REWARD LEARNING AS DOUBLY NONPARAMETRIC BANDITS: OPTIMAL DESIGN AND SCALING LAWS

Anonymous authors
Paper under double-blind review

ABSTRACT

Specifying reward functions for complex tasks like object manipulation or driving is challenging to do by hand. Reward learning seeks to address this by learning a reward model using human feedback on selected query policies. This shifts the burden of reward specification to the optimal design of the queries. We propose a theoretical framework for studying reward learning and the associated optimal experiment design problem. Our framework models rewards and policies as nonparametric functions belonging to subsets of Reproducing Kernel Hilbert Spaces (RKHSs). The learner receives (noisy) oracle access to a true reward and must output a policy that performs well under the true reward. For this setting, we first derive non-asymptotic excess risk bounds for a simple plug-in estimator based on ridge regression. We then solve the query design problem by optimizing these risk bounds with respect to the choice of query set and obtain a finite sample statistical rate, which depends primarily on the eigenvalue spectrum of a certain linear operator on the RKHSs. Despite the generality of these results, our bounds are stronger than previous bounds developed for more specialized problems. We specifically show that the well-studied problem of Gaussian process (GP) bandit optimization is a special case of our framework, and that our bounds either improve or are competitive with known regret guarantees for the Matérn kernel.

1 INTRODUCTION

Specifying the reward function accurately for a desired objective, or *reward engineering*, is challenging to perform by hand, as the consequences of even small errors can be drastic (Hadfield-Menell et al., 2017). To address this, reward learning seeks to learn a predictive model of the reward function from data, which is obtained from carefully selected queries to human annotators. The learned reward model is then used as the optimization objective for policy learning. Reward learning has achieved significant empirical success in domains such as text summarization (Stiennon et al., 2020; Böhm et al., 2019), robot locomotion (Daniel et al., 2014), predicting driving styles (Kuderer et al., 2015), and Atari game playing (Christiano et al., 2017).

Despite their success, reward learning methods still lack theoretical grounding. Moreover, their behavior can be brittle even on simple tasks, due to the difficulty of choosing appropriate queries and due to feedback loops from adaptive querying (Freire et al., 2020). Indeed, an ablation study in Christiano et al. (2017) suggests that random queries can outperform or be competitive with adaptive query procedures. To address these issues, we provide a theoretical framework for analyzing reward learning, framing it as a *doubly nonparametric experimental design* problem. This framework helps elucidate the role of query selection (Chaloner & Verdinelli, 1995) and also enables us to derive scaling laws—how the sizes of the policy and reward models affect the query complexity—for reward learning (Kaplan et al., 2020).

Proposed framework. In our framework, we suppose we are given a reward class C_r and policy class C_π. Our goal is to find a policy $\hat{\pi} \in C_\pi$ that performs well according to an unknown true reward $r^* \in C_r$. To do this, we query policies $\pi \in C_\pi$, observing noisy estimates of their true reward, and use this information to choose the eventual policy $\hat{\pi}$.

To be compatible with modern nonparametric learning methods (i.e. neural nets), we view C_r and C_π as subsets of Reproducing Kernel Hilbert Spaces (RKHS). The learner therefore optimizes a non-

Algorithm	Regret $\mathfrak{R}_T$	Non-vacuous regime
GP-UCB (Srinivas et al., 2010)	$\tilde{O}(T^{\frac{2\nu+d(3d+3)}{4\nu+d(2d+2)}})$	$\nu > \frac{d^2+d}{2}$
GP-TS (Chowdhury & Gopalan, 2017)	$\tilde{O}(T^{\frac{2\nu+d(3d+3)}{4\nu+d(2d+2)}})$	$\nu > \frac{d^2+d}{2}$
Our work	$\tilde{O}(T^{\frac{4\nu+d(4d+6)}{6\nu+d(4d+7)}})$	$\nu > \frac{3}{2}$
π-GP-UCB (Janz et al., 2020)	$\tilde{O}(T^{\frac{2\nu+d(2d+3)}{4\nu+d(2d+4)}})$	$\nu > 1$

Table 1. Our algorithm specializes to the case of kernel multi-armed bandits and yields strong bounds. For a d-dimensional Matérn kernel with smoothness ν, we outperform both GP-UCB and GP-TS unless $\nu \gtrsim d^2$. The only work to achieve better bounds for small ν is π-GP UCB, which was designed specifically for the Matérn kernel.

parametric reward function over a nonparametric space of policies, making the task "doubly" nonparametric. In contrast, previous work considers a nonparametric function class or reward class, but typically not both. For instance, nonparametric zeroth order or bandit optimization (Srinivas et al., 2010; Mockus, 2012; Wang et al., 2018) considers a nonparametric function on a *finite-dimensional* input space. Conversely, nonparametric supervised learning (Wahba, 1990; Hofmann et al., 2008) minimizes a *known* loss function over a nonparametric input space.

The doubly nonparametric nature of our task poses new challenges. The (possibly) infinite-dimensional RKHS requires the learner to select which subspace to explore given a finite number of queries. Furthermore, the unknown reward function makes it challenging for the learner to reason about the information gained from the selected query policies. We address these challenges by deriving a risk upper bound for a family of plug-in estimators based on ridge regression, and then optimizing this bound to solve the optimal design task.

In addition to the optimal design problem, our framework allows us to study scaling laws with respect to the reward (or policy) class by varying the rate of decay of their corresponding eigenspectrum. This decay rate determines the effective dimensionality of a RKHS (Zhang, 2002), and provides a natural proxy for varying the the size of the reward or policy class. Qualitatively, our main results show that the excess risk asymptotically vanishes as long as the policy class grows at a slower rate relative to the reward class.

Sharpness of analysis. Our risk bounds apply to reward and policy classes of arbitrary or even infinite dimensionality. Despite this generality, we show they provide stronger guarantees than previous known bounds for the specialized settings of compact policy sets and kernel multi-armed bandits.

In Section 4.3, we look at a special case of our problem when the policy set C_π is a compact subspace and thus has finite rank. For these instances, we show that our learning algorithm obtains a better excess risk $O(n^{-\frac{\beta}{\beta+2}})$ versus a rate of $O(n^{-\frac{\beta-1}{2(\beta+1)}})$ obtained by the adaptive GP-UCB algorithm (Srinivas et al., 2010), where $\beta > 0$ is a power law decay rate.

