax lower bounds, the following additional condition is imposed on the reference function f_0 of which the data analyst has perfect information.

(A2') There exist constants $c'_0, C'_0 > 0$ such that $M(L_{f_0}(\epsilon), \delta) \geq C'_0 \mu_{f_0}(\epsilon) \delta^{-d}$ for all $\epsilon, \delta \in (0, c'_0]$, where $M(L_{f_0}(\epsilon), \delta)$ is the maximum number of disjoint ℓ_2 balls of radius δ that can be packed into $L_{f_0}(\epsilon)$.

We denote $\Theta'_{\mathbf{C}'}$ as the class of functions that satisfy (A2') with respect to parameters $\mathbf{C}' = (c'_0, C'_0) > 0$. Intuitively, (A2') can be regarded as a converse of (A2).

We are now ready to state our main negative result, which shows, from an information-theoretic perspective, that the upper bound in Theorem 1 is not improvable.

Theorem 2. Suppose $\alpha, c_0, C_0, c'_0, C'_0 > 0$ and $\kappa = \infty$. Denote $\mathbf{C} = (c_0, C_0)$ and $\mathbf{C}' = (c'_0, C'_0)$. For any $f_0 \in \Theta_{\mathbf{C}} \cap \Theta'_{\mathbf{C}'}$, define

$$\varepsilon_n^L(f_0) := \sup \left\{ \epsilon > 0 : \epsilon^{-(2+d/\alpha)} \mu_{f_0}(\epsilon) \geq n \right\}. \quad (11)$$

Then there exists a constant $M > 0$ depending on α, d, \mathbf{C} and \mathbf{C}' such that, for any $f_0 \in \Sigma_\kappa^\alpha(M/2) \cap \Theta_{\mathbf{C}} \cap \Theta_{\mathbf{C}'}$,

$$\inf_{\hat{x}_n} \sup_{\substack{f \in \Sigma_\kappa^\alpha(M), \\ \|f - f_0\|_\infty \leq 2\varepsilon_n^L(f_0)}} \Pr_f \left[\mathcal{L}(\hat{x}_n; f) \geq \varepsilon_n^L(f_0) \right] \geq \frac{1}{3}. \quad (12)$$

Remark 4. We note in passing that for any f_0 and n it always holds that $\varepsilon_n^L(f_0) \leq \varepsilon_n^U(f_0)$.

Remark 5. If the distribution function $\mu_{f_0}(\epsilon)$ satisfies (9) (i.e. it does not change too abruptly) in Remark 3, then $\varepsilon_n^L(f_0) \geq \varepsilon_n^U(f_0)/[\log n]^{O(1)}$. Consequently, the upper and lower bounds for these functions match up to logarithmic factors.

The following proposition derives an explicit expression for $\varepsilon_n^L(f_0)$ for reference functions whose distribution functions have a polynomial growth, which matches the upper bound in Proposition 2 up to $\log n$ factors. The proof of this Proposition is straightforward and is omitted.

Proposition 3. Suppose $\mu_{f_0}(\epsilon) \gtrsim \epsilon^\beta$ for some $\beta \in [0, 2 + d/\alpha]$. Then $\varepsilon_n^L(f_0) = \Omega(n^{-\alpha/(2\alpha+d-\alpha\beta)})$.

602 The following proposition additionally shows the existence
 603 of $f_0 \in \Sigma_\infty^\alpha(M) \cap \Theta_C \cap \Theta_{C'}$ that satisfies $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$ for
 604 any values of $\alpha > 0$ and $\beta \in [0, d/\alpha]$. Its proof is given in
 605 the Appendix.

606 **Proposition 4.** *Fix arbitrary $\alpha, M > 0$ and $\beta \in [0, d/\alpha]$.
 607 There exists $f_0 \in \Sigma_\kappa^\alpha(M) \cap \Theta_C \cap \Theta_{C'}$ for $\kappa = \infty$ and constants
 608 $C = (c_0, C_0)$, $C' = (c'_0, C'_0)$ that depend only on α, β, M and
 609 d such that $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$.*

610 Theorem 2 and Proposition 3 show that the $n^{-\alpha/(2\alpha+d-\alpha\beta)}$
 611 upper bound on local minimax convergence rate established in
 612 Theorem 1 is not improvable up to logarithmic factors of n .
 613 Such information-theoretic lower bounds on the convergence
 614 rates hold *even if the data analyst has perfect information of*
 615 f_0 , the reference function on which the $n^{-\alpha/(2\alpha+d-\alpha\beta)}$ local
 616 rate is based. Our results also imply an $n^{-\alpha/(2\alpha+d)}$ minimax
 617 lower bound over all α -Hölder smooth functions, showing that
 618 without additional assumptions, noisy optimization of smooth
 619 functions is as difficult as reconstructing the unknown function
 620 in sup-norm.

621 Our proof of Theorem 2 also differs from those of existing
 622 minimax lower bounds for active nonparametric models [50].
 623 The classical approach is to invoke Fano's inequality and to
 624 upper bound the KL divergence between different underlying
 625 functions f and g using $\|f - g\|_\infty$, corresponding to the
 626 point $x \in \mathcal{X}$ that leads to the largest KL divergence. Such
 627 an approach, however, does not produce tight lower bounds
 628 for our problem. To overcome such difficulties, we borrow
 629 the lower bound analysis for bandit pure exploration problems
 630 in [51]. In particular, our analysis considers the query distribution
 631 of any active query algorithm $\mathcal{A} = (\varphi_1, \dots, \varphi_n, \phi_n)$
 632 under the reference function f_0 and bounds the perturbation in
 633 query distributions between f_0 and f using Le Cam's lemma.
 634 Afterwards, an adversarial function choice f can be made
 635 based on the query distributions of the considered algorithm \mathcal{A} .
 636 We defer the complete proof of Theorem 2 to Section V-C.

637 Theorem 2 applies to any global optimization method that
 638 makes *active* queries, corresponding to the query model in
 639 Definition 2. The following theorem, on the other hand, shows
 640 that for passive algorithms (Definition 1) the $n^{-\alpha/(2\alpha+d)}$
 641 optimization rate is not improvable even with additional level-set
 642 assumptions imposed on f_0 . This demonstrates an explicit
 643 gap between passive and adaptive query models in global
 644 optimization problems.

645 **Theorem 3.** *Suppose $\alpha, c_0, C_0, c'_0, C'_0 > 0$ and $\kappa = \infty$.
 646 Denote $C = (c_0, C_0)$ and $C' = (c'_0, C'_0)$. Then there exist
 647 constants $M > 0$ depending on α, d, C, C' and N depending
 648 on M such that, for any $f_0 \in \Sigma_\kappa^\alpha(M/2) \cap \Theta_C \cap \Theta_{C'}$ satisfying
 649 $\varepsilon_n^L(f_0) \leq \tilde{\varepsilon}_n^L := [\log n/n]^{\alpha/(2\alpha+d)}$,*

$$650 \inf_{\hat{x}_n} \sup_{\substack{f \in \Sigma_\kappa^\alpha(M), \\ \|f - f_0\|_\infty \leq 2\tilde{\varepsilon}_n^L}} \Pr_f \left[\mathcal{L}(\hat{x}_n; f) \geq \tilde{\varepsilon}_n^L \right] \geq \frac{1}{3} \quad \text{for all } n \geq N. \quad (13)$$

651 Intuitively, the apparent gap demonstrated by Theorems 2
 652 and 3 between the active and passive query models stems from

653 the observation that, a passive algorithm \mathcal{A} only has access
 654 to uniformly sampled query points x_1, \dots, x_n and therefore
 655 cannot focus on a small level-set of f in order to improve
 656 query efficiency. In addition, for functions that grow faster
 657 when moving away from their minima (implying a larger
 658 value of β), the gap between passive and active query models
 659 becomes bigger as active queries can more effectively exploit
 660 the restricted level-sets of such functions.

IV. OUR ALGORITHM

662 In this section we describe a concrete algorithm that attains
 663 the upper bound in Theorem 1. We start with a cleaner
 664 algorithm that operates under the slightly stronger condition
 665 that $\kappa = \infty$ in (A1), meaning that f is α -Hölder smooth on the
 666 entire domain \mathcal{X} . The generalization to $\kappa > 0$ being a constant
 667 is given in Section IV-C with an additional pre-processing step.

668 Let $G_n \in \mathcal{X}$ be a *finite* grid of points in \mathcal{X} . We assume the
 669 finite grid G_n satisfies the following two mild conditions:

670 (B1) Points in G_n are sampled i.i.d. from an unknown distribution
 671 P_X on \mathcal{X} ; furthermore, the density p_X associated
 672 with P_X satisfies $\underline{p}_0 \leq p_X(x) \leq \bar{p}_0$ for all $x \in \mathcal{X}$, where
 673 $0 < \underline{p}_0 \leq \bar{p}_0 < \infty$ are universal constants;
 674 (B2) $|G_n| \gtrsim n^3$ and $\log |G_n| = O(\log n)$.

675 **Remark 6.** *Although typically the choices of the grid points
 676 G_n belong to the data analyst, in some applications the choices
 677 of design points are not completely unconstrained. For exam-
 678 ple, in material synthesis experiments described previously
 679 some environmental parameter settings (e.g., temperature and
 680 pressure) might not be allowed due to budget or physical con-
 681 straints. Thus, we choose to consider less restrictive conditions
 682 imposed on the design grid G_n , allowing it to be more flexible
 683 in real-world applications.*

684 **Remark 7.** *Condition (B2) ensures that the grid G_n is
 685 sufficiently dense, such that even with the smallest bandwidth
 686 our algorithm possibly uses $(h_t(x) = 1/n^2$, see (18)), each
 687 $x \in G_n$ has abundant neighboring points in G_n , so that the
 688 local polynomial estimates in (15) are well-defined.*

689 For any subset $S \subseteq G_n$ and a “weight” function
 690 $\varrho : G_n \rightarrow \mathbb{R}^+$, define the extension $S^\circ(\varrho)$ of S with respect
 691 to ϱ as

$$692 S^\circ(\varrho) := \bigcup_{x \in S} B_{\varrho(x)}^\infty(x; G_n) \quad \text{where} \\ 693 B_{\varrho(x)}^\infty(x; G_n) = \{z \in G_n : \|z - x\|_\infty \leq \varrho(x)\}. \quad (14)$$

694 The algorithm can then be formulated as two levels of iter-
 695 ations, with the outer loop shrinking the “active set” S_τ and
 696 the inner loop collecting data in order to reduce the lengths
 697 of the confidence intervals on the points in the active set.
 698 A pseudocode description of our proposed algorithm is given
 699 in Figure 1.

A. Local Polynomial Regression

700 We use local polynomial regression [5] to obtain the esti-
 701 mate \hat{f} . In particular, for any $x \in G_n$ and a bandwidth

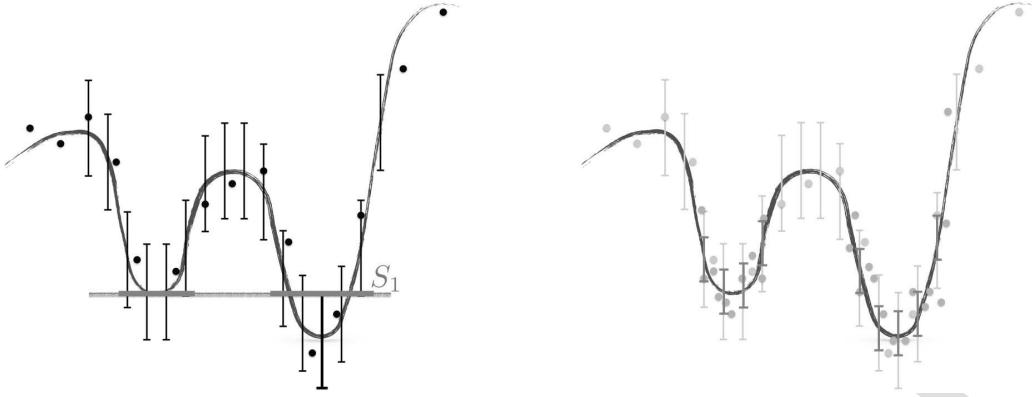


Fig. 1. An informal illustration of Algorithm 1. Solid blue curves depict the underlying function f to be optimized, black and red solid dots denote the query points and their responses $\{(x_t, y_t)\}$, and black and red vertical line segments correspond to uniform confidence intervals on function evaluations constructed using the current batch of data observed. The left figure illustrates the first epoch of our algorithm, where query points are uniformly sampled from the entire domain \mathcal{X} . Afterwards, sub-optimal locations based on the constructed confidence intervals are removed, and a shrunken “candidate set” S_1 is obtained. The algorithm then proceeds to the second epoch, illustrated in the right figure, where query points (in red) are sampled only from the restricted candidate set and shorter confidence intervals (also in red) are constructed and updated. The procedure is repeated until $O(\log n)$ epochs are completed.

