IMPROVED REGRET BOUNDS FOR GAUSSIAN PROCESS BANDITS UNDER HEAVY-TAILED NOISE

Anonymous authors
Paper under double-blind review

ABSTRACT

We study Gaussian process bandit optimization of an unknown objective function in a reproducing kernel Hilbert space (RKHS). We assume that the observation noise has a finite $(1+\epsilon)$-moment for some $\epsilon \in (0, 1]$. Unlike the sub-Gaussian noise assumption, this moment condition allows heavy-tailed noise and does not require finite variance when $\epsilon < 1$. For a horizon T and maximum information gain γ_T, existing regret bounds in this setting scale as either $\widetilde{\mathcal{O}}(T^{\frac{2+\epsilon}{2(1+\epsilon)}}\gamma_T)$ or $\widetilde{\mathcal{O}}(T^{\frac{1}{1+\epsilon}}\gamma_T)$, both linear in γ_T. To reduce this dependence on γ_T, we propose the heavy-tailed chaining phased elimination (HT-ChainPE) algorithm, which combines regularized D-optimal design with a robust chained estimator. HT-ChainPE achieves cumulative regret of order $\widetilde{\mathcal{O}}(T^{\frac{1}{1+\epsilon}}(1 \vee \gamma_T)^{\frac{1+3\epsilon}{2(1+\epsilon)}})$. For $\epsilon < 1$, this bound retains the $T^{1/(1+\epsilon)}$ dependence achieved in prior work while reducing the exponent of γ_T below one. We also establish a worst-case cumulative regret lower bound on $[0, 1]^d$ of $\Omega(T^{\frac{1}{1+\epsilon}}L_\delta^{\frac{\epsilon}{1+\epsilon}}(\frac{\log(T/L_\delta)}{\log\log(T/L_\delta)})^{\frac{d\epsilon}{1+\epsilon}})$ for the squared-exponential (SE) kernel, where $L_\delta = (1-2\delta)\log(\frac{1-\delta}{\delta})$ and $\delta \in (0, 1/3)$. For the Matérn kernel with smoothness ν on $[0, 1]^d$, we obtain a lower bound of $\Omega(T^{\frac{\nu+d\epsilon}{\nu(1+\epsilon)+d\epsilon}}L_\delta^{\frac{\nu\epsilon}{\nu(1+\epsilon)+d\epsilon}})$. Finally, experiments on synthetic and real-world benchmarks show that HT-ChainPE performs competitively against existing methods.

1 INTRODUCTION

Bayesian optimization (BO) seeks a maximizer of an unknown function whose evaluations are expensive and may be noisy (Garnett, 2023). It is widely used in applications such as hyperparameter tuning (Snoek et al., 2012), robotic gait optimization (Calandra et al., 2016), and chemical reaction optimization (Shields et al., 2021). BO commonly uses a Gaussian process (GP) surrogate to predict function values and quantify uncertainty, with a kernel that allows observations at one point to inform predictions at others. We study this problem in the GP bandit framework, where a learner sequentially selects query points and observes noisy function values, balancing exploration and exploitation Srinivas et al. (2010).

Most theoretical guarantees for GP bandits assume sub-Gaussian observation noise. Real-world observations, however, often exhibit heavy-tailed behavior, as seen in financial returns Cont (2001) and network traffic Crovella & Bestavros (1997). Intuitively, most observations may remain near their typical values, while occasional observations deviate substantially, such as a sudden price change or an unusually long network delay. These extreme deviations occur more frequently than sub-Gaussian models predict, so even a few observations can distort estimates and invalidate confidence bounds. One way to model this behavior is to assume that the observation noise has a finite $(1+\epsilon)$-moment for some $\epsilon \in (0, 1]$ Bubeck et al. (2013). This condition allows infinite variance when $\epsilon < 1$.

Several bandit algorithms have been developed for this noise model. Building on robust estimation for heavy-tailed linear bandits (Medina & Yang, 2016; Shao et al., 2018), Chowdhury & Gopalan (2019) introduced TGP-UCB and ATA-GP-UCB algorithms for GP bandits under heavy-tailed noise. These algorithms use truncation to limit the influence of extreme observations, with thresholds that control the resulting bias. ATA-GP-UCB achieves cumulative regret of order $\widetilde{\mathcal{O}}(\gamma_T T^{\frac{1}{1+\epsilon}})$,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Table 1: Comparison of cumulative-regret bounds under Assumption 1

Method	General	Squared-exponential kernel	Matérn-ν kernel
Upper bounds			
TGP-UCB (2019)	$\widetilde{\mathcal{O}}\Big(T^{\frac{2+\epsilon}{2(1+\epsilon)}}\gamma_T\Big)$	$\widetilde{\mathcal{O}}\Big(T^{\frac{2+\epsilon}{2(1+\epsilon)}}(\log T)^{d+1}\Big)$	$\widetilde{\mathcal{O}}\Big(T^{\frac{2+\epsilon}{2(1+\epsilon)}+\frac{d}{2\nu+d}}(\log T)^{\frac{2\nu}{2\nu+d}}\Big)$
ATA-GP-UCB (2019)	$\widetilde{\mathcal{O}}\Big(T^{\frac{1}{1+\epsilon}}\gamma_T\Big)$	$\widetilde{\mathcal{O}}\Big(T^{\frac{1}{1+\epsilon}}(\log T)^{d+1}\Big)$	$\widetilde{\mathcal{O}}\Big(T^{\frac{1}{1+\epsilon}+\frac{d}{2\nu+d}}(\log T)^{\frac{2\nu}{2\nu+d}}\Big)$
CA-TGP-UCB (2024)			
MoM-GP-UCB (2024)			
HT-ChainPE (Ours)	$\widetilde{\mathcal{O}}\Big(T^{\frac{1}{1+\epsilon}}(1\vee\gamma_T)^{\frac{1+3\epsilon}{2(1+\epsilon)}}\Big)$	$\widetilde{\mathcal{O}}\Big(T^{\frac{1}{1+\epsilon}}(\log T)^{\frac{(d+1)(1+3\epsilon)}{2(1+\epsilon)}}\Big)$	$\widetilde{\mathcal{O}}\Big(T^{\frac{1}{1+\epsilon}+\frac{d(1+3\epsilon)}{2(1+\epsilon)(2\nu+d)}}(\log T)^{\frac{\nu(1+3\epsilon)}{(1+\epsilon)(2\nu+d)}}\Big)$
Lower bounds			
Chowdhury & Gopalan (2019)	–	$\Omega\Big(T^{\frac{1}{1+\epsilon}}(\log T)^{\frac{d\epsilon}{2(1+\epsilon)}}\Big)$	$\Omega\Big(T^{\frac{\nu+d\epsilon}{\nu(1+\epsilon)+d\epsilon}}\Big)$
Ours	–	$\Omega\Big(T^{\frac{1}{1+\epsilon}}L_\delta^{\frac{\epsilon}{1+\epsilon}}\Big(\frac{\log(T/L_\delta)}{\log\log(T/L_\delta)}\Big)^{\frac{d\epsilon}{1+\epsilon}}\Big)$	$\Omega\Big(T^{\frac{\nu+d\epsilon}{\nu(1+\epsilon)+d\epsilon}}L_\delta^{\frac{\nu\epsilon}{\nu(1+\epsilon)+d\epsilon}}\Big)$

Here, we denote $L_\delta := (1-2\delta)\log(\frac{1-\delta}{\delta})$ with $\delta \in (0, \frac{1}{3})$. Factors depending only on B and the noise scale are suppressed in the lower bounds, while the δ-dependence of our bound is retained. Chowdhury & Gopalan (2019) establish lower bounds on expected regret, whereas ours establish lower bounds on high-probability regret guarantees. Table 5 gives a detailed comparison of the lower bounds. For the upper bounds, the SE column uses $\gamma_T = \mathcal{O}((\log T)^{d+1})$ Srinivas et al. (2010), and the Matérn column uses $\gamma_T = \mathcal{O}(T^{\frac{d}{2\nu+d}}(\log T)^{\frac{2\nu}{2\nu+d}})$ Vakili et al. (2021b). Since $\frac{1+3\epsilon}{2(1+\epsilon)} < 1$ for $\epsilon < 1$, HT-ChainPE improves the dependence on γ_T.

suppressing moment-dependent factors, where γ_T is the maximum information gain and measures kernel complexity (Srinivas et al., 2010). Park & Lee (2024) subsequently introduced CA-TGP-UCB and MoM-GP-UCB, which reduce computational complexity while preserving the regret order of ATA-GP-UCB. Despite the reduction in computational complexity, their regret bounds still scale linearly with γ_T.

Notably, recent progress in heavy-tailed linear bandits suggests that tighter complexity dependence may be possible. Tajdini et al. (2025) improved the heavy-tailed linear-bandit regret rate of Shao et al. (2018) from $\widetilde{\mathcal{O}}(dT^{\frac{1}{1+\epsilon}})$ to $\widetilde{\mathcal{O}}(d^{\frac{1+3\epsilon}{2(1+\epsilon)}}T^{\frac{1}{1+\epsilon}})$. The improved dimension exponent is strictly below one when $\epsilon < 1$, under which the noise may have infinite variance. In GP bandits, maximum information gain γ_T plays a related role by quantifying the complexity of learning through the kernel. These results motivate the following question:

Can GP bandits under heavy-tailed noise achieve a regret bound of order $\widetilde{\mathcal{O}}\big(\gamma_T^{\frac{1+3\epsilon}{2(1+\epsilon)}}T^{\frac{1}{1+\epsilon}}\big)$*?*

We answer this question with heavy-tailed chaining phased elimination (HT-ChainPE). Our main contributions are as follows.

- HT-ChainPE (Algorithm 1) achieves a high-probability cumulative-regret bound of $\widetilde{\mathcal{O}}\big(v^{\frac{1}{1+\epsilon}}(1\vee\gamma_T)^{\frac{1+3\epsilon}{2(1+\epsilon)}}T^{\frac{1}{1+\epsilon}}\big)$. To the best of our knowledge, this is the first general-kernel regret bound with sublinear dependence on γ_T for GP bandits under heavy-tailed noise. This improvement holds when $\epsilon < 1$, which allows infinite variance. Table 1 and Section 5 give detailed comparisons of regret bounds. Moreover, Table 2 and Appendix C compare its computational costs with those of existing methods.
- We establish high-probability cumulative-regret lower bounds for the SE and Matérn kernels. Compared with the expected-regret lower bound of Chowdhury & Gopalan (2019), our SE bound has a larger logarithmic exponent, up to a $\log\log T$ factor. Our Matérn bound has the same horizon exponent and makes the dependence on the failure probability explicit. Further discussion and a detailed comparison of the lower bounds are provided in Section 5.2 and Table 5, respectively.

- We evaluate HT-ChainPE on synthetic benchmarks and real-world stock-market and light-sensor datasets, reporting favorable empirical regret and competitive wall-clock time relative to existing methods. Detailed results are presented in Appendix E.

2 Related Work

Srinivas et al. (2010) initiated a line of research on Gaussian process (GP) bandits by introducing GP-UCB. Subsequent work refined regret bounds for GP and kernel bandits (Chowdhury & Gopalan, 2017; Janz et al., 2020; Vakili et al., 2021b; Salgia et al., 2021; 2024; Iwazaki & Takeno, 2025) and studied simple regret, contextual bandits, and batched feedback (Vakili et al., 2021a; Valko et al., 2013; Zenati et al., 2022; Li & Scarlett, 2022). These studies primarily assume sub-Gaussian noise or consider specialized noiseless and variance-dependent settings; their guarantees do not directly extend to observations with only a finite $(1+\epsilon)$-moment.

Heavy-tailed bandits were first studied with finite action sets (Bubeck et al., 2013; Agrawal et al., 2021; Lee et al., 2020; Lee & Lim, 2022) and later extended to linear bandits (Medina & Yang, 2016; Shao et al., 2018; Xue et al., 2020; Kang & Kim, 2023; Wang et al., 2025; Tajdini et al., 2025). Heavy-tailed rewards have also been studied in Lipschitz bandits (Lu et al., 2019) and graphical bandits (Gou et al., 2023), and heavy-tailed losses in adversarial bandits (Cheng et al., 2024). Chowdhury & Gopalan (2019) extended GP bandits to heavy-tailed noise using truncation-based estimators, and Park & Lee (2024) subsequently achieved the same regret order with lower computational complexity. For $\epsilon < 1$, we improve the dependence on γ_T in the regret bounds of these prior works, reducing its exponent from 1 to $\frac{1+3\epsilon}{2(1+\epsilon)}$.

Chowdhury & Gopalan (2019) established expected-regret lower bounds for heavy-tailed GP bandits with SE and Matérn kernels. Our high-probability lower bound for the SE kernel has a larger logarithmic exponent, up to a $\log\log T$ factor. For the Matérn kernel, our high-probability lower bound matches their exponent of T and makes the dependence on the confidence level explicit.

3 Problem Formulation

We study sequential optimization of an unknown reward function $f : \mathcal{X} \to \mathbb{R}$ over a compact input domain $\mathcal{X} \subseteq \mathbb{R}^d$, where $d \in \mathbb{N}$. The learner interacts with the environment for T rounds. Let $\mathcal{F}_t := \sigma(x_1, y_1, \ldots, x_t, y_t)$ denote the history after round t, with $\mathcal{F}_0$ the trivial sigma-field. At round t, the learner queries $x_t \in \mathcal{X}$ based on $\mathcal{F}_{t-1}$ and observes $y_t = f(x_t) + \eta_t$, where η_t denotes the observation noise. The kernel regularity conditions below ensure that a maximizer of f exists. Fix $x^\star \in \arg\max_{x\in\mathcal{X}} f(x)$. We evaluate a policy through its cumulative pseudo-regret $R_T := \sum_{t=1}^{T} \big(f(x^\star) - f(x_t)\big)$.

Kernel model and complexity. Let k be a positive semidefinite kernel on $\mathcal{X}$, let $\mathcal{H}_k$ be its reproducing kernel Hilbert space (RKHS), and define the canonical feature map $\phi(x) := k(x, \cdot)$. For $B > 0$, write $\mathcal{H}_k(\mathcal{X}, B) := \{g \in \mathcal{H}_k : \|g\|_{\mathcal{H}_k} \le B\}$. When $\mathcal{X}$ is infinite, we consider kernels k that are continuous on $\mathcal{X} \times \mathcal{X}$.

Assumption 1 (RKHS regularity Srinivas et al. (2010))**.** *The function f belongs to $\mathcal{H}_k(\mathcal{X}, B)$, and $k(x, x) \le 1$ for every $x \in \mathcal{X}$.*

By the reproducing property, Assumption 1 implies $|f(x)| \le B$ for every $x \in \mathcal{X}$. When k is continuous, the feature map ϕ and every function in $\mathcal{H}_k$ are continuous. Compactness of $\mathcal{X}$ therefore guarantees the existence of a maximizer of f. Our kernel-specific results use the squared-exponential (SE) kernel with length scale $\ell_{\mathrm{SE}} > 0$ and the Matérn kernel with smoothness $\nu > 0$ and inverse length scale $\lambda_0 > 0$. Let K_ν denote the modified Bessel function of the second kind. These normalized kernels take the form (Williams & Rasmussen, 2006):

$$k_{\mathrm{SE}}(x, x') := \exp\Big(-\frac{\|x - x'\|_2^2}{2\ell_{\mathrm{SE}}^2}\Big), \quad k_\nu(x, x') := \frac{2^{1-\nu}}{\Gamma(\nu)}\Big(2\pi\lambda_0\|x - x'\|_2\Big)^\nu K_\nu\Big(2\pi\lambda_0\|x - x'\|_2\Big). \tag{1}$$

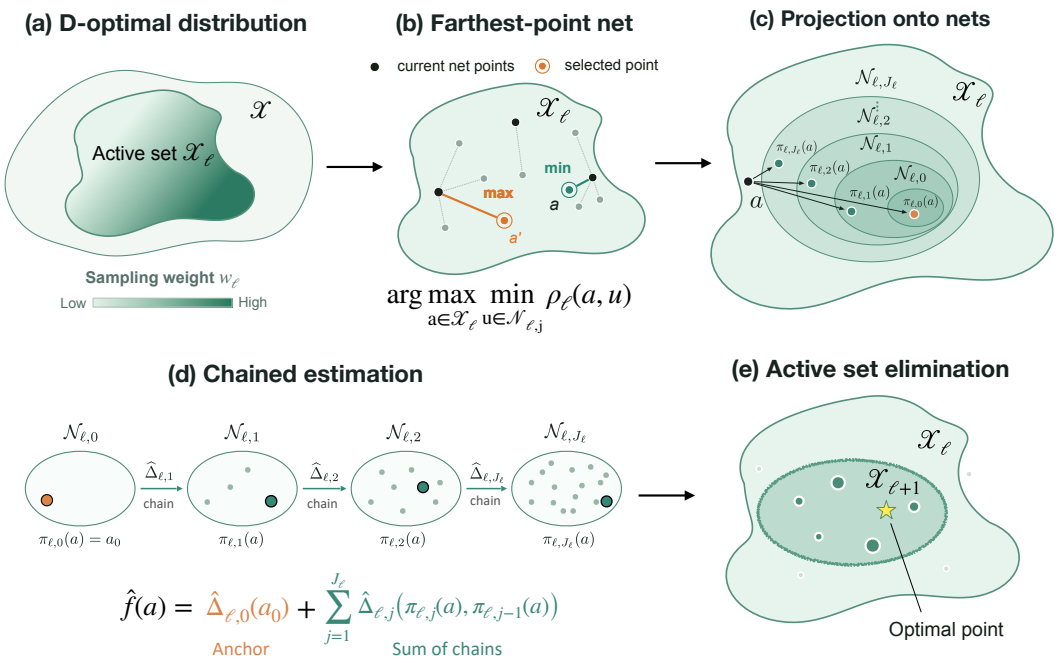


Figure 1: Overview of the HT-ChainPE algorithm (Algorithm 1).

Definition 2 (Maximum information gain (Srinivas et al., 2010)). *For each $(x_1, \dots, x_T) \in \mathcal{X}^T$, let $K_T := [k(x_i, x_j)]_{i,j=1}^T$. Let I denote the identity operator on the space under consideration. The maximum information gain of k after T observations is $\gamma_T := \sup_{(x_1,\dots,x_T)\in\mathcal{X}^T} \frac{1}{2} \log\det(I + K_T)$.*

The log-determinant aggregates information across the eigen-directions of K_T, while the supremum ranges over the possible sequences of T query points. Appendix B recalls the standard GP surrogate and its relation to information gain. We now formalize the heavy-tailed noise assumption on the observations.

Assumption 3 (Heavy-tailed noise (Bubeck et al., 2013)). *Let $\epsilon \in (0, 1]$. The variables $(\eta_t)_{t\geq 1}$ are i.i.d., and η_t is independent of $\sigma(\mathcal{F}_{t-1}, x_t)$ for every t. Moreover, $\mathbb{E}[\eta_t \mid \mathcal{F}_{t-1}, x_t] = 0$ and, for some $v > 0$, $\mathbb{E}[|\eta_t|^{1+\epsilon} \mid \mathcal{F}_{t-1}, x_t] \leq v$ almost surely.*

When $\epsilon < 1$, Assumption 3 admits infinite noise variance. Since $|a + b|^{1+\epsilon} \leq 2^\epsilon(|a|^{1+\epsilon} + |b|^{1+\epsilon})$, we have $\mathbb{E}[|y_t|^{1+\epsilon} \mid \mathcal{F}_{t-1}, x_t] \leq 2^\epsilon(B^{1+\epsilon} + v) =: \bar{v}$ almost surely. Throughout, $C_\epsilon \geq 1$ denotes a constant depending only on ϵ that may be enlarged when necessary.

4 Algorithm

The heavy-tailed chaining phased elimination (HT-ChainPE) algorithm (Algorithm 1) eliminates candidate maximizers in phases, beginning with $\mathcal{X}_1 = \mathcal{X}$. In phase ℓ, $\mathcal{X}_\ell$ contains the remaining candidates, and the goal is to estimate f uniformly over $\mathcal{X}_\ell$ with sufficiently small error. Figure 1 illustrates the steps of one phase. The algorithm first selects a sampling distribution w_ℓ to control uncertainty across $\mathcal{X}_\ell$ and constructs nested nets that cover $\mathcal{X}_\ell$ at increasingly fine scales. For each candidate, we form a chain from a common anchor through the candidate's projections onto successive nets. Using observations at query points sampled from w_ℓ, the algorithm estimates the anchor reward and the reward differences along the chain. Their sum gives an estimated reward for each candidate. The algorithm then eliminates candidates whose estimates fall sufficiently below the highest estimate. In the next phase, it halves the target error. HT-ChainPE requires exact optimization, sampling, operator, and set operations as specified below. Section 4.2 presents practical HT-ChainPE (Algorithm 4).

Algorithm 1 Heavy-tailed chaining phased elimination (HT-ChainPE)

Require: Nonempty compact domain $\mathcal{X}$, horizon T, kernel k with $\phi(x) = k(x, \cdot)$, RKHS-norm bound B, moment exponent ϵ, noise-moment bound v, and failure probability δ

Initialize $\mathcal{X}_1 \leftarrow \mathcal{X}$, $t \leftarrow 0$, and $\ell \leftarrow 1$
Fix the design certificate factor $\kappa = 1$ and set $\bar{v} \leftarrow 2^{\epsilon}(B^{1+\epsilon} + v)$
Use phase targets $B2^{-\ell}$, failure probabilities $\delta_\ell := 6\delta/(\pi^2\ell^2)$, and regularization $\lambda_T := 8/T$
while $t < T$ **do**
  Choose w_ℓ to maximize $\log\det_F(I + \lambda_T^{-1}A_w)$ over distributions w on $\mathcal{X}_\ell$
  $A_{w_\ell} \leftarrow \mathbb{E}_{x\sim w_\ell}[\phi(x) \otimes \phi(x)]$, $V_\ell \leftarrow A_{w_\ell} + \lambda_T I$, $d_{w_\ell} \leftarrow \mathrm{Tr}(V_\ell^{-1}A_{w_\ell})$
  if $B\sqrt{\lambda_T \kappa d_{w_\ell}} > B2^{-\ell-2}$ **then**
    Sample and query the remaining $T - t$ points i.i.d. from w_ℓ; **stop**
  end if
  Set ρ_ℓ by equation 4 and choose the chain depth as

$$J_\ell := \max\left\{1, \left\lceil \frac{1+\epsilon}{\epsilon}\log_2\left(2^{\ell+3}\cdot C_\epsilon\sqrt{\kappa}\,(1 \vee d_{w_\ell})^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}}\right)\right\rceil\right\}. \tag{3}$$

  $(\{\mathcal{N}_{\ell,j}, \pi_{\ell,j}\}_{j=0}^{J_\ell}, G_\ell, R_\ell) \leftarrow$ BUILDFPSCHAIN$(\mathcal{X}_\ell, \rho_\ell, J_\ell)$ (Algorithm 2)
  Let n_ℓ be the smallest positive integer with $\beta_\ell(n_\ell) \le B2^{-\ell}$ and set $\beta_\ell \leftarrow \beta_\ell(n_\ell)$
  if $n_\ell > T - t$ **then**
    Sample and query the remaining $T - t$ points i.i.d. from w_ℓ; **stop**
  end if
  $\widehat{f}_\ell \leftarrow$ CHAINEDESTIMATE$(\mathcal{X}_\ell, w_\ell, V_\ell, n_\ell, \epsilon, \bar{v}, \delta_\ell, \{\mathcal{N}_{\ell,j}, \pi_{\ell,j}\}_{j=0}^{J_\ell})$ (Algorithm 3)
  $\mathcal{X}_{\ell+1} \leftarrow \overline{\left\{a \in \mathcal{X}_\ell : \widehat{f}_\ell(a) \ge \sup_{u\in\mathcal{X}_\ell}\widehat{f}_\ell(u) - 2\beta_\ell\right\}}$
  Update $t \leftarrow t + n_\ell$ and $\ell \leftarrow \ell + 1$
end while

4.1 HEAVY-TAILED CHAINING PHASED ELIMINATION

D-optimal sampling distribution. We choose a sampling distribution that controls estimation uncertainty throughout $\mathcal{X}_\ell$. Fix $\lambda_T := 8/T$. For a Borel probability distribution w on $\mathcal{X}_\ell$, define $A_w := \mathbb{E}_{x\sim w}[\phi(x) \otimes \phi(x)]$ and $V_w := A_w + \lambda_T I$, where $u \otimes v$ is the rank-one operator $z \mapsto \langle v, z\rangle_{\mathcal{H}_k} u$. HT-ChainPE chooses the regularized D-optimal distribution

$$w_\ell \in \operatorname*{arg\,max}_{w \text{ on } \mathcal{X}_\ell}\big(\log\det_F(I + \lambda_T^{-1}A_w)\big). \tag{2}$$

The symbol $\det_F$ denotes the Fredholm determinant Gohberg et al. (2012). To relate the design objective to uncertainty at individual candidates, define the leverage score $h_w(a) := \|V_w^{-\frac{1}{2}}\phi(a)\|_{\mathcal{H}_k}^2$ for $a \in \mathcal{X}_\ell$. The distribution w determines which feature directions the observations cover, and a larger $h_w(a)$ indicates poorer coverage of $\phi(a)$. The average leverage score $d_w := \mathbb{E}_{x\sim w}[h_w(x)]$ is the effective dimension. Lemma 9 guarantees that a maximizing distribution w_ℓ exists and satisfies $\sup_{a\in\mathcal{X}_\ell} h_{w_\ell}(a) = d_{w_\ell}$. This certificate shows that even the least covered candidate has a leverage score no larger than the design average. Intuitively, w_ℓ distributes sampling weight across feature directions so that no candidate has unusually high uncertainty.

Farthest-point chain (Algorithm 2). Algorithm 2 builds nested nets over $\mathcal{X}_\ell$ using the pseudometric ρ_ℓ. This pseudometric describes how the sampling distribution w_ℓ affects the error in estimating $f(a) - f(b)$:

$$\rho_\ell(a,b)^2 := \mathbb{E}_{x\sim w_\ell}\big[\big\langle\phi(a) - \phi(b), V_\ell^{-1}\phi(x)\big\rangle_{\mathcal{H}_k}^2\big], \qquad a, b \in \mathcal{X}_\ell. \tag{4}$$

The algorithm first chooses a common anchor $a_0 \in \mathcal{X}_\ell$ and sets $\mathcal{N}_{\ell,0} := \{a_0\}$. At each level $j \ge 1$, it starts with $\mathcal{N}_{\ell,j-1}$ and adds the candidate farthest from the current net Gonzalez (1985):

$$\operatorname*{arg\,max}_{a\in\mathcal{X}_\ell}\ \min_{u\in\mathcal{N}_{\ell,j}}\rho_\ell(a,u). \tag{5}$$

It repeats this selection until the net contains 2^{2^j} points or every candidate has zero distance to the net. For each $a \in \mathcal{X}_\ell$, let $\pi_{\ell,j}(a)$ be a nearest point to a in $\mathcal{N}_{\ell,j}$ under ρ_ℓ. The projections

$\pi_{\ell,0}(a), \ldots, \pi_{\ell,J_\ell}(a)$ form a chain from the common anchor to the finest net. As the nets are nested, $\rho_\ell(a, \pi_{\ell,j}(a))$ cannot increase with j. The weighted chain cost G_ℓ controls the estimation error accumulated across levels, while R_ℓ bounds the remaining distance from a candidate to the finest net. The choice of J_ℓ ensures that this remaining distance contributes at most $BR_\ell \le B2^{-\ell-2}$ to the error (Lemma 14).

Chained estimation (Algorithm 3). In each phase, the algorithm draws one batch of query points from w_ℓ and uses the observations to estimate all candidates. To reduce the effect of extreme observations, it clips the kernel-weighted terms used to estimate the anchor reward and the increments along each candidate's chain. We add the anchor estimate and the increments along the chain to estimate the reward at the candidate's projection onto the finest net. We define $\mathrm{clip}_\tau(z) := \mathrm{sign}(z)\min\{|z|, \tau\}$. Let a_0 be the unique point in $\mathcal{N}_{\ell,0}$. Set $u_0 := \log(\frac{4}{\delta_\ell})$ and $u_j := 8\cdot 2^j + \log(\frac{4J_\ell}{\delta_\ell})$ for $j = 1, \ldots, J_\ell$. The anchor threshold is $\tau_0(a_0) := \bar{v}^{\frac{1}{1+\epsilon}} \|A_{w_\ell}^{\frac{1}{2}} V_\ell^{-1} \phi(a_0)\|_{\mathcal{H}_k} (\frac{n_\ell}{u_0})^{\frac{1}{1+\epsilon}}$. For a level-$j$ edge (a, b), the threshold is $\tau_j(a,b) := \bar{v}^{\frac{1}{1+\epsilon}} \rho_\ell(a,b) (\frac{n_\ell}{u_j})^{\frac{1}{1+\epsilon}}$, where $j = 1, \ldots, J_\ell$. For an edge (a, b) at level $j \ge 1$, the estimated increment is

$$\widehat{\Delta}_{\ell,j}(a,b) := \frac{1}{n_\ell} \sum_{i=1}^{n_\ell} \mathrm{clip}_{\tau_j(a,b)} \left(y_i \langle \phi(a) - \phi(b), V_\ell^{-1} \phi(x_i) \rangle_{\mathcal{H}_k} \right). \tag{6}$$

The anchor estimate $\widehat{\Delta}_{\ell,0}(a_0)$ is computed analogously, using $\phi(a_0)$ in place of $\phi(a) - \phi(b)$ and the threshold $\tau_0(a_0)$.

Definition 4 (Chained estimator)**.** *For each $a \in \mathcal{X}_\ell$, the chained estimator is defined as* $\widehat{f}_\ell(a) := \widehat{\Delta}_{\ell,0}\big(\pi_{\ell,0}(a)\big) + \sum_{j=1}^{J_\ell} \widehat{\Delta}_{\ell,j}\big(\pi_{\ell,j}(a), \pi_{\ell,j-1}(a)\big)$.

The estimator has three sources of error: estimating the anchor and chain increments from clipped observations, regularization in V_ℓ, and stopping the chain at the finest net. Lemma 5 bounds their combined effect over $\mathcal{X}_\ell$. Appendix B.4 proves Lemma 5.

Lemma 5 (Uniform confidence bound for the chained estimator)**.** *Under Assumptions 1 and 3, condition on the observations before phase ℓ. Let $\mathcal{X}_\ell$ be nonempty and compact, and let w_ℓ satisfy $\sup_{a\in\mathcal{X}_\ell} h_{w_\ell}(a) \le \kappa d_{w_\ell}$ for $\kappa \ge 1$. Fix the chain produced by Algorithm 2 on $(\mathcal{X}_\ell, \rho_\ell)$ and the sample count n before collecting the phase observations. Then, the chained estimator satisfies $\sup_{a\in\mathcal{X}_\ell} |\widehat{f}_\ell(a) - f(a)| \le \beta_\ell(n)$ with conditional probability at least $1 - \delta_\ell$, where δ_ℓ is defined in Algorithm 1. The confidence radius satisfies*

$$\beta_\ell(n) = \widetilde{\mathcal{O}}_\epsilon\Big(\underbrace{\bar{v}^{\frac{1}{1+\epsilon}} n^{-\frac{\epsilon}{1+\epsilon}} \big(1 + \sqrt{\kappa d_{w_\ell}} + G_\ell \sqrt{J_\ell}\big)}_{\text{stochastic estimation error}} + \underbrace{B\sqrt{\lambda_T \kappa d_{w_\ell}}}_{\text{regularization bias}} + \underbrace{BR_\ell}_{\text{terminal approximation bias}} \Big), \tag{7}$$

where $\widetilde{\mathcal{O}}_\epsilon$ suppresses constants depending only on ϵ and logarithmic factors in δ_ℓ^{-1} and J_ℓ.

The stochastic term in equation 7 includes finite-sample fluctuations and clipping bias in the anchor and increment estimates. For a fixed sampling distribution and chain, increasing n reduces this term, while regularization and chain depth control the two deterministic biases.

Phase length and elimination. The algorithm first determines whether the current phase can reach its target error and then eliminates candidates that are clearly suboptimal. Before constructing the chain or collecting observations, it checks that the regularization bias is at most one quarter of the phase target $B2^{-\ell}$. The chosen chain depth J_ℓ gives the same bound for the terminal approximation bias. The algorithm then selects the smallest sample count n_ℓ for which $\beta_\ell(n_\ell) \le B2^{-\ell}$ and sets $\beta_\ell := \beta_\ell(n_\ell)$. The phase failure probabilities are chosen so that $\sum_{\ell\ge1} \delta_\ell = \delta$. If the bias check fails or n_ℓ exceeds the remaining budget, the algorithm samples from w_ℓ until the horizon ends. After a completed phase, the algorithm retains candidates whose estimated rewards are within $2\beta_\ell$ of the largest estimate:

$$\mathcal{X}_{\ell+1} := \overline{\big\{a \in \mathcal{X}_\ell : \widehat{f}_\ell(a) \ge \sup_{u\in\mathcal{X}_\ell} \widehat{f}_\ell(u) - 2\beta_\ell\big\}}. \tag{8}$$

The threshold accounts for estimation error at both points being compared. On the confidence event in Lemma 5, an optimizer remains active, and every candidate retained for the next phase has reward at least $f(x^\star) - 4\beta_\ell$. Thus, every query in the next phase has regret at most $4\beta_\ell$. Taking the closure preserves this guarantee and keeps the next active set compact.

Table 2: Cumulative model-update arithmetic and continuous optimization calls. m_t is the Nyström dictionary size at round t. L_T is the number of phases initiated by time T. $S < \infty$ is the maximum number of design iterations per phase.

Method	Model-update arithmetic	Optimization calls
TGP-UCB (2019)	$\mathcal{O}\big(\sum_{t=1}^{T}(t^2+t^3+t^2+1)\big)$	T
ATA-GP-UCB (2019)	$\mathcal{O}\big(\sum_{t=1}^{T}(t+2m_t^3+3tm_t^2+tm_t)\big)$	T
CA-TGP-UCB (2024)	$\mathcal{O}\big(\sum_{t=1}^{T}(t^2+t^3+3t^2+1)\big)$	T
MoM-GP-UCB (2024)	$\mathcal{O}\big(\sum_{i=1}^{\lceil T/\log T\rceil}(i^2+i^3+i^2\log T+\log T)\big)$	$\mathcal{O}(T/\log T)$
Practical HT-ChainPE (Ours)	$\mathcal{O}(T+S\sum_{\ell=1}^{L_T}\lvert\mathcal{C}_\ell\rvert^3+\sum_{\ell=1}^{L_T}n_\ell J_\ell\lvert\mathcal{C}_\ell\rvert)$	$\mathcal{O}(\sum_{\ell=1}^{L_T}\lvert\mathcal{C}_\ell\rvert)$

4.2 Practical heavy-tailed chaining phased elimination

HT-ChainPE computes a D-optimal sampling distribution and builds chains directly over the continuous active set $\mathcal{X}_\ell$, which is generally intractable. Practical HT-ChainPE (Algorithm 4) uses a finite set of cover centers $\mathcal{C}_\ell \subseteq \mathcal{X}_\ell$. At phase ℓ, Algorithm 5 selects these centers so that every $a \in \mathcal{X}_\ell$ is within $2^{-\ell-2}$ of a center under the kernel pseudometric $d_k(a,b) := \|\phi(a)-\phi(b)\|_{\mathcal{H}_k}$. Compactness of $\mathcal{X}_\ell$ and continuity of d_k ensure that such a finite cover exists. For each $c \in \mathcal{C}_\ell$, let $\mathcal{V}_\ell(c) \subseteq \mathcal{X}_\ell$ be a closed region. These regions must cover $\mathcal{X}_\ell$ and satisfy $d_k(a,c) \le 2^{-\ell-2}$ for every $a \in \mathcal{V}_\ell(c)$. Closed d_k-balls of radius $2^{-\ell-2}$ provide one such construction. Algorithm 5 uses closed Voronoi cells Aurenhammer (1991) (Appendix C). Practical HT-ChainPE then computes a certified sampling distribution on $\mathcal{C}_\ell$ and uses the center estimates to select the regions retained for the next phase.

Finite distribution and estimation (Algorithm 6). Algorithm 6 applies multiplicative D-optimal updates on $\mathcal{C}_\ell$, increasing the weights of centers whose leverage exceeds the current average Harman & Trnovská (2009). For any fixed $\kappa_{\max} > 1$, full-support initialization yields a distribution w_ℓ satisfying $\max_{c\in\mathcal{C}_\ell} h_{w_\ell}(c) \le \kappa_{\max} d_{w_\ell}$ after finitely many updates (Corollary 15). Algorithms 2 and 3 then construct chains on $\mathcal{C}_\ell$ and estimate center rewards from observations sampled under w_ℓ.

Cellwise elimination. To decide which cells to retain, we compare their center estimates with a margin. The margin accounts for estimation error and reward variation within each cell. For every $a \in \mathcal{V}_\ell(c)$, the reproducing property and Cauchy-Schwarz inequality give

$$|f(a)-f(c)| \le \|f\|_{\mathcal{H}_k}\|\phi(a)-\phi(c)\|_{\mathcal{H}_k} \le B2^{-\ell-2}. \tag{9}$$

For this practical variant, let $\beta_\ell := \beta_\ell(n_\ell) + B2^{-\ell-2}$ denote the total radius, where $\beta_\ell(n_\ell)$ bounds estimation error at the centers. The added cover error transfers this bound to every point in a cell. We choose n_ℓ so that $\beta_\ell \le B2^{-\ell}$. Retain the cells whose centers satisfy $\widehat{f}_\ell(c) \ge \max_{u\in\mathcal{C}_\ell}\widehat{f}_\ell(u) - 2\beta_\ell$, and let $\mathcal{X}_{\ell+1}$ be their union. Because the retained regions are closed subsets of the compact set $\mathcal{X}_\ell$, their finite union $\mathcal{X}_{\ell+1}$ is compact.

Computational complexity. Table 2 compares cumulative model-update arithmetic and continuous optimization calls for practical HT-ChainPE and the baselines over T rounds. We count kernel evaluations outside these calls as unit-cost operations and exclude arithmetic within each call. Each fresh dense matrix factorization is charged cubic cost in its dimension. Practical HT-ChainPE computes a sampling distribution and kernel weights once per phase, then reuses them to process the phase observations. Its arithmetic bound comprises sampling and recording observations, design and preprocessing, and chained estimation, with costs determined by $|\mathcal{C}_\ell|$, J_ℓ, and n_ℓ. The baselines instead update their models from growing histories at each round or episode. Appendix C.7 gives the detailed analysis.

5 Theoretical Results

This section presents high-probability cumulative regret guarantees for HT-ChainPE and its practical variant. We complement these upper bounds with algorithm-independent lower bounds for the squared-exponential (SE) and Matérn RKHS classes.

5.1 CUMULATIVE REGRET UPPER BOUND

Theorem 6 (Cumulative regret upper bound of HT-ChainPE)**.** *Let $\delta \in (0,1)$. Suppose Assumptions 1 and 3 hold. Run HT-ChainPE (Algorithm 1) with regularized D-optimal distributions on its compact active sets. Then, with probability at least $1-\delta$, HT-ChainPE satisfies*

$$R_T \le \widetilde{\mathcal{O}}\big(\bar{v}^{\frac{1}{1+\epsilon}}(1 \vee \gamma_T)^{\frac{1+3\epsilon}{2(1+\epsilon)}} T^{\frac{1}{1+\epsilon}}\big). \tag{10}$$

The proof is presented in Appendix B.5. The bound in equation 10 retains the $T^{\frac{1}{1+\epsilon}}$ horizon dependence of ATA-GP-UCB, CA-TGP-UCB, and MoM-GP-UCB, while improving the linear γ_T dependence shared by all four baselines in Table 1 to $(1 \vee \gamma_T)^{\frac{1+3\epsilon}{2(1+\epsilon)}}$. To the best of our knowledge, this is the first general-kernel cumulative-regret bound for GP bandits under heavy-tailed noise with information-gain dependence of order $\gamma_T^{\frac{1+3\epsilon}{2(1+\epsilon)}}$. This exponent is below one for $\epsilon < 1$, where the moment assumption permits infinite variance, and equals one for $\epsilon = 1$, where it ensures finite variance.

