Active Learning for Gaussian Process Regression Under Self-Induced Boltzmann Weights

Jixiang Qing* **Henry Moss*** **Matthias Sachs***

Abstract

We consider the active learning problem where the goal is to learn an unknown function with low prediction error under an unknown Boltzmann distribution induced by the function itself. This self-induced weighting arises naturally in problems such as potential energy surface (PES) modeling in computational chemistry, yet poses unique challenges as the target distribution is unknown and its partition function is intractable. We propose `AB-SID-iVAR`, a Gaussian Process-based acquisition function that approximates the intractable Bayesian target distribution in closed form while avoiding partition function estimation, and is applicable to both discrete and continuous input domains. We also analyze a Thompson sampling alternative (`TS-SID-iVAR`) as a higher variance Monte Carlo variant. Despite the unknown target, under mild conditions, we establish that the terminal prediction error vanishes with high probability, and provide a tighter average-case guarantee. We demonstrate consistent improvements over existing approaches in this setting on synthetic benchmarks and real-world PES modeling and drug discovery tasks. Our code is available at `https://github.com/TsingQAQ/SIDAL`.

1. Introduction

Learning an *unknown function* from data is a fundamental problem in statistics and machine learning. When function evaluations are expensive, sequential experimental design [Pukelsheim, 2006, Fedorov, 2013, Huan et al., 2024] (also known as Active Learning (AL) [Settles, 2009, Rainforth et al., 2024]) provides a principled approach to optimize data acquisition based on a predictive surrogate. Gaussian processes [Rasmussen and Williams, 2006] are a popular choice ([Huang et al., 2026, Smith et al., 2023, Qing et al., 2022]) with convergence guarantees under mild regularity conditions [Krause et al., 2008, Srinivas et al., 2009, Vakili et al., 2021].

Given a surrogate, an *acquisition function* characterizes the utility criterion motivated by the learning objective, such as minimizing prediction error over input space weighted by a *target distribution* P, with uniform distribution as the most common scenario. Under this scheme, available knowledge about P naturally shapes the acquisition strategy: when P is fully specified, weighted uncertainty reduction applies [Seo et al., 2000, Kirsch et al., 2021, Smith et al., 2023]; when P is unknown, one may instead spec-

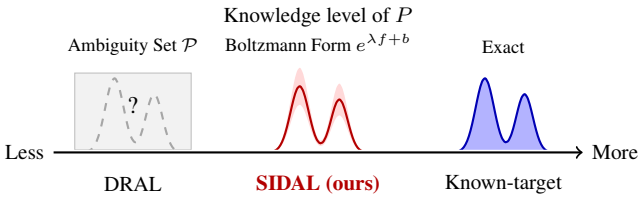


Figure 1: Spectrum of AL methods by knowledge of the target distribution P_f, where the goal is to minimize prediction error $\mathbb{E}_{x\sim P_f}[(f(x)-\hat{f}(x))^2]$. **Left**: DRAL assumes P_f lies in an ambiguity set. **Middle**: SIDAL (ours) knows P_f takes Boltzmann form $\exp(\lambda f+b)$ but f itself is unknown. **Right**: Known-target has exact P_f.

*MARS: Mathematics for AI in Real-world Systems, School of Mathematical Sciences, Lancaster University, Lancaster, LA1 4YF, United Kingdom. Correspondence to: `{j.qing, m.sachs}@lancaster.ac.uk`

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

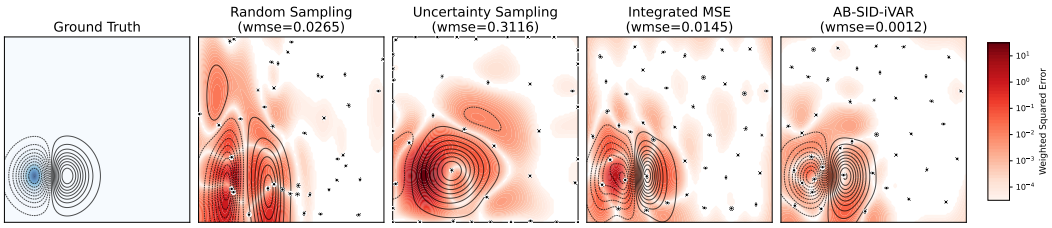


Figure 2: Comparison of active learning methods on the (continuous input) Gramacy 2D function [Gramacy and Lee, 2008], aiming to minimize prediction error under a self-induced Boltzmann distribution $P_f \propto \exp(-f)$. Left: ground truth with high-probability regions shaded in blue; accurate prediction is mainly required in these regions. Right: weighted squared error (heatmap, log scale) after 40 queries (black dots: •); contours show the learned GP posterior mean. SID-unaware baselines include Random Sampling (RS), Uncertainty Sampling (US), and Integrated MSE (IMSE). Our method `AB-SID-iVAR` balances exploration and concentration by placing samples in high-probability regions, achieving substantially lower weighted MSE (wmse) than baselines.

ify possible P through an ambiguity set $\mathcal{P}$ under a Distributionally Robust Active Learning (DRAL) framework [Frogner et al., 2021, Takeno et al., 2025].

Viewing the aforementioned variants on the unified spectrum (Figure 1) of knowledge about P, we argue that many applications naturally fall between established methods. For instance, in molecular modeling and free energy estimation, the distribution of interest takes the Boltzmann form [Frenkel and Smit, 2023] weighted by the unknown energy function itself. We categorize this setting as **Active Learning under a Self-Induced Distribution (SIDAL)**, where the target distribution P depends on the unknown function f. Specifically, P often takes a Boltzmann form (Assumption 3.1). To the best of our knowledge, this setting has not been formally studied in the active learning literature. While related heuristics have emerged in scientific domains, they require careful tuning and lack theoretical guarantees; we defer a detailed discussion to Appendix A.

Contributions Aiming to fill the aforementioned gap, we propose an SID-aware AL algorithm under the GP surrogate model, which samples to reduce the weighted mean-squared-error (MSE) sequentially. Building upon the theoretical analysis of [Srinivas et al., 2009, Cai and Scarlett, 2024, Takeno et al., 2025],we provide sublinear convergence guarantees for the terminal prediction error in both high-probability and average (a.k.a., in-expectation) forms. Our contributions can be summarized as follows:

- We formalize the novel problem of active learning under self-induced distributions (Section 3.1), capturing scenarios where the target distribution depends on the unknown function itself.
- We propose a simple-to-implement acquisition function `AB-SID-iVAR` (fully defined by Eq. 3 and Table 1) that is free of partition function estimation and applies to both discrete and continuous input domains. A Thompson sampling variant (`TS-SID-iVAR`) offering a bias-variance tradeoff is also analyzed.
- We provide sublinear convergence guarantees under high-probability and average settings (Section 4.2), explicitly accounting for Monte Carlo approximation error in continuous domains, and show (Section 4.3) that existing heuristics admit convergence analysis under our framework.
- We demonstrate competitive empirical performance on both synthetic (Figure 2, 3) and real-world problems in PES modeling and drug discovery (Figure 4, 5).

2. Preliminaries

2.1. Gaussian Processes

We model the unknown function using Gaussian Process (GP) regression, a natural choice for active learning due to its closed-form posterior updates and calibrated uncertainty estimates. Given a

positive semi-definite kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, which is normalized as $k(x, x) \leq 1$, and data $\mathcal{D}_t = \{(x_i, y_i)\}_{i=1}^t$, the predictive distribution is:

$$\mu_t(x) = k(x, \mathbf{X}_t)(K_t + \tau^2 I)^{-1}\mathbf{y}_t,\ k_t(x, x') = k(x, x') - k(x, \mathbf{X}_t)(K_t + \tau^2 I)^{-1}k(\mathbf{X}_t, x'), \quad (1)$$

where $\mathbf{X}_t = (x_1, \ldots, x_t)$ denotes the sequence of queried inputs, $\mathbf{y}_t = (y_1, \ldots, y_t)^\top$ the corresponding noisy observations (subjecting to zero-mean Gaussian observation noise with variance τ^2), $K_t = k(\mathbf{X}_t, \mathbf{X}_t) \in \mathbb{R}^{t \times t}$ the kernel (Gram) matrix, and $\sigma_t^2(x) := k_t(x, x)$ the predictive variance at x given the data $\mathcal{D}_t$.

2.2. Regularity Assumptions

We assume the regularity of f is consistent with the kernel k. We primarily consider the Bayesian setting where $f \sim \mathcal{GP}(0, k)$, and follow this assumption throughout unless stated otherwise.

Assumption 2.1 (Kernel Regularity)**.** The kernel k is stationary and four times differentiable.

This is satisfied by common kernels such as squared exponential and Matérn-ν with $\nu > 2$. Under this assumption, sample paths of $f \sim \mathcal{GP}(0, k)$ are Lipschitz continuous with high probability [Ghosal and Roy, 2006, Srinivas et al., 2009], and the posterior mean μ_t is similarly Lipschitz [Vakili et al., 2021, Takeno et al., 2025]. Together, these ensure the *prediction error* $f(x) - \mu_t(x)$ is Lipschitz:

Lemma 2.2 (Prediction Error Lipschitz Continuity [Takeno et al., 2025])**.** *Under Assumption 2.1 with compact input domain $\mathcal{X} \subset [0, r]^d$, $x \mapsto f(x) - \mu_t(x)$ is L-Lipschitz with probability at least $1 - \delta$, where $L = b\sqrt{\log(4ad/\delta)}$ for kernel-dependent constants a, b.*

This enables a high-probability uniform prediction error guarantee using GP as a model of f for continuous input:

Lemma 2.3 (GP Confidence Bound)**.** *Pick $\delta \in (0, 1)$ and $t \in \mathbb{N}$. Then with probability at least $1 - \delta$, for all $x \in \mathcal{X}$:*

$$|f(x) - \mu_{t-1}(x)| \leq \beta_t^{1/2}\sigma_{t-1}(x) + \epsilon_t,$$

where for finite $|\mathcal{X}| < \infty$: $\beta_t = 2\log(|\mathcal{X}|/\delta)$, $\epsilon_t = 0$; for continuous $\mathcal{X} \subset [0, r]^d$ under Assumption 2.1: $\beta_t = 2d\log(tdr + 1) + 2\log(2/\delta)$, $\epsilon_t = L/t$.

The non-negative ϵ_t in the continuous case is induced by approximating $f - \mu_{t-1}$ on $[0, r]^d$ using a Cartesian grid of spacing $1/(dt)$, and β_t grows logarithmically with t due to a union bound over $(tdr + 1)^d$ grid points. The proof follows from Srinivas et al. [2009] Lemma 5.1 for discrete $\mathcal{X}$, and Takeno et al. [2025] when using the Lipschitz constant L from Lemma 2.2.

3. Problem Setting and Performance Analysis

We query f to obtain noisy observations $y = f(x) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \tau^2)$. Given a *surrogate model* $\hat{f}(x)$ built from the current dataset $\mathcal{D}_t$, our goal is to sequentially select T query points that minimize the expected squared error under a distribution induced by f itself, which we term the *Self-Induced Distribution* (SID) P_f:

$$\mathcal{L}(\hat{f}) = \mathbb{E}_{x \sim P_f}\left[\left(\hat{f}(x) - f(x)\right)^2\right], \quad (2)$$

where we have the following assumption regarding the SID:

Assumption 3.1 (Boltzmann Self-Induced Distribution)**.** The target distribution $P_f(x)$ takes the Boltzmann distribution form[2] $P_f(x) = Z_f^{-1}\exp(\lambda f(x) + b(x))$ where: (i) $\lambda \neq 0$; (ii) the *bias* function $b : \mathcal{X} \to \mathbb{R}$ is known and bounded; (iii) the *partition function* $Z_f < \infty$.

We use the GP predictive mean $\hat{f}(x) = \mu_t(x)$ as our surrogate. While μ_t is not the Bayes-optimal estimator for Eq. 2 (which otherwise requires intractable integrations), this gap vanishes as the posterior concentrates (Appendix C.1).

[2]We parameterize using λ directly; this corresponds to $\lambda = -1/\tilde{\mathcal{T}}$ in the physical Boltzmann distribution with temperature parameter $\tilde{\mathcal{T}} > 0$.

Different from standard AL or DRAL as elaborated in Figure 1, one key challenge here is that the SID P_f itself depends on the unknown function f, creating a circular dependence dilemma. To establish that learning under SID is feasible, we derive upper bounds on $\mathcal{L}(\hat{f})$ by relating the unknown P_f to the computable surrogate P_{μ_t}. These upper bounds, presented in both high-probability and average forms, inform our algorithm design in the next section.

To first establish a high-probability bound, we decompose the density ratio between P_f and P_{μ_t} into a pointwise Boltzmann factor difference and a partition function difference:

Lemma 3.2 (Log-Density Ratio)**.** *Let $g_1, g_2 : \mathcal{X} \to \mathbb{R}$ and define $P_{g_i} \propto e^{\lambda g_i + b}$ for $i = 1, 2$. Define $\Delta(x) := |g_1(x) - g_2(x)|$ and $\Delta_{\max} := \sup_{x \in \mathcal{X}} \Delta(x)$. Then:*

$$|\log P_{g_1}(x) - \log P_{g_2}(x)| \leq |\lambda|(\Delta(x) + \Delta_{\max}).$$

We state this lemma abstractly for reuse later (in Lemma C.4). Substituting $(g_1, g_2) = (f, \mu_t)$ and combining with the confidence bound (Lemma 2.3), we convert expectations under the unknown P_f to the computable P_{μ_T}:

Lemma 3.3 (High-Probability MSE to Variance Bound)**.** *Pick $\delta \in (0, 1)$. Under Assumption 3.1, let β_T, ϵ_T be as in Lemma 2.3. Then with probability at least $1 - \delta$:*

$$\mathcal{L}(\mu_T) \leq e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T} + \epsilon_T)} \cdot \left(\beta_T \cdot \mathbb{E}_{x \sim P_{\mu_T}}\left[\sigma_T^2(x)\right] + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right),$$

where $\sigma_{\max,T} = \max_{x \in \mathcal{X}} \sigma_T(x)$.

The exponential prefactor arises from bounding the partition function ratio, which requires uniform control of $f - \mu_T$ over all of $\mathcal{X}$ and is challenging to avoid in high-probability analysis [Cai and Scarlett, 2024]. As we shall see in Section 4.2, while convergence is unaffected, this introduces additional logarithmic factors in the terminal prediction error bound. To circumvent estimating the partition function ratio entirely, we take expectation over $f|\mathcal{D}_T$, which allows Z_f/Z_{μ_T} to be absorbed via Jensen's inequality. The Gaussian posterior structure then admits closed-form control via Moment Generating Function (MGF) identities, yielding the following average upper bound:

Lemma 3.4 (Posterior-Averaged MSE to Variance Bound)**.** *Under Assumption 3.1, the terminal MSE averaged via posterior $f|\mathcal{D}_T$ satisfies:*

$$\mathbb{E}_{f|\mathcal{D}_T}\left[\mathbb{E}_{P_f}\left[(f - \mu_T)^2\right]\right] \leq e^{2\lambda^2}(1 + 4\lambda^2) \cdot \mathbb{E}_{P_{\mu_T}}[\sigma_T^2(x)].$$

Without the β_T factor, this yields a tighter terminal error bound that better reflects finite-query and average-case performance. Notably, both upper bounds reflect the intrinsic difficulty of the problem: a more concentrated target distribution makes accurate prediction in high-probability regions more challenging.

4. Methodology

4.1. Proposed Algorithms

Both upper bounds involve weighted posterior variance, suggesting variance reduction as the algorithmic objective. To this end, our proposed integrated Variance Reduction (iVAR) minimizes expected posterior variance over a tractable unnormalized SID surrogate $\tilde{p}_{t-1}^{\mathrm{u}}$ (we use $\tilde{p}_{t-1}$ for its normalized version):

$$x_t = \arg\min_{x \in \bar{\mathcal{X}}_t} \int_{x^\star \in \mathcal{X}} \tilde{p}_{t-1}^{\mathrm{u}}(x^\star)\sigma_t^2(x^\star \mid x_t = x)\, dx^\star, \tag{3}$$

where a *potential set* $\bar{\mathcal{X}}_t$ restricts the search to uncertain regions to ensure sufficient exploration. We propose `AB-SID-iVAR`, with the surrogate $\tilde{p}_{t-1}^{u}$ defined in Table 1.

`AB-SID` (Approximate Bayesian SID) approximates the intractable Bayesian SID $\mathbb{E}_{f|\mathcal{D}_{t-1}}[P_f]$ in closed form via a zero-order Taylor expansion (Appendix B.1).[3] Alternatively, one can approximate

[3] A first-order correction is available (Appendix B.1.1) with marginally better accuracy (Figure 6), but requires $\mathcal{O}(|\mathcal{X}|^2)$ additional covariance computations for discrete $\mathcal{X}$.

Table 1: SID-iVAR variants. **AB-SID**: closed-form zero-order Taylor approximation of $\mathbb{E}_{f|\mathcal{D}_{t-1}}[P_f]$ (Appendix B.1.1); biased, low-variance. **TS-SID**: single posterior sample as MC estimate; unbiased, high-variance. Both share constraint set $\bar{\mathcal{X}}_t$.

	AB-SID	**TS-SID**
$\tilde{p}^{\mathrm{u}}_{t-1}$	$e^{\lambda\mu_{t-1}+\frac{\lambda^2}{2}\sigma^2_{t-1}+b}$	$e^{\lambda\tilde{f}_{t-1}+b}$, $\tilde{f}_{t-1} \overset{\text{draw}}{\sim} f\|\mathcal{D}_{t-1}$
$\bar{\mathcal{X}}_t = \{x : \sigma^2_{t-1}(x) \geq \mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}]\}$		

Algorithm 1 AB/TS-SID-iVAR

Require: $\mathcal{X}$, GP prior $(0, k)$, budget T, $\mathcal{D}_0$
1: **for** $t = 1, \ldots, T$ **do**
2: Compute μ_{t-1}, σ^2_{t-1} from $\mathcal{D}_{t-1}$
3: Construct $\tilde{p}^{\mathrm{u}}_{t-1}$ via AB or TS-SID
4: Construct $\bar{\mathcal{X}}_t$ (Table 1)
5: $x_t \leftarrow \arg\min_{x\in\bar{\mathcal{X}}_t} \int \tilde{p}^{\mathrm{u}}_{t-1}(x^\star)\sigma^2_t(x^\star \mid x)\, dx^\star$
6: Observe $y_t = f(x_t) + \epsilon_t$; update $\mathcal{D}_t$
7: **end for**
8: **return** μ_T

this integral via a single Monte Carlo sample, yielding `TS-SID` (Thompson Sampling SID), which offers a bias-variance tradeoff (Table 1). Both methods avoid estimation of the partition function.

While both algorithms apply directly in discrete domains, continuous $\mathcal{X} \subset [0, r]^d$ requires approximating both the constraint set $\bar{\mathcal{X}}_t$ and the integral in Eq. 3. We use MCMC to sample from $\tilde{p}^{\mathrm{u}}_{t-1}$ and approximate the integral; for `TS-SID`, we additionally approximate sample trajectories $\tilde{f}_{t-1}$ via pathwise conditioning [Wilson et al., 2021].[4] Importantly, our theoretical analysis is agnostic to the specific MCMC sampler (provided that the variance-threshold estimator satisfies the conditional unbiasedness and concentration conditions specified in Appendix C.5).

Algorithm Summary. The overall procedure is summarized in Algorithm 1. The key design choices are: (1) the SID surrogate's form $\tilde{p}_{t-1}$ that approximates the unknown target distribution, and (2) the potential set $\bar{\mathcal{X}}_t$ that ensures exploration of uncertain regions. For continuous $\mathcal{X}$, the integral in Line 5 and the threshold in $\bar{\mathcal{X}}_t$ are approximated via MCMC samples from $\tilde{p}_{t-1}$.

