
Understanding the Eluder Dimension

Gene Li

Toyota Technological
Institute at Chicago
gene@ttic.edu

Pritish Kamath

Google Research
pritish@alum.mit.edu

Dylan J. Foster

Microsoft Research
dylanfoster@microsoft.com

Nathan Srebro

Toyota Technological
Institute at Chicago
nati@ttic.edu

Abstract

We provide new insights on eluder dimension, a complexity measure that has been extensively used to bound the regret of algorithms for online bandits and reinforcement learning with function approximation. First, we study the relationship between the eluder dimension for a function class and a generalized notion of *rank*, defined for any monotone “activation” $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, which corresponds to the minimal dimension required to represent the class as a generalized linear model. It is known that when σ has derivatives bounded away from 0, σ -rank gives rise to an upper bound on eluder dimension for any function class; we show however that eluder dimension can be exponentially smaller than σ -rank. We also show that the condition on the derivative is necessary; namely, when σ is the relu activation, the eluder dimension can be exponentially larger than σ -rank. For binary-valued function classes, we obtain a characterization of the eluder dimension in terms of star number and threshold dimension, quantities which are relevant in active learning and online learning respectively.

1 Introduction

Russo and Van Roy [37] introduced the notion of **eluder dimension** for a function class and used it to analyze algorithms (based on the *Upper Confidence Bound (UCB)* and *Thompson Sampling* paradigms) for the multi-armed bandit problem with function approximation. Since then, eluder dimension has been extensively used to construct and analyze the regret of algorithms for contextual bandits and reinforcement learning (RL) with function approximation [see, e.g., 46, 36, 44, 6, 14, 19, 12, 16, 26, 24]. Even though the eluder dimension has become a central technique for reinforcement learning theory, little is known about when exactly it is bounded. This paper makes progress toward filling this gap in our knowledge.

Russo and Van Roy [37] established upper bounds on eluder dimension for (i) function classes for which inputs have finite cardinality (the “tabular” setting), (ii) linear functions over \mathbb{R}^d of bounded norm, and (iii) generalized linear functions over \mathbb{R}^d of bounded norm, with any activation that has derivatives bounded away from 0. Apart from these function classes (and those that can be embedded into these), understanding of eluder dimension has been limited. Indeed, one might wonder whether a function class has “small” eluder dimension only if it can be realized as a class of (generalized) linear functions! This leads us to our first motivating question:

Question 1. *Are all function classes with small eluder dimension essentially generalized linear models?*

Answering this question has substantial ramifications on the scope of prior work. An answer of “yes” would imply that the results in the aforementioned work which gives regret guarantees in terms of eluder dimension do not go beyond already-established regret guarantees for generalized linear bandit or RL settings [see, e.g. 17, 32, 45, 31]. An answer of “no” can be construed as a positive result for RL theory, as it would indicate that existing (and future) results which use the eluder dimension apply to a richer set of function classes than generalized linear models.

To answer [Question 1](#), we first formally define what it means for a function class to be written as a generalized linear model (GLM). Informally, for an activation $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ and a function class $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \mathbb{R})$, we define the **σ -rank** to be the smallest dimension d needed to express every function in \mathcal{F} as a generalized linear function in \mathbb{R}^d with activation σ (see [Definition 3](#)). Intuitively, the σ -rank captures the best possible upper bound on eluder dimension that the results from Russo and Van Roy [37] can give for a given \mathcal{F} by treating it as a GLM with activation σ . We ask how the eluder dimension of any function class relates to its σ -rank for various activations σ . We show that the answer to [Question 1](#) is indeed “no”, i.e., there exists a function class with eluder dimension d but the σ -rank for any monotone σ is at least $\exp(\Omega(d))$. Thus, while Russo and Van Roy [37] show that the set of function classes with small eluder dimension is a *superset* of the set of GLM function classes, we (roughly speaking) show that the set of function classes with small eluder dimension is *strictly larger* than the set of GLM function classes.

We also prove that the requirement from Russo and Van Roy [37] that the derivative of the activation is bounded away from 0 is necessary in order to bound the eluder dimension of GLMs. The upper bound in the paper [37] becomes vacuous when the activation function has zero derivative; we show a lower bound which indicates this requirement cannot be dropped. Namely, when σ is the relu activation, we show that eluder dimension can be *exponentially larger* than σ -rank.