The factor $1 \vee \gamma_T$ is required for a bound uniform over kernels satisfying only $k(x,x) \le 1$, since their information gain can be arbitrarily small. For the normalized SE and Matérn kernels, $\gamma_T \ge \frac{1}{2}\log 2$, and hence $1 \vee \gamma_T \le \frac{2}{\log 2}\gamma_T$; thus, the maximum with one can be absorbed into the constant. For the d-dimensional linear kernel, Theorem 6 and $\gamma_T^{\text{linear}} = \mathcal{O}(d \log T)$ Srinivas et al. (2010) give $R_T \le \widetilde{\mathcal{O}}(\bar{v}^{\frac{1}{1+\epsilon}} d^{\frac{1+3\epsilon}{2(1+\epsilon)}} T^{\frac{1}{1+\epsilon}})$. This recovers, up to logarithmic factors, the cumulative-regret rate of Tajdini et al. (2025) for linear bandits under heavy-tailed noise.

Matérn kernel. For the Matérn kernel, ATA-GP-UCB, CA-TGP-UCB, and MoM-GP-UCB satisfy the bound $R_T \le \widetilde{\mathcal{O}}\big(T^{\frac{1}{1+\epsilon}+\frac{d}{2\nu+d}}(\log T)^{\frac{2\nu}{2\nu+d}}\big)$, which is sublinear when $\nu > d/(2\epsilon)$. Using $\gamma_T^{\text{Matérn}} = \mathcal{O}(T^{\frac{d}{2\nu+d}}(\log T)^{\frac{2\nu}{2\nu+d}})$ (Vakili et al., 2021b), Theorem 6 instead gives

$$R_T \le \widetilde{\mathcal{O}}\big(T^{\frac{1}{1+\epsilon}+\frac{d(1+3\epsilon)}{2(1+\epsilon)(2\nu+d)}}(\log T)^{\frac{\nu(1+3\epsilon)}{(1+\epsilon)(2\nu+d)}}\big), \tag{11}$$

which is sublinear under the weaker condition $\nu > \frac{d(1+\epsilon)}{4\epsilon}$. At $\epsilon = 1$, their rates and sublinearity conditions coincide.

For finite subsets of $[0,1]^d$, Tajdini et al. (2025) establish the Matérn-specific expected-regret rate $\widetilde{\mathcal{O}}(T^{1-\frac{2\nu\epsilon}{(1+\epsilon)(2\nu+d)}})$. Its T-exponent is smaller than ours by $d/(2(2\nu+d))$. However, extending their expected-regret bound to the continuous domain $[0,1]^d$ requires the additional discretization assumption of Vakili et al. (2021a, Assumption 4); see Appendix C.1 of Tajdini et al. (2025). For fixed B, this assumption yields a finite discretization with log-cardinality $\mathcal{O}(d \log T)$. A union bound over this set contributes a factor $\mathcal{O}((d \log T)^{\frac{\epsilon}{1+\epsilon}})$ to the stochastic confidence radius and an explicit $d^{\frac{\epsilon}{1+\epsilon}}$ factor to the resulting regret bound, up to logarithmic terms. In contrast, our result gives a high-probability bound directly on compact query sets without this discretization assumption.

SE kernel. Substituting $\gamma_T^{\text{SE}} = \mathcal{O}((\log T)^{d+1})$ (Srinivas et al., 2010) into the corresponding general-kernel bounds yields $\widetilde{\mathcal{O}}(T^{\frac{1}{1+\epsilon}}(\log T)^{d+1})$ for ATA-GP-UCB, CA-TGP-UCB, and MoM-GP-UCB, and $\widetilde{\mathcal{O}}(T^{\frac{1}{1+\epsilon}}(\log T)^{\frac{(d+1)(1+3\epsilon)}{2(1+\epsilon)}})$ for HT-ChainPE. The logarithmic factor contributed by information gain has a smaller exponent for $\epsilon < 1$ and the same exponent as the three baselines at $\epsilon = 1$; other logarithmic factors remain suppressed by $\widetilde{\mathcal{O}}$. Compared with TGP-UCB, HT-ChainPE also reduces the horizon exponent from $\frac{2+\epsilon}{2(1+\epsilon)}$ to $\frac{1}{1+\epsilon}$.

The following corollary shows that practical HT-ChainPE preserves the horizon and information-gain dependence in Theorem 6, up to the certificate factor $\sqrt{\kappa_{\max}}$. Appendix C gives the proof of Corollary 7.

Corollary 7 (Regret of practical HT-ChainPE)**.** *Under Assumptions 1 and 3, on the compact input domain of Theorem 6, run practical HT-ChainPE (Algorithm 4) with $\kappa_{\max} > 1$. Use full-support initialization and choose the iteration limit so that certification succeeds in every initiated phase. Then, with probability at least $1-\delta$, practical HT-ChainPE satisfies $R_T \le \widetilde{\mathcal{O}}(\bar{v}^{\frac{1}{1+\epsilon}}\sqrt{\kappa_{\max}}(1 \vee \gamma_T)^{\frac{1+3\epsilon}{2(1+\epsilon)}} T^{\frac{1}{1+\epsilon}})$.*

5.2 KERNEL-SPECIFIC CUMULATIVE-REGRET LOWER BOUNDS

We next establish kernel-specific cumulative-regret lower bounds for the SE and Matérn RKHS classes. The theorem below shows that, for any adaptive algorithm, there exist an admissible function and observation law such that the regret exceeds the stated threshold with probability greater than $\delta \in (0, 1/3)$. For a fixed algorithm, mean function f, and observation law P, let $\mathbb{P}_T^{f,P}$ denote the induced distribution of $(x_1, y_1, \ldots, x_T, y_T)$. Table 5 in Appendix D compares the rates, domains, and probability guarantees. The proofs of lower bounds are presented in Appendix D.

Theorem 8 (Kernel-specific cumulative-regret lower bounds)**.** *Fix $d \geq 1$, $B, v > 0$, $\epsilon \in (0, 1]$, and $\delta \in (0, \frac{1}{3})$, and set $L_\delta := (1 - 2\delta) \log \frac{1-\delta}{\delta}$. Let $\mathcal{P}_{\epsilon,v}(f)$ denote the heavy-tailed channel class in Appendix D.1. For every algorithm and all sufficiently large T, the following hold, where $c > 0$ depends only on d, ϵ, and the fixed kernel parameters.*

(i) ***SE kernel.*** *For k_{SE} with any fixed length scale, there exist $f \in \mathcal{H}_{k_{\mathrm{SE}}}([0,1]^d, B)$ and $P \in \mathcal{P}_{\epsilon,v}(f)$ such that*

$$\mathbb{P}_T^{f,P}\Big(R_T > c\, v^{\frac{1}{1+\epsilon}} T^{\frac{1}{1+\epsilon}} L_\delta^{\frac{\epsilon}{1+\epsilon}} \Big(\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big) \Big/ \log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big)\Big)^{\frac{d\epsilon}{1+\epsilon}}\Big) > \delta. \tag{12}$$

(ii) ***Matérn kernel.*** *For k_ν with fixed smoothness $\nu > 0$ and any fixed length scale, there exist $f \in \mathcal{H}_{k_\nu}([0,1]^d, B)$ and $P \in \mathcal{P}_{\epsilon,v}(f)$ such that*

$$\mathbb{P}_T^{f,P}\Big(R_T > c\, v^{\frac{\nu}{\nu(1+\epsilon)+d\epsilon}} B^{\frac{d\epsilon}{\nu(1+\epsilon)+d\epsilon}} T^{\frac{\nu+d\epsilon}{\nu(1+\epsilon)+d\epsilon}} L_\delta^{\frac{\nu\epsilon}{\nu(1+\epsilon)+d\epsilon}}\Big) > \delta. \tag{13}$$

Matérn kernel. For fixed B and v, the high-probability lower bound in equation 13 is $\Omega(T^{\frac{\nu+d\epsilon}{\nu(1+\epsilon)+d\epsilon}} L_\delta^{\frac{\nu\epsilon}{\nu(1+\epsilon)+d\epsilon}})$, whereas Chowdhury & Gopalan (2019) prove the heavy-tailed expected-regret lower bound $\Omega(T^{\frac{\nu+d\epsilon}{\nu(1+\epsilon)+d\epsilon}})$. Thus, the two bounds have the same T-dependence, while equation 13 additionally makes the δ-dependence explicit. At $\epsilon = 1$, our Matérn lower bound becomes $\Omega(T^{\frac{\nu+d}{2\nu+d}} L_\delta^{\frac{\nu}{2\nu+d}})$. This has the same T-exponent as the Gaussian-noise Matérn lower bounds of Scarlett et al. (2017) and Cai & Scarlett (2021). Since $L_\delta \asymp \log(1/\delta)$, the $\log(1/\delta)$ exponent of our bound at $\epsilon = 1$ matches that of Cai & Scarlett (2021).

SE kernel. For fixed problem parameters and δ, equation 12 is $\Omega(T^{\frac{1}{1+\epsilon}} (\log T / \log\log T)^{\frac{d\epsilon}{1+\epsilon}})$. Compared with the heavy-tailed expected lower bound of $\Omega(T^{\frac{1}{1+\epsilon}} (\log T)^{\frac{d\epsilon}{2(1+\epsilon)}})$ Chowdhury & Gopalan (2019), it retains the $T^{\frac{1}{1+\epsilon}}$ dependence and increases the logarithmic exponent from $\frac{d\epsilon}{2(1+\epsilon)}$ to $\frac{d\epsilon}{1+\epsilon}$, up to the $\log\log T$ denominator. At $\epsilon = 1$, equation 12 gives $\Omega\big(v^{1/2} T^{1/2} L_\delta^{1/2} (\frac{\log(B^2 T/(vL_\delta))}{\log\log(B^2 T/(vL_\delta))})^{d/2}\big)$. For fixed B, v, and δ, its T, $\log T$, and $\log\log T$ exponents match those of the Gaussian-noise lower bound of Iwazaki (2026). The bounds are established on different domains, with Iwazaki (2026) considering the unit sphere $\mathbb{S}^d = \{x \in \mathbb{R}^{d+1} : \|x\|_2 = 1\}$ and our result applying to $[0,1]^d$.

6 CONCLUSION

We propose heavy-tailed chaining phased elimination (HT-ChainPE) for Gaussian process (GP) bandits with finite $(1+\epsilon)$-moment noise. On compact query sets, HT-ChainPE achieves the high-probability cumulative-regret bound $\widetilde{\mathcal{O}}(\bar{v}^{\frac{1}{1+\epsilon}} (1 \vee \gamma_T)^{\frac{1+3\epsilon}{2(1+\epsilon)}} T^{\frac{1}{1+\epsilon}})$. This is the first regret bound in this setting with sublinear dependency on γ_T when $\epsilon < 1$. We also establish high-probability regret lower bounds for squared-exponential and Matérn kernels. These lower bounds improve previous results by sharpening the squared-exponential rate and extending the Matérn lower bound to a high-probability guarantee.

AI USE STATEMENT

We used AI tools to polish grammar and wording, identify additional relevant references, and draft portions of the manuscript. We also used these tools to assist with code implementation, debugging,

and synthetic dataset generation. We reviewed all AI-generated code and take full responsibility for the content and results presented in this paper.

Ethics statement

This work studies Gaussian process bandit optimization through theoretical analysis and experiments. The study does not involve human subjects, personal data, or deployment in decision-making systems. We have not identified any direct ethical concerns beyond the general risks associated with the use of optimization methods.

Reproducibility statement

Sections 3–5 specify the assumptions, algorithms, and main results. Appendices A, B, C, and D provide complete proofs and details of the practical algorithm. Appendix E describes the benchmarks, implementation settings, and evaluation protocol. We provide the implementation code as part of the supplementary material.

References

Shubhada Agrawal, Sandeep K Juneja, and Wouter M Koolen. Regret minimization in heavy-tailed bandits. In *Conference on Learning Theory*, pp. 26–62. PMLR, 2021.

Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3):337–404, 1950.

Franz Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM computing surveys (CSUR)*, 23(3):345–405, 1991.

Sébastien Bubeck, Nicolò Cesa-Bianchi, and Gábor Lugosi. Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717, 2013.

Xu Cai and Jonathan Scarlett. On lower bounds for standard and robust gaussian process bandit optimization. In *International Conference on Machine Learning*, pp. 1216–1226. PMLR, 2021.

Roberto Calandra, André Seyfarth, Jan Peters, and Marc Peter Deisenroth. Bayesian optimization for learning gaits under uncertainty: An experimental comparison on a dynamic bipedal walker. *Annals of Mathematics and Artificial Intelligence*, 76(1):5–23, 2016.

Claudio Carmeli, Ernesto De Vito, Alessandro Toigo, and Veronica Umanitá. Vector valued reproducing kernel hilbert spaces and universality. *Analysis and Applications*, 8(01):19–61, 2010.

Duo Cheng, Xingyu Zhou, and Bo Ji. Taming heavy-tailed losses in adversarial bandits and the best-of-both-worlds setting. *Advances in Neural Information Processing Systems*, 37:48083–48129, 2024.

Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17. JMLR.org, 2017.

Sayak Ray Chowdhury and Aditya Gopalan. Bayesian optimization under heavy-tailed payoffs. In *Annual Conference on Neural Information Processing Systems*, pp. 13790–13801, 2019.

Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, 1(2):223, 2001.

Mark E Crovella and Azer Bestavros. Self-similarity in world wide web traffic: Evidence and possible causes. *IEEE/ACM Transactions on networking*, 5(6):835–846, 1997.

David A. Freedman. On tail probabilities for martingales. *The Annals of Probability*, 3(1):100–118, 1975. doi: 10.1214/aop/1176996452.

Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.

Alison L Gibbs and Francis Edward Su. On choosing and bounding probability metrics. *International statistical review*, 70(3):419–435, 2002.

Israel Gohberg, Seymour Goldberg, and Nahum Krupnik. *Traces and determinants of linear operators*, volume 116. Birkhäuser, 2012.

Teofilo F Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical computer science*, 38:293–306, 1985.

Yutian Gou, Jinfeng Yi, and Lijun Zhang. Stochastic graphical bandits with heavy-tailed rewards. In *Uncertainty in Artificial Intelligence*, pp. 734–744. PMLR, 2023.

Radoslav Harman and Mária Trnovská. Approximate d-optimal designs of experiments on the convex hull of a finite set of information matrices. *Mathematica Slovaca*, 59(6):693–704, 2009.

Shogo Iwazaki. Tighter regret lower bound for gaussian process bandits with squared exponential kernel in hypersphere. *arXiv preprint arXiv:2602.17940*, 2026.

Shogo Iwazaki and Shion Takeno. Improved regret analysis in Gaussian process bandits: Optimality for noiseless reward, RKHS norm, and non-stationary variance. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 26642–26672. PMLR, 2025.

David Janz, David Burt, and Javier Gonzalez. Bandit optimisation of functions in the matérn kernel rkhs. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2486–2495. PMLR, 26–28 Aug 2020.

Minhyun Kang and Gi-Soo Kim. Heavy-tailed linear bandit with huber regression. In *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pp. 1027–1036. PMLR, 2023.

Kyungjae Lee and Sungbin Lim. Minimax optimal bandits for heavy tail rewards. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.

Kyungjae Lee, Hongjun Yang, Sungbin Lim, and Songhwai Oh. Optimal algorithms for stochastic multi-armed bandits with heavy tailed rewards. In *Advances in Neural Information Processing Systems*, 2020.

Zihan Li and Jonathan Scarlett. Gaussian process bandit optimization with few batches. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera (eds.), *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 92–107. PMLR, 28–30 Mar 2022.

Wei-Liem Loh. On latin hypercube sampling. *The annals of statistics*, 24(5):2058–2080, 1996.

Shiyin Lu, Guanghui Wang, Yao Hu, and Lijun Zhang. Optimal algorithms for lipschitz bandits with heavy-tailed rewards. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4154–4163. PMLR, 2019.

Andres Muñoz Medina and Scott Yang. No-regret algorithms for heavy-tailed linear bandits. In *Proceedings of the 33nd International Conference on Machine Learning*, volume 48, pp. 1642–1650. JMLR.org, June 2016.

Hyeonjun Park and Kyungjae Lee. Contextually adaptive algorithms for gaussian process bandit optimization under heavy-tailed noise. In *ECAI 2024: 27th European Conference on Artificial Intelligence, 19–24 October 2024, Santiago de Compostela, Spain–Including 13th Conference on Prestigious Applications of Intelligent Systems (PAIS 2024)*, pp. 3171–3178. SAGE Publications Pvt. Ltd 1 Oliver's Yard, 55 City Road, London, EC1Y 1SP, 2024.

Yury Polyanskiy and Yihong Wu. *Information theory: From coding to learning*. Cambridge university press, 2025.

Reinhold Remmert. *Theory of complex functions*, volume 122. Springer Science & Business Media, 1991.

Sudeep Salgia, Sattar Vakili, and Qing Zhao. A domain-shrinking based bayesian optimization algorithm with order-optimal regret performance. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 28836–28847. Curran Associates, Inc., 2021.

Sudeep Salgia, Sattar Vakili, and Qing Zhao. Random exploration in Bayesian optimization: Order-optimal regret and computational efficiency. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 43112–43141. PMLR, 2024.

Jonathan Scarlett, Ilija Bogunovic, and Volkan Cevher. Lower bounds on regret for noisy Gaussian process bandit optimization. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 1723–1742. PMLR, 07–10 Jul 2017.

Han Shao, Xiaotian Yu, Irwin King, and Michael R. Lyu. Almost optimal algorithms for linear stochastic bandits with heavy-tailed payoffs. In *Annual Conference on Neural Information Processing Systems*, pp. 8430–8439, December 2018.

Benjamin J Shields, Jason Stevens, Jun Li, Marvin Parasram, Farhan Damani, Jesus I Martinez Alvarado, Jacob M Janey, Ryan P Adams, and Abigail G Doyle. Bayesian reaction optimization as a tool for chemical synthesis. *Nature*, 590(7844):89–96, 2021.

Barry Simon. *Trace ideals and their applications*. Number 120 in Mathematical Surveys and Monographs. American Mathematical Soc., 2005.

Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems 25*, pp. 2960–2968, December 2012.

Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. of the 27th International Conference on Machine Learning*, pp. 1015–1022. Omnipress, 2010.

Artin Tajdini, Jonathan Scarlett, and Kevin Jamieson. Improved regret bounds for linear bandits with heavy-tailed rewards. *arXiv preprint arXiv:2506.04775*, 2025.

Michel Talagrand. *Upper and Lower Bounds for Stochastic Processes: Decomposition Theorems*, volume 60 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer, 2 edition, 2021.

Lloyd N Trefethen. *Approximation theory and approximation practice, extended edition*. SIAM, 2019.

Sattar Vakili, Nacime Bouziani, Sepehr Jalali, Alberto Bernacchia, and Da shan Shiu. Optimal order simple regret for gaussian process bandits. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021a.

Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in gaussian process bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 82–90. PMLR, 2021b.

Michal Valko, Nathaniel Korda, Rémi Munos, Ilias N. Flaounas, and Nello Cristianini. Finite-time analysis of kernelised contextual bandits. In *Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 2013.

Jing Wang, Yu-Jie Zhang, Peng Zhao, and Zhi-Hua Zhou. Heavy-tailed linear bandits: Huber regression with one-pass update. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 65595–65617. PMLR, 2025.

Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

Bo Xue, Guanghui Wang, Yimu Wang, and Lijun Zhang. Nearly optimal regret for stochastic linear bandits with heavy-tailed payoffs. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pp. 2936–2942. ijcai.org, 2020.

Houssam Zenati, Alberto Bietti, Eustache Diemert, Julien Mairal, Matthieu Martin, and Pierre Gaillard. Efficient kernelized ucb for contextual bandits. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera (eds.), *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 5689–5720. PMLR, 28–30 Mar 2022.

Antoni Zygmund. *Trigonometric series*, volume 1. Cambridge university press, 2002.

CONTENTS

Notation **15**

A Algorithms **16**

A.1 FPS chain 16

A.2 Chained estimator 17

B Proof of the Upper Bound **18**

B.1 Certified experimental design 18

B.2 Variance geometry and chaining 23

B.3 Constructive farthest-point chain 25

B.4 Robust chained estimation 27

B.5 Elimination and regret 30

C Finite-Cover Practical Algorithm and Corollaries **33**

C.1 Practical phase procedure 33

C.2 Kernel-cover construction 33

C.3 Certified design algorithm 34

C.4 Finite certification 34

C.5 Cover confidence and cellwise elimination 35

C.6 Regret of the practical version 37

C.7 Computational complexity 38

D Proof of the Lower Bound **40**

D.1 Heavy-tailed observation channel 41

D.2 Adaptive change of measure 42

D.3 Localized hard-family reduction 43

D.4 SE kernel 45

D.5 Matérn kernel on a Euclidean box 53

D.6 Derivation of the cumulative-regret lower bounds 55

E Experiments **60**

E.1 Experimental setup 60

E.2 Real-world benchmarks 61

E.3 Evaluation 62

E.4 Experimental results 62

E.5 Wall-clock time 65

NOTATION

Table 3: Main notation used throughout the paper.

Symbol	Definition
$\mathcal{X}, d, T, k, \mathcal{H}_k, \phi, B$	Input domain and dimension, horizon, kernel and RKHS, canonical feature map $\phi(x) = k(x, \cdot)$, and RKHS radius in Assumption 1.
$k_t(x), K_t, y_{1:t}, \mu_t, \sigma_t^2, \gamma_T$	Kernel vector, Gram matrix, observation vector, and surrogate mean and variance recalled in Appendix B; maximum information gain is defined in Section 3.
$f, x^\star, R_T$	Unknown reward function, an optimizer, and cumulative pseudo-regret.
$\mathcal{F}_t, x_t, \eta_t, y_t$	History filtration, selected query point, observation noise, and observed reward defined in Section 3.
$P, \mathbb{P}_T^{f,P}$	Observation law and the induced distribution of the query-observation sequence defined in Section 5.2.
$I, \otimes, \mathrm{Tr}, \|\cdot\|_1, \preceq$	Identity map of the required dimension, rank-one operator, operator trace, trace norm, and Loewner order.
$\epsilon, v, \bar{v}$	Moment parameter $\epsilon \in (0, 1]$, noise-moment bound v in Assumption 3, and the derived reward-moment bound $\bar{v} = 2^\epsilon(B^{1+\epsilon} + v)$.
$\delta, \ell, \mathcal{X}_\ell$	Overall failure probability, phase index, and active set at the start of phase ℓ in Algorithm 1.
$\delta_\ell, \lambda_T, J_\ell$	Phase failure probability, fixed regularization, and terminal depth defined in Algorithm 1 and equation 3.
$w, \lambda, A_w, V_w, d_w, h_w, \Phi_\lambda(w)$	Probability distribution, regularization level, covariance operator, regularized operator, effective dimension, leverage score, and regularized D-optimal objective $\Phi_\lambda(w) := \log\det_F(I+\lambda^{-1}A_w)$; see equation 14, equation 15, equation 17, and equation 16.
$\kappa, \kappa_{\max}$	Certificate factor in equation 18: $\kappa = 1$ in Algorithm 1 and $\kappa = \kappa_{\max} > 1$ in the practical version.
$w_\ell, V_\ell, f_{\ell,\lambda_T}$	Phase-ℓ sampling distribution, regularized operator $V_\ell := V_{w_\ell}$, and population target $f_{\ell,\lambda_T}(a) := \langle\phi(a), V_\ell^{-1}A_{w_\ell}f\rangle_{\mathcal{H}_k}$.
ρ, ρ_ℓ	Generic pseudometric and phase-ℓ variance pseudometric in equation 4.
$d_k, \mathcal{C}_\ell, \mathcal{V}_\ell(c)$	Kernel pseudometric in equation 150, phase-ℓ cover centers, and their associated closed regions.
$J, \mathcal{N}_j, \pi_j, G, R$	Depth supplied to Algorithm 2; level-j net, projection, weighted chain cost, and terminal radius returned by the algorithm. Adding ℓ denotes the phase-ℓ version.
$n, \beta_\ell(n), n_\ell, \beta_\ell, \widehat{\Delta}_{\ell,j}, \widehat{f}_\ell$	Candidate sample count, candidate confidence radius, phase length, phase confidence radius, clipped anchor ($j = 0$) or increment ($j \geq 1$) estimate, and chained estimator in Definition 4.
$\mathcal{S}, \overline{A}$	Generic compact candidate set and closure of a subset A in the current active set.
$\det_F$	Fredholm determinant used in the D-optimal selection rule equation 2.
$(\mathcal{A}_j), \mathcal{A}_j(a)$	Admissible partition sequence and the cell containing a, defined by equation 67.
$\Gamma_{\epsilon,2}(\mathcal{X}, \rho), C_\epsilon$	Generic chaining functional in Definition 11 and a common constant depending only on ϵ, enlarged when necessary in the chaining, farthest-point, and concentration bounds.
$k_{\mathrm{SE}}, \ell_{\mathrm{SE}}, k_\nu, \nu, \lambda_0, K_\nu$	Squared-exponential (SE) kernel and its length scale, and Matérn kernel, smoothness, inverse length scale, and modified Bessel function of the second kind defined in Section 3.

Table 4: Algorithm-local notation used in Algorithm 1 and Appendix A.

Symbol	Definition
$t, m, c_i, c_{\ell,i}, S, s, r, w^{(s)}, w_i^{(s)}$	Elapsed-round counter, number of cover centers or design candidates, design candidate, phase-ℓ cover center, iteration limit, iteration index, feature-span dimension, distribution after s multiplicative updates, and its candidate weights.
K_j, z_k	Level-j net-size budget and selected center in Algorithm 2.
clip_τ	Scalar clipping $\mathrm{clip}_\tau(z) = \mathrm{sign}(z)\min\{\lvert z\rvert, \tau\}$ at threshold $\tau \geq 0$.
$a_0, E_j, u_j, \tau_j, Z_{i,j}, \widehat{\Delta}_j$	Anchor, increment edge set ($j \geq 1$), log budget, clipping threshold, weighted observation, and clipped estimate defined in Appendix A.2; level 0 uses only the anchor argument.

A ALGORITHMS

Design quantities. For $u, v \in \mathcal{H}_k$, write $u \otimes v$ for the rank-one operator $z \mapsto \langle v, z\rangle_{\mathcal{H}_k} u$. For a probability distribution w on a candidate set $\mathcal{S} \subseteq \mathcal{X}$ and a regularization level $\lambda > 0$, define

$$A_w := \mathbb{E}_{x\sim w}[\phi(x) \otimes \phi(x)], \tag{14}$$

$$V_w := A_w + \lambda I, \tag{15}$$

$$h_w(a) := \left\|V_w^{-\frac{1}{2}}\phi(a)\right\|_{\mathcal{H}_k}^2, \qquad a \in \mathcal{S}, \tag{16}$$

$$d_w := \mathbb{E}_{x\sim w}\left\|V_w^{-\frac{1}{2}}\phi(x)\right\|_{\mathcal{H}_k}^2. \tag{17}$$

The distribution is certified with factor κ if

$$\sup_{a\in\mathcal{S}} h_w(a) \leq \kappa d_w. \tag{18}$$

For Algorithm 1, take $\mathcal{S} = \mathcal{X}_\ell$ and let w_ℓ be its D-optimal distribution. For the practical version, take $\mathcal{S} = \mathcal{C}_\ell$ and let w_ℓ be the certified distribution on the cover. Write $V_\ell := V_{w_\ell}$. For either version, equation 4 defines ρ_ℓ on the corresponding candidate set $\mathcal{S}$.

A.1 FPS CHAIN

Algorithm 2 Variance-pseudometric farthest-point chain

Require: Nonempty compact set $\mathcal{X}$, continuous pseudometric ρ, moment parameter ϵ, depth J, and fixed measurable tie-breaking rules

$K_0 \leftarrow 1$ and $K_j \leftarrow 2^{2^j}$ for $j = 1, \ldots, J$
Choose $z_1 \in \mathcal{X}$ and set $\mathcal{N}_0 \leftarrow \{z_1\}$
for $j = 1, \ldots, J$ **do**
  $\mathcal{N}_j \leftarrow \mathcal{N}_{j-1}$
  while $|\mathcal{N}_j| < K_j$ and $\sup_{a\in\mathcal{X}} \min_{u\in\mathcal{N}_j} \rho(a, u) > 0$ **do**
    $k \leftarrow |\mathcal{N}_j| + 1$
    Choose $z_k \in \arg\max_{a\in\mathcal{X}} \min_{u\in\mathcal{N}_j} \rho(a, u)$
    $\mathcal{N}_j \leftarrow \mathcal{N}_j \cup \{z_k\}$
  end while
end for
for $j = 0, \ldots, J$ **do**
  For every $a \in \mathcal{X}$, choose $\pi_j(a) \in \arg\min_{u\in\mathcal{N}_j} \rho(a, u)$, with $\pi_j(a) = a$ when $a \in \mathcal{N}_j$
end for
$G \leftarrow \sup_{a\in\mathcal{X}} \Big(\sum_{j=1}^{J} 2^{\frac{2j\epsilon}{1+\epsilon}} \rho\big(\pi_j(a), \pi_{j-1}(a)\big)^2\Big)^{\frac{1}{2}}$
$R \leftarrow \sup_{a\in\mathcal{X}} \rho\big(a, \pi_J(a)\big)$
return $(\{\mathcal{N}_j, \pi_j\}_{j=0}^J, G, R)$

The distance to a finite net is continuous, so every farthest point exists by compactness. Each added point has positive distance from the current net and is therefore distinct. The loop makes finitely many additions because J and all K_j are finite. If the covering radius becomes zero, the same net is used at all remaining levels. This also handles finite candidate sets and pseudometrics that identify distinct points.

A.2 CHAINED ESTIMATOR

Let a_0 be the unique point in $\mathcal{N}_0$. For levels $j = 1, \dots, J$, define the increment edges by

$$E_j := \{(\pi_j(a), \pi_{j-1}(a)) : a \in \mathcal{X}\}. \tag{19}$$

Even when $\mathcal{X}$ is infinite, each E_j is finite because it is contained in $\mathcal{N}_j \times \mathcal{N}_{j-1}$. Fixed tie-breaking makes the projections and the resulting finite-valued estimator Borel measurable. For $n \geq 1$ observations and failure probability $\delta \in (0, 1)$, set

$$u_j := \begin{cases} \log(\frac{4}{\delta}), & j = 0, \\ 8 \cdot 2^j + \log(\frac{4J}{\delta}), & j = 1, \dots, J. \end{cases} \tag{20}$$

The anchor and increment clipping thresholds are

$$\tau_0(a_0) := \bar{v}^{\frac{1}{1+\epsilon}} \|A_w^{\frac{1}{2}} V_w^{-1} \phi(a_0)\|_{\mathcal{H}_k} \Big(\frac{n}{u_0}\Big)^{\frac{1}{1+\epsilon}}, \tag{21}$$

$$\tau_j(a, b) := \bar{v}^{\frac{1}{1+\epsilon}} \rho(a, b) \Big(\frac{n}{u_j}\Big)^{\frac{1}{1+\epsilon}}, \qquad (a, b) \in E_j, \quad j = 1, \dots, J. \tag{22}$$

The anchor factor $\|A_w^{1/2} V_w^{-1} \phi(a_0)\|_{\mathcal{H}_k}$ is the square root of $\mathbb{E}_{x \sim w}[\langle \phi(a_0), V_w^{-1} \phi(x) \rangle_{\mathcal{H}_k}^2]$. Thus, both thresholds use the second moment of the corresponding kernel weight.

Algorithm 3 Phase-wise sampling and chained estimation

Require: $(\mathcal{X}, w, V_w, n, \epsilon, \bar{v}, \delta)$, a frozen chain $\{\mathcal{N}_j, \pi_j\}_{j=0}^J$, and the edges and thresholds in equation 19, equation 21, and equation 22

for $i = 1, \dots, n$ **do**
  Draw $x_i \sim w$, query it, and observe y_i
end for
$Z_{i,0}(a_0) \leftarrow \langle \phi(a_0), V_w^{-1} \phi(x_i) \rangle_{\mathcal{H}_k} y_i$ for $i = 1, \dots, n$
$\widehat{\Delta}_0(a_0) \leftarrow \dfrac{1}{n} \displaystyle\sum_{i=1}^{n} \operatorname{sign}(Z_{i,0}(a_0)) \min\{|Z_{i,0}(a_0)|, \tau_0(a_0)\}$
for $j = 1, \dots, J$ **do**
  for each $(a, b) \in E_j$ **do**
    $Z_{i,j}(a, b) \leftarrow \langle \phi(a) - \phi(b), V_w^{-1} \phi(x_i) \rangle_{\mathcal{H}_k} y_i$ for $i = 1, \dots, n$
    $\widehat{\Delta}_j(a, b) \leftarrow \dfrac{1}{n} \displaystyle\sum_{i=1}^{n} \operatorname{sign}(Z_{i,j}(a, b)) \min\{|Z_{i,j}(a, b)|, \tau_j(a, b)\}$
  end for
end for
For every $a \in \mathcal{X}$, set

$$\widehat{f}(a) := \widehat{\Delta}_0(a_0) + \sum_{j=1}^{J} \widehat{\Delta}_j\big(\pi_j(a), \pi_{j-1}(a)\big).$$

return the function $a \mapsto \widehat{f}(a)$ on $\mathcal{X}$

B Proof of the Upper Bound

This section proves the compact-domain upper bound by combining the D-optimal distribution, variance geometry, finite chain, robust estimator, and elimination analysis. For reuse in the practical corollaries, the geometry, confidence, and regret calculations allow any common certificate factor $\kappa \geq 1$; Algorithm 1 uses $\kappa = 1$.

The argument has five components. First, a regularized D-optimal design is shown to satisfy the certification condition in equation 18 (Lemma 9), and Lemma 10 bounds its effective dimension at regularization at least $8T^{-1}$. The finite multiplicative implementation and its corollaries are treated separately in Appendix C. Second, this certificate places the variance geometry inside a Hilbert ellipsoid controlled by the effective dimension (Lemma 13). Third, the ellipsoid bound is converted into the farthest-point guarantees used by HT-ChainPE (Lemma 14). Fourth, clipped anchor and increment estimates yield a uniform phase-wise confidence event (Lemma 5). Finally, elimination preserves an optimizer and the phase-wise regret bounds aggregate to the main result (Proof of Theorem 6).

Fix a nonempty compact metric input domain $\mathcal{X}$, a bounded continuous kernel k, and $\lambda > 0$. For a Borel probability distribution w on $\mathcal{X}$, use A_w, V_w, d_w, and h_w as defined in equation 14, equation 15, equation 17, and equation 16. The operator-valued expectation defining A_w is interpreted as a Bochner integral. In particular,

$$\mathbb{E}_{x\sim w}[h_w(x)] = d_w. \tag{23}$$

Recall that $\epsilon \in (0, 1]$ (Assumption 3). Notation used throughout the paper is listed in Table 3.

Kernel surrogate and information gain. For $t \geq 1$ and $x \in \mathcal{X}$, define $k_t(x) := (k(x_1, x), \ldots, k(x_t, x))^\top$, $K_t := [k(x_i, x_j)]_{i,j=1}^t$, and $y_{1:t} := (y_1, \ldots, y_t)^\top$. For a regularization parameter $\lambda > 0$, define the kernel surrogate mean and variance (Williams & Rasmussen, 2006; Srinivas et al., 2010) by

$$\mu_t(x) := k_t(x)^\top (K_t + \lambda I)^{-1} y_{1:t}, \tag{24}$$

$$\sigma_t^2(x) := k(x, x) - k_t(x)^\top (K_t + \lambda I)^{-1} k_t(x). \tag{25}$$

At $t = 0$, set $\sigma_0^2(x) := k(x, x)$. Under a zero-mean Gaussian process prior with covariance k and Gaussian observation noise of variance λ, equation 24 and equation 25 coincide with the posterior mean and variance. These quantities describe the standard kernel surrogate used for comparison. For unit observation-noise variance, set $\lambda = 1$ in equation 25. The sequential information-gain identity of Srinivas et al. (2010) states that every fixed query sequence satisfies

$$\frac{1}{2} \log\det(I + K_T) = \frac{1}{2} \sum_{t=1}^{T} \log\big(1 + \sigma_{t-1}^2(x_t)\big). \tag{26}$$

Thus, equation 26 relates information gain to the successive posterior variances.

B.1 Certified experimental design

For $\lambda > 0$, define the regularized D-optimal objective $\Phi_\lambda(w) := \log\det_F(I + \lambda^{-1} A_w)$, which gives equation 2 at $\lambda = \lambda_T$.

Lemma 9 (Optimal-design certificate)**.** *The objective in equation 2 has a maximizer $w^\star$, and every maximizer satisfies*

$$\sup_{a\in\mathcal{X}} h_{w^\star}(a) = d_{w^\star}. \tag{27}$$

Proof. **Proof outline.** The proof first establishes that the objective in equation 2 is well defined and attains its maximum. It then computes the directional derivative toward a point mass and applies first-order optimality to obtain equation 27. Throughout this proof, w denotes a Borel design measure, A_w and $V_w = A_w + \lambda I$ are its covariance and regularized covariance operators, and d_w and $h_w(a)$ are its average and pointwise leverage scores.