4.2. Theoretical Analysis

We now show that iteratively optimizing over the surrogate $\tilde{p}_{t-1}$ yields vanishing prediction error under the true SID P_f, and establish finite-sample upper bounds on the terminal prediction error $\mathcal{L}(\mu_T)$, as well as the rate at which it converges to zero as $T \to \infty$. Our analysis depends on the maximum information gain $\gamma_T = \max_{x_1,\ldots,x_T} I(y_{1:T}; f)$, which is sublinear in T for standard kernels [Srinivas et al., 2009]. We state results for continuous $\mathcal{X}$ directly, as discrete $\mathcal{X}$ is a simplified case with $\epsilon_T = \epsilon_{\mathrm{MC}} = 0$.

High-Probability Guarantee

Theorem 4.1 (High-Probability Terminal MSE Bound)**.** *Under Assumptions 2.1 and 3.1, let $\mathcal{X} \subset [0, r]^d$ be compact. Consider `AB-SID` or `TS-SID` selecting $x_t \in \bar{\mathcal{X}}_t$ (Table 1), where the constraint threshold at each iteration is estimated via M Monte Carlo samples. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$:*

$$\mathcal{L}(\mu_T) \leq e^{2(\kappa+2)|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\beta_T\gamma_T}{T} + \beta_T\epsilon_{\mathrm{MC}} + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right), \tag{4}$$

where $\kappa = 1$ for `AB-SID`, $\kappa = 2$ for `TS-SID`, $C_1 = 2/\log(1 + \tau^{-2})$, β_T, ϵ_T are the coefficients of the uniform confidence event defined in Appendix C.5 with total GP failure budget $\delta/2$, and the cumulative MC error $\epsilon_{\mathrm{MC}} = \frac{2\log(4/\delta)}{3T} + \sqrt{\frac{\log(4/\delta)}{2MT}}$. For finite $\mathcal{X}$ with exact threshold computation, $\epsilon_T = \epsilon_{\mathrm{MC}} = 0$.

The upper bound in Eq. 4 consists of three components: (i) the information gain term $\frac{\gamma_T}{T}$; (ii) a cumulative MC error ϵ_{MC} from approximating the constraint threshold via MC sampling; and (iii) terms involving ϵ_T from extending pointwise GP confidence bounds to the entire continuous domain. The probability δ is combined with the GP confidence event and the MC concentration event via a union bound. Unlike prior work [Takeno et al., 2025] which assumes exact computation, our analysis

[4] The constraint set $\bar{\mathcal{X}}_t$ is always non-empty under the approximation.

explicitly accounts for this MC error. The exponential prefactor, as elaborated earlier, arises from bounding the density ratio between P_f and P_{μ_T}, with `TS-SID` incurring a larger factor ($\kappa = 2$) due to the additional Thompson sampling approximation.

Further analyzing each term in the upper bound yields the following convergence rate.

Corollary 4.2 (High-Probability Terminal MSE Convergence Rate)**.** *Under the conditions of Theorem 4.1, suppose that $\gamma_T = \mathcal{O}((\log T)^q)$ for some fixed $q \geq 0$. Then: (i) M is constant $\Rightarrow$ $\mathcal{L}(\mu_T) = \mathcal{O}(T^{-1/2+\epsilon})$; (ii) $M = \Omega(T) \Rightarrow \mathcal{L}(\mu_T) = \mathcal{O}(T^{-1+\epsilon})$, for arbitrarily small $\epsilon > 0$. The ϵ absorbs the polylogarithmic factors and the subpolynomial prefactor $e^{\mathcal{O}(\sqrt{\log T})} = o(T^{\epsilon})$.*

The assumed form of γ_T holds, for example, for the squared-exponential kernel on bounded domains. The two cases reflect a trade-off between computational cost and convergence rate: with constant M, the cumulative MC error term $\sqrt{\frac{\log(4/\delta)}{2MT}} = \mathcal{O}(1/\sqrt{T})$ dominates; increasing M tightens the bound but does not change the $T^{-1/2}$ rate. With $M = \Omega(T)$, the MC error becomes $\mathcal{O}(1/T)$, recovering the near-optimal rate. In both cases, the bound vanishes as $T \to \infty$: although the prefactor $e^{\mathcal{O}(\sqrt{\log T})}$ grows without bound, it grows slower than any polynomial, so the polynomial decay in T still dominates.

Average Guarantee To complement the high-probability bound on realized terminal MSE, we provide a posterior-averaged MSE bound with a Bayesian terminal MSE bound.

Theorem 4.3 (High-Probability Posterior-Averaged Terminal MSE Bound)**.** *Under Assumptions 2.1 and 3.1, let $\mathcal{X} \subset [0, r]^d$ be compact and consider `AB-SID` or `TS-SID` selecting $x_t \in \bar{\mathcal{X}}_t$ (Table 1), where the constraint threshold at each iteration is estimated via M Monte Carlo samples. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$:*

$$\mathbb{E}_{f|\mathcal{D}_T}\left[\mathbb{E}_{P_f}\left[(f - \mu_T)^2\right]\right] \leq e^{2\lambda^2}(1 + 4\lambda^2)e^{2(1+\kappa)|\lambda|(\beta_T^{1/2} + \epsilon_T)} \cdot \left(\frac{C_1 \gamma_T}{T} + \epsilon_{\mathrm{MC}}\right), \tag{5}$$

where κ, β_T, ϵ_T, and ϵ_{MC} are as in Theorem 4.1, and the outer expectation is over terminal posterior $f|\mathcal{D}_T$. For finite $\mathcal{X}$ with exact threshold computation, $\epsilon_T = \epsilon_{\mathrm{MC}} = 0$.

Here, δ denotes the failure probability over the random data-acquisition process, as in Theorem 4.1 As detailed in Appendix C.7, the average upper bound achieves the same asymptotic rates as Corollary 4.2, but with tighter logarithmic factors. Specifically, the main term is γ_T/T rather than $\beta_T \gamma_T/T$, and the prefactor exponent is also reduced from $2(\kappa + 2)|\lambda|$ to $2(\kappa + 1)|\lambda|$.

Corollary 4.4 (Bayesian Terminal MSE Bound)**.** *Under the conditions of Theorem 4.3, for every $\delta \in (0, 1)$, the following Bayesian terminal MSE bound holds:*

$$\mathbb{E}_f\left[\mathbb{E}_{\mathcal{D}_T|f}\left[\mathcal{L}(\mu_T)\right]\right] \leq e^{2\lambda^2}(1+4\lambda^2)e^{2(1+\kappa)|\lambda|(\beta_T^{1/2} + \epsilon_T)} \cdot \left(\frac{C_1 \gamma_T}{T} + \epsilon_{\mathrm{MC}}\right) + e^{2\lambda^2}(1+4\lambda^2)\delta. \tag{6}$$

Here, δ is an arbitrary confidence parameter inherited from Theorem 4.3, and the additional term accounts for the contribution from the event where its high-probability bound fails. The resulting Bayesian MSE bound holds unconditionally for any $\delta \in (0, 1)$.

Corollary 4.5 (Bayesian Terminal MSE Convergence Rate)**.** *Under the conditions of Corollary 4.4, suppose that $\gamma_T = \mathcal{O}((\log T)^q)$ for some fixed $q \geq 0$. Then the Bayesian terminal MSE $\overline{\mathcal{L}}_T := \mathbb{E}_f[\mathbb{E}_{\mathcal{D}_T|f}[\mathcal{L}(\mu_T)]]$ satisfies: (i) M is constant $\Rightarrow$ $\overline{\mathcal{L}}_T = \mathcal{O}(T^{-1/2+\epsilon})$; (ii) $M = \Omega(T) \Rightarrow \overline{\mathcal{L}}_T = \mathcal{O}(T^{-1+\epsilon})$, for arbitrarily small $\epsilon > 0$.*

Note that although Corollaries 4.2 and 4.5 achieve the same asymptotic rates, as illustrated in the proof given in Appendix C.7, the Bayesian analysis yields a more favorable finite-sample bound, with γ_T/T instead of $\beta_T \gamma_T/T$ and a smaller exponential coefficient, up to the additional failure-event term.

Robustness to Boltzmann Parameter Miss-specification. In practice the parameters used for data acquisition may not exactly coincide with those used in test-time evaluation (e.g., in molecular active learning [Kulichenko et al., 2023]). We therefore consider bounded misspecification of both the scaling parameter λ and the bias function b. As shown in Proposition C.6 in Appendix C.9, the

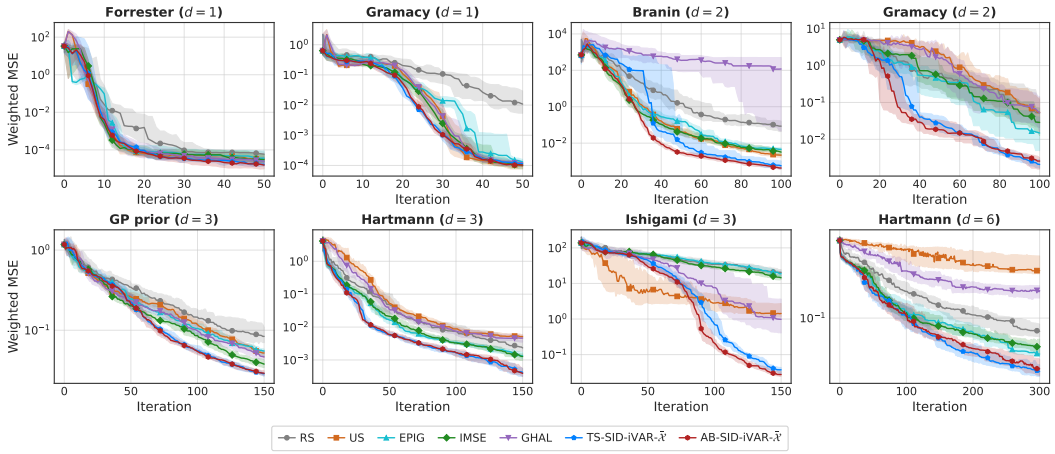


Figure 3: Weighted MSE convergence on synthetic benchmarks across dimensions $d = 1$ to $d = 6$. Each panel shows median performance over 30 random seeds (shaded regions: interquartile range). `AB-SID-iVAR-`$\bar{\mathcal{X}}$ consistently achieves lower weighted MSE than baselines (RS, US, IMSE), with the gap widening in higher dimensions.

high-probability guarantee of Theorem 4.1 can still hold with an additional multiplicative constant that depends on the degree of parameter misspecification and the range of f over $\mathcal{X}$.

Other Performance Metrics. While our main focus is $\mathcal{L}(\hat{f})$ defined in Eq. (2), the analysis extends to any metric controlled by posterior uncertainty:

Remark 4.6 (General Performance Metrics). Suppose a metric $\mathcal{M}_T$ satisfies

$$\mathbb{E}_{P_f}[\mathcal{M}_T] \leq C \cdot \mathbb{E}_{P_{\mu_T}}[g(\sigma_T(x))]$$

for some non-decreasing $g : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ with $g(0) = 0$ and g continuous at 0, and $C > 0$. Then $\mathbb{E}_{P_f}[\mathcal{M}_T] \to 0$ as $T \to \infty$.

4.3. Analyzing Existing Heuristics

The heuristic active learning approaches mentioned in the Related Works section and detailed in Appendix D.2, differ in various ad-hoc practical choices, but share a common algorithmic structure (which we term Generic Heuristic Active Learning, `GHAL`), targeting variance reduction under a Boltzmann surrogate as in our methods. This connection, together with the analytical tools developed in Section 3, allows us to make a first attempt at analyzing their convergence properties.

Specifically, at each iteration, starting from a point outside the feasible region $\bar{\mathcal{X}}'_t := \{x : \sigma^2_{t-1}(x) \geq \eta\}$ with threshold $\eta > 0$, a Markov chain is simulated targeting a surrogate distribution of generic form $\tilde{p}_{t-1} \propto \exp\big(\lambda\mu_{t-1}(x) + \lambda h(\sigma^2_{t-1}(x)) + b(x)\big)$, where $h(\sigma^2)$ is a smooth transformation of σ^2, until the chain crosses into the feasible region, at which point x_t is queried; if the Markov chain does not cross into the feasible region, the sample of maximum uncertainty is queried. If we assume that exploration by the Markov chain is exhaustive so that $\sigma^2_{t-1}(x_t) \geq \min\{\eta, \max_{x\in\mathcal{X}} \sigma^2_{t-1}(x)\}$, the conditions of the following convergence guarantee are satisfied:

Proposition 4.7. *For a fixed threshold $\eta \in (0, 1]$, define $\bar{\mathcal{X}}'_t := \{x : \sigma^2_{t-1}(x) \geq \eta\}$. Consider an active learning algorithm that selects x_t from the set $\bar{\mathcal{X}}'_t$ if $\bar{\mathcal{X}}'_t \neq \phi$ and $x_t = \arg\max_x \sigma_{t-1}(x)$ otherwise. Then, under Assumptions 2.1 and 3.1, and assuming $\gamma_T = \mathcal{O}((\log T)^q)$ for some fixed $q \geq 0$ as in Corollary 4.2, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$:*

$$\mathcal{L}(\mu_T) = \mathcal{O}(T^{-1+\epsilon}) \tag{7}$$

for arbitrarily small $\epsilon > 0$ and sufficiently large T such that $\eta > C_1\gamma_T/T$.

To the best of our knowledge, this is the first convergence analysis under the SIDAL framework for heuristic active learning methods tackling this problem.

5. Experiments

We validate our proposed methods against Random Sampling (RS), Uncertainty Sampling (US), Integrated Mean Squared Error (IMSE) [Sacks et al., 1989, Seo et al., 2000], Expected Predictive Information Gain (EPIG) [Smith et al., 2023], and GHAL detailed in Appendix D.2. We omit comparing with DRAL as it optimizes for worst-case performance over an ambiguity set rather than exploiting the known Boltzmann form. We use a Matérn-5/2 GP surrogate by default. For our methods, we set $\lambda = 1$ and $b(x) = 0$ as default (see Appendix D.8 for sensitivity to other settings). Each experiment starts with a single initial sample and is repeated 30 times with different random seeds. Full experimental setup is deferred to Appendix D.1.

5.1. Synthetic Experiments

We first evaluate AL performance on standard synthetic benchmarks spanning $d = 1$ to $d = 6$. Results are illustrated in Figure 3. AB-SID-iVAR consistently achieves lower weighted MSE than baselines, with the gap becoming more pronounced in higher dimensions. The Thompson sampling variant TS-SID-iVAR is competitive but exhibits higher variance across seeds, particularly on multimodal targets such as Branin. US shows a clear failure mode: by concentrating queries in high-uncertainty regions regardless of their probability under P_f, it can perform worse than random sampling. IMSE is a strong baseline as it also minimizes integrated variance, but learns uniformly and thus underperforms SID-aware approaches. GHAL, while SID-aware and competitive in some tasks, exhibits high variance and inconsistent performance due to its sensitivity to hyperparameters η and h; in contrast, our methods avoid task-specific hyperparameter tuning.

Ablation Study. To understand the contribution of different components, we conduct an ablation study under two settings: (1) using μ_{t-1} as a naive surrogate for SID (i.e., $e^{\lambda \mu_{t-1}}$), tested with both our acquisition (denoted μ-SID-iVAR) and EPIG (denoted μ-SID-EPIG) to verify that combining any acquisition function with a plug-in SID surrogate is insufficient, and (2) removing the constraint set (denoted SID-iVAR-ϕ). Results are shown in Figure 8 and analyzed in detail in Appendix D.4. While these variants achieve comparable performance on benchmarks dominated by a broad basin (e.g., Hartmann), both plug-in variants fail in scenarios where the target distribution has well-separated modes (Branin) or strongly oscillatory structure (Ishigami), as the point estimate μ_{t-1} discards posterior uncertainty and commits prematurely to whichever mode is discovered first. This validates that the Bayesian treatment of f in our surrogate $\tilde{p}_{t-1}$, alongside the constraint set $\tilde{\mathcal{X}}_t$, enables its robust performance.

Finally, we report computational cost comparisons in Appendix D.6. In discrete domains, our methods are slightly slower than IMSE (<10 seconds per iteration). In continuous domains, our methods incur higher per-iteration cost than IMSE (which uses uniform random samples) due to MCMC sampling and constrained acquisition optimization, but remain practical (<30 seconds per iteration), a sensitivity regarding MCMC particle size at Figure 9 in Appendix D.7 also shows that the empirical performance stays the same while using $N = 250d$ particles.

5.2. Potential Energy Surface Modeling

Learning the potential energy surface (PES) $E(\mathbf{x})$ of an atomistic system is a natural SIDAL instance: downstream molecular dynamics samples configurations from the Boltzmann distribution $P_E(\mathbf{x}) \propto \exp(-E(\mathbf{x})/k_{\mathrm{B}}\mathcal{T})$ (i.e., $\lambda = -(k_{\mathrm{B}}\mathcal{T})^{-1} < 0$, $b \equiv 0$, where $\mathcal{T} > 0$ denotes the thermal temperature and k_{B} is the Boltzmann constant), so prediction accuracy is only required where this self-induced density places mass, typically a small, a priori unknown fraction of configuration space corresponding to physically accessible states. Ground-truth queries to E using, e.g., density functional theory [Szabo and Ostlund, 2012], typically cost hours to weeks per evaluation, making sample efficiency essential.

We benchmark the proposed methods on four PES fitting problems, the PES of (i) a hydrogen molecule on a copper surface ($\mathrm{H_2}$ on Cu), (ii) a hydrogen atom on a copper cluster ($\mathrm{H_2}$ on Cu Cluster), (iii) a highly symmetric silicon cluster (Si Crystal), and (iv) a water molecule on a platinum surface ($\mathrm{H_2O}$ on Pt). Details, including coordinate parameterizations and λ values, are deferred to Appendix D.5.1.

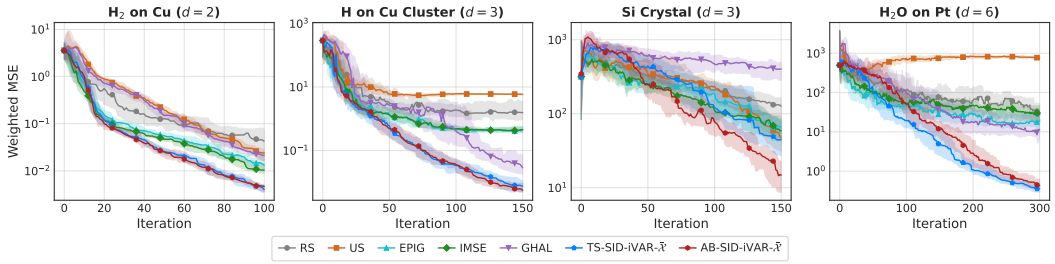


Figure 4: Weighted MSE convergence on four PES fitting tasks of increasing dimensionality: H_2 on Cu (d=2), H on Cu_{13} (d=3), Si Crystal (d=3), and H_2O on Pt (d=6). `AB-SID-iVAR` consistently achieves the lowest error across systems.