In Section 5, we specialize our general results to the well-studied problem of Gaussian process bandit optimization (Williams & Rasmussen, 2006), also known as kernel multi-armed bandit (MAB). Specifically, for the class of Matérn kernels with parameter ν in d dimensions, we show that our algorithm achieves a regret bound of $\tilde{O}(T^{\frac{4\nu+d(4d+6)}{6\nu+d(4d+7)}})$ which is strictly better than those achieved by the GP-UCB and GP-Thompson Sampling (GP-TS) (Chowdhury & Gopalan, 2017) algorithms and comparable with π-GP UCB (Janz et al., 2020); see Table 1 for details. GP-UCB and GP-TS are only yield sub-linear regret bounds when the smoothness of the kernel $\nu > d^2$—thus in high dimensions, these bounds essentially become vacuous. The π-GP UCB algorithm was designed specifically to overcome this issue. Our proposed algorithm achieves sublinear regret for all $\nu > 3/2$.

Our Contributions. We propose doubly-nonparameteric bandits as a framework for theoretically studying the reward learning problem. Within this framework, we obtain finite sample risk bounds for a ridge regression based plug-in estimator and derive scaling laws for reward learning. As a corollary of our risk bounds, we are able to provide sharper regret bounds for a class of kernel MAB problems compared to several existing algorithms, showing that the doubly-nonparametric lens of reward learning is fruitful even for "singly-nonparametric" tasks.

2 FRAMEWORK: DOUBLY NONPARAMETRIC BANDITS

Our framework considers non-parametric policy learning with non-parametric reward models. We let $\pi \in \mathbb{H}_\pi$ denote an arbitrary policy and $r \in \mathbb{H}_r$ denote an arbitrary reward function, where $\mathbb{H}_\pi$ and $\mathbb{H}_r$ are Reproducing Kernel Hilbert Spaces. For technical reasons, we assume the corresponding kernel functions $\mathcal{K}_\pi$ and $\mathcal{K}_r$ both satisfy the Hilbert-Schmidt condition (see Appendix A for details).

We let $F(\pi, r) \in \mathbb{R}$ denote the reward obtained by selecting policy π under reward function r and consider the case where the evaluation functional F is linear in both π and r. In other words, $F(\pi, r) = \langle r, M\pi \rangle_{\mathbb{H}_r}$ where $M : \mathbb{H}_\pi \mapsto \mathbb{H}_r$ is a known linear mapping from the policy space to the reward space. Since $\mathbb{H}_\pi$ and $\mathbb{H}_r$ may be infinite-dimensional, linearity is only a weak restriction– e.g. the map $f \mapsto f(x)$ is linear in f for any RKHS.

To incorporate problem structure, we let r^* denote the true reward function and assume that $r^* \in C_r$ for some known set $C_r \subseteq \mathbb{H}_r$ such that $\|r^*\|_{\mathbb{H}_r} = 1$. We further assume that policies π are restricted to lie in some C_π which is a subset of the unit ball in $\mathbb{H}_\pi$ (for instance, C_π might incorporate physical constraints on implementable policies). Thus, given the true reward r^*, the optimal policy is $\pi^* \in \operatorname{argmax}_{\pi \in C_\pi} F(\pi, r^*)$. This proposed framework, which allows for infinite-dimensional policy as well as reward classes, allows us to study how both the policy and reward space affect the difficulty of learning.

Query access to reward r^*. The true reward function r^* is unknown to the learner but is accessible via queries to an oracle (e.g. a human expert), which provide noisy zeroth-order evaluations of the reward r^*. When queried with a policy $\pi \in C_\pi$, the oracle provides a response

$$\text{Oracle } \mathcal{O}_{r^*} : \pi \mapsto F(\pi, r^*) + \epsilon \quad \text{where} \quad \epsilon \sim \mathcal{N}(0, \tau^2) \ , \tag{1}$$

with τ^2 denoting the variance of the response. There are two possible query models: passive queries (Atkinson, 1996; Sebastiani & Wynn, 2000), where the learner selects all queries at the same time, and active queries (Bubeck et al., 2011; Lattimore & Szepesvári, 2020), where the learner is allowed to select queries sequentially. Our focus in this work will be on the passive query model, but in many cases we will outperform existing active query algorithms.

Problem statement. Given passive access to the oracle $\mathcal{O}_{r^*}$, the objective of the learner is to output a policy $\hat{\pi} \in C_\pi$ that has small excess risk Δ, defined as

$$\Delta(\hat{\pi}; r^*) := F(\pi^*, r^*) - F(\hat{\pi}, r^*) \ . \tag{2}$$

We think of queries to the oracle as expensive, and are interested in achieving low excess risk with as few queries as possible.

Representations in $\ell_2(\mathbb{N})$. By Mercer's theorem, we can represent any RKHS as a subset of $\ell_2(\mathbb{N})$. Formally, the policy and the reward spaces are isomorphic to the ellipsoids

$$\mathbb{H}_\pi := \left\{ \sum_{j=1}^{\infty} \kappa_{\pi,j} \phi_{\pi,j} \,\middle|\, (\kappa_{\pi,j})_{j=1}^{\infty} \in \ell^2(\mathbb{N}) \text{ with } \sum_{j=1}^{\infty} \frac{\kappa_{\pi,j}^2}{\mu_{\pi,j}^2} < \infty \right\} \text{ and}$$

$$\mathbb{H}_r := \left\{ \sum_{j=1}^{\infty} \kappa_{r,j} \phi_{r,j} \,\middle|\, (\kappa_{r,j})_{j=1}^{\infty} \in \ell^2(\mathbb{N}) \text{ with } \sum_{j=1}^{\infty} \frac{\kappa_{r,j}^2}{\mu_{r,j}^2} < \infty \right\} \ ,$$

for appropriately chosen eigenfunctions $\phi_{\pi,j}$ and $\phi_{r,j}$ (Wainwright, 2019). These are defined with respect to a base measure $\mathbb{P}$ over the input domain; see Appendix A for details. With a slight abuse of notation, going forward, we will use π and r to denote the corresponding coefficients $(\kappa_{\pi,j})$ and $(\kappa_{r,j})$ in the expansion above. [1] With this, the inner products associated with $\mathbb{H}_\pi$ and $\mathbb{H}_r$ simplify

$$\langle \pi_1, \pi_2 \rangle_{\mathbb{H}_\pi} := \sum_{j=1}^{\infty} \frac{\pi_{1,j}\pi_{2,j}}{\mu_{\pi,j}} \quad \text{and} \quad \langle r_1, r_2 \rangle_{\mathbb{H}_r} := \sum_{j=1}^{\infty} \frac{r_{1,j} r_{2,j}}{\mu_{r,j}} \ . \tag{3}$$