In a second line of inquiry, we study a *combinatorial* version of eluder dimension. The original definition of Russo and Van Roy [37] is defined for real-valued function classes, but one can specialize the definition to binary-valued function classes, leading to a so-called **combinatorial eluder dimension**. Thus, our second motivating question is:

Question 2. *Can we bound the combinatorial eluder dimension, perhaps in terms of more familiar learning-theoretic quantities?*

One might wonder: if the combinatorial eluder dimension is just a special case of the scale-sensitive version, why study it at all? Our reasons are threefold.

- The first and most immediate reason is that we are able to show new characterizations of eluder dimension once we move to the combinatorial definition. We elucidate a fundamental connection between eluder dimension and two other well-studied learning-theoretic quantities: (1) star number, a quantity that characterizes the label complexity of *pool-based active learning* [21], (2) threshold dimension, a quantity that characterizes the regret of *online learning* [3].¹ We believe that this new result may help us better understand how different learning tasks relate to each other.
- The second reason is that the combinatorial eluder dimension (or a multi-class variant of it) has already been studied for policy-based learning for contextual bandits and RL [see, e.g., 19, 34]. Thus, understanding the combinatorial eluder dimension may shed light on the challenges of policy-based RL.
- Our last reason has more philosophical bent. Historically, the discovery of VC dimension placed statistical learning theory on solid footing. Specifically, the fundamental theorem of statistical learning allows us to precisely characterize the statistical complexity of PAC learnability in terms of the combinatorial VC dimension. The insights from understanding the role of VC dimension in classification have led researchers to develop *scale-sensitive* complexity measures such as fat shattering dimension to provide sharper guarantees on learning. While we do not claim that the eluder dimension is fundamental to online RL and bandit settings, we believe that a better combinatorial understanding can lead to a better understanding of online decision making. In some sense, we are “working backwards” from the original scale-sensitive definition of eluder dimension to understand its combinatorial properties.

¹Finiteness of the threshold dimension is equivalent to finiteness of Littlestone dimension [40, 22, 3].

1.1 Main contributions

In this work, we provide several results which show when eluder dimension can be bounded (or is unbounded). Our results can be separated into two categories: (1) those pertaining to the (scale-sensitive) eluder dimension (as originally defined by Russo and Van Roy [37]); (2) those pertaining to the combinatorial eluder dimension, defined for binary-valued function classes.

First, we investigate the relationship between eluder dimension and our notion of *generalized rank* that captures the realizability of any function class as a GLM. In Section 2, we formally introduce the eluder dimension (as well as a related quantity of the scale-sensitive star number). In Section 3, we formalize the notion of “generalized rank”. In Section 4, we provide several results.

1. We show that eluder dimension can be *exponentially smaller* than σ -rank for any monotone activation σ , not just those with derivatives bounded away from 0 (Theorem 6).
2. We show that the condition on the derivative being bounded away from 0 is necessary for σ -rank to be an upper bound on eluder dimension. Namely, when σ is the relu activation, we show that eluder dimension can be *exponentially larger* than σ -rank (Theorem 7).²

In Section 5, we specialize the eluder dimension to the binary-valued setting and present our results on the combinatorial eluder dimension.

1. We show that eluder dimension is finite if and only if both star number and threshold dimension are finite. Specifically, in Theorem 8 we show the following:

$$\max\{\text{Sdim}(\mathcal{F}), \text{Tdim}(\mathcal{F})\} \leq \text{Edim}(\mathcal{F}) \leq \exp(O(\max\{\text{Sdim}(\mathcal{F}), \text{Tdim}(\mathcal{F})\})).$$

Furthermore, we demonstrate that both inequalities can be tight (see Theorem 9 and discussion above it).

2. We investigate the comparison between eluder dimension and sign-rk and prove stronger separations than in the scale-sensitive case; namely we show examples where one quantity is finite while the other is infinite (Theorem 10).

1.2 Related work

In this section, we highlight several related works.

Bounds on eluder dimension. Several papers provide bounds on eluder dimension for various function classes. The original bounds on tabular, linear, and generalized linear functions were proved by Russo and Van Roy [37] (and later generalized by Osband and Van Roy [36]). Mou et al. [34] provide several bounds for the combinatorial eluder dimension, mostly focusing on linear function classes. When the function class lies in an RKHS, Huang et al. [25] show that the eluder dimension is equivalent to the notion of *information gain* [41, 38], which can be seen as an infinite dimensional generalization of the fact that the eluder dimension for linear functions over \mathbb{R}^d is $\tilde{\Theta}(d)$. In concurrent work, Dong et al. [12] also prove an exponential lower bound on the eluder dimension for ReLU networks.