Well-definedness. For any $a, b \in \mathcal{X}$, the reproducing property gives

$$\begin{aligned}\|\phi(a) - \phi(b)\|_{\mathcal{H}_k}^2 &= \langle \phi(a) - \phi(b), \phi(a) - \phi(b) \rangle_{\mathcal{H}_k} \\ &= \langle \phi(a), \phi(a) \rangle_{\mathcal{H}_k} + \langle \phi(b), \phi(b) \rangle_{\mathcal{H}_k} - 2\langle \phi(a), \phi(b) \rangle_{\mathcal{H}_k} \\ &= k(a,a) + k(b,b) - 2k(a,b). \end{aligned} \tag{28}$$

Thus, the continuity of k implies that $\phi : \mathcal{X} \to \mathcal{H}_k$ is continuous. For $u, v, g \in \mathcal{H}_k$, let $(u \otimes v)g := \langle v, g \rangle_{\mathcal{H}_k} u$ denote the associated rank-one operator. Then,

$$u \otimes u - v \otimes v = (u - v) \otimes u + v \otimes (u - v), \tag{29}$$

and the trace norm of a rank-one operator satisfies $\|u \otimes v\|_1 = \|u\|_{\mathcal{H}_k} \|v\|_{\mathcal{H}_k}$. Hence,

$$\begin{aligned}\|u \otimes u - v \otimes v\|_1 &\le \|u\|_{\mathcal{H}_k} \|u - v\|_{\mathcal{H}_k} + \|u - v\|_{\mathcal{H}_k} \|v\|_{\mathcal{H}_k} \\ &= \big(\|u\|_{\mathcal{H}_k} + \|v\|_{\mathcal{H}_k}\big) \|u - v\|_{\mathcal{H}_k}. \end{aligned} \tag{30}$$

The bound in equation 30 shows that $x \mapsto \phi(x) \otimes \phi(x)$ is continuous in trace norm. Assumption 1 gives

$$\|\phi(x) \otimes \phi(x)\|_1 = \|\phi(x)\|_{\mathcal{H}_k}^2 = k(x,x) \le 1. \tag{31}$$

Hence, the trace-norm Bochner integral in equation 14 exists and A_w is trace class. For every $g \in \mathcal{H}_k$, substituting equation 14 into its quadratic form gives

$$\begin{aligned}\langle g, A_w g \rangle_{\mathcal{H}_k} &= \mathbb{E}_{x \sim w}[\langle g, (\phi(x) \otimes \phi(x)) g \rangle_{\mathcal{H}_k}] \\ &= \mathbb{E}_{x \sim w}[|\langle g, \phi(x) \rangle_{\mathcal{H}_k}|^2] \ge 0. \end{aligned} \tag{32}$$

Thus, A_w is positive. Let $(e_r)_{r \ge 1}$ be an orthonormal basis of the closed linear span of $\phi(\mathcal{X})$. For every r, the function $x \mapsto |\langle e_r, \phi(x) \rangle_{\mathcal{H}_k}|^2$ is measurable and nonnegative by equation 32, so Tonelli's theorem permits exchanging the sum and the integral. The rank-one trace identity and Parseval's identity Zygmund (2002) then give

$$\begin{aligned}\operatorname{Tr}(A_w) &= \sum_{r \ge 1} \langle e_r, A_w e_r \rangle_{\mathcal{H}_k} = \sum_{r \ge 1} \mathbb{E}_{x \sim w} \big| \langle e_r, \phi(x) \rangle_{\mathcal{H}_k} \big|^2 \\ &= \mathbb{E}_{x \sim w} \Big[\sum_{r \ge 1} \big| \langle e_r, \phi(x) \rangle_{\mathcal{H}_k} \big|^2 \Big] = \mathbb{E}_{x \sim w} \|\phi(x)\|_{\mathcal{H}_k}^2 \\ &= \mathbb{E}_{x \sim w} \operatorname{Tr}\big(\phi(x) \otimes \phi(x)\big) = \mathbb{E}_{x \sim w}[k(x,x)] \le 1. \end{aligned} \tag{33}$$

The fourth equality follows from Parseval's identity. Let $(\xi_r)_{r \ge 1}$ be the nonnegative eigenvalues of A_w, counted with multiplicity. The trace bound in equation 33 gives $\sum_{r \ge 1} \xi_r = \operatorname{Tr}(A_w) \le 1$. Therefore,

$$\begin{aligned}1 \le \det_F(I + \lambda^{-1} A_w) &= \prod_{r \ge 1} \Big(1 + \frac{\xi_r}{\lambda}\Big) \\ &\le \prod_{r \ge 1} \exp\Big(\frac{\xi_r}{\lambda}\Big) = \exp\Big(\lambda^{-1} \operatorname{Tr}(A_w)\Big) \le \exp(\lambda^{-1}) < \infty. \end{aligned} \tag{34}$$

The first inequality in equation 34 follows because every factor is at least one, and the second applies $1 + t \le e^t$ to each factor. Thus, the Fredholm determinant in equation 2 is finite and strictly positive, and its logarithm is well defined.

Existence of a maximizer. It remains to prove that the objective attains its maximum. By equation 30, the map $x \mapsto \phi(x) \otimes \phi(x)$ is trace-norm continuous and hence uniformly continuous on compact $\mathcal{X}$. Suppose that w_n converges weakly to w. Fix $\eta > 0$. Uniform continuity gives finitely many open sets $U_1, \dots, U_m$ covering $\mathcal{X}$ and points $x_i \in U_i$ such that

$$x \in U_i \quad \Longrightarrow \quad \big\|\phi(x) \otimes \phi(x) - \phi(x_i) \otimes \phi(x_i)\big\|_1 \le \eta. \tag{35}$$

Let $\psi_1, \dots, \psi_m$ be a continuous partition of unity subordinate to this cover. Thus, each ψ_i is nonnegative, $\sum_{i=1}^m \psi_i(x) = 1$, and $\psi_i(x) > 0$ only when $x \in U_i$. Using $\sum_{i=1}^m \psi_i(x) = 1$ gives

$$\phi(x) \otimes \phi(x) - \sum_{i=1}^m \psi_i(x) \big(\phi(x_i) \otimes \phi(x_i)\big) = \sum_{i=1}^m \psi_i(x) \Big(\phi(x) \otimes \phi(x) - \phi(x_i) \otimes \phi(x_i)\Big), \tag{36}$$

and every nonzero summand satisfies equation 35. The triangle inequality and the nonnegativity of the partition weights therefore yield

$$\sup_{x\in\mathcal{X}} \Big\| \phi(x)\otimes\phi(x) - \sum_{i=1}^{m} \psi_i(x)\big(\phi(x_i)\otimes\phi(x_i)\big) \Big\|_1 \le \eta. \tag{37}$$

Since each ψ_i is bounded and continuous, weak convergence gives

$$\mathbb{E}_{x\sim w_n}[\psi_i(x)] \longrightarrow \mathbb{E}_{x\sim w}[\psi_i(x)] \tag{38}$$

for every $i \in \{1,\dots,m\}$. Integrating equation 37 with respect to w_n and w shows that each covariance operator is within η of the corresponding finite sum. Therefore,

$$\|A_{w_n} - A_w\|_1 \le 2\eta + \left\| \sum_{i=1}^{m} \Big(\mathbb{E}_{x\sim w_n}[\psi_i(x)] - \mathbb{E}_{x\sim w}[\psi_i(x)] \Big) \big(\phi(x_i)\otimes\phi(x_i)\big) \right\|_1. \tag{39}$$

The last term in equation 39 is bounded by

$$\sum_{i=1}^{m} \big\|\phi(x_i)\otimes\phi(x_i)\big\|_1 \Big| \mathbb{E}_{x\sim w_n}[\psi_i(x)] - \mathbb{E}_{x\sim w}[\psi_i(x)] \Big|, \tag{40}$$

which converges to zero. It follows that

$$\limsup_{n\to\infty} \|A_{w_n} - A_w\|_1 \le \inf_{\eta>0} 2\eta = 0. \tag{41}$$

Hence, $\|A_{w_n} - A_w\|_1 \to 0$. The Fredholm determinant is continuous in trace norm (Simon, 2005), so $\Phi_\lambda(w_n) \to \Phi_\lambda(w)$. The Borel probability measures on compact $\mathcal{X}$ form a weakly compact set, so Φ_λ attains its maximum.

Directional derivative. The directional derivative toward a point mass is computed next. Fix $a \in \mathcal{X}$, and let δ_a denote the point mass at a. For $t \in [0,1]$, set $w_t := (1-t)w + t\delta_a$. Linearity of the Bochner integral gives

$$A_{w_t} = A_w + t\big(\phi(a)\otimes\phi(a) - A_w\big). \tag{42}$$

Substituting equation 42 and $V_w = A_w + \lambda I$ gives

$$\begin{aligned} I + \lambda^{-1}A_{w_t} &= I + \lambda^{-1}A_w + t\lambda^{-1}\big(\phi(a)\otimes\phi(a) - A_w\big) \\ &= \lambda^{-1}V_w + t\lambda^{-1}\big(\phi(a)\otimes\phi(a) - A_w\big) \\ &= \lambda^{-1}V_w\Big[I + tV_w^{-1}\big(\phi(a)\otimes\phi(a) - A_w\big)\Big] \\ &= \big(I + \lambda^{-1}A_w\big)\Big[I + tV_w^{-1}\big(\phi(a)\otimes\phi(a) - A_w\big)\Big]. \end{aligned} \tag{43}$$

The operator $V_w^{-1}\big(\phi(a)\otimes\phi(a) - A_w\big)$ is trace class because V_w^{-1} is bounded, $\phi(a)\otimes\phi(a)$ is rank one, and A_w is trace class. This is the ideal property of trace-class operators: multiplication on either side by a bounded operator preserves trace class. The factorization in equation 43 and multiplicativity of the Fredholm determinant (Gohberg et al., 2012, Chapter VII, Theorem 3.3) imply

$$\det_F\big(I + \lambda^{-1}A_{w_t}\big) = \det_F\big(I + \lambda^{-1}A_w\big) \times \det_F\Big(I + tV_w^{-1}\big(\phi(a)\otimes\phi(a) - A_w\big)\Big). \tag{44}$$

The operators $I + \lambda^{-1}A_{w_t}$ and $I + \lambda^{-1}A_w$ are strictly positive, so their Fredholm determinants are positive; the remaining determinant in equation 44 is therefore their positive ratio. Taking logarithms and subtracting the term corresponding to w gives

$$\begin{aligned} \Phi_\lambda(w_t) - \Phi_\lambda(w) &= \log\det_F\big(I + \lambda^{-1}A_{w_t}\big) - \log\det_F\big(I + \lambda^{-1}A_w\big) \\ &= \log\det_F\Big(I + tV_w^{-1}\big(\phi(a)\otimes\phi(a) - A_w\big)\Big). \end{aligned} \tag{45}$$

The first-order expansion of the Fredholm log-determinant (Gohberg et al., 2012, Chapter VII, Theorem 4.2) yields

$$\Phi_\lambda(w_t) - \Phi_\lambda(w) = t\,\mathrm{Tr}\Big(V_w^{-1}\big[\phi(a)\otimes\phi(a) - A_w\big]\Big) + o(t). \tag{46}$$

Dividing by t and letting $t \downarrow 0$ gives

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\Phi_\lambda(w_t)\Big|_{t=0+} &= \mathrm{Tr}\Big(V_w^{-1}\big[\phi(a)\otimes\phi(a) - A_w\big]\Big)\\ &= \mathrm{Tr}\Big(V_w^{-1}\big[\phi(a)\otimes\phi(a)\big]\Big) - \mathrm{Tr}(V_w^{-1}A_w)\\ &= h_w(a) - d_w. \end{aligned} \tag{47}$$

The first term is $h_w(a)$ by the rank-one trace identity. For the second term, boundedness of V_w^{-1} and the Bochner integral in equation 14 give

$$\mathrm{Tr}(V_w^{-1}A_w) = \mathbb{E}_{x\sim w}\langle\phi(x), V_w^{-1}\phi(x)\rangle_{\mathcal{H}_k} = d_w, \tag{48}$$

This identifies the second trace term in equation 47 with d_w.

First-order optimality. It remains to apply first-order optimality to a maximizer. At a maximizer $w^\star$, the derivative in equation 47 is nonpositive for every a, so $h_{w^\star}(a) \le d_{w^\star}$. Conversely, equation 23 gives $\sup_{a\in\mathcal{X}} h_{w^\star}(a) \ge d_{w^\star}$. This proves equation 27. □

Lemma 10 (Effective dimension and information gain)**.** *Under Assumption 1, every Borel probability measure w and every $\lambda \ge 8T^{-1}$ satisfy*

$$d_w \le \gamma_T. \tag{49}$$

Proof. **Proof outline.** The proof derives a lower bound on the random-design information gain from the full sum of expected posterior variances and then compares the resulting spectral sum with the effective dimension at regularization $8T^{-1}$. The quantities γ_T and d_w denote the maximum information gain over T queries and the effective dimension of the design distribution w, respectively.

Random-design information gain. Let $X_1, \dots, X_T$ be independent samples from w, and define

$$S_{t-1} := \sum_{i=1}^{t-1} \phi(X_i)\otimes\phi(X_i). \tag{50}$$

The posterior variance associated with $X_1, \dots, X_{t-1}$ is

$$\sigma_{t-1}^2(x) = \big\langle\phi(x), (I + S_{t-1})^{-1}\phi(x)\big\rangle_{\mathcal{H}_k}. \tag{51}$$

Assumption 1 gives $0 \le \sigma_{t-1}^2(x) \le k(x,x) \le 1$. Definition 2, equation 26, and $\log(1+u) \ge \frac{u}{2}$ for $u \in [0,1]$ therefore give

$$\begin{aligned}\gamma_T &\ge \frac{1}{2}\sum_{t=1}^{T} \mathbb{E}\big[\log(1+\sigma_{t-1}^2(X_t))\big]\\ &\ge \frac{1}{4}\sum_{t=1}^{T} \mathbb{E}\big[\sigma_{t-1}^2(X_t)\big]. \end{aligned} \tag{52}$$

Expected posterior variance. Conditioning on $X_1, \dots, X_{t-1}$ and using equation 14 yield

$$\mathbb{E}\big[\sigma_{t-1}^2(X_t) \mid X_1, \dots, X_{t-1}\big] = \mathrm{Tr}\left(A_w(I + S_{t-1})^{-1}\right). \tag{53}$$

For positive operators, $\succeq$ denotes the Loewner order. For every positive operator Q and $u \in \mathcal{H}_k$, we have $I + Q \succeq I$. Hence, $I + Q$ is invertible and has a positive square root. For a scalar c, let $\mathrm{Re}(c)$ denote its real part. For any $z \in \mathcal{H}_k$, completing the square gives

$$\begin{aligned}&2\,\mathrm{Re}\langle u, z\rangle_{\mathcal{H}_k} - \|z\|_{\mathcal{H}_k}^2 - \langle z, Qz\rangle_{\mathcal{H}_k}\\ &= 2\,\mathrm{Re}\langle u, z\rangle_{\mathcal{H}_k} - \langle z, (I+Q)z\rangle_{\mathcal{H}_k}\\ &= \langle u, (I+Q)^{-1}u\rangle_{\mathcal{H}_k} - \Big\|(I+Q)^{\frac{1}{2}}\big(z - (I+Q)^{-1}u\big)\Big\|_{\mathcal{H}_k}^2. \end{aligned} \tag{54}$$

The squared norm is nonnegative, so the expression on the left is at most $\langle u, (I+Q)^{-1}u\rangle_{\mathcal{H}_k}$. At $z = (I+Q)^{-1}u$, the squared norm is zero and this upper bound is attained. Therefore,

$$\langle u, (I+Q)^{-1}u\rangle_{\mathcal{H}_k} = \sup_{z\in\mathcal{H}_k} \Big\{ 2\,\mathrm{Re}\langle u, z\rangle_{\mathcal{H}_k} - \|z\|^2_{\mathcal{H}_k} - \langle z, Qz\rangle_{\mathcal{H}_k} \Big\}. \tag{55}$$

For each fixed z, the expression inside the supremum is affine in Q because its only Q-dependent term is $-\langle z, Qz\rangle_{\mathcal{H}_k}$. A pointwise supremum of affine functions is convex, so $Q \mapsto \langle u, (I+Q)^{-1}u\rangle_{\mathcal{H}_k}$ is convex. Jensen's inequality and $\mathbb{E}[S_{t-1}] = (t-1)A_w$ imply

$$\mathbb{E}\big[(I+S_{t-1})^{-1}\big] \succeq \big(I + (t-1)A_w\big)^{-1}. \tag{56}$$

Taking the trace against the positive trace-class operator A_w in equation 53 gives

$$\mathbb{E}\big[\sigma^2_{t-1}(X_t)\big] \geq \mathrm{Tr}\left(A_w(I+(t-1)A_w)^{-1}\right). \tag{57}$$

Spectral comparison. Let $(\zeta_r)_{r\geq 1}$ be the eigenvalues of A_w. Since A_w is positive and trace class, the spectral theorem expresses the trace of a function of A_w as the sum of that function over $(\zeta_r)_{r\geq 1}$. In particular, for every $t \geq 1$,

$$\mathrm{Tr}\left(A_w(I+(t-1)A_w)^{-1}\right) = \sum_{r\geq 1} \frac{\zeta_r}{1+(t-1)\zeta_r}. \tag{58}$$

Multiplying equation 52 by four, applying equation 57 for each t, and substituting this trace identity give

$$4\gamma_T \geq \sum_{r\geq 1}\sum_{t=1}^{T} \frac{\zeta_r}{1+(t-1)\zeta_r}. \tag{59}$$

For each r, the function $s \mapsto \frac{\zeta_r}{1+s\zeta_r}$ is nonincreasing on $[0,T]$. Hence, its left Riemann sum over the unit intervals dominates its integral:

$$\sum_{t=1}^{T} \frac{\zeta_r}{1+(t-1)\zeta_r} \geq \int_0^T \frac{\zeta_r}{1+s\zeta_r}\,\mathrm{d}s. \tag{60}$$

The integrand is the derivative of $\log(1+s\zeta_r)$, including when $\zeta_r = 0$, so

$$\int_0^T \frac{\zeta_r}{1+s\zeta_r}\,\mathrm{d}s = \log(1+T\zeta_r). \tag{61}$$

Summing the Riemann-sum bound and the integral identity over r gives

$$4\gamma_T \geq \sum_{r\geq 1} \log(1+T\zeta_r). \tag{62}$$

For $x \geq 0$, an elementary derivative check gives $\log(1+x) \geq \frac{4x}{x+8}$. Applying this scalar inequality with $x = T\zeta_r$ and summing over r yield

$$\sum_{r\geq 1} \log(1+T\zeta_r) \geq 4\sum_{r\geq 1} \frac{T\zeta_r}{T\zeta_r + 8}. \tag{63}$$

The spectral theorem also gives

$$\begin{aligned}\mathrm{Tr}\left(A_w(A_w + 8T^{-1}I)^{-1}\right) &= \sum_{r\geq 1} \frac{\zeta_r}{\zeta_r + 8T^{-1}} \\ &= \sum_{r\geq 1} \frac{T\zeta_r}{T\zeta_r + 8}.\end{aligned} \tag{64}$$

Substituting the scalar logarithm bound and the spectral trace identity into the information-gain lower bound gives

$$4\gamma_T \geq 4\,\mathrm{Tr}\left(A_w(A_w + 8T^{-1}I)^{-1}\right). \tag{65}$$

If $\lambda \geq 8T^{-1}$, the spectral theorem and equation 65 give

$$\begin{aligned} d_w &= \mathrm{Tr}\left(A_w(A_w+\lambda I)^{-1}\right) \\ &\leq \mathrm{Tr}\left(A_w(A_w + 8T^{-1}I)^{-1}\right) \leq \gamma_T. \end{aligned} \tag{66}$$

This proves equation 49. □

B.2 VARIANCE GEOMETRY AND CHAINING

Fix a phase ℓ and let w_ℓ be a Borel probability distribution on a nonempty compact candidate set $\mathcal{X}_\ell$ satisfying equation 18 with factor κ. Algorithm 1 satisfies this condition with $\kappa = 1$ by Lemma 9. The statements in this subsection and the following two subsections also apply to a finite candidate set by substituting $\mathcal{C}_\ell$ for $\mathcal{X}_\ell$.

Definition 11 (Generic chaining functional)**.** *Let $(\mathcal{X}, \rho)$ be a pseudometric space. Following Talagrand (2021), a sequence of partitions $(\mathcal{A}_j)_{j\geq 0}$ is admissible if*

$$\mathcal{A}_0 = \{\mathcal{X}\}, \qquad |\mathcal{A}_j| \leq 2^{2^j} \quad (j \geq 1), \tag{67}$$

and $\mathcal{A}_{j+1}$ refines $\mathcal{A}_j$. Writing $\mathcal{A}_j(a)$ for the cell containing a, define the ϵ-indexed functional

$$\Gamma_{\epsilon,2}(\mathcal{X}, \rho) := \inf_{(\mathcal{A}_j)_{j\geq 0}} \sup_{a\in\mathcal{X}} \Big(\sum_{j\geq 0} 2^{\frac{2j\epsilon}{1+\epsilon}} \operatorname{diam}_\rho\big(\mathcal{A}_j(a)\big)^2 \Big)^{\frac{1}{2}}. \tag{68}$$

The following Hilbert-space formulation is equivalent to the sequence-space result of Talagrand (2021, Corollary 4.1.6) under the canonical coefficient isometry.

Theorem 12 (Hilbert-ellipsoid chaining bound Talagrand (2021))**.** *Let $\mathcal{H}$ be a Hilbert space, let $(u_r)_{r\geq 1}$ be an orthonormal system in $\mathcal{H}$, and let $(a_r)_{r\geq 1}$ be a bounded sequence of nonnegative numbers. Define*

$$\mathcal{E}(a) := \Big\{ \sum_{r\geq 1} a_r \alpha_r u_r : \sum_{r\geq 1} \alpha_r^2 \leq 1 \Big\}. \tag{69}$$

Then, for every $\epsilon \in (0, 1]$, there exists $C_\epsilon > 0$ depending only on ϵ such that

$$\Gamma_{\epsilon,2}\big(\mathcal{E}(a), \|\cdot\|_{\mathcal{H}}\big) \leq C_\epsilon \sup_{t>0} t \,\big|\{r \geq 1 : a_r \geq t\}\big|^{\frac{\epsilon}{1+\epsilon}}. \tag{70}$$

Lemma 13 (Certified variance geometry and chaining)**.** *Define the ρ_ℓ-diameter of $\mathcal{X}_\ell$ by $\operatorname{diam}_{\rho_\ell}(\mathcal{X}_\ell) := \sup_{a,b\in\mathcal{X}_\ell} \rho_\ell(a, b)$. If w_ℓ is certified with factor κ, then ρ_ℓ is a pseudometric on $\mathcal{X}_\ell$ and*

$$\operatorname{diam}_{\rho_\ell}(\mathcal{X}_\ell) \leq 2\sqrt{\kappa d_{w_\ell}}. \tag{71}$$

Moreover,

$$\Gamma_{\epsilon,2}(\mathcal{X}_\ell, \rho_\ell) \leq C_\epsilon \sqrt{\kappa}\, d_{w_\ell}^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}}. \tag{72}$$

Proof. **Proof outline.** The proof represents ρ_ℓ as a Hilbert-space distance and uses certification to bound its diameter. It then embeds the image into a Hilbert ellipsoid and controls the chaining functional through the ellipsoid's singular values. Recall that A_{w_ℓ} is the covariance operator of w_ℓ, $V_\ell = A_{w_\ell} + \lambda_T I$ is its regularized covariance operator, d_{w_ℓ} is its effective dimension, and ρ_ℓ is its induced variance pseudometric.

Pseudometric representation and diameter. Expanding equation 4 and using equation 14 gives

$$\rho_\ell(a, b)^2 = \big\langle \phi(a) - \phi(b), V_\ell^{-1} A_{w_\ell} V_\ell^{-1} \big(\phi(a) - \phi(b)\big) \big\rangle_{\mathcal{H}_k}. \tag{73}$$

Since A_{w_ℓ} is positive, equation 73 implies

$$\rho_\ell(a, b) = \big\| A_{w_\ell}^{\frac{1}{2}} V_\ell^{-1} \big(\phi(a) - \phi(b)\big) \big\|_{\mathcal{H}_k}. \tag{74}$$

The identity in equation 74 expresses ρ_ℓ as the Hilbert-space distance between the images of two candidate points, so it is a pseudometric on $\mathcal{X}_\ell$. It need not be a metric because the map need not be injective. The map $a \mapsto A_{w_\ell}^{1/2} V_\ell^{-1} \phi(a)$ is continuous because ϕ is continuous and the operator is bounded. Its image of $\mathcal{X}_\ell$ is compact, which justifies the farthest-point maximizations used below. Throughout this proof, $\preceq$ denotes the Loewner order on self-adjoint operators. Since $0 \preceq A_{w_\ell} \preceq V_\ell$, congruence by V_ℓ^{-1} gives

$$V_\ell^{-1} A_{w_\ell} V_\ell^{-1} \preceq V_\ell^{-1}. \tag{75}$$

Substituting $V_\ell^{-1} A_{w_\ell} V_\ell^{-1} \preceq V_\ell^{-1}$ into equation 73 and applying the certification condition in equation 18 give

$$\begin{aligned}\rho_\ell(a,b) &\le \|V_\ell^{-\frac{1}{2}}\big(\phi(a)-\phi(b)\big)\|_{\mathcal{H}_k} \\ &\le \|V_\ell^{-\frac{1}{2}}\phi(a)\|_{\mathcal{H}_k} + \|V_\ell^{-\frac{1}{2}}\phi(b)\|_{\mathcal{H}_k} \\ &\le 2\sqrt{\kappa d_{w_\ell}}. \end{aligned} \tag{76}$$

Taking the supremum over $a, b \in \mathcal{X}_\ell$ proves equation 71.

Ellipsoid embedding. The factorization in equation 73 also gives

$$\rho_\ell(a,b) = \Big\| A_{w_\ell}^{\frac{1}{2}} V_\ell^{-\frac{1}{2}} V_\ell^{-\frac{1}{2}} \big(\phi(a)-\phi(b)\big)\Big\|_{\mathcal{H}_k}. \tag{77}$$

Hence, the map in equation 74 preserves every ρ_ℓ-distance. The certification condition in equation 18 gives

$$\|V_\ell^{-\frac{1}{2}}\phi(a)\|_{\mathcal{H}_k}^2 = h_{w_\ell}(a) \le \kappa d_{w_\ell} \tag{78}$$

for every $a \in \mathcal{X}_\ell$. Therefore, the image of equation 74 is contained in the ellipsoid

$$\mathcal{E}_\ell := \Big\{ A_{w_\ell}^{\frac{1}{2}} V_\ell^{-\frac{1}{2}} z : \|z\|_{\mathcal{H}_k} \le \sqrt{\kappa d_{w_\ell}} \Big\}. \tag{79}$$

Spectral chaining bound. Since A_{w_ℓ} is trace class and $V_\ell^{-\frac{1}{2}}$ is bounded, $V_\ell^{-\frac{1}{2}} A_{w_\ell} V_\ell^{-\frac{1}{2}}$ is positive and trace class. For an operator T, write T^* for its Hilbert-space adjoint. Moreover,

$$\big(A_{w_\ell}^{\frac{1}{2}} V_\ell^{-\frac{1}{2}}\big)^* \big(A_{w_\ell}^{\frac{1}{2}} V_\ell^{-\frac{1}{2}}\big) = V_\ell^{-\frac{1}{2}} A_{w_\ell} V_\ell^{-\frac{1}{2}}. \tag{80}$$

Hence, $A_{w_\ell}^{\frac{1}{2}} V_\ell^{-\frac{1}{2}}$ is Hilbert-Schmidt and therefore compact. Let $(s_r)_{r\ge1}$ denote its singular values in nonincreasing order, repeated according to multiplicity. Since $A_{w_\ell} \preceq V_\ell$, congruence by $V_\ell^{-\frac{1}{2}}$ gives $V_\ell^{-\frac{1}{2}} A_{w_\ell} V_\ell^{-\frac{1}{2}} \preceq I$, so $0 \le s_r \le 1$. Cyclicity of the trace and equation 23 give

$$\sum_{r\ge1} s_r^2 = \mathrm{Tr}\big(V_\ell^{-\frac{1}{2}} A_{w_\ell} V_\ell^{-\frac{1}{2}}\big) = \mathrm{Tr}(V_\ell^{-1} A_{w_\ell}) = d_{w_\ell}. \tag{81}$$

For any admissible sequence on $\mathcal{E}_\ell$, the nonempty inverse images of its cells under equation 74 form an admissible sequence on $\mathcal{X}_\ell$. The number of cells does not increase, refinement is preserved, and each inverse-image cell has no larger diameter because the map preserves distances. Definition 11 therefore gives

$$\Gamma_{\epsilon,2}(\mathcal{X}_\ell, \rho_\ell) \le \Gamma_{\epsilon,2}(\mathcal{E}_\ell, \|\cdot\|_{\mathcal{H}_k}). \tag{82}$$

The ellipsoid $\mathcal{E}_\ell$ has semiaxes $(\sqrt{\kappa d_{w_\ell}}\, s_r)_{r\ge1}$. Since $\epsilon \in (0,1]$, we have $\frac{1+\epsilon}{\epsilon} \ge 2$. For every $t > 0$,

$$\begin{aligned} t^{\frac{1+\epsilon}{\epsilon}} \big|\{r : \sqrt{\kappa d_{w_\ell}}\, s_r \ge t\}\big| &\le \sum_{r:\sqrt{\kappa d_{w_\ell}}\, s_r \ge t} \big(\sqrt{\kappa d_{w_\ell}}\, s_r\big)^{\frac{1+\epsilon}{\epsilon}} \\ &\le (\kappa d_{w_\ell})^{\frac{1+\epsilon}{2\epsilon}} \sum_{r\ge1} s_r^{\frac{1+\epsilon}{\epsilon}}. \end{aligned} \tag{83}$$

Combining equation 82, Theorem 12, and equation 83 yields

$$\Gamma_{\epsilon,2}(\mathcal{X}_\ell, \rho_\ell) \le C_\epsilon \sqrt{\kappa d_{w_\ell}} \Big(\sum_{r\ge1} s_r^{\frac{1+\epsilon}{\epsilon}} \Big)^{\frac{\epsilon}{1+\epsilon}}. \tag{84}$$

Since $\frac{1+\epsilon}{\epsilon} \ge 2$ and $0 \le s_r \le 1$,

$$\sum_{r\ge1} s_r^{\frac{1+\epsilon}{\epsilon}} \le \sum_{r\ge1} s_r^2 = d_{w_\ell}. \tag{85}$$

Combining equation 84 and equation 85 proves equation 72. □

B.3 CONSTRUCTIVE FARTHEST-POINT CHAIN

Lemma 14 (Phase-wise farthest-point bounds)**.** *In phase ℓ, the outputs of Algorithm 2 satisfy*

$$G_\ell \le C_\epsilon \sqrt{J_\ell}\sqrt{\kappa}\, d_{w_\ell}^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}}, \tag{86}$$

$$R_\ell \le 2C_\epsilon 2^{-\frac{J_\ell \epsilon}{1+\epsilon}} \sqrt{\kappa}\, d_{w_\ell}^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}}. \tag{87}$$

Consequently,

$$BR_\ell \le B2^{-\ell-2}. \tag{88}$$

Proof. **Proof outline.** The proof establishes a generic farthest-first approximation, uses generic chaining to control each net level, and then bounds the realized chain cost G_ℓ and terminal bias R_ℓ. The phase-ℓ quantities J_ℓ, G_ℓ, and R_ℓ denote the chain depth, weighted chain cost, and terminal covering radius, respectively. The quantities d_{w_ℓ} and κ denote the design effective dimension and the allowed certificate factor, respectively.

Farthest-first approximation. The traversal is the standard farthest-first construction of Gonzalez (1985). Consider a finite pseudometric space $(\mathcal{X}, \rho)$, or a compact pseudometric space for which the relevant suprema are attained and farthest points exist. For an integer $K \ge 1$, with $K \le |\mathcal{X}|$ when $\mathcal{X}$ is finite, define the optimal K-center radius

$$r_K^\star := \inf_{\substack{\mathcal{S}\subseteq\mathcal{X} \\ 1\le|\mathcal{S}|\le K}} \sup_{a\in\mathcal{X}} \min_{u\in\mathcal{S}} \rho(a,u). \tag{89}$$

Let $z_1, \dots, z_K$ be the first K distinct centers chosen without repetition by farthest-first traversal, and let

$$r_K := \sup_{a\in\mathcal{X}} \min_{1\le i\le K} \rho(a, z_i) \tag{90}$$

be their covering radius. The farthest-first covering radius satisfies

$$r_K \le 2r_K^\star. \tag{91}$$

Suppose instead that $r_K > 2r_K^\star$. Choose r satisfying $r_K^\star < r < \frac{r_K}{2}$ and a set of at most K centers whose radius-r balls cover $\mathcal{X}$. Choose $z_{K+1} \in \mathcal{X}$ whose distance from $\{z_1, \dots, z_K\}$ is greater than $2r$. Such a point exists by the definition of r_K, and it can be taken as an attained maximizer when the supremum is attained. For $1 \le i < j \le K$, the farthest-first rule implies

$$\rho(z_i, z_j) \ge \min_{1\le h<j} \rho(z_j, z_h) \ge r_K > 2r, \tag{92}$$

because the covering radii decrease as more centers are added. The choice of z_{K+1} also gives $\rho(z_i, z_{K+1}) > 2r$ for every $i \le K$. Thus, the $K+1$ points $z_1, \dots, z_{K+1}$ are pairwise farther than $2r$. By the triangle inequality, a ball of radius r cannot contain two of these points. Therefore, K radius-r balls cannot cover all $K+1$ points, contradicting the chosen K-center cover. This proves equation 91.

Covering radii from generic chaining. The generic chaining functional yields level-wise covering-radius bounds. For $a \in \mathcal{X}_\ell$, let $\pi_{\ell,j}(a) \in \arg\min_{u\in\mathcal{N}_{\ell,j}} \rho_\ell(a,u)$ be a nearest point in the level-j FPS net returned in phase ℓ by Algorithm 2. The infimum in Definition 11 need not be attained. However, for every $\alpha > 0$, there is an admissible sequence whose value in equation 68 is at most $\Gamma_{\epsilon,2}(\mathcal{X}_\ell, \rho_\ell) + \alpha$. For this sequence and every $a \in \mathcal{X}_\ell$,

$$\Big(\sum_{i\ge 0} 2^{\frac{2i\epsilon}{1+\epsilon}} \operatorname{diam}_{\rho_\ell}\big(\mathcal{A}_i(a)\big)^2\Big)^{\frac{1}{2}} \le \Gamma_{\epsilon,2}(\mathcal{X}_\ell, \rho_\ell) + \alpha. \tag{93}$$

For each $a \in \mathcal{X}_\ell$, let $\mathcal{A}_i(a)$ denote the cell of the partition $\mathcal{A}_i$ containing a. After squaring, every summand is nonnegative, so retaining only the level-j term gives

$$2^{\frac{2j\epsilon}{1+\epsilon}} \operatorname{diam}_{\rho_\ell}\big(\mathcal{A}_j(a)\big)^2 \le (\Gamma_{\epsilon,2}(\mathcal{X}_\ell, \rho_\ell) + \alpha)^2. \tag{94}$$

Dividing by $2^{\frac{2j\epsilon}{1+\epsilon}}$ and then taking square roots gives

$$\operatorname{diam}_{\rho_\ell}\big(\mathcal{A}_j(a)\big) \le 2^{-\frac{j\epsilon}{1+\epsilon}}(\Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell)+\alpha). \tag{95}$$

For any nonempty cell $A \in \mathcal{A}_j$ and any $a \in A$, we have $A = \mathcal{A}_j(a)$. Therefore,

$$\operatorname{diam}_{\rho_\ell}(A) \le 2^{-\frac{j\epsilon}{1+\epsilon}}(\Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell)+\alpha). \tag{96}$$

Choose one representative from each nonempty cell of $\mathcal{A}_j$. Every point in a cell is within its cell diameter of the representative, and equation 67 gives at most 2^{2^j} representatives. Hence, the optimal 2^{2^j}-center radius is at most $2^{-\frac{j\epsilon}{1+\epsilon}}(\Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell)+\alpha)$. Algorithm 2 uses the first 2^{2^j} farthest-first centers unless the covering radius becomes zero earlier. If the level-j net reaches its size budget, applying equation 91 and then letting $\alpha \downarrow 0$ gives

$$\sup_{a\in\mathcal{X}_\ell} \rho_\ell(a,\pi_{\ell,j}(a)) \le 2^{1-\frac{j\epsilon}{1+\epsilon}}\Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell). \tag{97}$$

If additions stop before the size budget is reached, the covering radius is zero. Then $\rho_\ell(a,\pi_{\ell,j}(a)) = 0$ for every $a \in \mathcal{X}_\ell$, so equation 97 still holds.

Realized chain bounds. The level-wise radius bounds control the realized chain cost and terminal radius. For every $a \in \mathcal{X}_\ell$ and $j \ge 1$, the triangle inequality gives

$$\begin{aligned}\rho_\ell\big(\pi_{\ell,j}(a),\pi_{\ell,j-1}(a)\big) &\le \rho_\ell\big(\pi_{\ell,j}(a),a\big) + \rho_\ell\big(a,\pi_{\ell,j-1}(a)\big)\\ &\le \sup_{b\in\mathcal{X}_\ell}\rho_\ell\big(b,\pi_{\ell,j}(b)\big) + \sup_{b\in\mathcal{X}_\ell}\rho_\ell\big(b,\pi_{\ell,j-1}(b)\big).\end{aligned} \tag{98}$$

For $j \ge 2$, substituting equation 97 into equation 98 yields

$$2^{\frac{j\epsilon}{1+\epsilon}}\rho_\ell\big(\pi_{\ell,j}(a),\pi_{\ell,j-1}(a)\big) \le 2(1+2^{\frac{\epsilon}{1+\epsilon}})\Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell). \tag{99}$$

For $j = 1$, the level-zero net has one point, so

$$\sup_{a\in\mathcal{X}_\ell}\rho_\ell\big(a,\pi_{\ell,0}(a)\big) \le \operatorname{diam}_{\rho_\ell}(\mathcal{X}_\ell). \tag{100}$$

Every admissible sequence satisfies $\mathcal{A}_0 = \{\mathcal{X}_\ell\}$, so $\mathcal{A}_0(a) = \mathcal{X}_\ell$ and its level-zero summand in equation 68 is $\operatorname{diam}_{\rho_\ell}(\mathcal{X}_\ell)^2$. Since all other summands are nonnegative, taking the supremum over a and then the infimum over admissible sequences gives

$$\operatorname{diam}_{\rho_\ell}(\mathcal{X}_\ell) \le \Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell). \tag{101}$$

Combining $\operatorname{diam}_{\rho_\ell}(\mathcal{X}_\ell) \le \Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell)$ with equation 97 at $j = 1$ verifies equation 99 at the first level. Recall from Algorithm 2 that

$$\begin{aligned}G_\ell &= \sup_{a\in\mathcal{X}_\ell}\Big[\sum_{j=1}^{J_\ell} 2^{\frac{2j\epsilon}{1+\epsilon}}\rho_\ell\big(\pi_{\ell,j}(a),\pi_{\ell,j-1}(a)\big)^2\Big]^{\frac12},\\ R_\ell &= \sup_{a\in\mathcal{X}_\ell}\rho_\ell\big(a,\pi_{\ell,J_\ell}(a)\big).\end{aligned} \tag{102}$$

Squaring equation 99, summing over $j = 1,\dots,J_\ell$, and using the definition of G_ℓ give

$$G_\ell \le 2(1+2^{\frac{\epsilon}{1+\epsilon}})\sqrt{J_\ell}\,\Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell). \tag{103}$$

Applying equation 97 with $j = J_\ell$ and using the definition of R_ℓ give

$$R_\ell \le 2^{1-\frac{J_\ell\epsilon}{1+\epsilon}}\Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell). \tag{104}$$

The estimate in equation 72, applied in phase ℓ with certificate factor κ, bounds $\Gamma_{\epsilon,2}(\mathcal{X}_\ell,\rho_\ell)$. Substituting equation 72 into equation 103 and equation 104, and enlarging C_ϵ to absorb $2(1+2^{\frac{\epsilon}{1+\epsilon}})$, proves equation 86 and equation 87.

Terminal bias. It remains to verify that the depth schedule makes the terminal bias smaller than the phase budget. By the definition of J_ℓ,

$$J_\ell \geq \frac{1+\epsilon}{\epsilon} \log_2 \left(2^{\ell+3} \cdot C_\epsilon \sqrt{\kappa} \, (1 \vee d_{w_\ell})^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}} \right). \tag{105}$$

Multiplying equation 105 by $\frac{\epsilon}{1+\epsilon}$ and exponentiating with base 2 give

$$2^{-\frac{J_\ell \epsilon}{1+\epsilon}} \leq \frac{2^{-\ell-3}}{C_\epsilon \sqrt{\kappa} \, (1 \vee d_{w_\ell})^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}}}. \tag{106}$$

Substituting equation 106 into equation 87 gives

$$\begin{aligned} BR_\ell &\leq 2C_\epsilon B 2^{-\frac{J_\ell \epsilon}{1+\epsilon}} \sqrt{\kappa} \, d_{w_\ell}^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}} \\ &\leq B 2^{-\ell-2} \cdot \frac{d_{w_\ell}^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}}}{(1 \vee d_{w_\ell})^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}}} \leq B 2^{-\ell-2}. \end{aligned} \tag{107}$$

This proves equation 88. □

B.4 Robust chained estimation

Proof of Lemma 5. **Proof outline.** The proof conditions on the phase history, establishes concentration for each clipped edge, combines the anchor and multiscale increments, and then controls the terminal approximation and regularization bias.