Also let $S_r := \text{diag}(\mu_{r,j}^{-1})$ and $S_\pi := \text{diag}(\mu_{\pi,j}^{-1})$ be diagonal matrices comprising the inverse of the eigenvalues of $\mathbb{H}_r$ and $\mathbb{H}_\pi$. With this notation, if we view the map M as a (infinite-dimensional) matrix, its Hermitian adjoint[2] is equal to $M^* = S_\pi^{-1} M^\top S_r$.

[1] While the eigenfunctions ϕ_π and ϕ_r can be different, this representation can still be used by modifying the map M appropriately. This is detailed in Appendix A.

[2] Recall the Hermitian adjoint of M satisfies $\langle r, M\pi \rangle_{\mathbb{H}_r} = \langle M^* r, \pi \rangle_{\mathbb{H}_\pi}$

Algorithm 1: Policy Learning via Reward Learning

Input: Number of queries n, policy set C_π, oracle $\mathcal{O}_{r^*}$
Select n policies $\mathcal{Q} = \{\pi_1, \ldots, \pi_n\}$ and receive noisy reward evaluations $y_i = \mathcal{O}_{r^*}(\pi_i)$.
Estimate $\hat{r}$ using observed responses $\{(\pi_1, y_1), \ldots, (\pi_n, y_n)\}$ using ridge regression (5).
Obtain plug-in policy $\hat{\pi}_{\mathsf{plug}} \in \operatorname{argmax}_{\pi \in C_\pi} F(\pi, \hat{r})$.
Output: Policy $\hat{\pi}_{\mathsf{plug}}$

In order for the evaluation functional $F(\pi, r^*)$ to be finite for all $\pi \in \mathbb{H}_\pi$, the operator norm $\|S_r^{\frac{1}{2}} M S_\pi^{-\frac{1}{2}}\|_{\mathsf{op}}$ must be bounded (see Appendix A). We will see later that the decay of this operator's singular values is closely related to the difficulty of learning in our setting.

3 Algorithm: Policy Learning via Reward Learning

Given the setup above, we now describe a meta-algorithm, policy Learning via reward learning (Algorithm 1), for the non-parametric policy learning problem. The algorithm is a three-stage procedure: it (i) selects a subset of policies $\mathcal{Q}$ to query for reward feedback, (ii) uses the responses to learn a reward estimate $\hat{r}$, and (iii) optimizes this learnt estimate to output the policy $\hat{\pi}_{\mathsf{plug}}$, that is,

$$\hat{\pi}_{\mathsf{plug}} \in \operatorname*{argmin}_{\pi \in C_\pi} \langle \hat{r}, M\pi \rangle_{\mathbb{H}_r} \ . \tag{4}$$

Such general plug-in procedure have been studied in the statistics (Van der Vaart, 2000) and the machine learning (Devroye et al., 2013) literature. We analyze the excess risk of this estimator for our doubly-nonparametric setup and use this risk bound to select our query set $\mathcal{Q}$.

We now discuss the two key design choices in our algorithm: the choice of the reward estimation procedure as well as the choice of query set $\mathcal{Q}$.

Reward learning via ridge regression. We estimate the reward $\hat{r}$ via ridge regression in the RKHS $\mathbb{H}_r$ (Friedman et al., 2001; Shawe-Taylor et al., 2004). Suppose that in the first step of the algorithm, we have already queried the oracle on n policies and let $\{(\pi_i, y_i)\}_{i=1}^n$ represent the query-response pairs. For a regularization parameter $\lambda_{\mathsf{reg}} > 0$, the ridge regression estimate of the reward function is

$$\hat{r} \in \operatorname*{argmin}_{r \in \mathbb{H}_r} \frac{1}{n} \sum_{i=1}^n (y_i - \langle r, M\pi_i \rangle_{\mathbb{H}_r})^2 + \lambda_{\mathsf{reg}} \|r\|_{\mathbb{H}_r}^2 \ . \tag{5}$$

The parameter λ_{reg}, which is usually set as a function of n, controls the bias-variance trade-off in estimating r^*—smaller values of λ_{reg} reduce bias while larger values help reduce variance.

Excess risk bound for fixed query set. Observe that the plug-in estimator $\hat{\pi}_{\mathsf{plug}}(\mathcal{Q})$ is implicitly a function of the query set $\mathcal{Q}$. Ideally, we want to choose the set $\mathcal{Q}$ which minimizes the expected risk of the plugin estimator. This requires us to solve the optimization problem

$$\mathcal{Q} = \operatorname*{argmin}_{S : |S| \leq n} \mathbb{E}[\Delta(\hat{\pi}_{\mathsf{plug}}(S); r^*)] \ . \tag{6}$$

However, solving the above precisely requires knowledge about the underlying reward function r^*, and the combinatorial nature of the optimization problem makes it hard to find an exact solution. To address this, we first upper bound the excess risk of the plug-in policy $\hat{\pi}_{\mathsf{plug}}$ in terms of the query set $\mathcal{Q} = \{\pi_1, \ldots, \pi_n\}$. The following theorem[3] bounds the excess risk in terms of the spectrum of the spaces $\mathbb{H}_r$ and $\mathbb{H}_\pi$, as well as the covariance matrix of the queried policies $\Sigma_{\mathcal{Q}} := \frac{1}{n} \sum_{\pi \in \mathcal{Q}} \pi\pi^\top$.

Theorem 1 (Excess risk of plug-in)**.** *For any query set $\mathcal{Q}$ consisting of n policies and regularization parameter $\lambda_{\mathsf{reg}} > 0$, the excess risk of the plug-in estimator $\hat{\pi}_{\mathsf{plug}}$ is upper bounded as*

$$\mathbb{E}[\Delta(\hat{\pi}_{\mathsf{plug}}; r^*)] \leq 2\|M^*(r^* - \hat{r})\|_{\mathbb{H}_\pi} \ . \tag{7}$$

[3]Throughout the paper, for clarity purposes, we denote by c a universal constant whose value changes across lines. All our proofs in the appendices explicitly track this constant.