Applications of eluder dimension. The main application of the eluder dimension is to design algorithms and prove regret guarantees for contextual bandits and reinforcement learning. A few examples include the papers [46, 36, 44, 6, 14, 19, 28, 16, 26, 24, 34]. While the majority of papers prove upper bounds via eluder dimension, Foster et al. [19] provided lower bounds for contextual bandits in terms of eluder dimension, if one is hoping for instance-dependent regret bounds. In addition, several works observe that eluder dimension sometimes does not characterize the sample complexity, as the guarantee via eluder dimension can be too loose [23, 20]. Beyond the online RL setting, eluder dimension has been applied to risk sensitive RL [15], Markov games [24, 29], representation learning [47], and active online learning [10].

Other complexity measures for RL. We also touch upon other complexity measures which have been suggested for RL. One category is Bellman/Witness rank approach [see e.g. 27, 11, 42], which is generalized to bilinear classes [13]. These complexity measures capture an interplay between

²This result was independently established by Dong et al. [12].

the MDP dynamics and the function approximator class; in contrast, eluder dimension is purely a property of the function approximator class and can be stated (and studied) without referring to an online RL problem. Jin et al. [28] define a *Bellman-Eluder dimension* which captures function classes which have small Bellman rank or eluder dimension. Lastly, Foster et al. [20] propose a *Decision-Estimation Coefficient* and prove that it is necessary and sufficient for sample-efficient interactive learning.

Notions of rank. The notion of rank we propose is a generalization of the classical notion of *sign rank*, also called *dimension complexity*. Sign rank has been studied extensively in combinatorics, learning theory, and communication complexity [see e.g. 1, 18, 5, 2, and references therein]. The norm requirements in our definition of σ -rank are related to the notion of *margin complexity* [5, 7, 30].

2 Eluder dimension and star number

Eluder dimension is a “sequential” notion of complexity for function classes, originally defined by Russo and Van Roy [37]. Informally speaking, it characterizes the longest sequence of *adversarially chosen* points one must observe in order to accurately estimate the function value at any other point. We consider a variant of the original definition, proposed by Foster et al. [19], that is never larger and is sufficient to analyze all the applications of eluder dimension in literature.

Definition 1. For any function class $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \mathbb{R})$, $f^* \in \mathcal{F}$, and scale $\varepsilon \geq 0$, the **exact eluder dimension** $\underline{\text{Edim}}_{f^*}(\mathcal{F}, \varepsilon)$ is the largest m such that there exists $(x_1, f_1), \dots, (x_m, f_m) \in \mathcal{X} \times \mathcal{F}$ satisfying for all $i \in [m]$:

$$|f_i(x_i) - f^*(x_i)| > \varepsilon, \quad \text{and} \quad \sum_{j < i} (f_i(x_j) - f^*(x_j))^2 \leq \varepsilon^2. \quad (1)$$

Then for all $\varepsilon > 0$:

- the **eluder dimension** is $\text{Edim}_{f^*}(\mathcal{F}, \varepsilon) = \sup_{\varepsilon' \geq \varepsilon} \underline{\text{Edim}}_{f^*}(\mathcal{F}, \varepsilon')$.
- $\underline{\text{Edim}}(\mathcal{F}, \varepsilon) := \sup_{f^* \in \mathcal{F}} \underline{\text{Edim}}_{f^*}(\mathcal{F}, \varepsilon)$ and $\text{Edim}(\mathcal{F}, \varepsilon) := \sup_{f^* \in \mathcal{F}} \text{Edim}_{f^*}(\mathcal{F}, \varepsilon)$.

This definition is never larger than the original definition of Russo and Van Roy [37], which asks for a witnessing *pair* of functions $f_i, f'_i \in \mathcal{F}$ (the above restricts $f'_i = f^*$). Hence, all lower bounds on our variant of eluder dimension immediately apply to the original definition. Moreover, all upper bounds on eluder dimension in this paper can also be shown to hold for the other definition (unless stated otherwise).

We also consider the closely related notion of *star number* defined by Foster et al. [19], which generalizes a combinatorial parameter introduced in the active learning literature by Hanneke and Yang [21] (we will denote it as *Sdim* for consistency). We study the combinatorial star number in more detail in [Section 5](#). The *only* difference between the definitions of eluder dimension and star number is that $\sum_{j < i}$ is replaced by $\sum_{j \neq i}$, which makes the star number a “non-sequential” notion of complexity.