Parameters: α, M, δ, n

Output: \hat{x}_n , the final prediction

Initialization: $S_0 = G_n, \varrho_0(x) \equiv \infty, T = \lfloor \log_2 n \rfloor, n_0 = \lfloor n/T \rfloor$;

for $\tau = 1, 2, \dots, T$ **do**

 Compute “extended” sample set $S_{\tau-1}^o(\varrho_{\tau-1})$ defined in (14);

for $t = (\tau - 1)n_0 + 1$ to τn_0 **do**

 Sample x_t uniformly at random from $S_{\tau-1}^o(\varrho_{\tau-1})$ and observe $y_t = f(x_t) + w_t$;

end

 For every $x \in S_{\tau-1}$, compute bandwidth $h_\tau(x)$ using (18) and build the confidence interval

$[\ell_\tau(x), u_\tau(x)]$ in (19);

$S_\tau := \{x \in S_{\tau-1} : \ell_\tau(x) \leq \min_{x' \in S_{\tau-1}} u_\tau(x')\},$

$\varrho_\tau(x) := \min\{\varrho_{\tau-1}(x), h_\tau(x)\},$

end

Final processing: for every $x \in S_T$ let $\hat{f}_{h_T, x}(\cdot)$ be the local polynomial estimates constructed in (15) at x .

Output $\hat{x}_n = \arg \min_{x \in S_T} \min_{x' \in B_{h_T(x)}^\infty(x; \mathcal{X})} \hat{f}_{h_T, x}(x')$.

Algorithm 1 The Main Algorithm

parameter $h > 0$, consider the least squares polynomial estimate

$$\hat{f}_h \in \operatorname{argmin}_{g \in \mathcal{P}_k} \sum_{t'=1}^t \mathbb{I}[x_{t'} \in B_h^\infty(x)] \cdot (y_{t'} - g(x_{t'}))^2, \quad (15)$$

where $B_h^\infty(x) := \{x' \in \mathcal{X} : \|x' - x\|_\infty \leq h\}$ and \mathcal{P}_k denotes all polynomials of degree k on \mathcal{X} .

To analyze the performance of \hat{f}_h evaluated at a certain point $x \in \mathcal{X}$, define the mapping

$$\psi_{x, h} : z \mapsto (1, \psi_{x, h}^1(z), \dots, \psi_{x, h}^k(z))$$

where $\psi_{x, h}^j : z \mapsto [\prod_{\ell=1}^j h^{-1}(z_{i_\ell} - x_{i_\ell})]_{i_1, \dots, i_j=1}^d$ is the degree- j polynomial mapping from \mathbb{R}^d to \mathbb{R}^{d^j} . Also define $\Psi_{t, h} := (\psi_{x, h}(x_{t'}))_{1 \leq t' \leq t, x_{t'} \in B_h(x)}$ as the $m \times D$ aggregated

design matrix, where $m = \sum_{t'=1}^t \mathbb{I}[x_{t'} \in B_h^\infty(x)]$ and $D = 1 + d + \dots + d^k$, $k = \lfloor \alpha \rfloor$. The estimate \hat{f}_h defined in (15) then admits the following closed-form expression:

$$\hat{f}_h(z) \equiv \psi_{x, h}(z)^\top (\Psi_{t, h}^\top \Psi_{t, h})^\dagger \Psi_{t, h}^\top Y_{t, h}, \quad (16)$$

where $Y_{t, h} = (y_{t'})_{1 \leq t' \leq t, x_{t'} \in B_h^\infty(x)}$ and A^\dagger is the Moore-Penrose pseudo-inverse of A .

The following lemma gives a finite-sample analysis of the error of $\hat{f}_h(x)$:

Lemma 1. Suppose that f satisfies (6) on $B_h^\infty(x; \mathcal{X})$, $\max_{z \in B_h^\infty(x; \mathcal{X})} \|\psi_{x, h}(z)\|_2 \leq b$ and $\frac{1}{m} \Psi_{t, h}^\top \Psi_{t, h} \geq \sigma I_{D \times D}$ for some $\sigma > 0$. Then for any $\delta \in (0, 1/2)$, with probability $1 - \delta$

$$|\hat{f}_h(x') - f(x')| \leq \underbrace{\frac{b^2}{\sigma} M d^k h^\alpha}_{\mathfrak{b}_{h, \delta}(x)} + \underbrace{b \sqrt{\frac{5D \ln(1/\delta)}{\sigma m}}}_{\mathfrak{s}_{h, \delta}(x)} =: \eta_{h, \delta}(x), \quad \forall x' \in B_h^\infty(x; \mathcal{X}). \quad (17)$$

Remark 8. $\mathfrak{b}_{h, \delta}(x)$, $\mathfrak{s}_{h, \delta}(x)$ and $\eta_{h, \delta}(x)$ depend on x because σ depends on $\Psi_{t, h}$, which further depends on the sample points in the neighborhood $B_h^\infty(x; \mathcal{X})$ of x .

In the rest of the paper we define $\mathfrak{b}_{h, \delta}(x) := (b^2/\sigma) M d^k h^\alpha$ and $\mathfrak{s}_{h, \delta}(x) := b \sqrt{5D \ln(1/\delta)/\sigma m}$ as the bias and standard deviation terms in the error of $\hat{f}_h(x)$, respectively. We also denote $\eta_{h, \delta}(x) := \mathfrak{b}_{h, \delta}(x) + \mathfrak{s}_{h, \delta}(x)$ as the overall error in $\hat{f}_h(x)$.

Notice that when bandwidth h increases, the bias term $\mathfrak{b}_{h, \delta}(x)$ increases because of the h^α term; on the other hand, with h increasing the local neighborhood $B_h^\infty(x; \mathcal{X})$ grows and would potentially contain more samples, implying a larger m and smaller standard deviation term $\mathfrak{s}_{h, \delta}(x)$. A careful selection of the bandwidth h balances $\mathfrak{b}_{h, \delta}(x)$ and $\mathfrak{s}_{h, \delta}(x)$ and yields appropriate confidence intervals on $f(x)$, and we turn our attention to this in the next section.

745 **B. Bandwidth Selection and Confidence Intervals**

746 Given the expressions of bias $\mathfrak{b}_{h,\delta}(x)$ and standard deviation
 747 $\mathfrak{s}_{h,\delta}(x)$ in (17), the bandwidth $h_\tau(x) > 0$ at epoch τ and point
 748 x is selected as

$$749 \quad h_\tau(x) := \frac{j_\tau(x)}{n^2} \text{ where } j_\tau(x) := \arg \max \left\{ j \in \mathbb{N}^+ : \mathfrak{b}_{j/n^2,\delta}(x) \leq \mathfrak{s}_{j/n^2,\delta}(x) \right\}. \quad (18)$$

751 More specifically, $h_\tau(x)$ is the largest positive value in an
 752 evenly spaced grid $\{j/n^2\}$ such that the bias of $\hat{f}_{h_\tau}(x)$ is
 753 smaller than its standard deviation. This bandwidth selection
 754 is in principle similar to the Lepski's method [52], with the
 755 exception that an upper bound on the bias for any bandwidth
 756 parameter is known and does not need to be estimated from
 757 data.

758 With the selection of bandwidth $h_\tau(x)$ at epoch τ and
 759 query point x , a confidence interval on $f(x')$ for all $x' \in$
 760 $B_{h_\tau(x)}^\infty(x; \mathcal{X})$ is constructed as

$$761 \quad \ell_\tau(x) := \max_{1 \leq t' \leq \tau} \sup_{x' \in B_{h_\tau(x)}^\infty(x; \mathcal{X})} \left\{ \hat{f}_{h_{t'}(x)}(x') - \eta_{h_{t'}(x), \delta}(x) \right\};$$

$$762 \quad u_\tau(x) := \min_{1 \leq t' \leq \tau} \inf_{x' \in B_{h_\tau(x)}^\infty(x; \mathcal{X})} \left\{ \hat{f}_{h_{t'}(x)}(x') + \eta_{h_{t'}(x), \delta}(x) \right\}. \quad (19)$$

763 Note that for any $x \in \mathcal{X}$, the lower confidence edge $\ell_\tau(x)$ is
 764 a non-decreasing function in τ and the upper confidence edge
 765 $u_\tau(x)$ is a non-increasing function in τ .

767 **C. Pre-processing**

768 We describe a pre-processing step that relaxes the smoothness
 769 condition from $\kappa = \infty$ to $\kappa = \Omega(1)$, meaning that only
 770 local smoothness of f around its minimum values is required.
 771 Let $n_0 = \lfloor n/\log n \rfloor$, x_1, \dots, x_{n_0} be points i.i.d. uniformly sam-
 772 pled from \mathcal{X} and y_1, \dots, y_{n_0} be their corresponding responses.
 773 For every grid point $x \in G_n$, perform the following:

- 774 1) Compute $\hat{f}_x(\cdot)$ as the local polynomial fits of all y_i
 775 corresponding to $\|x_i - x\|_\infty \leq n_0^{-1/2d} \log^3 n =: h_0$;
- 776 2) Compute $\bar{f}(x)$ as the sample average of all y_i corre-
 777 sponding to $\|x_i - x\|_\infty \leq h_0$;
- 778 3) Remove all $x \in G_n$ from S_0 if $\bar{f}(x) \geq$
 779 $\min_{z \in G_n} \inf_{z' \in B_{h_0}^\infty(z; \mathcal{X})} \hat{f}_z(z') + 1/\log n$.

780 **Remark 9.** The $1/\log n$ term in the removal condition $\bar{f}(x) \geq$
 781 $\min_{z \in G_n} \bar{f}(z) + 1/\log n$ is not important, and can be replaced
 782 with any sequence $\{\omega_n\}$ such that $\lim_{n \rightarrow \infty} \omega_n = 0$ and
 783 $\lim_{n \rightarrow \infty} \omega_n n^t = \infty$ for any $t > 0$. The readers are referred to
 784 the proof of Proposition 5 in the appendix for the motivation
 785 of this term as well as the selection of the pre-processing
 786 bandwidth h_0 .

787 To analyze the pre-processing step, we state the following
 788 proposition:

789 **Proposition 5.** Assume $f \in \Sigma_\kappa^\alpha(M)$ and let S'_0 be the screened
 790 grid after step 2 of the pre-processing procedure. Then for
 791 sufficiently large n , with probability $1 - O(n^{-1})$ we have

$$792 \quad B_{h_0}^\infty(x; \mathcal{X}) \cap L_f(\kappa/2) \neq \emptyset, \quad \forall x \in S'_0, \quad (20)$$

793 where $L_f(\kappa/2) = \{x \in \mathcal{X} : f(x) \leq f^* + \kappa/2\}$.

794 To interpret Proposition 5, note that for sufficiently large n ,
 795 $f \in \Sigma_\kappa^\alpha(M)$ implies f being α -Hölder smooth (i.e., f
 796 satisfies (6)) on $\bigcup_{x \in L_f(\kappa/2)} B_{h_0}^\infty(x; \mathcal{X})$, because $\kappa > 0$ is a
 797 constant and $h_0 \rightarrow 0$ as $n \rightarrow \infty$. Subsequently, the proposition
 798 shows that with high probability, the pre-processing step will
 799 remove all grid points in G_n in non-smooth regions of f ,
 800 while maintaining the global optimal solution. This justifies
 801 the pre-processing step for $f \in \Sigma_\kappa^\alpha(M)$, because f is smooth
 802 on the grid and its close neighborhood after pre-processing.

803 The proof of Proposition 5 uses the fact that the local
 804 mean estimation is large provided that all data points in the
 805 local mean estimator are large, regardless of their underlying
 806 smoothness. The complete proof of Proposition 5 is deferred
 807 to the Appendix.

808 **V. PROOFS OF MAIN THEOREMS**809 **A. Proof of Lemma 1**

810 Our proof closely follows the analysis of asymptotic con-
 811 vergence rates for series estimators in [53]. We further work
 812 out all constants in the error bounds to arrive at a com-
 813 pletely finite-sample result, which is then used to construct
 814 finite-sample confidence intervals.

815 We start with polynomial interpolation results for all
 816 Hölder smooth functions in $B_{h_t}^\infty(x; \mathcal{X})$.

