NON-REPLACEMENT FUNCTION SPACE SAMPLING FOR BAYESIAN OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Bayesian optimization (BO) is a probabilistic framework for optimizing expensive black-box functions through sequential decision making guided by an acquisition function. Existing acquisition strategies explicitly balance exploration and exploitation, often requiring carefully designed heuristics or hyper-parameter tuning. We introduce Non-Replacement Function Space Sampling (NRFS), achieving desired exploration-exploitation balance by prioritizing the functions that have not contributed to previous acquisition decisions, rather than relying on explicit tuning. By establishing a correspondence between each candidate and the set of functions that consider it as the corresponding optimizer, NRFS induces a principled and efficient search strategy in the design space. Empirical evaluations on a diverse set of benchmark problems demonstrate that NRFS outperforms state-of-the-art BO methods, particularly in challenging regimes that demand both global exploration and local refinement. Moreover, we provide theoretical guarantees that NRFS converges for any Lipschitz-continuous objective function, regardless of the complexity of its landscape.

1 INTRODUCTION

Bayesian Optimization (BO) is a sample-efficient framework for optimizing expensive black-box functions, with the goal of rapidly identifying the optimizer, i.e., the optimal input location in the design space. BO strives for the balance between exploration and exploitation to efficiently identify optimal solutions (Jalali et al., 2012; Candelieri, 2023), where exploration encourages querying in regions of high uncertainty to improve understanding of function responses globally, while exploitation favors regions with promising predicted values to quickly find optima (De Ath et al., 2021). Classical acquisition functions in BO, such as Expected Improvement (EI; Jones et al., 1998) and Probability of Improvement (PI; Kushner, 1964; Snoek et al., 2012), primarily focus on quantifying potential performance gains. Entropy-based approaches, including Predictive Entropy Search (PES; Hernández-Lobato et al., 2014), Max-value Entropy Search (MES; Wang & Jegelka, 2017) and Joint Entropy Search (JES; Hvarfner et al., 2022), focus on reducing uncertainty about the location of the global optimum. Other uncertainty reduction approaches, such as Expected Information Gain (EIG; Tsilifis et al., 2017) and step-wise uncertainty reduction (Villemonteix et al., 2009), aim to enhance surrogate models. Methods that emphasize performance gains may suffer from oversampling when the surrogate is biased or mis-specified (Wang & de Freitas, 2014), whereas uncertainty reduction approaches do not directly target identifying precise optima, leading to potential inefficiency. To balance these two objectives, hybrid methods such as Upper Confidence Bound (UCB; Auer et al., 2002) and Variational Entropy Search (VES; Cheng et al., 2025) combine performance gain and uncertainty reduction into a single reward to guide optimization. Other strategies adaptively switch between exploration and exploitation based on current observations (Bian et al., 2021).

These existing BO acquisition strategies share two common features: 1) they are typically based on human-defined subjective metrics that may deviate from finding true optima: No matter whether they are framed in terms of the surrogate model or optimizer location uncertainty, iterative improvement with respect to specified criteria, or a combination of both, the corresponding acquisitions can be biased instead of fully aligned with the ground-truth optimizer. 2) The estimation of the acquisition functions can also introduce biases based on the chosen surrogate model space. Researchers often interpret the surrogate as a collection of candidate functions, any of which could represent the true objective function. However, this assumption may not lead to fast identification of optima in

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

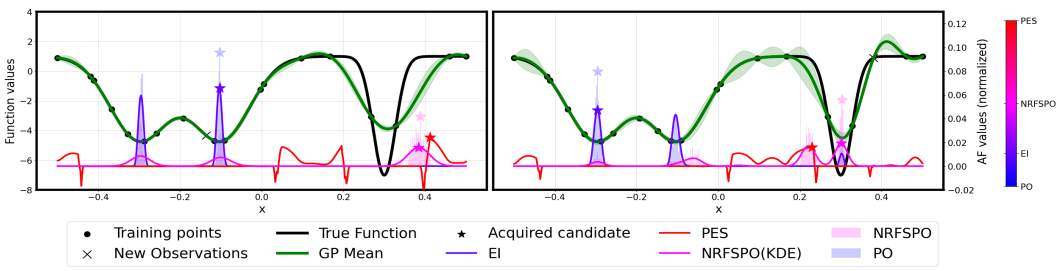

Figure 1: Comparison for different acquisition strategies over two iterations: the left plot shows the initial state, and the right plot shows the state after adding the NRFS acquisition candidate to the training samples. The colormap highlights the exploration/exploitation tendencies of different acquisition strategies, where red indicates exploration and blue indicates exploitation.

practice. For instance, the goal of BO is to identify the true optimum beyond the current best, while during BO iterations, many sampled functions could not reach the true optimum, particularly those that already assigned the current best observation as their optimizer. This can lead to redundancy and bias towards inefficient acquisition function estimation without focusing quickly on the underlying true objective function and its optimal solution(s).

To develop a more efficient BO acquisition strategy, we focus on two key aspects: (1) what is the most suitable objective target to guide the acquisition, and (2) what is the most effective method to estimate this objective target. We develop a probabilistic reasoning framework relying on "Probability of Optimality (PO)", which directly characterizes the likelihood of a candidate being the optimizer without depending on human-defined heuristics. Furthermore, we employ a non-replacement sampling strategy to estimate PO precisely.

Within this framework, we propose a new acquisition strategy, Non-Replacement Function Space Sampling (NRFS), which adopts a unified strategy that directly connects optimizer probability to function space coverage, consistently ensuring convergence to the global optimum. In our formulation, the surrogate model is treated as a pool of functions that contain the true objective. We iteratively identify the optimizers of functions sampled from this pool without replacement until the pool is fully depleted. The key advantage of this approach is that, even if the fitted surrogate is over represented in certain regions, as long as the surrogate contains the true objective, non-replacement sampling guarantees convergence once the pool is exhausted.

Assume that we have collected observations $\mathcal{D}$, to which we fit a GP and denote the corresponding function space as $\mathcal{F}^{\mathcal{D}}_{\mathcal{GP}}$. We then take each design point $\boldsymbol{x}$ to define a function *bucket*:

$$\mathcal{F}^{\mathcal{D}}_{\boldsymbol{x}} = \left\{ f \in \mathcal{F}^{\mathcal{D}}_{\mathcal{GP}} \;\middle|\; \boldsymbol{x} = \arg\min_{\boldsymbol{x}' \in \mathcal{X}} f(\boldsymbol{x}') \right\}, \tag{1}$$

which contains all functions with $\boldsymbol{x}$ being their optimizer. This definition guarantees that each function, including the objective function, belongs to at least one bucket. For functions with multiple global optima, we assign them randomly to one of their valid buckets to have one-to-one mappings between $\boldsymbol{x}$ and $\mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}}$. If we could identify the bucket containing the true objective function, the BO task would be solved. Our approach selects the bucket containing the largest number of functions, thereby maximizing the probability that the true objective resides within the bucket. Once a bucket is selected, we avoid all functions assigned to it from future consideration, ensuring that no function is selected more than once and thereby improving search efficiency.

Figure 1 illustrates an example demonstrating the effectiveness of our NRFS strategy directly focusing on the global optimizer without tuning exploration-exploitation trade-off explicitly as in existing methods. The objective is to locate the global optimum in the rightmost valley. In the initial state, EI and PO remain focused on local optima (left two valleys), where the surrogate already exhibits high confidence. This observation suggests that relying solely on PO, without incorporating NRFS, still suffers from the oversampling issue observed in EI. PES and NRFS demonstrate stronger performance on avoiding oversampling; however, after new acquisition, PES continues to target high-uncertainty areas. In contrast, NRFS shifts attention towards the promising region as shown by the pink curve in the right plot of Figure 1, enabling successful identification of the global optimum.

As we have briefly introduced the limitations of existing acquisition strategies in BO and how our proposed NRFS addresses them, we will present more details of our contributions based on the following organization: Section 2 provides a more detailed overview of commonly used acquisition functions in BO; Section 3 introduces our NRFS-based BO formulation; Section 4 reports empirical results, and Section 5 concludes the paper.

2 BACKGROUND

Let $\boldsymbol{x}$ be a d-dimensional vector of decision variables in the feasible design space $\mathcal{X} \subset \mathbb{R}^d$; and $f^*(\cdot) : \mathcal{X} \to \mathbb{R}$ a continuous black-box objective function, with any evaluation $f^*(\boldsymbol{x})$ being an expensive process with respect to time and/or cost. We aim to approach the global optimizer $\boldsymbol{x}^*_{f^*}$ within a limited number of evaluations by a given function evaluation budget:

$$\boldsymbol{x}^*_{f^*} = \arg\min_{\boldsymbol{x}\in\mathcal{X}} f^*(\boldsymbol{x}). \tag{2}$$

In BO, the unknown black-box objective function is often modeled by a GP, $p(f)$, characterized by its mean function $\mu(\cdot) : \mathcal{X} \to \mathbb{R}$ and covariance kernel function $k(\cdot,\cdot) : \mathcal{X}^2 \to \mathbb{R}$. BO sequentially selects a sequence of observed samples for their corresponding evaluations (Frazier, 2018). Given the observation data set until the n-th iteration $\mathcal{D}_n = \{\boldsymbol{X}_n, Y_n\}$, the GP posterior of f is updated (Rasmussen, 2003). We denote the posterior belief of f as $p(y \mid \boldsymbol{x}, \mathcal{D}_n) = P(y = f^*(\boldsymbol{x}) \mid \mathcal{D}_n)$. In each BO iteration, the next query point is chosen by optimizing an acquisition function:

$$\boldsymbol{x} = \arg\max_{\boldsymbol{x'}\in\mathcal{X}} u_n(\boldsymbol{x'}), \tag{3}$$

where $u_n(\boldsymbol{x})$ is the expected utility of evaluating $\boldsymbol{x}$ based on the updated GP posterior. The acquisition function should avoid oversampling and unnecessary exploration, which means that the resulting sequential queries should favor both the points with potential good values with respect to the objective and the informative points from the unexplored regions for learning better optimizer distribution.

Related work. A commonly used strategy in BO is to estimate the potential improvement obtained by evaluating a candidate point $\boldsymbol{x}$. EI is a widely used acquisition function that accounts for potential improvements in objective value. A well-known limitation of EI and its variants are their tendency to oversample near local optima, particularly when the probabilistic model's prior is the initialization is biased or mis-specified (Wang & de Freitas, 2014) . PI, which targets the likelihood of improvement, exhibits similar behavior (Kushner, 1964; Snoek et al., 2012). Oversampling usually happens when the acquisition is close to the evaluated training points, and the updated surrogate model will show minimal differences compared to the previous iteration. Consequently, the acquisition is repeatedly computed from a function realization set resembling the previous one, which results in nearly identical acquisition decisions iteratively.

Beyond acquisition functions that target improvement-based rewards, some strategies emphasize uncertainty reduction. These methods can be grouped into two categories. The first category aims to reduce uncertainty about the optimizer's location (Hernández-Lobato et al., 2014; Wang & Jegelka, 2017; Hvarfner et al., 2022). For example, PES (Hernández-Lobato et al., 2014) explicitly models the posterior distribution over the unknown minimizer, denoted by:

$$P(\boldsymbol{x} = \boldsymbol{x}^*_{f^*} \mid \mathcal{D}_n) \approx \mathbb{E}_{f\sim\mathcal{GP}}[P(\boldsymbol{x} = \arg\min_{\boldsymbol{x'}} f(\boldsymbol{x'}) \mid \mathcal{D}_n)]. \tag{4}$$

Since PES aims to reduce the uncertainty in the location of the true optimizer $\boldsymbol{x}^*_{f^*}$ by reducing the entropy of the posterior distribution defined in equation 4, the selected observation sequence is not necessarily close to $\boldsymbol{x}^*_{f^*}$, which may not provide promising candidates with optimal objective value under limited evaluation budget. In the extreme scenario where the optimizer's location is already known, PES still cannot determine the next acquisition point because the entropy contributions for optimal impossible candidates and the definitive optimizer are both 0. The second category aims at reducing the surrogate model uncertainty (Villemonteix et al., 2009; Tsilifis et al., 2017), which shrinks the set of possible functions over the entire search space. However, this strategy can also be inefficient as it needs extra effort distinguishing between functions with the same optimizer.

To mitigate oversampling in a single region and suboptimal suggestions from uncertainty reduction, researchers have explored hybrid approaches, such as UCB (Auer et al., 2002), Moment Generating

Function (MGF; Wang et al., 2017), Truncated Variance Reduction (TVR; Bogunovic et al., 2016), and VES (Cheng et al., 2025), balancing exploration and exploitation with hyperparameters that guide acquisitions toward regions of high uncertainty when the process gets trapped in local optima. These methods typically favor either exploitation-driven or uncertainty-reduction-driven decisions in specific iterations. Alternative hybrid methods, such as ϵ-EI (De Ath et al., 2021), enforce exploration in random iterations with a probability ϵ. Researchers also design schedules that decrease exploration as iterations progress. However, the optimal balance between exploration and exploitation is problem-dependent, making it challenging to predefine a universally effective schedule for diverse unknown objectives. To address this challenge, online tuning strategies are employed to dynamically adjust hyperparameters or schedules. Unfortunately, these strategies usually require at least 50 to 100 iterations to converge to optimal settings and can be even more computationally expensive for problems of high complexity.

3 Non-Replacement Function Space Sampling

3.1 Probability of Optimality (PO)

To avoid redundant and biased BO acquisitions due to subjective acquisition functions, our new acquisition strategy, NRFS, adopts PO as our utility function, which aims at maximizing the probability that the next acquisition corresponds to the true optimizer:

$$\boldsymbol{x} = \arg\max_{\boldsymbol{x'} \in \mathcal{X}} P(\boldsymbol{x'} = \boldsymbol{x}^*_{f^*}). \tag{5}$$

This formulation (5) expresses the objective of BO more directly and accurately, as it is derived from the ultimate goal in (2) and does not depend on any additional, subjectively defined rewards compared to the methods adopting specific evaluation criteria, such as EI or PES. However, as we do not have the underlying objective function f^*, the key challenge here becomes reliably estimating this probability by replacing f^* in (2) with f drawn from the surrogate function distribution, meanwhile the fixed optimizer location on the left-hand side becomes a random variable over optimizer locations. The objective is then to identify the most probable optimizer location with respect to the iteratively updated surrogate model space considering all inherent uncertainty.