Definition 2. For any function class $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \mathbb{R})$, $f^* \in \mathcal{F}$, and scale $\varepsilon \geq 0$, the **exact star number** $\underline{\text{Sdim}}_{f^*}(\mathcal{F}, \varepsilon)$ is the largest m such that there exists $(x_1, f_1), \dots, (x_m, f_m) \in \mathcal{X} \times \mathcal{F}$ satisfying for all $i \in [m]$:

$$|f_i(x_i) - f^*(x_i)| > \varepsilon, \quad \text{and} \quad \sum_{j \neq i} (f_i(x_j) - f^*(x_j))^2 \leq \varepsilon^2.$$

Then for all $\varepsilon > 0$:

- the **star number** is $\text{Sdim}_{f^*}(\mathcal{F}, \varepsilon) = \sup_{\varepsilon' \geq \varepsilon} \underline{\text{Sdim}}_{f^*}(\mathcal{F}, \varepsilon')$.
- $\underline{\text{Sdim}}(\mathcal{F}, \varepsilon) := \sup_{f^* \in \mathcal{F}} \underline{\text{Sdim}}_{f^*}(\mathcal{F}, \varepsilon)$ and $\text{Sdim}(\mathcal{F}, \varepsilon) := \sup_{f^* \in \mathcal{F}} \text{Sdim}_{f^*}(\mathcal{F}, \varepsilon)$.

It immediately follows from these definitions that the star number is never larger than eluder dimension. On the other hand, the star number can be arbitrarily smaller than eluder dimension.

Proposition 1. For all \mathcal{F} , $f^* \in \mathcal{F}$ and scale $\varepsilon \geq 0$, it holds that³

$$\underline{\text{Sdim}}_{f^*}(\mathcal{F}, \varepsilon) \leq \underline{\text{Edim}}_{f^*}(\mathcal{F}, \varepsilon) \leq \min \{|\mathcal{X}|, |\mathcal{F}| - 1\}.$$

Proposition 2 (simplified from Foster et al. [19, Prop 2.3]). For the class of threshold functions given as $\mathcal{F}_n^{\text{th}} := \{f_i : [n] \rightarrow \{0, 1\} \mid i \in [n+1]\}$, where $f_i(x) := \mathbb{1}\{x \geq i\}$, and for $f^* = f_{n+1}$, it holds for all $\varepsilon < 1$ that $\underline{\text{Sdim}}_{f^*}(\mathcal{F}, \varepsilon) = 2$ and $\underline{\text{Edim}}_{f^*}(\mathcal{F}, \varepsilon) = n$.

3 Generalized rank

Dimension complexity has been studied extensively in combinatorics, learning theory, and communication complexity [see e.g. 1, 18, 5, 2]. The classical notion of dimension complexity, also known as *sign rank*, corresponds to the smallest dimension required to embed the input space such that all hypotheses in the function class under consideration are realizable as halfspaces. We consider a generalized notion of rank that is specified for any particular activation $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, and captures the smallest dimension required to represent the function class as a GLM when σ is the activation. In what follows, we let $\mathcal{B}_d(R) := \{x \in \mathbb{R}^d \mid \|x\|_2 \leq R\}$.

Definition 3. For any $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, the **σ -rank** of a function class $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \mathbb{R})$ at scale $R > 0$, denoted as $\sigma\text{-rk}(\mathcal{F}, R)$, is the smallest dimension d for which there exists mappings $\phi : \mathcal{X} \rightarrow \mathcal{B}_d(1)$ and $w : \mathcal{F} \rightarrow \mathcal{B}_d(R)$ such that⁴

$$\text{for all } (x, f) \in \mathcal{X} \times \mathcal{F} : f(x) = \sigma(\langle w(f), \phi(x) \rangle), \quad (2)$$

or ∞ if no such d exists. For a collection of activation functions $\Sigma \subseteq (\mathbb{R} \rightarrow \mathbb{R})$, the **Σ -rank** is

$$\Sigma\text{-rk}(\mathcal{F}, R) := \min_{\sigma \in \Sigma} \sigma\text{-rk}(\mathcal{F}, R).$$

Examples. We present some examples of $\Sigma\text{-rk}$ that motivate our definition above.