In addition, letting $A = M\Sigma_{\mathcal{Q}}M^\top S_r + \lambda_{\mathsf{reg}} I$, the expected squared distance is equal to

$$\mathbb{E}[\|M^*(r^* - \hat{r})\|^2_{\mathbb{H}_\pi}] = \lambda^2_{\mathsf{reg}} \cdot \|M^* A^{-1} r^*\|^2_{\mathbb{H}_\pi} + \frac{\tau^2}{n} \cdot \mathrm{tr}\left[S_\pi (M^* A^{-1} M)\Sigma_{\mathcal{Q}}(M^* A^{-1} M)^\top\right]. \quad (8)$$

The proof follows a standard analysis of ridge regression and is deferred to Appendix B. Observe that in the above theorem, the query set $\pi \in \mathcal{Q}$ participates in the excess risk only via the covariance $\Sigma_{\mathcal{Q}}$. The risk bound is the sum of two term: the first corresponding to the bias and the second corresponding to the variance. In both these terms, $\Sigma_{\mathcal{Q}}$ appears as part of A^{-1}—thus query sets $\mathcal{Q}$ which induce a larger correlation with the map M will generally have lower excess risk. Choices of queries which are orthogonal to the right singular vectors of M will have a constant excess risk, since for those directions the matrix $A \approx \lambda_{\mathsf{reg}} I$.

As shown later in the appendix, in the special case when the policy set consists of the entire unit ball $C_\pi = \{\pi \in \mathbb{H}_\pi \mid \|\pi\|_{\mathbb{H}_\pi} \leq 1\}$, the excess risk bound can be improved by a quadratic factor: $\mathbb{E}[\Delta(\hat{\pi}_{\mathsf{plug}}; r^*)] \leq O\left(\|M^*(r^* - \hat{r})\|^2_{\mathbb{H}_\pi}\right)$.

4 Query selection and statistical guarantees

We now show how to select the query set $\mathcal{Q}$ effectively and study the excess risk of the corresponding plug-in estimator $\hat{\pi}_{\mathsf{plug}}$. We will start with the special case where the policy set C_π is the unit ball in $\mathbb{H}_\pi$ and the map M is diagonal, and then generalize to arbitrary policy sets. In both cases, low excess risk can be achieved by repeatedly querying (approximations of) the projections of top eigenvectors of M^*M onto the $\mathbb{H}_\pi$ space. For the special case when the map M is diagonal, this is reduces to querying the top eigenvectors of $\mathbb{H}_\pi$.

The excess risk will ultimately depend on the the eigenspectrum of the operator $S_\pi^{-\frac{1}{2}} M^\top S_r M S_\pi^{-\frac{1}{2}}$, which is similar to the operator M^*M. To interpret our results, we instantiate them for a power law spectrum with exponent $\beta > 0$, that is,

$$\sigma_j(S_\pi^{-\frac{1}{2}} M^\top S_r M S_\pi^{-\frac{1}{2}}) \asymp j^{-\beta}. \quad (9)$$

Such power law spectra have been observed in a variety of practical settings, for instance, in the Hessian of trained deep neural networks (Ghorbani et al., 2019).

4.1 Warm-up: C_π = unit ball, M = diagonal

In order to get some intuition, we study the special case where the policy set C_π consists of the entire unit ball in the space $\mathbb{H}_\pi$ and the map M is diagonal with $M = \mathrm{diag}(\nu_j)$. Further, let us denote the operator $\tilde{M} = S_r^{1/2} M S_\pi^{-1/2}$.

For this special case, our sampling algorithm (Algorithm 2) simply selects the top J eigenvectors of the space $\mathbb{H}_\pi$ to query. To see why, observe that for a diagonal map M, the right singular vectors of the operator $\tilde{M}$ are the same as the eigenvectors of the policy space $\mathbb{H}_\pi$. Therefore, the choice of policy π_j in our algorithm is simply the scaled eigenfunction $\sqrt{\mu_{\pi,j}} \cdot \phi_{\pi,j}$. Having selected these J queries, the algorithm queries each one of the $\frac{n}{J}$ times and uses this as query set $\mathcal{Q}$.

The intuition for this choice of query set $\mathcal{Q}$ is that since we are in the passive setup with no knowledge of r^*, any policy $\pi \in C_\pi$ can be an optimal policy. By querying the top J ones out of these, we can obtain a good enough approximation to the performance of any policy in the unit ball. The particular choice of the parameter J depends on the number of queries n available. Since the oracle responses are noisy, to reduce variance in the responses along those directions, our algorithm performs multiple queries along the same direction. This simple query strategy enjoys the following excess risk bound.

Proposition 1 (Risk bound for C_π = unit ball.)**.** *For any $J \leq n$ and regularization parameter $\lambda_{\mathsf{reg}} > 0$, consider the plug-in estimator obtained via the passive sampling algorithm which explores the first J eigenfunctions of $\mathbb{H}_\pi$. The excess risk satisfies*

$$\mathbb{E}[\Delta(\hat{\pi}_{\mathsf{plug}}; r^*)] \leq c \cdot \left(1 + \frac{\tau^2}{n\lambda^2_{\mathsf{reg}}}\right) \cdot \max\left\{\sup_{j \leq J} \frac{\lambda^2_{\mathsf{reg}} J^2 \zeta_j}{\zeta_j^2 + \lambda^2_{\mathsf{reg}} J^2}, \sup_{j > J} \zeta_j\right\}, \quad (10)$$

Algorithm 2: Passive querying strategy

Input: Number of queries n, map M, policy set C_π, exploration parameter J

Construct linear map $\tilde{M} = S_r^{\frac{1}{2}} M S_\pi^{-\frac{1}{2}}$ and compute eigenvectors $\{\phi_{\tilde{M},j}\}_j$ of operator $\tilde{M}^\top \tilde{M}$

Set policy $\pi_j = \Phi_\pi S_\pi^{-\frac{1}{2}} \Phi_\pi^\top \phi_{\tilde{M},j}$ for all $j \leq J$

Obtain policy $\tilde{\pi}_j \in C_\pi$ such that $\tilde{\pi}_j \tilde{\pi}_j^\top \succeq c_\pi \pi_j \pi_j^\top$

Form query set $\mathcal{Q} = \{\tilde{\pi}_1^{(n/J)}, \ldots, \tilde{\pi}_{n^\alpha}^{(n/J)}\}$ where $a^{(b)} = \{a, \ldots, a\}\}$ repeated b times

Output: Query set $\mathcal{Q}$

where the quantity $\zeta_j = \frac{\nu_j^2 \mu_{\pi,j}}{\mu_{r,j}}$ *for some universal constant* $c > 0$.