Direct estimation of the probability $P(\boldsymbol{x} = \boldsymbol{x}^*_{f^*})$ is nontrivial when each $\boldsymbol{x} \in \mathcal{X}$ is treated as an observed input location. Following the strategy in Hennig & Schuler (2012), we reformulate PO as:

$$P(\boldsymbol{x} = \boldsymbol{x}^*_{f^*}) = \int_{f:\mathcal{X}\to\mathcal{Y}} p(f) \prod_{\substack{\tilde{\boldsymbol{x}}\in\mathcal{X} \\ \tilde{\boldsymbol{x}}\neq\boldsymbol{x}}} H[f(\tilde{\boldsymbol{x}}) - f(\boldsymbol{x})]\, df, \tag{6}$$

where $\mathcal{Y}$ is the objective space and $H[\cdot]$ is the Heaviside step function, and $\prod_{\substack{\tilde{\boldsymbol{x}}\in\mathcal{X} \\ \tilde{\boldsymbol{x}}\neq\boldsymbol{x}}} H[f(\tilde{\boldsymbol{x}}) - f(\boldsymbol{x})]$ acts as an indicator whether $\boldsymbol{x}$ is the global optimizer of the function f. The above equation (6) denotes the expected probability that $\boldsymbol{x}$ is the optimizer over all sampled functions within the surrogate model space. With the definition of the function bucket in (1), we can estimate the size of $\mathcal{F}^{\mathcal{D}}_{\boldsymbol{x}}$ as follows:

$$\mid \mathcal{F}^{\mathcal{D}}_{\boldsymbol{x}} \mid = \mid \mathcal{F}^{\mathcal{D}}_{\mathcal{GP}} \mid \int_{f:\mathcal{X}\to\mathcal{Y}} p(f) \prod_{\substack{\tilde{\boldsymbol{x}}\in\mathcal{X} \\ \tilde{\boldsymbol{x}}\neq\boldsymbol{x}}} H[f(\tilde{\boldsymbol{x}}) - f(\boldsymbol{x})]\, df, \tag{7}$$

where $\mid \mathcal{F}^{\mathcal{D}}_{\mathcal{GP}} \mid$ represents the size of the surrogate model space. Based on (6) and (7), we establish the connection between PO and the function space coverage ratio (6):

$$P(\boldsymbol{x} = \boldsymbol{x}^*_{f^*}) = \mid \mathcal{F}^{\mathcal{D}}_{\boldsymbol{x}} \mid \,/\, \mid \mathcal{F}^{\mathcal{D}}_{\mathcal{GP}} \mid, \tag{8}$$

which transforms a nontrivial probability into a quantity that can be estimated by sampling functions from the surrogate GP. Note that $\mid \mathcal{F}^{\mathcal{D}}_{\mathcal{GP}} \mid$ is typically estimated by how many functions are sampled from the surrogate model space, which is usually fixed. Consequently, optimizing PO is equivalent to identifying the maximizer of the numerator $\mid \mathcal{F}^{\mathcal{D}}_{\boldsymbol{x}} \mid$.

3.2 Coverage contribution estimation

Equation (8) suggests that increasing the probability of selecting the true optimizer can be achieved by enlarging the set of functions included in the future function bucket. However, (8) only describes

a one-step strategy conditioned on some observation data $\mathcal{D}$. Since BO is inherently sequential, incorporating the acquisitions $\mathcal{D}_n$ from previous iterations can further help mitigate its potential myopic behavior. Thus, we extend (8) into a cumulative form, defined as:

$$P(\boldsymbol{x}^*_{f^*} \in \boldsymbol{X}_{n+1}) = P(\boldsymbol{x}^*_{f^*} \in \boldsymbol{X}_n) + P(\boldsymbol{x}_{n+1} = \boldsymbol{x}^*_{f^*} \mid \boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n)P(\boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n). \tag{9}$$

Note that $P(\boldsymbol{x}^*_{f^*} \in \boldsymbol{X}_n)$ and $P(\boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n)$ are fixed after n iterations, optimization of (9) can be achieved by maximization on the term $P(\boldsymbol{x}_{n+1} = \boldsymbol{x}^*_{f^*} \mid \boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n)$. Unlike the one-step maximization target defined in (8), $P(\boldsymbol{x}_{n+1} = \boldsymbol{x}^*_{f^*} \mid \boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n)$ is conditioned on the event $\boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n$. Intuitively, this is reasonable: if $\boldsymbol{x}^*_{f^*} \in \boldsymbol{X}_n$, the BO process would already be completed. Thus, (9) formalizes that continuing BO requires conditioning on $\boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n$.

To the best of our knowledge, no prior work has used $\boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n$ as a condition to guide BO, since this condition appears to merely shrink the design space. However, we observe that $\boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n$ also influences the distribution of the optimum value in the objective space, thereby exerting a broader impact on the function space. From $\boldsymbol{x}^*_{f^*} \notin \boldsymbol{X}_n$, we obtain a corresponding condition for the optimum distribution: $f^*(\boldsymbol{x}^*_{f^*}) = Y^* \notin \boldsymbol{Y}_n$. For a minimization problem, we have $Y^* < Y^*_n = \min\{\boldsymbol{Y}_n\}$. This inequality condition is crucial, as it indicates that the true optimum must improve upon the current best. Consequently, any sampled function that cannot achieve a value better than the current best cannot be the true objective function. When sampling the function space to estimate (9), such functions should be excluded, as their probability $p(f)$ is zero. Thus to maximize (9), we must consider how many functions from $\mathcal{F}^{\mathcal{D}}_{\boldsymbol{x}}$ are from impossible regions and avoid including them from future consideration in the objective function space. We partition the candidate's function bucket into two subsets:

$$\mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}} = \mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}, f(\boldsymbol{x}) \geq Y^*_n} \cup \mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}, f(\boldsymbol{x}) < Y^*_n} \tag{10}$$

For all functions in $\mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}, f(\boldsymbol{x}) \geq Y^*_n}$, their optima do not surpass the current best observation Y^*_n. As discussed earlier, the true objective function cannot belong to this set. Hence, functions from set $\mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}, f(\boldsymbol{x}) \geq Y^*_n}$ contribute only ineffective coverage and the effective coverage can be computed as:

$$|\mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}}|_{eff} = |\mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}, f(\boldsymbol{x}) < Y^*_n}| + |\mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}, f(\boldsymbol{x}) \geq Y^*_n}| = |\mathcal{F}^{\mathcal{D}_n}_{\boldsymbol{x}, f(\boldsymbol{x}) < Y^*_n}|. \tag{11}$$

To accurately estimate the coverage contribution of each candidate point $\boldsymbol{x}$, we restrict sampling optimizers that overwhelm the current best for all $\boldsymbol{x} \in \mathcal{X}$. This constraint ensures unbiased estimation of coverage improvement across all possible candidates. Under this setup we sample from a Truncated Gaussian Process (TGP), defined as $\mathcal{TGP}(\mu, k, t)$, where t is the threshold. For minimization problems, we sample from $\mathcal{TGP}^-(\mu, k, Y^*_n)$, indicating that the sampling is restricted to the function space lying below the threshold Y^*_n. The final utility function is estimated by a variant of (6) with the samples from $\mathcal{TGP}^-(\mu, k, Y^*_n)$, shown as:

$$\mathbb{E}_{f \sim \mathcal{TGP}^-(\mu, k, Y^*_n)} \left[P\left(\boldsymbol{x} = \boldsymbol{x}^*_f \mid \mathcal{D}_n\right) \right] = \int_{\mathcal{TGP}^-} \prod_{\substack{\tilde{\boldsymbol{x}} \in \mathcal{X} \\ \tilde{\boldsymbol{x}} \notin \{\boldsymbol{X}_n \cup \boldsymbol{x}\}}} H[f(\tilde{\boldsymbol{x}}) - f(\boldsymbol{x})] p(f) df. \tag{12}$$

In summary, we maximize (9) by identifying the maximizer of (12), which illustrates how NRFS avoids the exploration-exploitation dilemma based on joint distribution of probability of optimality and function distribution. Compared with the original probability of optimality without any function distribution conditioning, our method by NRFS mitigates oversampling by filtering out all functions that have been identified with an optimizer whose $p(f) = 0$ in future steps. Meanwhile, the threshold condition Y^*_n also prevents redundancy by avoiding acquisitions in regions where the optimizer cannot exist ($\geq Y^*_n$), thereby enhancing exploration efficiency.

3.3 One-Step-Look-Ahead Variant Development

To further increase sample efficiency in BO, here we develop a one-step-look-ahead (OSLA) variant of (12). After T evaluations, the probability that the optimizer has been identified is

$$R_T = P(\boldsymbol{x}^*_{f^*} \in \boldsymbol{X}_T) = 1 - \prod_{t=1}^{T} \big(1 - P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_t)\big). \tag{13}$$

We refer to R_T as *the cumulative PO*. The corresponding incremental contribution of each iteration t is

$$r_t = P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_t) \prod_{s=1}^{t-1} \big(1 - P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_s)\big), \tag{14}$$

i.e., the probability of discovering the optimizer at iteration t given it has not been selected before. At iteration n, maximizing the defined cumulative PO can be estimated by maximizing the following one-step-look-ahead value function:

$$V_n \approx r_n + r_{n+1} = \prod_{s=1}^{n-1} \Big(1 - P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_s)\Big) \times \Big[P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_n) + \big(1 - P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_n)\big) P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_{n+1})\Big]. \tag{15}$$

Since $\prod_{s=1}^{n-1} \Big(1 - P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_s)\Big)$ is fixed after iteration $n-1$, optimization of the value function reduces to maximizing $P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_n) + \Big(1 - P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_n)\Big) P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}_{n+1})$. We therefore define the corresponding one-step-look-ahead utility function:

$$u_n(\boldsymbol{x}) = P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x} \mid \mathcal{D}_n) + \Big(1 - P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x} \mid \mathcal{D}_n)\Big) \times \mathbb{E}_{y|\mathcal{D}_n,\boldsymbol{x}} \Big[\max_{\boldsymbol{x}' \in \mathcal{X}} P\big(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}' \mid \mathcal{D}_n \cup (\boldsymbol{x}, y)\big)\Big], \tag{16}$$

where $P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x} \mid \mathcal{D}_n)$ and $\max_{\boldsymbol{x}' \in \mathcal{X}} P\big(\boldsymbol{x}^*_{f^*} = \boldsymbol{x}' \mid \mathcal{D}_n \cup (\boldsymbol{x}, y)\big)$ are estimated by (12).

In contrast to one-step-look-ahead variants of other acquisition functions, which typically require tuning a hyperparameter to balance immediate and future rewards, our one-step-look-ahead NRFS directly sets this hyperparameter to $1 - P(\boldsymbol{x}^*_{f^*} = \boldsymbol{x} \mid \mathcal{D}_n)$. This choice avoids potential probability gaps across different problem settings, eliminating the need to search for task-dependent hyperparameters.

Another advantage is that the one-step-look-ahead NRFS variant has the potential to achieve the maximum convergence rate. Traditionally, the convergence rate is defined by how quickly an error term, or a probability gap, decays to zero (Ryzhov, 2016; Bull, 2011). In our setting, this probability gap corresponds to the cumulative regret $R_T^{\text{reg}} = 1 - R_T$. Note that maximizing the value function V_n at each step is equivalent to minimizing cumulative regret, so each acquisition step can be interpreted as directly advancing the convergence rate.

3.4 Theoretical analysis and convergence rates

Estimating the probability of optimality defined in (12) and (16) requires comparing function values across the entire design space, resulting in a computational complexity that scales exponentially with respect to design space dimension d, which can become computationally expensive in high dimensional settings, particularly compared with methods that require only the design space input, like EI. To address this, we derive a simplified formulation by considering single sampled optimizer location (i.e., $M = 1$). Using an estimation strategy analogous to that employed in EI or PI, we approximate the expected acquisition values under the GP posterior and select query points based on these expectations. In analytical NRFS (NRFSA), the expected acquisition value is defined as the expected optimum value, which is shown as:

$$\alpha_n^{\text{NRFS}}(\boldsymbol{x}) := \mathbb{E}[y^* | \boldsymbol{x}, \mathcal{D}_n] = \int_{-\infty}^{Y_n^*} y^* p_{\text{trunc}}(y^*) dy^* = \mu_n(\boldsymbol{x}) - \sigma_n(\boldsymbol{x}) \frac{\varphi\big(z_n^*(\boldsymbol{x})\big)}{\Phi\big(z_n^*(\boldsymbol{x})\big)}, \tag{17}$$

where $p_{\text{trunc}}(y^*)$ denotes the normalized truncated distribution ensuring that the optimum value lies in the region beyond the current best observation. Here we use $\alpha(\boldsymbol{x})$ to denote $\alpha_n^{\text{NRFS}}(\boldsymbol{x})$ for any n. More details of the derivation of (18) within a GP surrogate can be found in Appendix A.1. Remarkably, the analytical form of NRFS can be expressed in the classical UCB form, with the inverse Mills ratio serving as the corresponding β coefficient. In contrast, existing UCB methods typically define β_t as a function of the acquisition iteration t, with β_t increasing over iterations to ensure convergence. Such iteration dependent schedules are inherently heuristic, as the appropriate schedule generally depends on the unknown characteristics of the underlying black-box function and therefore cannot be determined a priori. In contrast, analytical NRFS yields a purely data-driven β coefficient that is independent of the acquisition iteration. Moreover, this coefficient provides a deterministic convergence guarantee, as established in the following analysis.

Theorem 1. *Assume that the design space $\mathcal{X} \subset \mathbb{R}^d$ is compact and bounded for some $d \geq 1$, the objective function f^* is an element of a Reproducing Kernel Hilbert Space (RKHS) with kernel $k(\cdot,\cdot)$, which is a continuous and positive definite kernel function and has the Lipschitz continuity property. The observation dataset $\mathcal{D}_T$ generated by $\alpha_T^{NRFS}(\boldsymbol{x})$ has the following property: as $T \to \infty$, the simple regret $r_T^{\text{reg}} \to 0$.*

To prove **Theorem** 1, we first establish supporting lemmas.

Lemma 1. *Let $\mathcal{X}$ be a compact design space and consider a noise-free GP with a continuous kernel. For any fixed positive constant $\varepsilon > 0$, suppose acquisitions are selected only from locations $\boldsymbol{x} \in \mathcal{X}$ satisfying $\sigma_T(\boldsymbol{x}) > \varepsilon$. Then, the total number of such acquisitions is finite.*

Lemma 2. *Let $\mathcal{X}$ be a compact design space and consider a noise-free GP with a continuous kernel. For any fixed positive constant $\varepsilon > 0$, suppose acquisitions are selected only from locations $\boldsymbol{x} \in \mathcal{X}$ satisfying $\mu_T(\boldsymbol{x}) < y_{T-1}^* - \varepsilon$. Then, the total number of such acquisitions is finite.*

Both **Lemma** 1 and **Lemma** 2 can be proved by contradiction, since violating either lemma would require an unbounded design space. Detailed proofs are provided in Appendix A.2.1 and Appendix A.2.2. These two lemmas characterize two fundamental properties of the BO process. **Lemma** 1 states that it is impossible to perform infinitely many exploitation steps in regions where the predicted mean remains better than the current best value. **Lemma** 2 states that infinitely many exploration steps cannot be performed in regions where the posterior variance remains above a constant threshold.

Lemma 3. *Let the query points $\{\boldsymbol{x}_t\}_{t=1}^T$ be sequentially selected to minimize $\alpha_T^{NRFS}(\boldsymbol{x})$. As $T \to \infty$, the posterior variance of the GP over the entire design space $\mathcal{X}$ satisfies*

$$\sup_{\boldsymbol{x} \in \mathcal{X}} \sigma_T(\boldsymbol{x}) \longrightarrow 0.$$

Lemma 3 shows that NRFS has sufficient exploration power to ensure that the posterior uncertainty at the true optimizer $\boldsymbol{x}^*$ vanishes as $T \to \infty$. The proof of **Lemma** 3 relies on **Lemmas** 1 and 2, as shown in Appendix A.2.3.