Conditioning and population target. Condition on all observations before phase ℓ. The sampling distribution w_ℓ, its induced variance pseudometric ρ_ℓ, the level-j nets $(\mathcal{N}_{\ell,j})_{j=0}^{J_\ell}$, the nearest-point projections $(\pi_{\ell,j})_{j=0}^{J_\ell}$, and the phase length are then fixed. The estimator $\widehat{f}_\ell$ is subsequently computed from the new observations. The operators A_{w_ℓ} and $V_\ell = A_{w_\ell} + \lambda_T I$ are the covariance and regularized covariance operators of w_ℓ, respectively. The parameters $\lambda_T = 8T^{-1}$, $\bar{v}$, ϵ, and δ_ℓ are the regularization level, observation moment bound, moment parameter, and phase failure probability. Throughout this proof, C_ϵ denotes a positive constant depending only on ϵ and may be enlarged. For n observations, define the confidence radius used by Algorithms 1 and 4 as

$$\begin{aligned} \beta_\ell(n) := C_\epsilon \bar{v}^{\frac{1}{1+\epsilon}} n^{-\frac{\epsilon}{1+\epsilon}} \left[1 + \sqrt{\kappa d_{w_\ell}} \Big[\log\Big(\frac{4}{\delta_\ell}\Big) \Big]^{\frac{\epsilon}{1+\epsilon}} + G_\ell \Big(\sqrt{J_\ell} + \Big[\log\Big(\frac{4J_\ell}{\delta_\ell}\Big) \Big]^{\frac{\epsilon}{1+\epsilon}} \Big) \right] \\ + B\sqrt{\lambda_T \kappa d_{w_\ell}} + BR_\ell. \end{aligned} \tag{108}$$

For $a \in \mathcal{X}_\ell$, define the regularized population target by

$$f_{\ell,\lambda_T}(a) := \langle \phi(a), V_\ell^{-1} A_{w_\ell} f \rangle_{\mathcal{H}_k}. \tag{109}$$

Single-edge concentration. For a fixed $j \in \{1, \ldots, J_\ell\}$ and an edge $(a, b) \in E_j$, where E_j is the level-j edge set in equation 19, abbreviate the scalar observation in Algorithm 3 as

$$Z_i := Z_{i,j}(a,b) = \langle \phi(a) - \phi(b), V_\ell^{-1} \phi(x_i) \rangle_{\mathcal{H}_k} y_i, \qquad i = 1, \ldots, n,$$

where $x_i \sim w_\ell$ is the i-th query point in phase ℓ and $y_i = f(x_i) + \eta_i$ is its observation. Conditional on x_i and the available history, the coefficient $\langle \phi(a) - \phi(b), V_\ell^{-1} \phi(x_i) \rangle_{\mathcal{H}_k}$ is fixed and the noise has conditional mean zero. Taking conditional expectation and using the conditional mean-zero property reduces $\mathbb{E}[Z_i]$ to $\mathbb{E}_{x \sim w_\ell}[\langle \phi(a) - \phi(b), V_\ell^{-1} \phi(x) \rangle_{\mathcal{H}_k} f(x)]$. By the reproducing property, $f(x) = \langle \phi(x), f \rangle_{\mathcal{H}_k}$, so equation 14 gives $\mathbb{E}_{x \sim w_\ell}[\phi(x) f(x)] = A_{w_\ell} f$. Linearity of the expectation and the inner product therefore gives

$$\mathbb{E}[Z_i] = f_{\ell,\lambda_T}(a) - f_{\ell,\lambda_T}(b). \tag{110}$$

By Assumption 3 and the monotonicity of conditional moments,

$$\begin{aligned} \mathbb{E}[|Z_i|^{1+\epsilon}] &\leq \bar{v} \, \mathbb{E}_{x \sim w_\ell} \big| \langle \phi(a) - \phi(b), V_\ell^{-1} \phi(x) \rangle_{\mathcal{H}_k} \big|^{1+\epsilon} \\ &\leq \bar{v} \left(\mathbb{E}_{x \sim w_\ell} \langle \phi(a) - \phi(b), V_\ell^{-1} \phi(x) \rangle_{\mathcal{H}_k}^2 \right)^{\frac{1+\epsilon}{2}} \\ &= \bar{v} \, \rho_\ell(a,b)^{1+\epsilon}. \end{aligned} \tag{111}$$

Fix a logarithmic failure budget $u > 0$. If $\rho_\ell(a, b) = 0$, equation 111 gives $Z_i = 0$ almost surely, so the desired bound is immediate. Otherwise, define the clipping threshold at sample size n by

$$\tau_n := \bar{v}^{\frac{1}{1+\epsilon}} \rho_\ell(a, b) \Big(\frac{n}{u}\Big)^{\frac{1}{1+\epsilon}}. \tag{112}$$

Let $\mathbf{1}\{\cdot\}$ denote the indicator. Because $|Z_i|\mathbf{1}\{|Z_i| > \tau_n\} \le |Z_i|^{1+\epsilon}\tau_n^{-\epsilon}$, the clipping bias satisfies

$$\begin{aligned}\Big|\mathbb{E}\big[\operatorname{sign}(Z_i)\min\{|Z_i|, \tau_n\}\big] - \mathbb{E}[Z_i]\Big| &\le \mathbb{E}\big[|Z_i|\mathbf{1}\{|Z_i| > \tau_n\}\big] \\ &\le \bar{v}\,\rho_\ell(a, b)^{1+\epsilon}\tau_n^{-\epsilon} \\ &= \bar{v}^{\frac{1}{1+\epsilon}} \rho_\ell(a, b)\Big(\frac{u}{n}\Big)^{\frac{\epsilon}{1+\epsilon}}. \end{aligned} \tag{113}$$

The clipped observation also satisfies

$$\mathbb{E}\big[\min\{|Z_i|, \tau_n\}^2\big] \le \bar{v}\,\rho_\ell(a, b)^{1+\epsilon}\tau_n^{1-\epsilon}, \tag{114}$$

because $\min\{|Z_i|, \tau_n\}^2 \le \tau_n^{1-\epsilon}|Z_i|^{1+\epsilon}$. Its centered version has absolute value at most $2\tau_n$. Applying the conditional Bernstein inequality (Freedman, 1975) therefore shows that, with probability at least $1 - 2e^{-u}$,

$$\begin{aligned}&\Big|\frac{1}{n}\sum_{i=1}^{n}\operatorname{sign}(Z_i)\min\{|Z_i|, \tau_n\} - \mathbb{E}\big[\operatorname{sign}(Z_i)\min\{|Z_i|, \tau_n\}\big]\Big| \\ &\le \sqrt{\frac{2\bar{v}\,\rho_\ell(a, b)^{1+\epsilon}\tau_n^{1-\epsilon}u}{n}} + \frac{4\tau_n u}{3n} \\ &= \Big(\sqrt{2} + \frac{4}{3}\Big)\bar{v}^{\frac{1}{1+\epsilon}}\rho_\ell(a, b)\Big(\frac{u}{n}\Big)^{\frac{\epsilon}{1+\epsilon}}. \end{aligned} \tag{115}$$

Combining equation 113 and equation 115, and using $1 + \sqrt{2} + \frac{4}{3} < 4$, yields

$$\Big|\frac{1}{n}\sum_{i=1}^{n}\operatorname{sign}(Z_i)\min\{|Z_i|, \tau_n\} - \mathbb{E}[Z_i]\Big| \le 4\bar{v}^{\frac{1}{1+\epsilon}}\rho_\ell(a, b)\Big(\frac{u}{n}\Big)^{\frac{\epsilon}{1+\epsilon}} \tag{116}$$

with probability at least $1 - 2e^{-u}$.

Anchor estimate. Let a_0 be the unique point in $\mathcal{N}_{\ell,0}$, and write $Z_i := \langle\phi(a_0), V_\ell^{-1}\phi(x_i)\rangle_{\mathcal{H}_k} y_i$ for its weighted observations. The expectation and moment calculations in equation 110 and equation 111, with $\phi(a_0)$ in place of the feature difference, give

$$\mathbb{E}[Z_i] = f_{\ell,\lambda_T}(a_0) \tag{117}$$

and

$$\mathbb{E}[|Z_i|^{1+\epsilon}] \le \bar{v}\|A_{w_\ell}^{\frac{1}{2}}V_\ell^{-1}\phi(a_0)\|_{\mathcal{H}_k}^{1+\epsilon} \le \bar{v}(\kappa d_{w_\ell})^{\frac{1+\epsilon}{2}}. \tag{118}$$

The last inequality follows from $V_\ell^{-1}A_{w_\ell}V_\ell^{-1} \preceq V_\ell^{-1}$ and the certificate $h_{w_\ell}(a_0) \le \kappa d_{w_\ell}$. Apply the clipping-bias and Bernstein calculations in equation 113 and equation 115 to these observations, using the anchor threshold in equation 21 with $w = w_\ell$ and $u_0 = \log(4/\delta_\ell)$. If $\|A_{w_\ell}^{1/2}V_\ell^{-1}\phi(a_0)\|_{\mathcal{H}_k} = 0$, the moment bound gives $Z_i = 0$ almost surely, so the same conclusion holds. Thus, with probability at least $1 - \delta_\ell/2$,

$$\big|\widehat{\Delta}_{\ell,0}(a_0) - f_{\ell,\lambda_T}(a_0)\big| \le 4\bar{v}^{\frac{1}{1+\epsilon}}n^{-\frac{\epsilon}{1+\epsilon}}\sqrt{\kappa d_{w_\ell}}\Big[\log\Big(\frac{4}{\delta_\ell}\Big)\Big]^{\frac{\epsilon}{1+\epsilon}}. \tag{119}$$

Uniform multiscale increments. The multiscale increments are controlled simultaneously. Even for an uncountable $\mathcal{X}_\ell$, there are only finitely many distinct edges at each level, so the following union bounds do not require a union over all active points. For $j \ge 1$, the edge definition in equation 19 and the net cardinalities give

$$|E_j| \le |\mathcal{N}_{\ell,j}||\mathcal{N}_{\ell,j-1}| \le 2^{2^j + 2^{j-1}}. \tag{120}$$

The concentration inequality in equation 116 bounds the failure probability for each edge by $2e^{-u_j}$. Applying a union bound over E_j and substituting equation 120 and $u_j = 8 \cdot 2^j + \log(\frac{4J_\ell}{\delta_\ell})$ from equation 20 give

$$\begin{aligned} 2|E_j|e^{-u_j} &\le 2\exp\Big(\frac{3}{2}\cdot 2^j \log 2 - 8\cdot 2^j - \log\Big(\frac{4J_\ell}{\delta_\ell}\Big)\Big) \\ &= \frac{\delta_\ell}{2J_\ell}\exp\Big(-\Big(8-\frac{3}{2}\log 2\Big)2^j\Big) \\ &\le \frac{\delta_\ell}{2J_\ell}. \end{aligned} \tag{121}$$

The bound $\frac{\delta_\ell}{2J_\ell}\exp(-(8-\frac{3}{2}\log 2)2^j) \le \frac{\delta_\ell}{2J_\ell}$ holds because $8-\frac{3}{2}\log 2 > 0$. Summing over $j = 1, \dots, J_\ell$ shows that all realized increments are controlled with failure probability at most $\frac{\delta_\ell}{2}$. For a fixed $a \in \mathcal{X}_\ell$, subadditivity of $x \mapsto x^{\frac{\epsilon}{1+\epsilon}}$ and equation 20 give

$$u_j^{\frac{\epsilon}{1+\epsilon}} \le 8^{\frac{\epsilon}{1+\epsilon}} 2^{\frac{j\epsilon}{1+\epsilon}} + \Big[\log\Big(\frac{4J_\ell}{\delta_\ell}\Big)\Big]^{\frac{\epsilon}{1+\epsilon}}. \tag{122}$$

Cauchy-Schwarz inequality and the geometric-series identity $\sum_{j\ge 1} 2^{-\frac{2j\epsilon}{1+\epsilon}} = (2^{\frac{2\epsilon}{1+\epsilon}}-1)^{-1}$ imply

$$\begin{aligned} &\sum_{j=1}^{J_\ell} u_j^{\frac{\epsilon}{1+\epsilon}} \rho_\ell\big(\pi_{\ell,j}(a), \pi_{\ell,j-1}(a)\big) \\ &\le 8^{\frac{\epsilon}{1+\epsilon}}\sqrt{J_\ell}\Big[\sum_{j=1}^{J_\ell} 2^{\frac{2j\epsilon}{1+\epsilon}} \rho_\ell\big(\pi_{\ell,j}(a), \pi_{\ell,j-1}(a)\big)^2\Big]^{\frac{1}{2}} \\ &\quad + \Big[\log\Big(\frac{4J_\ell}{\delta_\ell}\Big)\Big]^{\frac{\epsilon}{1+\epsilon}} \Big(\sum_{j=1}^{J_\ell} 2^{-\frac{2j\epsilon}{1+\epsilon}}\Big)^{\frac{1}{2}} \Big[\sum_{j=1}^{J_\ell} 2^{\frac{2j\epsilon}{1+\epsilon}} \rho_\ell\big(\pi_{\ell,j}(a), \pi_{\ell,j-1}(a)\big)^2\Big]^{\frac{1}{2}} \\ &\le C_\epsilon G_\ell \left\{\sqrt{J_\ell} + \Big[\log\Big(\frac{4J_\ell}{\delta_\ell}\Big)\Big]^{\frac{\epsilon}{1+\epsilon}}\right\}. \end{aligned} \tag{123}$$

The definition of G_ℓ in Algorithm 2 gives the final inequality in equation 123. Combining equation 116 and equation 123 bounds the accumulated increment error uniformly over $a \in \mathcal{X}_\ell$.

Error decomposition and deterministic bias. Recall that $\widehat{\Delta}_{\ell,j}(a,b)$ is the clipped estimate of the population increment $f_{\ell,\lambda_T}(a) - f_{\ell,\lambda_T}(b)$ on a level-j edge. The stochastic bounds in equation 119 and equation 123 form the first two terms in the decomposition

$$\begin{aligned} \widehat{f}_\ell(a) - f(a) &= \Big[\widehat{\Delta}_{\ell,0}(a_0) - f_{\ell,\lambda_T}(a_0)\Big] \\ &\quad + \sum_{j=1}^{J_\ell}\Big[\widehat{\Delta}_{\ell,j}\big(\pi_{\ell,j}(a), \pi_{\ell,j-1}(a)\big) - f_{\ell,\lambda_T}\big(\pi_{\ell,j}(a)\big) + f_{\ell,\lambda_T}\big(\pi_{\ell,j-1}(a)\big)\Big] \\ &\quad + \Big[f_{\ell,\lambda_T}\big(\pi_{\ell,J_\ell}(a)\big) - f_{\ell,\lambda_T}(a)\Big] + \big[f_{\ell,\lambda_T}(a) - f(a)\big]. \end{aligned} \tag{124}$$

It remains to control the last two terms in equation 124. By equation 110, the population anchor and increments telescope to $f_{\ell,\lambda_T}(\pi_{\ell,J_\ell}(a))$. Since $V_\ell = A_{w_\ell} + \lambda_T I$, Cauchy-Schwarz inequality and $V_\ell \succeq \lambda_T I$ give

$$\begin{aligned} |f(a) - f_{\ell,\lambda_T}(a)| &= \lambda_T\big|\langle V_\ell^{-\frac{1}{2}}\phi(a), V_\ell^{-\frac{1}{2}} f\rangle_{\mathcal{H}_k}\big| \\ &\le B\sqrt{\lambda_T}\|V_\ell^{-\frac{1}{2}}\phi(a)\|_{\mathcal{H}_k} \le B\sqrt{\lambda_T \kappa d_{w_\ell}}. \end{aligned} \tag{125}$$

For any $a, b \in \mathcal{X}_\ell$, the reproducing property and equation 14 give

$$\begin{aligned} |f_{\ell,\lambda_T}(a) - f_{\ell,\lambda_T}(b)| &= \Big|\mathbb{E}_{x\sim w_\ell}\Big[f(x)\langle\phi(a)-\phi(b), V_\ell^{-1}\phi(x)\rangle_{\mathcal{H}_k}\Big]\Big| \\ &\le \Big(\mathbb{E}_{x\sim w_\ell}|f(x)|^2\Big)^{\frac{1}{2}}\rho_\ell(a,b) \le B\rho_\ell(a,b). \end{aligned} \tag{126}$$

Taking $b = \pi_{\ell,J_\ell}(a)$ in equation 126 bounds the terminal error by BR_ℓ.

Combining the bounds. It remains to combine the four terms in equation 124. A union bound shows that the anchor and multiscale events hold simultaneously with probability at least $1-\delta_\ell$. The regularization bias is at most $B\sqrt{\lambda_T \kappa d_{w_\ell}}$ by equation 125, and the terminal error is at most BR_ℓ by equation 126. On the intersection of the anchor and multiscale events, the triangle inequality yields

$$\sup_{a\in\mathcal{X}_\ell} |\widehat{f}_\ell(a) - f(a)| \le \beta_\ell(n) \tag{127}$$

after enlarging the constant in equation 108. The extra 1 in equation 108 is nonnegative and ensures a uniform lower bound on the phase length. □

B.5 Elimination and regret

Proof of Theorem 6. **Proof outline.** We establish uniform confidence on the compact active set, show that elimination preserves an optimizer and compactness, and then sum the phase-wise regret bounds. To reuse the calculation for the practical version, we retain a general certificate factor $\kappa \ge 1$ in the schedules and bounds. For Algorithm 1, the D-optimal certificate gives $\kappa = 1$.

Well-defined phases. The initial set $\mathcal{X}_1 = \mathcal{X}$ is nonempty and compact. On any such active set, Lemma 9 supplies a maximizing distribution. Continuity of ρ_ℓ ensures that the farthest points in Algorithm 2 exist. The finite nets and finite chain depth imply that $\widehat{f}_\ell$ takes only finitely many values on $\mathcal{X}_\ell$, so its supremum is attained. The set inside the closure in equation 8 is therefore nonempty. Its closure is a compact subset of $\mathcal{X}_\ell$, so induction makes every initiated phase well-defined under the stated oracle assumptions, even outside the confidence event.

Uniform confidence event. For every completed phase, let

$$\mathcal{G}_\ell := \Big\{ \sup_{a\in\mathcal{X}_\ell} |\widehat{f}_\ell(a) - f(a)| \le \beta_\ell \Big\}. \tag{128}$$

The phase length is fixed before the phase observations, and Lemma 5 bounds the conditional failure probability by δ_ℓ. Since $\sum_{\ell\ge1} \delta_\ell = \delta$, a union bound gives a common event of probability at least $1-\delta$ on which all completed phases satisfy $\mathcal{G}_\ell$.

Elimination invariant. Work on the common confidence event and suppose $x^\star \in \mathcal{X}_\ell$. For every $u \in \mathcal{X}_\ell$,

$$\widehat{f}_\ell(x^\star) \ge f(x^\star) - \beta_\ell \ge f(u) - \beta_\ell \ge \widehat{f}_\ell(u) - 2\beta_\ell. \tag{129}$$

Taking the supremum over u shows that $x^\star$ belongs to the set in equation 8 before taking its closure. For any point a in that set, the retention rule and the confidence bounds give

$$\begin{aligned} f(a) &\ge \widehat{f}_\ell(a) - \beta_\ell \\ &\ge \sup_{u\in\mathcal{X}_\ell} \widehat{f}_\ell(u) - 3\beta_\ell \ge \widehat{f}_\ell(x^\star) - 3\beta_\ell \ge f(x^\star) - 4\beta_\ell. \end{aligned} \tag{130}$$

Because f is continuous, the same inequality holds at every limit point of the retained set. Thus, taking the closure introduces no larger gaps, and

$$\sup_{a\in\mathcal{X}_{\ell+1}} \big(f(x^\star) - f(a)\big) \le 4\beta_\ell \le B2^{2-\ell}. \tag{131}$$

Since $x^\star \in \mathcal{X}_1$, induction shows that it is never eliminated.

Phase lengths and phase count. Fix a phase ℓ that reaches the FPS construction. For $n \ge 1$, write the candidate phase confidence radius as

$$\beta_\ell(n) := C_\epsilon \bar{v}^{\frac{1}{1+\epsilon}} n^{-\frac{\epsilon}{1+\epsilon}} \left[1 + \sqrt{\kappa d_{w_\ell}} \Big[\log\Big(\frac{4}{\delta_\ell}\Big) \Big]^{\frac{\epsilon}{1+\epsilon}} + G_\ell \Big(\sqrt{J_\ell} + \Big[\log\Big(\frac{4J_\ell}{\delta_\ell}\Big) \Big]^{\frac{\epsilon}{1+\epsilon}} \Big) \right] + B\sqrt{\lambda_T \kappa d_{w_\ell}} + BR_\ell. \tag{132}$$

The phase length and total confidence radius are

$$\begin{aligned} n_\ell &:= \min\Big\{ n \in \mathbb{N} : \beta_\ell(n) \le B2^{-\ell} \Big\}, \\ \beta_\ell &:= \beta_\ell(n_\ell), \end{aligned} \tag{133}$$

as used by Algorithm 1. Since the phase reaches the FPS construction, the bias check in Algorithm 1 gives

$$B\sqrt{\lambda_T \kappa d_{w_\ell}} \le B2^{-\ell-2}. \tag{134}$$

Lemma 14 gives $BR_\ell \le B2^{-\ell-2}$. Hence,

$$B\sqrt{\lambda_T \kappa d_{w_\ell}} + BR_\ell \le B2^{-\ell-1}. \tag{135}$$

The stochastic term in $\beta_\ell(n)$ decreases to zero as $n \to \infty$, so the minimum defining n_ℓ exists and $\beta_\ell \le B2^{-\ell}$. The bound on G_ℓ in equation 86 gives

$$\begin{aligned} &1 + \sqrt{\kappa d_{w_\ell}}\Big(\log\frac{4}{\delta_\ell}\Big)^{\frac{\epsilon}{1+\epsilon}} + G_\ell\left[\sqrt{J_\ell} + \Big(\log\frac{4J_\ell}{\delta_\ell}\Big)^{\frac{\epsilon}{1+\epsilon}}\right] \\ &\le C_\epsilon\sqrt{\kappa}\,(1 \vee d_{w_\ell})^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}}\left[1 + \Big(\log\frac{4}{\delta_\ell}\Big)^{\frac{\epsilon}{1+\epsilon}} + J_\ell + \sqrt{J_\ell}\Big(\log\frac{4J_\ell}{\delta_\ell}\Big)^{\frac{\epsilon}{1+\epsilon}}\right]. \end{aligned} \tag{136}$$

The deterministic terms in $\beta_\ell(n)$ sum to at most $B2^{-\ell-1}$. It is therefore sufficient to make the stochastic term at most $B2^{-\ell-2}$, leaving a margin of $B2^{-\ell-2}$ that will accommodate the cover error in the practical corollary. Substituting equation 136 into the definition of $\beta_\ell(n)$ and solving the stochastic-error requirement for n gives equation 137. In equation 137, $\widetilde{\mathcal{O}}_\epsilon$ suppresses constants depending only on ϵ, polynomial factors in ℓ, and logarithmic factors in δ^{-1}, κ, $1 \vee d_{w_\ell}$, and B.

$$n_\ell = \widetilde{\mathcal{O}}_\epsilon\Big(\bar{v}^{\frac{1}{\epsilon}}\kappa^{\frac{1+\epsilon}{2\epsilon}}(1 \vee d_{w_\ell})^{\frac{1+3\epsilon}{2\epsilon}}(B2^{-\ell})^{-\frac{1+\epsilon}{\epsilon}}\Big). \tag{137}$$

The confidence radius β_ℓ contains the nonnegative unit term $C_\epsilon \bar{v}^{\frac{1}{1+\epsilon}} n_\ell^{-\frac{\epsilon}{1+\epsilon}}$. Since $\beta_\ell \le B2^{-\ell}$, it follows that

$$C_\epsilon \bar{v}^{\frac{1}{1+\epsilon}} n_\ell^{-\frac{\epsilon}{1+\epsilon}} \le B2^{-\ell}. \tag{138}$$

Solving $C_\epsilon \bar{v}^{\frac{1}{1+\epsilon}} n_\ell^{-\frac{\epsilon}{1+\epsilon}} \le B2^{-\ell}$ for n_ℓ gives

$$n_\ell \ge C_\epsilon^{\frac{1+\epsilon}{\epsilon}} \bar{v}^{\frac{1}{\epsilon}}(B2^{-\ell})^{-\frac{1+\epsilon}{\epsilon}}. \tag{139}$$

Since Assumption 3 implies $\bar{v}^{\frac{1}{\epsilon}} B^{-\frac{1+\epsilon}{\epsilon}} \ge 2$, equation 139 yields $n_\ell \ge 2^{\frac{(1+\epsilon)\ell}{\epsilon}}$ after taking $C_\epsilon \ge 1$. Let L_T denote the final phase initiated by time T within this proof. If $L_T \ge 2$, Algorithm 1 can initiate phase L_T only after phase $L_T - 1$ has queried n_{L_T-1} points and advanced the phase index. This remains true when phase L_T stops at the bias check. The horizon constraint and the phase-length lower bound following equation 139 therefore give

$$T \ge n_{L_T-1} \ge 2^{\frac{(1+\epsilon)(L_T-1)}{\epsilon}}. \tag{140}$$

Taking base-two logarithms yields

$$L_T \le 1 + \frac{\epsilon}{1+\epsilon}\log_2 T \le 1 + \log_2 T. \tag{141}$$

The case $L_T = 1$ is immediate, so

$$L_T = \mathcal{O}_\epsilon(\log(2+T)). \tag{142}$$

Regret aggregation. Recall that $R_T = \sum_{t=1}^{T}(f(x^\star) - f(x_t))$ denotes cumulative pseudo-regret. Let $\mathcal{I}_T$ be the set of phases that pass the regularization-bias check in Algorithm 1, and let $m_\ell \le n_\ell$ be the number of query points evaluated in phase $\ell \in \mathcal{I}_T$. Assumption 1 gives instantaneous regret at most $2B$ in the first phase. For $\ell \ge 2$, equation 131 and $\beta_{\ell-1} \le B2^{1-\ell}$ give instantaneous regret at most $B2^{3-\ell}$. Thus, the regret incurred in the phases in $\mathcal{I}_T$ is at most

$$8\sum_{\ell \in \mathcal{I}_T} m_\ell B2^{-\ell}. \tag{143}$$

Since $\frac{1}{1+\epsilon} + \frac{\epsilon}{1+\epsilon} = 1$ and $m_\ell \le n_\ell$,

$$m_\ell B 2^{-\ell} = m_\ell^{\frac{1}{1+\epsilon}} m_\ell^{\frac{\epsilon}{1+\epsilon}} B 2^{-\ell} \le m_\ell^{\frac{1}{1+\epsilon}} n_\ell^{\frac{\epsilon}{1+\epsilon}} B 2^{-\ell}. \tag{144}$$

Substituting equation 137 and suppressing the logarithmic factors specified in equation 136 give

$$m_\ell B 2^{-\ell} \le \widetilde{\mathcal{O}}_\epsilon\Big(m_\ell^{\frac{1}{1+\epsilon}} \bar{v}^{\frac{1}{1+\epsilon}} \sqrt{\kappa}\, (1 \vee d_{w_\ell})^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}} \Big). \tag{145}$$

Lemma 10 gives $1 \vee d_{w_\ell} \le 1 \vee \gamma_T$ in every initiated phase. We retain $1 \vee \gamma_T$ because Assumption 1 allows arbitrarily small kernel diagonal values and gives no positive lower bound on γ_T that is uniform over the allowed kernels. Hölder's inequality, $\sum_{\ell \in \mathcal{I}_T} m_\ell \le T$, and $|\mathcal{I}_T| \le L_T$ then give

$$\begin{aligned}
\sum_{\ell \in \mathcal{I}_T} m_\ell B 2^{-\ell} &\le \widetilde{\mathcal{O}}_\epsilon\left(\Big(\sum_{\ell \in \mathcal{I}_T} m_\ell \Big)^{\frac{1}{1+\epsilon}} \left[\sum_{\ell \in \mathcal{I}_T} \bar{v}^{\frac{1}{\epsilon}} \kappa^{\frac{1+\epsilon}{2\epsilon}} (1 \vee d_{w_\ell})^{\frac{1+3\epsilon}{2\epsilon}} \right]^{\frac{\epsilon}{1+\epsilon}} \right) \\
&\le \widetilde{\mathcal{O}}_\epsilon\Big(\bar{v}^{\frac{1}{1+\epsilon}} \sqrt{\kappa}\, (1 \vee \gamma_T)^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}} T^{\frac{1}{1+\epsilon}} L_T^{\frac{\epsilon}{1+\epsilon}} \Big).
\end{aligned} \tag{146}$$

If every initiated phase passes the bias check, equation 143, equation 146, and equation 142 give

$$R_T \le \widetilde{\mathcal{O}}_\epsilon\Big(\bar{v}^{\frac{1}{1+\epsilon}} \sqrt{\kappa}\, (1 \vee \gamma_T)^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}} T^{\frac{1}{1+\epsilon}} \Big). \tag{147}$$

Otherwise, only the final phase fails the check and is omitted from $\mathcal{I}_T$. Its instantaneous regret is at most $B 2^{3-L_T}$, using $2B$ when $L_T = 1$ and equation 131 when $L_T \ge 2$. Failure of the regularization-bias check in Algorithm 1 means

$$B\sqrt{\lambda_T \kappa d_{w_{L_T}}} > B 2^{-L_T-2}, \qquad \text{or equivalently} \qquad B 2^{-L_T} < 4B\sqrt{\lambda_T \kappa d_{w_{L_T}}}. \tag{148}$$

Substituting $\lambda_T = \frac{8}{T}$ and multiplying by $8T$ give

$$\begin{aligned}
BT \cdot 2^{3-L_T} &< 32BT\sqrt{\frac{8\kappa d_{w_{L_T}}}{T}} \\
&= 32\sqrt{8}\, B\sqrt{\kappa d_{w_{L_T}} T} \\
&\le 32\sqrt{8}\, \bar{v}^{\frac{1}{1+\epsilon}} \sqrt{\kappa}\, (1 \vee \gamma_T)^{\frac{1}{2}+\frac{\epsilon}{1+\epsilon}} T^{\frac{1}{1+\epsilon}},
\end{aligned} \tag{149}$$

The final inequality in equation 149 uses $\bar{v}^{\frac{1}{1+\epsilon}} \ge B$, $1 + \epsilon \le 2$, and $d_{w_{L_T}} \le \gamma_T$ from Lemma 10. Adding the terminal-phase bound in equation 149 to equation 143 and equation 146 yields equation 147 in the remaining case. Finally, set $\kappa = 1$ for the D-optimal distributions in Algorithm 1 and use $\frac{1}{2} + \frac{\epsilon}{1+\epsilon} = \frac{1+3\epsilon}{2(1+\epsilon)}$ in equation 147. This proves equation 10. □

C FINITE-COVER PRACTICAL ALGORITHM AND COROLLARIES

Algorithm 4 Finite-cover HT-ChainPE (practical version)

Require: Compact domain $\mathcal{X}$, horizon T, parameters $(B, \epsilon, \bar{v}, \delta, k)$, $\kappa_{\max} > 1$, and $S \in \mathbb{N} \cup \{\infty\}$
Initialize $\mathcal{X}_1 \leftarrow \mathcal{X}$, $t \leftarrow 0$, $\ell \leftarrow 1$, and $\kappa \leftarrow \kappa_{\max}$
Use (δ_ℓ, λ_T) from Algorithm 1
while $t < T$ **do**
  $(\mathcal{C}_\ell, \{\mathcal{V}_\ell(c)\}) \leftarrow$ BUILDKERNELCOVER$(\mathcal{X}_\ell, 2^{-\ell-2})$ (Algorithm 5)
  $(w_\ell, V_\ell, d_{w_\ell}) \leftarrow$ CERTIFIEDDESIGN$(\mathcal{C}_\ell, \lambda_T, \kappa_{\max}, S)$ (Algorithm 6)
  if the distribution routine returns failure **then**
    return certificate failure
  end if
  if $B\sqrt{\lambda_T \kappa_{\max} d_{w_\ell}} > B2^{-\ell-2}$ **then**
    Sample and query the remaining $T - t$ points from w_ℓ and **return**
  end if
  Set ρ_ℓ by equation 4 and choose J_ℓ as in Algorithm 1
  $(\{\mathcal{N}_{\ell,j}, \pi_{\ell,j}\}_{j=0}^{J_\ell}, G_\ell, R_\ell) \leftarrow$ BUILDFPSCHAIN$(\mathcal{C}_\ell, \rho_\ell, J_\ell)$ (Algorithm 2)
  Use equation 108 and equation 151 to choose n_ℓ and β_ℓ
  if $n_\ell > T - t$ **then**
    Sample and query the remaining $T - t$ points from w_ℓ and **return**
  end if
  $\widehat{f}_\ell \leftarrow$ CHAINEDESTIMATE$(\mathcal{C}_\ell, w_\ell, V_\ell, n_\ell, \delta_\ell, \{\mathcal{N}_{\ell,j}, \pi_{\ell,j}\}_{j=0}^{J_\ell})$ (Algorithm 3)
  Update $\mathcal{X}_{\ell+1}$ by equation 152
  $t \leftarrow t + n_\ell$ and $\ell \leftarrow \ell + 1$
end while

This section gives the finite-cover implementation and derives its statistical guarantees from the compact-domain analysis. Compared with Algorithm 1, it computes a certified approximate D-optimal distribution and the chain on a finite kernel cover, adds the cover error to the confidence radius, and eliminates entire cells.

C.1 PRACTICAL PHASE PROCEDURE

Define the kernel pseudometric by

$$d_k(a, b) := \|\phi(a) - \phi(b)\|_{\mathcal{H}_k} = \big(k(a, a) + k(b, b) - 2k(a, b)\big)^{\frac{1}{2}}, \qquad a, b \in \mathcal{X}. \tag{150}$$

Algorithm 5 returns the cover centers $\mathcal{C}_\ell$ and their associated closed regions $\mathcal{V}_\ell(c)$. The practical algorithm uses the phase schedule and bias check in Algorithm 1, with chain depth given by equation 3. It uses the exact confidence radius in equation 108. It sets $\kappa = \kappa_{\max} > 1$ and constructs the chain on $\mathcal{C}_\ell$. It changes the phase-length rule to

$$\begin{aligned} n_\ell &:= \min\big\{n \in \mathbb{N} : \beta_\ell(n) + B2^{-\ell-2} \le B2^{-\ell}\big\}, \\ \beta_\ell &:= \beta_\ell(n_\ell) + B2^{-\ell-2}. \end{aligned} \tag{151}$$

After a completed phase, it retains cells rather than individual centers:

$$\mathcal{X}_{\ell+1} := \bigcup_{\substack{c \in \mathcal{C}_\ell: \\ \widehat{f}_\ell(c) \ge \max_{u \in \mathcal{C}_\ell} \widehat{f}_\ell(u) - 2\beta_\ell}} \mathcal{V}_\ell(c). \tag{152}$$

C.2 KERNEL-COVER CONSTRUCTION

Appendix C proves that Algorithm 5 terminates after finitely many iterations and returns a d_k-cover of radius at most $2^{-\ell-2}$.

Algorithm 5 Kernel-pseudometric cover

Require: Nonempty compact active set $\mathcal{X}_\ell$, kernel pseudometric d_k, and covering radius $2^{-\ell-2}$
Choose $c_{\ell,1} \in \mathcal{X}_\ell$, set $\mathcal{C}_\ell \leftarrow \{c_{\ell,1}\}$, and set $m \leftarrow 1$
while $\sup_{a\in\mathcal{X}_\ell} \min_{c\in\mathcal{C}_\ell} d_k(a,c) > 2^{-\ell-2}$ **do**
  Choose $c_{\ell,m+1} \in \arg\max_{a\in\mathcal{X}_\ell} \min_{c\in\mathcal{C}_\ell} d_k(a,c)$
  Set $\mathcal{C}_\ell \leftarrow \mathcal{C}_\ell \cup \{c_{\ell,m+1}\}$ and $m \leftarrow m+1$
end while
Set $\mathcal{V}_\ell(c) \leftarrow \{a \in \mathcal{X}_\ell : d_k(a,c) = \min_{u\in\mathcal{C}_\ell} d_k(a,u)\}$ for every $c \in \mathcal{C}_\ell$
return $(\mathcal{C}_\ell, \{\mathcal{V}_\ell(c) : c \in \mathcal{C}_\ell\})$

C.3 Certified design algorithm

Algorithm 6 Certified multiplicative design

Require: Finite candidate set $\{c_1,\dots,c_m\}$, regularization $\lambda > 0$, certification factor $\kappa > 1$, and iteration limit $S \in \mathbb{N} \cup \{\infty\}$
Choose $w^{(0)}$ with $w_i^{(0)} > 0$ and $\sum_{i=1}^m w_i^{(0)} = 1$
$r \leftarrow \dim \operatorname{span}\{\phi(c_1),\dots,\phi(c_m)\}$
for $s = 0,\dots,S-1$ **do**
  Compute $V_{w^{(s)}}$ and $h_{w^{(s)}}(c_i)$ for $i = 1,\dots,m$ using equation 15 and equation 16
  Compute $d_{w^{(s)}}$ using equation 17
  if $\max_{1\le i\le m} h_{w^{(s)}}(c_i) \le \kappa d_{w^{(s)}}$ **then**
    return $\left(w^{(s)}, V_{w^{(s)}}, d_{w^{(s)}}\right)$
  end if
  $w_i^{(s+1)} \leftarrow \frac{w_i^{(s)}\left(h_{w^{(s)}}(c_i)+r-d_{w^{(s)}}\right)}{r}$ for $i = 1,\dots,m$
end for
return failure

C.4 Finite certification

Lemma 9 shows that a maximizer is certified with factor $\kappa = 1$ in equation 18. Algorithm 6 only needs to return an iterate satisfying that condition.

Corollary 15 (Finite certification of the multiplicative distribution)**.** *Fix a finite candidate set $\mathcal{C} = \{c_1,\dots,c_m\}$, and let $w^\star$ maximize Φ_λ over probability distributions on $\mathcal{C}$. With full-support initialization, Algorithm 6 generates full-support probability distributions $w^{(s)}$. If the feature span is nonzero, then*

$$d_{w^{(s)}} \to d_{w^\star}, \qquad \frac{\max_{a\in\mathcal{C}} h_{w^{(s)}}(a)}{d_{w^{(s)}}} \to 1. \tag{153}$$

Consequently, for every $\kappa > 1$, certification succeeds after finitely many sweeps when $S = \infty$ or S is sufficiently large. If the feature span is zero, the algorithm returns at its first certificate check.

Proof. **Proof outline.** The proof first verifies that the update preserves the probability simplex and full support. It then identifies the update with a D-optimal multiplicative update and uses its convergence to prove finite certification. Write $w_i^{(s)}$ for the weight assigned to candidate c_i after s multiplicative updates, and recall that r is the dimension of their feature span. If $r = 0$, all leverage scores and d_w vanish, so the algorithm returns at its first certificate check. Henceforth, suppose $r \ge 1$.

Feasibility of the iterates. Let $\mathcal{S} := \operatorname{span}\{\phi(c_1),\dots,\phi(c_m)\}$. The identity in equation 23 gives

$$\begin{aligned}\sum_{i=1}^m w_i^{(s+1)} &= \frac{1}{r}\sum_{i=1}^m w_i^{(s)}\left(h_{w^{(s)}}(c_i) + r - d_{w^{(s)}}\right)\\ &= \frac{d_{w^{(s)}} + r - d_{w^{(s)}}}{r} = 1.\end{aligned} \tag{154}$$

Moreover,

$$r - d_{w^{(s)}} = \lambda \operatorname{Tr}((V_{w^{(s)}}^{-1})|_{\mathcal{S}}) > 0. \tag{155}$$

We write $(V_{w^{(s)}}^{-1})|_{\mathcal{S}}$ for the restriction of $V_{w^{(s)}}^{-1}$ to the finite-dimensional space $\mathcal{S}$. Thus every update factor is positive, so positive initialization preserves full support without an additional normalization step.