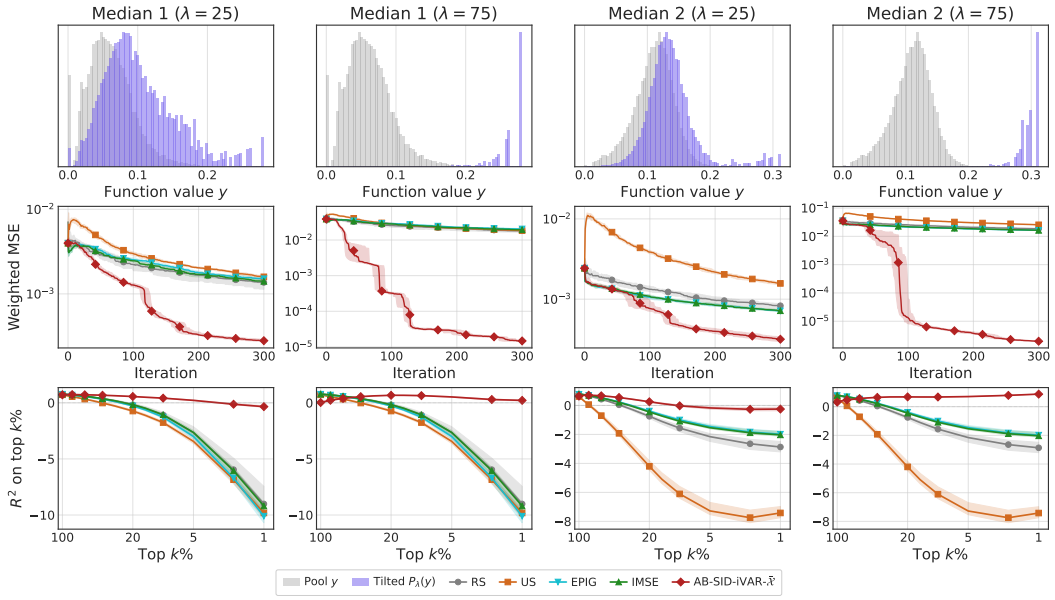


Figure 5: Active learning for molecular property prediction on GuacaMol scoring functions (Median 1, Median 2). The target distribution $P_\lambda(x) \propto \exp(\lambda f(x))$ emphasizes high-scoring compounds, relevant for drug discovery where accurate predictions matter most for promising candidates. **Top:** pool distribution (gray) vs. Boltzmann-tilted density (red). **Middle:** weighted MSE convergence over 300 iterations; `AB-SID-iVAR` achieves one to two orders of magnitude lower error than baselines. **Bottom:** R^2 on top $k\%$ compounds ranked by Boltzmann weight; baselines degrade to negative R^2 as k shrinks while `AB-SID-iVAR` remains near zero. Pool contains approximately 20,000 molecules with Tanimoto kernel GP.

The result summarized in Figure 4 shows that `AB-SID-iVAR` consistently achieves the lowest weighted MSE, the performance of GHAL, on the other hand, varies between problems, indicating its sensitivity to the specification of hyperparameters.

5.3. Molecular Drug Discovery

Finally, we evaluate our methods on a discrete pool-based active learning problem using real molecular dataset. In practice, evaluating molecular properties is expensive hence drug discovery campaigns prioritize accurate predictions for promising candidates over uniform accuracy across the library. This naturally matches the SIDAL setting, where the Boltzmann density concentrates on high-scoring compounds that matter most for downstream screening. We use a GP with a Tanimoto kernel over Morgan fingerprints as the surrogate, suited for the discrete molecular domain. This also tests whether our method remains effective beyond the kernel regularity assumed in our theory (Assumption 2.1).

We evaluate on the GuacaMol benchmark [Brown et al., 2019] (≈20,000 molecules) using two drug-likeness scoring functions (Median 1, Median 2), with $\lambda \in \{25, 75\}$ [5] controlling target concentration (Figure 5, top row). Setup details are deferred to Appendix D.5.2. Increasing λ from 25 to 75 sharpens the target, concentrating mass on the highest-scoring molecules. This exponential tilting appears broadly, including guided sampling in diffusion models [Song et al., 2023] and reward-tilted preference alignment [Rafailov et al., 2023]. Based on the convergence curves, `AB-SID-iVAR` demonstrates strong performance, converging roughly an order of magnitude lower at $\lambda = 25$, and this benefit becomes more pronounced at $\lambda = 75$. This reflects the increased cost of learning uniformly when only a narrow slice of the input space is of interest.

Further inspecting final model quality, the bottom row evaluates R^2 [6] on the top-$k\%$ of molecules ranked by Boltzmann weight. At $k = 100\%$ represent uniform learning objective, all methods achieve comparable R^2, reflecting uniform-average performance. As k decreases toward the elite subset, however, SID-unaware baselines degrade to strongly negative R^2, indicating predictions worse than a constant baseline on the molecules that matter most. In comparison, `AB-SID-iVAR` demonstrates significantly better prediction accuracy in high-value regions. The pairwise plots in Figure 11 confirm this finding: dark (high-weight) points align with the diagonal only for `AB-SID-iVAR`, while baselines show systematic spread or bias. Together, these results demonstrate that SID-unaware methods can appear reasonable under global metrics yet fail on the subset relevant for downstream screening, and λ provides a principled mechanism to specify *which* output region should be learned well, without requiring prior knowledge of where these regions lie in input space.

6. Conclusions

We formalized the problem of active learning under self-induced distributions (SIDAL), where the target distribution depends on the unknown function being learned. We proposed two uncertainty reduction acquisition functions based on Gaussian Process surrogates that approximate the intractable Bayesian target in closed form while avoiding partition function estimation. Our theoretical analysis establishes that the terminal prediction error vanishes with high probability under mild conditions, with tighter logarithmic factors available under an average-case analysis; a detailed elaboration of limitations and future work is deferred to Appendix E. Empirically, our methods consistently outperform SID-unaware baselines across synthetic benchmarks and real-world problems, with the gap widening as the target concentrates or the input dimension increases.

Acknowledgments

This work was supported by Research England under the Expanding Excellence in England (E3) funding stream, awarded to MARS: Mathematics for AI in Real-world Systems in the School of Mathematical Sciences at Lancaster University. The authors also gratefully acknowledge the use of Lancaster University's High End Computing (HEC) facility, which provided computational resources for this work. The authors are also grateful to Yueqi Cao for helpful discussions on the theoretical analysis presented in this paper, and to Zhuoyue Huang and Tom Rainforth for valuable discussions that helped clarify the positioning of this work within the broader active learning literature.

References

Tara Akhound-Sadegh, Jarrid Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, et al. Iterated denoising energy matching for sampling from boltzmann densities. *arXiv preprint arXiv:2402.06121*, 2024.

Tara Akhound-Sadegh, Jungyoon Lee, Avishek Joey Bose, Valentin De Bortoli, Arnaud Doucet, Michael M Bronstein, Dominique Beaini, Siamak Ravanbakhsh, Kirill Neklyudov, and Alexander Tong. Progressive inference-time annealing of diffusion models for sampling from boltzmann densities. *arXiv preprint arXiv:2506.16471*, 2025.

Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. Botorch: A framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems*, 33:21524–21538, 2020.

[5] We note the large λ is offset by the narrow score range $f \in [0.1, 0.3]$ (Figure 5, top row).

[6] We use $R^2 = 1 - \sum_i (y_i - \hat{y}_i)^2 / \sum_i (y_i - \bar{y})^2$ computed on the top-$k\%$ subset, where $\bar{y}$ is the subset mean.

Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M Elena, Dávid P Kovács, Janosh Riebesell, Xavier R Advincula, Mark Asta, Matthew Avaylon, William J Baldwin, et al. A foundation model for atomistic materials chemistry. *The Journal of chemical physics*, 163(18), 2025.

Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking models for de novo molecular design. *Journal of chemical information and modeling*, 59(3):1096–1108, 2019.

Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on scientific computing*, 16(5):1190–1208, 1995.

Xu Cai and Jonathan Scarlett. Kernelized normalizing constant estimation: bridging bayesian quadrature and bayesian optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 11150–11158, 2024.

Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(3):411–436, 2006.

Blake R Duschatko, Jonathan Vandermause, Nicola Molinari, and Boris Kozinsky. Uncertainty driven active learning of coarse grained free energy models. *npj Computational Materials*, 10(1):9, 2024.

Valerii Vadimovich Fedorov. *Theory of optimal experiments*. Elsevier, 2013.

David A Freedman. On tail probabilities for martingales. *the Annals of Probability*, pages 100–118, 1975.

Daan Frenkel and Berend Smit. *Understanding molecular simulation: from algorithms to applications*. Elsevier, 2023.

Charlie Frogner, Sebastian Claici, Edward Chien, and Justin Solomon. Incorporating unlabeled data into distributionally-robust learning. *J. Mach. Learn. Res.*, 22(1), January 2021. ISSN 1532-4435.

Subhashis Ghosal and Anindya Roy. Posterior consistency of gaussian process prior for nonparametric binary regression. 2006.

Robert B Gramacy and Herbert KH Lee. Gaussian processes and limiting linear models. *Computational Statistics & Data Analysis*, 53(1):123–136, 2008.

Xun Huan, Jayanth Jagalur, and Youssef Marzouk. Optimal experimental design: Formulations and computations. *Acta Numerica*, 33:715–840, 2024.

Zhuoyue Huang, Freddie Bickford Smith, and Tom Rainforth. Loss-driven bayesian active learning. *arXiv preprint arXiv:2604.11995*, 2026.

Jonas Hübotter, Lenart Treven, Yarden As, and Andreas Krause. Transductive active learning: Theory and applications. *Advances in Neural Information Processing Systems*, 37:124686–124755, 2024.

Carl Hvarfner, Erik Orm Hellsten, and Luigi Nardi. Vanilla bayesian optimization performs great in high dimensions. *arXiv preprint arXiv:2402.02229*, 2024.

Ajay Jasra, David A Stephens, Arnaud Doucet, and Theodoros Tsagaris. Inference for lévy-driven stochastic volatility models via adaptive sequential monte carlo. *Scandinavian Journal of Statistics*, 38(1):1–22, 2011.

Ryosuke Jinnouchi, Kazutoshi Miwa, Ferenc Karsai, Georg Kresse, and Ryoji Asahi. On-the-fly active learning of interatomic potentials for large-scale atomistic simulations. *The Journal of Physical Chemistry Letters*, 11 (17):6946–6955, 2020.

Andreas Kirsch, Tom Rainforth, and Yarin Gal. Test distribution-aware active learning: A principled approach against distribution shift and outliers. *arXiv preprint arXiv:2106.11719*, 2021.

Dieter Kraft. A software package for sequential quadratic programming. *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt fur Luft- und Raumfahrt*, 1988.

Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(2), 2008.

Maksim Kulichenko, Kipton Barros, Nicholas Lubbers, Ying Wai Li, Richard Messerly, Sergei Tretiak, Justin S Smith, and Benjamin Nebgen. Uncertainty-driven dynamics for active learning of interatomic potentials. *Nature computational science*, 3(3):230–239, 2023.

Zhenwei Li, James R Kermode, and Alessandro De Vita. Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces. *Physical review letters*, 114(9):096405, 2015.

RuiKang OuYang, Bo Qiang, and José Miguel Hernández-Lobato. Bnem: A boltzmann sampler based on bootstrapped noised energy matching. *arXiv preprint arXiv:2409.09787*, 2024.

Evgeny V Podryabinkin and Alexander V Shapeev. Active learning of linearly parametrized interatomic potentials. *Computational Materials Science*, 140:171–180, 2017.

Friedrich Pukelsheim. *Optimal design of experiments*. SIAM, 2006.

Jixiang Qing, Nicolas Knudde, Federico Garbuglia, Domenico Spina, Ivo Couckuyt, and Tom Dhaene. Adaptive sampling with automatic stopping for feasible region identification in engineering design. *Engineering with Computers*, 38(Suppl 3):1955–1972, 2022.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741, 2023.

Tom Rainforth, Adam Foster, Desi R Ivanova, and Freddie Bickford Smith. Modern bayesian experimental design. *Statistical Science*, 39(1):100–114, 2024.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 11 2006. ISBN 9780262256834. doi: 10.7551/mitpress/3206.001.0001. URL `https://doi.org/10.7551/mitpress/3206.001.0001`.

Jerome Sacks, William J Welch, Toby J Mitchell, and Henry P Wynn. Design and analysis of computer experiments. *Statistical science*, 4(4):409–423, 1989.

Christian Schäfer and Nicolas Chopin. Sequential monte carlo on large binary sampling spaces. *Statistics and Computing*, 23(2):163–184, 2013.

Sebastian Schnur and Axel Groß. Properties of metal–water interfaces studied from first principles. *New Journal of Physics*, 11(12):125003, 2009.

Sambu Seo, Marko Wallat, Thore Graepel, and Klaus Obermayer. Gaussian process regression: Active data selection and test point rejection. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks (IJCNN)*, volume 3, pages 241–246, 2000.

Burr Settles. Active learning literature survey. 2009.

Freddie Bickford Smith, Andreas Kirsch, Sebastian Farquhar, Yarin Gal, Adam Foster, and Tom Rainforth. Prediction-oriented bayesian active learning. In *International conference on artificial intelligence and statistics*, pages 7331–7348. PMLR, 2023.

B Smits and MF Somers. The quantum dynamics of h2 on cu (111) at a surface temperature of 925 k: Comparing state-of-the-art theory to state-of-the-art experiments 2. *The Journal of Chemical Physics*, 158(1), 2023.

Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation. In *International Conference on Machine Learning*, pages 32483–32498. PMLR, 2023.

Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.

Wojciech G. Stark, Julia Westermayr, Oscar A. Douglas-Gallardo, James Gardner, Scott Habershon, and Reinhard J. Maurer. Machine learning interatomic potentials for reactive hydrogen dynamics at metal surfaces based on iterative refinement of reaction probabilities. *The Journal of Physical Chemistry C*, 127(50):24168–24182, 2023. doi: 10.1021/acs.jpcc.3c06648.

Wojciech G Stark, Cas van der Oord, Ilyes Batatia, Yaolong Zhang, Bin Jiang, Gábor Csányi, and Reinhard J Maurer. Benchmarking of machine learning interatomic potentials for reactive hydrogen dynamics at metal surfaces. *Machine Learning: Science and Technology*, 5(3):030501, 2024.

Wojciech G. Stark, Connor L. Box, Matthias Sachs, Nils Hertl, and Reinhard J. Maurer. Nonadiabatic reactive scattering of hydrogen on different surface facets of copper. *Physical Review B*, 112(3):035422, 2025. doi: 10.1103/h7vd-94pk.

Frank H Stillinger and Thomas A Weber. Computer simulation of local order in condensed phases of silicon. *Physical review B*, 31(8):5262, 1985.

Attila Szabo and Neil S Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 2012.

Shion Takeno, Yoshito Okura, Yu Inatsu, Tatsuya Aoyama, Tomonari Tanaka, Satoshi Akahane, Hiroyuki Hanada, Noriaki Hashimoto, Taro Murayama, Hanju Lee, et al. Distributionally robust active learning for gaussian process regression. *arXiv preprint arXiv:2502.16870*, 2025.

Austin Tripp, Sergio Bacallado, Sukriti Singh, and José Miguel Hernández-Lobato. Tanimoto random features for scalable molecular machine learning. *Advances in Neural Information Processing Systems*, 36:33656–33686, 2023.

Sattar Vakili, Henry Moss, Artem Artemev, Vincent Dutordoir, and Victor Picheny. Scalable thompson sampling using sparse gaussian process models. *Advances in neural information processing systems*, 34:5631–5643, 2021.

Cas van der Oord, Matthias Sachs, Dávid Péter Kovács, Christoph Ortner, and Gábor Csányi. Hyperactive learning for data-driven interatomic potentials. *npj Computational Materials*, 9(1):168, 2023.

Jonathan Vandermause, Steven B Torrisi, Simon Batzner, Yu Xie, Lixin Sun, Alexie M Kolpak, and Boris Kozinsky. On-the-fly active learning of interpretable bayesian force fields for atomistic rare events. *npj Computational Materials*, 6(1):20, 2020.

James T Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Peter Deisenroth. Pathwise conditioning of gaussian processes. *Journal of Machine Learning Research*, 22(105):1–47, 2021.

Mingli Yang, Koblar A Jackson, Christof Koehler, Thomas Frauenheim, and Julius Jellinek. Structure and shape variations in intermediate-size copper clusters. *The Journal of chemical physics*, 124(2), 2006.

Kai Yu, Jinbo Bi, and Volker Tresp. Active learning via transductive experimental design. In *Proceedings of the 23rd international conference on Machine learning*, pages 1081–1088, 2006.

Appendix Contents

A Related Works **14**

B Methodology Details **15**

B.1 Acquisition Functions 15

B.1.1 AB-SID 15

B.1.2 TS-SID 17

B.1.3 Acquisition Illustrations 17

B.1.4 MCMC approximation of $\tilde{p}_t$ 17

C Theoretic Analysis **17**

C.1 Non-Bayes Optimality 17

C.2 Lemma 3.2 18

C.3 Lemma 3.3 18

C.4 Lemma 3.4 19

C.5 Theorem 4.1 and Corollary 4.2 20

C.6 Theorem 4.3 25

C.7 Rate Analysis for Posterior-Averaged and Bayesian MSE Bounds 26

C.8 Proposition 4.7 26

C.9 Robustness to Boltzmann Parameter Misspecification 27

D Experimental Details and Additional Experimental Investigations **29**

D.1 Experimental Setup Details 29

D.2 Generic Adaptation of Heuristics Active Learning Algorithms 29

D.3 Discrete Experimental Results 30

D.4 Ablation Study 30

D.5 Real World Problems 32

D.5.1 Potential Energy Surface Modeling 32

D.5.2 Molecular Drug Discovery 32

D.6 Runtime Report and Compute Resources 33

D.7 Sensitivity Analysis of MCMC sample size 33

D.8 Sensitivity Analysis of λ and b 33

D.9 Molecular Drug Discovery 35

E Limitations and Future Work **35**

A Related Works

Several lines of work consider active learning with structured targets rather than uniform error across the input domain. These methods differ in how the target is specified: a known finite set of test points [Yu et al., 2006, Hübotter et al., 2024], a continuous target distribution P [Kirsch et al., 2021, Smith et al., 2023], or an ambiguity set containing possible P when it cannot be exactly specified [Frogner et al., 2021, Takeno et al., 2025], hence optimizing for worst-case performance

over the set. Recently, a concurrent work by Huang et al. [2026] develops a general decision-theoretic framework for loss-driven Bayesian active learning, showing that the inner minimisation over the Bayes-optimal downstream action can be eliminated analytically for weighted Bregman-divergence losses. In particular, they consider exponentially weighted squared-error losses under GP regression, yielding an exponential–Gaussian weighting factor that also appears in our unconstrained AB-SID-iVAR acquisition criterion after the zeroth-order Taylor approximation. However, these expressions arise from different constructions: their exponential weighting is directly specified through the downstream loss, whereas ours is derived by approximating the posterior expectation of a normalised self-induced Boltzmann distribution. The latter involves an intractable expectation over a function-dependent partition function, so the exact Bayes-optimal estimator, although formally characterised in Appendix C.1, is generally not analytically tractable. We instead adopt the GP posterior mean as a tractable surrogate, introduce an explicit exploration constraint, and establish convergence guarantees for its terminal prediction error under the true self-induced distribution.

Outside the AL literature, related heuristics have emerged in scientific domains [Li et al., 2015, Podryabinkin and Shapeev, 2017, Vandermause et al., 2020, Jinnouchi et al., 2020, van der Oord et al., 2023, Kulichenko et al., 2023, Duschatko et al., 2024], where a surrogate model guides Markov Chain Monte Carlo (MCMC) sampling from a tilted Boltzmann distribution and triggers a query when predictive uncertainty exceeds a pre-specified threshold. These methods often require careful tuning of hyperparameters and lack theoretical guarantees. In Section 4.3, we formalize a generalized version of such heuristics within our framework and for the first time analyze their convergence properties.