1. **Threshold activation.** $\text{sign}(z)$ yields the classic notion of *sign-rank* (equivalent to *dimension complexity*, as already mentioned). In this case, the scale R is irrelevant, so we denote $\text{sign-rk}(\mathcal{F}) := \text{sign-rk}(\mathcal{F}, R)$ for any R . Note that this quantity is meaningful only for $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \{-1, 1\})$.
2. **Identity activation.** For $\text{id}(z) := z$, $\text{id-rk}(\mathcal{F}, R)$ is the smallest dimension needed to represent each $f \in \mathcal{F}$ as a (norm-bounded) linear function. We abbreviate $\text{rk} := \text{id-rk}$, as this corresponds to the standard notion of rank of the matrix $(f(x))_{x, f}$ (albeit with the additional norm constraint).
3. **Monotone activations.** For $L \geq \mu \geq 0$, \mathcal{M}_μ^L consists of all activations σ such that for all $z < z'$, it holds that $\mu \leq \frac{\sigma(z') - \sigma(z)}{z' - z} \leq L$ (for differentiable σ , this is equivalent to $\mu \leq \sigma'(z) \leq L$ for all $z \in \mathbb{R}$).⁵ An important special case is when $\mu = 0$, and a particular $\sigma \in \mathcal{M}_0^1$ of interest is the rectified linear unit (ReLU) defined as $\text{relu}(z) := \max\{z, 0\}$. For ease of notation, we denote $\mathcal{M}_\mu := \mathcal{M}_\mu^1$. While we will always be explicit about the Lipschitz constant, note that the scale of the Lipschitz constant L (and μ) is interchangeable with the scale of R . In particular,

$$\mathcal{M}_\mu^L\text{-rk}(\mathcal{F}, R) = \mathcal{M}_{\mu/L}\text{-rk}(\mathcal{F}, RL). \quad (3)$$

4. **All activations.** Σ^{all} consists of all activations σ . We mention this notion of rank only in passing, and we will not focus on it for the rest of the paper.

We present a result which relates the aforementioned quantities (proof in [Appendix A](#)).

³For the definition of eluder dimension considered by [37], an upper bound of $\min\{|\mathcal{X}|, \binom{|\mathcal{F}|}{2}\}$ holds, which can be tight. This upper bound holds because the witnessing pair of functions (f_i, f'_i) has to be distinct for each i .

⁴Note that only the product of the scales of ϕ and w is relevant. The definition remains equivalent if we let $\phi : \mathcal{X} \rightarrow \mathcal{B}_d(R_\phi)$ and $w : \mathcal{F} \rightarrow \mathcal{B}_d(R_w)$ for any R_ϕ and R_w such that $R = R_\phi \cdot R_w$.

⁵To prove upper bounds on eluder dimension, it suffices for this condition to hold only when $|z| \leq R$, [see e.g. 37]. Since we fix σ in our definition first and then consider σ -rank at different scales R , this weaker condition complicates our definitions. Note that at any specific scale R , we can always modify σ to satisfy the required constraint everywhere by extending it linearly whenever $|z| > R$.

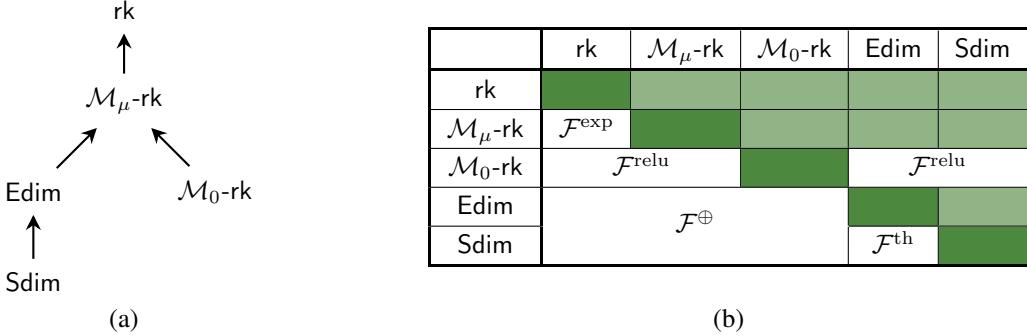


Figure 1: (a) Each arrow $M_1 \rightarrow M_2$ indicates that $M_1(\mathcal{F}) \lesssim M_2(\mathcal{F})$ for all \mathcal{F} , where the dependence on R and ε is suppressed for clarity (see Propositions 1, 3, 4 for precise bounds). Whenever $M_2 \rightarrow M_1$ arrow is missing, there is an example of a class \mathcal{F} where $M_1(\mathcal{F}) \ll M_2(\mathcal{F})$. (b) An entry \mathcal{F} in (M_1, M_2) means that $M_1(\mathcal{F}) \ll M_2(\mathcal{F})$. Green cells indicate that $M_1(\mathcal{F}) \gtrsim M_2(\mathcal{F})$ for all \mathcal{F} .