We defer the proof of the above proposition to Appendix B. The choice of the exploration parameter J allows us to trade-off between the two terms inside the maximum. Typically, the second term will be maximized at $j = J + 1$. For the first term, the supremum depends on the choice of λ_{reg} — for small values of λ_{reg}, the sup is achieved at $j = 1$ while for larger values, it is achieved at $j = J$. In order to gain more intuition about this bound, we instantiate this for the power law decay.

Corollary 1 (Risk bound for power-law decay)**.** *Suppose that eigenvalues of the police space* $\mathbb{H}_\pi$ *decay as* $j^{-\beta_\pi}$*, reward space* $\mathbb{H}_r$ *as* $j^{-\beta_r}$ *and the singular values of map* M *as* $j^{-\beta_M}$*. This satisfies the power law assumption with exponent* $\beta = \beta_\pi + \beta_M - \beta_r$*. The plug-in estimator with exploration parameter* $J = n^{\frac{1}{\beta+2}}$ *and regularization* $\lambda_{\mathsf{reg}} = n^{-\frac{\beta+1}{\beta+2}}$ *satisfies*

$$\mathbb{E}[\Delta(\hat{\pi}_{\mathsf{plug}}; r^*)] \leq c n^{-\frac{\beta}{\beta+2}} \ . \tag{11}$$

The above bound shows that our algorithm can learn in the framework as long as $\beta > 0$ or equivalently $\beta_\pi + \beta_M > \beta_r$, with better rates for larger values of β. Thus, for a fixed size of reward class β_r, the learning rate improves as the policy class grows smaller (β_π increases) – this is intuitive since we are required to search over a smaller policy space. On the other hand, for a fixed policy class β_π, our excess risk rate gets better as the reward class grows in size (β_r increases) – this is because a larger set of reward functions have similar optimal policies and hence learning gets easier.

4.2 General policy sets

We now describe our choice of query sets $\mathcal{Q}$ for general policy sets C_π. Our strategy, described in Algorithm 2, differs from the above special case in that we need to take into account the interaction of the policy space $\mathbb{H}_\pi$ with the map M. Specifically, we show in Appendix B that the upper bound in Theorem 1 can be diagonalized for this general case via a transformation.

Let us denote the operator $\tilde{M} = S_r^{1/2} M S_\pi^{-1/2}$. Our transformation reveals that the relevant directions to query for this general case corresponds to the columns of $\Phi_\pi S_\pi^{-1/2} \Phi_\pi^\top V_M$ where , then V_M are the eigenvectors of the self-adjoint operator $\tilde{M}^\top \tilde{M}$ – and it is precisely a subset of these directions that our algorithm queries.

In order to be able to query these policies, we require the set C_π to contain some policies which align well with them. We formally state this regularity assumption below.

Assumption 1 (Regularity assumption on C_π)**.** *For any eigenfunction* $\phi_{\tilde{M},j}$ *of the operator* $\tilde{M}^\top \tilde{M}$*, consider the policy* $\pi_j = \Phi_\pi S_\pi^{-1/2} \Phi_\pi^\top \phi_{\tilde{M},j}$*. There exists a policy* $\tilde{\pi}_j$ *in policy set* C_π *such that* $\tilde{\pi}_j \tilde{\pi}_j^\top \succeq c_\pi \pi_j \pi_j^\top$*, for some constant* c_π*.*

The above assumption requires that for every choice of the policy π_j in Algorithm 2, the set C_π has the another policy $\tilde{\pi}_j$ which is collinear with it. This assumption can be relaxed in various ways (for instance via convexification) but we omit this as it is not needed for our results. Given this assumption, the following theorem, a generalization of Proposition 1, provides a bound on the excess risk for the plug-in estimate for general policy sets C_π.

Theorem 2 (Risk bound for general policy sets C_π.)**.** *For any* $J \leq n$*, regularization parameter* $\lambda_{\mathsf{reg}} > 0$ *and set* C_π *satisfying Assumption 1, let* $\hat{\pi}_{\mathsf{plug}}$ *be the estimator output by Algorithm 1. The*

squared excess risk satisfies

$$(\mathbb{E}[\Delta(\hat{\pi}_{\mathsf{plug}}; r^*)])^2 \leq c \cdot \left(1 + \frac{\tau^2}{n\lambda_{\mathsf{reg}}^2}\right) \cdot \max\left\{\sup_{j \leq J} \frac{\lambda_{\mathsf{reg}}^2 J^2 \zeta_j}{\zeta_j^2 + \lambda_{\mathsf{reg}}^2 J^2}, \sup_{j > J} \zeta_j\right\} ,$$

where the values ζ_j correspond to the jth eigen values of the operator $\tilde{M}^*\tilde{M}$ with $\tilde{M} = S_r^{\frac{1}{2}} M S_\pi^{\frac{1}{2}}$.

We defer the proof of this theorem to Appendix B. The proof of this theorem goes via a transformation which diagonalizes the excess risk bound and reduces the problem to a similar setup as that of Proposition 1. Additionally, Assumption 1 allows us to generalize the results to arbitrary policy sets C_π. Note that the above upper bounds the square of the excess risk. As discussed in Section 3, one can obtain a quadratic improvement in this rate if the set C_π is the entire unit ball in $\mathbb{H}_\pi$. We specialize the above bound for the power law decay assumption in the following corollary.