B Methodology Details

B.1 Acquisition Functions

Starting from one-step look-ahead MSE minimization, we make two approximations: (i) replace squared error with posterior variance (according to lemma 2.3), and (ii) treat P_f as fixed when evaluating variance reduction, avoiding the need to recompute partition functions for each hypothetical observation. This reduces the acquisition function to minimizing $\mathbb{E}_{f|D_{t-1}}\mathbb{E}_{P_f}[\sigma_t^2 \mid x_t = x]$, which requires approximating $\mathbb{E}_{f|D_{t-1}}[P_f(x^\star)]$. Substituting the Boltzmann form:

$$\mathbb{E}_{f|D_{t-1}}[P_f(x^\star)] = \mathbb{E}_{f|D_{t-1}}\left[\frac{e^{\lambda f(x^\star)+b(x^\star)}}{Z_f}\right]. \tag{8}$$

B.1.1 AB-SID

For notation simplicity, let $A = e^{\lambda f(x^\star)+b(x^\star)}$ and $B = Z_f$. In order to compute $\mathbb{E}[A/B]$, We approximate $\mathbb{E}[A/B]$ via Taylor expansion around $B = \mathbb{E}[B]$:

$$\frac{A}{B} = \frac{A}{\mathbb{E}[B]} - \frac{A(B-\mathbb{E}[B])}{(\mathbb{E}[B])^2} + \frac{A(B-\mathbb{E}[B])^2}{(\mathbb{E}[B])^3} - \cdots.$$

Taking expectation:

$$\mathbb{E}\left[\frac{A}{B}\right] = \frac{\mathbb{E}[A]}{\mathbb{E}[B]} - \frac{\mathbb{E}[A(B-\mathbb{E}[B])]}{(\mathbb{E}[B])^2} + \frac{\mathbb{E}[A(B-\mathbb{E}[B])^2]}{(\mathbb{E}[B])^3} - \cdots.$$

For the second term, note that:

$$\mathbb{E}[A(B-\mathbb{E}[B])] = \mathbb{E}[AB] - \mathbb{E}[A]\mathbb{E}[B] = \mathrm{Cov}(A,B),$$

therefore,

$$\mathbb{E}\left[\frac{A}{B}\right] = \underbrace{\frac{\mathbb{E}[A]}{\mathbb{E}[B]}}_{T_0} - \underbrace{\frac{\mathrm{Cov}(A,B)}{(\mathbb{E}[B])^2}}_{T_1} + \frac{\mathbb{E}[A(B-\mathbb{E}[B])^2]}{(\mathbb{E}[B])^3} - \cdots$$

Zero-Order Term T_0 For the first term T_0, substituting A and B's expression back:

$$\mathbb{E}_{f|D}[P_f(x^\star)] \approx \frac{\mathbb{E}_{f|D}[e^{\lambda f(x^\star)+b(x^\star)}]}{\mathbb{E}_{f|D}[Z_f]}$$

For the numerator, again using the Gaussian MGF identity, we have:

$$\mathbb{E}_{f|D}[e^{\lambda f(x^\star)+b(x^\star)}] = \exp\left(\lambda\mu_t(x^\star) + \frac{\lambda^2}{2}\sigma_t^2(x^\star) + b(x^\star)\right)$$

Similarly, for the partition function:

$$\mathbb{E}_{f|D}[Z_f] = \sum_{x'} \exp\left(\lambda\mu_t(x') + \frac{\lambda^2}{2}\sigma_t^2(x') + b(x')\right)$$

First-Order T_1 For the first-order term

$$T_1 = -\frac{\mathrm{Cov}(A, B)}{(\mathbb{E}[B])^2}$$

For $\mathrm{Cov}(A, B)$, since $B = Z_f = \sum_{x'} e^{\lambda f(x')+b(x')}$, by bi-linearity of covariance:

$$\mathrm{Cov}(A, B) = \sum_{x'} \mathrm{Cov}\left(e^{\lambda f(x^\star)+b(x^\star)}, e^{\lambda f(x')+b(x')}\right)$$

To evaluate this we first introduce the following helper lemma.

Lemma B.1 (Exponential Covariance under Joint Gaussian)**.** *Let (U, V) be jointly Gaussian with means (μ_U, μ_V), variances (σ_U^2, σ_V^2), and covariance κ. Then:*

$$\mathrm{Cov}(e^{\lambda U}, e^{\lambda V}) = \mathbb{E}[e^{\lambda U}]\mathbb{E}[e^{\lambda V}]\left(e^{\lambda^2\kappa} - 1\right)$$

Proof. By the joint MGF of Gaussian random variables:

$$\mathbb{E}[e^{\lambda U+\lambda V}] = \exp\left(\lambda(\mu_U + \mu_V) + \frac{\lambda^2}{2}(\sigma_U^2 + \sigma_V^2 + 2\kappa)\right)$$

The product of marginal MGFs:

$$\mathbb{E}[e^{\lambda U}] \cdot \mathbb{E}[e^{\lambda V}] = \exp\left(\lambda(\mu_U + \mu_V) + \frac{\lambda^2}{2}(\sigma_U^2 + \sigma_V^2)\right)$$

Therefore:

$$\mathrm{Cov}(e^{\lambda U}, e^{\lambda V}) = \mathbb{E}[e^{\lambda U} \cdot e^{\lambda V}] - \mathbb{E}[e^{\lambda U}]\mathbb{E}[e^{\lambda V}] = \mathbb{E}[e^{\lambda U}]\mathbb{E}[e^{\lambda V}]\left(e^{\lambda^2\kappa} - 1\right)$$

□

Applying Lemma B.1 with $U = f(x^\star) + \frac{b(x^\star)}{\lambda}$, $V = f(x') + \frac{b(x')}{\lambda}$, and hence $\kappa = k_{t-1}(x^\star, x')$:

$$\mathrm{Cov}\left(e^{\lambda f(x^\star)+b(x^\star)}, e^{\lambda f(x')+b(x')}\right) = \mathbb{E}\left[e^{\lambda f(x^\star)+b(x^\star)}\right]\mathbb{E}\left[e^{\lambda f(x')+b(x')}\right]\left(e^{\lambda^2 k_{t-1}(x^\star,x')} - 1\right)$$

Summing over x':

$$\mathrm{Cov}(A, B) = \mathbb{E}[A]\sum_{x'} \mathbb{E}\left[e^{\lambda f(x')+b(x')}\right]\left(e^{\lambda^2 k_{t-1}(x^\star,x')} - 1\right)$$

Since $\mathbb{E}[B] = \sum_{x'} \mathbb{E}\left[e^{\lambda f(x')+b(x')}\right]$, the first-order term becomes:

$$\begin{aligned}
T_1 &= -\frac{\mathbb{E}[A]}{(\mathbb{E}[B])^2}\sum_{x'} \mathbb{E}\left[e^{\lambda f(x')+b(x')}\right]\left(e^{\lambda^2 k_{t-1}(x^\star,x')} - 1\right) \\
&= -\frac{\mathbb{E}[A]}{\mathbb{E}[B]}\sum_{x'} \frac{\mathbb{E}\left[e^{\lambda f(x')+b(x')}\right]}{\mathbb{E}[B]}\left(e^{\lambda^2 k_{t-1}(x^\star,x')} - 1\right) \\
&= -\frac{\mathbb{E}[A]}{\mathbb{E}[B]} \cdot \frac{\sum_{x'} \tilde{p}_{t-1}^{\mathrm{u}}(x')\left(e^{\lambda^2 k_{t-1}(x^\star,x')} - 1\right)}{\sum_{x''} \tilde{p}_{t-1}^{\mathrm{u}}(x'')}
\end{aligned}$$

where $\tilde{p}_{t-1}^{\mathrm{u}}(x') = e^{\left(\lambda\mu_{t-1}(x') + \frac{\lambda^2}{2}\sigma_{t-1}^2(x') + b(x')\right)}$, the same unnormalized surrogate as in Table 1.

B.1.2 TS-SID

Another approach is to approximate Eq. 8 by a Monte Carlo estimate,

$$\mathbb{E}_{f|D_{t-1}}[P_f(x^\star)] \approx \frac{1}{N}\sum_{i=1}^{N}\left[\frac{e^{\lambda f_i(x^\star)+b(x^\star)}}{Z_{f_i}}\right].$$

Specifically we use $N = 1$, which corresponds to Thompson sampling.

B.1.3 Acquisition Illustrations

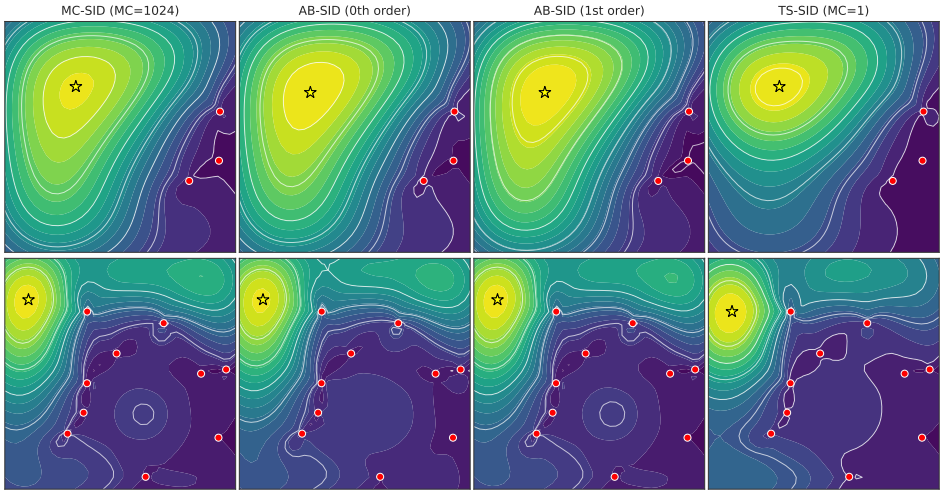


Figure 6: Acquisition function comparison on a 2D GP prior. We use Monte Carlo with 1024 samples (MC-SID, leftmost) as a reference approximation of Eq. 8. The zero-order `AB-SID`approximation closely matches the reference contours regardless of training data size (●), while the first-order correction provides marginal improvement especially when data is abundant. `TS-SID` (single posterior sample) shows higher variability. Stars (☆) indicate the acquisition maximum.

In order to better understand the approximation efficacy of the zero order and first order approximation of Eq. 8, we provide a two dimensional visual comparison of different acquisition function variants in Figure 6, which helps justifying the usage of our zero order approximation.

B.1.4 MCMC approximation of $\tilde{p}_t$

We leverage Sequential Monte Carlo (SMC) [Del Moral et al., 2006] as our importance sampler for P_f, using annealed tempering with Random Walk Metropolis (RWM) rejuvenation steps. The temperature schedule is adaptively determined via bisection to maintain a target effective sample size (ESS) ratio [Jasra et al., 2011, Schäfer and Chopin, 2013].

C Theoretic Analysis

C.1 Non-Bayes Optimality

We note that the Bayes-optimal estimator minimizes the posterior expected loss:

$$R(\hat{f}) = \mathbb{E}_{f|D_t}\left[\sum_{x\in\mathcal{X}} P_f(x)\cdot(\hat{f}(x)-f(x))^2\right]$$

Taking the derivative with respect to $\hat{f}$:

$$\frac{dR}{d\hat{f}} = \mathbb{E}_{f|D_t}[2P_f\cdot(\hat{f}-f)] = 0$$

Algorithm 2 Sequential Monte Carlo (SMC) Sampling

Require: Log unnormalized density $\log \tilde{p}(x)$, bounds $[l, u]$, particles N, ESS threshold α, RWM steps K
Ensure: Weighted samples $\{(x_i, w_i)\}_{i=1}^N$ approximating $p(x) \propto \tilde{p}(x)$

$x_i \sim \text{Uniform}(l, u)$ for $i = 1, \ldots, N$ `// Initialize particles`
$\log w_i \leftarrow 0$ for all i `// Uniform weights`
$\tau \leftarrow 0$ `// Initial temperature`
while $\tau < 1$ **do**
  $\tau_{\text{prev}} \leftarrow \tau$
  $\tau \leftarrow \min(\tau', 1)$ where τ' satisfies $\text{ESS}(\tau') \approx \alpha \cdot N$ `// Adaptive tempering`
  $\log w_i \leftarrow \log w_i + (\tau - \tau_{\text{prev}}) \cdot \log \tilde{p}(x_i)$ for all i `// Reweight`
  if ESS $< \alpha \cdot N$ **then**
    $\{x_i\} \leftarrow \text{SystematicResample}(\{x_i\}, \{\log w_i\})$ `// Resample`
    $\log w_i \leftarrow 0$ for all i
  end if
  for $k = 1, \ldots, K$ **do**
    $x_i \leftarrow \text{RWM}(x_i, \tau \cdot \log \tilde{p})$ for all i `// Rejuvenate`
  end for
end while
return $\{(x_i, \exp(\log w_i))\}_{i=1}^N$

This yields:

$$\hat{f}^\star(x') = \frac{\mathbb{E}_{f|D_t}[P_f(x') \cdot f(x')]}{\mathbb{E}_{f|D_t}[P_f(x')]} \neq \mu_t(x')$$

Since $P_f(x') = g(f(x'))/Z_f$ depends on f through both $g(f(x'))$ and the partition function $Z_f = \sum_x g(f(x))$, the numerator and denominator involve the intractable expectation. Hence we use $\mu_t(x)$ as a computationally tractable approximation.

C.2 Lemma 3.2

Proof. Expanding the log-density difference:

$$\log P_{g_1}(x) - \log P_{g_2}(x) = \lambda(g_1(x) - g_2(x)) + \log Z_{g_2} - \log Z_{g_1}$$

where $Z_{g_i} = \int_{x'} e^{\lambda g_i(x') + b(x')} dx'$. The first term by definition is:

$$|\lambda(g_1(x) - g_2(x))| = |\lambda| \Delta(x)$$

For the partition function ratio:

$$Z_{g_1} = \int_{x'} e^{\lambda g_2(x') + b(x')} e^{\lambda(g_1(x') - g_2(x'))} dx'$$

Since $|g_1(x') - g_2(x')| \leq \Delta_{\max}$ for all x' by definition:

$$e^{-|\lambda|\Delta_{\max}} Z_{g_2} \leq Z_{g_1} \leq e^{|\lambda|\Delta_{\max}} Z_{g_2}$$

Taking logarithms:

$$|\log Z_{g_1} - \log Z_{g_2}| \leq |\lambda| \Delta_{\max} \tag{9}$$

Combining via triangle inequality:

$$|\log P_{g_1}(x) - \log P_{g_2}(x)| \leq |\lambda| \Delta(x) + |\lambda| \Delta_{\max}$$

□

C.3 Lemma 3.3

Proof. By Lemma 2.3, with probability at least $1 - \delta$, at iteration T we have:

$$|f(x) - \mu_T(x)| \leq \beta_T^{1/2} \sigma_T(x) + \epsilon_T, \quad \forall x \in \mathcal{X}$$

Applying Lemma 3.2 with $(g_1, g_2) = (f, \mu_T)$ we have $\Delta(x) = |f(x) - \mu_T(x)| \le \beta_T^{1/2}\sigma_T(x) + \epsilon_T$ and $\Delta_{\max} \le \beta_T^{1/2}\sigma_{\max,T} + \epsilon_T$. Additionally with log partition function difference inequality (Eq. 9):

$$\begin{aligned}\log \frac{P_f(x)}{P_{\mu_T}(x)} &\le |\lambda||f(x) - \mu_T(x)| + |\lambda|\Delta_{\max} \\ &\le |\lambda|(\beta_T^{1/2}\sigma_T(x) + \epsilon_T) + |\lambda|(\beta_T^{1/2}\sigma_{\max,T} + \epsilon_T)\end{aligned}$$

Using $\sigma_T(x) \le \sigma_{\max,T}$ and exponentiating:

$$\frac{P_f(x)}{P_{\mu_T}(x)} \le e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T} + \epsilon_T)}$$

Bounding the MSE:

$$\begin{aligned}\mathbb{E}_{P_f}[(f - \mu_T)^2] &= \mathbb{E}_{P_{\mu_T}}\left[\frac{P_f(x)}{P_{\mu_T}(x)}(f(x) - \mu_T(x))^2\right] \\ &\le e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T} + \epsilon_T)} \cdot \mathbb{E}_{P_{\mu_T}}\left[(\beta_T^{1/2}\sigma_T(x) + \epsilon_T)^2\right] \\ &= e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T} + \epsilon_T)} \cdot \mathbb{E}_{P_{\mu_T}}\left[\beta_T\sigma_T^2(x) + 2\beta_T^{1/2}\sigma_T(x)\epsilon_T + \epsilon_T^2\right] \\ &\le e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T} + \epsilon_T)} \cdot \left(\beta_T \cdot \mathbb{E}_{P_{\mu_T}}\left[\sigma_T^2(x)\right] + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right)\end{aligned}$$

where the last inequality holds as $\sigma_T \le 1$.

□

C.4 Lemma 3.4

We first establish a helper lemma for expectations involving Gaussian random variables.

Lemma C.1. *Let (ξ, ζ) be jointly Gaussian with $\mathbb{E}[\xi] = \mathbb{E}[\zeta] = 0$, $\mathbb{V}(\xi) = u^2$, $\mathbb{V}(\zeta) = w^2$, and $\mathrm{Cov}(\xi, \zeta) = c$. Then:*

$$\mathbb{E}[\xi^2 e^{\lambda\zeta}] = (u^2 + \lambda^2 c^2)e^{\lambda^2 w^2/2}$$

Proof. Using the conditional distribution $\xi|\zeta \sim \mathcal{N}\left(\frac{c}{w^2}\zeta, u^2 - \frac{c^2}{w^2}\right)$:

$$\mathbb{E}[\xi^2|\zeta] = \mathbb{V}[\xi|\zeta] + (\mathbb{E}[\xi|\zeta])^2 = u^2 - \frac{c^2}{w^2} + \frac{c^2}{w^4}\zeta^2$$

By the tower property:

$$\begin{aligned}\mathbb{E}[\xi^2 e^{\lambda\zeta}] = \mathbb{E}\left[\mathbb{E}[\xi^2 e^{\lambda\zeta}|\zeta]\right] = \mathbb{E}\left[e^{\lambda\zeta}\mathbb{E}[\xi^2|\zeta]\right] &= \mathbb{E}\left[e^{\lambda\zeta}\left(u^2 - \frac{c^2}{w^2} + \frac{c^2}{w^4}\zeta^2\right)\right] \\ &= \left(u^2 - \frac{c^2}{w^2}\right)\mathbb{E}[e^{\lambda\zeta}] + \frac{c^2}{w^4}\mathbb{E}[\zeta^2 e^{\lambda\zeta}]\end{aligned}$$

Using the Gaussian MGF, $\mathbb{E}[e^{\lambda\zeta}] = e^{\lambda^2 w^2/2}$ and $\mathbb{E}[\zeta^2 e^{\lambda\zeta}] = (w^2 + \lambda^2 w^4)e^{\lambda^2 w^2/2}$ $\left(\text{i.e., } \frac{d\mathbb{E}[e^{\lambda\zeta}]}{d\lambda^2}\right)$, the result follows. □

Now we prove Lemma 3.4.

Proof. Let $r_t(x) = f(x) - \mu_{t-1}(x)$ denote the residual. Under the GP posterior, $r_t(x) \sim \mathcal{N}(0, \sigma_{t-1}^2(x))$. Using $f(x) = \mu_{t-1}(x) + r_t(x)$ and noting that $Z_f = Z_{\mu_{t-1}} \cdot \mathbb{E}_{P_{\mu_{t-1}}}[e^{\lambda r_t(x')}]$, the weighted MSE becomes:

$$L_t := \mathbb{E}_{P_f}[r_t(x)^2] = \int \frac{e^{\lambda\mu_{t-1}(x) + b(x)} \cdot e^{\lambda r_t(x)}}{Z_{\mu_{t-1}} \cdot \mathbb{E}_{P_{\mu_{t-1}}}[e^{\lambda r_t(x')}]} r_t(x)^2 dx = \frac{\mathbb{E}_{P_{\mu_{t-1}}}[e^{\lambda r_t(x)} r_t(x)^2]}{\mathbb{E}_{P_{\mu_{t-1}}}[e^{\lambda r_t(x')}]}$$

Applying Jensen's inequality to the denominator (using convexity of e^x):

$$L_t \leq \mathbb{E}_{P_{\mu_{t-1}}}[e^{\lambda r_t(x)} r_t(x)^2] \cdot e^{-\lambda \mathbb{E}_{P_{\mu_{t-1}}}[r_t(x')]}$$

Define $\bar{r}_t := \mathbb{E}_{P_{\mu_{t-1}}}[r_t(x')]$. Since $e^{-\lambda \bar{r}_t}$ does not depend on x, we have:

$$L_t \leq \mathbb{E}_{P_{\mu_{t-1}}}[e^{\lambda r_t(x)} r_t(x)^2] \cdot e^{-\lambda \bar{r}_t} = \mathbb{E}_{P_{\mu_{t-1}}}[r_t(x)^2 e^{\lambda(r_t(x) - \bar{r}_t)}]$$

Taking the posterior expectation over $f|D_{t-1}$ and applying Fubini's theorem:

$$\mathbb{E}_{f|D_{t-1}}[L_t] \leq \mathbb{E}_{P_{\mu_{t-1}}}\left[\mathbb{E}_{f|D_{t-1}}\left[r_t(x)^2 e^{\lambda(r_t(x) - \bar{r}_t)}\right]\right]$$

Define $\zeta(x) := r_t(x) - \bar{r}_t$. Under the posterior, $(r_t(x), \zeta(x))$ are jointly Gaussian with:

$$\begin{aligned}
\mathbb{V}(\zeta(x)) &= \sigma^2_{t-1}(x) + V^2_{t-1} - 2C_{t-1}(x) \\
\mathrm{Cov}(r_t(x), \zeta(x)) &= \sigma^2_{t-1}(x) - C_{t-1}(x)
\end{aligned}$$

where

$$\begin{aligned}
V^2_{t-1} :&= \mathbb{E}_{f|D_{t-1}}\left[\left(\mathbb{E}_{x' \sim P_{\mu_{t-1}}}[r_t(x')]\right)^2\right] = \mathbb{E}_{f|D_{t-1}}\left[\mathbb{E}_{(x',x'') \sim P^2_{\mu_{t-1}}}[r_t(x') r_t(x'')]\right] \\
&= \mathbb{E}_{(x',x'') \sim P^2_{\mu_{t-1}}}[k_{t-1}(x', x'')]
\end{aligned}$$

and (similarly)

$$\begin{aligned}
C_{t-1}(x) &:= \mathrm{Cov}(r_t(x), \bar{r}_t) = \mathbb{E}_{f|D_{t-1}}[r_t(x)\bar{r}_t] = \mathbb{E}_{f|D_{t-1}}\left[r_t(x)\mathbb{E}_{x' \sim P_{\mu_{t-1}}}[r_t(x')]\right] \\
&= \mathbb{E}_{x' \sim P_{\mu_{t-1}}}[k_{t-1}(x, x')]
\end{aligned}$$

using $\mathbb{E}_{f|D_{t-1}}[r_t(x)] = 0$ and $\mathbb{E}_{f|D_{t-1}}[r_t(x) r_t(x')] = k_{t-1}(x, x')$ and Fubini's theorem.