Lemma 4. *Under the NRFS acquisition $\alpha(\boldsymbol{x})$, the gap between the current best and the global optimum satisfies*

$$r_T^{\text{reg}} = \min_{t=1,\ldots,T} \{y_t - Y^*\} = O\big(\sigma_T^2(\boldsymbol{x}^*)\big).$$

Lemma 4 shows that the convergence rate of NRFS is of the same order as the rate at which the posterior uncertainty at the global optimizer $\boldsymbol{x}^*$ decreases. The detailed proof can be found in Appendix A.2.4. Combining **Lemma** 3 and **Lemma** 4, we first establish that the posterior uncertainty converges to zero at every point in the design space, including the true optimizer $\boldsymbol{x}^*$. We then show that the simple regret decreases at the same asymptotic rate as the posterior variance, as detailed in Appendix A.2.4.

Compared to previous convergence analyses typically relying on strong assumptions, such as bounded maximum information gain (MIG) for GP-UCB (Srinivas et al., 2012), sufficient exploration conditions (Vazquez & Bect, 2010), and appropriate specification of the GP prior for EI (Bull, 2011), our analysis requires only minimal conditions: Lipschitz continuity of the objective, a compact design space and bounded surrogate hyperparameters to establish convergence. Under these basic assumptions, we show that NRFS converges surely for objective functions with arbitrary landscapes. Moreover, the algebraic proof does not depend on any prior theoretical results, highlighting the conceptual novelty and generality of our analysis.

4 Empirical results

4.1 Experimental setup

We follow the practice in posterior sampling from BOTorch (Wilson et al., 2020; Balandat et al., 2020), to estimate (12) and (16). Specifically, we first sample M (100 for synthetic problems and 1000 for real application) functions from $\mathcal{TGP}^-(\mu, k, Y_n^*)$ to obtain a function set $\{f_j(\cdot)\}_{j=1}^M$. For each sampled function $f_j(\cdot)$, we identify its optimizer location and then aggregate these optimizers

to estimate the distribution over the design space. In continuous domains, NRFS can be implemented by applying any density estimation techniques. In our implementation, we apply Parzen estimator (Silverman, 2018), resulting in a continuous density function. The acquisition is selected as the most probable optimizer location under this estimated distribution. More acquisition computation and optimization details could be found in Appendix A.3.

4.2 Search behavior

To understand the search behavior of NRFS intuitively, we test an illustrative example with a Gaussian mixture (GM) objective function to minimize: $-\mathcal{N}(-\mu_1, \sigma_1^2) - \mathcal{N}(\mu_2, \sigma_1^2) - 0.55\mathcal{N}(\mu_1, \sigma_2^2) + 1$. We set $\mu_1 = 0.3$, $\mu_2 = -0.1$, $\sigma_1 = 0.05$ and $\sigma_2 = \frac{0.3}{\sqrt{2}}$ (Figure 1). The objective function has two local minima and one global minimum far from local minima. The search behavior comparison has been discussed in Section 1 and more empirical performance comparison can be found in Section 4.3.

Besides the search behavior analysis on 1D GM example, we apply NRFS on other benchmark objective functions which are commonly used in evaluating BO methods, including 1) BraninRcos2 (Al-Roomi, 2015); 2) Himmelblau's (Himmelblau et al., 2018) and 3) Forrester (Forrester et al., 2008) to check the convergence. Figure 2 shows that, with sufficient iterations, NRFS converges to all global optima, regardless of whether the optimum is unique, multiple, or infinite in number. The distribution of training samples shows high density near the true optima and lower density elsewhere, indicating that acquisition is guided by the potential to identify global optima in given regions. Oversampling occurs rarely and primarily when the search is already close to an optimum. At the same time, later iterations still allocate some acquisitions in high uncertainty regions, which avoids missing possible global optimum which has extreme objective value. Additional behavioral analysis on other objective functions is provided in Appendix A.4.

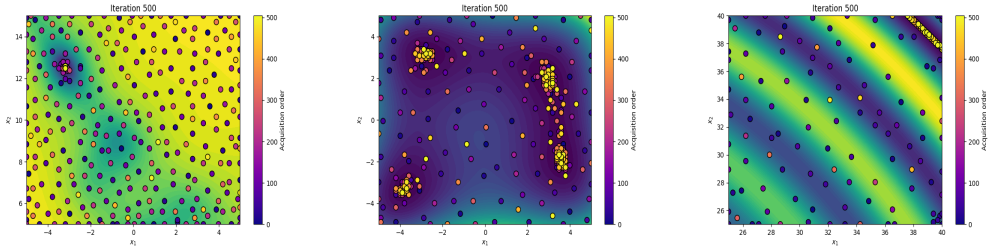

(a) BraninRcos2 (b) Himmelblau's (c) Forrester

Figure 2: Convergence performance on different objective functions with multiple local optima: (a) BraninRcos2; (b) Himmelblau's; and (c) Forrester. Darker regions indicate lower values, which are more desirable for minimization task. Dots mark the iterative acquisitions.

4.3 Acquisition function evaluation

We compare NRFS and its one-step-look-ahead variant against several baselines including EI (Jones et al., 1998), PES (Villemonteix et al., 2009; Hernández-Lobato et al., 2014), ϵ-EI, UCB (Srinivas et al., 2009), Tree-structured Parzen Estimator (TPE; Bergstra et al., 2011; Watanabe, 2023), PI (Kushner, 1964; Snoek et al., 2012), Knowledge Gradient (KG; Frazier et al., 2008), and Latin Hypercube random sampling (RS; McKay et al., 2000). EI, ϵ-EI, PI, UCB, PES, KG are implemented within the BOTorch framework (Balandat et al., 2020), while TPE is based on Optuna (Akiba et al., 2019). Detailed reviews of these previous acquisition functions can be found in Appendix A.5.

We have evaluated all acquisition functions on four known objective functions: the GM function introduced in Section 4.2, the Forrester function (Forrester et al., 2008), a Rosenbrock variant (Al-Roomi, 2015; Rosenbrock, 1960), and the Shekel function (Molga & Smutnicki, 2005). These known objective functions require a balance of exploration and exploitation to effectively locate the global optimum. Visualizations of the objective landscapes and analyses of search behaviors under different BO strategies are provided in Appendix A.4. In addition to the four known objective functions, we consider two real-world case studies aimed at identifying the optimal composition of six elemental materials (Hastings et al., 2025). The first case targets the highest stacking fault energy (SFE), while the second focuses on maximizing heat capacity (HC). In both cases, the mapping from composition to property is unknown.

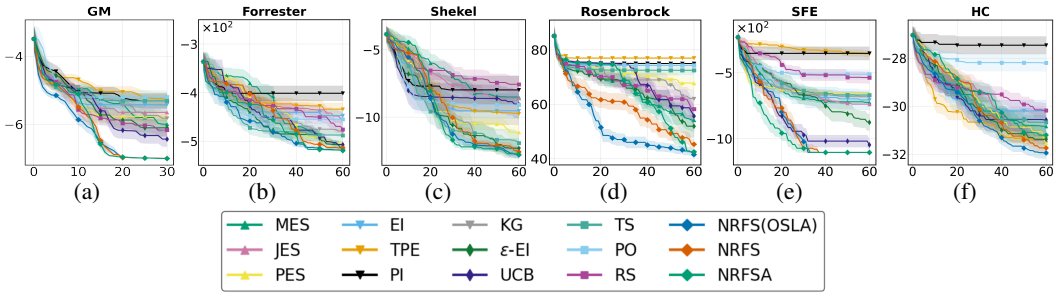

Figure 3: Performance comparison of 20 independent trials for (a) 2D Gaussian mixture model; (b) Forrester; (c) Shekel; (d) Rosenbrock; (e) SFE and (f) HC

In the cases of GM, the Rosenbrock function and SFE, our NRFS consistently outperforms all other competing methods, not only consistently achieving better objective function values across iterations but also exhibiting greater stability. A clearer presentation of their standard deviation (std) trends in Appendix A.6 further highlight their reliability. In the Forrester case, ϵ-EI and UCB also demonstrate competitive performance compared to NRFS based methods. This is because the global optimal region is relatively large and smooth, unlike the other three examples. As a result, ϵ-EI and UCB can more easily acquire points in the global optimal region via random sampling or uncertainty reduction. However, achieving the performance shown in Figure 3b requires sweeping ϵ from 0.1 to 0.9 to identify the best-performing hyperparameter, resulting in significantly higher evaluation cost compared to NRFS based methods. For UCB, even when using a self-adjusting schedule $\beta = 0.5d \cdot \log(n+1) \cdot c$ (Srinivas et al., 2009), selecting the constant c remains challenging due to the lack of prior knowledge about the scale and smoothness of the objective function, making it difficult to define a reasonable hyperparameter range. In contrast, our NRFS-based BO methods do not require hyperparameter tuning compared to these methods. In the Shekel and HC cases, NRFS-based methods are initially outperformed by other acquisition functions in the early iterations. However, as the number of BO iterations increases, they achieve the best overall performance. The strong early performance of value improvement based methods can be attributed to their rapid convergence toward local minima. Nevertheless, as the optimization progresses, these methods eventually fail to provide meaningful acquisitions due to their inherently unbalanced acquisition strategies.

In addition to the sequential NRFS strategy, we also evaluate a batch variant NRFS on the four synthetic objective functions. The result and implementation detail can be found in Appendix A.7. We extend our study to high-dimensional variants of the Forrester function. We benchmark the same 14 acquisition strategies as in Figure 3 across 5D, 10D, 20D, and 50D problems. The experimental details can be found in Appendix A.8. We have also performed benchmarking on other commonly used objective functions, limited budget real-world case studies, together with robustness test against noise at different levels for all of the acquisition methods mentioned above. Most of these results in Appendix A.9 and A.10 have demonstrated consistent superior performance by NRFS. Beyond empirical performance comparison, we further analyze the acquisition behaviors of EI, PES, and NRFS on the SFE task, highlighting how NRFS can facilitate real-world materials discovery. Detailed search behavior analyses can be found in Appendix A.11. Appendix A.3 provides the complexity analysis and computational cost comparison for different acquisition methods.

5 Conclusion & Future Research

We have proposed a novel BO strategy, NRFS, which acquires the query candidate that has the maximum probability to be the true optimizer. For acquisition efficiency, we convert it to the estimation of function space coverage by focusing on surrogate functions whose optima are likely to be the true optimum. We provide strong empirical evidence that NRFS can converge to global optima for a diverse family of benchmark objective functions and demonstrates superior empirical advantages over existing BO baselines. Moreover, NRFS not only provides a new BO strategy, but also opens a new research direction on how one should utilize the surrogate models in BO more efficiently. The computation of other acquisition functions in other BO methods can also be based on the sampled functions from NRFS. A natural extension of NRFS is to multi-objective optimization, where truncation is defined by the Pareto front rather than a single threshold, making sampling substantially more challenging.

AI USE STATEMENT

In this work, we used generative AI tools for grammar checking and sentence rephrasing.

REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2623–2631, 2019.

Ali R. Al-Roomi. Unconstrained Single-Objective Benchmark Functions Repository, 2015. URL `https://www.al-roomi.org/benchmarks/unconstrained`.

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47:235–256, 2002.

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.

James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. *Advances in neural information processing systems*, 24, 2011.

Chao Bian, Xiaofang Wang, Changjun Liu, Xinyu Xie, and Liu Haitao. Impact of exploration-exploitation trade-off on UCB-based Bayesian optimization. In *IOP Conference Series: Materials Science and Engineering*, volume 1081, pp. 012023. IOP Publishing, 2021.

Ilija Bogunovic, Jonathan Scarlett, Andreas Krause, and Volkan Cevher. Truncated variance reduction: A unified approach to Bayesian optimization and level-set estimation. *Advances in neural information processing systems*, 29, 2016.

Adam D Bull. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12(10), 2011.

Antonio Candelieri. Mastering the exploration-exploitation trade-off in Bayesian optimization. *arXiv preprint arXiv:2305.08624*, 2023.

Nuojin Cheng, Leonard Papenmeier, Stephen Becker, and Luigi Nardi. A unified framework for entropy search and expected improvement in Bayesian optimization. *arXiv preprint arXiv:2501.18756*, 2025.

Nicolas Chopin. Fast simulation of truncated gaussian distributions. *Statistics and Computing*, 21 (2):275–288, 2011.

George De Ath, Richard M Everson, Alma AM Rahat, and Jonathan E Fieldsend. Greed is good: Exploration and exploitation trade-offs in Bayesian optimisation. *ACM Transactions on Evolutionary Learning and Optimization*, 1(1):1–22, 2021.

Laurence Charles Ward Dixon and Giorgio Philip Szegő. The global optimization problem: an introduction. *Towards Global Optimiation 2*, pp. 1–15, 1978.

Alexander Forrester, Andras Sobester, and Andy Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons, 2008.

Peter I Frazier. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

Peter I Frazier, Warren B Powell, and Savas Dayanik. A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, 47(5):2410–2439, 2008.

Trevor Hastings, Mrinalini Mulukutla, Danial Khatamsaz, Daniel Salas, Wenle Xu, Daniel Lewis, Nicole Person, Matthew Skokan, Braden Miller, James Paramore, et al. Accelerated multi-objective alloy discovery through efficient Bayesian methods: Application to the FCC high entropy alloy space. *Acta Materialia*, pp. 121173, 2025.

Philipp Hennig and Christian J Schuler. Entropy search for information-efficient global optimization. *The Journal of Machine Learning Research*, 13(1):1809–1837, 2012.

José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in neural information processing systems*, pp. 918–926, 2014.

David M Himmelblau et al. *Applied nonlinear programming*. McGraw-Hill, 2018.

Carl Hvarfner, Frank Hutter, and Luigi Nardi. Joint entropy search for maximally-informed bayesian optimization. *Advances in Neural Information Processing Systems*, 35:11494–11506, 2022.

Ali Jalali, Javad Azimi, and Xiaoli Fern. Exploration vs exploitation in Bayesian optimization. *CoRR*, 2012.

Donald R Jones. Large-scale multi-disciplinary mass optimization in the auto industry. In *MOPTA 2008 Conference (20 August 2008)*, volume 64, 2008.

Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.

Harold J Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. 1964.

Yifang Li and Sujit K Ghosh. Efficient sampling methods for truncated multivariate normal and student-t distributions subject to linear inequality constraints. *Journal of Statistical Theory and Practice*, 9(4):712–732, 2015.

Michael D McKay, Richard J Beckman, and William J Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 2000.

Marcin Molga and Czesław Smutnicki. Test functions for optimization needs. *Test functions for optimization needs*, 101(48):32, 2005.

Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pp. 63–71. Springer, 2003.

Howard H Rosenbrock. An automatic method for finding the greatest or least value of a function. *The computer journal*, 3(3):175–184, 1960.

Ilya O Ryzhov. On the convergence rates of expected improvement methods. *Operations Research*, 64(6):1515–1528, 2016.

Kenan Šehić, Alexandre Gramfort, Joseph Salmon, and Luigi Nardi. Lassobench: A high-dimensional hyperparameter optimization benchmark suite for lasso. In *International Conference on Automated Machine Learning*, pp. 2–1. PMLR, 2022.