Corollary 2 (Risk bound for power-law decay)**.** *Suppose that eigenspectrum of the operator $S_\pi^{-\frac{1}{2}} M^\top S_r M S_\pi^{-\frac{1}{2}}$ satisfy the power law assumption with exponent $\beta > 0$, that is, $\sigma_j(S_\pi^{-\frac{1}{2}} M^\top S_r M S_\pi^{-\frac{1}{2}}) \asymp j^{-\beta}$. The plug-in estimator $\hat{\pi}_{\mathsf{plug}}$ with parameter $J = n^{\frac{1}{\beta+2}}$ and regularization $\lambda_{\mathsf{reg}} = n^{-\frac{\beta+1}{\beta+2}}$ satisfies*

$$\mathbb{E}[\Delta(\hat{\pi}_{\mathsf{plug}}; r^*)] \leq c n^{-\frac{\beta}{2(\beta+2)}} . \tag{12}$$

The above bound indicates that for the general case, learning is possible if the spectrum decay has parameter $\beta > 0$. To get such a spectrum decay with the operator defined in the above corollary, one sufficient condition is that the map M does not flip the larger eigenvectors of $\mathbb{H}_\pi$ towards the smaller eigenvectors of $\mathbb{H}_r$ – such a misaligned scenario would require learning a very accurate representation of the reward to learn a good policy and will make learning harder.

4.3 Comparison with UCB-style adaptive algorithms

We next turn to evaluating the sharpness of Theorem 2. Existing frameworks for studying "singly"-nonparametric setups require the input domain to be compact. In our doubly-nonparametric setup, the input space is the policy set C_π which is often non-compact (i.e. the unit ball is not compact in infinite dimensions). We address this for singly-nonparametrics algorithm by taking a finite-dimensional approximation.

Even though our proposed method is passive, it achieves better rates than well-known *adaptive* sampling algorithms. Specifically, in the power law setting of Section 4.1, the analysis of GP-UCB algorithm (Srinivas et al., 2010) provides a rate of $O(n^{-\frac{\beta-1}{2(\beta+1)}})$, which is strictly worse than the $O(n^{-\frac{\beta}{\beta+1}})$ obtained by our analysis in Corollary 1.

We refer the reader to Proposition 2 in Appendix D for an exact statement. The proof adapts the analysis from Srinivas et al. (2010), which hinges on a quantity called the information gain, which we bound for our setup. While we are comparing upper bounds for the two algorithms, we believe that our improved bound is due to a better algorithm and not an analysis gap. As we show in Section 5, our proposed passive sampling strategy can even be competitive in the no-regret kernel multi-armed bandit setups for which the GP-UCB algorithm was been designed.

5 New bounds for kernel multi-armed bandits

In the previous subsection, we saw that our passive sampling algorithm actually outperforms existing adaptive sampling algorithms for the reward learning task we care about. Here we take this a step further—we specialize our algorithm to the case of kernel MABs, and show that it outperforms standard algorithms for that setting and is competitive with a specialized algorithm for Matérn kernels.

We consider the task of maximizing an unknown function $f^* : \mathcal{X} \mapsto \mathbb{R}$ over its domain $\mathcal{X} \subset \mathbb{R}^d$. In the kernel multi-armed bandit (MAB) setup, this unknown function f belongs to an RKHS $\mathbb{H}$, equipped with a positive-definite kernel[4] $\mathcal{K}$, such that $\|f^*\|_{\mathbb{H}} = 1$. Let us further restrict our attention

[4]We require that the kernel $\mathcal{K}$ be a Mercer's kernel satisfying $\mathcal{K}(x, x) = c$ for all $x \in \mathcal{X}$.

to the space of input points $\mathcal{X} = \{x \in \mathbb{R}^d \mid \|x\|_2 \leq 1\}$. The learner is allowed to access this function via a noisy zeroth-order oracle

$$\mathcal{O}_{f^*} : x \mapsto f^*(x) + \eta \quad \text{where } \eta \sim \mathcal{N}(0, \tau^2) \ . \tag{13}$$

Going forward we will assume that $\tau = 1$. The above oracle is similar to the reward oracle $\mathcal{O}_{r^*}$, except that the query points x belong to a finite dimensional space and f^* is a non-linear function of the query point x. The goal in MAB is to minimize the T-step regret $\mathfrak{R}_T := T \cdot \max_{x \in \mathcal{X}} f^*(x) - \sum_{t=1}^{T} f^*(x_t)$, where x_t is the datapoint queried in the tth round. There have been several algorithms proposed to solve this problem including general purpose UCB algorithms (Srinivas et al., 2010; Chowdhury & Gopalan, 2017), Thompson sampling approaches (Chowdhury & Gopalan, 2017), and special-purpose algorithms for specific kernels (Janz et al., 2020).

We next show that kernel MAB can be cast as a special case of our non-parametric policy learning framework. The resulting regret bounds, derived from an application of Theorem 3, are better than several general purpose algorithms (GP-UCB, IGP-UCB, GP-TS) and comparable to those specialized for the Matérn kernel (π-GP-UCB).

In order to reduce kernel MAB to our framework, we need to introduce three elements – the policy space $\mathbb{H}_\pi$, the reward space $\mathbb{H}_r$ and the map M. We would like spaces $\mathbb{H}_r$ and $\mathbb{H}_\pi$ such that (1) the resulting objective $F(r, \pi)$ is linear in this space, (2) the resulting rewards and policies have unit norm in their respective space, and (3) we have a good understanding of the eigenvalues of the resulting operator. This last point ensures that we can employ our upper bounds from Section 4.

Before we define these, we let $N_{\mathsf{cov}}(\epsilon)$ denote an ϵ-net of the input space $\mathcal{X}$ under the ℓ_2 norm and denote its elements by $\mathcal{C}_\epsilon$. We define the kernel matrix $K \in \mathbb{R}^{N_{\mathsf{cov}} \times N_{\mathsf{cov}}}$ on points selected in the cover as $K(i, j) = \mathcal{K}(x_i, x_j)$ for all $(x_i, x_j) \in \mathcal{C}_\epsilon \times \mathcal{C}_\epsilon$.