D-optimal representation. Identify $\mathcal{S}$ with $\mathbb{R}^r$ and define

$$H_i := I + \lambda^{-1}\phi(c_i)\phi(c_i)^\top, \qquad M_s := \sum_{i=1}^{m} w_i^{(s)} H_i = I + \lambda^{-1} A_{w^{(s)}}|_{\mathcal{S}}. \tag{156}$$

A direct trace calculation gives

$$\operatorname{Tr}(M_s^{-1} H_i) = r - d_{w^{(s)}} + h_{w^{(s)}}(c_i). \tag{157}$$

By equation 157, the standard matrix-valued D-optimal multiplicative update $w_i^{(s+1)} = \frac{w_i^{(s)} \operatorname{Tr}(M_s^{-1} H_i)}{r}$ coincides with the update in Algorithm 6.

D-optimal convergence. Because every H_i is positive definite, Harman & Trnovská (2009, Section 5, Theorem 4) gives

$$\det(M_s) \longrightarrow \det(M_\star), \qquad M_\star := \sum_{i=1}^{m} w_i^\star H_i. \tag{158}$$

The convex hull $\operatorname{conv}\{H_1, \dots, H_m\}$ is the feasible matrix set and is compact, while the log-determinant is strictly concave on positive-definite matrices. Therefore $M_\star$ is the unique optimal matrix and $M_s \to M_\star$. Continuity of inversion and the trace yields $d_{w^{(s)}} \to d_{w^\star}$.

Certificate convergence and termination. For each i, first-order optimality at $M_\star$ in the direction $H_i - M_\star$ implies

$$\operatorname{Tr}\big(M_\star^{-1}(H_i - M_\star)\big) \le 0, \tag{159}$$

and hence $\operatorname{Tr}(M_\star^{-1} H_i) \le r$. On the other hand, averaging these traces with the optimal weights gives

$$\sum_{i=1}^{m} w_i^\star \operatorname{Tr}(M_\star^{-1} H_i) = \operatorname{Tr}\Big(M_\star^{-1} \sum_{i=1}^{m} w_i^\star H_i\Big) = \operatorname{Tr}(M_\star^{-1} M_\star) = r. \tag{160}$$

Therefore,

$$\max_{1 \le i \le m} \operatorname{Tr}(M_\star^{-1} H_i) = r. \tag{161}$$

The identity in equation 157, continuity, and finiteness of the candidate set yield

$$\max_{1 \le i \le m} h_{w^{(s)}}(c_i) - d_{w^{(s)}} \longrightarrow 0. \tag{162}$$

Because $r \ge 1$, at least one candidate feature is nonzero, and optimality gives $d_{w^\star} > 0$. Since $d_{w^{(s)}}$ converges to $d_{w^\star} > 0$, division by $d_{w^{(s)}}$ shows that the certificate ratio converges to one. If $\kappa > 1$, convergence of the certificate ratio shows that the stopping inequality holds for all sufficiently large s, proving finite return. □

C.5 Cover confidence and cellwise elimination

Recall the kernel pseudometric d_k from equation 150. At the beginning of phase ℓ, choose a finite set $\mathcal{C}_\ell \subseteq \mathcal{X}_\ell$ satisfying

$$\sup_{a \in \mathcal{X}_\ell} \min_{c \in \mathcal{C}_\ell} d_k(a, c) \le 2^{-\ell-2}. \tag{163}$$

Indeed, $d_k(a,b)$ is the Hilbert-space distance between $\phi(a)$ and $\phi(b)$. Continuity of ϕ and compactness of $\mathcal{X}_\ell$ imply that $\phi(\mathcal{X}_\ell)$ is compact in $\mathcal{H}_k$. The Hilbert-space balls of radius $2^{-\ell-2}$ centered at points of $\phi(\mathcal{X}_\ell)$ form an open cover of $\phi(\mathcal{X}_\ell)$. Compactness gives a finite subcover, and choosing the corresponding preimages as centers yields a finite set $\mathcal{C}_\ell$ satisfying equation 163. Algorithm 5 constructs the cover by the following farthest-first procedure. Choose any $c_{\ell,1} \in \mathcal{X}_\ell$. Given $c_{\ell,1}, \dots, c_{\ell,m}$, if their covering radius exceeds $2^{-\ell-2}$, choose

$$c_{\ell,m+1} \in \arg\max_{a \in \mathcal{X}_\ell} \min_{1 \le i \le m} d_k(a, c_{\ell,i}). \tag{164}$$

The maximizer exists because the objective is continuous on compact $\mathcal{X}_\ell$. If the procedure did not terminate, the selected centers would form an infinite $2^{-\ell-2}$-separated set, contradicting total boundedness. Thus, it stops after finitely many steps, and the resulting centers form $\mathcal{C}_\ell$. For each $c \in \mathcal{C}_\ell$, Algorithm 5 defines the following closed region:

$$\mathcal{V}_\ell(c) := \Big\{ a \in \mathcal{X}_\ell : d_k(a,c) = \min_{u \in \mathcal{C}_\ell} d_k(a,u) \Big\}. \tag{165}$$

Every $a \in \mathcal{X}_\ell$ has a nearest center in the finite set $\mathcal{C}_\ell$, so these regions cover $\mathcal{X}_\ell$. For $a \in \mathcal{V}_\ell(c)$, equation 165 and equation 163 give $d_k(a,c) \le 2^{-\ell-2}$. If several centers attain the minimum, a belongs to each corresponding region and remains active whenever one of them is retained.

Algorithm 4 runs the distribution, FPS, and estimator subroutines on $\mathcal{C}_\ell$ with $\kappa = \kappa_{\max}$. After the certified design returns, apply the regularization-bias check in Algorithm 4; if it fails, query the remaining points from w_ℓ and stop. For a phase that passes the bias check, compute J_ℓ and the chain on $\mathcal{C}_\ell$, and select (n_ℓ, β_ℓ) by equation 151. Because only cover centers are sampled and estimated, the cover-level design certificate suffices. The decision to retain $\mathcal{V}_\ell(c)$ is based on its center estimate $\widehat{f}_\ell(c)$; overlapping regions retain separate center estimates. After a completed phase, the active set is updated according to equation 152. The additive term $B2^{-\ell-2}$ in equation 151 accounts for reward variation within each region.

Corollary 16 (Cover confidence and cellwise elimination)**.** *Suppose that the certified-design routine succeeds on $\mathcal{C}_\ell$ in every initiated phase. For every completed phase ℓ, conditionally on the observations collected before phase ℓ, the cover condition equation 163 and the phase-length rule equation 151 give*

$$\sup_{c \in \mathcal{C}_\ell} \sup_{a \in \mathcal{V}_\ell(c)} \big| \widehat{f}_\ell(c) - f(a) \big| \le \beta_\ell \tag{166}$$

with probability at least $1 - \delta_\ell$. On the common confidence event, $x^\star$ remains active and

$$\sup_{a \in \mathcal{X}_{\ell+1}} \big(f(x^\star) - f(a) \big) \le 4\beta_\ell \le B2^{2-\ell}. \tag{167}$$

Proof. **Proof outline.** We first extend the confidence bound from the cover centers to their associated regions. We then use this bound to establish optimizer preservation and the active-set gap. The quantity $\beta_\ell(n)$ is the candidate confidence radius on the cover, whereas β_ℓ in equation 151 also includes the cell approximation error $B2^{-\ell-2}$.

Confidence transfer. Apply Lemma 5 with its compact candidate set replaced by $\mathcal{C}_\ell$ and certificate factor $\kappa = \kappa_{\max}$. A finite set is compact, so the same confidence bound applies without a new estimation argument. Conditionally on the observations before phase ℓ, with probability at least $1 - \delta_\ell$,

$$\sup_{c \in \mathcal{C}_\ell} \big| \widehat{f}_\ell(c) - f(c) \big| \le \beta_\ell(n_\ell). \tag{168}$$

For $c \in \mathcal{C}_\ell$ and $a \in \mathcal{V}_\ell(c)$, the reproducing property gives the first two equalities below. Cauchy-Schwarz, $\|f\|_{\mathcal{H}_k} \le B$, and the cover condition give the remaining inequalities.

$$\begin{aligned} |f(a) - f(c)| &= \big| \langle f, \phi(a) \rangle_{\mathcal{H}_k} - \langle f, \phi(c) \rangle_{\mathcal{H}_k} \big| \\ &= \big| \langle f, \phi(a) - \phi(c) \rangle_{\mathcal{H}_k} \big| \\ &\le \|f\|_{\mathcal{H}_k} \|\phi(a) - \phi(c)\|_{\mathcal{H}_k} \\ &\le B d_k(a,c) \\ &\le B2^{-\ell-2}. \end{aligned} \tag{169}$$

On the event in equation 168, fix $c \in \mathcal{C}_\ell$ and $a \in \mathcal{V}_\ell(c)$. The confidence guarantee in equation 168 bounds $|\widehat{f}_\ell(c) - f(c)|$ by $\beta_\ell(n_\ell)$, while equation 169 bounds $|f(c) - f(a)|$ by $B2^{-\ell-2}$. The triangle inequality therefore gives

$$\begin{aligned}\left|\widehat{f}_\ell(c) - f(a)\right| &\le \left|\widehat{f}_\ell(c) - f(c)\right| + |f(c) - f(a)| \\ &\le \beta_\ell(n_\ell) + B2^{-\ell-2} = \beta_\ell, \end{aligned} \tag{170}$$

where the equality is the definition of β_ℓ in equation 151. Taking the supremum over $a \in \mathcal{V}_\ell(c)$ and then over $c \in \mathcal{C}_\ell$ proves equation 166.

Continuous elimination. Let $x^\star \in \arg\max_{a \in \mathcal{X}} f(a)$ be an optimizer. Suppose that $x^\star \in \mathcal{X}_\ell$, and choose $c^\star \in \mathcal{C}_\ell$ such that $x^\star \in \mathcal{V}_\ell(c^\star)$. For every $u \in \mathcal{C}_\ell$, equation 166 and optimality of $x^\star$ give

$$\begin{aligned}\widehat{f}_\ell(c^\star) &\ge f(x^\star) - \beta_\ell \ge f(u) - \beta_\ell \\ &\ge \widehat{f}_\ell(u) - 2\beta_\ell. \end{aligned} \tag{171}$$

Thus, equation 152 retains $\mathcal{V}_\ell(c^\star)$ and hence retains $x^\star$. For any $a \in \mathcal{X}_{\ell+1}$, there is a retained $c \in \mathcal{C}_\ell$ with $a \in \mathcal{V}_\ell(c)$. The confidence guarantee in equation 166 gives $f(a) \ge \widehat{f}_\ell(c) - \beta_\ell$. Since c is retained, the condition in equation 152 gives

$$\widehat{f}_\ell(c) \ge \max_{u \in \mathcal{C}_\ell} \widehat{f}_\ell(u) - 2\beta_\ell \ge \widehat{f}_\ell(c^\star) - 2\beta_\ell. \tag{172}$$

Subtracting β_ℓ gives $\widehat{f}_\ell(c) - \beta_\ell \ge \widehat{f}_\ell(c^\star) - 3\beta_\ell$. Moreover, $x^\star \in \mathcal{V}_\ell(c^\star)$, so equation 166 gives $\widehat{f}_\ell(c^\star) \ge f(x^\star) - \beta_\ell$. Subtracting $3\beta_\ell$ gives $\widehat{f}_\ell(c^\star) - 3\beta_\ell \ge f(x^\star) - 4\beta_\ell$. Combining $f(a) \ge \widehat{f}_\ell(c) - \beta_\ell$, $\widehat{f}_\ell(c) - \beta_\ell \ge \widehat{f}_\ell(c^\star) - 3\beta_\ell$, and $\widehat{f}_\ell(c^\star) - 3\beta_\ell \ge f(x^\star) - 4\beta_\ell$ gives

$$\begin{aligned} f(a) &\ge \widehat{f}_\ell(c) - \beta_\ell \\ &\ge \widehat{f}_\ell(c^\star) - 3\beta_\ell \ge f(x^\star) - 4\beta_\ell. \end{aligned} \tag{173}$$

This proves equation 167. Since d_k is continuous, each cell in equation 165 is closed in $\mathcal{X}_\ell$. The definition in equation 152 is a finite union of these cells, so induction shows that every active set is compact and admits the next finite cover.

□

C.6 Regret of the practical version

Proof of Corollary 7. Apply Lemmas 13 and 14 to the finite compact candidate set $\mathcal{C}_\ell$ with $\kappa = \kappa_{\max}$. Lemma 5 then controls all cover-center estimates. Corollary 16 transfers this confidence to the cells and preserves an optimizer with active-set gap at most $B2^{2-\ell}$. Its conditional failure probabilities sum to at most δ.

For a phase passing the bias check, the regularization bias and terminal bias are each at most $B2^{-\ell-2}$. The cover adds $B2^{-\ell-2}$ to the confidence radius. Thus,

$$B\sqrt{\lambda_T \kappa_{\max} d_{w_\ell}} + BR_\ell + B2^{-\ell-2} \le 3 \cdot B2^{-\ell-2}. \tag{174}$$

Making the stochastic term at most $B2^{-\ell-2}$ therefore suffices in equation 151. This is exactly the stochastic requirement used to obtain equation 137, so that phase-length bound holds with $\kappa = \kappa_{\max}$. The nonnegative unit term in the radius also gives equation 139, hence the same logarithmic phase count.

Every sampled center belongs to the compact active set. The cellwise gap bound therefore gives the same per-round regret bound as equation 131. Lemma 10 applies to the discrete distribution w_ℓ as a Borel distribution on the original domain and gives $d_{w_\ell} \le \gamma_T$. The aggregation in equation 143–equation 146, and the failed-bias-check estimate in equation 149, now apply with $\kappa = \kappa_{\max}$. Combining these bounds completes the proof. □

C.7 COMPUTATIONAL COMPLEXITY

We compare the time complexity of practical HT-ChainPE (Algorithm 4) with the corresponding complexities of the baseline algorithms. We count each kernel evaluation as unit cost and continuous-optimization calls separately because their arithmetic cost depends on the domain and numerical solver. For the baselines, each dense kernel-matrix factorization is assumed to be computed from scratch.

HT-ChainPE. Let L_T denote the final phase initiated by time T. In phase ℓ, the kernel-cover construction (Algorithm 5) repeatedly finds a point farthest from the current cover $\mathcal{C}_\ell$. It makes one optimization call for each added point and one final call to certify the covering radius, giving $\mathcal{O}(|\mathcal{C}_\ell|)$ calls in that phase and $\mathcal{O}(\sum_{\ell=1}^{L_T} |\mathcal{C}_\ell|)$ calls overall. After selecting the cover, forming its kernel matrix requires evaluating $k(c, c')$ for pairs $c, c' \in \mathcal{C}_\ell$. There are at most $|\mathcal{C}_\ell|^2$ such pairs, so this costs $\mathcal{O}(|\mathcal{C}_\ell|^2)$ kernel evaluations outside the optimization calls. The matrix is then cached for the remaining computations in the phase.

The certified-design algorithm (Algorithm 6) uses this finite matrix and performs at most S multiplicative sweeps. Each sweep involves dense matrix operations costing $\mathcal{O}(|\mathcal{C}_\ell|^3)$ arithmetic. The final factorization is reused to precompute the influence weights $\langle \phi(a), V_\ell^{-1}\phi(c)\rangle_{\mathcal{H}_k}$ and chaining distances $\rho_\ell(a, b)$ for $a, b, c \in \mathcal{C}_\ell$ in $\mathcal{O}(|\mathcal{C}_\ell|^3)$ dense arithmetic. For an edge (a, b), the estimator uses the weight $\langle \phi(a) - \phi(b), V_\ell^{-1}\phi(c)\rangle_{\mathcal{H}_k}$. After the regularization-bias check and the computation of J_ℓ and ρ_ℓ, the farthest-point-chain algorithm (Algorithm 2) constructs the chain on $\mathcal{C}_\ell$. After this precomputation, accessing all pairwise ρ_ℓ values costs $\mathcal{O}(|\mathcal{C}_\ell|^2)$. Maintaining the nearest representatives $\pi_{\ell,j}(a)$ over $j = 1, \ldots, J_\ell$ costs $\mathcal{O}(J_\ell|\mathcal{C}_\ell|)$.

After collecting n_ℓ observations, the chained estimator (Algorithm 3) processes at most $|\mathcal{C}_\ell|$ distinct edges per level. With these weights cached, processing all observations therefore costs $\mathcal{O}(n_\ell J_\ell|\mathcal{C}_\ell|)$. Finally, scanning $\mathcal{C}_\ell$ and recording the retained regions as a symbolic union costs $\mathcal{O}(|\mathcal{C}_\ell|)$ arithmetic operations.

For a finite design cap S, summing these costs over the phases gives

$$\mathcal{O}\Bigg(\underbrace{T}_{\text{per-round sampling}} + \underbrace{S \sum_{\ell=1}^{L_T} |\mathcal{C}_\ell|^3}_{\substack{\text{certified-design} \\ \text{matrix operations}}} + \underbrace{\sum_{\ell=1}^{L_T} n_\ell J_\ell |\mathcal{C}_\ell|}_{\substack{\text{chained} \\ \text{estimation}}} \Bigg). \tag{175}$$

The $\mathcal{O}(T)$ term covers recording the observations and handling the final incomplete phase. The cubic design term and the estimation term also bound the kernel-matrix, chain-construction, and elimination costs described above.

TGP-UCB. At round t, TGP-UCB chooses a query point using a posterior mean computed from past truncated rewards and the exact posterior standard deviation. It then observes y_t and truncates this reward. The truncated reward enters the new posterior mean. The posterior variance depends only on the query points and kernel. Building the dense $t \times t$ kernel matrix uses $\mathcal{O}(t^2)$ kernel evaluations. A fresh factorization costs $\mathcal{O}(t^3)$. The posterior-mean solve costs $\mathcal{O}(t^2)$, and scalar truncation costs $\mathcal{O}(1)$. Hence the model-update costs are

$$\mathcal{O}\Big(\underbrace{t^2}_{\text{kernel matrix}} + \underbrace{t^3}_{\text{factorization}} + \underbrace{t^2}_{\text{mean solve}} + \underbrace{1}_{\text{truncation}} \Big) = \mathcal{O}(t^3),$$

$$\sum_{t=1}^{T} \mathcal{O}(t^3) = \mathcal{O}(T^4). \tag{176}$$

After factorization, evaluating the variance at a candidate point costs $\mathcal{O}(t^2)$. The T acquisition maximizations are counted separately.

ATA-GP-UCB. ATA-GP-UCB extends the directional truncation method of Shao et al. (2018) to GP bandits over a continuous domain. It projects kernel features onto a finite-dimensional Nyström

subspace. It then applies directional truncation in that subspace. At round t, it chooses a query point using an approximate posterior mean and standard deviation. After observing y_t, it constructs a Nyström dictionary by sampling a subset of the query points $\{x_1, \ldots, x_t\}$. It then recomputes the feature embeddings of all query points using the approximate feature map defined by this dictionary. It forms the regularized feature covariance and its inverse square root. It also recomputes the truncation directions and truncates past rewards in those directions. The resulting rewards determine the new approximate posterior mean. Let $m_t \le t$ be the Nyström dictionary size at round t. Constructing the dictionary costs $\mathcal{O}(t)$. Factoring its $m_t \times m_t$ kernel matrix costs $\mathcal{O}(m_t^3)$. Recomputing the embeddings of the t past query points costs $\mathcal{O}(tm_t^2)$. Forming the regularized feature covariance costs another $\mathcal{O}(tm_t^2)$. Computing its inverse square root costs $\mathcal{O}(m_t^3)$. Computing the truncation directions costs $\mathcal{O}(tm_t^2)$. Truncating t rewards in m_t directions costs $\mathcal{O}(tm_t)$. Thus the model-update costs are

$$\mathcal{O}\Big(\underbrace{t}_{\text{dictionary}} + \underbrace{2m_t^3}_{\substack{\text{kernel factorization and} \\ \text{inverse square root}}} + \underbrace{3tm_t^2}_{\substack{\text{embeddings, covariance,} \\ \text{and truncation directions}}} + \underbrace{tm_t}_{\text{history truncation}} \Big) = \mathcal{O}(tm_t^2),$$

$$\sum_{t=1}^{T} \mathcal{O}(tm_t^2) = \mathcal{O}(T^2 m_{\max}^2) \le \mathcal{O}(T^4). \tag{177}$$

Each acquisition-score evaluation costs at most $\mathcal{O}(t^2)$, and the T maximizations are counted separately. In the three synthetic benchmarks of Park & Lee (2024, Table 1), ATA-GP-UCB was about 8–106 times slower than TGP-UCB.

CA-TGP-UCB. At round t, CA-TGP-UCB chooses a query point using a posterior mean based on past truncated rewards and the exact posterior standard deviation. After observing y_t, it computes a query-dependent truncation weight and threshold. It replaces y_t by zero if the absolute weighted reward exceeds that threshold. The resulting reward vector determines the new posterior mean. Building the dense kernel matrix costs $\mathcal{O}(t^2)$ kernel evaluations. A fresh factorization costs $\mathcal{O}(t^3)$. The weight, threshold, and mean solves cost at most $\mathcal{O}(t^2)$ each. Thus the model-update costs are

$$\mathcal{O}\Big(\underbrace{t^2}_{\text{kernel matrix}} + \underbrace{t^3}_{\text{factorization}} + \underbrace{3t^2}_{\substack{\text{weight, threshold,} \\ \text{and mean solves}}} + \underbrace{1}_{\text{truncation}} \Big) = \mathcal{O}(t^3),$$

$$\sum_{t=1}^{T} \mathcal{O}(t^3) = \mathcal{O}(T^4). \tag{178}$$

There is one acquisition maximization per round. Unlike ATA-GP-UCB, this algorithm does not rebuild a feature embedding or re-truncate the entire reward history in multiple directions.

MoM-GP-UCB. For fixed failure probability, each episode selects one query point and collects $\mathcal{O}(\log T)$ rewards there. The horizon thus contains $\mathcal{O}(T/\log T)$ episodes and the same number of acquisition maximizations. After each episode, one reward is added to each of the $\mathcal{O}(\log T)$ reward vectors. These vectors share one kernel matrix over the episode query points. The algorithm factors that matrix once and solves for each posterior mean. It takes the pointwise median of these predictions and updates the common posterior variance. Summing the cubic factorizations over all episodes costs $\mathcal{O}((T/\log T)^4)$. The quadratic solves for the $\mathcal{O}(\log T)$ means per episode sum to $\mathcal{O}(T^3/\log^2 T)$. Thus the cumulative model-update cost is

$$\sum_{i=1}^{\lceil T/\log T \rceil} \mathcal{O}\Big(\underbrace{i^2}_{\text{kernel matrix}} + \underbrace{i^3}_{\text{factorization}} + \underbrace{i^2 \log T}_{\text{mean solves}} + \underbrace{\log T}_{\text{reward-vector updates}} \Big) = \mathcal{O}(T^4/\log^4 T).$$

Computing the pointwise median at candidate points is part of acquisition optimization.

D PROOF OF THE LOWER BOUND

This section proves the kernel-specific lower bounds (Theorem 8). We first construct a heavy-tailed channel and derive its adaptive KL bound; together, they establish the localized hard-family reduction (Lemma 17). We then construct the required families for the squared-exponential (SE) and Matérn kernels (Lemmas 18 and 19).

Table 5: Comparison of cumulative-regret lower bounds for the SE and Matérn kernels.

Work	Noise	Cumulative-regret lower bound	Guarantee and domain
SE kernel			
Scarlett et al. (2017)	Gaussian	$\Omega\big(T^{\frac{1}{2}}(\log T)^{\frac{d}{4}}\big)$	Expected $[0,1]^d$
Cai & Scarlett (2021)	Gaussian	$\Omega\big(T^{\frac{1}{2}}(\log(\frac{1}{\delta}))^{\frac{1}{2}}\big(\log(\frac{T}{\log(1/\delta)})\big)^{\frac{d}{4}}\big)$	High probability $[0,1]^d$
Iwazaki (2026)	Gaussian	$\Omega\Big(T^{\frac{1}{2}}\big(\frac{\log(\frac{T}{\log(1/\delta)})}{\log\log(\frac{T}{\log(1/\delta)})}\big)^{\frac{d}{2}}\Big)$	High probability $\mathbb{S}^d$
Chowdhury & Gopalan (2019)	Heavy-tailed	$\Omega\big(T^{\frac{1}{1+\epsilon}}(\log T)^{\frac{d\epsilon}{2(1+\epsilon)}}\big)$	Expected $[0,1]^d$
This work	Heavy-tailed	$\Omega\Big(T^{\frac{1}{1+\epsilon}}L_\delta^{\frac{\epsilon}{1+\epsilon}}\big(\frac{\log(T/L_\delta)}{\log\log(T/L_\delta)}\big)^{\frac{d\epsilon}{1+\epsilon}}\Big)$	High probability $[0,1]^d$
Matérn-ν kernel			
Scarlett et al. (2017)	Gaussian	$\Omega\big(T^{\frac{\nu+d}{2\nu+d}}\big)$	Expected $[0,1]^d$
Cai & Scarlett (2021)	Gaussian	$\Omega\big(T^{\frac{\nu+d}{2\nu+d}}(\log(\frac{1}{\delta}))^{\frac{\nu}{2\nu+d}}\big)$	High probability $[0,1]^d$
Iwazaki (2026)	Gaussian	–	High probability $\mathbb{S}^d$
Chowdhury & Gopalan (2019)	Heavy-tailed	$\Omega\big(T^{\frac{\nu+d\epsilon}{\nu(1+\epsilon)+d\epsilon}}\big)$	Expected $[0,1]^d$
This work	Heavy-tailed	$\Omega\big(T^{\frac{\nu+d\epsilon}{\nu(1+\epsilon)+d\epsilon}}L_\delta^{\frac{\nu\epsilon}{\nu(1+\epsilon)+d\epsilon}}\big)$	High probability $[0,1]^d$

Write $L_\delta := (1-2\delta)\log(\frac{1-\delta}{\delta})$. Factors depending only on B and the noise scale are suppressed, while the δ-dependence of each high-probability bound is retained. Theorem 8 gives the parameter dependence for this work. A high-probability lower bound means that the threshold is exceeded with probability greater than δ for some instance, ruling out a $(1-\delta)$-confidence guarantee below that threshold. Since $R_T \geq 0$, it also implies the corresponding expected-regret lower bound multiplied by δ. Expected and high-probability rates should therefore be compared at fixed δ, while explicit δ-dependence is comparable only between the high-probability rows.

When $\epsilon = 1$, the Matérn exponent becomes $\frac{\nu+d}{2\nu+d}$, matching the sub-Gaussian-noise rows. The SE rate at $\epsilon = 1$ has the same T and logarithmic exponents as Iwazaki (2026), but the domains differ.

Table 6: Notation used in the lower-bound proof.

Symbol	Definition
$\mathcal{P}_{\epsilon,v}(f), P_{\mu,\Delta}, A_\Delta, \pi_\Delta$	Conditional finite-moment channel class in equation 179, its three-point member, and the amplitude and nonzero-mass parameters in equation 180.
$L_\delta, \Delta, \mathcal{Z}_\Delta, M_\Delta, R_z, g_z$	Binary testing information, lower-bound resolution, hard-family index set and size, its partition cells, and its localized functions.
$\mathbb{P}_T^h, \mathbb{E}_h$	Transcript law and expectation generated by the channel $x \mapsto P_{h(x),\Delta}$ in Appendix D.2.
C_{ov}, C_*	Squared-overlap and master-reduction constants in Appendix D.3.

D.1 HEAVY-TAILED OBSERVATION CHANNEL

For a mean function f, define

$$\mathcal{P}_{\epsilon,v}(f) := \Big\{(P_x)_{x\in\mathcal{X}} : \int_{\mathbb{R}} y\,P_x(dy) = f(x), \int_{\mathbb{R}} |y-f(x)|^{1+\epsilon} P_x(dy) \le v \text{ for every } x \in \mathcal{X}\Big\}, \tag{179}$$

where $x \mapsto P_x(E)$ is measurable for every Borel set $E \subseteq \mathbb{R}$. Under any $P \in \mathcal{P}_{\epsilon,v}(f)$, observations are generated by $y_t \mid (\mathcal{F}_{t-1}, x_t = x) \sim P_x$ using a fresh conditional draw at each round. Thus, $\eta_t = y_t - f(x_t)$ is conditionally centered and satisfies the moment condition in Assumption 3.

For a resolution $\Delta > 0$, define

$$A_\Delta := \Big(\frac{v}{128\Delta}\Big)^{\frac{1}{\epsilon}}, \qquad \pi_\Delta := \frac{32\Delta}{A_\Delta}. \tag{180}$$

Whenever

$$\Delta \le \frac{v^{\frac{1}{1+\epsilon}}}{128}, \tag{181}$$

we have $\pi_\Delta \le 1$. Indeed, equation 181 implies $4(32\Delta)^{1+\epsilon} \le v$, which is equivalent to $32\Delta \le A_\Delta$. For $|\mu| \le 16\Delta$, let $P_{\mu,\Delta}$ be the distribution on $\{-A_\Delta, 0, A_\Delta\}$ given by

$$\begin{aligned} P_{\mu,\Delta}(A_\Delta) &= \frac{\pi_\Delta}{2} + \frac{\mu}{2A_\Delta}, \\ P_{\mu,\Delta}(-A_\Delta) &= \frac{\pi_\Delta}{2} - \frac{\mu}{2A_\Delta}, \\ P_{\mu,\Delta}(0) &= 1 - \pi_\Delta. \end{aligned} \tag{182}$$

Because $\pi_\Delta A_\Delta = 32\Delta$, the restriction on μ gives $\frac{|\mu|}{2A_\Delta} \le \frac{\pi_\Delta}{4}$. The two nonzero masses in equation 182 therefore lie in $[\frac{\pi_\Delta}{4}, \frac{3\pi_\Delta}{4}]$, the mass at zero is nonnegative, and the three masses sum to one. Finally,

$$\begin{aligned} \mathbb{E}_{P_{\mu,\Delta}}[Y] &= A_\Delta\Big(\frac{\pi_\Delta}{2} + \frac{\mu}{2A_\Delta}\Big) - A_\Delta\Big(\frac{\pi_\Delta}{2} - \frac{\mu}{2A_\Delta}\Big) = \mu, \\ \mathbb{E}_{P_{\mu,\Delta}}[|Y|^{1+\epsilon}] &= \pi_\Delta A_\Delta^{1+\epsilon} = \frac{v}{4}. \end{aligned} \tag{183}$$

The equation 181 also gives $|\mu|^{1+\epsilon} \le \frac{v}{4}$. The elementary inequality $|a-b|^{1+\epsilon} \le 2^\epsilon(|a|^{1+\epsilon} + |b|^{1+\epsilon})$ and equation 183 yield

$$\begin{aligned} \mathbb{E}_{P_{\mu,\Delta}}\big[|Y-\mu|^{1+\epsilon}\big] &\le 2^\epsilon\Big(\frac{v}{4} + |\mu|^{1+\epsilon}\Big) \\ &\le 2^\epsilon\Big(\frac{v}{4} + \frac{v}{4}\Big) = 2^{\epsilon-1}v \le v. \end{aligned} \tag{184}$$

The bound $2^{\epsilon-1}v \le v$ uses $\epsilon \le 1$. The equation 184 verifies the required centered, rather than merely raw, moment condition.

The channel also has a quadratic information bound. For $|\mu|, |\widetilde{\mu}| \le 16\Delta$, the two nonzero masses of $P_{\widetilde{\mu},\Delta}$ are at least $\frac{\pi_\Delta}{4}$. At A_Δ and $-A_\Delta$, the mass differences between $P_{\mu,\Delta}$ and $P_{\widetilde{\mu},\Delta}$ are $\frac{\mu-\widetilde{\mu}}{2A_\Delta}$ and $-\frac{\mu-\widetilde{\mu}}{2A_\Delta}$, respectively. The two distributions have the same mass at zero, so the zero term in the chi-squared divergence vanishes. Therefore,

$$\begin{aligned} D_{\mathrm{KL}}\big(P_{\mu,\Delta} \,\|\, P_{\widetilde{\mu},\Delta}\big) &\le \chi^2\big(P_{\mu,\Delta} \,\|\, P_{\widetilde{\mu},\Delta}\big) \\ &= \frac{(\mu-\widetilde{\mu})^2}{4A_\Delta^2 P_{\widetilde{\mu},\Delta}(A_\Delta)} + \frac{(\mu-\widetilde{\mu})^2}{4A_\Delta^2 P_{\widetilde{\mu},\Delta}(-A_\Delta)} \\ &\le 2\frac{(\mu-\widetilde{\mu})^2/(4A_\Delta^2)}{\pi_\Delta/4} = \frac{2(\mu-\widetilde{\mu})^2}{\pi_\Delta A_\Delta^2} \\ &= \frac{128^{\frac{1}{\epsilon}}}{16} v^{-\frac{1}{\epsilon}} \Delta^{\frac{1}{\epsilon}-1}(\mu-\widetilde{\mu})^2. \end{aligned} \tag{185}$$

The first inequality in equation 185 follows from the relation in Gibbs & Su (2002, Theorem 5): $D_{\mathrm{KL}}(P\|Q) \le \log(1+\chi^2(P\|Q)) \le \chi^2(P\|Q)$.

D.2 ADAPTIVE CHANGE OF MEASURE

For a measurable mean function h whose range lies in $[-16\Delta, 16\Delta]$, each mass in equation 182 is an affine function of $h(x)$. Thus, $x \mapsto P_{h(x),\Delta}(E)$ is measurable for every Borel set $E \subseteq \mathbb{R}$, so the distributions $P_{h(x),\Delta}$ form a measurable family of probability measures on $\mathbb{R}$. Let $\mathbb{P}_T^h$ and $\mathbb{E}_h$ denote the transcript law and expectation generated by this channel. All mean functions used in the lower-bound construction belong to $\mathcal{H}_k$. When $\mathcal{X}$ is infinite, continuity of k on $\mathcal{X} \times \mathcal{X}$ implies that every RKHS function is continuous and hence Borel measurable. When $\mathcal{X}$ is finite, every function from $\mathcal{X}$ to $\mathbb{R}$ is Borel measurable. For any two such mean functions f and $\widetilde{f}$, the conditional query-selection kernel of a fixed algorithm is identical under both instances. The conditional and sequential chain rules for relative entropy therefore give

$$D_{\mathrm{KL}}\big(\mathbb{P}_T^f \,\|\, \mathbb{P}_T^{\widetilde{f}}\big) = \mathbb{E}_f\left[\sum_{t=1}^{T} D_{\mathrm{KL}}\big(P_{f(x_t),\Delta} \,\|\, P_{\widetilde{f}(x_t),\Delta}\big)\right]. \tag{186}$$

To make the cancellation explicit, let $Q_t(\cdot \mid \mathcal{F}_{t-1})$ denote the algorithm's conditional distribution of x_t. For a fixed realized past, the conditional joint kernels under f and $\widetilde{f}$ are

$$Q_t(dx \mid \mathcal{F}_{t-1})P_{f(x),\Delta}(dy) \quad \text{and} \quad Q_t(dx \mid \mathcal{F}_{t-1})P_{\widetilde{f}(x),\Delta}(dy), \tag{187}$$

respectively. The factor $Q_t(dx \mid \mathcal{F}_{t-1})$ is identical because the algorithm uses the same query-selection rule under both instances. On the common support of the two joint laws, their Radon-Nikodym derivative satisfies

$$\frac{d\big\{Q_t(dx \mid \mathcal{F}_{t-1})P_{f(x),\Delta}(dy)\big\}}{d\big\{Q_t(dx \mid \mathcal{F}_{t-1})P_{\widetilde{f}(x),\Delta}(dy)\big\}}(x,y) = \frac{dP_{f(x),\Delta}}{dP_{\widetilde{f}(x),\Delta}}(y). \tag{188}$$

The common query-selection factor therefore disappears from the likelihood ratio. Taking its logarithm and integrating under the first joint law is the conditional chain rule for relative entropy, which gives

$$\begin{aligned}
&D_{\mathrm{KL}}\big(Q_t(dx \mid \mathcal{F}_{t-1})P_{f(x),\Delta}(dy) \,\|\, Q_t(dx \mid \mathcal{F}_{t-1})P_{\widetilde{f}(x),\Delta}(dy)\big) \\
&= D_{\mathrm{KL}}\big(Q_t(\cdot \mid \mathcal{F}_{t-1}) \,\|\, Q_t(\cdot \mid \mathcal{F}_{t-1})\big) + \int_{\mathcal{X}} D_{\mathrm{KL}}\big(P_{f(x),\Delta} \,\|\, P_{\widetilde{f}(x),\Delta}\big)Q_t(dx \mid \mathcal{F}_{t-1}) \\
&= \mathbb{E}_f\Big[D_{\mathrm{KL}}\big(P_{f(x_t),\Delta} \,\|\, P_{\widetilde{f}(x_t),\Delta}\big) \mid \mathcal{F}_{t-1}\Big],
\end{aligned} \tag{189}$$

The first KL term compares the conditional distributions of x_t and is zero because they are identical. The integral is the reward-kernel KL averaged over the conditional distribution of x_t under f. Taking expectation under f, summing over t, and applying the tower property proves equation 186.

For $a, b \in (0,1)$, define the binary Kullback-Leibler divergence by

$$\mathrm{kl}(a,b) := a\log\frac{a}{b} + (1-a)\log\frac{1-a}{1-b}. \tag{190}$$

We also record the testing consequence used in the localized hard-family reduction. Let $\mathbb{P}$ and $\mathbb{Q}$ be probability measures on a common measurable space, and let $\mathcal{A}$ be a measurable event such that $\mathbb{P}(\mathcal{A}) \geq 1-\delta$ and $\mathbb{Q}(\mathcal{A}) \leq \delta$ for $\delta \in (0, \frac{1}{3})$. The measurable indicator map $\mathbf{1}_{\mathcal{A}} : \omega \mapsto \mathbf{1}\{\omega \in \mathcal{A}\}$ has distributions $\mathrm{Bern}(\mathbb{P}(\mathcal{A}))$ and $\mathrm{Bern}(\mathbb{Q}(\mathcal{A}))$ under $\mathbb{P}$ and $\mathbb{Q}$, respectively. The data-processing inequality (Polyanskiy & Wu, 2025, Theorem 2.17) therefore gives

$$D_{\mathrm{KL}}(\mathbb{P} \,\|\, \mathbb{Q}) \geq \mathrm{kl}\big(\mathbb{P}(\mathcal{A}), \mathbb{Q}(\mathcal{A})\big). \tag{191}$$

For $a > b$, direct differentiation gives

$$\begin{aligned}
\frac{\partial}{\partial a}\,\mathrm{kl}(a,b) &= \log\frac{a(1-b)}{b(1-a)} > 0, \\
\frac{\partial}{\partial b}\,\mathrm{kl}(a,b) &= \frac{b-a}{b(1-b)} < 0.
\end{aligned} \tag{192}$$

Therefore,

$$D_{\mathrm{KL}}(\mathbb{P} \,\|\, \mathbb{Q}) \geq \mathrm{kl}(1-\delta, \delta) = (1-2\delta)\log\frac{1-\delta}{\delta} =: L_\delta. \tag{193}$$

The lower bound in equation 193 holds at the boundary by continuity, while it is immediate if the right-hand measure assigns zero probability to an event of positive probability under the left-hand measure.