By Cauchy-Schwarz inequality on the positive semi-definite kernel (i.e., $|k_{t-1}(x, x')| \leq \sigma_{t-1}(x)\sigma_{t-1}(x') \leq \sigma_{t-1}(x)$), we have $|C_{t-1}(x)| \leq \sigma_{t-1}(x)$, so:

$$\begin{aligned}
&|\sigma^2_{t-1}(x) - C_{t-1}(x)| \leq |\sigma^2_{t-1}(x)| + |C_{t-1}(x)| \leq 2\sigma_{t-1}(x) \\
&\mathbb{V}(\zeta(x)) \leq 4
\end{aligned}$$

Applying Lemma C.1 with $\xi = r_t(x)$ and $\zeta = \zeta(x)$ as defined above, we have $u^2 = \sigma^2_{t-1}(x)$, $w^2 = \mathbb{V}(\zeta(x))$, and $c = \sigma^2_{t-1}(x) - C_{t-1}(x)$:

$$\begin{aligned}
\mathbb{E}_{f|D_{t-1}}[r_t(x)^2 e^{\lambda\zeta(x)}] &= (\sigma^2_{t-1}(x) + \lambda^2(\sigma^2_{t-1}(x) - C_{t-1}(x))^2) e^{\lambda^2 \mathbb{V}(\zeta(x))/2} \\
&\leq (\sigma^2_{t-1}(x) + 4\lambda^2 \sigma^2_{t-1}(x)) e^{2\lambda^2} \\
&= (1 + 4\lambda^2)\sigma^2_{t-1}(x) \cdot e^{2\lambda^2}
\end{aligned}$$

Finally, taking expectation over $x \sim P_{\mu_{t-1}}$:

$$\mathbb{E}_{f|D_{t-1}}[L_t] \leq e^{2\lambda^2}(1 + 4\lambda^2) \cdot \mathbb{E}_{P_{\mu_{t-1}}}[\sigma^2_{t-1}(x)]$$

□

C.5 Theorem 4.1 and Corollary 4.2

We first recall a standard result on information gain.

Lemma C.2 (Information Gain Bound [Srinivas et al., 2009, Lemma 5.4])**.** *For any sequence of GP queries $x_1, \ldots, x_T$:*

$$\sum_{t=1}^{T} \sigma^2_{t-1}(x_t) \leq C_1 \gamma_T,$$

where $C_1 = 2/\log(1 + \tau^{-2})$ and τ^2 is the noise variance.

Uniform confidence event. Throughout this subsection we work on the event that the bound of Lemma 2.3 holds simultaneously for all residuals $f - \mu_{t-1}$, $t \le T+1$. By a union bound over these $T+1$ rounds, using a common discretizations grid of spacing $1/(dT)$, this event has probability at least $1-\delta_{\mathrm{GP}}$ when Lemma 2.3 is invoked with failure probability $\delta_{\mathrm{GP}}/(T+1)$. As this only adds a $\log(T+1)$ term inside β_T and L, we retain the notation β_T, ϵ_T, with $\beta_T = \mathcal{O}(\log T)$ and $\epsilon_T = \mathcal{O}(\sqrt{\log T}/T)$. In Theorems 4.1 and 4.3, β_T, ϵ_T are the coefficients of this event with total GP budget $\delta_{\mathrm{GP}} = \delta/2$ (split equally between the events for f and for the posterior samples in the case of `TS-SID`). Lemma 3.3 is applied on the same event.

Lemma C.3 (Variance Bound for `AB-SID-iVAR`)**.** *Consider `AB-SID-iVAR` selecting $x_t \in \bar{\mathcal{X}}_t$ where $\bar{\mathcal{X}}_t$ is defined in Table 1. If the constraint threshold at each iteration is estimated via M Monte Carlo samples, then for any $\delta_{\mathrm{GP}}, \delta_{\mathrm{MC}} \in (0,1)$, with probability at least $1-\delta_{\mathrm{GP}}-\delta_{\mathrm{MC}}$:*

$$\mathbb{E}_{P_{\mu_T}}[\sigma_T^2] \le e^{4|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1 \gamma_T}{T} + \frac{2\log(2/\delta_{\mathrm{MC}})}{3T} + \sqrt{\frac{\log(2/\delta_{\mathrm{MC}})}{2MT}} \right),$$

where β_T, ϵ_T are as in Lemma 2.3 with failure probability δ_{GP}. For finite $\mathcal{X}$ with exact threshold computation, the $\log(2/\delta_{\mathrm{MC}})$ terms vanish.

Proof. The proof proceeds in four parts: (1) relating the surrogate distribution to the target, (2) establishing a martingale structure for MC errors, (3) bounding the cumulative variance, and (4) handling the distribution shift.

Since $\tilde{p}^{\mathrm{u}}_{t-1} = e^{\lambda \mu_{t-1} + \frac{\lambda^2}{2}\sigma^2_{t-1} + b}$ and its normalized version satisfies $\tilde{p}_{t-1}(x) \propto P_{\mu_{t-1}}(x) \cdot e^{\frac{\lambda^2}{2}\sigma^2_{t-1}}$, we have:

$$\mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] = \frac{\mathbb{E}_{P_{\mu_{t-1}}}[\sigma^2_{t-1} \cdot e^{\frac{\lambda^2}{2}\sigma^2_{t-1}}]}{\mathbb{E}_{P_{\mu_{t-1}}}[e^{\frac{\lambda^2}{2}\sigma^2_{t-1}}]} \ge \mathbb{E}_{P_{\mu_{t-1}}}[\sigma^2_{t-1}], \tag{10}$$

where the last inequality uses $\mathrm{Cov}_{P_{\mu_{t-1}}}(\sigma^2_{t-1}, e^{\frac{\lambda^2}{2}\sigma^2_{t-1}}) \ge 0$: let $x, x' \overset{\text{i.i.d.}}{\sim} P_{\mu_{t-1}}$, $h(x) = \sigma^2_{t-1}(x)$, $w(x) = e^{\frac{\lambda^2}{2}\sigma^2_{t-1}(x)}$. Since w is increasing in h, $(h(x)-h(x'))$ and $(w(x)-w(x'))$ have the same sign, so $\mathrm{Cov}(h,w) = \frac{1}{2}\mathbb{E}[(h(x)-h(x'))(w(x)-w(x'))] \ge 0$.

The MC estimate of the $\bar{\mathcal{X}}$'s threshold at round t:

$$\hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] := \frac{1}{M}\sum_{i=1}^{M} \sigma^2_{t-1}(x^{(i)}), \quad x^{(i)} \overset{\text{i.i.d.}}{\sim} \tilde{p}_{t-1}. \tag{11}$$

Let $\mathcal{F}_{t-1}$ denote the σ-algebra generated by $D_{t-1} = \{(x_i, y_i)\}_{i=1}^{t-1}$ together with all Monte Carlo samples used up to round $t-1$. Since $\tilde{p}_{t-1}$ is fully determined by μ_{t-1} and σ^2_{t-1}, which are $\mathcal{F}_{t-1}$-measurable, and the round-t samples $x^{(i)}$ are drawn afresh, the estimation error

$$\xi_t := \hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] - \mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] = \frac{1}{M}\sum_{i=1}^{M} \sigma^2_{t-1}(x^{(i)}) - \mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] \tag{12}$$

satisfies $\mathbb{E}[\xi_t | \mathcal{F}_{t-1}] = 0$, forming a martingale difference sequence.

Since $\sigma^2_{t-1}(x) \in [0,1]$ by kernel normalization, the increments are bounded $|\xi_t| \le 1$. For the conditional variance, given $\mathcal{F}_{t-1}$, the expectation term is a constant and the samples $x^{(i)} \overset{\text{i.i.d.}}{\sim} \tilde{p}_{t-1}$, so by the sample mean variance formula:

$$\mathrm{Var}(\xi_t | \mathcal{F}_{t-1}) = \mathrm{Var}\left(\frac{1}{M}\sum_{i=1}^{M} \sigma^2_{t-1}(x^{(i)}) \Big| \mathcal{F}_{t-1} \right) = \frac{1}{M}\mathrm{Var}_{x\sim\tilde{p}_{t-1}}(\sigma^2_{t-1}(x)) \le \frac{1}{4M},$$

where the last inequality uses $\mathrm{Var}(Y) \le 1/4$ for any random variable $Y \in [0,1]$. Hence, the cumulative error $S_T := \sum_{t=1}^{T} \xi_t$ has total conditional variance $V_T := \sum_{t=1}^{T} \mathrm{Var}(\xi_t | \mathcal{F}_{t-1}) \le \frac{T}{4M}$.

By Freedman's inequality [Freedman, 1975] (and abuse of notation of ϵ), for any $\epsilon > 0$:

$$\Pr\left(|S_T| \ge \epsilon\right) \le 2\exp\left(-\frac{\epsilon^2}{2(V_T + \epsilon/3)} \right) \le 2\exp\left(-\frac{\epsilon^2}{2(T/(4M) + \epsilon/3)} \right).$$

To obtain a high-probability bound, we seek ϵ such that the r.h.s. is at most δ_{MC}. This requires:

$$\frac{\epsilon^2}{2(T/(4M) + \epsilon/3)} \geq \log(2/\delta_{MC}).$$

Rearranging yields the quadratic inequality, we have:

$$\epsilon^2 - \frac{2\log(2/\delta_{MC})}{3}\epsilon - \frac{T\log(2/\delta_{MC})}{2M} \geq 0.$$

The positive root of the corresponding equality is:

$$\epsilon^+ = \frac{\log(2/\delta_{MC})}{3} + \sqrt{\frac{\log^2(2/\delta_{MC})}{9} + \frac{T\log(2/\delta_{MC})}{2M}}.$$

Therefore, with probability at least $1 - \delta_{MC}$:

$$|S_T| \leq \frac{\log(2/\delta_{MC})}{3} + \sqrt{\frac{\log^2(2/\delta_{MC})}{9} + \frac{T\log(2/\delta_{MC})}{2M}} \leq \frac{2\log(2/\delta_{MC})}{3} + \sqrt{\frac{T\log(2/\delta_{MC})}{2M}},$$

where the last (optional) inequality uses $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ to obtain a cleaner rate analysis.

Using Eq. 10 and the algorithm's selection rule $\sigma^2_{t-1}(x_t) \geq \hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}]$,

$$\begin{aligned}
\mathbb{E}_{P_{\mu_{t-1}}}[\sigma^2_{t-1}] &\leq \mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] \\
&= \hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] - \left(\hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] - \mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}]\right) \\
&\leq \sigma^2_{t-1}(x_t) - \left(\hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] - \mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}]\right) \\
&\leq C_1\gamma_T + \left|\sum_{t=1}^{T}\left(\hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}] - \mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}]\right)\right|.
\end{aligned}$$

Summing over $t = 1, \ldots, T$ and applying Lemma C.2 and Part 2:

$$\sum_{t=1}^{T}\mathbb{E}_{P_{\mu_{t-1}}}[\sigma^2_{t-1}] \leq C_1\gamma_T + |S_T| \leq C_1\gamma_T + \frac{2\log(2/\delta_{MC})}{3} + \sqrt{\frac{T\log(2/\delta_{MC})}{2M}} \tag{13}$$

By monotonicity of posterior variance, $\sigma^2_T(x) \leq \sigma^2_{t-1}(x)$ for all $x \in \mathcal{X}$ and $t \leq T$:

$$T \cdot \mathbb{E}_{P_{\mu_T}}[\sigma^2_T] \leq \sum_{t=1}^{T}\mathbb{E}_{P_{\mu_T}}[\sigma^2_{t-1}].$$

To relate $\mathbb{E}_{P_{\mu_T}}[\sigma^2_{t-1}]$ to $\mathbb{E}_{P_{\mu_{t-1}}}[\sigma^2_{t-1}]$, we bound the density ratio. By Lemma 2.3, with probability at least $1 - \delta_{\mathrm{GP}}$, simultaneously for all $t \leq T$ and $x \in \mathcal{X}$ (on the uniform confidence event defined above):

$$|\mu_T(x) - \mu_{t-1}(x)| \leq |f(x) - \mu_T(x)| + |f(x) - \mu_{t-1}(x)| \leq 2\beta_T^{1/2} + 2\epsilon_T,$$

where we used $\sigma_T(x), \sigma_{t-1}(x) \leq 1$.

Applying Lemma 3.2 with $g_1 = \mu_T$, $g_2 = \mu_{t-1}$, $\Delta(x) \leq 2\beta_T^{1/2} + 2\epsilon_T$, and $\Delta_{\max} \leq 2\beta_T^{1/2} + 2\epsilon_T$:

$$\frac{P_{\mu_T}(x)}{P_{\mu_{t-1}}(x)} \leq e^{4|\lambda|(\beta_T^{1/2} + \epsilon_T)}.$$

Therefore:

$$\mathbb{E}_{P_{\mu_T}}[\sigma^2_{t-1}] = \mathbb{E}_{P_{\mu_{t-1}}}\left[\sigma^2_{t-1} \cdot \frac{P_{\mu_T}}{P_{\mu_{t-1}}}\right] \leq e^{4|\lambda|(\beta_T^{1/2} + \epsilon_T)} \cdot \mathbb{E}_{P_{\mu_{t-1}}}[\sigma^2_{t-1}].$$

Summing over t and applying Eq. 13:

$$T \cdot \mathbb{E}_{P_{\mu_T}}[\sigma_T^2] \leq e^{4|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \sum_{t=1}^{T} \mathbb{E}_{P_{\mu_{t-1}}}[\sigma_{t-1}^2]$$
$$\leq e^{4|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(C_1\gamma_T + \frac{2\log(2/\delta_{MC})}{3} + \sqrt{\frac{T\log(2/\delta_{MC})}{2M}} \right).$$

Dividing by T:

$$\mathbb{E}_{P_{\mu_T}}[\sigma_T^2] \leq e^{4|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\gamma_T}{T} + \frac{2\log(2/\delta_{MC})}{3T} + \sqrt{\frac{\log(2/\delta_{MC})}{2MT}} \right).$$

For finite $\mathcal{X}$ with exact threshold computation, $\hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma_{t-1}^2] = \mathbb{E}_{\tilde{p}_{t-1}}[\sigma_{t-1}^2]$ for all t, so $\epsilon_{\mathrm{MC}} = \delta_{\mathrm{MC}} = 0$.

□

Lemma C.4 (Variance Bound for TS-SID-iVAR)**.** *Consider* `TS-SID-iVAR` *selecting $x_t \in \bar{\mathcal{X}}_t$ where $\bar{\mathcal{X}}_t$ is defined in Table 1. If the constraint threshold at each iteration is estimated via M Monte Carlo samples, then for any $\delta_{\mathrm{GP}}, \delta_{\mathrm{MC}} \in (0,1)$, with probability at least $1 - \delta_{\mathrm{GP}} - \delta_{\mathrm{MC}}$:*

$$\mathbb{E}_{P_{\mu_T}}[\sigma_T^2] \leq e^{6|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\gamma_T}{T} + \frac{2\log(2/\delta_{\mathrm{MC}})}{3T} + \sqrt{\frac{\log(2/\delta_{\mathrm{MC}})}{2MT}} \right),$$

where β_T, ϵ_T are as in Lemma 2.3 with failure probability δ_{GP}. For finite $\mathcal{X}$ with exact threshold computation, the $\log(2/\delta_{\mathrm{MC}})$ terms vanish.

Proof. The proof follows a similar structure to Lemma C.3, with an additional step to handle the Thompson sampling randomness.

At each round t, `TS-SID` draws a posterior sample $\tilde{f}_{t-1} \sim f|D_{t-1}$ and defines $\tilde{p}_{t-1} \propto e^{\lambda \tilde{f}_{t-1}+b}$. By Lemma 2.3 applied to the posterior samples $\{\tilde{f}_{t-1}\}_{t\leq T}$ (which, conditionally on D_{t-1}, share the law of $f \mid D_{t-1}$) on the analogous uniform confidence event, with δ_{GP} split equally between this event and the one for f used in the distribution-shift step below, with probability at least $1 - \delta_{\mathrm{GP}}$:

$$|\tilde{f}_{t-1}(x) - \mu_{t-1}(x)| \leq \beta_T^{1/2}\sigma_{t-1}(x) + \epsilon_T \leq \beta_T^{1/2} + \epsilon_T$$

for all $x \in \mathcal{X}$ and $t \leq T$. Applying Lemma 3.2:

$$\frac{P_{\mu_{t-1}}(x)}{P_{\tilde{f}_{t-1}}(x)} \leq e^{2|\lambda|(\beta_T^{1/2}+\epsilon_T)}. \tag{14}$$

This implies:

$$\mathbb{E}_{P_{\tilde{f}_{t-1}}}[\sigma_{t-1}^2] \geq e^{-2|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \mathbb{E}_{P_{\mu_{t-1}}}[\sigma_{t-1}^2].$$

The algorithm estimates $\mathbb{E}_{P_{\tilde{f}_{t-1}}}[\sigma_{t-1}^2]$ via MC samples $x^{(i)} \overset{\text{i.i.d.}}{\sim} P_{\tilde{f}_{t-1}}$. Conditioned on $\tilde{\mathcal{F}}_{t-1}$ (which includes both D_{t-1} and $\tilde{f}_{t-1}$), the distribution $P_{\tilde{f}_{t-1}}$ is fixed, so the estimation errors

$$\xi_t := \hat{\mathbb{E}}_{P_{\tilde{f}_{t-1}}}[\sigma_{t-1}^2] - \mathbb{E}_{P_{\tilde{f}_{t-1}}}[\sigma_{t-1}^2]$$

form a martingale difference sequence with respect to $\{\tilde{\mathcal{F}}_t\}$. Following the same argument as in Lemma C.3 Part 2, with probability at least $1 - \delta_{MC}$:

$$|S_T| \leq \frac{2\log(2/\delta_{MC})}{3} + \sqrt{\frac{T\log(2/\delta_{MC})}{2M}}.$$

The same argument as Lemma C.3 Part 3 applies with $P_{\tilde{f}_{t-1}}$ in place of $\tilde{p}_{t-1}$, yielding:

$$\sum_{t=1}^{T} \mathbb{E}_{P_{\tilde{f}_{t-1}}}[\sigma_{t-1}^2] \leq C_1\gamma_T + |S_T|.$$

Unlike AB-SID, here we need convert from $P_{\tilde{f}_{t-1}}$ to $P_{\mu_{t-1}}$. Using the density ratio defined in Eq. 14:

$$\mathbb{E}_{P_{\mu_{t-1}}}[\sigma_{t-1}^2] = \mathbb{E}_{P_{\tilde{f}_{t-1}}}\left[\sigma_{t-1}^2 \cdot \frac{P_{\mu_{t-1}}}{P_{\tilde{f}_{t-1}}}\right] \leq e^{2|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \mathbb{E}_{P_{\tilde{f}_{t-1}}}[\sigma_{t-1}^2].$$

Summing over t:

$$\sum_{t=1}^{T} \mathbb{E}_{P_{\mu_{t-1}}}[\sigma_{t-1}^2] \leq e^{2|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot (C_1\gamma_T + |S_T|)\,. \tag{15}$$

This additional factor of $e^{2|\lambda|(\beta_T^{1/2}+\epsilon_T)}$ arises from the Thompson sampling approximation and is absent in AB-SID.