Reward space $\mathbb{H}_r$. Given the RKHS $\mathbb{H}$ as well as the elements of the cover $\mathcal{C}_\epsilon$, we view the reward function as a map from $\mathcal{C}_\epsilon$ to $\mathbb{R}$, or equivalently as a vector in $\mathbb{R}^{N_{\mathsf{cov}}(\epsilon)}$. More precisely, letting $\tilde{f} = [f(x_1), \ldots, f(x_{N_{\mathsf{cov}}})]$ denote the vector of evaluations of a function f, we define

$$\mathbb{H}_r := \text{span}\{\tilde{f} \mid f \in \mathbb{H}\} \quad \text{with} \quad \langle \tilde{f}_1, \tilde{f}_2 \rangle_{\mathbb{H}_r} := \tilde{f}_1^\top K^{-1} \tilde{f}_2, \ . \tag{14}$$

With this notation, we define the true reward $r^* := \tilde{f}^* = [f^*(x_1), \ldots, f^*(x_{N_{\mathsf{cov}}})]$.

Policy Space $\mathbb{H}_\pi$. Similarly to rewards, we will embed policies in $\mathbb{R}^{N_{\mathsf{cov}}}$. For any point $x \in \mathcal{C}_\epsilon$, let $k_x = [\mathcal{K}(x, x_1), \ldots, \mathcal{K}(x, x_{N_{\mathsf{cov}}})]$ denote the corresponding vector in $\mathbb{R}^{N_{\mathsf{cov}}}$ obtained by evaluating the kernel $\mathcal{K}$ over the cover. Then, the space

$$\mathbb{H}_\pi := \text{span}\{k_x \mid x \in \mathcal{C}_\epsilon\} \quad \text{with} \quad \langle k_1, k_2 \rangle_{\mathbb{H}_\pi} := \langle k_1, K^{-2} k_2 \rangle \ . \tag{15}$$

The choice of the above norm ensures that $\langle k_i, k_j \rangle_{\mathbb{H}_\pi} = \langle K^{-1} k_i, K^{-1} k_j \rangle = \delta_{i,j}$ for all $(x_i, x_j) \in \mathcal{C}_\epsilon \times \mathcal{C}_\epsilon$. Thus in particular, $\mathbb{H}_\pi$ contains an orthonormal embedding of the set of vectors $\{k_x\}_{x \in \mathcal{C}_\epsilon}$.

Map M. Both the reward space $\mathbb{H}_r$ and policy space $\mathbb{H}_\pi$ can be associated with $\mathbb{R}^{N_{\mathsf{cov}}}$. Under this transformation, the evaluation $f^*(x)$ for any $x \in \mathcal{C}_\epsilon$ corresponds to the standard inner product with $F(r^*, \pi_x) = f^*(x) = (\tilde{f}^*)^\top K^{-1} k_x = \langle r^*, k_x \rangle_{\mathbb{H}_r}$. This indicates that we should take the map M to be the identity. Furthermore, as a simple application of Mercer's theorem it follows that this map M is a bounded linear operator.

We make an additional assumption on the kernel function $\mathcal{K}$, requiring it to be Lipschitz in its input arguments. This assumption is often satisfied, in particular for the Matérn kernel when $\nu > 3/2$.

Assumption 2 (Lipschitz Kernel $\mathcal{K}$)**.** *The Kernel $\mathcal{K}$ associated with the Hilbert space $\mathbb{H}$ is $L_\mathcal{K}$-Lipschitz with respect to the ℓ_2-norm for some $L_\mathcal{K} > 0$:*

$$|\mathcal{K}(x, y) - \mathcal{K}(x, x)| \leq L_\mathcal{K} \|x - y\|_2 \quad \textit{for all } x \in \mathcal{X}, y \in \mathcal{X} \ .$$

Furthermore, the kernel also satisfies $\mathcal{K}(x, x) = 1$ for all points $x \in \mathcal{X}$.

Applying Theorem 2 under the above assumption, we obtain the following excess risk bound for the plug-in estimator evaluated on the unknown function f^*.

Theorem 3 (Excess risk for Kernel MAB)**.** *Suppose that the eigenvalues of a $L_\mathcal{K}$-Lipschitz kernel $\mathcal{K}$ satisfy the power-law decay $\mu_j \asymp j^{-\beta}$. Let $\hat{x}_{\mathsf{plug}}$ be the output of Algorithm 1 using n queries to*

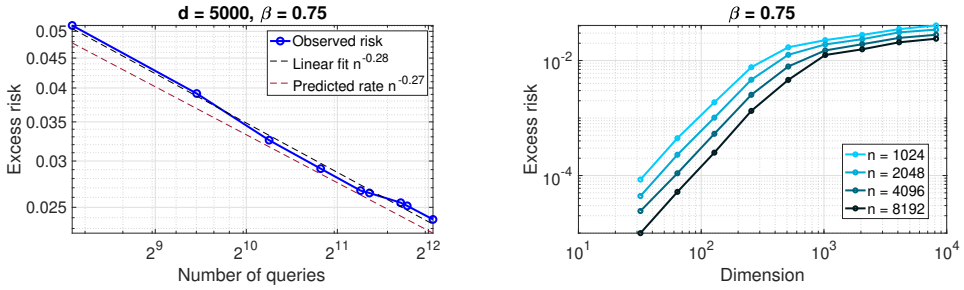

(a) (b)

Figure 1. (a) Corroborating upper bound from Corollary 1. Our theoretical bounds predict a rate of $n^{-0.27}$ and the experiment shows an almost matching rate of $n^{-0.28}$. (b) As the dimension d is increased, the excess risk curves asymptote at different levels for different n. This shows that our algorithm achieves non-vacuous error for the doubly-nonparametric set in the regime $d \to \infty$.

the oracle $\mathcal{O}_{f^*}$. Then, for any value of $\beta > 1 + \frac{2}{d} + \log(\frac{1}{\delta})$ and $\epsilon \in (0,1)$, the excess risk satisfies

$$\max_{x:\|x\|_2 \le 1} f^*(x) - f^*(\hat{x}_{\mathsf{plug}}) \lesssim N_{\mathsf{cov}}^{\frac{1}{\beta+2}}(\epsilon) \cdot n^{\frac{-\beta}{2(\beta+2)}} + N_{\mathsf{cov}}^{\frac{1-\beta}{2}}(\epsilon) + \sqrt{L_{\mathcal{K}}\epsilon}\,,$$

with probability at least $1-\delta$.

For Matérn kernels, it is known that the eigenvalues decay with parameter $\beta = 1 + \frac{2\nu}{d}$ (Janz et al., 2020). Using this, we can obtain the following corollary.