Michael D Shields and Jiaxin Zhang. The generalization of latin hypercube sampling. *Reliability Engineering & System Safety*, 148:96–108, 2016.

Bernard W Silverman. *Density estimation for statistics and data analysis*. Routledge, 2018.

Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25, 2012.

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.

Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias W Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE transactions on information theory*, 58(5):3250–3265, 2012.

Shion Takeno, Yu Inatsu, and Masayuki Karasuyama. Randomized gaussian process upper confidence bound with tighter Bayesian regret bounds. In *International Conference on Machine Learning*, pp. 33490–33515. PMLR, 2023.

Panagiotis Tsilifis, Roger G Ghanem, and Paris Hajali. Efficient Bayesian experimentation using an expected information gain lower bound. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1): 30–62, 2017.

Emmanuel Vazquez and Julien Bect. Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and inference*, 140(11):3088–3095, 2010.

Julien Villemonteix, Emmanuel Vazquez, and Eric Walter. An informational approach to the global optimization of expensive-to-evaluate functions. *Journal of Global Optimization*, 44(4):509, 2009.

Hao Wang, Bas van Stein, Michael Emmerich, and Thomas Back. A new acquisition function for Bayesian optimization based on the moment-generating function. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 507–512. IEEE, 2017.

Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient Bayesian optimization. In *International conference on machine learning*, pp. 3627–3635. PMLR, 2017.

Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka. Batched large-scale Bayesian optimization in high-dimensional spaces. In *International Conference on Artificial Intelligence and Statistics*, pp. 745–754. PMLR, 2018.

Ziyu Wang and Nando de Freitas. Theoretical analysis of Bayesian optimisation with unknown Gaussian process hyper-parameters. *arXiv preprint arXiv:1406.7758*, 2014.

Shuhei Watanabe. Tree-structured Parzen Estimator: Understanding its algorithm components and their roles for better empirical performance. *arXiv preprint arXiv:2304.11127*, 2023.

James Wilson, Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *International conference on machine learning*, pp. 10292–10302. PMLR, 2020.

A APPENDIX

A.1 NRFS ANALYTICAL FORM DERIVATION

As discussed in Section 3.2 of the main text, only function values exceeding the current best value can correspond to a potential optimum at $\boldsymbol{x}$. Accordingly, the expected optimum value at $\boldsymbol{x}$ is given by

$$\alpha_n^{\text{NRFS}}(\boldsymbol{x}) := \mathbb{E}[y^*|\boldsymbol{x}, \mathcal{D}_n] = \int_{-\infty}^{Y_n^*} y^* p_{\text{trunc}}(y^*) dy^*, \tag{18}$$

where $p_{\text{trunc}}(y^*)$ denotes the normalized truncated distribution ensuring that the optimum value lies in the region beyond the current best observation. Here we use $\alpha(\boldsymbol{x})$ to denote $\alpha_n^{\text{NRFS}}(\boldsymbol{x})$ for any n. Notice that this formulation can be applied to any surrogate model not necessarily GPs. If we apply GP as the surrogate model, the truncated density is given by

$$p_{\text{trunc}}(y^*) \;=\; \frac{\varphi(z^*(\boldsymbol{x}))}{\sigma_n(\boldsymbol{x})\Phi(z_n^*(\boldsymbol{x}))}, \tag{19}$$

where $z_n^*(\boldsymbol{x}) = \dfrac{Y_n^* - \mu_n(\boldsymbol{x})}{\sigma_n(\boldsymbol{x})}$ and $z^*(\boldsymbol{x}) = \dfrac{y^* - \mu_n(\boldsymbol{x})}{\sigma_n(\boldsymbol{x})}$. Combining (18) and (19), the analytical form of NRFS can be rewritten as:

$$\begin{aligned}
\alpha_n^{\text{NRFS}}(\boldsymbol{x}) :&= \int_{-\infty}^{Y_n^*} y^* \frac{\varphi(z^*(\boldsymbol{x}))}{\sigma_n(\boldsymbol{x})\Phi(z_n^*(\boldsymbol{x}))} dy^* \\
&= \frac{1}{\Phi(z_n^*(\boldsymbol{x}))} \int_{-\infty}^{Y_n^*} (z^*(\boldsymbol{x})\sigma_n(\boldsymbol{x}) + \mu_n(\boldsymbol{x})) \frac{\varphi(z^*(\boldsymbol{x}))}{\sigma_n(\boldsymbol{x})} d\left(z^*(\boldsymbol{x})\sigma_n(\boldsymbol{x}) + \mu_n(\boldsymbol{x})\right) \\
&= \frac{1}{\Phi(z_n^*(\boldsymbol{x}))} \int_{-\infty}^{z_n^*} (z^*(\boldsymbol{x})\sigma_n(\boldsymbol{x}) + \mu_n(\boldsymbol{x})) \varphi(z^*(\boldsymbol{x})) \, dz^*(\boldsymbol{x}) \\
&= \frac{1}{\Phi(z_n^*(\boldsymbol{x}))} \left(\sigma_n(\boldsymbol{x}) \int_{-\infty}^{z_n^*} z^*(\boldsymbol{x}) \varphi(z^*(\boldsymbol{x})) \, dz^*(\boldsymbol{x}) + \mu_n(\boldsymbol{x}) \int_{-\infty}^{z_n^*} \varphi(z^*(\boldsymbol{x})) \, dz^*(\boldsymbol{x}) \right) \\
&= \frac{1}{\Phi(z_n^*(\boldsymbol{x}))} \left(-\sigma_n(\boldsymbol{x}) \varphi(z_n^*(\boldsymbol{x})) + \mu_n(\boldsymbol{x}) \Phi(z_n^*(\boldsymbol{x})) \right) \\
&= \mu_n(\boldsymbol{x}) - \frac{\varphi(z_n^*(\boldsymbol{x}))}{\Phi(z_n^*(\boldsymbol{x}))} \sigma_n(\boldsymbol{x})
\end{aligned} \tag{20}$$

Minimizing the acquisition function in (20) enables identification of the global minimizer. For the maximization setting, the corresponding acquisition function is given by:

$$\alpha_n^{\text{NRFS}}(\boldsymbol{x}) := \mu_n(\boldsymbol{x}) + \sigma_n(\boldsymbol{x}) \frac{\varphi\big(z_n^*(\boldsymbol{x})\big)}{1 - \Phi\big(z_n^*(\boldsymbol{x})\big)}. \tag{21}$$

Maximizing the acquisition function in (21) enables identification of the global maximizer.

From (20) and (21), we observe that the analytical form of NRFS resembles UCB-type acquisition functions. Based on (20), we identify four key properties that will be repeatedly used in the subsequent proofs. Assume a minimization problem, and let $\alpha(\boldsymbol{x}; \mu, \sigma, y^*)$ denote the acquisition function evaluated at $\boldsymbol{x}$, parameterized by the GP predictive mean μ, standard deviation σ, and threshold objective value y^*.

Property 1. *For a fixed input $\boldsymbol{x}$, fixed posterior standard deviation σ and fixed current best y^*, the acquisition function $\alpha(\boldsymbol{x})$ is monotonically increasing with respect to the posterior mean μ. That is, for any $\mu_1 \leq \mu_2$,*

$$\alpha(\boldsymbol{x}; \mu_1, \sigma, y^*) \leq \alpha(\boldsymbol{x}; \mu_2, \sigma, y^*),$$

Property 2. *For a fixed input $\boldsymbol{x}$, fixed posterior mean μ and fixed current best y^*, the acquisition function $\alpha(\boldsymbol{x})$ is monotonically decreasing with respect to the posterior standard deviation σ. That is, for any $\sigma_1 \leq \sigma_2$,*

$$\alpha(\boldsymbol{x}; \mu, \sigma_1, y^*) \geq \alpha(\boldsymbol{x}; \mu, \sigma_2, y^*).$$

Properties 1 and 2 follow directly from the fact that assigning more weight beyond a fixed threshold, either by increasing the mean or by increasing the variance, pushes the expected value farther away from the current best threshold.

Property 3. *For a fixed input $\boldsymbol{x}$, fixed posterior mean μ and fixed posterior standard deviation σ, function $\alpha(\boldsymbol{x}) - y^*$ is monotonically decreasing with respect to the previous best observation y^*. That is, for any $y_1^* \leq y_2^*$,*

$$\alpha(\boldsymbol{x}; \mu, \sigma, y_1^*) - y_1^* \geq \alpha(\boldsymbol{x}; \mu, \sigma, y_2^*) - y_2^*.$$

Property 3 follows from the fact that the first derivative of the inverse Mills ratio $r(z)$ takes values in the range $[0, 1]$ which can be proved as follows.

For any $z = \frac{y^* - \mu}{\sigma}$, we have

$$\mathbb{E}[Z \mid Z \leq z] < z \implies -r(z) < z \implies r(z) + z > 0, \tag{22}$$

and

$$\mathrm{Var}(X \mid X \leq z) = 1 - z \cdot r(z) - r(z)^2 > 0. \tag{23}$$

Meanwhile $r'(z) = -r(z)(r(z) + z)$, thus

$$r(z)(r(z) + z) > 0 \implies r'(z) < 0 \tag{24}$$

and

$$1 - z \cdot r(z) - r(z)^2 > 0 \implies r'(z) > -1 \tag{25}$$

$\alpha(\boldsymbol{x}; \mu, \sigma, y^*) - y^* = \mu - \sigma r(z) - y^*$. For fixed $\boldsymbol{x}$, both μ and σ are constants with respect to y^*. Therefore, $\frac{\partial z}{\partial y^*} = \frac{1}{\sigma}$. Hence,

$$\frac{\partial \alpha(\boldsymbol{x}; \mu, \sigma, y^*)}{\partial y^*} = -\sigma r'(z) \frac{\partial z}{\partial y^*} - 1 = -r'(z) - 1. \tag{26}$$

For the inverse Mills ratio, it is known that$-1 < r'(z) < 0, \qquad z \in \mathbb{R}$. Consequently,

$$-1 < \frac{\partial \alpha(\boldsymbol{x}; \mu, \sigma, y^*)}{\partial y^*} < 0. \tag{27}$$

Therefore, $\alpha(\boldsymbol{x}; \mu, \sigma, y^*)$ is strictly decreasing with respect to y^*. **Property** 3 indicates an instinct that the acquisition function assigns higher priority to regions whose thresholds are worse. Consequently, these regions may offer greater potential for improving upon the incumbent optimum.

Property 4. *Given a minimization problem, for any input $\boldsymbol{x}$, the acquisition function $\alpha(\boldsymbol{x})$ is upper bounded by the current best objective value y_T^*. That is,*

$$\alpha(\boldsymbol{x}) \leq y_T^*$$

Property 4 follows directly from the fact that the expectation of the possible optimum is below the current best threshold for a minimization problem.

Property 1 and 2 indicate that, if a future observation location has either a higher expected value or greater uncertainty, the acquisition function will preferentially select that region. **Property** 4 provides an upper bound that previous methods may overlook: the true optimum must lie below the current best observed value. Therefore, the optimum lower bound must exceed the current best; otherwise, as illustrated in Figure 1, when the prior is far from the true objective value, the confident region may fail to include the true optimum at its location. In such cases, applying an underestimated upper bound to estimate acquisition functions such as UCB can prevent the true optimum from ever entering the confidence region, violating the standard GP-UCB condition $r_T^{\mathrm{reg}} = \min_{t=1,\ldots,T}\{y_t - Y^*\} \leq 2\beta \min_{t=1,\ldots,T} \sigma(\boldsymbol{x}_t)$, where the coefficient β scales the posterior standard deviation and controls the exploration–exploitation trade-off (Takeno et al., 2023; Srinivas et al.,

2009). However, such heuristic schedules are difficult to tune across different problem settings and may result in overly conservative exploration when β is large, or excessive exploration around local optima when β is small. In contrast, NRFS determines the confidence scaling directly from the GP posterior, adapting to the candidate's mean and variance. As a result, NRFS enforces a necessary condition for global optimality which is the confidence region must at least cover the current best observation.

In BO, two popular approaches are commonly used to characterize convergence rates: (1) the cumulative regret defined as $R_T^{reg} = \sum_{t=1}^{T} f^*(\boldsymbol{x}_t) - f^*(\boldsymbol{x}^*)$, and (2) the simple regret defined as $r_T^{\text{reg}} = \min_{t=1,\dots,T} f^*(\boldsymbol{x}_t) - f^*(\boldsymbol{x}^*)$. In our analysis, we focus on demonstrating that the simple regret converges to zero as the number of iterations increases, which directly aligns with our objective of showing that the method is able to identify the global optimum.

A.2 Proof of Lemmas

A.2.1 Proof of **Lemma** 1

Proof. Let $\boldsymbol{x}_T \in \mathcal{X}$ be an acquisition point satisfying $\sigma_T(\boldsymbol{x}_T) > \varepsilon$, and let $X_T = \{\boldsymbol{x}_1, \dots, \boldsymbol{x}_T\}$ denote the set of previously observed locations. Consider the RBF kernel $k_T(\boldsymbol{x}, \boldsymbol{x}') = \sigma_f^2 \exp\left(-\frac{\|\boldsymbol{x}-\boldsymbol{x}'\|^2}{2\ell^2}\right)$, where the kernel hyperparameters are bounded, ensuring that the surrogate GP remains well conditioned and avoids degenerate kernel parameterizations.

For a noise-free GP, the posterior variance conditioned on all observations in X_T is $\sigma_T^2(\boldsymbol{x} \mid X_T) = k_T(\boldsymbol{x}, \boldsymbol{x}) - k_T(\boldsymbol{x}, X_T) K_T^{-1} k_T(X_T, \boldsymbol{x})$, where $K_T = k_T(X_T, X_T)$. Conditioning on additional observations cannot increase the posterior variance. Hence, conditioning only on any single previously observed point provides an upper bound on the posterior variance conditioned on the entire set X_T. In particular, let $\boldsymbol{x}_{\text{NN}}(\boldsymbol{x}) = \underset{\boldsymbol{x}_i \in X_T}{\arg\min} \|\boldsymbol{x} - \boldsymbol{x}_i\|$ be the nearest previously observed point. Then

$$\begin{aligned}\sigma_T^2(\boldsymbol{x} \mid X_T) &\le \sigma_T^2(\boldsymbol{x} \mid \boldsymbol{x}_{\text{NN}}) = k_T(\boldsymbol{x}, \boldsymbol{x}) - \frac{k_T(\boldsymbol{x}, \boldsymbol{x}_{\text{NN}})^2}{k_T(\boldsymbol{x}_{\text{NN}}, \boldsymbol{x}_{\text{NN}})} \\ &= \sigma_f^2 \left[1 - \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{x}_{\text{NN}}\|^2}{\ell^2}\right)\right] \le \sigma_{f,\max}^2 \left[1 - \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{x}_{\text{NN}}\|^2}{\ell_{\min}^2}\right)\right].\end{aligned}$$

Therefore, for any point satisfying $\sigma_T(\boldsymbol{x}) > \varepsilon$, we must have

$$\varepsilon^2 < \sigma_{f,\max}^2 \left[1 - \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{x}_{\text{NN}}\|^2}{\ell_{\min}^2}\right)\right].$$

Rearranging gives $\|\boldsymbol{x} - \boldsymbol{x}_{\text{NN}}\| > \delta_\varepsilon$, where $\delta_\varepsilon = \ell_{\min} \sqrt{-\log\left(1 - \frac{\varepsilon^2}{\sigma_{f,\max}^2}\right)} \approx \frac{\ell_{\min}\varepsilon}{\sigma_{f,\max}} > 0$. when ε is sufficiently small.