817 **Lemma 2.** Suppose f satisfies (6) on $B_h^\infty(x; \mathcal{X})$. Then there
 818 exists $\tilde{f}_x \in \mathcal{P}_k$ such that

$$819 \quad \sup_{z \in B_h^\infty(x; \mathcal{X})} |f(z) - \tilde{f}_x(z)| \leq M d^k h^\alpha. \quad (21)$$

821 *Proof.* Consider

$$822 \quad \tilde{f}_x(z) := f(x) + \sum_{j=1}^k \sum_{\alpha_1+\dots+\alpha_d=j} \frac{\partial^j f(x)}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} \prod_{\ell=1}^d (z_\ell - x_\ell)^{\alpha_\ell}. \quad (22)$$

823 By Taylor expansion with Lagrangian remainders, there exists
 824 $\xi \in (0, 1)$ such that

$$825 \quad |\tilde{f}_x(z) - f(z)| \leq \sum_{\alpha_1+\dots+\alpha_d=k} |f^{(\alpha)}(x + \xi(z-x)) - f^{(\alpha)}(x)| \cdot \prod_{\ell=1}^d |z_\ell - x_\ell|^{\alpha_\ell}.$$

827 Because f satisfies (6) on $B_h^\infty(x; \mathcal{X})$, we have that $|f^{(\alpha)}(x +$
 828 $\xi(z-x)) - f^{(\alpha)}(x)| \leq M \cdot \|z - x\|_\infty^{\alpha-k}$. Also note that $|z_\ell - x_\ell| \leq$
 829 $\|z - x\|_\infty \leq h$ for all $z \in B_h^\infty(x; \mathcal{X})$. The lemma is thus
 830 proved. \square

831 Using (16), the local polynomial estimate \hat{f}_h can be written
 832 as $\hat{f}_h(z) \equiv \psi_{x,h}(z)^\top \hat{\theta}_h$, where

$$833 \quad \hat{\theta}_h = (\Psi_{t,h}^\top \Psi_{t,h})^{-1} \Psi_{t,h}^\top Y_{t,h}. \quad (23)$$

834 In addition, because $\tilde{f}_x \in \mathcal{P}_k$, there exists $\tilde{\theta} \in \mathbb{R}^D$
 835 such that $\tilde{f}_x(z) \equiv \psi_{x,h}(z)^\top \tilde{\theta}$. Denote also that $F_{t,h} :=$
 836 $(f(x_{t'}))_{1 \leq t' \leq t, x_{t'} \in B_h^\infty(x)}$, $\Delta_{t,h} := (f(x_{t'}) -$
 $\tilde{f}_x(x_{t'}))$

837 $1 \leq t' \leq t, x_{t'} \in B_h^\infty(x)$ and $W_{t,h} := (w_{t'})_{1 \leq t' \leq t, x_{t'} \in B_h^\infty(x)}$. (23) can
838 then be re-formulated as

$$839 \quad \hat{\theta}_h = (\Psi_{t,h}^\top \Psi_{t,h})^{-1} \Psi_{t,h}^\top [\Psi_{t,h} \tilde{\theta} + \Delta_{t,h} + W_{t,h}] \quad (24)$$

$$840 \quad = \tilde{\theta} + \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \left[\frac{1}{m} \Psi_{t,h}^\top (\Delta_{t,h} + W_{t,h}) \right]. \quad (25)$$

841 Because $\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \geq \sigma I_{D \times D}$ and $\sup_{z \in B_h^\infty(x)} \|\psi_{x,h}(z)\|_2 \leq b$, we have that
842

$$843 \quad \|\hat{\theta}_h - \tilde{\theta}\|_2 \leq \frac{b}{\sigma} \|\Delta_{t,h}\|_\infty + \left\| \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \frac{1}{m} \Psi_{t,h}^\top W_t \right\|_2. \quad (26)$$

844 Invoking Lemma 2 we have $\|\Delta_{t,h}\|_\infty \leq M d^k h^\alpha$. In addition,
845 because $W_t \sim \mathcal{N}_m(0, I_{m \times n})$, we have that

$$846 \quad \left\| \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \frac{1}{m} \Psi_{t,h}^\top W_t \right\| \sim \mathcal{N}_D \left(0, \frac{1}{m} \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \right). \quad (27)$$

848 Applying concentration inequalities for quadratic forms of
849 Gaussian random vectors (Lemma 10), with probability $1 - \delta$
850 it holds that

$$851 \quad \left\| \left[\frac{1}{m} \Psi_{t,h}^\top \Psi_{t,h} \right]^{-1} \frac{1}{m} \Psi_{t,h}^\top W_t \right\|_2 \leq \sqrt{\frac{5D \log(1/\delta)}{\sigma m}}. \quad (28)$$

852 We then have that with probability $1 - \delta$ that

$$853 \quad \|\hat{\theta}_h - \tilde{\theta}\|_2 \leq \frac{b}{\sigma_h} M d^k h^\alpha + \sqrt{\frac{5D \log(1/\delta)}{\sigma m}}. \quad (29)$$

854 Finally, noting that for all $x' \in B_h^\infty(x; \mathcal{X})$, $\|\psi_{x,h}(x')\|_2 \leq b$
855 by definition, we have that

$$856 \quad |\hat{f}_h(x') - f(x')| = |\hat{f}_h(x') - \tilde{f}_x(x')| \\ 857 \quad = |\psi_{x,h}(x')^\top (\hat{\theta}_h - \tilde{\theta})| \leq b \|\hat{\theta}_h - \tilde{\theta}\|_2,$$

858 which completes the proof of Lemma 1.

859 B. Proof of Theorem 1

860 In this section we prove Theorem 1. We prove the theorem
861 by considering every reference function $f_0 \in \Sigma_\kappa^\alpha(M) \cap \Theta_C$
862 separately. For simplicity, we assume $\kappa = \infty$ throughout
863 the proof. The $0 < \kappa < \infty$ can be handled by replacing
864 \mathcal{X} with S_0 which is the grid after the pre-processing step
865 described in Section IV-C. We also suppress dependency on
866 $d, \alpha, M, \mathbf{C}, \underline{p}_0, \bar{p}_0$ in $O(\cdot), \Omega(\cdot), \Theta(\cdot), \gtrsim, \lesssim$ and \asymp notations.
867 We further suppress logarithmic terms of n in $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$
868 notations.

869 The following lemma is our main lemma, which shows
870 that the active set S_τ in our proposed algorithm shrinks
871 geometrically before it reaches a certain level. To simplify
872 notations, denote $\tilde{c}_0 := 10c_0$ and (A2) then hold for all
873 $\epsilon, \delta \in [0, \tilde{c}_0]$ for all $f_0 \in \Theta_C$.

874 **Lemma 3.** For $\tau = 1, \dots, T$ define $\varepsilon_\tau := \max\{\tilde{c}_0 \cdot 2^{-\tau}, C_3[\varepsilon_n^U(f_0) + n^{-1/2} \log^2 n]\}$, where $C_3 > 0$ is a constant depending only on $d, \alpha, M, \underline{p}_0, \bar{p}_0$ and \mathbf{C} . Denote also $\rho_\tau^* := \max_{x \in S_\tau} \varrho_\tau(x)$. Then for sufficiently large n , with

875 probability $1 - O(n^{-1})$ the following holds uniformly for all
876 outer iterations $\tau = 1, \dots, T$:

$$877 \quad B_{\rho_\tau^*}^\infty(x; \mathcal{X}) \cap L_f(\varepsilon_\tau) \neq \emptyset. \quad (30)$$

881 Lemma 3 shows that the level ε_τ in $L_f(\varepsilon_\tau)$ that
882 contains $S_{\tau-1}$ shrinks geometrically, until the condition $\varepsilon_\tau \geq C_3[\varepsilon_n^U(f_0) + n^{-1/2} \log^2 n]$ is violated. If the condition is
883 never violated, then at the end of the last epoch τ^* we
884 have $\varepsilon_{\tau^*} = O(n^{-1})$ because $\tau^* = \log n$. On the other
885 hand, because $S_\tau \subseteq S_{\tau-1}$ always holds, we have $\varepsilon_{\tau^*} \lesssim [C_3[\varepsilon_n^U(f_0) + n^{-1/2} \log^2 n]]^{1/2}$. Combining both cases we have that
886 $\varepsilon_{\tau^*} \lesssim [C_3[\varepsilon_n^U(f_0) + n^{-1/2} \log^2 n]]^{1/2} n + n^{-1}$. Theorem 1 is thus
887 proved.

888 In the rest of this section we prove Lemma 3. We need
889 several technical lemmas and propositions. Except for Proposition
890 6 that is straightforward, the proofs of the other technical
891 lemmas are deferred to the end of this section.

892 Denote $x_n^* := \arg\min_{x \in G_n} f(x)$ as the point on the grid G_n
893 with the smallest objective value. The following proposition
894 shows that with high probability, the confidence intervals
895 constructed in the algorithm are truthful and the successive
896 rejection procedure will never exclude the true optimizer of f
897 on G_n .

898 **Proposition 6.** Suppose $\delta = 1/n^4 |G_n|$. Then with probability
899 $1 - O(n^{-1})$ the following hold:

- 900 1) $f(x') \in [\ell_t(x), u_t(x)]$ for all $1 \leq t \leq n$ and $x \in G_n$,
901 $x' \in B_{h_t(x)}^\infty(x; \mathcal{X})$;
- 902 2) $x_n^* \in S_\tau$ for all $0 \leq \tau \leq n$.

903 *Proof.* The first property is true by applying the union bound
904 over all $t = 1, \dots, n$ and $x \in G_n$. The second property
905 then follows, because $\ell_t(x_n^*) \leq f(x_n^*)$ and $\min_{x \in S_{\tau-1}} u_t(x) \geq f(x_n^*)$ for all τ . \square

906 The following lemma shows that every small box centered
907 around a certain sample point $x \in G_n$ contains a sufficient
908 number of sample points whose least eigenvalue can be
909 bounded with high probability under the polynomial mapping
910 $\psi_{x,h}$ defined in Section III-B.

911 **Lemma 4.** For any $x \in G_n$, $1 \leq m \leq n$ and $h > 0$, let $K_{h,m}^1(x), \dots, K_{h,m}^n(x)$ be n independent point sets, where
912 each point set consists of m points sampled i.i.d. uniformly at
913 random from $B_h^\infty(x; G_n) = G_n \cap B_h^\infty(x; \mathcal{X})$. With probability
914 $1 - O(n^{-1})$ the following holds true uniformly for all $x \in G_n$,
915 $h \in \{j/n^2 : j \in \mathbb{N}, j \leq n^2\}$ and $K_{h,m}^\ell(x)$, $\ell \in [n]$ as $n \rightarrow \infty$:

- 916 1) $\sup_{h>0} \sup_{z \in B_h^\infty(x)} \|\psi_{x,h}(z)\|_2 \asymp \Theta(1)$;
- 917 2) $|B_h^\infty(x; G_n)| \asymp h^d |G_n|$;
- 918 3) $\sigma_{\min}(K_{h,m}^\ell(x)) \asymp \Theta(1)$ for all $m \geq \Omega(\log^2 n)$ and
919 $m \leq |G_n|$, where $\sigma_{\min}(K_{h,m}^\ell(x))$ is the least eigenvalue
920 of $\frac{1}{m} \sum_{z \in K_{h,m}^\ell(x)} \psi_{x,h}(z) \psi_{x,h}(z)^\top$.

921 **Remark 10.** It is possible to improve the concentration result
922 in (48) using the strategies adopted in [35] based on sharper
923 Bernstein type concentration inequalities. Such improvements
924 are, however, not important in establishing the main results of
925 this paper.

The next lemma shows that, the bandwidth h_t selected at the end of each outer iteration τ is near-optimal, being sandwiched between two quantities determined by the size of the active sample grid $\tilde{S}_{\tau-1} := S_{\tau-1}^\circ(\rho_{\tau-1})$.

Lemma 5. *There exist constants $C_1, C_2 > 0$ depending only on $d, \alpha, M, \underline{p}_0, \bar{p}_0$ and \mathbf{C} such that with probability $1 - O(n^{-1})$, the following holds for every outer iteration $\tau \in \{1, \dots, T\}$ and all $x \in S_{\tau-1}$:*

$$C_1[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} - \tau/n \leq \varrho_\tau(x) \leq C_2[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} \log n + \tau/n, \quad (31)$$

where $\tilde{v}_{\tau-1} := |G_n|/|\tilde{S}_{\tau-1}|$.

We are now ready to state the proof of Lemma 3, which is based on an inductive argument over the epochs $\tau = 1, \dots, T$.