D.3 LOCALIZED HARD-FAMILY REDUCTION

At resolution Δ, let $\mathcal{Z}_\Delta \subseteq \mathcal{X}$ be a finite set of hard-family centers and set $M_\Delta := |\mathcal{Z}_\Delta|$. Let $(R_z)_{z\in\mathcal{Z}_\Delta}$ be a measurable partition of $\mathcal{X}$ with $z \in R_z$. For $C_{\mathrm{ov}} > 0$, we call $(g_z, R_z)_{z\in\mathcal{Z}_\Delta}$ a localized hard family with overlap constant C_{ov} if

$$\|g_z\|_{\mathcal{H}_k} \le \frac{B}{7}, \qquad \|g_z\|_\infty \le 2\Delta, \qquad \max_{x\in\mathcal{X}} g_z(x) = g_z(z) = 2\Delta, \tag{194}$$

$$x \notin R_z \Longrightarrow g_z(x) < \Delta, \tag{195}$$

$$\sum_{z\in\mathcal{Z}_\Delta} \sup_{x\in R_u} |g_z(x)|^2 \le C_{\mathrm{ov}}\Delta^2 \qquad \text{for every } u \in \mathcal{Z}_\Delta, \tag{196}$$

where C_{ov} is independent of B, v, T, δ, and Δ. Set

$$C_* := 36 \cdot \frac{128^{\frac{1}{\epsilon}}}{16} \cdot C_{\mathrm{ov}}. \tag{197}$$

Lemma 17 (Localized hard-family reduction)**.** *For a localized hard family with $M_\Delta \ge 2$, if $\Delta \le v^{\frac{1}{1+\epsilon}}/128$ and*

$$T < \frac{M_\Delta - 1}{C_*} v^{\frac{1}{\epsilon}} \Delta^{-\frac{1+\epsilon}{\epsilon}} L_\delta, \tag{198}$$

then some $f \in \mathcal{H}_k(\mathcal{X}, B)$ and $P \in \mathcal{P}_{\epsilon,v}(f)$ satisfy

$$\mathbb{P}_T^{f,P}\Big(R_T(f) > \frac{T\Delta}{2}\Big) > \delta. \tag{199}$$

Proof. **Proof outline.** The proof compares a base function with amplified alternatives, converts low regret into a separating event, and then contradicts the horizon condition using the adaptive KL bound.

Admissible alternatives and regret separation. Fix $z \in \mathcal{Z}_\Delta$ and, for each $z' \ne z$, define within this proof the amplified alternative

$$\widetilde{g}_{z'} := g_z + 6g_{z'}. \tag{200}$$

The triangle inequality and equation 194 give

$$\|g_z\|_{\mathcal{H}_k} \le B, \qquad \|\widetilde{g}_{z'}\|_{\mathcal{H}_k} \le B, \qquad \|g_z\|_\infty \le 2\Delta, \qquad \|\widetilde{g}_{z'}\|_\infty \le 14\Delta. \tag{201}$$

The factor 32 in the definition of π_Δ determines the mean range allowed by the channel. Indeed, equation 180 gives $\pi_\Delta A_\Delta = 32\Delta$, and hence

$$|\mu| \le 16\Delta = \frac{\pi_\Delta A_\Delta}{2} \quad \Longrightarrow \quad \frac{|\mu|}{2A_\Delta} \le \frac{\pi_\Delta}{4}. \tag{202}$$

Therefore the two nonzero masses in equation 182 are at least $\pi_\Delta/4$, as required in the KL bound equation 185. Moreover, equation 194 gives

$$\begin{aligned} &\|g_z\|_\infty \le 2\Delta < 16\Delta, \\ &\|\widetilde{g}_{z'}\|_\infty \le \|g_z\|_\infty + 6\|g_{z'}\|_\infty \le 2\Delta + 12\Delta = 14\Delta < 16\Delta. \end{aligned} \tag{203}$$

Thus, the ranges of both g_z and $\widetilde{g}_{z'}$ lie in $[-16\Delta, 16\Delta]$, so the channels $x \mapsto P_{g_z(x),\Delta}$ and $x \mapsto P_{\widetilde{g}_{z'}(x),\Delta}$ are well defined.

Under g_z, every $x \notin R_z$ has regret $2\Delta - g_z(x) > \Delta$ by equation 194 and equation 195. For the alternative, $z' \in \mathcal{Z}_\Delta \subseteq \mathcal{X}$, so its maximum is at least its value at z'. The sup-norm bound in equation 194 gives $|g_z(z')| \le 2\Delta$ and hence $g_z(z') \ge -2\Delta$. The center condition in equation 194 gives $g_{z'}(z') = 2\Delta$. Using $\widetilde{g}_{z'} = g_z + 6g_{z'}$, we obtain

$$\begin{aligned} \max_{u\in\mathcal{X}} \widetilde{g}_{z'}(u) &\ge \widetilde{g}_{z'}(z') \\ &= g_z(z') + 6g_{z'}(z') \\ &\ge -2\Delta + 6(2\Delta) \\ &= 10\Delta. \end{aligned} \tag{204}$$

If $x \in R_z$, then $x \notin R_{z'}$ because $z' \neq z$ and $(R_u)_{u \in \mathcal{Z}_\Delta}$ is a partition of $\mathcal{X}$. Applying equation 195 to $g_{z'}$ gives $g_{z'}(x) < \Delta$. The sup-norm bound in equation 194 gives $g_z(x) \leq |g_z(x)| \leq \|g_z\|_\infty \leq 2\Delta$. Substituting these two bounds into the definition of $\widetilde{g}_{z'}$ yields

$$\begin{aligned} \widetilde{g}_{z'}(x) &= g_z(x) + 6g_{z'}(x) \\ &\leq 2\Delta + 6g_{z'}(x) \\ &< 2\Delta + 6\Delta \\ &= 8\Delta. \end{aligned} \tag{205}$$

Thus, every point in R_z has regret greater than 2Δ under $\widetilde{g}_{z'}$.

Separating event. Define an event that separates the base instance from every alternative. Suppose that every mean function among g_z and $(\widetilde{g}_{z'})_{z' \neq z}$ satisfies

$$\mathbb{P}_T^h\Big(R_T(h) \leq \frac{T\Delta}{2}\Big) \geq 1 - \delta. \tag{206}$$

Let $\mathbf{1}\{\cdot\}$ denote the indicator, and define

$$N_u(T) := \sum_{t=1}^{T} \mathbf{1}\{x_t \in R_u\}. \tag{207}$$

Since every query point outside R_z incurs regret greater than Δ under g_z, whereas every query point inside R_z incurs regret greater than 2Δ under $\widetilde{g}_{z'}$, we have

$$R_T(g_z) \leq \frac{T\Delta}{2} \quad \Longrightarrow \quad N_z(T) > \frac{T}{2}, \tag{208}$$

whereas

$$R_T(\widetilde{g}_{z'}) \leq \frac{T\Delta}{2} \quad \Longrightarrow \quad N_z(T) < \frac{T}{4}. \tag{209}$$

The equation 206 hence implies

$$\mathbb{P}_T^{g_z}\big(N_z(T) \geq \frac{T}{2}\big) \geq 1 - \delta, \qquad \mathbb{P}_T^{\widetilde{g}_{z'}}\big(N_z(T) \geq \frac{T}{2}\big) \leq \delta. \tag{210}$$

Fix $z' \neq z$ and apply equation 193 with $\mathbb{P} = \mathbb{P}_T^{g_z}$, $\mathbb{Q} = \mathbb{P}_T^{\widetilde{g}_{z'}}$, and $\mathcal{A} = \{N_z(T) \geq T/2\}$. The inequalities $\mathbb{P}(\mathcal{A}) \geq 1 - \delta$ and $\mathbb{Q}(\mathcal{A}) \leq \delta$ verify the hypotheses of equation 193. Data processing through the indicator of $\mathcal{A}$ gives the KL-to-binary-KL inequality, and the binary-KL monotonicity used in equation 193 lowers the binary KL to $\mathrm{kl}(1-\delta, \delta)$:

$$\begin{aligned} &D_{\mathrm{KL}}\big(\mathbb{P}_T^{g_z} \,\|\, \mathbb{P}_T^{\widetilde{g}_{z'}}\big) \\ &\geq \mathrm{kl}\Big(\mathbb{P}_T^{g_z}\big(N_z(T) \geq T/2\big), \mathbb{P}_T^{\widetilde{g}_{z'}}\big(N_z(T) \geq T/2\big)\Big) \\ &\geq \mathrm{kl}(1-\delta, \delta) \\ &= (1-2\delta)\log\frac{1-\delta}{\delta} \\ &= L_\delta. \end{aligned} \tag{211}$$

Since $z' \neq z$ was arbitrary, equation 211 holds for every alternative $\widetilde{g}_{z'}$.

Information contradiction. It remains to upper-bound the information available to distinguish these instances. The condition $\Delta \leq v^{\frac{1}{1+\epsilon}}/128$ in Lemma 17 ensures that $\pi_\Delta \leq 1$ and that the channel satisfies the centered $(1+\epsilon)$-moment bound. Moreover, equation 201 shows that the ranges of g_z and $\widetilde{g}_{z'}$ lie in $[-16\Delta, 16\Delta]$. Hence, the channel of Appendix D.1 is valid for both mean functions. For every $z' \neq z$, the partition gives

$$\begin{aligned} \sum_{t=1}^{T} \mathbb{E}_{g_z}\big[|g_{z'}(x_t)|^2\big] &= \sum_{u \in \mathcal{Z}_\Delta} \mathbb{E}_{g_z}\left[\sum_{t=1}^{T} \mathbf{1}\{x_t \in R_u\}|g_{z'}(x_t)|^2\right] \\ &\leq \sum_{u \in \mathcal{Z}_\Delta} \mathbb{E}_{g_z}[N_u(T)] \sup_{x \in R_u} |g_{z'}(x)|^2. \end{aligned} \tag{212}$$

Fix $z' \neq z$. For each t, the channel bound equation 185 can be applied with $\mu = g_z(x_t)$ and $\widetilde{\mu} = \widetilde{g}_{z'}(x_t)$. Then, we have

$$\begin{aligned}
&D_{\mathrm{KL}}\big(\mathbb{P}_T^{g_z} \,\|\, \mathbb{P}_T^{\widetilde{g}_{z'}}\big) \\
&= \mathbb{E}_{g_z}\left[\sum_{t=1}^{T} D_{\mathrm{KL}}\big(P_{g_z(x_t),\Delta} \,\|\, P_{\widetilde{g}_{z'}(x_t),\Delta}\big)\right] \\
&\leq \frac{128^{\frac{1}{\epsilon}}}{16} v^{-\frac{1}{\epsilon}} \Delta^{\frac{1}{\epsilon}-1} \sum_{t=1}^{T} \mathbb{E}_{g_z}\Big[\big(g_z(x_t) - \widetilde{g}_{z'}(x_t)\big)^2\Big] \\
&= 36 \cdot \frac{128^{\frac{1}{\epsilon}}}{16} v^{-\frac{1}{\epsilon}} \Delta^{\frac{1}{\epsilon}-1} \sum_{t=1}^{T} \mathbb{E}_{g_z}\big[|g_{z'}(x_t)|^2\big] \\
&\leq 36 \cdot \frac{128^{\frac{1}{\epsilon}}}{16} v^{-\frac{1}{\epsilon}} \Delta^{\frac{1}{\epsilon}-1} \sum_{u \in \mathcal{Z}_\Delta} \mathbb{E}_{g_z}[N_u(T)] \sup_{x \in R_u} |g_{z'}(x)|^2. \qquad (213)
\end{aligned}$$

The first equality in equation 213 is the adaptive chain rule equation 186. The first inequality applies equation 185 inside the expectation for each t. The definition $\widetilde{g}_{z'} = g_z + 6g_{z'}$ gives $g_z(x_t) - \widetilde{g}_{z'}(x_t) = -6g_{z'}(x_t)$, so the next equality uses $|g_z(x_t) - \widetilde{g}_{z'}(x_t)|^2 = 36|g_{z'}(x_t)|^2$. The final inequality in equation 213 follows from equation 212. Summing equation 211 and equation 213 over $z' \neq z$, applying equation 196, and using $\sum_u N_u(T) = T$ almost surely give

$$\begin{aligned}
(M_\Delta - 1)L_\delta &\leq 36 \cdot \frac{128^{\frac{1}{\epsilon}}}{16} v^{-\frac{1}{\epsilon}} \Delta^{\frac{1}{\epsilon}-1} \sum_{u} \mathbb{E}_{g_z}[N_u(T)] \sum_{z' \neq z} \sup_{x \in R_u} |g_{z'}(x)|^2 \\
&\leq 36 \cdot \frac{128^{\frac{1}{\epsilon}}}{16} \cdot C_{\mathrm{ov}} v^{-\frac{1}{\epsilon}} \Delta^{\frac{1+\epsilon}{\epsilon}} \sum_{u} \mathbb{E}_{g_z}[N_u(T)] \\
&= C_* v^{-\frac{1}{\epsilon}} \Delta^{\frac{1+\epsilon}{\epsilon}} T. \qquad (214)
\end{aligned}$$

equation 214 contradicts equation 198. Therefore equation 206 cannot hold for every alternative, which proves equation 199. □

D.4 SE KERNEL

On $\mathbb{R}^d$, let

$$k_{\mathrm{SE}}(x, y) := \exp\Big(-\frac{\|x - y\|_2^2}{2\ell_{\mathrm{SE}}^2}\Big), \qquad \ell_{\mathrm{SE}} > 0, \qquad (215)$$

and use its restriction to $[0, 1]^d$.

Lemma 18 (SE hard family). *Let $B > 0$ be the RKHS-norm radius in $\mathcal{H}_{k_{\mathrm{SE}}}(\mathcal{X}, B)$, and let $\Delta > 0$ be the hard-family resolution in equation 194. For sufficiently large $\frac{B}{\Delta}$, there is a localized hard family with*

$$M_\Delta \asymp \left[\frac{\log(\frac{B}{\Delta})}{\log\log(\frac{B}{\Delta})}\right]^d. \qquad (216)$$

Proof. **Proof outline.** The proof has four stages. First, we construct a one-dimensional polynomial that equals one at a chosen center and decreases away from it. We multiply this polynomial by a Gaussian and use its coefficient bounds to control the SE RKHS norm. Taking products across coordinates gives a d-dimensional localized function, and placing copies at separated grid points produces the required hard family.

One-dimensional localization. We first build a Chebyshev factor that is bounded by one, equals one at ξ, and becomes small outside a prescribed radius. Multiplying such factors over dyadic radii gives a polynomial that decays as t moves away from ξ. We then bound its degree and its coefficients

in powers of $t-\xi$. The degree bound controls the size of the polynomial, while the coefficient bound is substituted into equation 248 to control the Gaussian-RKHS norm.

Chebyshev building block. Fix $r > 1$. Throughout this proof, $C_r \geq 1$ denotes a constant depending only on r and the fixed length scale ℓ_{SE}, enlarged when necessary. Let $(T_m)_{m\geq 0}$ denote the Chebyshev polynomials of the first kind Trefethen (2019). They are defined recursively on $\mathbb{R}$ by $T_0(t) := 1$, $T_1(t) := t$, and

$$T_{m+1}(t) := 2tT_m(t) - T_{m-1}(t), \qquad m \geq 1. \tag{217}$$

The proof uses the equivalent representations

$$\begin{aligned} T_m(t) &= \cos(m \arccos t), \qquad t \in [-1, 1], \\ T_m(t) &= \cosh(m \operatorname{arcosh} t), \qquad t \geq 1. \end{aligned} \tag{218}$$

For $h \in (0, 2]$ and $\xi \in [-1, 1]$, define

$$P_{m,h,\xi}(t) := \left[\frac{T_m\big(1 + \frac{h^2}{16} - \frac{(t-\xi)^2}{16}\big)}{T_m(1 + \frac{h^2}{16})} \right]^2. \tag{219}$$

We first show that $P_{m,h,\xi}(\xi) = 1$, $0 \leq P_{m,h,\xi}(t) \leq 1$, and $P_{m,h,\xi}(t) \leq 4\exp(-mh/2)$ whenever $|t - \xi| \geq h$. The map $t \mapsto 1 + \frac{h^2}{16} - \frac{(t-\xi)^2}{16}$ is a polynomial of degree two in t. Since T_m has degree m, the composition $t \mapsto T_m(1 + \frac{h^2}{16} - \frac{(t-\xi)^2}{16})$ has degree at most $2m$. Squaring this polynomial gives $\deg(P_{m,h,\xi}) \leq 4m$. At $t = \xi$, the equality $T_m(1 + \frac{h^2}{16} - \frac{(t-\xi)^2}{16}) = T_m(1 + \frac{h^2}{16})$ gives $P_{m,h,\xi}(\xi) = 1$.

Single-factor bounds. For fixed $t, \xi \in [-1, 1]$, write $u := 1 + \frac{h^2}{16} - \frac{(t-\xi)^2}{16}$ and $v := 1 + \frac{h^2}{16}$. To prove $P_{m,h,\xi}(t) \in [0, 1]$, it remains to show that $|T_m(u)| \leq T_m(v)$. Since $|t - \xi| \leq 2$ and $h \in (0, 2]$,

$$0 \leq \frac{(t-\xi)^2}{16} \leq \frac{1}{4}, \qquad 0 < \frac{h^2}{16} \leq \frac{1}{4}. \tag{220}$$

Therefore,

$$\frac{3}{4} \leq u \leq v \leq \frac{5}{4}, \qquad v > 1. \tag{221}$$

The function arcosh is the inverse hyperbolic cosine. If $u \leq 1$, then $u \in [\frac{3}{4}, 1]$, and the cosine representation of the Chebyshev polynomial gives

$$|T_m(u)| = |\cos(m \arccos u)| \leq 1 \leq \cosh(m \operatorname{arcosh} v) = T_m(v). \tag{222}$$

If $u > 1$, then $1 < u \leq v$. Both arcosh on $[1, \infty)$ and $\cosh$ on $[0, \infty)$ are increasing, so

$$0 < T_m(u) = \cosh(m \operatorname{arcosh} u) \leq \cosh(m \operatorname{arcosh} v) = T_m(v). \tag{223}$$

Thus, $|T_m(u)| \leq T_m(v)$ in both cases. Since $T_m(v) = \cosh(m \operatorname{arcosh} v) > 0$, dividing $|T_m(u)| \leq T_m(v)$ by $T_m(v)$ and squaring yields

$$0 \leq P_{m,h,\xi}(t) \leq 1 \qquad \text{for every } t \in [-1, 1]. \tag{224}$$

For $h \in (0, 2]$, we have $\frac{h}{4\sqrt{2}} \leq 1$. The identity $\operatorname{arcosh}(1 + x) = 2\operatorname{arsinh}(\sqrt{x/2})$ and the integral representation of arsinh give

$$\begin{aligned} \operatorname{arcosh}\Big(1 + \frac{h^2}{16}\Big) &= 2\operatorname{arsinh}\Big(\frac{h}{4\sqrt{2}}\Big) \\ &= 2\int_0^{\frac{h}{4\sqrt{2}}} \frac{1}{\sqrt{1+s^2}}\, ds \\ &\geq 2 \cdot \frac{1}{\sqrt{2}} \cdot \frac{h}{4\sqrt{2}} = \frac{h}{4}. \end{aligned} \tag{225}$$

The inequality uses $(1+s^2)^{-1/2} \geq 1/\sqrt{2}$ for $0 \leq s \leq h/(4\sqrt{2}) \leq 1$. Since $\cosh(x) \geq \frac{1}{2}e^x$ for every $x \geq 0$, we obtain

$$T_m\Big(1+\frac{h^2}{16}\Big) = \cosh\Big(m \operatorname{arcosh}\Big(1+\frac{h^2}{16}\Big)\Big) \geq \frac{1}{2}\exp\Big(\frac{mh}{4}\Big). \tag{226}$$

If $|t-\xi| \geq h$, then $(t-\xi)^2 \geq h^2$, so $u = 1 + \frac{h^2}{16} - \frac{(t-\xi)^2}{16} \leq 1$. Since $|t-\xi| \leq 2$ and $h^2 \geq 0$, we also have $u \geq 1 - \frac{(t-\xi)^2}{16} \geq \frac{3}{4}$. Thus, $u \in [\frac{3}{4}, 1]$, and the cosine representation gives $|T_m(u)| \leq 1$. Combining $|T_m(u)| \leq 1$ with $T_m(v) = T_m(1+\frac{h^2}{16}) \geq \frac{1}{2}\exp(\frac{mh}{4})$ gives

$$|t-\xi| \geq h \quad \Longrightarrow \quad P_{m,h,\xi}(t) = \left[\frac{T_m(u)}{T_m(v)}\right]^2 \leq \left[\frac{1}{\frac{1}{2}\exp(mh/4)}\right]^2 = 4\exp\Big(-\frac{mh}{2}\Big). \tag{227}$$

Dyadic product. Choose $A > 0$ so large that $4e^{-A/2} \leq 2^{-r}$. For an integer $N \geq 1$, set within this proof

$$J_N := \lfloor \log_2(2N) \rfloor, \qquad h_j := \frac{2^j}{N}, \qquad m_j := \Big\lceil \frac{A}{h_j} \Big\rceil, \tag{228}$$

Define

$$p_{N,\xi}(t) := \prod_{j=0}^{J_N} P_{m_j,h_j,\xi}(t). \tag{229}$$

Every factor lies in $[0,1]$ and equals one at ξ. Suppose that $|t-\xi| \geq N^{-1}$. Since $h_0 = N^{-1}$, at least the factor indexed by $j=0$ satisfies $h_j \leq |t-\xi|$. For every index j satisfying $h_j \leq |t-\xi|$, the definition of m_j gives $m_j h_j = \lceil A/h_j \rceil h_j \geq A$. Applying equation 227 with $m = m_j$ and $h = h_j$ yields

$$P_{m_j,h_j,\xi}(t) \leq 4\exp(-m_j h_j/2) \leq 4e^{-A/2} \leq 2^{-r}. \tag{230}$$

Thus, every factor with $h_j \leq |t-\xi|$ is at most 2^{-r}. Since $h_j = 2^j/N$, the condition $h_j \leq |t-\xi|$ is equivalent to $j \leq \log_2(N|t-\xi|)$. Moreover, $|t-\xi| \leq 2$ implies $\lfloor \log_2(N|t-\xi|) \rfloor \leq J_N$. Hence, the relevant indices are $j = 0, \ldots, \lfloor \log_2(N|t-\xi|) \rfloor$, whose number is $\lfloor \log_2(N|t-\xi|) \rfloor + 1 > \log_2(N|t-\xi|)$. All remaining factors are at most one, so

$$\begin{aligned} 0 \leq p_{N,\xi}(t) &\leq \big(2^{-r}\big)^{\lfloor \log_2(N|t-\xi|) \rfloor + 1} \\ &\leq 2^{-r\log_2(N|t-\xi|)} = (N|t-\xi|)^{-r} \\ &\leq C_r(1+N|t-\xi|)^{-r}. \end{aligned} \tag{231}$$

The bound $(N|t-\xi|)^{-r} \leq C_r(1+N|t-\xi|)^{-r}$ uses $N|t-\xi| \geq 1$, hence $1+N|t-\xi| \leq 2N|t-\xi|$, and $C_r \geq 2^r$. For $|t-\xi| < N^{-1}$, equation 231 follows from $p_{N,\xi} \leq 1$ after increasing C_r.

Degree and coefficient bounds. Moreover,

$$\begin{aligned} \deg(p_{N,\xi}) &\leq 4\sum_{j=0}^{J_N} m_j \\ &\leq 4A\sum_{j=0}^{J_N} h_j^{-1} + 4(J_N+1) \\ &\leq 8AN + 4(J_N+1) \\ &\leq C_r N. \end{aligned} \tag{232}$$

The bound $\deg(p_{N,\xi}) \leq 4\sum_{j=0}^{J_N} m_j$ follows from $\deg(P_{m_j,h_j,\xi}) \leq 4m_j$. The bound $4\sum_{j=0}^{J_N} m_j \leq 4A\sum_{j=0}^{J_N} h_j^{-1} + 4(J_N+1)$ uses $m_j = \lceil A/h_j \rceil \leq Ah_j^{-1} + 1$. For the third inequality, $h_j^{-1} = N2^{-j}$ and $\sum_{j=0}^{J_N} 2^{-j} \leq 2$. Finally, $J_N + 1 \leq 2N$ for $N \geq 1$, and A depends only on r, so choosing $C_r \geq 8A + 8$ gives $8AN + 4(J_N+1) \leq C_r N$. It remains to bound the polynomial coefficients.

Since $p_{N,\xi}$ is bounded by one on $[-1, 1]$, every Chebyshev coefficient of $p_{N,\xi}$ has absolute value at most two Trefethen (2019).

The initial polynomials $T_0(t) = 1$ and $T_1(t) = t$ both have power-basis coefficient sum one. Multiplication by t shifts every monomial up by one degree without changing its coefficient, so the coefficient sum of $2tT_j(t)$ is twice that of $T_j(t)$. The recurrence $T_{j+1}(t) = 2tT_j(t) - T_{j-1}(t)$ and the triangle inequality therefore show that the coefficient sum of T_{j+1} is at most twice the coefficient sum of T_j plus the coefficient sum of T_{j-1}. If the coefficient sums of T_j and T_{j-1} are at most 3^j and 3^{j-1}, respectively, then the coefficient sum of T_{j+1} is at most $2 \cdot 3^j + 3^{j-1} \le 3^{j+1}$. Induction thus shows that the sum of the absolute power-basis coefficients of T_j is at most 3^j.

Because the absolute value of every Chebyshev coefficient of $p_{N,\xi}$ is at most two and the power-basis coefficient sum of T_j is at most 3^j, the degree-j Chebyshev term contributes at most $2 \cdot 3^j$ to the power-basis coefficient sum of $p_{N,\xi}$. Applying the triangle inequality across all Chebyshev terms bounds the power-basis coefficient sum of $p_{N,\xi}$ by $2\sum_{j=0}^{\deg(p_{N,\xi})} 3^j$. Since $\deg(p_{N,\xi}) \ge 1$, replace C_r by $\max\{C_r, 9\}$, which still depends only on r. Then

$$2 \sum_{j=0}^{\deg(p_{N,\xi})} 3^j = 3^{\deg(p_{N,\xi})+1} - 1 \le 9^{\deg(p_{N,\xi})} \le C_r^{\deg(p_{N,\xi})}. \tag{233}$$

Therefore, the power-basis coefficient sum of $p_{N,\xi}$ is at most $C_r^{\deg(p_{N,\xi})}$. For $m = 0, \dots, \deg(p_{N,\xi})$, define

$$u_m := \frac{p_{N,\xi}^{(m)}(\xi)}{m!}. \tag{234}$$

Since $p_{N,\xi}$ has finite degree, its Taylor expansion at ξ is

$$p_{N,\xi}(t) = \sum_{m=0}^{\deg(p_{N,\xi})} u_m (t - \xi)^m. \tag{235}$$

By the binomial theorem,

$$t^k = \sum_{m=0}^{k} \binom{k}{m} \xi^{k-m} (t - \xi)^m. \tag{236}$$

Since $|\xi| \le 1$, the sum of the absolute coefficients in this expansion satisfies

$$\sum_{m=0}^{k} \binom{k}{m} |\xi|^{k-m} = (1 + |\xi|)^k \le 2^k \le 2^{\deg(p_{N,\xi})}. \tag{237}$$

Thus, the coefficient multiplying t^k contributes at most $2^{\deg(p_{N,\xi})}$ times its absolute value to the coefficient sum in powers of $t - \xi$. Since the coefficient sum of $p_{N,\xi}$ in powers of t is at most $C_r^{\deg(p_{N,\xi})}$, summing these contributions gives

$$\sum_{m=0}^{\deg(p_{N,\xi})} |u_m| \le 2^{\deg(p_{N,\xi})} C_r^{\deg(p_{N,\xi})} = (2C_r)^{\deg(p_{N,\xi})}. \tag{238}$$

Increasing C_r by a factor of two and reusing the same notation, we obtain

$$\max_{0 \le m \le \deg(p_{N,\xi})} |u_m| \le \sum_{m=0}^{\deg(p_{N,\xi})} |u_m| \le C_r^{\deg(p_{N,\xi})}. \tag{239}$$

Gaussian RKHS embedding. For a fixed $z \in \mathbb{R}$, define

$$\psi_{m,z}(x) := \exp\Big(-\frac{(x - z)^2}{2\ell_{\mathrm{SE}}^2}\Big) \frac{(x - z)^m}{\ell_{\mathrm{SE}}^m \sqrt{m!}}. \tag{240}$$

The Gaussian kernel admits the product identity

$$k_{\mathrm{SE}}(x,y) = \exp\Big(-\frac{(x-z)^2}{2\ell_{\mathrm{SE}}^2}\Big)\exp\Big(-\frac{(y-z)^2}{2\ell_{\mathrm{SE}}^2}\Big)\exp\Big(\frac{(x-z)(y-z)}{\ell_{\mathrm{SE}}^2}\Big). \tag{241}$$

Expanding the last exponential gives

$$k_{\mathrm{SE}}(x,y) = \sum_{m=0}^{\infty} \psi_{m,z}(x)\psi_{m,z}(y). \tag{242}$$

The coefficient representation induced by equation 242 is injective. The proof has three links. First, the Weierstrass M-test and the locally uniform limit theorem show that the associated power series is entire. Second, the identity theorem and uniqueness of Taylor coefficients show that only the zero coefficient sequence represents the zero function. Finally, this injectivity identifies the RKHS norm with the ℓ_2 norm of the unique feature-coefficient sequence.

To verify injectivity, suppose that an ℓ_2 sequence (c_m) represents the zero function, and define

$$F(\zeta) := \sum_{m=0}^{\infty} \frac{c_m \zeta^m}{\sqrt{m!}}, \qquad \zeta \in \mathbb{C}. \tag{243}$$

For every $R > 0$ and $|\zeta| \le R$, the Cauchy-Schwarz inequality gives

$$\begin{aligned}\sum_{m=0}^{\infty} \frac{|c_m||\zeta|^m}{\sqrt{m!}} &\le \sum_{m=0}^{\infty} \frac{|c_m| R^m}{\sqrt{m!}} \\ &\le \left(\sum_{m=0}^{\infty} |c_m|^2\right)^{1/2} \left(\sum_{m=0}^{\infty} \frac{R^{2m}}{m!}\right)^{1/2} \\ &= \|(c_m)\|_{\ell_2} \exp(R^2/2) < \infty. \end{aligned} \tag{244}$$

For fixed R, the terms of the series defining F are bounded by a summable sequence that is independent of ζ on the closed disk $|\zeta| \le R$. The Weierstrass M-test therefore shows that the series defining F converges uniformly on that disk. Since $R > 0$ is arbitrary, the convergence is locally uniform on $\mathbb{C}$. Every partial sum is a polynomial and hence entire, so the locally uniform limit theorem for holomorphic functions implies that F is entire.

For every real ζ, the sequence (c_m) represents the zero function. Since $(z + \ell_{\mathrm{SE}}\zeta) - z = \ell_{\mathrm{SE}}\zeta$, the definition of $\psi_{m,z}$ gives

$$\psi_{m,z}(z + \ell_{\mathrm{SE}}\zeta) = \exp\Big(-\frac{(\ell_{\mathrm{SE}}\zeta)^2}{2\ell_{\mathrm{SE}}^2}\Big)\frac{(\ell_{\mathrm{SE}}\zeta)^m}{\ell_{\mathrm{SE}}^m\sqrt{m!}} = \exp(-\zeta^2/2)\frac{\zeta^m}{\sqrt{m!}}. \tag{245}$$

Therefore, using the definition of F,

$$\begin{aligned} 0 = \sum_{m=0}^{\infty} c_m \psi_{m,z}(z + \ell_{\mathrm{SE}}\zeta) &= \exp(-\zeta^2/2)\sum_{m=0}^{\infty} \frac{c_m \zeta^m}{\sqrt{m!}} \\ &= \exp(-\zeta^2/2) F(\zeta). \end{aligned} \tag{246}$$

Because $\exp(-\zeta^2/2) > 0$, we have $F(\zeta) = 0$ for every real ζ. Since $F(\zeta) = 0$ for every real ζ, we have $F(1/n) = 0$ for every $n \ge 1$. These distinct zeros converge to $0 \in \mathbb{C}$, so the zeros of F have an accumulation point in $\mathbb{C}$. The identity theorem for holomorphic functions (Remmert, 1991) now applies because F is entire and its zero set has an accumulation point in the connected domain $\mathbb{C}$. It follows that F is identically zero on $\mathbb{C}$. The coefficient of ζ^m in the Taylor expansion of F at the origin is $c_m/\sqrt{m!}$. Uniqueness of Taylor coefficients therefore gives $c_m = 0$ for every $m \ge 0$. Thus no nonzero ℓ_2 sequence represents the zero function, which proves the claimed injectivity.

In the standard feature-space construction of RKHSs (Carmeli et al., 2010, Proposition 1), the squared RKHS norm of a function is the infimum of the squared ℓ_2 norms of all feature-coefficient sequences representing it. Injectivity makes this representing sequence unique, so the infimum

equals the squared ℓ_2 norm of that sequence. Let $D \geq 1$ be an integer and let $P(x) = \sum_{m=0}^{D} u_m (x-z)^m$ be a polynomial with real coefficients. The definition of $\psi_{m,z}$ gives

$$\begin{aligned}\exp\Big(-\frac{(x-z)^2}{2\ell_{\mathrm{SE}}^2}\Big)P(x) &= \exp\Big(-\frac{(x-z)^2}{2\ell_{\mathrm{SE}}^2}\Big)\sum_{m=0}^{D} u_m (x-z)^m \\ &= \sum_{m=0}^{D} u_m \ell_{\mathrm{SE}}^m \sqrt{m!}\,\psi_{m,z}(x). \end{aligned} \tag{247}$$

Hence, the unique feature coefficient multiplying $\psi_{m,z}$ is $u_m \ell_{\mathrm{SE}}^m \sqrt{m!}$. Taking the squared ℓ_2 norm of these coefficients yields

$$\begin{aligned}\left\|\exp\Big(-\frac{(x-z)^2}{2\ell_{\mathrm{SE}}^2}\Big)P(x)\right\|^2_{\mathcal{H}_{k_{\mathrm{SE}}}(\mathbb{R})} &= \sum_{m=0}^{D}\Big(u_m \ell_{\mathrm{SE}}^m \sqrt{m!}\Big)^2 \\ &= \sum_{m=0}^{D} m!\,\ell_{\mathrm{SE}}^{2m} u_m^2. \end{aligned} \tag{248}$$

For $0 \leq m \leq D$, every factor in $m!$ is at most D, so $m! \leq D^m \leq D^D$. Moreover, the binomial theorem gives $2^D = (1+1)^D \geq \binom{D}{0} + \binom{D}{1} = D+1$. If $\max_{0\leq m\leq D} |u_m| \leq C_r^D$, then equation 248, $m! \leq D^D$, and $D+1 \leq 2^D$ give a norm bound of $(C_r D)^{D/2}$ after increasing C_r by a fixed factor depending only on ℓ_{SE}.

For $z \in [0,1]$, define the function $G_{N,z}$ on $\mathbb{R}$ by

$$G_{N,z}(x) := \exp\Big(-\frac{(x-z)^2}{2\ell_{\mathrm{SE}}^2}\Big) p_{N,2z-1}(2x-1). \tag{249}$$

For $x \in [0,1]$, we have $2x-1 \in [-1,1]$, so every factor in $p_{N,2z-1}(2x-1)$ lies in $[0,1]$. The Gaussian factor also lies in $(0,1]$, giving $0 \leq G_{N,z}(x) \leq 1$ on $[0,1]$. At $x = z$, the Gaussian factor equals one and the polynomial is evaluated at its center $2z-1$. Since $p_{N,2z-1}(2z-1) = 1$, this gives $G_{N,z}(z) = 1$.

For $x, z \in [0,1]$, apply equation 231 with $t = 2x-1$ and $\xi = 2z-1$ to obtain the decay bound. Since $|(2x-1)-(2z-1)| = 2|x-z|$,

$$\begin{aligned} p_{N,2z-1}(2x-1) &\leq C_r(1+2N|x-z|)^{-r} \\ &\leq C_r(1+N|x-z|)^{-r}. \end{aligned} \tag{250}$$

Multiplying by the Gaussian factor, which is at most one, gives $G_{N,z}(x) \leq C_r(1+N|x-z|)^{-r}$. For the norm bound, apply equation 235 with $\xi = 2z-1$ and $t = 2x-1$. Since $(2x-1)-(2z-1) = 2(x-z)$,

$$\begin{aligned} p_{N,2z-1}(2x-1) &= \sum_{m=0}^{\deg(p_{N,2z-1})} u_m\big((2x-1)-(2z-1)\big)^m \\ &= \sum_{m=0}^{\deg(p_{N,2z-1})} u_m\big(2(x-z)\big)^m \\ &= \sum_{m=0}^{\deg(p_{N,2z-1})} 2^m u_m (x-z)^m. \end{aligned} \tag{251}$$

For every $m \leq \deg(p_{N,2z-1})$, equation 239 gives

$$2^m |u_m| \leq 2^{\deg(p_{N,2z-1})} C_r^{\deg(p_{N,2z-1})} = (2C_r)^{\deg(p_{N,2z-1})}. \tag{252}$$

Thus, the transformed coefficients satisfy the same bound after doubling C_r and reusing the notation. Applying equation 248 and using $\deg(p_{N,2z-1}) \leq C_r N$ therefore give

$$\|G_{N,z}\|_{\mathcal{H}_{k_{\mathrm{SE}}}(\mathbb{R})} \leq (C_r N)^{C_r N}. \tag{253}$$

Combining the pointwise bounds on $[0,1]$ with the full-space RKHS norm bound gives, for $x, z \in [0,1]$,

$$\begin{aligned} &0 \le G_{N,z}(x) \le 1, \\ &G_{N,z}(z) = 1, \\ &G_{N,z}(x) \le C_r(1 + N|x - z|)^{-r}, \\ &\|G_{N,z}\|_{\mathcal{H}_{k_{\mathrm{SE}}}(\mathbb{R})} \le (C_r N)^{C_r N}. \end{aligned} \tag{254}$$

Tensorization. Take $r = 2$ in the one-dimensional construction. For a center $z = (z_1, \dots, z_d) \in [0,1]^d$, define the function on $\mathbb{R}^d$ by

$$G^{(d)}_{N,z}(x) := \prod_{i=1}^{d} G_{N,z_i}(x_i), \qquad x = (x_1, \dots, x_d) \in \mathbb{R}^d. \tag{255}$$

The full-space SE kernel is the tensor product of its one-dimensional factors. Let $C \ge 1$ denote a constant depending only on d and ℓ_{SE}, enlarged when necessary. The pointwise bounds in equation 254 and the tensor-product feature representation give, for $x \in [0,1]^d$,

$$0 \le G^{(d)}_{N,z}(x) \le 1, \qquad G^{(d)}_{N,z}(z) = 1, \qquad \|G^{(d)}_{N,z}\|_{\mathcal{H}_{k_{\mathrm{SE}}}(\mathbb{R}^d)} \le (CN)^{CN}, \tag{256}$$

Packing construction. For $\frac{B}{\Delta}$ sufficiently large, choose a small fixed $c_1 > 0$ and set

$$N_\Delta := \left\lfloor c_1 \frac{\log(\frac{B}{\Delta})}{\log\log(\frac{B}{\Delta})} \right\rfloor. \tag{257}$$

Then $N_\Delta \le c_1 \log(B/\Delta)/\log\log(B/\Delta)$. Choose $c_1 \le 1/(4C)$, where C is the constant in equation 256. After fixing c_1, choose a fixed constant $R_0 \ge 49$ such that $N_\Delta \ge 1$ and $CN_\Delta \le \log(B/\Delta)$ whenever $B/\Delta \ge R_0$. Such an R_0 exists because $N_\Delta \ge 1$ for all sufficiently large B/Δ. Moreover, if $\log\log(B/\Delta) \ge 1/4$, then

$$CN_\Delta \le C \frac{c_1 \log(\frac{B}{\Delta})}{\log\log(\frac{B}{\Delta})} \le \frac{\log(\frac{B}{\Delta})}{4\log\log(\frac{B}{\Delta})} \le \log\Big(\frac{B}{\Delta}\Big), \tag{258}$$

where the second inequality follows from $c_1 \le 1/(4C)$ and the third inequality holds by $\log\log(B/\Delta) \ge 1/4$. For the rest of the proof, suppose that $B/\Delta \ge R_0$. Since the logarithm is increasing, $CN_\Delta \le \log(B/\Delta)$ also gives $\log(CN_\Delta) \le \log\log(B/\Delta)$. Therefore,

$$CN_\Delta \log(CN_\Delta) \le C \frac{c_1 \log(\frac{B}{\Delta})}{\log\log(\frac{B}{\Delta})} \log\log\Big(\frac{B}{\Delta}\Big) = Cc_1 \log\Big(\frac{B}{\Delta}\Big) \le \frac{1}{4}\log\Big(\frac{B}{\Delta}\Big). \tag{259}$$

Using equation 256, we obtain

$$\begin{aligned} \|G^{(d)}_{N_\Delta,z}\|_{\mathcal{H}_{k_{\mathrm{SE}}}(\mathbb{R}^d)} &\le (CN_\Delta)^{CN_\Delta} \\ &= \exp\big(CN_\Delta \log(CN_\Delta)\big) \\ &\le \exp\Big(\frac{1}{4}\log\Big(\frac{B}{\Delta}\Big)\Big) = \Big(\frac{B}{\Delta}\Big)^{1/4}. \end{aligned} \tag{260}$$

Since $B/\Delta \ge R_0 \ge 49$,

$$\Big(\frac{B}{\Delta}\Big)^{1/4} \le \Big(\frac{B}{7\Delta}\Big)^{1/2}, \qquad \frac{B}{7\Delta} \ge 7 \ge 4. \tag{261}$$

Hence,

$$\|G^{(d)}_{N_\Delta,z}\|_{\mathcal{H}_{k_{\mathrm{SE}}}(\mathbb{R}^d)} \le \Big(\frac{B}{7\Delta}\Big)^{1/2}, \qquad \frac{B}{7\Delta} \ge 4. \tag{262}$$

Thus,

$$\begin{aligned} \|2\Delta G^{(d)}_{N_\Delta,z}\|_{\mathcal{H}_{k_{\mathrm{SE}}}(\mathbb{R}^d)} &\le 2\Delta\Big(\frac{B}{7\Delta}\Big)^{\frac{1}{2}} \\ &= 2\Big(\frac{B\Delta}{7}\Big)^{\frac{1}{2}} \le \frac{B}{7}. \end{aligned} \tag{263}$$

It remains to verify localization, bounded overlap, and the size of the family. Choose a sufficiently large fixed separation constant $C_{\mathrm{sep}} \geq 2$. Increase R_0 so that $B/\Delta \geq R_0$ also implies $N_\Delta \geq 2C_{\mathrm{sep}}$, which is possible because $N_\Delta \to \infty$ as $B/\Delta \to \infty$. Set $m_\Delta := \lfloor N_\Delta / C_{\mathrm{sep}} \rfloor \geq 2$ and partition $[0,1]^d$ into m_Δ^d equal cubes with centers z. In each coordinate, take the subintervals to be closed on the left and open on the right, except for the last subinterval, which includes 1. This makes the cubes disjoint and assigns every point of $[0,1]^d$ to exactly one cube. Let R_z denote the cube centered at z and set

$$g_z := 2\Delta G^{(d)}_{N_\Delta, z}. \tag{264}$$

By equation 256 and the definition $g_z = 2\Delta G^{(d)}_{N_\Delta,z}$,

$$0 \leq g_z(x) \leq 2\Delta, \qquad g_z(z) = 2\Delta. \tag{265}$$

Hence $\|g_z\|_\infty = \max_{x\in[0,1]^d} g_z(x) = g_z(z) = 2\Delta$. Since the RKHS norm on $[0,1]^d$ is the infimum over full-space extensions, restriction from $\mathbb{R}^d$ to $[0,1]^d$ cannot increase the norm (Aronszajn, 1950); hence, equation 263 gives $\|g_z\|_{\mathcal{H}_{k_{\mathrm{SE}}}([0,1]^d)} \leq \|g_z\|_{\mathcal{H}_{k_{\mathrm{SE}}}(\mathbb{R}^d)} \leq B/7$. These are exactly the three conditions in equation 194.