The same argument as Lemma C.3 Part 4 gives:

$$\mathbb{E}_{P_{\mu_T}}[\sigma_{t-1}^2] \leq e^{4|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \mathbb{E}_{P_{\mu_{t-1}}}[\sigma_{t-1}^2].$$

Summing over t and applying (15):

$$\begin{aligned} T \cdot \mathbb{E}_{P_{\mu_T}}[\sigma_T^2] &\leq e^{4|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \sum_{t=1}^{T} \mathbb{E}_{P_{\mu_{t-1}}}[\sigma_{t-1}^2] \\ &\leq e^{4|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot e^{2|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot (C_1\gamma_T + |S_T|) \\ &= e^{6|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot (C_1\gamma_T + |S_T|)\,. \end{aligned}$$

Dividing by T and substituting the bound from Part 2:

$$\mathbb{E}_{P_{\mu_T}}[\sigma_T^2] \leq e^{6|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\gamma_T}{T} + \frac{2\log(2/\delta_{MC})}{3T} + \sqrt{\frac{\log(2/\delta_{MC})}{2MT}}\right).$$

For finite $\mathcal{X}$ with exact threshold computation, $\hat{\mathbb{E}}_{\tilde{p}_{t-1}}[\sigma_{t-1}^2] = \mathbb{E}_{P_{\tilde{f}_{t-1}}}[\sigma_{t-1}^2]$ for all t, so $\epsilon_{\mathrm{MC}} = \delta_{\mathrm{MC}} = 0$. □

We are now able to prove Theorem 4.1.

Proof. Set $\delta_{\mathrm{GP}} = \delta_{\mathrm{MC}} = \delta/2$. By union bound, the following holds with probability at least $1-\delta$. By Lemma 3.3 with failure probability $\delta/2$:

$$\mathbb{E}_{P_f}[(f-\mu_T)^2] \leq e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T}+\epsilon_T)} \cdot \left(\beta_T \cdot \mathbb{E}_{P_{\mu_T}}[\sigma_T^2] + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right).$$

Using $\sigma_{\max,T} \leq 1$ (kernel normalization):

$$\mathbb{E}_{P_f}[(f-\mu_T)^2] \leq e^{2|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\beta_T \cdot \mathbb{E}_{P_{\mu_T}}[\sigma_T^2] + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right). \tag{16}$$

By Lemma C.3 (for AB-SID) or Lemma C.4 (for TS-SID) with failure probability $\delta/2$:

$$\mathbb{E}_{P_{\mu_T}}[\sigma_T^2] \leq e^{2(1+\kappa)|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\gamma_T}{T} + \frac{2\log(4/\delta)}{3T} + \sqrt{\frac{\log(4/\delta)}{2MT}}\right),$$

where $\kappa = 1$ for AB-SID(prefactor $e^{4|\lambda|(\cdot)}$) and $\kappa = 2$ for TS-SID (prefactor $e^{6|\lambda|(\cdot)}$). Denoting the *cumulative MC error* $\epsilon_{\mathrm{MC}} := \frac{2\log(4/\delta)}{3T} + \sqrt{\frac{\log(4/\delta)}{2MT}}$ and substituting into (16):

$$\mathbb{E}_{P_f}[(f-\mu_T)^2] \leq e^{2|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\beta_T \cdot e^{2(1+\kappa)|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\gamma_T}{T} + \epsilon_{\mathrm{MC}}\right) + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right).$$

Using $e^{2|\lambda|(\cdot)} \cdot e^{2(1+\kappa)|\lambda|(\cdot)} = e^{2(\kappa+2)|\lambda|(\cdot)}$:

$$\mathbb{E}_{P_f}[(f-\mu_T)^2] \le e^{2(\kappa+2)|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\beta_T\gamma_T}{T} + \beta_T\epsilon_{\mathrm{MC}}\right) + e^{2|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right).$$

Since $e^{2|\lambda|(\cdot)} \le e^{2(\kappa+2)|\lambda|(\cdot)}$ for $\kappa \ge 0$:

$$\mathbb{E}_{P_f}[(f-\mu_T)^2] \le e^{2(\kappa+2)|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\beta_T\gamma_T}{T} + \beta_T\epsilon_{\mathrm{MC}} + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right).$$

□

This enables the following convergence rate analysis:

Proof. Since $\beta_T = \mathcal{O}(\log T)$, we have $\sqrt{\beta_T} = \mathcal{O}(\sqrt{\log T})$. For any $\epsilon > 0$: $e^{c\sqrt{\log T}} = o(T^\epsilon)$ since $\sqrt{\log T} = o(\log T)$. Thus the prefactor satisfies $e^{\mathcal{O}(\sqrt{\beta_T})} = o(T^\epsilon)$.

For the MC error:

- If M is constant: $\beta_T\epsilon_{\mathrm{MC}} = \mathcal{O}\left(\frac{\log T}{\sqrt{T}}\right) = o(T^{-1/2+\epsilon})$.
- If $M = \Omega(T)$: $\beta_T\epsilon_{\mathrm{MC}} = \mathcal{O}\left(\frac{\log T}{T}\right) = o(T^{-1+\epsilon})$.

Combining the prefactor with the dominant term, for any $\epsilon > 0$:

- If M is constant: $o(T^\epsilon) \cdot o(T^{-1/2+\epsilon}) = \mathcal{O}(T^{-1/2+2\epsilon})$.
- If $M = \Omega(T)$: $o(T^\epsilon) \cdot o(T^{-1+\epsilon}) = \mathcal{O}(T^{-1+2\epsilon})$.

Since $\epsilon > 0$ is arbitrary, the result follows. □

C.6 Theorem 4.3

Proof of Theorem 4.3. By Lemma 3.4, conditional on the terminal dataset $\mathcal{D}_T$, we have

$$\mathbb{E}_{f|\mathcal{D}_T}\left[\mathbb{E}_{P_f}[(f-\mu_T)^2]\right] \le e^{2\lambda^2}(1+4\lambda^2) \cdot \mathbb{E}_{P_{\mu_T}}[\sigma_T^2].$$

By Lemma C.3 (for `AB-SID`) or Lemma C.4 (for `TS-SID`), with probability at least $1-\delta$:

$$\mathbb{E}_{P_{\mu_T}}[\sigma_T^2] \le e^{2(1+\kappa)|\lambda|(\beta_T^{1/2}+\epsilon_T)} \cdot \left(\frac{C_1\gamma_T}{T} + \epsilon_{\mathrm{MC}}\right),$$

where $\kappa = 1$ for `AB-SID` and $\kappa = 2$ for `TS-SID`. Substituting completes the proof. □

Proof of Corollary 4.4. For notation simplicity, let

$$R_T(\mathcal{D}_T) := \mathbb{E}_{f|\mathcal{D}_T}[\mathcal{L}(\mu_T)], \qquad C_\lambda := e^{2\lambda^2}(1+4\lambda^2),$$

and let $B_T(\delta)$ denote the right-hand side of Theorem 4.3. By Lemma 3.4 and the kernel normalization $k(x,x) \le 1$, we have

$$0 \le R_T(\mathcal{D}_T) \le C_\lambda \mathbb{E}_{P_{\mu_T}}[\sigma_T^2] \le C_\lambda.$$

For any fixed $\delta \in (0,1)$, define

$$\mathcal{E}_\delta := \{R_T(\mathcal{D}_T) \le B_T(\delta)\}.$$

By Theorem 4.3, $\mathbb{P}(\mathcal{E}_\delta) \ge 1-\delta$. Splitting the expectation over this event and its complement,

$$\begin{aligned}\mathbb{E}_{\mathcal{D}_T}[R_T(\mathcal{D}_T)] &= \mathbb{E}[R_T(\mathcal{D}_T)\mathbf{1}_{\mathcal{E}_\delta}] + \mathbb{E}[R_T(\mathcal{D}_T)\mathbf{1}_{\mathcal{E}_\delta^c}] \\ &\le B_T(\delta)\mathbb{P}(\mathcal{E}_\delta) + C_\lambda\mathbb{P}(\mathcal{E}_\delta^c) \\ &\le B_T(\delta) + C_\lambda\delta.\end{aligned}$$

Finally, by the law of total expectation under the joint Bayesian distribution of f and the adaptive sampling history,

$$\mathbb{E}_{\mathcal{D}_T}[R_T(\mathcal{D}_T)] = \mathbb{E}_f\left[\mathbb{E}_{\mathcal{D}_T|f}[\mathcal{L}(\mu_T)]\right].$$

Here, $\mathcal{D}_T \mid f$ denotes the sampling history generated by the algorithm for a fixed f, including observation noise and algorithmic randomness. This completes the proof. □

C.7 Rate Analysis for Posterior-Averaged and Bayesian MSE Bounds

Posterior-averaged bound. The posterior-averaged bound (Theorem 4.3) achieves the same asymptotic rates as the high-probability bound (Corollary 4.2), under polylogarithmic information-gain growth $\gamma_T = \mathcal{O}((\log T)^q)$ for some fixed $q \geq 0$.
The prefactor $e^{2\lambda^2}(1 + 4\lambda^2)$ is constant in T. For the remaining prefactor, since $\beta_T = \mathcal{O}(\log T)$, we have $e^{2(1+\kappa)|\lambda|(\beta_T^{1/2}+\epsilon_T)} = e^{\mathcal{O}(\sqrt{\log T})} = o(T^\epsilon)$ for any $\epsilon > 0$.
For the MC error:

- If M is constant: $\epsilon_{\mathrm{MC}} = \mathcal{O}(1/\sqrt{T})$.
- If $M = \Omega(T)$: $\epsilon_{\mathrm{MC}} = \mathcal{O}(1/T)$.

Combining the prefactor with the information-gain and MC terms yields $\mathcal{O}(T^{-1/2+\epsilon})$ for constant M and $\mathcal{O}(T^{-1+\epsilon})$ for $M = \Omega(T)$.
Compared to the high-probability realized-MSE bound, the posterior-averaged bound has tighter dependence on the confidence parameter: the main term scales as γ_T/T instead of $\beta_T\gamma_T/T$, and the prefactor exponent is reduced from $2(\kappa + 2)|\lambda|$ to $2(\kappa + 1)|\lambda|$.

Bayesian terminal MSE bound. For Corollary 4.4, choose the auxiliary confidence parameter $\delta_T = (T + 1)^{-2}$. The additive term $C_\lambda\delta_T$ is then $\mathcal{O}(T^{-2})$. Evaluating the confidence-dependent quantities at δ_T gives

$$\beta_T = \mathcal{O}(\log T), \qquad \epsilon_T = \mathcal{O}\left(\frac{\sqrt{\log T}}{T}\right),$$

and, for continuous domains,

$$\epsilon_{\mathrm{MC}} = \mathcal{O}\left(\frac{\log T}{T} + \sqrt{\frac{\log T}{MT}}\right).$$

The exponential prefactor remains $e^{\mathcal{O}(\sqrt{\log T})} = o(T^\epsilon)$ for every $\epsilon > 0$. The fully averaged Bayesian terminal MSE is therefore bounded by

$$\mathcal{O}\left[e^{\mathcal{O}(\sqrt{\log T})}\left(\frac{\gamma_T}{T} + \frac{\log T}{T} + \sqrt{\frac{\log T}{MT}}\right) + T^{-2}\right].$$

Under the same polylogarithmic information-gain condition, this recovers the rates $\mathcal{O}(T^{-1/2+\epsilon})$ for constant M and $\mathcal{O}(T^{-1+\epsilon})$ for $M = \Omega(T)$, now for the unconditional Bayesian expected MSE. For finite domains with exact threshold computation, $\epsilon_T = \epsilon_{\mathrm{MC}} = 0$.
Compared with the high-probability realized-MSE bound in Theorem 4.1, the posterior-averaged bound has γ_T/T rather than $\beta_T\gamma_T/T$ as its main term and reduces the coefficient in the exponential prefactor from $2(\kappa + 2)|\lambda|$ to $2(\kappa + 1)|\lambda|$.

C.8 Proposition 4.7

Analogous to our previous theoretical analysis, we first establish an upper bound for fixed-threshold methods, then derive the convergence rate. The key observation is that any algorithm querying points with $\sigma^2_{t-1}(x_t) \geq \eta$ must exit this regime within $\mathcal{O}(\gamma_T/\eta)$ iterations by the information gain bound, after which the posterior variance remains uniformly below η by monotonicity.

Lemma C.5 (High-Probability MSE Upper Bound for Fixed-Threshold Methods)**.** *Consider any active learning algorithm that selects x_t satisfying $\sigma^2_{t-1}(x_t) \geq \eta$ when $\bar{\mathcal{X}}'_t := \{x : \sigma^2_{t-1}(x) \geq \eta\} \neq \phi$, and selects $x_t = \arg\max_{x\in\mathcal{X}} \sigma_{t-1}(x)$ otherwise, where $\eta \in (0, 1]$ is a fixed threshold. Under Assumptions 2.1 and 3.1, let $\mathcal{X} \subset [0, r]^d$ be compact. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$:*

$$L(\mu_T) \leq e^{2|\lambda|(\beta_T^{1/2}\sqrt{\eta}+\epsilon_T)} \cdot \left(\beta_T\eta + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right)$$

provided $\eta > C_1\gamma_T/T$, where β_T, ϵ_T are as in Lemma 2.3 with failure probability δ, and $C_1 = 2/\log(1 + \tau^{-2})$. For finite $\mathcal{X}$, $\epsilon_T = 0$.

Proof. According to the algorithm, we define the AL iteration as the following two phases:

$$T^{(1)} := \{t \le T : \max_{x\in\mathcal{X}} \sigma^2_{t-1}(x) \ge \eta\}, \quad T^{(2)} := \{t \le T : \max_{x\in\mathcal{X}} \sigma^2_{t-1}(x) < \eta\}.$$

In Phase 1, by the algorithm's selection rule, any query point satisfies $\sigma^2_{t-1}(x_t) \ge \eta$. By Lemma C.2:

$$|T^{(1)}| \cdot \eta \le \sum_{t\in T^{(1)}} \sigma^2_{t-1}(x_t) \le \sum_{t=1}^{T} \sigma^2_{t-1}(x_t) \le C_1\gamma_T.$$

Thus $|T^{(1)}| \le C_1\gamma_T/\eta < T$ under the stated condition on η, ensuring Phase 2 is non-empty.

Once Phase 2 begins, $\max_{x\in\mathcal{X}} \sigma^2_{t-1}(x) < \eta$ by definition. Since posterior variance is monotonically non-increasing with additional observations, $\sigma_{\max,T} < \sqrt{\eta}$ and $\mathbb{E}_{P_{\mu_T}}[\sigma^2_T(x)] < \eta$.

By Lemma 2.3, with probability at least $1-\delta$, the confidence bound holds uniformly. Applying Lemma 3.3:

$$L(\mu_T) \le e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T}+\epsilon_T)} \cdot \left(\beta_T \cdot \mathbb{E}_{P_{\mu_T}}[\sigma^2_T(x)] + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right).$$

Substituting $\sigma_{\max,T} < \sqrt{\eta}$ and $\mathbb{E}_{P_{\mu_T}}[\sigma^2_T] < \eta$ yields the result. □

Now we start to prove proposition 4.7.

Proof. By Lemma C.5, Phase 1 contains at most $|T^{(1)}| \le C_1\gamma_T/\eta$ iterations. In Phase 2, the algorithm selects $x_t = \arg\max_{x\in\mathcal{X}} \sigma_{t-1}(x)$, so $\sigma^2_{\max,t}$ is non-increasing for $t \in T^{(2)}$. Applying the information gain bound to Phase 2:

$$|T^{(2)}| \cdot \sigma^2_{\max,T} \le \sum_{t\in T^{(2)}} \sigma^2_{t-1}(x_t) \le C_1\gamma_T.$$

Since $|T^{(2)}| = T - |T^{(1)}| \ge T - C_1\gamma_T/\eta$, we obtain:

$$\sigma^2_{\max,T} \le \frac{C_1\gamma_T}{T - C_1\gamma_T/\eta}.$$

For any fixed $\eta \in (0,1]$, since $\gamma_T = o(T)$ for standard kernels, we have $T - C_1\gamma_T/\eta \ge T/2$ for sufficiently large T, yielding $\sigma^2_{\max,T} \le 2C_1\gamma_T/T$.

Applying Lemma 3.3 with $\sigma_{\max,T} \le \sqrt{2C_1\gamma_T/T}$ and $\mathbb{E}_{P_{\mu_T}}[\sigma^2_T] \le \sigma^2_{\max,T} \le 2C_1\gamma_T/T$:

$$L(\mu_T) \le e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T}+\epsilon_T)} \cdot \left(\beta_T \cdot \frac{2C_1\gamma_T}{T} + 2\beta_T^{1/2}\epsilon_T + \epsilon_T^2\right).$$

For the prefactor, since $\beta_T = \mathcal{O}(\log T)$, $\gamma_T = \mathcal{O}((\log T)^c)$ for kernel-dependent constant c, and $\epsilon_T = \mathcal{O}(1/T)$:

$$\beta_T^{1/2}\sigma_{\max,T} = \mathcal{O}\left(\sqrt{\log T} \cdot \sqrt{\frac{(\log T)^c}{T}}\right) = \mathcal{O}\left(\frac{(\log T)^{(c+1)/2}}{\sqrt{T}}\right) = o(1).$$

Thus $e^{2|\lambda|(\beta_T^{1/2}\sigma_{\max,T}+\epsilon_T)} = 1 + o(1)$. The dominant term is $\beta_T \cdot 2C_1\gamma_T/T = \mathcal{O}((\log T)^{c+1}/T)$, while the remaining terms satisfy $2\beta_T^{1/2}\epsilon_T + \epsilon_T^2 = \mathcal{O}(\sqrt{\log T}/T) = o((\log T)^{c+1}/T)$. Therefore $L(\mu_T) = \mathcal{O}(T^{-1+\epsilon})$ for arbitrarily small $\epsilon > 0$. □

C.9 Robustness to Boltzmann Parameter Misspecification

Proposition C.6 (Robustness to Boltzmann parameter misspecification)**.** *Let*

$$P_f^{(i)}(x) = \frac{1}{Z_f^{(i)}} \exp\big(\lambda_i f(x) + b_i(x)\big), \qquad i \in \{1,2\},$$

denote the self-induced distributions associated with parameters (λ_i, b_i). For parameters (λ_1, b_1), (λ_2, b_2) that has bounded difference

$$|\lambda_2 - \lambda_1| \le \epsilon_\lambda, \qquad \|b_2 - b_1\|_\infty \le \epsilon_b,$$

and additionally with bounded function range over $\mathcal{X}$:

$$R_f := \sup_{x\in\mathcal{X}} f(x) - \inf_{x\in\mathcal{X}} f(x) < \infty.$$

Then, for any surrogate $\hat{f}$,

$$\mathbb{E}_{x\sim P_f^{(2)}}\left[\left(\hat{f}(x) - f(x)\right)^2\right] \le \exp(\epsilon_\lambda R_f + 2\epsilon_b)\,\mathbb{E}_{x\sim P_f^{(1)}}\left[\left(\hat{f}(x) - f(x)\right)^2\right].$$

Consequently, under the conditions of Theorem 4.1, its high-probability guarantee and the convergence rates in Corollary 4.2 continue to hold under bounded parameter misspecification, up to the multiplicative constant above.