Proof. We use induction to prove this lemma. For the base case $\tau = 1$, because $\|f - f_0\|_\infty \leq \varepsilon_n^U(f_0)$ and $\varepsilon_n^U(f_0) \rightarrow 0$ as $n \rightarrow \infty$, it suffices to prove that $B_{\rho_1^*}^\infty(x; \mathcal{X}) \cap L_{f_0}(\tilde{c}_0/4) \neq \emptyset$ for all $x \in S_1$ and sufficiently large n . Because $\tilde{S}_0 = S_0 = G_n$, invoking Lemmas 5 and 1 we have that $|\eta_{h_t(x), \delta}(x)| = \tilde{O}(n^{-\alpha/(2\alpha+d)})$ for all $x \in G_n$ with high probability at the end of the first outer iteration $\tau = 1$. Therefore, for sufficiently large n we conclude that $\sup_{x \in G_n} |\eta_{h_t(x), \delta}(x)| \leq c_0/16$ and hence $B_{\rho_1^*}^\infty(x; \mathcal{X}) \cap L_{f_0}(\tilde{c}_0/4) \neq \emptyset$ for all $x \in S_1$.

We now prove the lemma for $\tau \geq 2$, assuming it holds for $\tau - 1$. We also assume that n (and hence n_0) is sufficiently large, such that the maximum CI length $\max_{x \in G} |\eta_{h_t(x), \delta}(x)|$ after the first outer iteration $\tau = 1$ is smaller than $c_0/2$.

Because $\|f - f_0\|_\infty \leq \varepsilon_n^U(f_0)$ and $\varepsilon_{\tau-1} \geq C_3 \varepsilon_n^U(f_0) \log^2 n$, for appropriately chosen constant C_3 that is not too small, we have that $\|f - f_0\|_\infty \leq \varepsilon_{\tau-1}$. By the inductive hypothesis we have

$$\forall x \in S_{\tau-1}, \quad B_{\rho_{\tau-1}^*}^\infty(x; \mathcal{X}) \cap L_f(\varepsilon_{\tau-1}) \neq \emptyset;$$

Equivalently,

$$\begin{aligned} S_{\tau-1} &\subseteq L_f^\circ(\varepsilon_{\tau-1}, \rho_{\tau-1}^*) \subseteq L_{f_0}^\circ(\varepsilon_{\tau-1} + \|f - f_0\|_\infty, \rho_{\tau-1}^*) \\ &\subseteq L_{f_0}^\circ(2\varepsilon_{\tau-1}, \rho_{\tau-1}^*). \end{aligned} \quad (32)$$

Subsequently,

$$\tilde{S}_{\tau-1} = S_{\tau-1}^\circ \subseteq L_{f_0}^\circ(2\varepsilon_{\tau-1}, 2\rho_{\tau-1}^*). \quad (33)$$

Let $\bigcup_{x \in H_n} B_{2\rho_{\tau-1}^*}^2(x)$ be the smallest covering set of $L_{f_0}(2\varepsilon_{\tau-1})$, meaning that $L_{f_0}(2\varepsilon_{\tau-1}) \subseteq \bigcup_{x \in H_n} B_{2\rho_{\tau-1}^*}^2(x)$, where $B_{2\rho_{\tau-1}^*}^2(x) = \{z \in \mathcal{X} : \|z - x\|_2 \leq 2\rho_{\tau-1}^*\}$ is the ℓ_2 ball of radius $2\rho_{\tau-1}^*$ centered at x . By (A2), we know that $|H_n| \lesssim 1 + [\rho_{\tau-1}^*]^{-d} \mu_{f_0}(2\varepsilon_{\tau-1})$. In addition, the enlarged level-set satisfies $L_{f_0}^\circ(2\varepsilon_{\tau-1}, 2\rho_{\tau-1}^*) \subseteq \bigcup_{x \in H_n} B_{4\rho_{\tau-1}^*}^\infty(x)$. Subsequently,

$$\mu_{f_0}^\circ(2\varepsilon_{\tau-1}, \rho_{\tau-1}^*) \lesssim |H_n| \cdot [\rho_{\tau-1}^*]^d \lesssim \mu_{f_0}(2\varepsilon_{\tau-1}) + [\rho_{\tau-1}^*]^d. \quad (34)$$

By Lemma 5, the monotonicity of $|\tilde{S}_{\tau-1}|$ and the fact that $\underline{p}_0 \leq p_X(z) \leq \bar{p}_0$ for all $z \in \mathcal{X}$, we have

$$\rho_{\tau-1}^* \lesssim [\mu_f^\circ(\varepsilon_{\tau-1}, \rho_{\tau-1}^*)]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log n \quad (35)$$

$$\leq [\mu_{f_0}^\circ(2\varepsilon_{\tau-1}, \rho_{\tau-1}^*)]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log n \quad (36)$$

$$\lesssim \left(\mu_{f_0}(2\varepsilon_{\tau-1}) + [\rho_{\tau-1}^*]^d \right)^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log n. \quad (37)$$

Re-arranging terms on both sides of (37) we have

$$\rho_{\tau-1}^* \lesssim \max \left\{ [\mu_{f_0}(2\varepsilon_{\tau-1})]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log n, n_0^{-1/(2\alpha+d)} \log n \right\}. \quad (38)$$

On the other hand, according to the selection procedure of the bandwidth $h_t(x)$, we have that $\eta_{h_t(x), \delta}(x) \lesssim b_{h_t(x), \delta}(x)$. Invoking Lemma 5 we have for all $x \in S_{\tau-1}$ that

$$\eta_{h_t(x), \delta}(x) \lesssim b_{h_t(x), \delta}(x) \lesssim [h_t(x)]^\alpha \quad (39)$$

$$\lesssim [\tilde{v}_{\tau-1}n_0]^{-\alpha/(2\alpha+d)} \log n \quad (40)$$

$$\lesssim [\tilde{v}_{\tau-2}n_0]^{-\alpha/(2\alpha+d)} \log n \quad (41)$$

$$\lesssim [\rho_{\tau-1}^*]^\alpha \log n. \quad (42)$$

Here (40) holds by invoking the upper bound on $h_t(x)$ in Lemma 5, (41) holds because $\tilde{v}_{\tau-1} \geq \tilde{v}_{\tau-2}$, and (42) holds by again invoking the lower bound on $\varrho_{\tau-1}(x)$ in Lemma 5. Combining Eqs. (38,42) we have

$$\max_{x \in S_{\tau-1}} \eta_{h_t(x), \delta}(x) \quad (43)$$

$$\lesssim \max \left\{ [\mu_{f_0}(2\varepsilon_{\tau-1})]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log^2 n, n_0^{-1/(2\alpha+d)} \log n \right\}. \quad (44)$$

Recall that $n_0 = n/\log n$ and $\varepsilon_n^U(f_0) \leq \varepsilon_{\tau-1}$, provided that C_3 is not too small. By definition, every $\varepsilon \geq \varepsilon_n^U(f_0)$ satisfies $\varepsilon^{-(2+d/\alpha)} \mu_{f_0}(\varepsilon) \leq n/\log^\omega n$ for some large constant $\omega > 5 + d/\alpha$. Subsequently,

$$[\mu_{f_0}(2\varepsilon_{\tau-1})]^{1/(2\alpha+d)} n_0^{-1/(2\alpha+d)} \log^2 n \quad (1002)$$

$$\lesssim 2\varepsilon_{\tau-1} n^{1/(2\alpha+d)} \log^{-\frac{\omega\alpha}{2\alpha+d}} n \cdot n_0^{-1/(2\alpha+d)} \log^2 n \quad (1003)$$

$$\lesssim \varepsilon_{\tau-1} / [\log n]^{\frac{(\omega-5-d/\alpha)\alpha}{2\alpha+d}}. \quad (1004)$$

Because $\omega > 5 + d/\alpha$, the right-hand side of (46) is asymptotically dominated⁶ by $\varepsilon_{\tau-1}$. In addition, $n_0^{-1/(2\alpha+d)} \log n$ is also asymptotically dominated by $\varepsilon_{\tau-1}$ because $\varepsilon_{\tau-1} \geq C_3 n^{-1/2} \log^\omega n$. Therefore, for sufficiently large n we have

$$\max_{x \in S_{\tau-1}} \eta_{h_t(x), \delta}(x) \leq \varepsilon_{\tau-1}/4. \quad (47)$$

Lemma 3 is thus proved. \square

⁶We say $\{a_n\}$ is asymptotically dominated by $\{b_n\}$ if $\lim_{n \rightarrow \infty} |a_n|/|b_n| = 0$.

1011 *1) Proof of Lemma 4:*

1012 *Proof.* We first show that the first property holds almost surely.
1013 Recall the definition of $\psi_{x,h}$, we have that $1 \leq \|\psi_{x,h}(z)\|_2 \leq$
1014 $D \cdot [\max_{1 \leq j \leq d} h^{-1} |z_j - x_j|]^k$. Because $\|z - x\|_\infty \leq h$ for
1015 all $z \in B_h^\infty(x)$, $\sup_{z \in B_h^\infty(x)} \|\psi_{x,h}(z)\|_2 \lesssim O(1)$ for all $h > 0$.
1016 Thus, $\sup_{h>0} \sup_{z \in B_h^\infty(x)} \|\psi_{x,h}(z)\|_2 \asymp \Theta(1)$ for all $x \in G_n$.

1017 For the second property, by Hoeffding's inequality
1018 (Lemma 9) and the union bound, with probability $1 - O(n^{-1})$
1019 we have that

$$1020 \max_{x,h} \left| \frac{|B_h^\infty(x; G_n)|}{|G_n|} - P_X(z \in B_h^\infty(x)) \right| \lesssim \sqrt{\frac{\log n}{|G_n|}}. \quad (48)$$

1021 In addition, note that $P_X(z \in B_h^\infty(x; \mathcal{X})) \geq$
1022 $\underline{p}_0 \lambda(B_h^\infty(x; \mathcal{X})) \gtrsim h^d$ and $P_X(z \in B_h^\infty(x; \mathcal{X})) \leq$
1023 $\bar{p}_0 \lambda(B_h^\infty(x; \mathcal{X})) \lesssim h^d$, where $\lambda(\cdot)$ denotes the Lebesgue
1024 measure on \mathcal{X} . Subsequently, $|B_h^\infty(x; G_n)|$ is lower bounded by
1025 $\Omega(h^d |G_n| - \sqrt{|G_n| \log n})$ and upper bounded by
1026 $O(h^d |G_n| + \sqrt{|G_n| \log n})$. The second property is then
1027 proved by noting that $h_d \gtrsim n^{-d}$ and $|G_n| \gtrsim n^{3d/\min(\alpha, 1)}$.

1028 We next prove the third property. Because $\underline{p}_0 \leq p_X(z) \in \bar{p}_0$
1029 for all $z \in \mathcal{X}$, we have that

$$1030 \underline{p}_0 \int_{B_h^\infty(x; \mathcal{X})} \psi_{x,h}(z) \psi_{x,h}(z)^\top dU_{x,h}(z) \\ 1031 \leq \mathbb{E} \left[\frac{1}{m} \sum_{z \in K_{h,m}^\ell} \psi_{x,h}(z) \psi_{x,h}(z)^\top \right] \quad (49)$$

$$1032 \leq \bar{p}_0 \int_{B_h^\infty(x; \mathcal{X})} \psi_{x,h}(z) \psi_{x,h}(z)^\top dU_{x,h}(z), \quad (50)$$

1033 where $U_{x,h}$ is the uniform distribution on $B_h^\infty(x; \mathcal{X})$. Note
1034 also that

$$1035 \int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z) \\ 1036 \leq \int_{B_h^\infty(x; \mathcal{X})} \psi_{x,h}(z) \psi_{x,h}(z)^\top dU_{x,h}(z) \quad (51)$$

$$1037 \leq 2^d \int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z) \quad (52)$$

1038 where U is the uniform distribution on $\mathcal{X} = [0, 1]^d$. The
1039 following proposition upper and lower bounds the eigenvalues
1040 of $\int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z)$, which is proved in the appendix.