Proposition 3. For all $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \mathbb{R})$, $R > 0$ and $\mu \in (0, 1]$, we have:

$$\text{rk}(\mathcal{F}, R) \geq \mathcal{M}_\mu\text{-rk}(\mathcal{F}, R) \geq \mathcal{M}_0\text{-rk}(\mathcal{F}, R) \geq \text{sign-rk}(\mathcal{F}) - 1,$$

where the last inequality is meaningful only for $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \{-1, 1\})$. Moreover, for each inequality above, there exists a function class \mathcal{F} which exhibits an infinite gap between the two quantities.

4 Eluder dimension versus generalized rank

In this section, we compare eluder dimension and star number with each notion of generalized rank: rk , $\mathcal{M}_\mu\text{-rk}$ (for $\mu > 0$) and $\mathcal{M}_0\text{-rk}$. Our results are summarized in Figure 1.

Eluder vs. rk and $\mathcal{M}_\mu\text{-rk}$. Russo and Van Roy [37] and Osband and Van Roy [36] provided upper bounds on eluder dimension for linear and generalized linear function classes. For completeness, we restate this result, with a slight improvement and include the proof with precise dependence on problem parameters in Appendix B. Intuitively, the generalized linear rank allows us to capture the tightest possible upper bound that the guarantees in the papers [37, 36] can provide on eluder dimension.

Proposition 4 (cf. [37], Prop. 6, 7; [36], Prop. 2, 4). For any function class $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \mathbb{R})$ and $\varepsilon > 0$:

- (i) For all $R > 0$, $\underline{\text{Edim}}(\mathcal{F}, \varepsilon) \leq \text{rk}(\mathcal{F}, R) \cdot O\left(\log \frac{R}{\varepsilon}\right)$.
- (ii) For all $L, \mu, R > 0$, $\underline{\text{Edim}}(\mathcal{F}, \varepsilon) \leq \mathcal{M}_\mu^L\text{-rk}(\mathcal{F}, R) \cdot O\left(\frac{L^2}{\mu^2} \log\left(\frac{RL}{\varepsilon}\right)\right)$.

This result has been used to prove upper bounds on eluder dimension of various function classes beyond GLMs; for example, the class of bounded degree polynomials, by taking the feature map $\phi(x)$ to be the vector of low degree monomials. The upper bound in Part (i) is in fact tight (up to constants) for the class of linear functions, as shown in Proposition 5 below. This trivially implies the optimality of the bound in Part (ii) up to the factor of $(L/\mu)^2$ which to the best of our knowledge is open.

Proposition 5 ([33]). For any $R > 0$, $\mathcal{X} := \mathcal{B}_d(1)$ and $\mathcal{F} := \{f_\theta : x \mapsto \langle \theta, x \rangle \mid \theta \in \mathcal{B}_d(R)\}$, it holds that:

$$\underline{\text{Edim}}(\mathcal{F}, \varepsilon) \geq \Omega\left(d \log\left(\frac{R}{\varepsilon}\right)\right).$$

For completeness, we include the proof in Appendix C.

Eluder vs. $\mathcal{M}_0\text{-rk}$. It turns out that eluder dimension and $\mathcal{M}_0\text{-rk}$ are incomparable. That is, there exists a function class for which eluder dimension is exponentially smaller than $\mathcal{M}_0\text{-rk}$ (and hence $\mathcal{M}_\mu\text{-rk}$ and rk by Proposition 3). Moreover, there exists a different function class for which eluder dimension (even the star number) is exponentially larger than relu-rk (and hence $\mathcal{M}_0\text{-rk}$).

First, we show that the eluder dimension can be exponentially smaller than $\mathcal{M}_0\text{-rk}$ for the class of parities over d bits. Thus, parities over d bits exhibits an example where the eluder dimension is exponentially smaller than the best possible bound one can derive using the existing results of [Proposition 4](#).