Thus, every acquisition satisfying $\sigma_T(\boldsymbol{x}_T) > \varepsilon$ must be separated from all previously observed locations by at least δ_ε, i.e., $\|\boldsymbol{x}_T - \boldsymbol{x}_i\| > \delta_\varepsilon, \forall\, \boldsymbol{x}_i \in X_T$.

Since $\mathcal{X}$ is compact and therefore has finite packing number for any fixed $\delta_\varepsilon > 0$, only finitely many points can be placed in $\mathcal{X}$ while maintaining pairwise separation greater than δ_ε. Consequently, only finitely many acquisition points can satisfy $\sigma_T(\boldsymbol{x}_T) > \varepsilon$. This proves the result. □

A.2.2 Proof of **Lemma** 2

Proof. We prove the result by mathematical induction. Consider first the base case $T = 1$, where the evaluated set contains only $\boldsymbol{x}_1$. For the next point $\boldsymbol{x}_2$, the acquisition rule in **Lemma** 2 requires $\mu_2(\boldsymbol{x}_2) - y_1^* < -\epsilon$. Since $y_1^* = y_1$, the Lipschitz continuity of the posterior mean implies $y_1 - \mu_2(\boldsymbol{x}_2) \le L_\mu \|\boldsymbol{x}_2 - \boldsymbol{x}_1\|$. Therefore, $\|\boldsymbol{x}_2 - \boldsymbol{x}_1\| > \frac{\epsilon}{L_\mu}$.

Now, suppose that for some $T \ge 2$, the evaluated points satisfy the pairwise separation property

$$\|\boldsymbol{x}_i - \boldsymbol{x}_j\| > \frac{\epsilon}{L_\mu}, \qquad \forall\, \boldsymbol{x}_i, \boldsymbol{x}_j \in X_T, \quad i \ne j.$$

Let $\boldsymbol{x}_{T+1}$ denote the next point selected by the acquisition function, and let $\boldsymbol{x}_{\mathrm{NN}} \in X_T$ be its nearest neighbor among the previously evaluated points. We show that $\|\boldsymbol{x}_{T+1} - \boldsymbol{x}_{\mathrm{NN}}\| > \frac{\epsilon}{L_\mu}$.

Suppose, for contradiction, that $\|\boldsymbol{x}_{T+1} - \boldsymbol{x}_{\mathrm{NN}}\| \leq \frac{\epsilon}{L_\mu}$. By the Lipschitz continuity of the posterior mean,

$$\mu_{T+1}(\boldsymbol{x}_{T+1}) - y_T^* \geq \mu_{T+1}(\boldsymbol{x}_{T+1}) - y_{\mathrm{NN}} \geq -L_\mu \|\boldsymbol{x}_{T+1} - \boldsymbol{x}_{\mathrm{NN}}\| \geq -\epsilon,$$

which is against the acquisition rule $\mu_{T+1}(\boldsymbol{x}_{T+1}) - y_T^* < -\epsilon$..

Hence, $\|\boldsymbol{x}_{T+1} - \boldsymbol{x}_{\mathrm{NN}}\| > \frac{\epsilon}{L_\mu}$, and therefore $\min_{\boldsymbol{x} \in X_T} \|\boldsymbol{x}_{T+1} - \boldsymbol{x}\| > \frac{\epsilon}{L_\mu}$.

We then augment the evaluated set according to $X_{T+1} = X_T \cup \{\boldsymbol{x}_{T+1}\}$. Because $\boldsymbol{x}_{T+1}$ is separated from every point in X_T by more than $\frac{\epsilon}{L}$, while the induction hypothesis guarantees the same separation among points already in X_T, the pairwise separation property is preserved:

$$\|\boldsymbol{x}_i - \boldsymbol{x}_j\| > \frac{\epsilon}{L_\mu}, \qquad \forall\, \boldsymbol{x}_i, \boldsymbol{x}_j \in X_{T+1}, \quad i \neq j.$$

Thus, by induction, every newly selected point is separated from all previously evaluated points by at least $\frac{\epsilon}{L_\mu}$. In other words, the acquisition rule induces a minimum separation sampling procedure with separation distance $\frac{\epsilon}{L_\mu}$.

Finally, because the design space is bounded, it can contain only finitely many points that are pairwise separated by a distance greater than $\frac{\epsilon}{L_\mu}$. Consequently, only finitely many points can satisfy the acquisition condition, and the acquisition process cannot continue infinitely. □

A.2.3 Proof of **Lemma** 3

Before proceeding with the proof, we first combine **Lemmas** 1 and 2 into a single auxiliary lemma, which facilitates the subsequent analysis.

Lemma 5. *Let $\mathcal{X}$ be a compact design space and consider a noise-free GP with a continuous kernel. For any fixed constant $\varepsilon,\ C_1,\ C_2 > 0$, suppose acquisitions are selected only from locations $\boldsymbol{x} \in \mathcal{X}$ satisfying $C_1(y_{T-1}^* - \mu_T(\boldsymbol{x})) + C_2\sigma_T(\boldsymbol{x}_T) > \varepsilon$. Then, the total number of such acquisitions is finite.*

For every acquisition, the condition $C_1\big(y_{T-1}^* - \mu_T(\boldsymbol{x})\big) + C_2\sigma_T(\boldsymbol{x}_T) > \varepsilon$ implies that at least one of the two terms must exceed $\frac{\varepsilon}{2}$. Therefore, either $\mu_T(\boldsymbol{x}) - y_{T-1}^* < -\frac{\varepsilon}{2C_1}$, or $\sigma_T(\boldsymbol{x}_T) > \frac{\varepsilon}{2C_2}$.

By **Lemma** 1 and **Lemma** 2, each of these conditions can hold for only finitely many acquisitions. Hence, the total number of acquisitions satisfying the above condition is finite.

Proof. We prove by contradiction that $\sup_{\boldsymbol{x} \in \mathcal{X}} \sigma_T(\boldsymbol{x}) \to \sigma_0 > 0$ is impossible. We consider a noise-free GP model, for which evaluating the objective at an observation location drives the posterior variance at this location to zero. Since the design space $\mathcal{X}$ is compact and $\sigma_T(\boldsymbol{x})$ is continuous, the supremum of σ_T over $\mathcal{X}$ is attained. If $\sup_{\boldsymbol{x} \in \mathcal{X}} \sigma_T(\boldsymbol{x}) \to \sigma_0 > 0$, then by compactness there exists a convergent subsequence $\boldsymbol{x}_{T_k} \to \boldsymbol{x}_0 \in \mathcal{X}$ such that $\sigma_{T_k}(\boldsymbol{x}_0) \to \sigma_0 > 0$. This implies that $\boldsymbol{x}_0$ is never selected by the acquisition function in all future iterations, i.e., $\alpha_T(\boldsymbol{x}_T) > \alpha_T(\boldsymbol{x}_0)$ holds for all sufficiently large T. We show that this persistently strict inequality leads to a contradiction.

The overall strategy is to rewrite $\alpha_T(\boldsymbol{x}_T) - y_{T-1}^*$ into the form required by **Lemma** 5 by applying appropriate bounds on the inverse Mills ratio.

First we prove that $\alpha_T(\boldsymbol{x}_0) - y_{T-1}^*$ has an upper bound as if variance on $\boldsymbol{x}_0$ is no smaller than σ_0^2. Based on **Property** 1 and **Property** 3 We have

$$\begin{aligned}
\alpha_T(\boldsymbol{x}_0) - y_{T-1}^* &= \alpha_T(\boldsymbol{x}_0; \mu_T(\boldsymbol{x}_0), \sigma_0, y_{T-1}^*) - y_{T-1}^* \\
&> \alpha_T(\boldsymbol{x}_0; \mu_{sup}, \sigma_0, y_{T-1}^*) - y_{T-1}^* \\
&> \alpha_T(\boldsymbol{x}_0; \mu_{sup}, \sigma_0, Y^*) - Y^* \\
&= -\sigma_0(r(z_{0,\inf}) + z_{0,\inf}),
\end{aligned} \tag{28}$$

where $z_{0,\inf} = \frac{Y^* - \mu_{\sup}}{\sigma_0}$. Thus, $\alpha_T(\boldsymbol{x}_0) - y_T^*$ is upper bounded by constant $\epsilon_{0,\sup} = -\sigma_0(r(z_{\inf}) + z_{\inf})$. This means all future acquisition values are at least lower than the current best with a certain constant $\epsilon_{0,\sup} = -\sigma_0(r(z_{0,\inf}) + z_{0,\inf})$.

We next establish bounds on the inverse Mills ratio r to express the acquisition function in the form required by **Lemma** 5. Using the derivative $r'(z) = -r(z)\big(r(z) + z\big)$ together with the established bound $-1 < r'(z) < 0$, we derive the corresponding bounds on $r(z)$ as a function of z.

Case 1 When $z_T > 0$, we have

$$r(z_T) < \frac{-z_T + \sqrt{z_T^2 + 4}}{2} < 1. \tag{29}$$

Thus, we have

$$\begin{aligned} \mu_T(\boldsymbol{x}_T) - y_{T-1}^* - \sigma_T(\boldsymbol{x}_T) &< \mu_T(\boldsymbol{x}_T) - r(z_T)\sigma_T(\boldsymbol{x}_T) - y_{T-1}^* \\ &= \alpha_T(\boldsymbol{x}_T) - y_{T-1}^* < \epsilon_{0,\inf}, \end{aligned} \tag{30}$$

Therefore, by **Lemma** 5, only finitely many acquisitions can satisfy this condition, and hence the acquisition process cannot continue indefinitely under this case.

Case 2 When $z_T < 0$, we have

$$r(z_T) < \frac{-z_T + \sqrt{z_T^2 + 4}}{2} < -z_T + 1. \tag{31}$$

Thus we have,

$$\begin{aligned} -\sigma_T(\boldsymbol{x}_T) &= \mu_T(\boldsymbol{x}_T) + (z_T - 1)\sigma_T(\boldsymbol{x}_T) - y_{T-1}^* \\ &< \mu_T(\boldsymbol{x}_T) - r(z_T)\sigma_T(\boldsymbol{x}_T) - y_{T-1}^* = \alpha_T(\boldsymbol{x}_T) - y_{T-1}^* < \epsilon_{0,\inf}. \end{aligned} \tag{32}$$

Therefore, the same conclusion follows from **Lemma** 1, which implies that only finitely many acquisitions can satisfy this condition. □

Lemma 3 shows that analytical NRFS has sufficient exploration power to ensure that the posterior uncertainty at the true optimizer $\boldsymbol{x}^*$ vanishes as $T \to \infty$. For a noise-free GP with a strictly positive-definite kernel, the posterior uncertainty vanishes if and only if the corresponding location has been previously observed. In particular, let $A:\ \boldsymbol{x} \in \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T\},\ B:\ \sigma_T(\boldsymbol{x}) = 0$. First, if $\boldsymbol{x}$ has been previously observed, i.e., A holds, then the noise-free GP interpolates the observation exactly, and hence $A \Rightarrow B$.

Conversely, suppose that $\boldsymbol{x}$ has not been observed, i.e., $\neg A$ holds. Since the kernel is strictly positive definite, the augmented kernel matrix remains positive definite, and its Schur complement gives $\sigma_T^2(\boldsymbol{x} \mid X_T) = k_T(\boldsymbol{x}, \boldsymbol{x}) - k_T(\boldsymbol{x}, X_T)K_T^{-1}k_T(X_T, \boldsymbol{x}) > 0$. Therefore, $\neg A \Rightarrow \neg B$. Since $\neg A \Rightarrow \neg B$ is logically equivalent to $B \Rightarrow A$, we obtain $A \Rightarrow B$ and $B \Rightarrow A$, and consequently $\boldsymbol{x} \in \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T\} \iff \sigma_T(\boldsymbol{x}) = 0$. By the continuity of the kernel function, approaching any previously observed location is equivalent to the posterior uncertainty converging to zero, i.e.,

$$\boldsymbol{x} \to \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T\} \iff \sigma_T(\boldsymbol{x}) \to 0. \tag{33}$$

If the standard deviation upper bound $\sigma_0 \to 0$ then it is natural to get the conclusion that the optimizer standard deviation also $\sigma(\boldsymbol{x}^*) \to 0$. Then based on (33), we can also conclude that $\boldsymbol{x}^* \to \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T\}$ and finish the proof of the convergence.

A.2.4 Proof of **Lemma** 4

Proof. We adopt an argument analogous to that used in the proof of (A.2.3) to characterize the convergence rate. Specifically, we consider the worst case scenario in which the global optimizer remains undiscovered and is located in the region of highest posterior uncertainty during the whole BO process. Under this assumption, we derive a upper bound on the number of iterations required to reach the minimum resolution around the global optimizer, and hence obtain a corresponding lower bound on the convergence rate.

We first formalize the notion of convergence. Given an arbitrarily small distance $d_{\min} > 0$, we consider two acquisition points $\boldsymbol{x}'$ and $\boldsymbol{x}''$ as equivalent if $||\boldsymbol{x}' - \boldsymbol{x}''|| < d_{\min}$. We refer to $d_{\min}$ as the *minimum resolution*.We then characterize the convergence rate by defining T_c as the minimum number of iterations such that at least one acquisition point $\boldsymbol{x}_t$, $1 \le t \le T_c$, satisfies $\|\boldsymbol{x}_t - \boldsymbol{x}^*\| < d_{\min}$. With *minimum resolution* $d_{\min}$, we define $\sigma_{d_{\min}} > 0$ as the corresponding lower bound on posterior uncertainty for unidentified designs. Any design whose posterior standard deviation satisfies $\sigma(\boldsymbol{x}) < \sigma_{d_{\min}}$ is treated as already acquired.

The threshold $\sigma_{d_{\min}}$ is estimated under the assumption that the design space is fully covered by a set of points with pairwise separation at least $d_{\min}$. For the RBF kernel, when $d_{\min}$ is sufficiently small relative to the kernel lengthscale ℓ, the kernel correlation can be approximated by

$$\exp\left(-\frac{d_{\min}^2}{2\ell^2}\right) = 1 - \frac{d_{\min}^2}{2\ell^2} + O\left(\frac{d_{\min}^4}{\ell^4}\right).$$

This provides a direct relationship between the prescribed spatial resolution $d_{\min}$ and the corresponding uncertainty threshold $\sigma_{d_{\min}}$.