1041 **Proposition 7.** *There exist constants $0 < \psi_0 \leq \Psi_0 < \infty$
1042 depending only on d, D such that*

$$1043 \psi_0 I_{D \times D} \leq \int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z) \leq \Psi_0 I_{D \times D}. \quad (53)$$

1044 Using Proposition 7 and Eqs. (51,52), we conclude that

$$1045 \Omega(1) \cdot I_{D \times D} \leq \mathbb{E} \left[\frac{1}{m} \sum_{z \in K_{h,m}^\ell} \psi_{x,h}(z) \psi_{x,h}(z)^\top \right] \leq O(1) \cdot I_{D \times D}. \quad (54)$$

1046 Applying matrix Chernoff bound (Lemma 11) and the union
1047 bound, we have that with probability $1 - O(n^{-1})$,
1048

$$1049 \max_{x,h,m,\ell} \left\| \frac{1}{m} \sum_{z \in K_{h,m}^\ell(x)} \psi_{x,h}(z) \psi_{x,h}(z)^\top \right\|_{\text{op}} \\ 1050 - \mathbb{E} \left[\psi_{x,h}(z) \psi_{x,h}(z)^\top \mid z \in B_h^\infty(x) \right] \lesssim \sqrt{\frac{\log n}{m}}. \quad 1051$$

1052 Combining Eqs. (54,55) and applying Weyl's inequality
1053 (Lemma 12) we have

$$1054 \Omega(1) - O(\sqrt{\log n/m}) \lesssim \sigma_{\min}(K_{h,m}^\ell(x)) \\ 1055 \lesssim O(1) - O(\sqrt{\log n/m}). \quad 1055$$

1056 The third property is therefore proved. \square 1056

1057 *2) Proof of Lemma 5: Proof.* We use induction to prove this
1058 lemma. For the base case of $\tau = 1$, we have $\tilde{S}_0 = S_0 = G_n$
1059 and therefore $\tilde{v}_{\tau-1} = 1$. Furthermore, applying Lemma 4 we
1060 have that for all $h = j/n^2$,

$$1061 \mathfrak{b}_{h,\delta}(x) \asymp h^\alpha, \quad \mathfrak{s}_{h,\delta}(x) \asymp \sqrt{\frac{\log n}{h^d n_0}}. \quad 1061$$

1062 Thus, for h selected according to (18) as the largest bandwidth
1063 of the form j/n^2 , $j \in \mathbb{N}$ such that $\mathfrak{b}_{h,\delta}(x) \leq \mathfrak{s}_{h,\delta}(x)$,
1064 both $\mathfrak{b}_{h,\delta}(x), \mathfrak{s}_{h,\delta}(x)$ are on the order of $n_0^{-1/(2\alpha+d)}$ up to
1065 logarithmic terms of n , and therefore one can pick appropriate
1066 constants $C_1, C_2 > 0$ such that $C_1 n_0^{-1/(2\alpha+d)} \leq \varrho_1(x) \leq$
1067 $C_2 n_0^{-1/(2\alpha+d)} \log n$ holds for all $x \in G_n$. 1067

1068 We next prove the lemma for $\tau > 1$, assuming it holds
1069 for $\tau - 1$. We first establish the lower bound part. Define
1070 $\rho_{\tau-1}^* := \min_{z \in S_{\tau-1}} \varrho_{\tau-1}(z)$. By inductive hypothesis, $\rho_{\tau-1}^* \geq$
1071 $C_1 [\tilde{v}_{\tau-2} n_0]^{-1/(2\alpha+d)} - (\tau - 1)/n$. Note also that $\tilde{v}_{\tau-1} \geq \tilde{v}_{\tau-2}$
1072 because $\tilde{S}_{\tau-1} \subseteq \tilde{S}_{\tau-2}$, which holds because $S_{\tau-1} \subseteq S_{\tau-2}$
1073 and $\varrho_{\tau-1}(z) \leq \varrho_{\tau-2}(z)$ for all z . Let h_t^* be the smallest
1074 number of the form j_t^*/n^2 , $j_t^* \in [n^2]$ such that $h_t^* \geq$
1075 $C_1 [\tilde{v}_{\tau-1} n_0]^{-1/(2\alpha+d)} - \tau/n$. We then have $h_t^* \leq \rho_{\tau-1}^*$ and
1076 therefore query points in epoch τ are uniformly distributed in
1077 $B_{h_t^*}^\infty(x; G_n)$. Subsequently, applying Lemma 4 we have with
1078 probability $1 - O(n^{-1})$ that

$$1079 \mathfrak{b}_{h_t^*,\delta}(x) \leq C' [h_t^*]^\alpha, \quad \mathfrak{s}_{h_t^*,\delta}(x) \geq C'' \sqrt{\frac{\log n}{[h_t^*]^d \tilde{v}_{\tau-1} n}}, \quad 1079$$

1080 where $C', C'' > 0$ are constants that depend on
1081 $d, \alpha, M, \underline{p}_0, \bar{p}_0$ and \mathbf{C} , but not C_1, C_2, τ or h_t^* . By choosing
1082 C_1 appropriately (depending on C' and C'') we can make
1083 $\mathfrak{b}_{h_t^*,\delta}(x) \leq \mathfrak{s}_{h_t^*,\delta}(x)$ holds for all $x \in S_{\tau-1}$, thus establishing
1084 $\varrho_\tau(x) \geq \min\{\varrho_{\tau-1}(x), h_t^*\} \geq C_1 [\tilde{v}_{\tau-1} n_0]^{-1/(2\alpha+d)} - \tau/n$. 1084

1085 We next prove the upper bound part. For any $h_t = j_t/n^2$
1086 where $j_t \in [n^2]$, invoking Lemma 4 we have that

$$1087 \mathfrak{b}_{h,\delta}(x) \geq \tilde{C}' h^\alpha, \quad \mathfrak{s}_{h,\delta}(x) \leq \tilde{C}'' \sqrt{\frac{\log n}{\min\{h, \rho_{\tau-1}^*\}^d \cdot \tilde{v}_{\tau-1} n}}, \quad 1087$$

1088 (58)

where \tilde{C}' and \tilde{C}'' are again constants depending on $d, \alpha, M, \underline{p}_0, \bar{p}_0$ and \mathbf{C} , but not C_1, C_2 . Note also that $\rho_{\tau-1}^* \geq C_1[\tilde{v}_{\tau-2}n_0]^{-1/(2\alpha+d)} - (\tau-1)/n \geq C_1[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} - \tau/n$, because $\tilde{v}_{\tau-1} \geq \tilde{v}_{\tau-2}$. By selecting constant $C_2 > 0$ carefully (depending on \tilde{C}', \tilde{C}'' and C_1), we can ensure $b_{h,\delta}(x) > s_{h,\delta}(x)$ for all $h \geq C_2[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} + \tau/n$. Therefore, $\varrho_\tau(x) \leq h_\tau(x) \leq C_2[\tilde{v}_{\tau-1}n_0]^{-1/(2\alpha+d)} + \tau/n$. \square

1096 C. Proof of Theorem 2

1097 In this section we prove the main negative result in Theorem 2. To simplify presentation, we suppress dependency on 1098 α, d, c_0 and C_0 in $\lesssim, \gtrsim, \asymp, O(\cdot)$ and $\Omega(\cdot)$ notations. However, 1099 we do *not* suppress dependency on \underline{C}_R or M in any of the 1100 above notations.

1101 Let $\varphi_0 : [-2, 2]^d \rightarrow \mathbb{R}^*$ be a non-negative function 1102 defined on \mathcal{X} such that $\varphi_0 \in \Sigma_{\kappa}^{[\alpha]}(1)$ with $\kappa = \infty$, 1103 $\sup_{x \in \mathcal{X}} \varphi_0(x) = \Omega(1)$ and $\varphi_0(z) = 0$ for all $\|z\|_2 \geq 1$. Here 1104 $[\alpha]$ denotes the smallest integer that upper bounds α . Such 1105 functions exist and are the cornerstones of the construction of 1106 information-theoretic lower bounds in nonparametric estimation 1107 problems [50]. One typical example is the ‘‘smoothstep’’ 1108 function (see for example [54])

$$1110 S_N(x) := \frac{1}{Z} x^{N+1} \sum_{n=0}^N \binom{N+n}{n} \binom{2N+1}{N-n} (-x)^n, \\ 1111 N = 0, 1, 2, \dots,$$

1112 where $Z > 0$ is a scaling parameter. The smoothstep function 1113 S_N is defined on $[0, 1]$ and satisfies the Hölder condition in (6) 1114 of order $\alpha = N$ on $[0, 1]$. It can be easily extended to $\tilde{S}_{N,d} : 1115 [-2, 2]^d \rightarrow \mathbb{R}$ by considering $\tilde{S}_{N,d}(x) := 1/Z - S_N(a\|x\|_1)$ 1116 where $\|x\|_1 = |x_1| + \dots + |x_d|$ and $a = 1/(2d)$. It is easy 1117 to verify that, with Z chosen appropriately, $\tilde{S}_{N,d} \in \Sigma_{\infty}^N(1)$, 1118 $\sup_{x \in \mathcal{X}} \tilde{S}_{N,d}(x) = 1/Z = \Omega(1)$ and $\tilde{S}_{N,d}(z) = 0$ for all 1119 $\|z\|_2 \geq 1$, where $M > 0$ is a constant.

1120 For any $x \in \mathcal{X}$ and $h > 0$, define $\varphi_{x,h} : \mathcal{X} \rightarrow \mathbb{R}^*$ as

$$1121 \varphi_{x,h}(z) := \mathbb{I}[z \in B_h^{\infty}(x)] \cdot \frac{Mh^{\alpha}}{2} \varphi_0\left(\frac{z-x}{h}\right). \quad (59)$$

1122 It is easy to verify that $\varphi_{x,h} \in \Sigma_{\infty}^{\alpha}(M/2)$, and furthermore 1123 $\sup_{z \in \mathcal{X}} \varphi_{x,h}(z) \asymp Mh^{\alpha}$ and $\varphi_{x,h}(z) = 0$ for all $z \notin B_h^{\infty}(x)$.

1124 Let $L_{f_0}(\varepsilon_n^L(f_0))$ be the level-set of f_0 at $\varepsilon_n^L(f_0)$. Let $H_n \subseteq 1125 L_{f_0}(\varepsilon_n^L(f_0))$ be the largest *packing* set such that $B_h^{\infty}(x)$ are 1126 disjoint for all $x \in H_n$, and $\bigcup_{x \in H_n} B_h^{\infty}(x) \subseteq L_{f_0}(\varepsilon_n^L(f_0))$. 1127 By (A2') and the definition of $\varepsilon_n^L(f_0)$, we have that

$$1128 |H_n| \geq M(L_{f_0}(\varepsilon_n^L(f_0)), 2\sqrt{d}h) \\ 1129 \gtrsim \mu_{f_0}(\varepsilon_n^L(f_0)) \cdot h^{-d} \geq [\varepsilon_n^L(f_0)]^{2+d/\alpha} \cdot nh^{-d}. \quad (60)$$

1130 For any $x \in H_n$, construct $f_x : \mathcal{X} \rightarrow \mathbb{R}$ as

$$1131 f_x(z) := f_0(z) - \varphi_{x,h}(z). \quad (61)$$

1132 Let $\mathcal{F}_n := \{f_x : x \in H_n\}$ be the class of functions indexed 1133 by $x \in H_n$. Let also $h \asymp (\varepsilon_n^L(f_0)/M)^{1/\alpha}$ such that $\|\varphi_{x,h}\|_{\infty} = 1134 2\varepsilon_n^L(f_0)$. We then have that $\|f_x - f_0\|_{\infty} \leq 2\varepsilon_n^L(f_0)$ and $f_x \in 1135 \Sigma_{\infty}^{\alpha}(M)$, because $f_0, \varphi_{x,h} \in \Sigma_{\infty}^{\alpha}(M/2)$.

1136 The next lemma shows that, with n adaptive queries to the 1137 noisy zeroth-order oracle $y_t = f(x_t) + w_t$, it is information 1138 theoretically not possible to identify a certain f_x in \mathcal{F}_n with 1139 high probability.

1140 **Lemma 6.** Suppose $|\mathcal{F}_n| \geq 2$. Let $\mathcal{A}_n = (\chi_1, \dots, \chi_n, \phi_n)$ 1141 be an active optimization algorithm operating with a sample 1142 budget n , which consists of samplers $\chi_{\ell} : \{(x_i, y_i)\}_{i=1}^{\ell-1} \mapsto x_{\ell}$ 1143 and an estimator $\phi_n : \{(x_i, y_i)\}_{i=1}^n \mapsto \hat{f}_x \in \mathcal{F}_n$, both can be 1144 deterministic or randomized functions. Then

$$1145 \inf_{\mathcal{A}_n} \sup_{f_x \in \mathcal{F}_n} \Pr_{\hat{f}_x} \left[\hat{f}_x \neq f_x \right] \geq \frac{1}{2} - \sqrt{\frac{n \cdot \sup_{f_x \in \mathcal{F}_n} \|f_x - f_0\|_{\infty}^2}{2|\mathcal{F}_n|}}. \quad (62)$$

1146 **Lemma 7.** There exists constant $M > 0$ depending on 1147 α, d, c_0, C_0 such that the right-hand side of (62) is lower 1148 bounded by $1/3$.