Proof. Define

$$h(x) := (\lambda_2 - \lambda_1)f(x) + b_2(x) - b_1(x).$$

The density ratio between the test-time and acquisition-time self-induced distributions can be written as

$$\frac{P_f^{(2)}(x)}{P_f^{(1)}(x)} = \frac{Z_f^{(1)}}{Z_f^{(2)}} e^{h(x)}.$$

Moreover,

$$Z_f^{(2)} = \int_{\mathcal{X}} e^{\lambda_1 f(x) + b_1(x)} e^{h(x)}\,dx \ge e^{\inf_{x\in\mathcal{X}} h(x)} Z_f^{(1)}.$$

Hence,

$$\frac{P_f^{(2)}(x)}{P_f^{(1)}(x)} \le \exp\left(h(x) - \inf_{x'\in\mathcal{X}} h(x')\right) \le \exp\left(\sup_{x\in\mathcal{X}} h(x) - \inf_{x\in\mathcal{X}} h(x)\right).$$

By the definition of h,

$$\begin{aligned}\sup_{x\in\mathcal{X}} h(x) - \inf_{x\in\mathcal{X}} h(x) &\le |\lambda_2 - \lambda_1|\left(\sup_{x\in\mathcal{X}} f(x) - \inf_{x\in\mathcal{X}} f(x)\right) \\ &\quad + \sup_{x\in\mathcal{X}}\big(b_2(x) - b_1(x)\big) - \inf_{x\in\mathcal{X}}\big(b_2(x) - b_1(x)\big) \\ &\le \epsilon_\lambda R_f + 2\epsilon_b.\end{aligned}$$

Therefore,

$$\frac{P_f^{(2)}(x)}{P_f^{(1)}(x)} \le \exp(\epsilon_\lambda R_f + 2\epsilon_b) \qquad \text{for all } x \in \mathcal{X}.$$

Changing measure from $P_f^{(2)}$ to $P_f^{(1)}$ now gives

$$\begin{aligned}\mathbb{E}_{x\sim P_f^{(2)}}\left[\left(\hat{f}(x) - f(x)\right)^2\right] &= \mathbb{E}_{x\sim P_f^{(1)}}\left[\frac{P_f^{(2)}(x)}{P_f^{(1)}(x)}\left(\hat{f}(x) - f(x)\right)^2\right] \\ &\le \exp(\epsilon_\lambda R_f + 2\epsilon_b)\,\mathbb{E}_{x\sim P_f^{(1)}}\left[\left(\hat{f}(x) - f(x)\right)^2\right].\end{aligned}$$

The multiplicative factor is independent of T for fixed ϵ_λ, ϵ_b, and R_f. Therefore, applying Theorem 4.1 to the final expectation preserves the convergence rates in Corollary 4.2. □

D Experimental Details and Additional Experimental Investigations

D.1 Experimental Setup Details

For all methods except the molecular drug discovery example, we model the objective function using a GP with a Matérn-5/2 kernel, estimate hyperparameters via maximum a posteriori (MAP) with the default priors from Hvarfner et al. [2024], and fix observation noise at $\tau^2 = 10^{-4}$, except for the GP prior problem where we assume the true kernel hyperparameters are known. We pre-normalize inputs to $[0, 1]^d$ and standardize outputs online. For acquisition function optimization over continuous domains, we use SLSQP [Kraft, 1988] to enforce the variance constraint defining $\bar{\mathcal{X}}_t$ [7], and L-BFGS-B [Byrd et al., 1995] otherwise, both via `BoTorch` [Balandat et al., 2020].

For continuous $\mathcal{X}$, we approximate the acquisition function integral via Sequential Monte Carlo (SMC), which has demonstrated strong performance for sampling Boltzmann-form distributions [Akhound-Sadegh et al., 2024, 2025]. Specifically, we use $N = 1000d$ particles, 10 temperature levels, and 10 random walk Metropolis steps per level. For fair comparison, the same number of reference points is used for all acquisition functions requiring integration over the input space. A sensitivity study to particle number in terms of performance is also reported in Appendix D.7.

D.2 Generic Adaptation of Heuristics Active Learning Algorithms

Several heuristic active learning approaches have been proposed in the computational materials science and potential energy modeling literature for efficient assembly of data for the training of machine-learned interatomic potentials (MLIP); See, e.g., [Li et al., 2015, Podryabinkin and Shapeev, 2017, Vandermause et al., 2020, Jinnouchi et al., 2020, van der Oord et al., 2023, Kulichenko et al., 2023, Duschatko et al., 2024]. These MLIPs are required to be accurate for configurations that are probable under the specific statistical ensembles targeted in the molecular dynamics simulation in which these MLIPs are intended to be used, but accuracy of the MLIP on configurations that are very improbable within this ensemble is of less importance. As such, the problem of fitting an MLIP resembles the here presented SIDAL problem and heuristic active learning algorithms designed for these problems are natural contenders to be considered in benchmark problems. Using the notation introduced in the main body of the paper, the generic form of a query of these algorithms consists of the following sampling procedure:

1. Run a single-chain Markov Chain Monte Carlo algorithm (e.g., a stochastic thermostat or constant-energy molecular dynamics) targeting a surrogate density $\tilde{p}_{t-1}$ of the generic form
$$\tilde{p}_{t-1}(x) \propto \exp\big(\lambda\mu_{t-1}(x) + \lambda h(\sigma^2_{t-1}(x)) + b(x)\big),$$
where $h(\sigma^2)$ is a smooth transformation of σ^2.
2. Once the Markov chain reaches a configuration where a method-dependent uncertainty measure $UC(x)$ exceeds a preset threshold, $\eta > 0$, the Markov chain terminates and the current sample is queried.

Algorithms proposed in the literature differ in how these two steps are implemented, specifically in terms of the choice of the surrogate density $\tilde{p}_{t-1}(x)$, and the uncertainty measure UC (see Table 2), are often motivated by problem-specific properties and constraints and thus are not necessarily suited for the generic SIDAL setting considered here.

We therefore adapt the existing heuristics in the following Generic Heuristic Active Learning (GHAL) approach to fit the here considered generic SIDAL setting:

1. Run single-chain MCMC targeting $\tilde{p}_{t-1}(x) \propto \exp(\lambda(\mu_{t-1}(x) - \tau\sigma_{t-1}(x)))$ initialized at a point in input space close to the current training set. (This is to ensure that the uncertainty threshold is not exceeded at initialization.)
2. If $\sigma^2_{t-1}(x) > \eta$, select this sample x. If the threshold is not exceeded after the maximum number of steps of the Markov chain, select the sample with highest predictive variance $\sigma^2_{t-1}(x)$.

[7] In practice, constrained acquisition optimization was numerically stable with feasible initializations. On the Branin benchmark, only 9 out of 3,000 selections by AB-SID-iVAR across 30 seeds exhibited constraint violations, all within the numerical tolerance of 10^{-8}, with no optimizer failures observed.

3. Query $y_t = f(x_t) + \epsilon_t$ and update the GP posterior

Note that for $\tau = 0$, this methods closely resembles the query strategy in the active learning loop of FLARE [Vandermause et al., 2020], whereas for $\tau > 0$, this approach resembles the Hyperactive Learning HAL method [van der Oord et al., 2023] up to the form of the uncertainty measure UC used in the thresholding step. Here, both $\tau > 0$ and the threshold $\eta > 0$ are hyper-parameters. To the best of our knowledge, there are no systematic apriori tuning strategies described in the literature. Therefore, we resort to randomly sample the values of these hyperparameters uniformly from a range of plausible values at the beginning of each active learning loop, i.e., $\tau \sim \mathcal{U}([0, 3])$ and $\eta \sim \mathcal{U}([1, 5])$.

We mention that both HAL and the here proposed methods modify the sampling distribution, but with fundamentally different mechanisms and goals. HAL uses a *single-chain reactive* approach: sampling from a heuristically biased distribution and triggering queries when local uncertainty exceeds a threshold. In contrast, our method uses *multi-particle proactive* optimization: maintaining N particles via SMC on the surrogate $\tilde{p}_{t-1}$ of the target distribution $P_f \propto \exp(\lambda f)$, then selecting points that maximize variance reduction under this distribution. The linear $+\tau\sigma$ term in HAL is a heuristic exploration bonus, whereas our $+\frac{\lambda^2}{2}\sigma^2$ term arises from the Gaussian MGF identity $\mathbb{E}[e^{\lambda f}] = e^{\lambda\mu + \frac{\lambda^2}{2}\sigma^2}$, naturally weighting uncertainty by its relevance to the rare-event distribution.

Table 2: Summary of heuristic active learning methods for interatomic potentials. Here, $\beta = -\lambda > 0$, so that $1/\beta$ corresponds, up to a constant pre-factor, to the temperature of the simulated system. f is a vector-valued quantity measuring the uncertainty in force predictions of the MLIP.

Method	**Sampling Distribution**	**Uncertainty measure**
FLARE [Vandermause et al., 2020]	$\exp(-\beta\mu)$	$UC = \sigma$
HAL [van der Oord et al., 2023]	$\exp(-\beta(\mu - \tau\sigma))$	$UC = \max \mathrm{softmax}(f)$
UDD-AL [Kulichenko et al., 2023]	$\exp(-\beta(\mu + E_{\text{bias}}(\sigma)))$	$UC = \sigma$

D.3 Discrete Experimental Results

We evaluate the same benchmark functions on discretized domains, where $\mathcal{X}$ consists of a finite grid of candidate points. This setting allows exact computation of the constraint threshold $\mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}]$ without Monte Carlo approximation, corresponding to the $\epsilon_{\mathrm{MC}} = 0$ case in our theoretical analysis.

Results are shown in Figure 7. The overall trends mirror the continuous setting: `AB-SID-iVAR-`$\bar{\mathcal{X}}$ consistently achieves the lowest weighted MSE across benchmarks, with the gap widening in higher dimensions. The relative ordering of methods remains largely unchanged from the continuous experiments (Figure 3), confirming that our findings are not artifacts of MCMC approximation quality.

Notably, the MC variant `TS-SID-iVAR` remains competitive in general but exhibits high variance on multimodal targets such as Branin. This occurs because occasionally, when the Thompson sample concentrates on a subset of modes, the constraint threshold $\mathbb{E}_{\tilde{p}_{t-1}}[\sigma^2_{t-1}]$ becomes small, resulting in a vast feasible region $\bar{\mathcal{X}}_t$ that fails to encourage sufficient exploration. This empirical finding does not contradict our theoretical guarantees (Theorem 4.1), which are asymptotic in high probability as $T \to \infty$. In discrete domains, this can be empirically mitigated by adopting a more conservative threshold such as the domain-averaged variance.

D.4 Ablation Study

We compare: (1) the surrogate form (`AB-SID` and `TS-SID` versus the naive plug-in μ-SID using $e^{\lambda\mu_{t-1}}$), tested with both our `iVAR` acquisition (denoted μ-`SID-iVAR`) and EPIG (denoted μ-`SID-EPIG`) to verify that the choice of acquisition function alone cannot compensate for a plug-in surrogate, and (2) the presence of the constraint set $\bar{\mathcal{X}}$ versus unconstrained optimization over full $\mathcal{X}$ (denoted ϕ). Results are shown in Figure 8.

On functions where the target distribution does not exhibit strongly separated modes (Forrester, Gramacy 1D, Hartmann), all variants perform comparably. However, when the target distribution

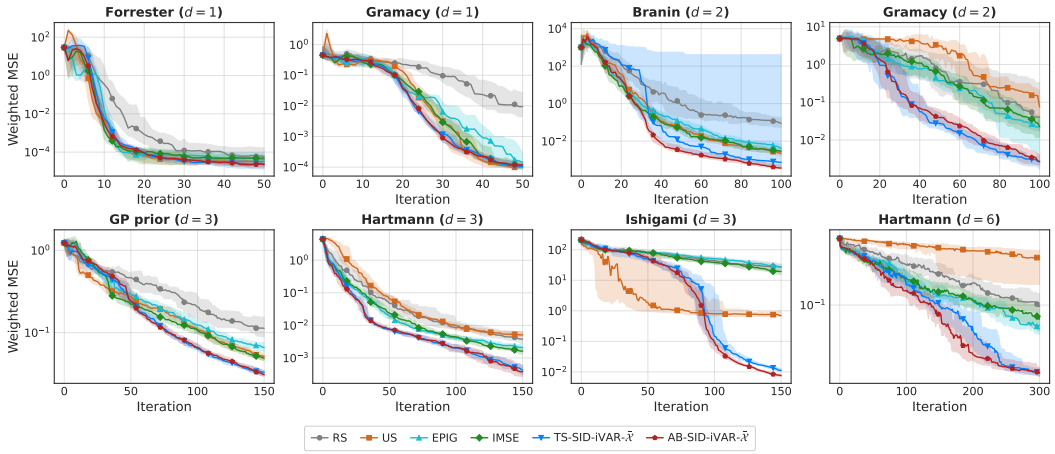


Figure 7: Weighted MSE convergence on synthetic benchmarks with discrete input domains across dimensions $d = 1$ to $d = 6$. Each panel shows median performance over 30 random seeds (shaded regions: interquartile range). Results mirror the continuous setting (Figure 3), with AB-SID-iVAR-$\bar{\mathcal{X}}$ consistently achieving the lowest weighted MSE.

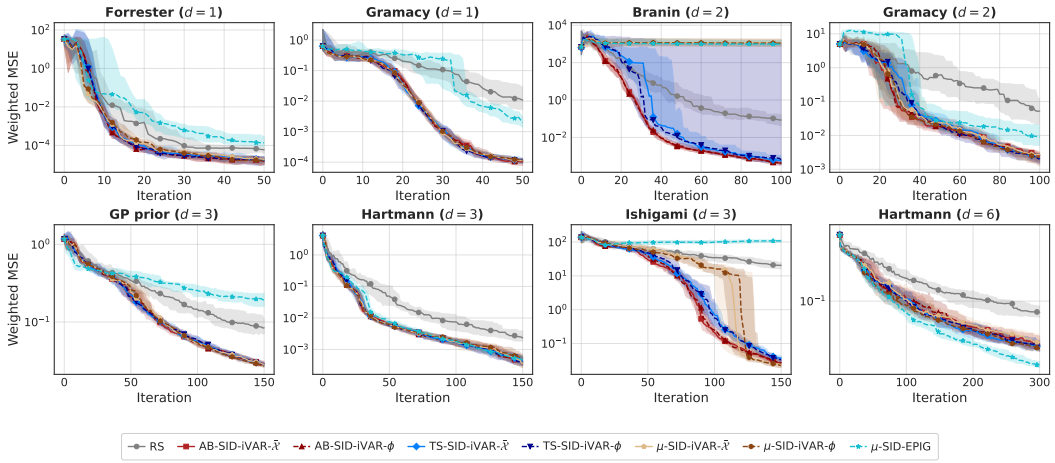


Figure 8: Ablation study comparing surrogate forms, acquisition strategies, and constraint sets. $\bar{\mathcal{X}}$: methods with the variance-threshold constraint; ϕ: unconstrained optimization over $\mathcal{X}$. The μ-SID family uses the plug-in surrogate $e^{\lambda\mu_{t-1}}$; μ-SID-EPIG replaces the iVAR criterion with EPIG on the same surrogate. Random sampling (RS) is included as a baseline. On problems dominated by a broad basin of comparable local minima (Hartmann $d = 3, 6$), plug-in μ-SID variants achieve competitive convergence. In more general settings, particularly when the target distribution has well-separated modes (Branin) or strongly oscillatory structure (Ishigami), AB/TS-SID-iVAR yield more robust convergence.

concentrates on well-separated regions or exhibits oscillatory structure, clear failure modes emerge. On Branin, where the target distribution P_f is bimodal, both μ-SID-iVAR and μ-SID-EPIG fail entirely as the plug-in surrogate $e^{\lambda\mu_{t-1}}$ commits to whichever mode is discovered first and cannot explore alternatives, resulting in weighted MSE that remains near its initial value throughout. On Ishigami, both plug-in variants converge substantially later than AB/TS-SID-iVAR. AB-SID and TS-SID avoid these failures by incorporating posterior uncertainty into the surrogate via the variance-inflation term $\frac{\lambda^2}{2}\sigma^2_{t-1}$, allowing probability mass to spread across uncertain regions.

The constraint set $\bar{\mathcal{X}}$ provides complementary robustness. On Branin and Ishigami, the unconstrained variants (ϕ) exhibit higher variance across seeds, as they may exploit the current surrogate's modes without sufficient exploration.

D.5 Real World Problems

D.5.1 Potential Energy Surface Modeling

In all four systems below we report λ; for $\lambda < 0$, the physical temperature relates to λ via $\mathcal{T} = -(k_{\mathrm{B}}\lambda)^{-1}$, where $k_{\mathrm{B}} = 8.617 \times 10^{-5}$ eV/K. Energies are reported in eV throughout.

$\mathbf{H_2}$ on $\mathbf{Cu(100)}$. We fit the potential energy surface of a hydrogen molecule, H_2, on top of a $\mathrm{Cu}(100)$ surface. The chemistry of primary interest is the dissociative adsorption of H_2 on the metal surface [Smits and Somers, 2023], which is well described by two reaction coordinates: the molecule's distance from the surface, z, and the interatomic distance between the two H atoms, d. We fit E as a function of (z, d). As ground truth we use a custom MACE model [Stark et al., 2024] trained on H_2 reactive surface configurations. The simulation cell follows the setup of [Stark et al., 2023, 2025]: a single H_2 molecule above a 3×3, six-layer $\mathrm{Cu}(100)$ slab consisting of 54 Cu atoms in face-centered cubic (FCC) lattice structure, within a periodic $7.749 \times 7.748 \times 49.132\,\text{Å}^3$ unit cell. We use a conservative weighting $|\lambda| = 1$ (treated as an abstract concentration parameter rather than a physical temperature), yielding broad coverage of both the bonded and dissociated states and the barrier between them.

H on $\mathbf{Cu_{13}}$. We consider the adsorption of a hydrogen probe atom on a frozen 13-atom copper cluster in the icosahedral global-minimum geometry [Yang et al., 2006] in vacuum. We parametrise the probe position in spherical coordinates (r, u, ϕ) centered at the cluster centroid,

$$\mathbf{x} = \left(r\sqrt{1-u^2}\cos\phi,\ r\sqrt{1-u^2}\sin\phi,\ ru\right) \in \mathbb{R}^3,$$

where $u = \cos\theta$ is used in place of the polar angle θ so that uniform sampling on (u, ϕ) induces the uniform measure on the sphere. We restrict $r \in [0.5, 6.0]\,\text{Å}$ to cover surface and near-surface configurations, and exploit the icosahedral symmetry of the cluster by restricting (u, ϕ) to a single fundamental domain $[0.7947, 1] \times [0, \pi/5]$. Energies are evaluated using the MACE-MPA-0 foundation model [Batatia et al., 2025]. We use $|\lambda| = 8.55\,\text{eV}^{-1}$, corresponding to $\mathcal{T} \approx 1358\,\text{K}$ (close to the melting point of copper), at which the Boltzmann target concentrates on the cluster's surface adsorption sites.