To obtain a tractable estimate of the lower bound, we retain the contribution from the nearest neighbor of $\boldsymbol{x}^*$ and treat it as the dominant local contribution. This approximation is justified by the assumption that the optimizer $\boldsymbol{x}^*$ has not yet been acquired and that all previously sampled points remain at least $d_{\min}$ away from $\boldsymbol{x}^*$. Consequently, the nearest neighbor correlation provides the largest kernel contribution among the previously acquired points, while contributions from more distant points decay according to the RBF kernel. Thus,

$$\sigma_T^2(\boldsymbol{x}^*) = \sigma_f^2 - k(\boldsymbol{x}^*, X_T)K_T^{-1}k(X_T, \boldsymbol{x}^*) \approx \sigma_f^2 - k(\boldsymbol{x}^*, X_{M(d)})K_{M(d)}^{-1}k(X_{M(d)}, \boldsymbol{x}^*) \tag{34}$$

where $M(d)$ is the maximum number of points that can be placed around a given point in a d dimensional design space given the prescribed separation $d_{\min}$ and $X_{M(d)}$ represents the set contains all the nearest neighbors to $\boldsymbol{x}^*$. This quantity characterizes the maximum number of nearest neighbor contributions that need to be considered when estimating the local posterior variance at $\boldsymbol{x}^*$. Although a general closed form expression for $M(d)$ remains unknown, $M(d)$ can be determined for any fixed dimension d, with $M(1) = 2$, $M(2) = 6$, and $M(3) = 12$. Therefore, when the input dimension is fixed throughout the optimization process, we treat $M(d)$ as a constant. We further derive a looser lower bound on $\sigma_T^2(\boldsymbol{x}^*)$ by assuming negligible correlations among the neighboring observations.

$$\begin{aligned}
\sigma_T^2(\boldsymbol{x}^*) \approx& \sigma_f^2 - k(\boldsymbol{x}^*, X_{M(d)})K_{M(d)}^{-1}k(X_{M(d)}, \boldsymbol{x}^*) && (35)\\
>& \sigma_f^2 - \frac{M(d)k(\boldsymbol{x}^*, \boldsymbol{x}_i)^2}{\sigma_f^2} && (36)\\
>& \sigma_{f,\min}^2\left[1 - M(d)\exp\left(-\frac{d_{\min}^2}{2\ell_{\max}^2}\right)\right] && (37)\\
=& \sigma_{\inf}^2(\boldsymbol{x}^*) && (38)
\end{aligned}$$

Thus, based on **Property** 1, 2 and 3

$$\begin{aligned}
\alpha_T(\boldsymbol{x}^*) - y_{T-1}^* &= \alpha_T(\boldsymbol{x}^*; \mu_T(\boldsymbol{x}^*), \sigma_T(\boldsymbol{x}^*), y_{T-1}^*) - y_{T-1}^*\\
&\le \alpha_T(\boldsymbol{x}^*; \mu_T(\boldsymbol{x}^*), \sigma_{\inf}(\boldsymbol{x}^*), y_{T-1}^*) - y_{T-1}^*\\
&< \alpha_T(\boldsymbol{x}^*; \mu_{sup}, \sigma_{\inf}(\boldsymbol{x}^*), y_{T-1}^*) - y_{T-1}^*\\
&< \alpha_T(\boldsymbol{x}^*; \mu_{sup}, \sigma_{\inf}(\boldsymbol{x}^*), Y^*) - Y^*\\
&= -\sigma_{\inf}(\boldsymbol{x}^*)(r(z_{*,\inf}) + z_{*,\inf}) && (39)
\end{aligned}$$

If the surrogate model is fitted appropriately, usually we have $Y^* < \mu_{sup}$, so $z_{*,\inf} = \frac{Y^* - \mu_{sup}}{\sigma_{\inf}(\boldsymbol{x}^*)} \to -\infty$ if $\sigma_{\inf}(\boldsymbol{x}^*)$ is a very small value. The inverse Mills ratio can be approximated by

$$r(z_{*,\inf}) = -z_{*,\inf} - \frac{\sigma_{\inf}(\boldsymbol{x}^*)}{\mu_{sup} - y^*} + \mathcal{O}(\sigma_{\inf}^3(\boldsymbol{x}^*)).$$

Thus,

$$\alpha_T(\boldsymbol{x}^*) - y^*_{T-1} < -\sigma_{\inf}(\boldsymbol{x}^*)(r(z_{*,\inf}) + z_{*,\inf})$$
$$\approx \frac{\sigma^2_{\inf}(\boldsymbol{x}^*)}{Y^* - \mu_{sup}} \tag{40}$$
$$= \epsilon_{*,\inf} \tag{41}$$

The simple regret satisfies

$$r_T^{\text{reg}} = y_T - Y^* \leq y_{NN} - Y^* < L_f \|\boldsymbol{x}_{NN} - \boldsymbol{x}^*\|.$$

Using (38), the minimum distance between the observed points and the optimizer can be translated into a corresponding lower bound on the posterior uncertainty at $\boldsymbol{x}^*$. In particular, if we require $\|\boldsymbol{x}_{NN} - \boldsymbol{x}^*\| \leq d_{\min}$, then achieving this requires a corresponding space filling distance, which can be characterized using either **Lemma** 1 or **Lemma** 5, which concludes the proof.

$$d_{nrfs} = \min\left\{\frac{\epsilon_{*,\inf}}{2L_\mu}, \frac{\ell_{\min}\epsilon_{*,\inf}}{2\sigma_{f,\max}}\right\} = \min\left\{\frac{1}{2L_\mu(Y^* - \mu_{sup})}, \frac{\ell_{\min}}{2\sigma_{f,\max}(Y^* - \mu_{sup})}\right\}\sigma^2_{\inf}(\boldsymbol{x}^*).$$

□

Consequently, the number of iterations required to fully cover the design space is bounded by $T_c \leq \left(\frac{\sqrt{d}}{d_{nrfs}}\right)^d$, which provides a sufficient number of iterations to guarantee that at least one acquisition point lies within a distance $d_{\min}$ of the global optimizer $\boldsymbol{x}^*$.

For a conventional space filling strategy or an uncertainty reduction strategy, the distance between the optimizer and its nearest observation can be related to the posterior uncertainty as

$$\sigma(\boldsymbol{x}^*) \approx \sigma_f\sqrt{1 - \exp\left(-\frac{\|\boldsymbol{x}_{\text{NN}} - \boldsymbol{x}^*\|^2}{\ell^2}\right)} \approx \frac{\sigma_f}{\ell}\|\boldsymbol{x}_{\text{NN}} - \boldsymbol{x}^*\|.$$

This relationship suggests that the convergence rate of such exploration based strategies is directly governed by the reduction in posterior uncertainty. In contrast, NRFS can exploit the objective structure while simultaneously resolving the optimizer neighborhood, leading to a convergence rate that is faster than that of conventional exploration based strategies.

A.3 Computational detail

For the discrete case, we sample function values from the truncated Gaussian distribution at each observation location using the inverse-CDF sampling method mentioned in the recent works, which have shown that truncated Gaussian distributions can be sampled efficiently and directly using efficient non-rejection algorithms (Chopin, 2011; Li & Ghosh, 2015). In practice, this is implemented in scipy.stats via truncnorm. We then select the maximum sampled value among all observation locations, compute the histogram of optimizer occurrences over the discrete space, and choose the location with the highest frequency as the next acquisition point.

For the continuous case, the procedure works by repeatedly drawing a single posterior-consistent function sample from the surrogate model. Posterior-consistent function is achieved by sampling a finite set of latent random coefficients and then using them to define a deterministic, differentiable function over the entire search space. For a truncated GP, these latent coefficients are drawn from the corresponding truncated posterior rather than from an unconstrained Gaussian posterior. Once these coefficients are sampled and fixed, they define a deterministic function realization that can be optimized in the same way as in the standard GP setting. Each sampled path is then treated as an ordinary objective function and numerically optimized over the search bounds, using multi-start gradient-based optimization. For each sampled function, the method records the input location where the function attains its optimum, along with the corresponding optimal value. Repeating this process many times produces a collection of sampled optimal inputs and sampled optimal outputs,

which together form an optimizer distribution. This sampling and optimization technique is implemented in BOTorch (Balandat et al., 2020) and is commonly used as a preparatory step for different acquisition functions.

Regarding the number of samples, 1000 samples are sufficient for all cases tested in both the main text and the appendix. In fact, the default setting in BOTorch is only 12 for some low-dimensional cases such as GM, so it is not necessary to use 1000 samples for the method to work effectively. We mainly use 1000 samples for our materials discovery case, where we aim to estimate the density distribution more accurately, especially since the evaluation cost is much higher than that of standard synthetic benchmark functions and acquisition function computation.

For all objective functions, the experimental runs were distributed across 8 Intel(R) Xeon(R) Gold 6248R CPUs.

For KG and one-step-look-ahead NRFS, if we denote the number of fantasy samples in the one-step-look-ahead step as M, then the computational complexity becomes d^2M, which incurs a high cost when computed over the entire design space. To mitigate this, we first use Latin Hypercube Sampling (LHS; Shields & Zhang, 2016) to generate evenly distributed acquisition evaluation points (100) in the design space. The acquisition function is then computed on these points to identify the maximizer. Finally, we select the nearest neighbor of this maximizer from all the grid candidates and use the selected candidate as the next evaluation point. This strategy reduces computational resource consumption to a level comparable with other methods such as PES.

Acquisition function	GM	Modified Rosenbrock	Forrester	Shekel	Branin	BraninRcos2
RS	0.0003	0.0004	0.0003	0.0003	0.0003	0.0003
TPE	0.004	0.021	0.018	0.021	0.019	0.018
ϵ-EI	1.073	0.721	0.703	0.654	0.699	0.711
EI	1.150	1.017	1.132	0.985	1.007	1.185
PI	0.372	0.942	0.739	0.724	0.846	0.730
UCB	0.256	0.776	0.592	0.603	0.702	0.659
NRFS	0.221	1.792	1.437	1.398	0.816	1.131
PES	8.014	11.245	10.458	9.671	8.931	10.120
KG	9.356	10.167	8.134	9.871	10.776	8.804
NRFS (OSLA)	5.432	7.114	5.437	5.557	5.981	5.102
NRFSA	1.098	1.020	1.117	1.225	1.008	1.216

Table 1: Run-time (seconds) per iteration on different objective functions with different BO methods

Table 1 reports the computation time per-iteration for each BO method with the corresponding acquisition function. The computational cost of NRFS is comparable to that of EI and remains stable across iterations. The additional overhead of the one-step-look-ahead NRFS arises primarily from repeated model updates. In our setup, we use $M = 5$ fantasy samples, resulting in the runtime approximately five times that of NRFS. The analytical form of NRFS incurs a computational cost comparable to EI and is among the most computationally efficient variants of NRFS.

Among all methods, random sampling (RS) is the fastest, as expected. TPE exhibits relatively lower computation time compared to EI, primarily due to the efficiency of kernel density estimation relative to GP updates. PES incurs significantly higher computational cost; this is consistent with what has been reported in previous work, as PES typically relies on expectation propagation (EP; Hennig & Schuler, 2012) to derive an analytical approximation of the utility function. For KG, the computational time increases over iterations: it starts at approximately 1 second per iteration and grows to around 15 seconds by iteration 60.

A.4 Search behaviors

In addition to the search behavior analysis on the illustrative Gaussian Mixture (GM) example in Section 4.2, we provide additional analyses based on the BO results on different objective functions discussed in Section 4.3. This extended comparison highlights the differences in search behavior dynamics among commonly used non-hybrid acquisition strategies: EI (Jones et al., 1998), PES (Hernández-Lobato et al., 2014), and our proposed NRFS.

These functions represent three types of multi-modal objective landscapes: 1) The modified Rosenbrock functions exhibit large, smooth local optima along with a sharp global optimum distant to local optima; 2) The 2D Forrester function contains large, smooth local optima and a similarly smooth

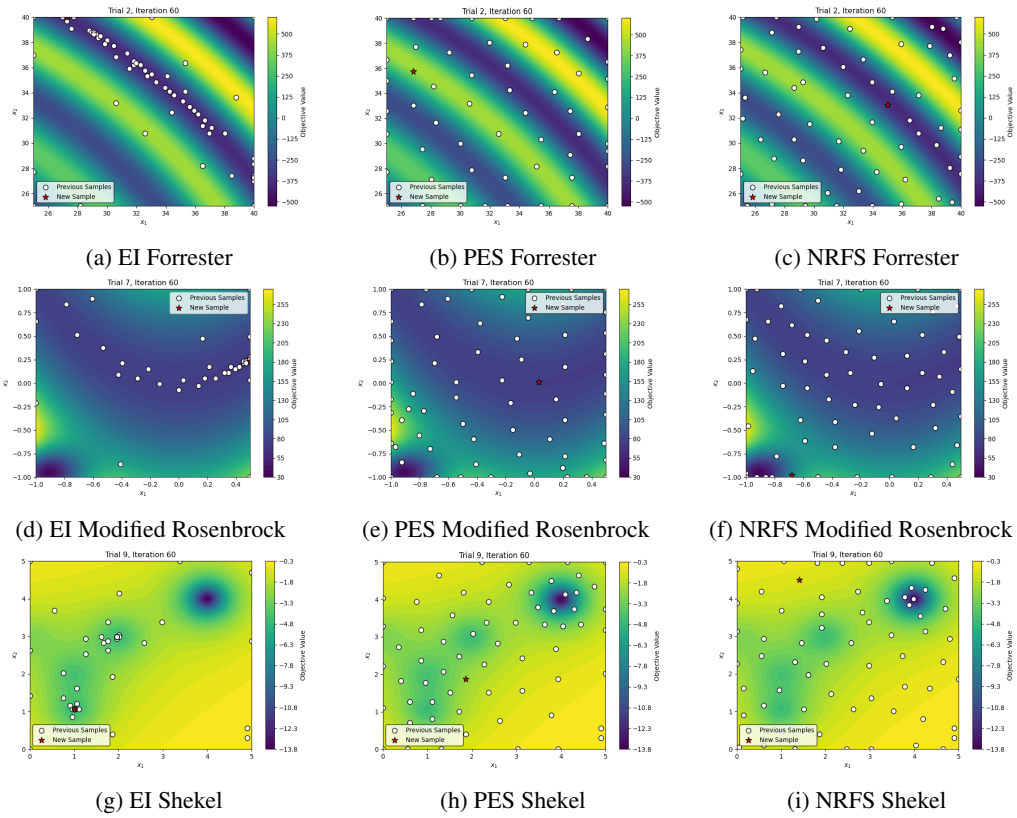

(a) EI Forrester (b) PES Forrester (c) NRFS Forrester

(d) EI Modified Rosenbrock (e) PES Modified Rosenbrock (f) NRFS Modified Rosenbrock

(g) EI Shekel (h) PES Shekel (i) NRFS Shekel

Figure 4: Search behavior comparison between EI-, PES- and NRFS-based BO. Red stars: Next acquisitions; White dots: Previous acquisitions.

global optimum; 3) The Shekel function features multiple sharp peaks, with all optima being narrow and distinct.

For the three minimization cases in Figure 4 above, EI tends to oversample regions near local optima, highly likely to continue until those regions are thoroughly exploited. In contrast, PES exhibits the opposite behavior: even after identifying a candidate in a region likely to contain the global optimum, it continues to explore points that reduce uncertainty about the optimizer's location, rather than concentrating on the regions with high potential. NRFS demonstrates a more balanced search strategy. Its acquisition distribution is closely aligned with the objective landscape, assigning higher acquisition frequency to regions closer to the optimum.