1149 Lemmas 6 and 7 are proved at the end of this section. 1150 Combining both lemmas and noting that for any distinct 1151 $f_x, f_{x'} \in \mathcal{F}_n$ and $z \in \mathcal{X}$, $\max\{\mathcal{L}(z; f_x), \mathcal{L}(z; f_{x'})\} \geq \varepsilon_n^L(f_0)$, 1152 we proved the minimax lower bound formulated in Theorem 2. 1153

1154 1) *Proof of Lemma 6:* Our proof is inspired by the negative 1155 result of multi-arm bandit pure exploration problems estab- 1156 lished in [51].

1157 *Proof.* For any $x \in H_n$, define

$$1158 n_x := \mathbb{E}_{f_0} \left[\sum_{i=1}^n \mathbb{I}[x \in B_h^{\infty}(x)] \right]. \quad (63)$$

1159 Because $B_h^{\infty}(x)$ are disjoint for $x \in H_n$, we have $\sum_{x \in H_n} n_x \leq 1160 n$. Also define, for every $x \in H_n$,

$$1161 \varphi_x := \Pr_{f_0} \left[\hat{f}_x = f_x \right]. \quad (64)$$

1162 Because $\sum_{x \in H_n} \varphi_x = 1$, by pigeonhole principle there is at 1163 most one $x \in H_n$ such that $\varphi_x > 1/2$. Let $x_1, x_2 \in H_n$ 1164 be the points that have the smallest and second smallest n_x . 1165 Then there exists $x \in \{x_1, x_2\}$ such that $\varphi_x \leq 1/2$ and 1166 $n_x \leq 2n/|\mathcal{F}_n|$. By Le Cam’s and Pinsker’s inequality (see, 1167 for example, [4]) we have that

$$1168 \Pr_{f_x} \left[\hat{f}_x = f_x \right] \leq \Pr_{f_0} \left[\hat{f}_x = f_x \right] + d_{\text{TV}}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n}) \quad (65)$$

$$1169 \leq \Pr_{f_0} \left[\hat{f}_x = f_x \right] + \sqrt{\frac{1}{2} \text{KL}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n})} \quad (66)$$

$$1170 = \varphi_x + \sqrt{\frac{1}{2} \text{KL}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n})} \quad (67)$$

$$1171 \leq \frac{1}{2} + \sqrt{\frac{1}{2} \text{KL}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n})}. \quad (68)$$

1172 It remains to upper bound KL divergence of the active 1173 queries made by \mathcal{A}_n . Using the standard lower bound analysis 1174 for active learning algorithms [50], [55] and the fact that

1175 $f_x \equiv f_0$ on $\mathcal{X} \setminus B_h^\infty(x)$, we have

$$1176 \quad \text{KL}(P_{f_0}^{\mathcal{A}_n} \| P_{f_x}^{\mathcal{A}_n}) = \mathbb{E}_{f_0, \mathcal{A}_n} \left[\log \frac{P_{f_0, \mathcal{A}_n}(x_{1:n}, y_{1:n})}{P_{f_x, \mathcal{A}_n}(x_{1:n}, y_{1:n})} \right] \quad (69)$$

$$1177 \quad = \mathbb{E}_{f_0, \mathcal{A}_n} \left[\log \frac{\prod_{i=1}^n P_{f_0}(y_i|x_i) P_{\mathcal{A}_n}(x_i|x_{1:(i-1)}, y_{1:(i-1)})}{\prod_{i=1}^n P_{f_x}(y_i|x_i) P_{\mathcal{A}_n}(x_i|x_{1:(i-1)}, y_{1:(i-1)})} \right] \quad (70)$$

$$1179 \quad = \mathbb{E}_{f_0, \mathcal{A}_n} \left[\log \frac{\prod_{i=1}^n P_{f_0}(y_i|x_i)}{\prod_{i=1}^n P_{f_x}(y_i|x_i)} \right] \quad (71)$$

$$1180 \quad = \mathbb{E}_{f_0, \mathcal{A}_n} \left[\sum_{x_i \in B_h^\infty(x)} \log \frac{P_{f_0}(y_i|x_i)}{P_{f_x}(y_i|x_i)} \right] \quad (72)$$

$$1181 \quad \leq n_x \cdot \sup_{z \in B_h^\infty(x; \mathcal{X})} \text{KL}(P_{f_0}(\cdot|z) \| P_{f_x}(\cdot|z)) \quad (73)$$

$$1182 \quad \leq n_x \cdot \|f_0 - f_x\|_\infty^2. \quad (74)$$

1183 Therefore,

$$1184 \quad \Pr_{f_x} \left[\hat{f}_x = f_x \right] \leq \frac{1}{2} + \sqrt{\frac{1}{4} n_x \varepsilon_n^2} \leq \frac{1}{2} + \sqrt{\frac{n \|f_x - f_0\|_\infty^2}{2 |\mathcal{F}_n|}}. \quad (75)$$

1185 \square

1186 2) Proof of Lemma 7:

1187 *Proof.* By construction, $n \sup_{f_x \in \mathcal{F}_x} \|f_x - f_0\|_\infty^2 \lesssim M^2 nh^{2\alpha}$ and $|\mathcal{F}_n| = |H_n| \gtrsim [\underline{C}_\varepsilon \varepsilon_n^\perp(f_0)]^{2+d/\alpha} nh^{-d}$. Note also that $h \asymp (\varepsilon/M)^{1/\alpha} \asymp (\underline{C}_\varepsilon \varepsilon_n^\perp(f_0)/M)^{1/\alpha}$ because $\|f_x - f_0\|_\infty = \varepsilon = \underline{C}_\varepsilon \varepsilon_n^\perp(f_0)$. Subsequently,

$$1191 \quad \frac{n \sup_{f_x \in \mathcal{F}_x} \|f_x - f_0\|_\infty^2}{2 |\mathcal{F}_n|} \lesssim \frac{n [\underline{C}_\varepsilon \varepsilon_n^\perp(f_0)]^2}{n [\underline{C}_\varepsilon \varepsilon_n^\perp(f_0)]^2 \cdot M^{d/\alpha}} = M^{-d/\alpha}. \quad (76)$$

1192 By choosing the constant $M > 0$ to be sufficiently large, the right-hand side of the above inequality is upper bounded by $1/36$. The lemma is thus proved. \square

1193 D. Proof of Theorem 3

1194 The proof of Theorem 3 is similar to the proof of Theorem 2, but is much more standard by invoking the Fano's inequality [4]. In particular, adapting the Fano's inequality on any finite function class \mathcal{F}_n constructed we have the following lemma:

1202 **Lemma 8** (Fano's inequality). *Suppose $|\mathcal{F}_n| \geq 2$, and $\{(x_i, y_i)\}_{i=1}^n$ are i.i.d. random variables. Then*

$$1204 \quad \inf_{\hat{f}_x} \sup_{f_x \in \mathcal{F}_n} \Pr_{f_x} \left[\hat{f}_x \neq f_x \right] \\ 1205 \quad \geq 1 - \frac{\log 2 + n \cdot \sup_{f_x, f_{x'} \in \mathcal{F}_n} \text{KL}(P_{f_x} \| P_{f_{x'}})}{\log |\mathcal{F}_n|}, \quad (77)$$

1206 where P_{f_x} denotes the distribution of (x, y) under the law of f_x .

1208 Let \mathcal{F}_n be the function class constructed in the previous proof of Theorem 2, corresponding to the largest packing set H_n of $L_{f_0}(\varepsilon_n^\perp)$ such that $B_h^\infty(x)$ for all $x \in H_n$ are disjoint, where $h \asymp (\varepsilon_n^\perp/M)^{1/\alpha}$ such that $\|\varphi_{x,h}\|_\infty = 2\varepsilon_n^\perp$ for

1212 all $x \in H_n$. Because f_0 satisfies (A2'), we have that $|\mathcal{F}_n| = 1213 |H_n| \gtrsim \mu_{f_0}(\varepsilon_n^\perp) h^{-d}$. Under the condition that $\varepsilon_n^\perp(f_0) \leq \varepsilon_n^\perp$, it 1214 holds that $\mu_{f_0}(\varepsilon_n^\perp) \geq [\varepsilon_n^\perp]^{2+d/\alpha} n$. Therefore,

$$1215 \quad |\mathcal{F}_n| \gtrsim [\varepsilon_n^\perp]^{2+d/\alpha} \cdot nh^{-d} \gtrsim [\varepsilon_n^\perp]^2 \cdot n M^{d/\alpha}. \quad (78)$$

1216 Because $\log(n/\varepsilon_n^\perp) \gtrsim \log n$ and $M > 0$ is a constant, we have 1217 that $\log |\mathcal{F}_n| \geq c \log n$ for all $n \geq N$, where $c > 0$ is a constant 1218 depending only on α, d and $N \in \mathbb{N}$ is a constant depending on M . 1219

1220 Let U be the uniform distribution on \mathcal{X} . Because $x \sim U$ 1221 and $f_x \equiv f_{x'}$ on $\mathcal{X} \setminus B_h^\infty(x)$, we have that

$$1222 \quad \text{KL}(P_{f_x} \| P_{f_{x'}}) = \frac{1}{2} \int_{\mathcal{X}} |f_x(z) - f_{x'}(z)|^2 dU(z) \quad (79)$$

$$1223 \quad \leq \frac{1}{2} \Pr_U [z \in B_h^\infty(x)] \cdot \|f_x - f_{x'}\|_\infty^2 \quad (80)$$

$$1224 \quad \leq \frac{1}{2} \lambda(B_h^\infty(x)) \cdot [\varepsilon_n^\perp]^2 \quad (81)$$

$$1225 \quad \lesssim h^d [\varepsilon_n^\perp]^2 \lesssim [\varepsilon_n^\perp]^{2+d/\alpha} / M^{d/\alpha}. \quad (82)$$

1226 By choosing M to be sufficiently large, the right-hand side of 1227 (77) can be lower bounded by an absolute constant. The 1228 theorem is then proved following the same argument as in the 1229 proof of Theorem 2.

1230 APPENDIX A 1231 SOME CONCENTRATION INEQUALITIES

1232 In this section, to ease readability of our paper, we provide 1233 some concentration inequalities and other standard results that 1234 we use extensively.

1235 **Lemma 9** ([56]). *Suppose X_1, \dots, X_n are i.i.d. random 1236 variables such that $a \leq X_i \leq b$ almost surely. Then for any 1237 $t > 0$,*

$$1238 \quad \Pr \left[\left| \frac{1}{n} \sum_{i=1}^n X_i - \mathbb{E}X \right| > t \right] \leq 2 \exp \left\{ - \frac{nt^2}{2(b-a)^2} \right\}.$$

1239 **Lemma 10** ([57]). *Suppose $x \sim \mathcal{N}_d(0, I_{d \times d})$ and let A be 1240 a $d \times d$ positive semi-definite matrix. Then for all $t > 0$, 1241*

$$1242 \quad \Pr \left[x^\top Ax > \text{tr}(A) + 2\sqrt{\text{tr}(A^2)t} + 2\|A\|_{\text{op}}t \right] \leq e^{-t}. \quad (83)$$

1244 **Lemma 11** ([58], simplified). *Suppose A_1, \dots, A_n are 1245 i.i.d. positive semidefinite random matrices of dimension d and 1246 $\|A_i\|_{\text{op}} \leq R$ almost surely. Then for any $t > 0$, 1247*

$$1248 \quad \Pr \left[\left\| \frac{1}{n} \sum_{i=1}^n A_i - \mathbb{E}A \right\|_{\text{op}} > t \right] \leq 2 \exp \left\{ - \frac{nt^2}{8R^2} \right\}.$$

1249 **Lemma 12** (Weyl's inequality). *Let A and $A + E$ 1250 be $d \times d$ matrices with $\sigma_1, \dots, \sigma_d$ and $\sigma'_1, \dots, \sigma'_d$ be 1251 their singular values, sorted in descending order. Then 1252 $\max_{1 \leq i \leq d} |\sigma_i - \sigma'_i| \leq \|E\|_{\text{op}}$. 1253*

1253 APPENDIX B
1254 ADDITIONAL PROOFS

1255 *Proof of Proposition 1.* Consider arbitrary $x^* \in \mathcal{X}$ such
1256 that $f(x^*) = \inf_{x \in \mathcal{X}} f(x)$. Then we have that $\mathcal{L}(\hat{x}_n; f) =$
1257 $f(\hat{x}_n) - f(x^*) \leq [f_n(\hat{x}_n) + \|f_n - f\|_\infty] - [f_n(x^*) - \|f_n -$
1258 $f\|_\infty] \leq 2\|f_n - f\|_\infty$, where the last inequality holds because
1259 $f_n(\hat{x}_n) \leq f_n(x^*)$ by optimality of \hat{x}_n . \square