We next verify equation 195. Each cube has side length m_Δ^{-1} and center z. If $x \notin R_z$, then some coordinate i satisfies

$$|x_i - z_i| \geq \frac{1}{2m_\Delta}. \tag{266}$$

Since $m_\Delta = \lfloor N_\Delta / C_{\mathrm{sep}} \rfloor \leq N_\Delta / C_{\mathrm{sep}}$,

$$|x_i - z_i| \geq \frac{1}{2m_\Delta} \geq \frac{C_{\mathrm{sep}}}{2N_\Delta}. \tag{267}$$

Multiplying $|x_i - z_i| \geq \frac{C_{\mathrm{sep}}}{2N_\Delta}$ by $N_\Delta > 0$ gives $N_\Delta |x_i - z_i| \geq \frac{C_{\mathrm{sep}}}{2}$. Since $s \mapsto (1+s)^{-2}$ is decreasing on $[0,\infty)$, the decay bound in equation 254 with $r = 2$ gives

$$\begin{aligned} G_{N_\Delta, z_i}(x_i) &\leq C_r (1 + N_\Delta |x_i - z_i|)^{-2} \\ &\leq C_r \Big(1 + \frac{C_{\mathrm{sep}}}{2}\Big)^{-2}. \end{aligned} \tag{268}$$

Every other one-dimensional factor lies in $[0,1]$ by equation 254. Therefore, the product definition of $G^{(d)}_{N_\Delta,z}$ gives

$$G^{(d)}_{N_\Delta,z}(x) \leq C_r \Big(1 + \frac{C_{\mathrm{sep}}}{2}\Big)^{-2}. \tag{269}$$

Choose the fixed constant C_{sep} large enough that

$$2C_r \Big(1 + \frac{C_{\mathrm{sep}}}{2}\Big)^{-2} < 1. \tag{270}$$

Multiplying the product bound by 2Δ then yields

$$g_z(x) = 2\Delta G^{(d)}_{N_\Delta,z}(x) < \Delta, \qquad x \notin R_z, \tag{271}$$

which proves equation 195.

Fix a grid center u, and let R_u be the cube centered at u. Write any other grid center z as $z = u + m_\Delta^{-1} j$ with $j \in \mathbb{Z}^d$. For $x \in R_u$, we have $|x_i - u_i| \leq (2m_\Delta)^{-1}$. If $|j_i| \geq 1$, then $z_i - u_i = j_i / m_\Delta$. The reverse triangle inequality and $|x_i - u_i| \leq (2m_\Delta)^{-1}$ give

$$\begin{aligned} |x_i - z_i| &= \Big| (x_i - u_i) - \frac{j_i}{m_\Delta} \Big| \\ &\geq \frac{|j_i|}{m_\Delta} - |x_i - u_i| \geq \frac{1}{m_\Delta}\Big(|j_i| - \frac{1}{2}\Big). \end{aligned} \tag{272}$$

Multiplying $|x_i - z_i| \geq m_\Delta^{-1}(|j_i| - \frac{1}{2})$ by $N_\Delta > 0$, adding one, and using $N_\Delta/m_\Delta \geq C_{\text{sep}} \geq 2$ give

$$1 + N_\Delta|x_i - z_i| \geq 1 + \frac{N_\Delta}{m_\Delta}\Big(|j_i| - \frac{1}{2}\Big) \geq 1 + |j_i|. \tag{273}$$

The second inequality follows from $\frac{N_\Delta}{m_\Delta}(|j_i| - \frac{1}{2}) \geq 2(|j_i| - \frac{1}{2}) = 2|j_i| - 1 \geq |j_i|$ for $|j_i| \geq 1$. The same bound is immediate when $j_i = 0$. For $|j_i| \geq 1$, the decay bound in equation 254 and $1 + N_\Delta|x_i - z_i| \geq 1 + |j_i|$ give $G_{N_\Delta,z_i}(x_i) \leq C_r(1 + |j_i|)^{-2}$. When $j_i = 0$, the unit bound gives $G_{N_\Delta,z_i}(x_i) \leq 1 = (1 + |j_i|)^{-2}$, so the same upper bound holds for every coordinate after enlarging C_r so that $C_r \geq 1$. Using $g_z(x) = 2\Delta \prod_{i=1}^d G_{N_\Delta,z_i}(x_i)$ and then squaring the coordinate-wise bounds gives

$$\sup_{x \in R_u} |g_z(x)|^2 \leq 4\Delta^2 C_r^{2d} \prod_{i=1}^d (1 + |j_i|)^{-4} \leq C\Delta^2 \prod_{i=1}^d (1 + |j_i|)^{-4}. \tag{274}$$

The constant C absorbs the fixed factor $4C_r^{2d}$. Write $j = (j_1, \ldots, j_d)$. For every integer $K \geq 1$, the distributive law gives

$$\sum_{j \in \{-K,\ldots,K\}^d} \prod_{i=1}^d (1 + |j_i|)^{-4} = \prod_{i=1}^d \Bigg(\sum_{j_i=-K}^{K} (1 + |j_i|)^{-4} \Bigg). \tag{275}$$

All terms are nonnegative, so both sides increase as K increases. Letting $K \to \infty$ gives

$$\sum_{j \in \mathbb{Z}^d} \prod_{i=1}^d (1 + |j_i|)^{-4} = \prod_{i=1}^d \Bigg(\sum_{j_i \in \mathbb{Z}} (1 + |j_i|)^{-4} \Bigg) < \infty. \tag{276}$$

Each one-dimensional sum is finite because its terms decay as a p-series with exponent $4 > 1$. The resulting bound is independent of u and Δ, proving equation 196. Finally, the definition of N_Δ and $m_\Delta \asymp N_\Delta$ give

$$M_\Delta = m_\Delta^d \asymp N_\Delta^d \asymp \left[\frac{\log(\frac{B}{\Delta})}{\log\log(\frac{B}{\Delta})} \right]^d, \tag{277}$$

which proves equation 216. □

D.5 Matérn kernel on a Euclidean box

Use the Fourier convention

$$\widehat{g}(\omega) := \int_{\mathbb{R}^d} g(x) e^{-2\pi i \omega^\top x}\, dx. \tag{278}$$

Under the Fourier convention in equation 278, the normalized Matérn kernel satisfying $k_\nu(x, x) = 1$ has spectral density (Williams & Rasmussen, 2006)

$$\widehat{k}_\nu(\omega) := \frac{\Gamma(\nu + \frac{d}{2})\lambda_0^{2\nu}}{\pi^{\frac{d}{2}}\Gamma(\nu)} \big(\lambda_0^2 + \|\omega\|_2^2\big)^{-(\nu + \frac{d}{2})}, \qquad \lambda_0 > 0. \tag{279}$$

The prefactor in equation 279 makes $\int_{\mathbb{R}^d} \widehat{k}_\nu(\omega)\, d\omega = 1$. With the convention in equation 278, the full-space RKHS norm is

$$\|g\|^2_{\mathcal{H}_{k_\nu}(\mathbb{R}^d)} = \frac{\pi^{\frac{d}{2}}\Gamma(\nu)}{\Gamma(\nu + \frac{d}{2})\lambda_0^{2\nu}} \int_{\mathbb{R}^d} |\widehat{g}(\omega)|^2 \big(\lambda_0^2 + \|\omega\|_2^2\big)^{\nu + \frac{d}{2}}\, d\omega. \tag{280}$$

Lemma 19 (Compactly supported Matérn hard family)**.** *Let $B > 0$ be the RKHS-norm radius in $\mathcal{H}_{k_\nu}(\mathcal{X}, B)$, and let $\Delta > 0$ be the hard-family resolution in equation 194. For sufficiently small $\frac{\Delta}{B}$, there is a localized hard family with*

$$M_\Delta \asymp \Big(\frac{B}{\Delta}\Big)^{\frac{d}{\nu}}. \tag{281}$$

Proof. **Proof outline.** In this proof, a bump is a smooth function that attains its maximum at a chosen center and is zero outside a small neighborhood of that center. The construction selects such a compactly supported profile, controls the RKHS norm of its rescaled copies as a function of their bandwidth, and places disjoint copies on a grid.

Bump construction and optimal set. First, choose a bump with a localized Δ-optimal set. Let $C_c^\infty((-1,1)^d)$ denote the smooth functions compactly supported in $(-1,1)^d$. Fix $\psi \in C_c^\infty((-1,1)^d)$ such that

$$0 \le \psi \le 1, \qquad \psi(0) = 1, \qquad \psi(u) < \frac{1}{2} \quad \text{when } \|u\|_\infty > \frac{1}{2}. \tag{282}$$

For a center z and a bandwidth $w \in (0,1]$, define within this proof

$$g_{z,w}(x) := 2\Delta\psi\Big(\frac{x-z}{w}\Big). \tag{283}$$

The function is nonnegative, bounded by 2Δ, and takes the value 2Δ at z. If $g_{z,w}(x) \ge \Delta$, then equation 282 implies $\|x-z\|_\infty \le \frac{w}{2}$. Thus, the entire Δ-optimal set is contained in $z + [-\frac{w}{2}, \frac{w}{2}]^d$.

RKHS norm of the scaled bump. The scaled bump satisfies the following RKHS norm bound. Let $C_\psi \ge 1$ denote a constant depending only on ψ, d, ν, and λ_0, enlarged when necessary. Substituting the definition of $g_{z,w}$ into equation 278 gives

$$\begin{aligned}
\widehat{g}_{z,w}(\omega) &= 2\Delta \int_{\mathbb{R}^d} \psi\Big(\frac{x-z}{w}\Big) e^{-2\pi i\omega^\top x}\, dx \\
&= 2\Delta w^d \int_{\mathbb{R}^d} \psi(y) e^{-2\pi i\omega^\top (z+wy)}\, dy \\
&= 2\Delta w^d e^{-2\pi i\omega^\top z} \int_{\mathbb{R}^d} \psi(y) e^{-2\pi i(w\omega)^\top y}\, dy.
\end{aligned} \tag{284}$$

For the second equality, set $y = (x-z)/w \in \mathbb{R}^d$. The inverse transformation is $x = z + wy$, whose Jacobian matrix is wI_d. Hence, $dx = |\det(wI_d)|\, dy = w^d\, dy$. By equation 278, the final integral is $\widehat{\psi}(w\omega)$. Therefore,

$$\widehat{g}_{z,w}(\omega) = 2\Delta w^d e^{-2\pi i\omega^\top z} \widehat{\psi}(w\omega). \tag{285}$$

Substituting equation 285 into equation 280 and using $|e^{-2\pi i\omega^\top z}| = 1$ yields the integral over ω in equation 286. Under the change of variables $u = w\omega$, the volume element becomes $d\omega = w^{-d}\, du$. Thus, we obtain

$$\begin{aligned}
\|g_{z,w}\|^2_{\mathcal{H}_{k_\nu}(\mathbb{R}^d)} &\le C_\psi \Delta^2 w^{2d} \int_{\mathbb{R}^d} |\widehat{\psi}(w\omega)|^2 \big(\lambda_0^2 + \|\omega\|_2^2\big)^{\nu+\frac{d}{2}}\, d\omega \\
&= C_\psi \Delta^2 w^d \int_{\mathbb{R}^d} |\widehat{\psi}(u)|^2 \Big(\lambda_0^2 + \frac{\|u\|_2^2}{w^2}\Big)^{\nu+\frac{d}{2}}\, du \\
&\le C_\psi \Delta^2 w^{d-2(\nu+\frac{d}{2})} = C_\psi \Delta^2 w^{-2\nu}.
\end{aligned} \tag{286}$$

For $w \le 1$, factorizing w^{-2} and using $w^2\lambda_0^2 \le \lambda_0^2$ give

$$\lambda_0^2 + \frac{\|u\|_2^2}{w^2} \le w^{-2}\big(\lambda_0^2 + \|u\|_2^2\big). \tag{287}$$

Raising this pointwise bound to the power $\nu + \frac{d}{2}$ contributes the factor $w^{-2(\nu+\frac{d}{2})}$ to the integral in equation 286. Here, $\widehat{\psi}$ denotes the Fourier transform of ψ under the convention in equation 278. Because ψ is smooth, integration by parts can be repeated arbitrarily many times. Its compact support removes all boundary terms and makes every derivative integrable. Hence, for every integer $N \ge 1$,

$$|\widehat{\psi}(u)| = O\big((1 + \|u\|_2)^{-N}\big). \tag{288}$$

Therefore, the integral with integrand $|\widehat{\psi}(u)|^2(\lambda_0^2 + \|u\|_2^2)^{\nu+\frac{d}{2}}$ is finite and independent of w and Δ. Enlarging C_ψ to absorb this integral gives the second inequality in equation 286. Finally, $d - 2(\nu + \frac{d}{2}) = -2\nu$, which gives the last equality in equation 286.

Bandwidth and packing. The bandwidth is chosen to satisfy the RKHS budget and determine the packing size. Choose a fixed c_w such that $C_\psi^{\frac{1}{2}} c_w^{-\nu} \le 1$. Set

$$w_\Delta := c_w \Big(\frac{7\Delta}{B}\Big)^{\frac{1}{\nu}}. \tag{289}$$

The bounds in equation 286 and equation 289 give

$$\|g_{z,w_\Delta}\|^2_{\mathcal{H}_{k_\nu}(\mathbb{R}^d)} \le C_\psi \Delta^2 \Big[c_w \Big(\frac{7\Delta}{B}\Big)^{\frac{1}{\nu}}\Big]^{-2\nu} = C_\psi c_w^{-2\nu} \frac{B^2}{49}. \tag{290}$$

Taking square roots and using $C_\psi^{1/2} c_w^{-\nu} \le 1$ yield

$$\|g_{z,w_\Delta}\|_{\mathcal{H}_{k_\nu}(\mathbb{R}^d)} \le C_\psi^{1/2} c_w^{-\nu} \frac{B}{7} \le \frac{B}{7}. \tag{291}$$

Restriction to $[0,1]^d$ cannot increase the RKHS norm (Aronszajn, 1950).

For Δ/B small enough, $w_\Delta \le \frac{1}{8}$, and hence

$$m_\Delta := \Big\lfloor \frac{1}{4w_\Delta} \Big\rfloor \ge 2. \tag{292}$$

Partition $[0,1]^d$ into m_Δ^d equal cubes. In each coordinate, take the subintervals to be closed on the left and open on the right, except for the last subinterval, which includes 1. This makes the cubes disjoint and assigns every point of $[0,1]^d$ to exactly one cube. Let $\mathcal{Z}_\Delta$ be the set of geometric centers of these cubes, define $M_\Delta := |\mathcal{Z}_\Delta| = m_\Delta^d$, and write R_z for the cube centered at $z \in \mathcal{Z}_\Delta$. For each center, set $g_z := g_{z,w_\Delta}$. Since $m_\Delta \le \frac{1}{4w_\Delta}$, each cell has half-side $1/(2m_\Delta) \ge 2w_\Delta$. Since $\psi(y) = 0$ outside $[-1,1]^d$,

$$g_z(x) \ne 0 \quad \Longrightarrow \quad \frac{x-z}{w_\Delta} \in [-1,1]^d \quad \Longrightarrow \quad x \in z + [-w_\Delta, w_\Delta]^d.$$

Hence, the support of g_z lies in the box $z + [-w_\Delta, w_\Delta]^d$. Since R_z has half-side at least $2w_\Delta$, this box is contained in the interior of R_z. Distinct cell interiors are disjoint, so the support inclusions above make the bumps disjoint. The identities in equation 282 and equation 283 also give

$$g_z(x) \ge \Delta \quad \Longrightarrow \quad \|x - z\|_\infty \le \frac{w_\Delta}{2}.$$

The box $z + [-w_\Delta/2, w_\Delta/2]^d$ lies in R_z. Hence, $x \notin R_z$ implies $g_z(x) < \Delta$, proving equation 195.

Fix a cell R_u. If $z \ne u$, then the support of g_z does not intersect R_u. Since $0 \le \psi \le 1$,

$$\begin{aligned} \sup_{x \in R_u} |g_z(x)| &= 0, \qquad z \ne u, \\ \sup_{x \in R_u} |g_u(x)| &\le 2\Delta. \end{aligned} \tag{293}$$

Therefore,

$$\sum_z \sup_{x \in R_u} |g_z(x)|^2 \le 4\Delta^2. \tag{294}$$

This proves equation 196. Finally, since $\lfloor x \rfloor \ge x - 1$ for every real x and $\frac{1}{4w_\Delta} \ge 2$,

$$\frac{1}{8w_\Delta} \le \frac{1}{4w_\Delta} - 1 \le m_\Delta \le \frac{1}{4w_\Delta}. \tag{295}$$

Therefore, $M_\Delta = m_\Delta^d \asymp w_\Delta^{-d} \asymp (\frac{B}{\Delta})^{\frac{d}{\nu}}$, proving equation 281. □

D.6 Derivation of the cumulative-regret lower bounds

Proof of Theorem 8. Recall from Section D.3 that $M_\Delta := |\mathcal{Z}_\Delta|$ is the number of centers in the localized hard family at resolution Δ, with one function g_z for each center $z \in \mathcal{Z}_\Delta$. For each kernel, $M_{\Delta_T} = |\mathcal{Z}_{\Delta_T}|$ therefore counts the functions in its hard family at the chosen resolution Δ_T.

Proof outline. For each kernel, we choose a resolution Δ_T, verify the assumptions of Lemma 17, and show directly that the right-hand side of equation 198 is larger than T. Lemma 17 then gives a regret gap $T\Delta_T/2$, whose parameter dependence yields the stated lower bound.

SE kernel. For a constant $c_{\mathrm{SE}} > 0$ chosen below, define

$$\Delta_T := c_{\mathrm{SE}} v^{\frac{1}{1+\epsilon}} \Big(\frac{L_\delta}{T}\Big)^{\frac{\epsilon}{1+\epsilon}} \left[\frac{\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big)}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big)}\right]^{\frac{d\epsilon}{1+\epsilon}}. \tag{296}$$

By Lemma 18, fix constants $R_{\mathrm{SE}} > e^e$ and $c > 0$, depending only on d and the SE length scale, such that

$$\frac{B}{\Delta} \geq R_{\mathrm{SE}} \quad \Longrightarrow \quad M_\Delta \geq 2, \qquad M_\Delta - 1 \geq c\left[\frac{\log(\frac{B}{\Delta})}{\log\log(\frac{B}{\Delta})}\right]^d. \tag{297}$$

Take $T > e^e v^{1/\epsilon} L_\delta / B^{(1+\epsilon)/\epsilon}$, so the logarithms in equation 296 are positive. For the hard-family construction and the observation channel, respectively, require

$$\frac{B}{\Delta_T} \geq R_{\mathrm{SE}} \quad \text{and} \quad \Delta_T \leq \frac{v^{\frac{1}{1+\epsilon}}}{128}. \tag{298}$$

To express the packing size in terms of T, also require

$$\log\Big(\frac{B}{\Delta_T}\Big) \geq \frac{\epsilon}{2(1+\epsilon)} \log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big) \quad \text{and} \quad \log\log\Big(\frac{B}{\Delta_T}\Big) \leq 2\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big). \tag{299}$$

For fixed B, v, δ, and c_{SE}, these requirements exclude only finitely many positive integers T. Indeed, equation 296 gives

$$\frac{B}{\Delta_T} = c_{\mathrm{SE}}^{-1} \left(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\right)^{\frac{\epsilon}{1+\epsilon}} \left[\frac{\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big)}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big)}\right]^{-\frac{d\epsilon}{1+\epsilon}}, \tag{300}$$

so $B/\Delta_T \to \infty$ and $\Delta_T / v^{1/(1+\epsilon)} \to 0$ as $T \to \infty$. To check the logarithmic limits, take logarithms in equation 296 to obtain

$$\begin{aligned}\log\Big(\frac{B}{\Delta_T}\Big) &= -\log c_{\mathrm{SE}} + \frac{\epsilon}{1+\epsilon}\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big) - \frac{d\epsilon}{1+\epsilon}\log\left[\frac{\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big)}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big)}\right] \\ &= \frac{\epsilon}{1+\epsilon}\log T + O(\log\log T).\end{aligned} \tag{301}$$

The last equality follows because the logarithm of the ratio inside square brackets is $O(\log\log T)$, whereas

$$\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}} T}{v^{\frac{1}{\epsilon}} L_\delta}\Big) = \log T + O(1). \tag{302}$$

Dividing the numerator and denominator by $\log T$ and using $\log\log T / \log T \to 0$, we obtain, as $T \to \infty$,

$$\frac{\log(B/\Delta_T)}{\log\big(B^{\frac{1+\epsilon}{\epsilon}} T / v^{\frac{1}{\epsilon}} L_\delta\big)} \longrightarrow \frac{\epsilon}{1+\epsilon}. \tag{303}$$

For the second limit, factoring out $\log T$ gives

$$\begin{aligned}\log\Big(\frac{B}{\Delta_T}\Big) &= \log T\left(\frac{\epsilon}{1+\epsilon} + O\Big(\frac{\log\log T}{\log T}\Big)\right) \\ &= \log T\Big(\frac{\epsilon}{1+\epsilon} + o(1)\Big).\end{aligned} \tag{304}$$

Taking another logarithm then gives

$$\log\log\Big(\frac{B}{\Delta_T}\Big) = \log\Big[\log T\Big(\frac{\epsilon}{1+\epsilon} + o(1)\Big)\Big] = \log\log T + \log\Big(\frac{\epsilon}{1+\epsilon}\Big) + o(1), \tag{305}$$

$$\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big) = \log\log T + o(1). \tag{306}$$

Since $\log\log T \to \infty$, the fixed additive term $\log(\epsilon/(1+\epsilon))$ does not affect the second ratio. Thus, as $T \to \infty$,

$$\frac{\log\log(\frac{B}{\Delta_T})}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)} \longrightarrow 1. \tag{307}$$

The logarithmic bounds in equation 299 imply

$$\frac{\log(\frac{B}{\Delta_T})}{\log\log(\frac{B}{\Delta_T})} \geq \frac{\epsilon}{4(1+\epsilon)} \cdot \frac{\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}. \tag{308}$$

Combining equation 297 and equation 308 gives

$$M_{\Delta_T} - 1 \geq c_0\left[\frac{\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}\right]^d, \tag{309}$$

where $c_0 := c[\epsilon/(4(1+\epsilon))]^d$. We use the observation channel with resolution $\Delta = \Delta_T$. The condition $\Delta_T \leq v^{1/(1+\epsilon)}/128$ in equation 298 is equation 181 at this resolution. Substituting $A_\Delta = (v/(128\Delta))^{1/\epsilon}$ into $\pi_\Delta = 32\Delta/A_\Delta$, both defined in equation 180, gives

$$\begin{aligned}\pi_{\Delta_T} &= 32 \cdot 128^{\frac{1}{\epsilon}} v^{-\frac{1}{\epsilon}} \Delta_T^{\frac{1+\epsilon}{\epsilon}} \\ &\leq 32 \cdot 128^{\frac{1}{\epsilon}} v^{-\frac{1}{\epsilon}} \Big(\frac{v^{\frac{1}{1+\epsilon}}}{128}\Big)^{\frac{1+\epsilon}{\epsilon}} \\ &= 32 \cdot 128^{\frac{1}{\epsilon} - \frac{1+\epsilon}{\epsilon}} v^{-\frac{1}{\epsilon}+\frac{1}{\epsilon}} = \frac{32}{128} = \frac{1}{4} \leq 1.\end{aligned} \tag{310}$$

Thus, for every $|\mu| \leq 16\Delta_T$, the probabilities in equation 182 are nonnegative and sum to one. The moment identities in equation 183 and equation 184 also give $\mathbb{E}[Y] = \mu$ and $\mathbb{E}[|Y-\mu|^{1+\epsilon}] \leq v$, so this channel satisfies the required noise-moment bound.

To verify the horizon condition, raising equation 296 to the power $-(1+\epsilon)/\epsilon$ gives

$$\Delta_T^{-\frac{1+\epsilon}{\epsilon}} = c_{\mathrm{SE}}^{-\frac{1+\epsilon}{\epsilon}} v^{-\frac{1}{\epsilon}} \frac{T}{L_\delta}\left[\frac{\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}\right]^{-d}. \tag{311}$$

Using equation 309 and equation 311, the right-hand side of equation 198 satisfies

$$\begin{aligned}\frac{v^{\frac{1}{\epsilon}}}{C_*}(M_{\Delta_T}-1)L_\delta\Delta_T^{-\frac{1+\epsilon}{\epsilon}} &\geq \frac{c_0}{C_*} v^{\frac{1}{\epsilon}} L_\delta \left[\frac{\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}\right]^d \Delta_T^{-\frac{1+\epsilon}{\epsilon}} \\ &= \frac{c_0}{C_*} c_{\mathrm{SE}}^{-\frac{1+\epsilon}{\epsilon}} T.\end{aligned} \tag{312}$$

Choose $c_{\mathrm{SE}} > 0$ so that

$$\frac{c_0}{C_*} c_{\mathrm{SE}}^{-\frac{1+\epsilon}{\epsilon}} > 1. \tag{313}$$

The equation 312 is then strictly larger than T, so equation 198 holds. Lemma 17 gives an admissible instance satisfying

$$\mathbb{P}_T^{f,P}\Big(R_T(f) > \frac{T\Delta_T}{2}\Big) > \delta. \tag{314}$$

Finally, the definition of Δ_T yields

$$\frac{T\Delta_T}{2} = \frac{c_{\mathrm{SE}}}{2} v^{\frac{1}{1+\epsilon}} T^{\frac{1}{1+\epsilon}} L_\delta^{\frac{\epsilon}{1+\epsilon}} \left[\frac{\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}{\log\log\Big(\frac{B^{\frac{1+\epsilon}{\epsilon}}T}{v^{\frac{1}{\epsilon}}L_\delta}\Big)}\right]^{\frac{d\epsilon}{1+\epsilon}}, \tag{315}$$

which proves equation 12 after absorbing the fixed factor $c_{\mathrm{SE}}/2$ into the theorem constant.

Matérn kernel. Let c_w be the fixed constant in equation 289. For a constant $c_{\mathrm{Mat}} > 0$ chosen below, define

$$\alpha := \nu(1+\epsilon) + d\epsilon,$$
$$\Delta_T := c_{\mathrm{Mat}} v^{\frac{\nu}{\alpha}} B^{\frac{d\epsilon}{\alpha}} T^{-\frac{\nu\epsilon}{\alpha}} L_\delta^{\frac{\nu\epsilon}{\alpha}}. \tag{316}$$

The bandwidth used in the Matérn hard-family construction is

$$w_{\Delta_T} := c_w\Big(\frac{7\Delta_T}{B}\Big)^{\frac{1}{\nu}}. \tag{317}$$

The resolution satisfies

$$\frac{\Delta_T}{B} = c_{\mathrm{Mat}}\left(\frac{v^{\frac{1}{\epsilon}}L_\delta}{B^{\frac{1+\epsilon}{\epsilon}}T}\right)^{\frac{\epsilon\nu}{\alpha}} \quad \text{and} \quad \left(\frac{\Delta_T}{v^{\frac{1}{1+\epsilon}}}\right)^{\frac{\alpha}{\epsilon}} = c_{\mathrm{Mat}}^{\frac{\alpha}{\epsilon}} \frac{B^d L_\delta^\nu}{T^\nu v^{\frac{d}{1+\epsilon}}}. \tag{318}$$

We consider horizons T satisfying

$$c_{\mathrm{Mat}}\left(\frac{v^{\frac{1}{\epsilon}}L_\delta}{B^{\frac{1+\epsilon}{\epsilon}}T}\right)^{\frac{\epsilon\nu}{\alpha}} \le \frac{1}{7(8c_w)^\nu} \quad \text{and} \quad c_{\mathrm{Mat}}^{\frac{\alpha}{\epsilon}} \frac{B^d L_\delta^\nu}{T^\nu v^{\frac{d}{1+\epsilon}}} \le 128^{-\frac{\alpha}{\epsilon}}. \tag{319}$$

The first condition ensures the validity of the Matérn hard-family construction (Lemma 19). The second ensures the feasibility of the heavy-tailed channel required by the localized hard-family reduction (Lemma 17). Both left-hand sides decrease to zero as T increases, so only finitely many integers T violate equation 319.

The first inequality in equation 319 and equation 317 give

$$\begin{aligned} w_{\Delta_T} = c_w\Big(7\frac{\Delta_T}{B}\Big)^{\frac{1}{\nu}} &\le c_w\left[7\frac{1}{7(8c_w)^\nu}\right]^{\frac{1}{\nu}} \\ &= c_w\Big(\frac{1}{(8c_w)^\nu}\Big)^{\frac{1}{\nu}} = \frac{1}{8}. \end{aligned} \tag{320}$$

Thus, the construction in Lemma 19 applies. Moreover, we have $m_{\Delta_T} = \lfloor 1/(4w_{\Delta_T})\rfloor \ge 2$ and $m_{\Delta_T} \ge 1/(8w_{\Delta_T})$. Here, the bound $m_{\Delta_T} \ge 1/(8w_{\Delta_T})$ uses $\lfloor x\rfloor \ge x-1 \ge x/2$ for $x \ge 2$ with $x = 1/(4w_{\Delta_T})$. Thus,

$$M_{\Delta_T} = (m_{\Delta_T})^d \ge 2^d \ge 2. \tag{321}$$

Since $M_{\Delta_T} \ge 2$, we have $M_{\Delta_T} - 1 \ge M_{\Delta_T}/2$. Combining this inequality with $M_{\Delta_T} = (m_{\Delta_T})^d$ and $m_{\Delta_T} \ge 1/(8w_{\Delta_T})$ gives

$$\begin{aligned} M_{\Delta_T} - 1 \ge \frac{M_{\Delta_T}}{2} &= \frac{1}{2}(m_{\Delta_T})^d \\ &\ge \frac{1}{2}\Big(\frac{1}{8w_{\Delta_T}}\Big)^d \\ &= \frac{1}{2(8c_w 7^{\frac{1}{\nu}})^d} B^{\frac{d}{\nu}} \Delta_T^{-\frac{d}{\nu}} \\ &= c_1 B^{\frac{d}{\nu}} \Delta_T^{-\frac{d}{\nu}}, \end{aligned} \tag{322}$$

where $c_1 := 1/[2(8c_w 7^{1/\nu})^d]$. The second inequality in equation 319 and the second equality in equation 318 imply

$$\frac{\Delta_T}{v^{\frac{1}{1+\epsilon}}} \leq \frac{1}{128}, \tag{323}$$

which verifies equation 181.

The definition of α gives

$$\frac{1+\epsilon}{\epsilon} + \frac{d}{\nu} = \frac{\nu(1+\epsilon)+d\epsilon}{\epsilon\nu} = \frac{\alpha}{\epsilon\nu}. \tag{324}$$

Combining equation 322 and equation 324, the right-hand side of equation 198 satisfies the complete chain of inequalities

$$\begin{aligned}
&\frac{v^{\frac{1}{\epsilon}}}{C_*}(M_{\Delta_T}-1)L_\delta \Delta_T^{-\frac{1+\epsilon}{\epsilon}} \\
&\geq \frac{c_1}{C_*} v^{\frac{1}{\epsilon}} B^{\frac{d}{\nu}} L_\delta \Delta_T^{-\left(\frac{1+\epsilon}{\epsilon}+\frac{d}{\nu}\right)} \\
&= \frac{c_1}{C_*} v^{\frac{1}{\epsilon}} B^{\frac{d}{\nu}} L_\delta \Delta_T^{-\frac{\alpha}{\epsilon\nu}} \\
&= \frac{c_1}{C_*} c_{\mathrm{Mat}}^{-\frac{\alpha}{\epsilon\nu}} T.
\end{aligned} \tag{325}$$

The last equality follows directly from

$$\Delta_T^{-\frac{\alpha}{\epsilon\nu}} = c_{\mathrm{Mat}}^{-\frac{\alpha}{\epsilon\nu}} v^{-\frac{1}{\epsilon}} B^{-\frac{d}{\nu}} T L_\delta^{-1}, \tag{326}$$

which is obtained by raising equation 316 to the power $-\alpha/(\epsilon\nu)$. Choose $c_{\mathrm{Mat}} > 0$ so that

$$\frac{c_1}{C_*} c_{\mathrm{Mat}}^{-\frac{\alpha}{\epsilon\nu}} > 1. \tag{327}$$

Then, the equation 325 is then strictly larger than T, so equation 198 holds. Lemma 17 gives an admissible instance satisfying

$$\mathbb{P}_T^{f,P}\left(R_T(f) > \frac{T\Delta_T}{2}\right) > \delta. \tag{328}$$

Finally,

$$\begin{aligned}
\frac{T\Delta_T}{2} &= \frac{c_{\mathrm{Mat}}}{2} v^{\frac{\nu}{\alpha}} B^{\frac{d\epsilon}{\alpha}} L_\delta^{\frac{\nu\epsilon}{\alpha}} T^{1-\frac{\nu\epsilon}{\alpha}} \\
&= \frac{c_{\mathrm{Mat}}}{2} v^{\frac{\nu}{\alpha}} B^{\frac{d\epsilon}{\alpha}} L_\delta^{\frac{\nu\epsilon}{\alpha}} T^{\frac{\nu+d\epsilon}{\alpha}},
\end{aligned} \tag{329}$$

where

$$1 - \frac{\nu\epsilon}{\alpha} = \frac{\alpha-\nu\epsilon}{\alpha} = \frac{\nu+d\epsilon}{\alpha}. \tag{330}$$

This proves equation 13 after absorbing the fixed factor $c_{\mathrm{Mat}}/2$ into the theorem constant. □

E EXPERIMENTS

We compare HT-ChainPE with TGP-UCB, ATA-GP-UCB, CA-TGP-UCB, and MoM-GP-UCB on a synthetic Rosenbrock benchmark and two benchmarks based on stock-market and light-sensor data. The synthetic experiments vary dimension, kernel, noise family, and moment parameter. Within each seed and setting, all methods share the same environment and pre-generated observations.

E.1 EXPERIMENTAL SETUP

Objective and candidate set. We use dimensions $d \in \{2, 5, 10, 20, 30\}$ and query points in $[0, 1]^d$. For $x \in [0, 1]^d$, let $z(x) := (15x_1 - 5, \ldots, 15x_d - 5)$ and define

$$L_d(z) := \sum_{i=1}^{d-1} \big[100(z_{i+1} - z_i^2)^2 + (1 - z_i)^2\big]. \tag{331}$$

For each seed s, the candidate set $\mathcal{X}_s$ contains $N = 3{,}000$ points: the optimizer $x^\star = (0.4, \ldots, 0.4)$, 300 local candidates, and 2,699 global candidates generated by randomized Latin hypercube sampling Loh (1996). The local candidates are Gaussian perturbations of $x^\star$ with coordinate-wise standard deviations 0.01, 0.025, 0.05, and 0.1 (75 per scale), clipped to $[0, 1]^d$. This optimum-informed set provides global coverage and finer resolution near $x^\star$. We normalize the reward by

$$a_s := \frac{1}{N-1} \sum_{u \in \mathcal{X}_s \setminus \{x^\star\}} L_d(z(u)) \quad \text{and} \quad f_s(x) := -\frac{L_d(z(x))}{a_s}. \tag{332}$$

Thus $f_s(x^\star) = 0$, and uniform sampling from $\mathcal{X}_s$ has expected per-round regret $(N - 1)/N$. All methods optimize over $\mathcal{X}_s$ and observe noisy rewards; the evaluator uses f_s to compute regret.