A.5 Baseline Methods

In this section, we review the existing acquisition functions and explain why they are used as our comparison baselines. We start with EI. As illustrated in Fig. 1, EI and PO exhibit similar performance, both being affected by the oversampling issue. Many prior studies (Hennig & Schuler, 2012) also use EI for GP global PO estimation, however, the underlying principles of PO and EI are fundamentally different. Consider a simple case with two candidate buckets. While most functions identify the first bucket as containing the optimizer, its expected improvement is lower than that of the second bucket because the second buckets contains optimum with extreme value. As a result, EI selects the second bucket due to the higher expected improvement. However, from an intuitive perspective, the first bucket should be preferred since it has a higher likelihood of containing the true objective function. This misstep arises because EI prioritizes potential magnitude over the likelihood of containing the true optimum which is an inherent property of the objective that remains fixed regardless of human-defined expectations. The difference between EI and PO in Fig. 1 is not pronounced because the objective values follow a Gaussian distribution with small variance, different from the earlier examples with those following extreme-value distributions. When the Gaussian

distribution has larger variance their difference becomes more apparent, especially in the design space boundary region (Hennig & Schuler, 2012).

For uncertainty reduction based methods, PES aims to reduce the number of remaining candidate buckets rather than directly selecting the most probable one. Methods with function uncertainty reduction ability like UCB, TVR and EIG (Tsilifis et al., 2017), primarily focus on decreasing the total number of possible functions across all buckets. This approach can be inefficient, as it devotes unnecessary effort to keep reducing the number of functions including the buckets that have already been considered.

Besides the most popular acquisition functions such as EI, PES and UCB. TPE is a variant of EI, and we include it in our comparison because it also incorporates both truncated modeling and density estimation over the design space. Given a minimization problem, TPE sets the threshold above the current best value, and uses candidate locations with objective values below this threshold to estimate the optimizer density. Our goal is to evaluate how its truncation strategy compares to ours. KG, a popular one-step-look-ahead strategy, that estimates the expected value of information from evaluating a candidate point, which we include as a baseline to compare with our one-step-look-ahead NRFS. PI, a replacement function space sampling strategy, that focuses solely on the likelihood of improving over the current best observation rather than targeting global optimality, is used as the baseline to compare the difference between Non-replacement and replacement sampling, objective improvement and optimality. RS is included as a non-informative baseline to assess optimization performance.

Acquisition function	Suboptimal convergence	Inefficient exploration	Heuristic design	Hyperparameter tuning	Complexity	Robustness
EI	×	✓	×	✓	✓	×
TPE	×	✓	×	×	✓	✓
ES	✓	×	×	✓	×	✓
PO	×	✓	✓	✓	✓	✓
TS	×	✓	✓	✓	✓	×
UCB	?	?	×	×	✓	×
RS	✓	×	✓	✓	✓	✓
KG	×	✓	×	✓	×	×
VES	✓	✓	×	×	×	×
NRFS	✓	✓	✓	✓	✓	✓
NRFS(OSLA)	✓	✓	✓	✓	×	×
NRFSA	✓	✓	✓	✓	✓	✓

Table 2: Comparison of different acquisition functions with respect to various drawbacks and challenges. Checkmark indicates that the method effectively addresses the corresponding issue; cross indicates that the method may suffer from the issue or perform poorly on the corresponding metric; question mark indicates that the outcome is inconclusive and depends on the properties of the black-box objective function.

The advantages and limitations of each acquisition function are summarized in Table 2. In summary, most these popular acquisition functions are driven by subjective rewards such as objective improvement, entropy reduction, variance minimization, or their combinations. While some of these strategies can yield strong empirical performance in certain cases, there still lacks a single unified policy strategy that performs efficiently across all scenarios. This leads to an awkward situation in practice: identifying the most appropriate acquisition strategy often requires testing multiple options, merging existing ones, or constructing new hybrids. Rather than investing effort in developing a general theoretical framework, the prevailing trend is to combine heuristics in the hope of achieving a better lower bound of the convergence rate, more focusing on *ad-hoc* engineering attempts by testing empirical performances. However, as long as acquisition is guided by subjective rewards, the process will remain biased towards human-designed targets rather than the true optimizer location, regardless of how much effort is spent on merging and tuning engineering heuristics.

A.6 PERFORMANCE CONSISTENCY

We further extract the standard deviation values to illustrate the performance consistency explicitly in this section.

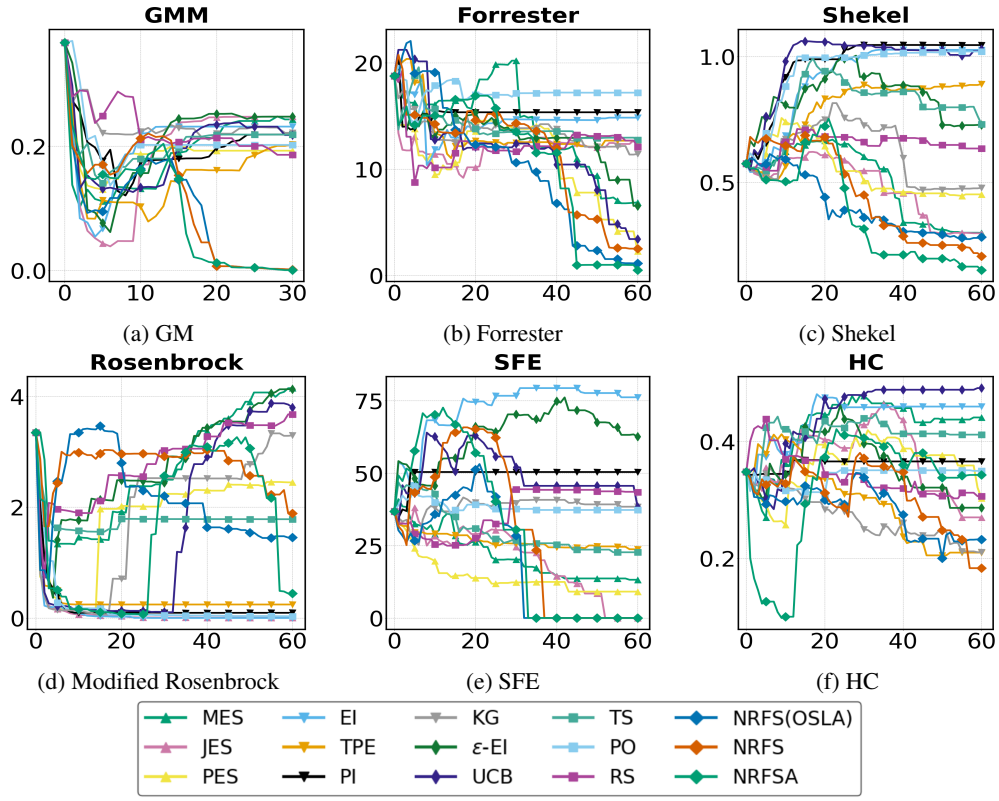

(a) GM (b) Forrester (c) Shekel

(d) Modified Rosenbrock (e) SFE (f) HC

Figure 5: Std performance comparison on different objective functions: (a) 2D Gaussian mixture model; (b) Forrester; (c) Shekel; (d) Modified Rosenbrock; (e) SFE and (f) HC

We observe that, across all six case studies with different objective functions as discussed in the main text, our NRFS exhibits the most stable performance among the cases with GM, Forrester and Shekel. For Modified Rosenbrock, as Fig. 5d shows, EI, TPE, and PI demonstrate more consistent behavior than NRFS, likely due to their tendency to oversample in local optima across all 20 trials. By combining the mean performance reported in the main text with the observed standard deviation values, we find that NRFS exhibits low variability only when it converges to the global optimum. This behavior suggests that the standard deviation can serve as a reliable indicator of proximity to the global optimum: since NRFS does not suffer from oversampling, a near-zero standard deviation implies that the current best is close to the global optimum.

A.7 BATCH ACQUISITION PERFORMANCE

For batch selection, we continue to follow a non-replacement strategy: once the candidate with the highest probability of being the optimizer is identified, we will not consider these sampled function realizations again from future optimizer location estimation. From the remaining realizations, we then select the next most probable optimizer location. This process is repeated until we obtain a full batch of acquisition points. To ensure fairness, we fix the evaluation budget to be identical to that of the sequential acquisition and assess how varying the batch size impacts optimization performance under the same evaluation limit.

As shown in Figure 6, sequential optimization usually has the best performance. The only exception is the Shekel function, where large batch methods perform better. This difference comes from the structure of the Shekel objective: the global optimum is located in a very sharp and narrow region. Once this region is identified, batch acquisition functions tend to place many points directly around

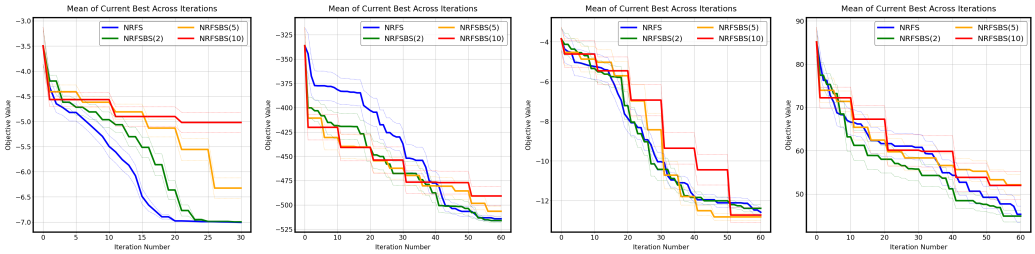

(a) Branin (b) Forrester (c) Shekel (d) RosenbrockModified

Figure 6: Batch performance comparison of 20 independent trials on different objective functions: (a) 2D Gaussian mixture model; (b) Forrester; (c) Shekel and (d) Modified Rosenbrock

that area because the optimizer distribution becomes highly concentrated there, allowing batches to exploit the optimum quickly. In contrast, sequential acquisition updates the surrogate and the current best after each evaluation, which smooths the optimizer distribution and reduces the likelihood of repeatedly sampling the same small region.

A.8 High dimensional performance

The initialization number $(10d)$ for high dimensional benchmark is adjusted according to dimensionality to ensure optimization efficiency. The other experimental setup is similar to the tasks mentioned in the main text.

As Figure 7 shows, the performance gap between acquisition strategies shrinks as the dimensionality increases. By the time the dimension reaches 50, all methods struggle to discover substantially better objective values. Nevertheless, NRFS and NRFS(OSLA) maintain superior performance compared to the other strategies.

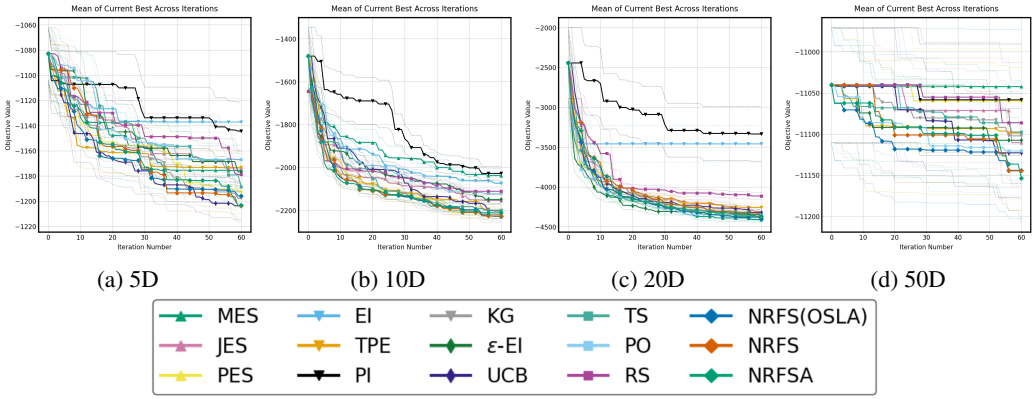

(a) 5D (b) 10D (c) 20D (d) 50D

Figure 7: Optimization performance comparison of 20 independent trials on Forrester functions with different dimensions: (a) 5D; (b) 10D; (c) 20D and (d) 50D

A.9 Additional benchmarking results

Besides the objective functions that require a balanced strategy, we also consider Branin (Dixon & Szegő, 1978) as a purely exploitation-driven case, and BraninRcos2 (Al-Roomi, 2015) as an exploration-driven case, to evaluate NRFS under these extremes. Branin contains three global optima, and identifying any one of them suffices to achieve optimal performance. In contrast, BraninRcos2 resembles the Shekel function, characterized by many local optima and a relatively narrower global optimum region. This implies that once the global optimum region is explored, the optimal solution is nearly recovered.

As Figure 9 shows, in the Branin case, NRFS appears to require more iterations to converge to the global optimum compared to EI. This slower convergence can be interpreted as the inherent trade-

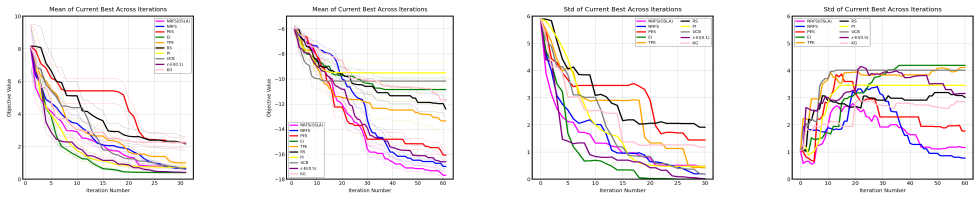

(a) Mean on Branin (b) Mean on BraninRcos2 (c) Std on Branin (d) Std on BraninRcos2

Figure 8: Performance comparison of BO methods on 2 extra objectives: (a, b) current best of 20 trials for Branin and BraninRcos2; (d, e) std of the current best on the same benchmarks.

off for acquiring all global optima rather than focusing on a single one which is similar to case Himmelblau's in Fig. 2b. This observation indicates that a current limitation of NRFS: it may not outperform some existing methods when the problem contains only one optimal region or multiple optimal regions with the same optimal values. In such cases, identifying one region is sufficient to guarantee the global optimum, thus there is no need to acquire candidates from other high potential regions. In the BraninRcos2 setting, the performance trend aligns with that observed in the Shekel function (Figure 3c, Section 4.3). Specifically, PES is able to identify the global optimum region more quickly in the initial iterations. However, once the optimal region is located, NRFS starts to outperform PES in subsequent acquisitions.

We further evaluate our method on three real-world benchmark tasks: (1) a 124-dimensional soft-constrained variant of the Mopta08 automotive design problem, originally introduced by Jones (2008); (2) the 180-dimensional Lasso-DNA objective selection task, proposed by Šehić et al. (2022), which involves optimizing hyperparameters of a sparse regression model used to recover underlying genetic signal patterns; and (3) a 60-dimensional rover trajectory-planning benchmark from Wang et al. (2018), where the optimizer must choose spline-parameterized way points to minimize traversal cost in the presence of obstacles.

It is important to note that these problems have at least 60 input dimensions. Based on the rate of convergence analysis, it would generally take on the order of 10^5 iterations or more to observe full convergence of an acquisition function. Therefore, the 500 iterations considered here correspond to the early stage of optimization, where methods that rapidly exploit local regions may temporarily show better performance. For example, on the Forrest, GM, Shekel and HC problems discussed in the main text, we observe a similar trend: initially, NRFS may be dominated by methods such as EI or TPE that can quickly exploit local regions. However, over future iterations, NRFS tends to outperform these methods, as it does not suffer from local trapping. This illustrates a fundamental dillema: achieving both rapid exploitation and a balanced exploration-exploitation strategy simultaneously is challenging.