1260 *Proof of Example 2.* Because $f_0 \in \Sigma_\kappa^2(M)$ is strongly convex,
1261 there exists $\sigma > 0$ such that $\nabla^2 f_0(x) \succeq \sigma I$ for all
1262 $x \in \mathcal{X}_{f_0, \kappa}$, where $\mathcal{X}_{f_0, \kappa} := L_{f_0}(\kappa)$ is the κ -level-set of f_0 .
1263 Let $x^* = \arg \min_{x \in \mathcal{X}} f_0(x)$, which is unique because f_0 is
1264 strongly convex. The smoothness and strong convexity of f_0
1265 implies that

$$1266 f_0^* + \frac{\sigma}{2} \|x - x^*\|_\infty^2 \leq f_0(x) \leq f_0^* + \frac{M}{2} \|x - x^*\|_\infty^2 \quad \forall x \in \mathcal{X}_{f_0, \kappa}. \quad (83)$$

1268 Subsequently, there exist constants $c_0, C_1, C_2 > 0$ depending
1269 only on σ, M, κ and d such that for all $\epsilon \in (0, c_0]$,

$$1270 B_{C_1 \sqrt{\epsilon}}^\infty(x^*; \mathcal{X}) \subseteq L_{f_0}(\epsilon) \subseteq B_{C_2 \sqrt{\epsilon}}^\infty(x^*; \mathcal{X}). \quad (84)$$

1271 The property $\mu_{f_0}(\epsilon) \lesssim \epsilon^\beta$ holds because $\mu(L_{f_0}(\epsilon)) \leq$
1272 $\mu(B_{C_1 \sqrt{\epsilon}}^\infty(x^*; \mathcal{X})) \lesssim \epsilon^{d/2}$. To prove (A2), note that
1273 $N(L_{f_0}(\epsilon), \delta) \leq N(B_{C_2 \sqrt{\epsilon}}^\infty(x^*; \mathcal{X}), \delta) \lesssim 1 + (\sqrt{\epsilon}/\delta)^d$. Because
1274 $\epsilon^{d/2} \lesssim \mu(L_{f_0}(\epsilon)) = \mu_{f_0}(\epsilon)$, we conclude that
1275 $N(L_{f_0}(\epsilon), \delta) \lesssim 1 + \delta^{-d} \mu_{f_0}(\epsilon)$ and (A2) is thus proved. \square

1276 *Proof of Proposition 4.* Consider $f_0 \equiv 0$ if $\beta = 0$ and $f_0(z) :=$
1277 $a_0 [z_1^p + \dots + z_d^p]$ for all $z = (z_1, \dots, z_d) \in [0, 1]^d$, where
1278 $a_0 > 0$ is a constant depending on α, M , and $p = d/\beta$ for
1279 $\beta \in (0, d/\alpha]$. The $\beta = 0$ case where $f_0 \equiv 0$ trivially holds.
1280 So we shall only consider the case of $\beta \in (0, d/\alpha]$.

1281 We first show $f_0 \in \Sigma_\kappa^\alpha(M)$ with $\kappa = \infty$, provided that a_0 is
1282 sufficiently small. For any $j \leq k = \lfloor \alpha \rfloor$ and $\alpha_1 + \dots + \alpha_d = j$,
1283 we have

$$1284 \frac{\partial^j}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} f_0(z) = \begin{cases} a_0 j! \cdot z_\ell^{p-j} & \text{if } \alpha_\ell = j, \ell \in [d]; \\ 0 & \text{otherwise.} \end{cases} \quad (85)$$

1286 Because $z_1, \dots, z_d \in [0, 1]$ and $p = d/\beta \geq \alpha \geq j$, it's clear
1287 that $0 \leq \frac{\partial^j}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} f_0(z) \leq a_0 j!$. In addition, for any
1288 $z, z' \in [0, 1]^d$ and $\alpha_\ell = k, \ell \in [d]$, we have

$$1289 \left| \frac{\partial^k}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} f_0(z) - \frac{\partial^k}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} f_0(z') \right| \leq a_0 k! \cdot |[z_\ell]^{p-k} - [z'_\ell]^{p-k}| \quad (86)$$

$$1291 \leq a_0 k! \cdot |z_\ell - z'_\ell|^{\min\{p-k, 1\}}, \quad (87)$$

1292 where the last inequality holds because x^t is $\min\{t, 1\}$ -Hölder
1293 continuous on $[0, 1]$ for $t \geq 0$. The $|z_\ell - z'_\ell|^{\min\{p-k, 1\}}$
1294 term can be further upper bounded by $\|z - z'\|_\infty^{\alpha-k}$, because
1295 $p = d/\beta \geq \alpha$. By selecting $a_0 > 0$ to be sufficiently small
1296 (depending on M) we have $f_0 \in \Sigma_\infty^\alpha(M)$.

1297 We next prove f_0 satisfies $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$ with parameter β
1298 depending on a_0 and p . For any $\epsilon > 0$, the level-set $L_{f_0}(\epsilon)$ can

1299 be expressed as $L_{f_0}(\epsilon) = \{z \in [0, 1]^d : z_1^p + \dots + z_d^p \leq \epsilon/a_0\}$.
1300 Subsequently,

$$1301 \left[0, \left(\frac{\epsilon}{a_0 d} \right)^{1/p} \right]^d \subseteq L_{f_0}(\epsilon) \subseteq \left[0, \left(\frac{\epsilon}{a_0} \right)^{1/p} \right]^d. \quad (88)$$

1302 Therefore,

$$1303 [\epsilon/(a_0 d)]^{dp} \leq \mu_{f_0}(\epsilon) \leq [\epsilon/a_0]^{dp}. \quad (89)$$

1304 Because a_0, d are constants and $dp = \beta$, we established
1305 $\mu_{f_0}(\epsilon) \asymp \epsilon^\beta$ for $\beta = dp$.

1306 Finally, note that for any $\epsilon > 0$, $L_{f_0}(\epsilon)$ is sandwiched
1307 between two cubics whose volumes only differ by a constant.
1308 This proves (A2) and (A2') on the covering and packing
1309 numbers of $L_{f_0}(\epsilon)$. \square

1310 *Proof of Proposition 5.* By the Chernoff bound and the
1311 union bound, with probability $1 - O(n^{-1})$ uniformly over all
1312 $x \in G_n$, there are $\Omega(\sqrt{n} \log^2 n)$ uniform samples in
1313 $B_{h_0}^\infty(x; \mathcal{X})$. Because $h_0 \leq \zeta$ for sufficiently large n_0 (ζ is
1314 defined in condition (A1)), by Lemma 1 it holds that

$$1315 |\check{f}_x(x') - f_x(x')| \lesssim h_0^\alpha + n_0^{-1/4} \lesssim n_0^{-\alpha/2d} + n_0^{-1/4}, \quad 1316 \forall x \in G_n, x' \in B_{h_0}^\infty(x; \mathcal{X}). \quad (90)$$

1317 Also, using the standard Gaussian concentration inequality,
1318 with probability $1 - O(n^{-1})$ we have

$$1319 \inf_{x' \in B_{h_0}^\infty(x; \mathcal{X})} f(x) - O(n_0^{-1/4}) \leq \bar{f}(x) \leq \sup_{x' \in B_{h_0}^\infty(x; \mathcal{X})} f(x) + O(n_0^{-1/4}) \quad \forall x \in G_n. \quad (91)$$

1320 Let x^* be the minimizer of f on \mathcal{X} and $x \in G_n$ such
1321 that $\|x - x^*\|_\infty \leq h_0$. By (90), we have with probability
1322 $1 - O(n^{-1})$ that $\inf_{x' \in B_{h_0}^\infty(x; \mathcal{X})} \check{f}_x(x') \leq f^* + O(n_0^{-\alpha/2d} +$
1323 $n_0^{-1/4}) \leq f^* + 1/2 \log n$, where $f^* = f(x^*)$. Now consider
1324 arbitrary $z \in G_n$ such that $B_{h_0}^\infty(z; \mathcal{X}) \cap L_f(\kappa/2) = \emptyset$,
1325 meaning that for all $z' \in \mathcal{X}$, $\|z' - z\|_\infty \leq h_0$, $f(z') >$
1326 $\kappa/2$. By (90), $\bar{f}(z) \geq \kappa/2 - O(n_0^{-1/4}) \geq \kappa/2 - 1/2 \log n$.
1327 Hence when n_0 is sufficiently large, $z \notin S'_0$, which is to be
1328 demonstrated. \square

1329 *Proof of Proposition 7.* The upper bound part of (53) trivially
1330 holds because the absolute values of every element in
1331 $\psi_{0,1}(z) \psi_{0,1}(z)^\top$ for $z \in \mathcal{X} = [0, 1]^d$ is upper bounded by
1332 $O(1)$. To prove the lower bound part, we only need to show
1333 $\int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z)$ is invertible. Assume the contrary.
1334 Then there exists $v \in \mathbb{R}^D \setminus \{0\}$ such that

$$1335 v^\top \left[\int_{\mathcal{X}} \psi_{0,1}(z) \psi_{0,1}(z)^\top dU(z) \right] v = \int_{\mathcal{X}} |\psi_{0,1}(z)^\top v|^2 dU(z) = 0. \quad (92)$$

1336 Therefore, $\langle \psi_{0,1}(z), v \rangle = 0$ almost everywhere on $z \in$
1337 $[0, 1]^d$. Because $h > 0$, by re-scaling with constants this

implies the existence of non-zero coefficient vector ξ such that

$$P(z_1, \dots, z_m) := \sum_{\alpha_1+\dots+\alpha_m \leq k} \xi_{\alpha_1, \dots, \alpha_m} z_1^{\alpha_1} \dots z_m^{\alpha_m} = 0$$

almost everywhere on $z \in [0, 1]^d$.

We next use induction to show that, for any degree- k polynomial P of s variables z_1, \dots, z_s that has at least one non-zero coefficient, the set $\{z_1, \dots, z_s \in [0, 1]^d : P(z_1, \dots, z_s) = 0\}$ must have zero measure. This would then result in the desired contradiction. For the base case of $s = 1$, the fundamental theorem of algebra asserts that $P(z_1) = 0$ can have at most k roots, which is a finite set and of measure 0.

We next consider the case where $P(z_1, \dots, z_s)$ takes on s variables. Re-organizing the terms we have

$$P(z_1, \dots, z_s) \equiv P_0(z_1, \dots, z_{s-1}) + z_s P_1(z_1, \dots, z_{s-1}) + \dots + z_s^k P_k(z_1, \dots, z_{s-1}), \quad (93)$$

where P_1, \dots, P_k are degree- k polynomials of z_1, \dots, z_{s-1} . Because P has a non-zero coefficient, at least one P_j must also have a non-zero coefficient. By the inductive hypothesis, the set $\{z_1, \dots, z_{s-1} : P_j(z_1, \dots, z_{s-1}) = 0\}$ has measure 0. On the other hand, if $P_j(z_1, \dots, z_{s-1}) \neq 0$, then invoking the fundamental theorem of algebra again on z_s we know that there are finitely many z_s such that $P(z_1, \dots, z_s) = 0$. Therefore, $\{z_1, \dots, z_s : P(z_1, \dots, z_s) = 0\}$ must also have measure zero. \square