Kernels. We use squared-exponential (SE), Matérn-$3/2$, and Matérn-$5/2$ kernels with unit marginal variance and length scale $\ell = \sqrt{d}$ in unit-cube coordinates. Writing $r := \|x - x'\|_2$, these kernels are

$$\begin{aligned} k_{\mathrm{SE}}(x, x') &:= \exp\!\left(-\frac{r^2}{2\ell^2}\right), \\ k_{3/2}(x, x') &:= \left(1 + \frac{\sqrt{3}r}{\ell}\right) \exp\!\left(-\frac{\sqrt{3}r}{\ell}\right), \\ k_{5/2}(x, x') &:= \left(1 + \frac{\sqrt{5}r}{\ell} + \frac{5r^2}{3\ell^2}\right) \exp\!\left(-\frac{\sqrt{5}r}{\ell}\right). \end{aligned} \tag{333}$$

Hyperparameters are fixed within each run and shared across methods. Under the inverse-length-scale convention, $2\pi\lambda_0 = \sqrt{2\nu}/\ell$ for the Matérn kernels.

Heavy-tailed observations: Pareto and Fréchet noise. At round t, the observed reward is $y_t = f_s(x_t) + \eta_t$. We set $p := 1 + \epsilon \in \{1.1, 1.5, 1.9\}$, the shape parameter $\alpha := p + 0.01$, and the scale to one. Let S_t be independent Rademacher variables, independent of the positive variables below. For Pareto noise, independently draw Z_t with $\Pr(Z_t > z) = z^{-\alpha}$ for $z \geq 1$ and set

$$\eta_t^{\mathrm{P}} := S_t \left(Z_t - \frac{\alpha}{\alpha - 1}\right). \tag{334}$$

For Fréchet noise, independently draw U_t uniformly on $(0, 1)$ and set

$$F_t := (-\log U_t)^{-1/\alpha} \quad \text{and} \quad \eta_t^{\mathrm{F}} := S_t \left(F_t - \Gamma\left(1 - \frac{1}{\alpha}\right)\right). \tag{335}$$

Here, $\Pr(F_t \leq z) = \exp(-z^{-\alpha})$ for $z > 0$ and $\mathbb{E}[F_t] = \Gamma(1 - 1/\alpha)$. Since $p < \alpha < 2$, both noises have zero mean, a finite p-th absolute moment, and infinite variance. The factorial design has $2 \times 3 \times 3 \times 5 = 90$ settings.

Table 7: Real-data benchmarks. Rows are observations and columns are arms. Stock uses the same rows for its kernel and rewards; Light uses separate training and test rows.

Benchmark	Query points	Kernel rows	Reward rows
Stock market	29	823	823
Light sensor	41	601	192

Common configuration and moment calibration. The synthetic horizon is $T = 20{,}000$, with nominal confidence parameter $\delta = 0.05$ wherever applicable. For each instance, we supply $B_{\exp} := \max_{x\in\mathcal{X}_s} |f_s(x)|$ as the radius parameter. Let v denote the supplied bound on $\mathbb{E}[|\eta_t|^p]$. For Pareto noise, we numerically evaluate

$$v_{\mathrm{P}} := \alpha \int_1^\infty \left| z - \frac{\alpha}{\alpha - 1} \right|^p z^{-\alpha-1}\, \mathrm{d}z. \tag{336}$$

For Fréchet noise, Minkowski's inequality and $\mathbb{E}[F_t^p] = \Gamma(1 - p/\alpha)$ give the upper bound

$$v_{\mathrm{F}} := \left[\Gamma\Big(1 - \frac{p}{\alpha}\Big)^{1/p} + \Gamma\Big(1 - \frac{1}{\alpha}\Big)\right]^p. \tag{337}$$

HT-ChainPE and the truncation-based baselines receive the payoff-moment bound

$$\bar{v}_{\exp} := (B_{\exp} + v^{1/p})^p, \tag{338}$$

while MoM-GP-UCB receives the central noise-moment bound v.

E.2 Real-world benchmarks

Data sources and finite-arm formulation. Each stock or sensor is a query point. All methods share a fixed empirical kernel within each benchmark, with horizon $T = 1{,}000$ and seeds 1-10. Table 7 summarizes the data.

Stock market. Following Chowdhury & Gopalan (2019), we use 823 adjusted closing-price observations per stock for 29 stocks from the Quandl End-of-Day US Stock Prices dataset, dated January 4, 2016 through April 10, 2019. A query returns a raw price sampled uniformly with replacement from the selected stock's observations. The kernel is the sample covariance of the price columns after each column is centered and divided by its sample standard deviation. The same 823 observations determine the kernel, rewards, empirical means, and calibration quantities.

Light sensor. Following Chowdhury & Gopalan (2019), we use readings from 41 light sensors at the CMU Intelligent Workplace, collected in November 2005. The dataset contains 601 training and 192 test observations per sensor. The kernel is the empirical covariance of the centered, standardized training columns. A query returns a raw test reading sampled uniformly with replacement. Test readings also determine the empirical means and calibration quantities.

Paired rewards and evaluation objective. Let $D \in \mathbb{R}^{m\times K}$ be the raw stock-price matrix or light-sensor test matrix. For each seed, independently sample $I_{t,a}$ uniformly from $\{1, \dots, m\}$ for every round-arm pair and generate

$$Y_{t,a} := D_{I_{t,a},a} \qquad \text{and} \qquad y_t := Y_{t,a_t}. \tag{339}$$

All methods share Y within a seed, so they observe the same reward when selecting the same arm at the same round. No synthetic noise is added. Define

$$\mu_a := \frac{1}{m}\sum_{i=1}^m D_{ia}, \qquad B_{\mathrm{real}} := \max_{1\le a\le K} \mu_a, \qquad f_{\mathrm{real}}(a) := \frac{\mu_a}{B_{\mathrm{real}}}. \tag{340}$$

Algorithms observe raw rewards, while the evaluator uses f_{real}. Each selected arm therefore incurs normalized pseudo-regret $1 - \mu_{a_t}/B_{\mathrm{real}}$.

Real-data calibration. We use $p = 2$ and the pooled empirical second moment

$$M_{2,\mathrm{real}} := \frac{1}{mK} \sum_{i=1}^{m} \sum_{a=1}^{K} D_{ia}^2. \tag{341}$$

We supply B_{real} as the radius proxy and $M_{2,\mathrm{real}}$ as the payoff-moment input to HT-ChainPE and the truncation-based baselines. MoM-GP-UCB receives $M_{2,\mathrm{real}}$ as its central-moment input. The kernel and calibration quantities remain fixed throughout each run.

E.3 Evaluation

Paired evaluation and metric. We report results over seeds 1-10. Within each synthetic seed and setting, methods share $\mathcal{X}_s$, the objective normalization, and the pre-generated noise sequence η_t. Their observed rewards can still differ because they query different points. Real-data methods instead share the round-arm table Y. Let $f = f_s$ for Rosenbrock and $f = f_{\mathrm{real}}$ for real data, and let $x^\star$ maximize f over the corresponding set. We report time-average pseudo-regret

$$\frac{R_t}{t} := \frac{1}{t} \sum_{\tau=1}^{t} \big[f(x^\star) - f(x_\tau) \big]. \tag{342}$$

In total, we run $90 \times 5 \times 10 = 4{,}500$ synthetic and $2 \times 5 \times 10 = 100$ real-data experiments.

Implementation and wall-clock measurement. We use float64 on Linux with an Intel Core i7-13700K CPU, an NVIDIA GeForce RTX 4070 Ti GPU, Python 3.12, NumPy 2.2.6, and PyTorch 2.6.0/CUDA 12.4. Real-data runs use NumPy on CPU; Rosenbrock runs use the configured CPU/CUDA backends. A monotonic timer measures each seed, with CUDA synchronization immediately before and after timed computation when applicable. Timing includes problem construction or data loading, paired observation generation, kernel and design setup, algorithm execution, and regret computation. It excludes imports, CUDA initialization, pre-run cache clearing, and serialization. Algorithm and noise-moment caches are cleared between seeds, with one CPU numerical thread per worker. We report means and sample standard deviations across seeds.

E.4 Experimental results

Rosenbrock with Pareto noise. Figures 2–4 compare the five algorithms under Pareto noise with the SE, Matérn-$3/2$, and Matérn-$5/2$ kernels, respectively. Each figure covers all five dimensions and three moment parameters.

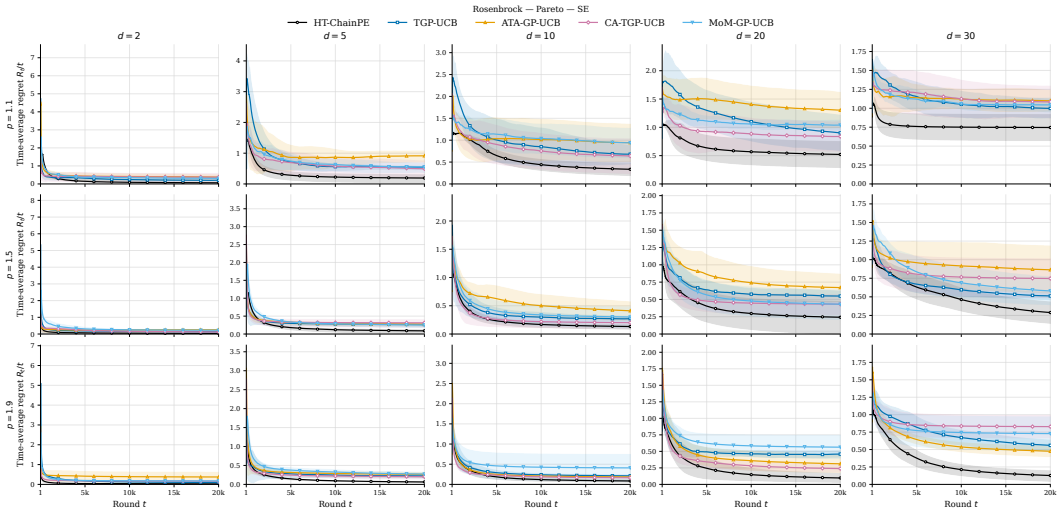


Figure 2: Time-average pseudo-regret on Rosenbrock with Pareto noise and the SE kernel. Columns correspond to $d = 2, 5, 10, 20, 30$, and rows to $p = 1.1, 1.5, 1.9$. Curves and shaded bands show the mean and one sample standard deviation over ten seeds.

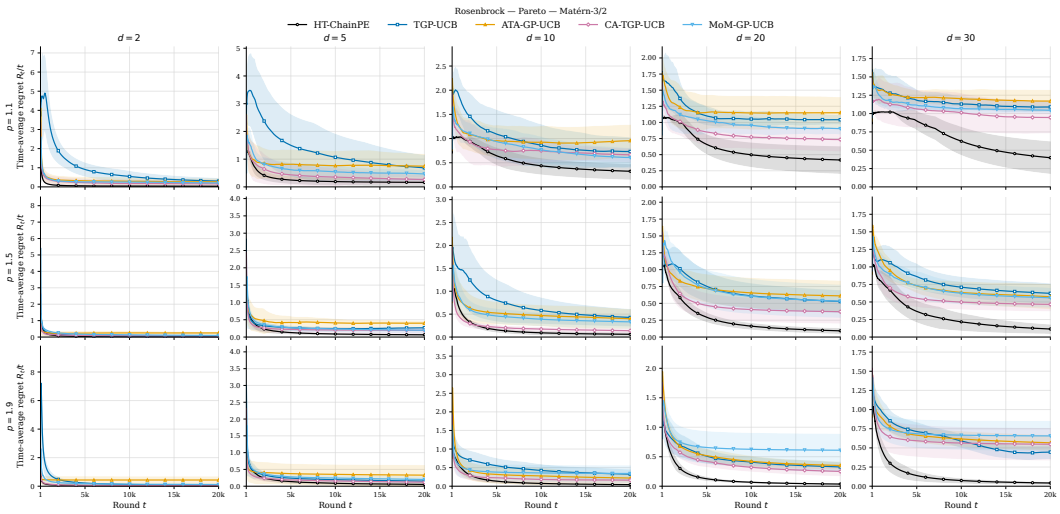


Figure 3: Time-average pseudo-regret on Rosenbrock with Pareto noise and the Matérn-$3/2$ kernel. Columns correspond to $d = 2, 5, 10, 20, 30$, and rows to $p = 1.1, 1.5, 1.9$. Curves and shaded bands show the mean and one sample standard deviation over ten seeds.

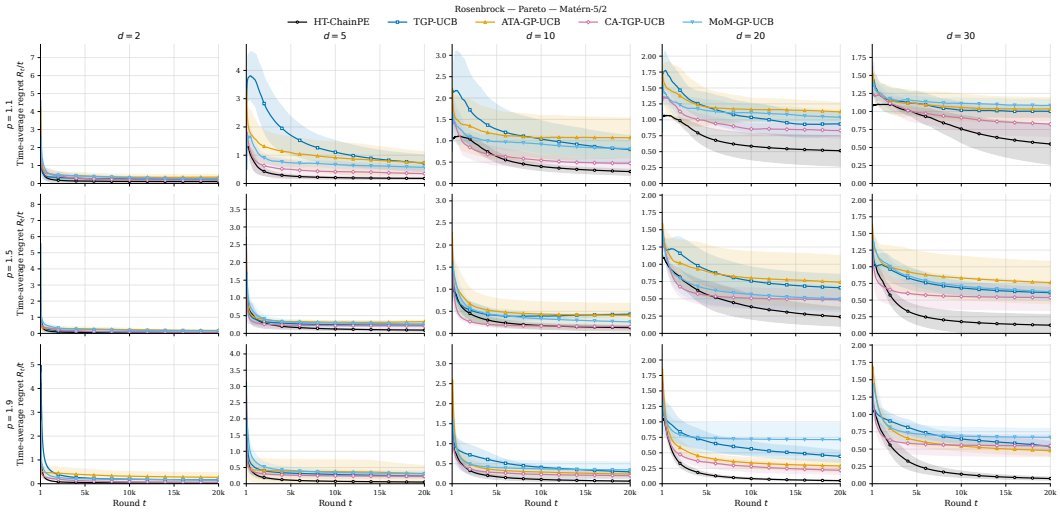


Figure 4: Time-average pseudo-regret on Rosenbrock with Pareto noise and the Matérn-$5/2$ kernel. Columns correspond to $d = 2, 5, 10, 20, 30$, and rows to $p = 1.1, 1.5, 1.9$. Curves and shaded bands show the mean and one sample standard deviation over ten seeds.

Rosenbrock with Fréchet noise. Figures 5–7 compare the five algorithms under Fréchet noise with the SE, Matérn-$3/2$, and Matérn-$5/2$ kernels, respectively. Each figure covers all five dimensions and three moment parameters.

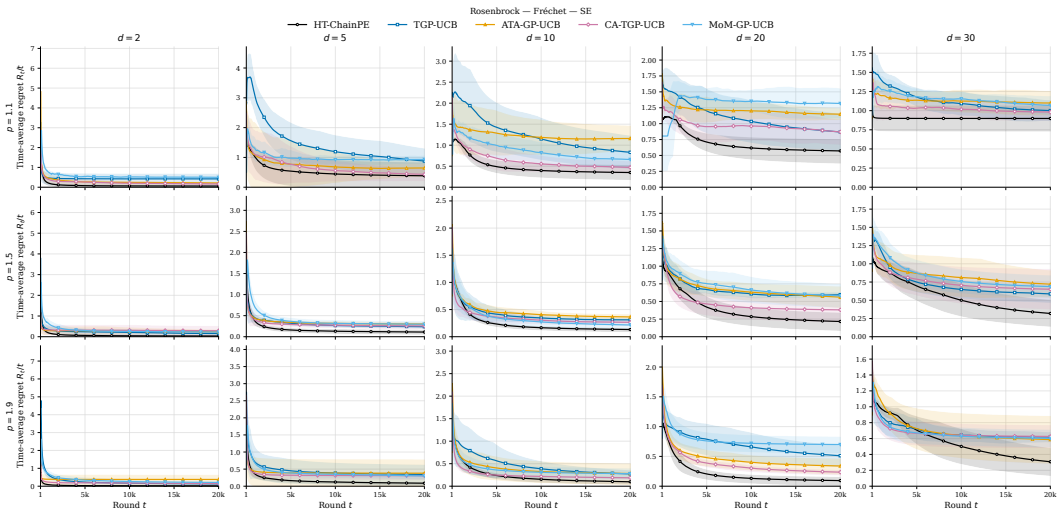


Figure 5: Time-average pseudo-regret on Rosenbrock with Fréchet noise and the SE kernel. Columns correspond to $d = 2, 5, 10, 20, 30$, and rows to $p = 1.1, 1.5, 1.9$. Curves and shaded bands show the mean and one sample standard deviation over ten seeds.

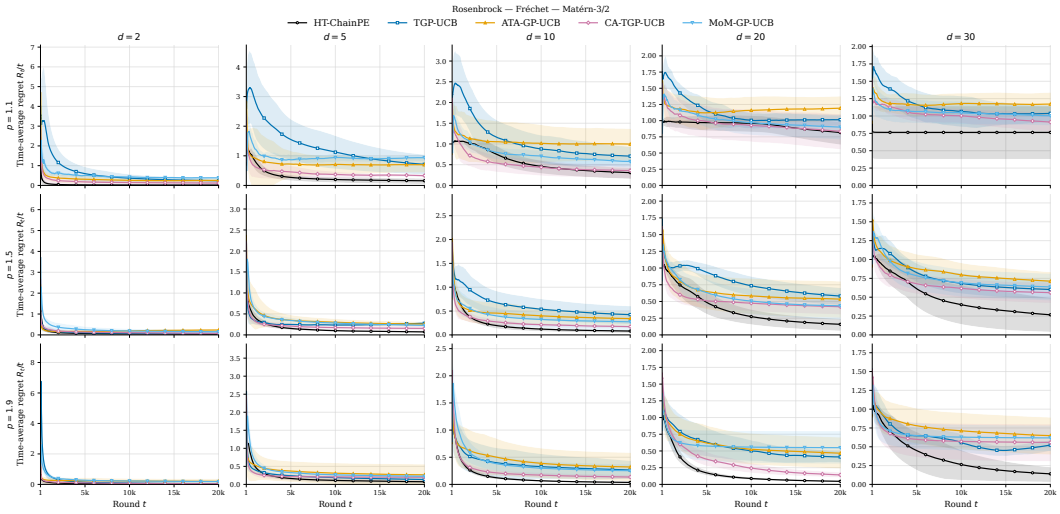


Figure 6: Time-average pseudo-regret on Rosenbrock with Fréchet noise and the Matérn-$3/2$ kernel. Columns correspond to $d = 2, 5, 10, 20, 30$, and rows to $p = 1.1, 1.5, 1.9$. Curves and shaded bands show the mean and one sample standard deviation over ten seeds.

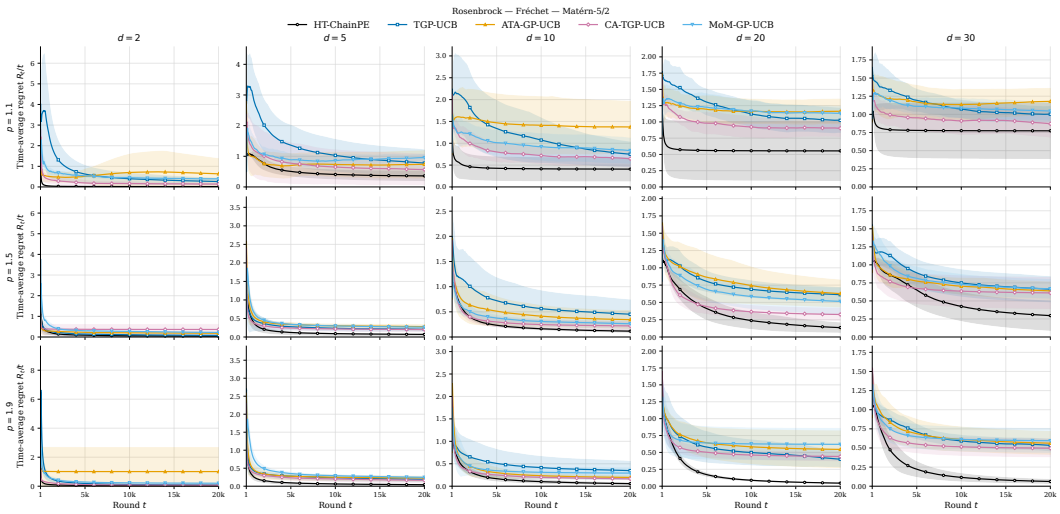


Figure 7: Time-average pseudo-regret on Rosenbrock with Fréchet noise and the Matérn-$5/2$ kernel. Columns correspond to $d = 2, 5, 10, 20, 30$, and rows to $p = 1.1, 1.5, 1.9$. Curves and shaded bands show the mean and one sample standard deviation over ten seeds.

Real-world datasets. Figure 8 reports results on the stock-market and light-sensor benchmarks using rewards sampled from recorded observations.

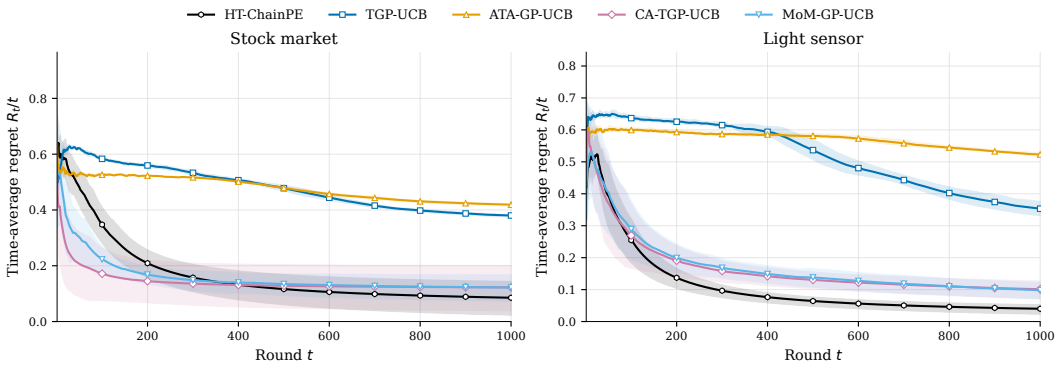


Figure 8: Time-average normalized pseudo-regret on the stock-market (left) and light-sensor (right) benchmarks. Curves and shaded bands show the mean and one sample standard deviation over ten seeds.

E.5 Wall-clock time

Tables 8–10 and Tables 11–13 report the per-seed wall-clock times for Rosenbrock with Pareto and Fréchet noise, respectively. The values come from the runs used for the regret plots. Bold entries indicate the lowest mean time in each row.

Table 8: Per-seed wall-clock time in seconds for Rosenbrock with Pareto noise and the SE kernel ($T = 20{,}000$). Entries show the mean $\pm$ sample standard deviation over seeds 1–10.

p	d	HT-ChainPE	TGP-UCB	ATA-GP-UCB	CA-TGP-UCB	MoM-GP-UCB
1.1	2	$\mathbf{0.740 \pm 0.090}$	133.552 ± 46.098	38.778 ± 8.887	39.711 ± 9.550	0.885 ± 0.210
1.1	5	$\mathbf{0.785 \pm 0.114}$	133.644 ± 46.299	75.844 ± 20.585	72.400 ± 18.998	0.864 ± 0.185
1.1	10	$\mathbf{1.934 \pm 0.239}$	134.560 ± 47.055	128.499 ± 31.574	103.180 ± 39.168	2.060 ± 0.562
1.1	20	3.065 ± 0.483	135.233 ± 47.216	152.220 ± 54.461	130.639 ± 47.078	$\mathbf{2.458 \pm 0.720}$
1.1	30	2.892 ± 0.459	135.784 ± 46.003	274.393 ± 59.120	73.153 ± 20.645	$\mathbf{2.008 \pm 0.572}$
1.5	2	$\mathbf{0.991 \pm 0.273}$	134.719 ± 46.612	30.835 ± 7.153	56.801 ± 21.371	1.142 ± 0.249
1.5	5	1.865 ± 0.312	135.450 ± 46.875	58.740 ± 16.639	61.020 ± 23.644	$\mathbf{1.493 \pm 0.468}$
1.5	10	2.873 ± 0.389	134.751 ± 46.691	111.062 ± 38.395	72.362 ± 27.572	$\mathbf{2.818 \pm 0.653}$
1.5	20	2.843 ± 0.412	134.819 ± 47.048	151.716 ± 52.688	76.170 ± 22.090	$\mathbf{2.377 \pm 0.795}$
1.5	30	3.461 ± 0.436	135.307 ± 46.882	174.069 ± 71.198	60.810 ± 19.626	$\mathbf{1.108 \pm 0.281}$
1.9	2	$\mathbf{1.457 \pm 0.645}$	134.979 ± 46.021	18.570 ± 2.752	53.738 ± 19.467	4.025 ± 0.937
1.9	5	$\mathbf{1.194 \pm 0.410}$	135.333 ± 46.928	53.354 ± 13.666	42.126 ± 13.852	1.253 ± 0.353
1.9	10	$\mathbf{2.856 \pm 0.345}$	134.660 ± 46.941	91.633 ± 27.996	67.097 ± 25.636	9.741 ± 1.070
1.9	20	$\mathbf{2.941 \pm 0.452}$	135.311 ± 46.238	226.443 ± 51.981	79.528 ± 29.508	9.153 ± 0.852
1.9	30	3.230 ± 0.275	135.274 ± 46.968	430.181 ± 115.987	59.093 ± 17.345	$\mathbf{2.260 \pm 0.678}$

Table 9: Per-seed wall-clock time in seconds for Rosenbrock with Pareto noise and the Matérn-$3/2$ kernel ($T = 20{,}000$). Entries show the mean $\pm$ sample standard deviation over seeds 1–10.

p	d	HT-ChainPE	TGP-UCB	ATA-GP-UCB	CA-TGP-UCB	MoM-GP-UCB
1.1	2	$\mathbf{0.821 \pm 0.113}$	135.062 ± 47.169	119.664 ± 50.799	96.740 ± 38.401	2.005 ± 0.617
1.1	5	5.149 ± 1.455	134.983 ± 46.501	343.502 ± 111.230	109.722 ± 33.073	$\mathbf{1.027 \pm 0.277}$
1.1	10	3.299 ± 0.452	134.795 ± 46.858	427.580 ± 106.959	133.889 ± 42.765	$\mathbf{1.059 \pm 0.264}$
1.1	20	3.040 ± 0.422	135.083 ± 46.525	427.122 ± 104.092	134.475 ± 49.079	$\mathbf{2.335 \pm 0.767}$
1.1	30	3.454 ± 0.605	135.417 ± 47.171	918.667 ± 272.782	118.482 ± 39.228	$\mathbf{2.655 \pm 0.703}$
1.5	2	6.516 ± 2.663	267.906 ± 89.482	214.392 ± 67.258	173.332 ± 46.444	$\mathbf{4.013 \pm 0.558}$
1.5	5	12.563 ± 4.725	125.654 ± 40.207	330.942 ± 80.716	58.246 ± 19.368	$\mathbf{2.742 \pm 0.731}$
1.5	10	3.066 ± 0.405	8.842 ± 2.023	57.599 ± 12.907	77.932 ± 27.442	$\mathbf{1.157 \pm 0.199}$
1.5	20	3.101 ± 0.457	16.092 ± 4.362	132.233 ± 62.432	86.950 ± 31.478	$\mathbf{0.801 \pm 0.186}$
1.5	30	$\mathbf{3.161 \pm 0.427}$	10.347 ± 3.342	172.424 ± 69.007	91.793 ± 28.398	10.212 ± 1.229
1.9	2	3.160 ± 2.049	84.665 ± 74.723	25.683 ± 3.995	166.754 ± 53.536	$\mathbf{2.117 \pm 0.606}$
1.9	5	52.482 ± 20.879	4.440 ± 1.010	38.138 ± 6.708	63.792 ± 15.427	$\mathbf{4.137 \pm 0.977}$
1.9	10	$\mathbf{5.461 \pm 2.039}$	6.585 ± 2.378	91.394 ± 30.020	65.882 ± 20.641	10.204 ± 1.321
1.9	20	3.085 ± 0.421	36.100 ± 11.275	568.831 ± 158.813	104.675 ± 33.210	$\mathbf{1.289 \pm 0.321}$
1.9	30	$\mathbf{3.458 \pm 0.496}$	131.128 ± 48.671	758.165 ± 259.477	61.968 ± 14.634	10.366 ± 1.394

Table 10: Per-seed wall-clock time in seconds for Rosenbrock with Pareto noise and the Matérn-$5/2$ kernel ($T = 20{,}000$). Entries show the mean $\pm$ sample standard deviation over seeds 1–10.

p	d	HT-ChainPE	TGP-UCB	ATA-GP-UCB	CA-TGP-UCB	MoM-GP-UCB
1.1	2	$\mathbf{0.832 \pm 0.136}$	135.091 ± 46.947	43.227 ± 15.044	96.365 ± 35.149	1.085 ± 0.349
1.1	5	1.137 ± 0.252	135.464 ± 46.830	179.926 ± 61.967	100.090 ± 38.975	$\mathbf{0.812 \pm 0.216}$
1.1	10	2.861 ± 0.388	134.989 ± 45.946	228.529 ± 62.193	124.638 ± 49.196	$\mathbf{1.181 \pm 0.257}$
1.1	20	3.134 ± 0.459	135.311 ± 47.068	421.321 ± 128.459	154.288 ± 59.543	$\mathbf{2.353 \pm 0.763}$
1.1	30	3.239 ± 0.462	135.382 ± 47.055	374.533 ± 95.960	142.123 ± 54.821	$\mathbf{2.590 \pm 0.792}$
1.5	2	43.778 ± 18.748	134.878 ± 47.224	256.768 ± 188.337	95.927 ± 35.659	$\mathbf{4.028 \pm 0.833}$
1.5	5	174.496 ± 52.111	134.921 ± 46.652	337.552 ± 62.443	61.119 ± 18.123	$\mathbf{4.198 \pm 0.881}$
1.5	10	3.410 ± 0.627	135.345 ± 47.190	61.046 ± 12.987	70.241 ± 26.867	$\mathbf{2.538 \pm 0.793}$
1.5	20	3.100 ± 0.529	133.872 ± 46.135	117.045 ± 46.595	73.461 ± 27.265	$\mathbf{2.110 \pm 0.606}$
1.5	30	3.272 ± 0.444	135.419 ± 46.839	198.914 ± 92.039	71.775 ± 23.508	$\mathbf{2.919 \pm 0.732}$
1.9	2	$\mathbf{3.013 \pm 2.464}$	134.713 ± 47.216	23.423 ± 4.597	67.516 ± 26.812	4.142 ± 0.893
1.9	5	3.895 ± 2.718	134.687 ± 46.653	34.048 ± 6.858	63.881 ± 19.577	$\mathbf{1.171 \pm 0.347}$
1.9	10	3.526 ± 0.501	134.557 ± 46.901	73.511 ± 25.925	66.452 ± 25.041	$\mathbf{1.122 \pm 0.328}$
1.9	20	$\mathbf{3.061 \pm 0.414}$	134.835 ± 47.134	277.807 ± 65.833	100.869 ± 37.515	9.409 ± 1.016
1.9	30	3.200 ± 0.576	134.985 ± 46.569	668.432 ± 198.805	69.347 ± 18.275	$\mathbf{1.291 \pm 0.354}$

Table 11: Per-seed wall-clock time in seconds for Rosenbrock with Fréchet noise and the SE kernel ($T = 20{,}000$). Entries show the mean $\pm$ sample standard deviation over seeds 1–10.

p	d	HT-ChainPE	TGP-UCB	ATA-GP-UCB	CA-TGP-UCB	MoM-GP-UCB
1.1	2	$\mathbf{0.705 \pm 0.110}$	135.247 ± 47.163	35.806 ± 7.945	100.929 ± 33.847	0.843 ± 0.209
1.1	5	$\mathbf{0.838 \pm 0.169}$	135.194 ± 46.983	100.012 ± 33.959	90.960 ± 33.348	0.956 ± 0.258
1.1	10	3.302 ± 3.005	134.796 ± 46.732	141.704 ± 44.051	99.507 ± 37.877	$\mathbf{0.650 \pm 0.124}$
1.1	20	2.894 ± 0.434	133.705 ± 46.257	208.796 ± 52.156	123.699 ± 43.363	$\mathbf{0.462 \pm 0.074}$
1.1	30	2.675 ± 0.394	135.311 ± 46.840	326.657 ± 79.448	107.075 ± 38.595	$\mathbf{0.942 \pm 0.237}$
1.5	2	$\mathbf{0.822 \pm 0.189}$	134.792 ± 46.950	23.642 ± 3.442	51.819 ± 19.305	1.114 ± 0.273
1.5	5	1.533 ± 0.917	134.779 ± 46.654	42.782 ± 7.900	60.304 ± 23.116	$\mathbf{1.110 \pm 0.216}$
1.5	10	$\mathbf{2.687 \pm 0.338}$	135.291 ± 45.897	136.658 ± 34.390	70.533 ± 26.182	10.285 ± 1.272
1.5	20	$\mathbf{2.770 \pm 0.377}$	134.774 ± 46.711	114.278 ± 47.112	80.273 ± 30.385	9.492 ± 0.925
1.5	30	3.045 ± 0.409	133.866 ± 46.451	166.773 ± 62.988	67.864 ± 19.972	$\mathbf{1.011 \pm 0.286}$
1.9	2	1.331 ± 0.656	135.226 ± 46.990	39.710 ± 11.884	54.454 ± 15.794	$\mathbf{0.908 \pm 0.242}$
1.9	5	23.554 ± 9.694	134.718 ± 46.760	29.542 ± 6.119	58.312 ± 22.549	$\mathbf{1.228 \pm 0.270}$
1.9	10	30.372 ± 23.926	135.355 ± 47.263	138.159 ± 46.252	68.522 ± 25.917	$\mathbf{9.935 \pm 1.157}$
1.9	20	2.993 ± 0.470	134.810 ± 46.858	220.134 ± 71.196	90.148 ± 33.684	$\mathbf{0.883 \pm 0.206}$
1.9	30	$\mathbf{3.410 \pm 0.571}$	134.291 ± 46.662	279.829 ± 93.466	68.983 ± 27.816	9.296 ± 0.912

Table 12: Per-seed wall-clock time in seconds for Rosenbrock with Fréchet noise and the Matérn-3/2 kernel ($T = 20{,}000$). Entries show the mean ± sample standard deviation over seeds 1–10.

p	d	HT-ChainPE	TGP-UCB	ATA-GP-UCB	CA-TGP-UCB	MoM-GP-UCB
1.1	2	$\mathbf{0.585 \pm 0.025}$	80.585 ± 18.959	122.770 ± 47.356	104.772 ± 39.778	0.967 ± 0.217
1.1	5	1.778 ± 0.574	134.542 ± 45.766	356.381 ± 84.928	118.565 ± 46.128	$\mathbf{0.782 \pm 0.196}$
1.1	10	2.485 ± 0.061	44.770 ± 0.152	299.861 ± 42.685	40.918 ± 3.151	$\mathbf{0.674 \pm 0.012}$
1.1	20	5.542 ± 1.286	135.526 ± 45.808	665.135 ± 126.893	134.527 ± 47.325	$\mathbf{1.138 \pm 0.289}$
1.1	30	$\mathbf{2.890 \pm 0.428}$	135.203 ± 47.023	804.389 ± 199.043	150.356 ± 49.433	9.617 ± 1.117
1.5	2	$\mathbf{0.939 \pm 0.275}$	135.470 ± 47.009	241.637 ± 84.314	195.598 ± 71.846	1.183 ± 0.309
1.5	5	3.199 ± 0.539	135.279 ± 47.349	53.258 ± 11.506	66.266 ± 25.781	$\mathbf{1.107 \pm 0.251}$
1.5	10	$\mathbf{3.026 \pm 0.403}$	7.847 ± 2.561	102.600 ± 31.824	88.847 ± 26.360	9.278 ± 1.000
1.5	20	$\mathbf{3.086 \pm 0.468}$	135.015 ± 46.555	102.047 ± 33.385	73.509 ± 27.706	8.702 ± 0.663
1.5	30	3.128 ± 0.460	15.311 ± 5.279	228.901 ± 88.626	92.222 ± 23.948	$\mathbf{1.103 \pm 0.249}$
1.9	2	2.293 ± 1.702	103.296 ± 52.793	51.916 ± 17.076	204.054 ± 85.737	$\mathbf{0.929 \pm 0.238}$
1.9	5	33.234 ± 13.708	5.495 ± 1.778	83.920 ± 24.031	59.918 ± 22.690	$\mathbf{1.133 \pm 0.261}$
1.9	10	3.202 ± 0.428	12.428 ± 5.197	141.659 ± 51.170	106.351 ± 41.690	$\mathbf{0.907 \pm 0.222}$
1.9	20	$\mathbf{3.144 \pm 0.624}$	78.619 ± 19.029	429.977 ± 148.699	148.853 ± 55.959	21.066 ± 3.674
1.9	30	$\mathbf{3.391 \pm 0.460}$	135.254 ± 46.552	1285.555 ± 507.272	64.238 ± 22.578	9.882 ± 1.057

Table 13: Per-seed wall-clock time in seconds for Rosenbrock with Fréchet noise and the Matérn-5/2 kernel ($T = 20{,}000$). Entries show the mean ± sample standard deviation over seeds 1–10.

p	d	HT-ChainPE	TGP-UCB	ATA-GP-UCB	CA-TGP-UCB	MoM-GP-UCB
1.1	2	$\mathbf{0.825 \pm 0.112}$	6.155 ± 4.864	51.029 ± 17.321	115.709 ± 43.503	1.132 ± 0.274
1.1	5	7.313 ± 5.693	5.372 ± 0.922	156.298 ± 40.216	87.626 ± 34.024	$\mathbf{1.012 \pm 0.293}$
1.1	10	$\mathbf{3.036 \pm 0.477}$	10.985 ± 2.669	77.591 ± 33.879	141.476 ± 55.307	10.116 ± 1.366
1.1	20	$\mathbf{2.791 \pm 0.376}$	133.999 ± 46.442	400.830 ± 82.017	155.639 ± 53.498	10.024 ± 1.232
1.1	30	3.000 ± 0.437	135.495 ± 46.935	393.395 ± 90.057	107.279 ± 36.983	$\mathbf{0.945 \pm 0.209}$
1.5	2	2.627 ± 1.716	36.695 ± 24.458	61.680 ± 21.235	35.902 ± 11.949	$\mathbf{1.093 \pm 0.252}$
1.5	5	82.776 ± 29.177	6.556 ± 1.129	66.063 ± 15.672	65.852 ± 26.078	$\mathbf{1.165 \pm 0.304}$
1.5	10	$\mathbf{3.509 \pm 1.063}$	7.585 ± 1.711	91.036 ± 26.568	77.891 ± 23.134	9.959 ± 1.119
1.5	20	$\mathbf{3.031 \pm 0.433}$	134.766 ± 46.365	124.876 ± 50.429	84.355 ± 30.001	10.088 ± 1.175
1.5	30	2.900 ± 0.411	17.022 ± 5.429	286.219 ± 74.294	72.130 ± 21.235	$\mathbf{1.216 \pm 0.278}$
1.9	2	29.840 ± 15.995	31.855 ± 22.939	40.466 ± 12.159	97.115 ± 25.309	$\mathbf{1.283 \pm 0.271}$
1.9	5	5.053 ± 2.536	3.740 ± 0.779	41.195 ± 7.975	50.317 ± 16.623	$\mathbf{1.213 \pm 0.285}$
1.9	10	$\mathbf{3.456 \pm 0.667}$	4.710 ± 0.918	120.169 ± 38.372	82.180 ± 29.391	10.239 ± 1.335
1.9	20	$\mathbf{3.223 \pm 0.501}$	9.163 ± 2.983	258.630 ± 108.048	67.686 ± 25.174	9.373 ± 0.810
1.9	30	$\mathbf{3.130 \pm 0.505}$	135.181 ± 46.886	298.147 ± 99.925	76.809 ± 27.437	9.358 ± 0.909