To address this gap in practical evaluation where the researchers prefer to boost the exploitation performance on the early stage, we propose a temporary strategy based on a tolerance level ϵ defined in the objective space. This tolerance accounts for the fact that true convergence is unattainable in real applications with limited budgets. Specifically, if the observed objective value falls within the tolerance level relative to the current best, we consider that the global optimum has been effectively found. This tolerance is applied inversely along the optimization direction relative to the current best, meaning that future function samples do not need to surpass the current best to contribute to convergence, achieving values within the tolerance is sufficient. This adjustment encourages the acquisition function to focus more on promising local regions rather than chasing unattainable global accuracy.

The tolerance level is selected as $\epsilon \propto C_0/t$, where t is the total evaluation budget and C_0 is a constant related to black-box objective scale. In boosted-NRFS (bNRFS), C_0 is the mean value of the initial observations and in fully boosted-NRFS (fbNRFS) C_0 is $+\infty$. This proportional relationship captures the intuitive: with an infinite budget, one could achieve arbitrarily high accuracy, while with 0 budget, all acquisitions are effectively treated as potential optimal possible candidates, and through comparison the current best is selected.

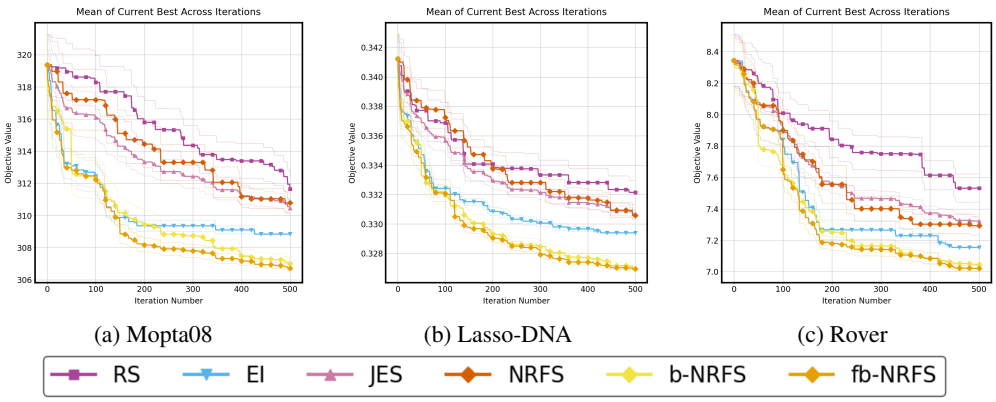

(a) Mopta08 (b) Lasso-DNA (c) Rover

Figure 9: Performance comparison of BO methods on three high dimesional objectives.

Empirically, the boosted NRFS consistently achieves the strongest optimization performance among acquisition strategies with high exploitation power in these three high dimensional applications. This suggests that, when rapid exploitation is the primary objective, boosted NRFS is a more effective choice than competing methods.Meanwhile, increasing the tolerance level can raise the risk that the acquisition becomes trapped in a local optimum. However, this effect is less pronounced in high dimensional settings, where such trapping typically occurs only after the first optimum has been identified, which is often beyond the available evaluation budget.

A.10 Robustness analysis

We have also evaluated robustness by adding varying levels of Gaussian noise to the four objective functions described in the main text. For hybrid acquisition strategies such as UCB and ϵ-EI, hyperparameters were further tuned to ensure optimal performance under different noise levels.

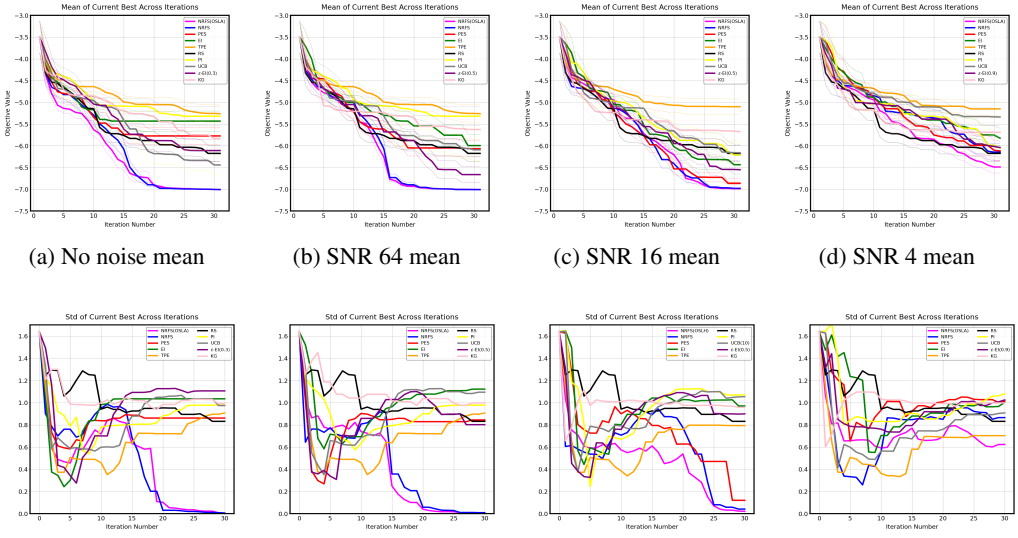

(a) No noise mean (b) SNR 64 mean (c) SNR 16 mean (d) SNR 4 mean

(e) No noise std (f) SNR 64 std (g) SNR 16 std (h) SNR 4 std

Figure 10: Performance comparison for the GM example with different signal-to-noise ratio (SNR) levels.

Based on the results, we conclude that NRFS is robust to different noise levels corresponding to SNR values above 16 across all four objective functions. It even demonstrates robustness with high noise (SNR = 4) in the Modified Rosenbrock and Forrester cases. For NRFS, adding noise to objective function responses broadens the effective function space, which is not necessarily detrimental. In

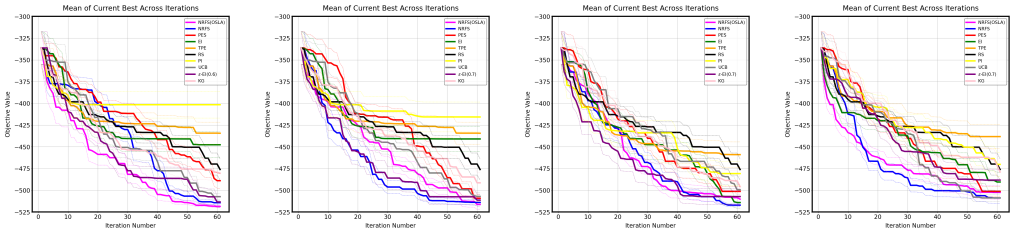

(a) No noise mean (b) SNR 64 mean (c) SNR 16 mean (d) SNR 4 mean

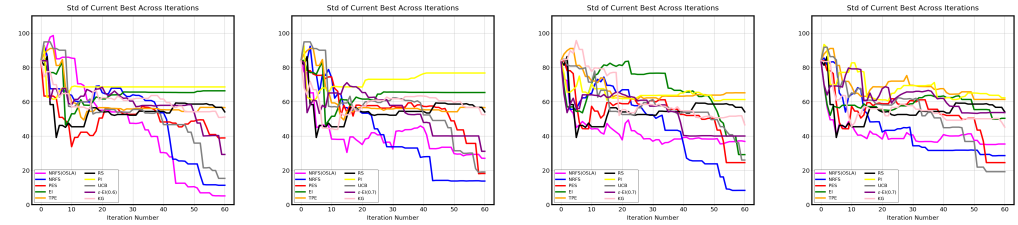

(e) No noise std (f) SNR 64 std (g) SNR 16 std (h) SNR 4 std

Figure 11: Performance comparison on Forrester function with different SNR levels.

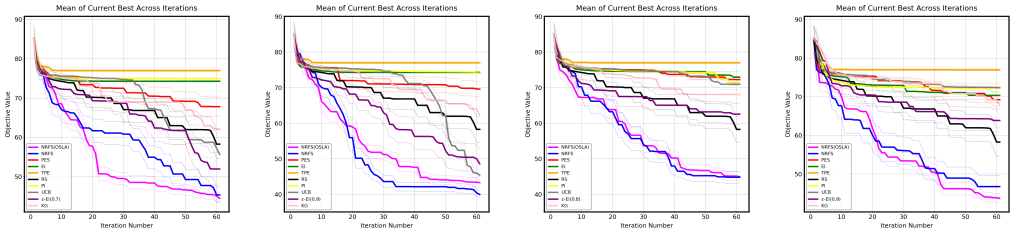

(a) No noise mean (b) SNR 64 mean (c) SNR 16 mean (d) SNR 4 mean

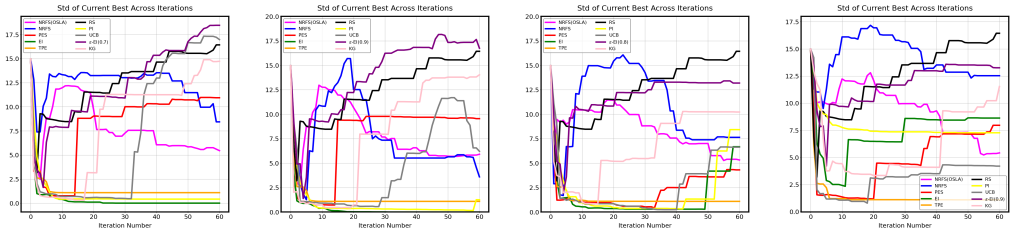

(e) No noise std (f) SNR 64 std (g) SNR 16 std (h) SNR 4 std

Figure 12: Performance comparison on modified Rosenbrock function with different SNR levels.

fact, as shown in Figures. 11b and 12b, small amount of noise can improve performance compared to the noise-free cases. The primary negative impact of noise arises from inaccurate estimation of the previous best values. When the noise level is high, overestimation of the incumbent solution may cause the true optimum to be bypassed, and our non-replacement sampling strategy may consequently skip the global optimum.

While one-step-look-ahead NRFS is occasionally outperformed by standard NRFS, as shown in Figures. 11b, 11d and 12b, this degradation in performance can be attributed to the fact that the one-step-look-ahead strategy may increase the risk of overestimating the current best.

Among the remaining methods, TPE demonstrates the most stable performance, showing no significant degradation or improvement across all four objective functions under varying noise levels. This robustness arises from TPE's use of a percentile-based threshold over top-performing observations, rather than relying solely on the previously observed best. In other words, with a given noise level,

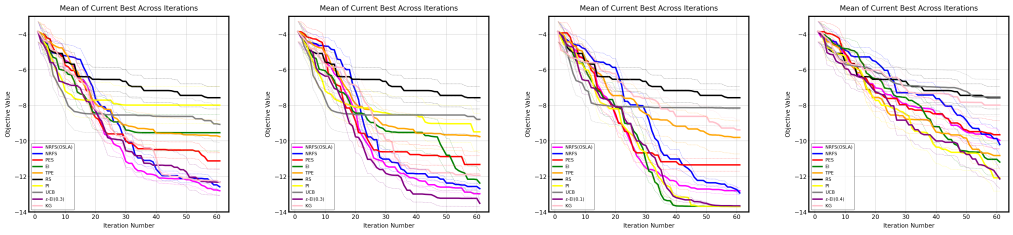

(a) No noise mean (b) SNR 64 mean (c) SNR 16 mean (d) SNR 4 mean

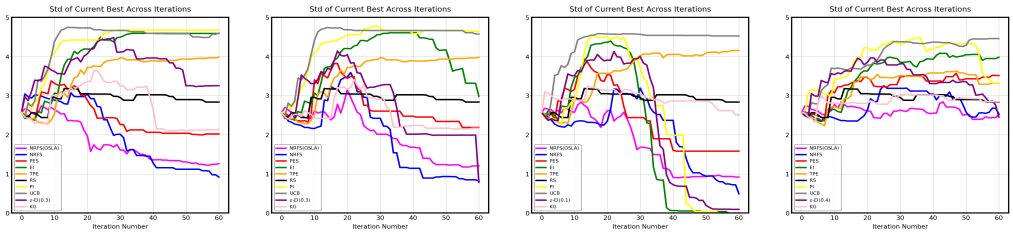

(e) No noise std (f) SNR 64 std (g) SNR 16 std (h) SNR 4 std

Figure 13: Performance comparison on modified Shekel function with different SNR levels.

the deviation in the lower bound of the top-performing set is smaller than that of the previous best values. For instance, if the mean of observations is used as the threshold, the deviation between the noisy threshold and the true threshold becomes zero. However, this stability also reduces its ability to escape local optimal regions, thereby lowering the chance of identifying the global optimum.

The other methods are affected by noise in different ways. An interesting observation is that, after introducing noise, methods such as EI and PI can escape from regions of local optima. This is because the surrogate model becomes smoother with noise compared to the one learned in the corresponding noise-free case. This results in wider modes, making it easier for these methods to enter or escape such regions. However, the limitation is clear: only a small amount of noise helps prevent from oversampling. When the noise level is too high, it still leads to degraded performance. Since we are dealing with a black-box objective function, it is not possible to determine whether the added noise is within a safe or effective range. As a result, introducing noise cannot be considered a stable or reliable strategy to avoid oversampling.

A.11 Real-world application acquisition comparison

The top 5% performing materials are distributed across three regions, with only one region containing the global optimizer. We compare the performance of our NRFS with the BO methods using EI and PES acquisition functions. Figure 14 illustrates the BO performance difference, where NRFS successfully finds the global minimum but EI and PES fail within 40 iterations. As illustrated, EI gets trapped in the local minimum; PES fails to identify the exact global minimum after getting close, suffering from its focus on information gain instead of acquisition in potential optimal region.

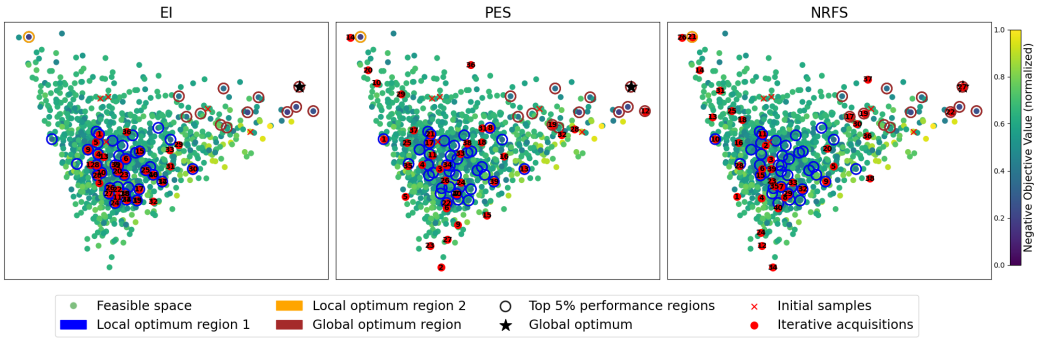

Figure 14: Acquisition behavior comparison of different strategies on the SFE material rediscovery task over 40 iterations. Darker candidates correspond to higher SFE values. Global optimizer is marked with a black star. Black numbers in red disks represent the iteration of acquisitions by three different strategies, among which only NRFS reaches the global optimizer at iteration 27 while EI and PES fails to identify it within 40 iterations.