$\mathbf{H_2O}$ on Pt. The energy landscape of an $\mathrm{H}_2\mathrm{O}$ molecule on a platinum surface features multiple physisorption wells, chemisorption minima, and dissociation pathways [Schnur and Groß, 2009]. We model the metal surface as a three-layer ABC-stacked $\mathrm{Pt}(111)$ slab in a 3×3 in-plane supercell with two-dimensional periodic boundary conditions ($\sim 8\,\text{Å}$ lateral spacing between water images). A single rigid $\mathrm{H}_2\mathrm{O}$ molecule is placed above the slab and parametrised by 6 degrees of freedom: the oxygen's lateral position (x, y) within the $p3m1$ fundamental triangle of the surface unit cell, its height $z \in [0.5, 5.0]\,\text{Å}$ above the topmost Pt layer, and the molecular orientation given by ZYZ Euler angles $(\alpha, \cos\beta, \gamma)$. As in the spherical parametrisation above, $\cos\beta$ replaces β for uniform sampling on $SO(3)$; γ is restricted to $[0, \pi)$ to account for the C_{2v} symmetry of $\mathrm{H}_2\mathrm{O}$. Energies are evaluated using the MACE-MPA-0 foundation model [Batatia et al., 2025]. We use $|\lambda| = 38.68\,\text{eV}^{-1}$, corresponding to room temperature $\mathcal{T} = 300\,\text{K}$, the standard condition for water-on-platinum experiments.

Si Crystal. We fit the potential energy of silicon in a highly symmetric diamond cubic structure, modeled by two atoms in a cubic unit cell with periodic boundary conditions. By translation invariance, the relevant degrees of freedom reduce to the displacement vector between the two atoms. As ground truth we use the Stillinger–Weber potential [Stillinger and Weber, 1985], a canonical empirical potential for tetrahedrally-bonded silicon. We use $|\lambda| = 11.60\,\text{eV}^{-1}$, corresponding to $\mathcal{T} = 1000\,\text{K}$, a typical Si annealing / chemical vapor deposition temperature, so that the Boltzmann distribution covers a wide range of thermalised crystal states. We emphasise that this setup is a minimal benchmark: the small cell, fixed lattice vectors, and absence of defects restrict the configuration space to thermal fluctuations around a single crystalline minimum.

D.5.2 Molecular Drug Discovery

The candidate pool consists of approximately 20k molecules from the GuacaMol benchmark [Brown et al., 2019], each represented by a 2048-bit Morgan fingerprint. We use a GP with a Tanimoto kernel

as our surrogate model and consider two commonly used multi-property drug-likeness scores (Median 1 and Median 2) as learning objectives, testing $\lambda \in \{25, 75\}$ to vary the concentration of the target distribution (Figure 5, top row). This setting lies outside our theoretical assumptions as the Tanimoto kernel is neither stationary nor differentiable; nevertheless, the problem is of practical interest and serves as a challenging test of whether `AB-SID-iVAR` remains effective beyond the provable regime. We omit TS-SID-iVAR as drawing posterior samples over the full 20k-point pool is computationally prohibitive, though efficient approximations exist for Tanimoto kernels [Tripp et al., 2023]. We also omit `GHAL` as it is designed for continuous domains.

D.6 Runtime Report and Compute Resources

Compute setup. All experiments were run on an internal SLURM cluster. Each run was allocated 4 CPU cores and 32 GB of RAM on internal computational nodes without usage of any GPUs.

Per-iteration runtime. We report the per-iteration runtime (in seconds) for each method in Table 3, averaged across seeds. For our methods (AB-$\bar{\mathcal{X}}$ and TS-$\bar{\mathcal{X}}$), we decompose the runtime into three components: GP fitting, acquisition optimization, and SMC sampling.

Several observations can be made from this analysis. First, GP fitting contributes negligibly to total runtime across all methods. Second, acquisition optimization dominates the computational cost for our methods. This is expected: optimizing the SID-iVAR objective requires evaluating the fantasy variance at all reference points for each candidate, and the $\bar{\mathcal{X}}$ constraint introduces an additional posterior evaluation per optimization step, effectively doubling the per-step cost compared to methods like IMSE, which use an unconstrained L-BFGS-B optimizer. Third, SMC sampling adds modest overhead in continuous mode but is not required for discrete domains where the candidate set is given.

Table 3: Runtime profile (seconds per iteration, mean $\pm$ std across seeds).

	d	IMSE	AB-$\bar{\mathcal{X}}$	AB-$\bar{\mathcal{X}}$ Breakdown			TS-$\bar{\mathcal{X}}$	TS-$\bar{\mathcal{X}}$ Breakdown		
				GP	Acq Opt	SMC		GP	Acq Opt	SMC
Continuous	1	0.30±0.04	2.3±0.6	0.06±0.01	2.2±0.6	0.62±0.16	3.6±1.2	0.05±0.01	3.5±1.1	1.9±0.7
	2	0.68±0.09	3.7±1.1	0.06±0.02	3.3±1.0	0.84±0.32	6.0±1.8	0.06±0.01	5.6±1.7	3.2±1.1
	3	1.5±0.2	6.1±1.2	0.08±0.01	5.2±1.2	1.2±0.3	12.8±2.3	0.09±0.02	11.9±2.2	7.2±1.4
	6	4.9±0.7	16.3±3.4	0.15±0.03	13.3±2.9	3.6±0.8	26.2±5.0	0.14±0.03	23.3±4.5	14.0±2.7
Discrete	1	2.4±0.6	3.7±0.5	0.06±0.01	3.6±0.5	NA	4.0±0.5	0.05±0.01	3.8±0.5	NA
	2	3.2±0.5	5.3±1.1	0.05±0.01	5.2±1.1	NA	4.6±2.4	0.06±0.02	4.4±2.4	NA
	3	2.6±0.7	5.9±1.3	0.08±0.02	5.7±1.3	NA	6.1±1.1	0.08±0.01	5.9±1.1	NA
	6	0.81±0.16	1.8±0.3	0.16±0.03	1.5±0.3	NA	2.0±0.4	0.15±0.03	1.7±0.3	NA

D.7 Sensitivity Analysis of MCMC sample size

We investigate the sensitivity of our methods to the SMC particle size N on GP prior functions across $d \in \{2, 4, 6, 8\}$, testing $N \in \{250d, 500d, 1000d, 1500d\}$. Figure 9 reports the weighted MSE convergence (top) and per-iteration MCMC runtime (bottom). Convergence is essentially unchanged across particle counts for both AB-$\bar{\mathcal{X}}$ and TS-$\bar{\mathcal{X}}$, while runtime scales roughly linearly with N. TS-$\bar{\mathcal{X}}$ incurs higher cost than AB-$\bar{\mathcal{X}}$ at matched N due to pathwise-conditioned posterior sampling. This justifies $N = 250d$ as a sufficient default in practice.

D.8 Sensitivity Analysis of λ and b

We investigate sensitivity to the temperature parameter λ and bias function $b(x)$ on functions drawn from a 2 dimensional GP prior across $\lambda \in \{-0.5, -1.0, -2.0, -5.0\}$ with either constant or function form bias $b(x)$ drawn from an independent GP prior.

The results are illustrated in Figure 10. As $|\lambda|$ increases, the target distribution $P_f \propto \exp(\lambda f + b)$ concentrates around the minima of f. Under weak concentration ($\lambda = -0.5$), all methods exhibit comparable performance. However, the performance benefit gets more obvious as $|\lambda|$ grows. This reflects the increasing inefficiency of uniform or uncertainty-driven sampling when the evaluation metric concentrates on a shrinking subset of the domain.

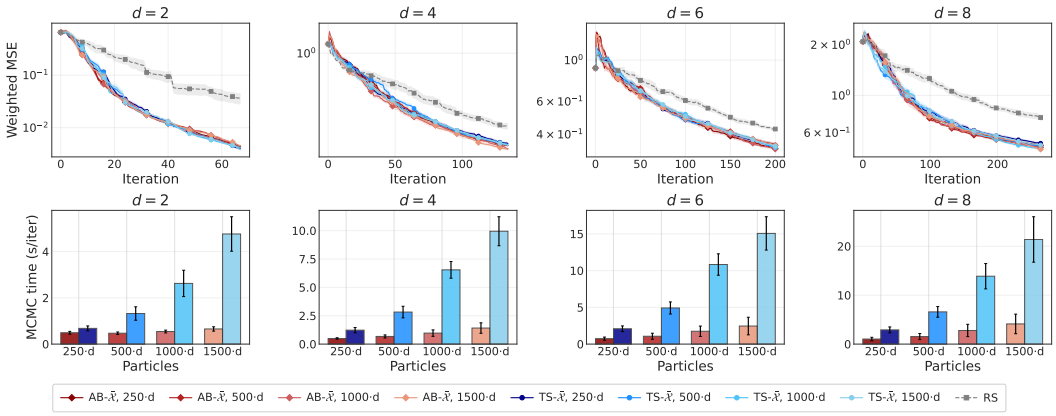


Figure 9: Sensitivity of SMC particle size $N \in \{250d, 500d, 1000d, 1500d\}$ on GP prior functions across $d \in \{2, 4, 6, 8\}$. Top: weighted MSE convergence; curves overlap across particle counts. Bottom: per-iteration MCMC time, scaling linearly with N.

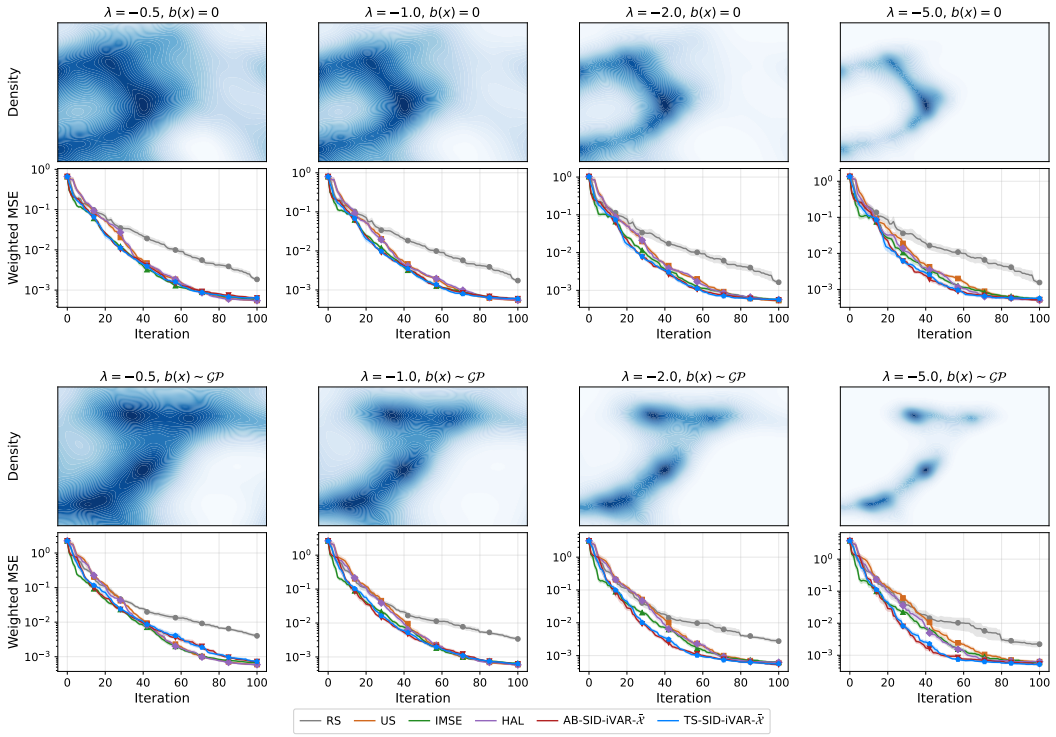


Figure 10: Sensitivity analysis of temperature λ and bias $b(x)$ on 2-dimensional GP prior samples. Top (two) row: constant bias $b(x) = 0$; bottom (two) row: function form bias $b(x) \sim \mathcal{GP}$. Columns correspond to $\lambda \in \{-0.5, -1.0, -2.0, -5.0\}$ from left to right. Density plots show the ground truth with high-probability regions under $P_f \propto \exp(\lambda f + b)$ shaded in blue. As $|\lambda|$ increases, the target concentrates and SID-aware methods show greater relative advantage.

The introduction of spatially-varying bias $b(x) \sim \mathcal{GP}$ (bottom row) produces more complex target distributions, yet the relative ordering among methods is preserved, confirming that our framework extends beyond the $b(x) = 0$ setting assumed in the main experiments.

D.9 Molecular Drug Discovery

To complement the quantitative R^2 analysis, Figure 11 visualizes prediction quality across the entire molecule pool. Each panel plots the GP posterior mean μ against the true score y, with points colored by their Boltzmann weight under P_λ. This visualization directly reveals *where* in function-value space each method succeeds or fails.

For `AB-SID-iVAR`, dark points (high Boltzmann weight) cluster tightly around the diagonal across all settings, indicating that the model has learned the high-value region accurately. In contrast, baseline methods exhibit two failure modes visible in the plots: (i) vertical spread of dark points, indicating high variance in predictions for task-relevant molecules, and (ii) systematic bias where dark points deviate from the diagonal, indicating the model has not adequately explored these regions. These deficiencies become more severe at $\lambda = 75$, where the target distribution concentrates on a narrower slice of high-scoring compounds. The pairwise plots thus provide a complementary diagnostic to the R^2-on-top-$k\%$ metric, confirming that `AB-SID-iVAR`'s advantage stems from concentrating sampling effort on the Boltzmann-weighted region of interest.

E Limitations and Future Work

Our theoretical bounds exhibit exponential dependence on $|\lambda|$ through the density-ratio argument (Lemmas 3.2–3.3). While this establishes consistency for any fixed λ, the global nature of this approach, requiring uniform control of the partition function ratio across the entire domain, introduces an exponential prefactor $e^{\mathcal{O}(|\lambda|\sqrt{\beta_T})}$ that does not appear in analogous Bayesian optimization results for the $\lambda \to \infty$ limit.

Computationally, our acquisition function relies on MCMC sampling to approximate the constraint set $\mathcal{X}_t$ and the integral in Eq. 3 over continuous domains. While Figure 9 shows that substantially smaller SMC sample sizes than those used in our experiments suffice without degrading performance, thereby reducing acquisition optimization cost, this still introduces overhead beyond methods that require only point evaluations of the surrogate. Since our algorithm is agnostic to the choice of MCMC sampler, we believe this cost can be further reduced via amortized Boltzmann distribution samplers (e.g., [OuYang et al., 2024, Akhound-Sadegh et al., 2025]).

Several directions are promising for future work. Empirically, our evaluation is limited to controlled scientific benchmarks, with continuous-domain experiments covering low to moderate input dimensionalities. Validation in more complex, end-to-end scientific applications remains an important future work. On the theoretical side, developing input-dependent analyses that directly track coverage of high-probability regions, bridging our analysis with Bayesian optimization regret bounds in the $\lambda \to \infty$ limit, and extending guarantees to non-Boltzmann SID forms are natural next steps. Additionally, our preliminary analysis of `GHAL` (Proposition 4.7) assumes exhaustive exploration of the input space by the Markov chain; relaxing this condition to cover more general settings remains a natural direction of future study. Finally, our acquisition function couples the AB/TS-SID surrogate with the iVAR criterion as an integrated design, motivated by and analyzed through the variance-based upper bounds in Lemmas 3.3-3.4. Studying alternative integrated designs under the SIDAL framework, along with their corresponding theoretical guarantees (which would require analysis distinct from Theorem 4.1), is a promising direction for future work.

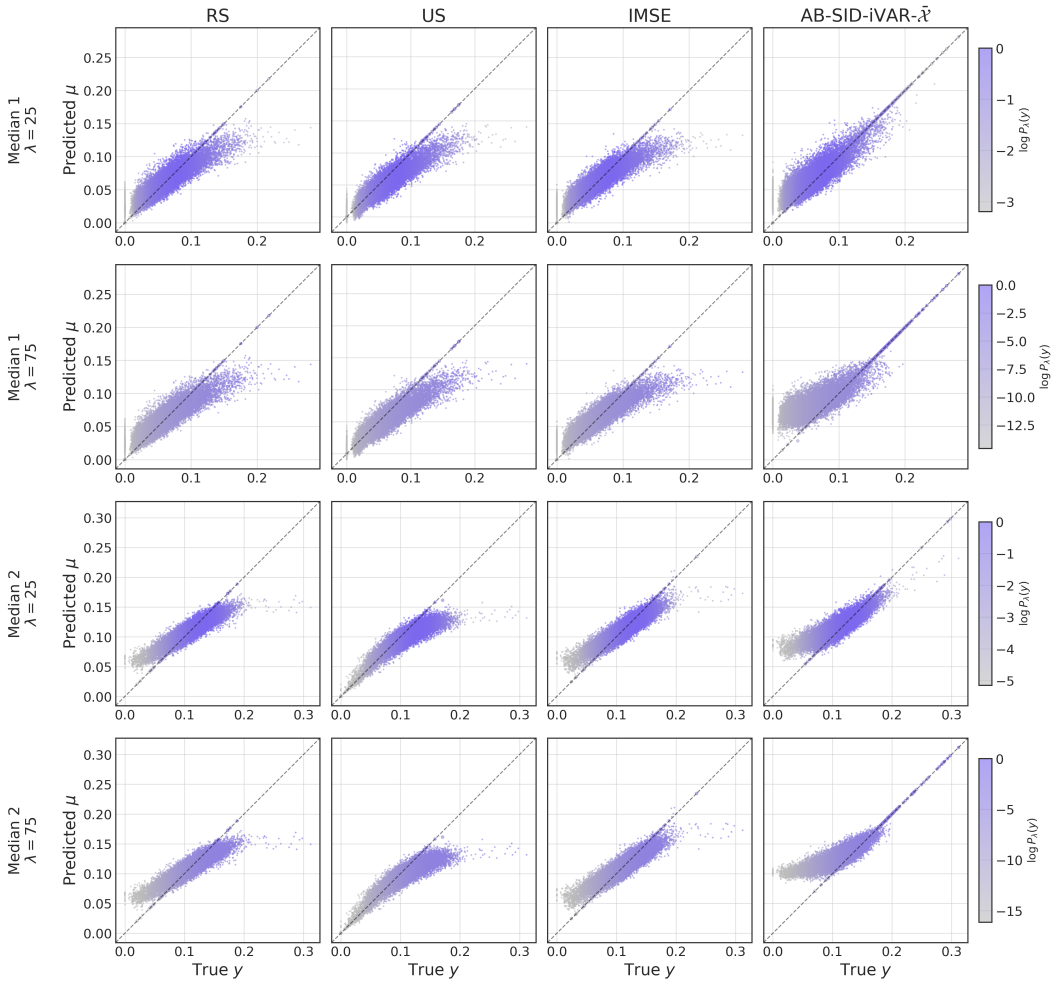


Figure 11: Pairwise comparison of predicted GP posterior mean μ versus true function value y across all ≈20k GuacaMol molecules. Rows: scoring functions (Median 1, Median 2) and Boltzmann temperatures ($\lambda = 25, 75$). Columns: acquisition strategies. Point color encodes log Boltzmann weight $\log P_\lambda(y) \propto \lambda y$ (darker = higher weight, i.e., more task-relevant molecules), we omit reporting the result of EPIG here for the sake of space, and it behaves almost identical to IMSE. Ideal performance corresponds to dark points clustering tightly along the $y = \mu$ diagonal. `AB-SID-iVAR` (rightmost column) achieves this across all settings, whereas baselines show systematic vertical spread or bias for high-weight molecules—particularly pronounced at $\lambda = 75$ where the target distribution is sharply concentrated.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes],

   Justification: The abstract and introduction faithfully reflect the paper's contribution and scope.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The limitations are discussed in Appendix.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: The assumptions are provided and the complete proofs are provided.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Experimental details are described in the paper and the code is provided for reproducible result.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The code is provided.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Experimental details are provided in the paper.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer:[Yes]

   Justification: All experiments are reported with sufficient repetitions.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The computer resources is reported in runtime report section.

9. **Code of ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

Answer: [Yes]

Justification: We have reviewed the NeurIPS Code of Ethics and confirm that this research conforms to it in all respects.

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A].

    Justification: This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: This work does not release any data or pre-trained models that pose risks of misuse.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All third-party assets used in this work are open-source and properly cited and used in accordance with their respective licenses.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: We release the source code implementing our proposed methods (AB-SID-iVAR and TS-SID-iVAR) along with scripts to reproduce all experiments. The code is provided at `https://github.com/TsingQAQ/SIDAL`

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This work does not involve crowdsourcing or research with human subjects.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: This work does not involve human subjects research.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs are not used as a component of the core methodology.