Corollary 3 (Regret bound for Matérn Kernel)**.** *Consider the family of Matérn kernels with parameter $\nu > \frac{3}{2}$ defined with the Euclidean norm over $\mathbb{R}^d$. The T-step regret of our algorithm is*

$$\mathfrak{R}_{\mathsf{mat},T} = \tilde{O}\left(T^{\frac{4\nu+d(6+4d)}{6\nu+d(7+4d)}}\right).$$

The above bound is for regret, which is an online notion, while our previous results are offline notions. We get from one to the other using a standard batch-to-online conversion bound based on an explore-then-commit strategy. Table 1 compares the above bound to the existing bounds in the literature. While the bounds for GP-UCB and GP-TS become vacuous for $\nu \lesssim d^2$, our bound from Corollary 3 is always sublinear in T. Our bound is slightly weaker than the corresponding regret bound obtained by π-GP UCB algorithm of Janz et al. (2020), which is analyzed specifically for the family of Matérn kernels.

6 Experimental Evaluation

In this section, we experimentally evaluate our proposed algorithm via a simulation study. We use these experiments to establish the dimension free nature of our results as well as to conjecture the optimality of our derived bounds.

Setup. In the simulation study, we work with d dimensional RKHSs $\mathbb{H}_r$ and $\mathbb{H}_\pi$. In order to simulate the nonparmeteric regime, we typically use value of n which are less or at most a constant times the dimension d. We set the matrices $S_\pi = \mathrm{diag}(j^{-1.75})$, $S_r = \mathrm{diag}(j^{-1})$ and the map $M = I$. With this, the effective decay parameter $\beta = \beta_\pi - \beta_r = 0.75$. We sampled the true reward r^* uniformly at random from the unit ball in $\mathbb{H}_r$. We further sampled the oracle noise $\epsilon \sim \mathcal{N}(0, 0.01)$. All plots were averaged over 10 runs.

Observations. Figure 1(a) shows the variation of excess risk as the number of queries n are varied from 256 to 4096 on a log-log plot. Our bounds in Corollary 1 for this setup predict that the excess risk should decay at a rate $O(n^{-0.27})$. By fitting a linear line through the plot, we found that observed risk to vary as $O(n^{-0.28})$. This plot is suggestive of the fact that our theoretical upper bounds might be tight in a minimax way over choices of decay parameter β.

In Figure 1(b), we plot the excess risk as we vary the dimension d from 32 to 8192 for four different choices of sample size, again, on a log-log scale. Increasing the number of queries decreases the excess risk for all dimensions consistently. The risk curves tend to asymptote at different error levels for different values of n. This corroborates our theoretical findings that our proposed algorithm provides non-vacuous bounds for the doubly-nonparametric setup in the regime $d \to \infty$.

ETHICS STATEMENT

Our main contribution is a theoretical framework to study reward learning and the associated optimal design problem. Since our contributions are mostly theoretical in nature, we do not anticipate any ethical issues arising in the near future.

REPRODUCIBILITY STATEMENT

On the theoretical side, we provide detailed proofs for all our results in the appendix and appropriately reference the intermediate results we might have used in the proofs. For the experimental aspect, we have attached our matlab code as a supplementary file and have provided all necessary hyper parameters details required to reproduce the experiments.

REFERENCES

Anthony C Atkinson. The usefulness of optimum experimental designs. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58, 1996.

Florian Böhm, Yang Gao, Christian M Meyer, Ori Shapira, Ido Dagan, and Iryna Gurevych. Better rewards yield better summaries: Learning to summarise without references. *arXiv preprint arXiv:1909.01214*, 2019.

Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12(5), 2011.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pp. 273–304, 1995.

Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, 2017.

Paul Christiano, Jan Leike, Tom B Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *arXiv preprint arXiv:1706.03741*, 2017.

Christian Daniel, Malte Viering, Jan Metz, Oliver Kroemer, and Jan Peters. Active reward learning. In *Robotics: Science and systems*, volume 98, 2014.

Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.

Pedro Freire, Adam Gleave, Sam Toyer, and Stuart Russell. In *Proceedings of the Workshop on Deep Reinforcement Learning at NeurIPS*, 2020.

Jerome Friedman, Trevor Hastie, Robert Tibshirani, et al. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.

Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via hessian eigenvalue density. In *International Conference on Machine Learning*, pp. 2232–2241. PMLR, 2019.

Dylan Hadfield-Menell, Smitha Milli, Pieter Abbeel, Stuart Russell, and Anca Dragan. Inverse reward design. *arXiv preprint arXiv:1711.02827*, 2017.

Thomas Hofmann, Bernhard Schölkopf, and Alexander J Smola. Kernel methods in machine learning. *The annals of statistics*, 36(3):1171–1220, 2008.

David Janz, David Burt, and Javier González. Bandit optimisation of functions in the matérn kernel rkhs. In *International Conference on Artificial Intelligence and Statistics*, 2020.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

Vladimir Koltchinskii and Karim Lounici. Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 23, 2017.

Erwin Kreyszig. *Introductory functional analysis with applications*, volume 1. wiley New York, 1978.

Markus Kuderer, Shilpa Gulati, and Wolfram Burgard. Learning driving styles for autonomous vehicles from demonstration. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2641–2646. IEEE, 2015.

Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

J. Mercer. Functions of positive and negative type, and their connection with the theory of integral equations. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 209, 1909.

Jonas Mockus. *Bayesian approach to global optimization: theory and applications*, volume 37. Springer Science & Business Media, 2012.

Paola Sebastiani and Henry P Wynn. Maximum entropy sampling and optimal bayesian experimental design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62, 2000.

John Shawe-Taylor, Nello Cristianini, et al. *Kernel methods for pattern analysis*. Cambridge university press, 2004.

Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias W Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *International Conference on Machine Learning*, 2010.

Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. Learning to summarize from human feedback. *arXiv preprint arXiv:2009.01325*, 2020.

Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

Grace Wahba. *Spline models for observational data*. SIAM, 1990.

Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

Yining Wang, Simon Du, Sivaraman Balakrishnan, and Aarti Singh. Stochastic zeroth-order optimization in high dimensions. In *International Conference on Artificial Intelligence and Statistics*, pp. 1356–1365. PMLR, 2018.

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

Tong Zhang. Effective dimension and generalization of kernel learning. In *NIPs*, volume 4, pp. 454–461. Citeseer, 2002.