REFERENCES

- [1] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*, vol. 1. Cambridge, MA, USA: MIT Press, 2006.
- [2] B. Reeja-Jayan, K. L. Harrison, K. Yang, C.-L. Wang, A. E. Yilmaz, and A. Manthiram, "Microwave-assisted low-temperature growth of thin films in solution," *Sci. Rep.*, vol. 2, Dec. 2012, Art. no. 1003.
- [3] N. Nakamura, J. Seepaul, J. B. Kadane, and B. Reeja-Jayan, "Design for low-temperature microwave-assisted crystallization of ceramic thin films," *Appl. Stochastic Models Bus. Ind.*, vol. 33, no. 3, pp. 314–321, 2017.
- [4] A. B. Tsybakov, *Introduction to Nonparametric Estimation* (Springer Series in Statistics). New York, NY, USA: Springer, 2009.
- [5] J. Fan and I. Gijbels, *Local Polynomial Modelling and its Applications*. Boca Raton, FL, USA: CRC Press, 1996.
- [6] A. D. Bull, "Convergence rates of efficient global optimization algorithms," *J. Mach. Learn. Res.*, vol. 12, pp. 2879–2904, Oct. 2011.
- [7] J. Scarlett, I. Bogunovic, and V. Cevher, "Lower bounds on regret for noisy Gaussian process bandit optimization," in *Proc. Annu. Conf. Learn. Theory (COLT)*, 2017, pp. 1723–1742.
- [8] E. Hazan, A. Klivans, and Y. Yuan, "Hyperparameter optimization: A spectral approach," 2017, *arXiv:1706.00764*. [Online]. Available: <https://arxiv.org/abs/1706.00764#>
- [9] A. S. Nemirovski and D. B. Yudin, *Problem Complexity and Method Efficiency in Optimization*. Hoboken, NJ, USA: Wiley, 1983.
- [10] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, "Online convex optimization in the bandit setting: Gradient descent without a gradient," in *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, 2005, pp. 385–394.
- [11] A. Agarwal, O. Dekel, and L. Xiao, "Optimal algorithms for online convex optimization with multi-point bandit feedback," in *Proc. Annu. Conf. Learn. Theory (COLT)*, 2010, pp. 28–40.
- [12] K. G. Jamieson, R. Nowak, and B. Recht, "Query complexity of derivative-free optimization," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2012, pp. 2672–2680.
- [13] A. Agarwal, D. P. Foster, D. Hsu, S. M. Kakade, and A. Rakhlin, "Stochastic convex optimization with bandit feedback," *SIAM J. Optim.*, vol. 23, no. 1, pp. 213–240, 2013.
- [14] S. Bubeck, Y. T. Lee, and R. Eldan, "Kernel-based methods for bandit convex optimization," in *Proc. 49th Annu. ACM SIGACT Symp. Theory Comput. (STOC)*, 2017, pp. 72–85.
- [15] A. H. G. R. Kan and G. T. Timmer, "Stochastic global optimization methods part I: Clustering methods," *Math. Program.*, vol. 39, no. 1, pp. 27–56, 1987.
- [16] A. H. G. R. Kan and G. T. Timmer, "Stochastic global optimization methods part II: Multi level methods," *Math. Program.*, vol. 39, no. 1, pp. 57–78, 1987.
- [17] S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári, "X-armed bandits," *J. Mach. Learn. Res.*, vol. 12, pp. 1655–1695, May 2011.
- [18] C. Malherbe, E. Contal, and N. Vayatis, "A ranking approach to global optimization," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2016.
- [19] C. Malherbe and N. Vayatis, "Global optimization of Lipschitz functions," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2017.
- [20] R. D. Kleinberg, "Nearly tight bounds for the continuum-armed bandit problem," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2005, pp. 697–704.
- [21] S. Minsker, "Estimation of extreme values and associated level sets of a regression function via selective sampling," in *Proc. Conf. Learn. Theory (COLT)*, 2013, pp. 105–121.
- [22] J.-B. Grill, M. Valko, and R. Munos, "Black-box optimization of noisy functions with unknown smoothness," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2015, pp. 667–675.
- [23] S. Minsker, "Non-asymptotic bounds for prediction problems and density estimation," Ph.D. dissertation, Georgia Inst. Technol., Atlanta, Georgia, 2012.
- [24] J. Kiefer and J. Wolfowitz, "Stochastic estimation of the maximum of a regression function," *Ann. Math. Stat.*, vol. 23, no. 3, pp. 462–466, 1952.
- [25] E. Purzen, "On estimation of a probability density and mode," *Ann. Math. Statist.*, vol. 39, no. 3, pp. 1065–1076, 1962.
- [26] H. Chen, "Lower rate of convergence for locating a maximum of a function," *Ann. Statist.*, vol. 16, no. 3, pp. 1330–1334, 1988.
- [27] Z. B. Zabinsky and R. L. Smith, "Pure adaptive search in global optimization," *Math. Program.*, vol. 53, no. 1, pp. 323–338, 1992.
- [28] M.-F. Balcan, A. Beygelzimer, and J. Langford, "Agnostic active learning," *J. Comput. Syst. Sci.*, vol. 75, no. 1, pp. 78–89, 2009.
- [29] S. Dasgupta, D. J. Hsu, and C. Monteleoni, "A general agnostic active learning algorithm," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2008, pp. 353–360.
- [30] S. Hanneke, "A bound on the label complexity of agnostic active learning," in *Proc. 24th Int. Conf. Mach. Learn. (ICML)*, 2007, pp. 353–360.
- [31] E. Even-Dar, S. Mannor, and Y. Mansour, "Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems," *J. Mach. Learn. Res.*, vol. 7, pp. 1079–1105, Jun. 2006.
- [32] W. Polonik, "Measuring mass concentrations and estimating density contour clusters—an excess mass approach," *Ann. Statist.*, vol. 23, no. 3, pp. 855–881, 1995.
- [33] P. Rigollet and R. Vert, "Optimal rates for plug-in estimators of density level sets," *Bernoulli*, vol. 15, no. 4, pp. 1154–1178, 2009.
- [34] A. Singh, C. Scott, and R. Nowak, "Adaptive Hausdorff estimation of density level sets," *Ann. Statist.*, vol. 37, no. 5B, pp. 2760–2782, 2009.
- [35] K. Chaudhuri, S. Dasgupta, S. Kpotufe, and U. V. Luxburg, "Consistent procedures for cluster tree estimation and pruning," *IEEE Trans. Inf. Theory*, vol. 60, no. 12, pp. 7900–7912, Dec. 2014.
- [36] S. Balakrishnan, S. Narayanan, A. Rinaldo, A. Singh, and L. Wasserman, "Cluster trees on manifolds," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2013, pp. 2679–2687.
- [37] Y. Nesterov and B. T. Polyak, "Cubic regularization of Newton method and its global performance," *Math. Program.*, vol. 108, no. 1, pp. 177–205, 2006.
- [38] E. Hazan, K. Levy, and S. Shalev-Shwartz, "Beyond convexity: Stochastic quasi-convex optimization," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2015, pp. 1594–1602.
- [39] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points—Online stochastic gradient for tensor decomposition," in *Proc. Annu. Conf. Learn. Theory (COLT)*, 2015, pp. 797–842.
- [40] N. Agarwal, Z. Allen-Zhu, B. Bullins, E. Hazan, and T. Ma, "Finding approximate local minima faster than gradient descent," in *Proc. 49th Annu. ACM SIGACT Symp. Theory Comput. (STOC)*, 2017, pp. 1195–1199.
- [41] Y. Carmon, O. Hinder, J. C. Duchi, and A. Sidford, "'Convex until proven guilty': Dimension-free acceleration of gradient descent on non-convex functions," 2017, *arXiv:1705.02766*. [Online]. Available: <https://arxiv.org/abs/1705.02766>

AQ:3

1477 [42] Y. Zhang, P. Liang, and M. Charikar, "A hitting time analysis of
1478 stochastic gradient Langevin dynamics," in *Proc. Annu. Conf. Learn.*
1479 *Theory (COLT)*, 2017, pp. 1–43.

1480 [43] Y. Zhu, S. Chatterjee, J. Duchi, and J. Lafferty, "Local minimax
1481 complexity of stochastic convex optimization," in *Proc. NIPS*, 2016,
1482 pp. 3431–3439.

1483 [44] J. Duchi and F. Ruan, "Asymptotic optimality in stochastic opti-
1484 mization," 2016, *arXiv:1612.05612*. [Online]. Available: <https://arxiv.org/abs/1612.05612>

1485 [45] A. Locatelli and A. Carpentier, "Adaptivity to smoothness in X-armed
1486 bandits," in *Proc. Conf. Learn. Theory (COLT)*, 2018, pp. 1463–1492.

1487 [46] A. W. van der Vaart, *Asymptotic Statistics*, vol. 3. Cambridge, U.K.:
1488 Cambridge Univ. Press, 1998.

1489 [47] C. Jin, L. T. Liu, R. Ge, and M. I. Jordan, "On the local minima of
1490 the empirical risk," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*,
1491 2018, pp. 1–10.

1492 [48] T. T. Cai and M. G. Low, "An adaptation theory for nonparametric
1493 confidence intervals," *Ann. Statist.*, vol. 32, no. 5, pp. 1805–1840, 2004.

1494 [49] A. P. Korostelev and A. B. Tsybakov, *Minimax Theory of Image
1495 Reconstruction*, vol. 82. Springer, 2012.

1496 [50] R. M. Castro and R. D. Nowak, "Minimax bounds for active learning,"
1497 *IEEE Trans. Inf. Theory*, vol. 54, no. 5, pp. 2339–2353, May 2008.

1498 [51] S. Bubeck, R. Munos, and G. Stoltz, "Pure exploration in multi-armed
1499 bandits problems," in *Proc. Int. Conf. Algorithmic Learn. Theory (ALT)*,
1500 2009, pp. 23–37.

1501 [52] O. V. Lepski, E. Mammen, and V. G. Spokoiny, "Optimal spatial
1502 adaptation to inhomogeneous smoothness: An approach based on kernel
1503 estimates with variable bandwidth selectors," *Ann. Statist.*, vol. 25, no. 3,
1504 pp. 929–947, 1997.

1505 [53] W. K. Newey, "Convergence rates and asymptotic normality for series
1506 estimators," *J. Econometrics*, vol. 79, no. 1, pp. 147–168, 1997.

1507 [54] D. S. Ebert, *Texturing & Modeling: A Procedural Approach*. San Mateo,
1508 CA, USA: Morgan Kaufmann, 2003.

1509 [55] R. M. Castro, "Adaptive sensing performance lower bounds for sparse
1510 signal detection and support estimation," *Bernoulli*, vol. 20, no. 4,
1511 pp. 2217–2246, 2014.

1512 [56] W. Hoeffding, "Probability inequalities for sums of bounded random
1513 variables," *J. Amer. Stat. Assoc.*, vol. 58, no. 301, pp. 13–30, 1963.

1514 [57] D. Hsu, S. M. Kakade, and T. Zhang, "A tail inequality for quadratic
1515 forms of subgaussian random vectors," *Electron. Commun. Probab.*,
1516 vol. 17, no. 52, pp. 1–6, 2012.

1517 [58] J. A. Tropp, "An introduction to matrix concentration inequalities,"
1518 *Found. Trends Mach. Learn.*, vol. 8, nos. 1–2, pp. 1–230, 2015.

Yining Wang received the B.Eng. degree in computer science and technology 1520
in 2014 from Tsinghua University, Beijing China, the M.S. degree in machine 1521
learning in 2017 from Carnegie Mellon University, Pittsburgh, PA, USA. 1522
He is currently a Ph.D. student in machine learning in the machine learning 1523
department at Carnegie Mellon University, Pittsburgh, PA, USA. His research 1524
interests are primarily in statistical machine learning, with emphasis on 1525
interactive methods, active learning, adaptive sampling. 1526

Sivaraman Balakrishnan is an Assistant Professor in the Department of 1527
Statistics and Data Science at Carnegie Mellon University. Prior to this he 1528
received his Ph.D. from the School of Computer Science at Carnegie Mellon 1529
University and was a postdoctoral researcher in the Department of Statistics at 1530
UC Berkeley. His Ph.D. work was supported by several fellowships including 1531
the Richard King Mellon Fellowship and a grant from the Gates Foundation. 1532
He is broadly interested in problems that lie at the interface between computer 1533
science and statistics. Some particular areas that have provided motivation 1534
for his past and current research include the applications of statistical 1535
methods in ranking problems, computational biology, clustering, topological 1536
data analysis, nonparametric statistics, robust statistics and non-convex 1537
optimization. 1538

Aarti Singh received the B.E. degree in electronics and communication 1539
engineering from the University of Delhi, New Delhi, India, in 2001, and 1540
the M.S. and Ph.D. degrees in electrical and computer engineering from the 1541
University of Wisconsin–Madison, Madison, WI, USA, in 2003 and 2008, 1542
respectively. She was a Postdoctoral Research Associate at the Program in 1543
Applied and Computational Mathematics, Princeton University, from 2008 to 1544
2009, before joining the School of Computer Science, Carnegie Mellon 1545
University, Pittsburgh, PA, USA, where she has been an Associate Professor 1546
since 2009. Her research interests include the intersection of machine learning, 1547
statistics and signal processing, and focus on designing statistically and 1548
computationally efficient algorithms that can leverage inherent structure of the 1549
data in the form of clusters, graphs, subspaces, and manifold using direct, 1550
compressive, and active queries. Her work is recognized by the NSF Career Award, 1551
the United States Air Force Young Investigator Award, A. Nico Habermann 1552
Faculty Chair Award, Harold A. Peterson Best Dissertation Award, and a best 1553
student paper award at Allerton. 1554