Theorem 6. *For $\mathcal{X} = \{-1, 1\}^d$ and $\mathcal{F}^\oplus := \{f_S : x \mapsto \prod_{i \in S} x_i \mid S \subseteq [d]\}$, it holds that*

- (i) $\mathcal{M}_0\text{-rk}(\mathcal{F}^\oplus, R) \geq 2^{d/2} - 1$ for all $R > 0$.
- (ii) $\underline{\text{Sdim}}(\mathcal{F}^\oplus, \varepsilon) \leq \underline{\text{Edim}}(\mathcal{F}^\oplus, \varepsilon) \leq d$ for all $\varepsilon \geq 0$.

Proof. Part (i). From [Proposition 3](#), we have that $\mathcal{M}_0\text{-rk}(\mathcal{F}^\oplus, R) \geq \text{sign-rk}(\mathcal{F}^\oplus) - 1$ for any $\sigma \in \mathcal{M}_0$. The proof is now complete by noting a well known result that $\text{sign-rk}(\mathcal{F}^\oplus) \geq 2^{d/2}$ [18].

Part (ii). For any $x \in \{-1, 1\}^d$ consider its 0-1 representation $\tilde{x} \in \mathbb{F}_2^d$ (representing +1 by 0 and -1 by 1). All functions in \mathcal{F}^\oplus can be simply viewed as linear functions over \mathbb{F}_2 . Namely, any parity function is indexed by a vector $a \in \mathbb{F}_2^d$, with $f_a(x) := (-1)^{\langle a, \tilde{x} \rangle}$.

Note that $\underline{\text{Edim}}(\mathcal{F}^\oplus, \varepsilon) = 0$ for all $\varepsilon \geq 2$. For any $\varepsilon < 2$, suppose $\underline{\text{Edim}}_{f^*}(\mathcal{F}^\oplus, \varepsilon) = m$, witnessed by $(x_1, f_{a_1}), \dots, (x_m, f_{a_m}) \in \{-1, 1\}^d$ and $f^* = f_{a^*}$. We have

- $f_{a_i}(x_i) \neq f_{a^*}(x_i)$, and
- $f_{a_i}(x_j) = f_{a^*}(x_j)$ for all $j < i$: since $\sum_{j < i} (f_{a_i}(x_j) - f_{a^*}(x_j))^2 < \varepsilon^2 < 4$ iff all terms are 0.

Thus, we have $\langle a_i - a^*, \tilde{x} \rangle = 0$ for all $\tilde{x} \in \mathbb{F}_2\text{-span}(\{\tilde{x}_1, \dots, \tilde{x}_{i-1}\})$. But $\langle a_i - a^*, \tilde{x}_i \rangle = 1$ and hence \tilde{x}_i is linearly independent of $\{\tilde{x}_1, \dots, \tilde{x}_{i-1}\}$ over \mathbb{F}_2^d . Thus, $\{\tilde{x}_1, \dots, \tilde{x}_m\}$ are all linearly independent over \mathbb{F}_2^d , and hence $m \leq d$. \square

The bound in part (ii) of [Theorem 6](#) was also calculated by Mou et al. [34, Prop. 3].

Next, we show a separation in the other direction for eluder dimension vs. $\mathcal{M}_0\text{-rk}$ using the ReLU function class. Thus, we cannot hope to remove the requirement for the activation function σ to be strictly monotonically increasing in [Proposition 4](#) (ii) for bounding the eluder dimension.

Theorem 7. *Let $R > 0$ and $\mathcal{X} = \mathcal{B}_d(1)$. Define*

$$\mathcal{F}^{\text{relu}} := \{f_{\theta, b} : x \mapsto \text{relu}(\langle \theta, x \rangle + b) \mid \|\theta\|^2 + b^2 \leq R^2\}.$$

It holds that

- (i) $\mathcal{M}_0\text{-rk}(\mathcal{F}^{\text{relu}}, R) \leq \text{relu-rk}(\mathcal{F}^{\text{relu}}, R) \leq d + 1$,
- (ii) $\underline{\text{Edim}}(\mathcal{F}^{\text{relu}}, \varepsilon) \geq \underline{\text{Sdim}}(\mathcal{F}^{\text{relu}}, \varepsilon) \geq \left(\frac{R}{4\varepsilon}\right)^{d/2}$ for all $\varepsilon \in (0, \frac{R}{4})$.

Proof. Part (i) is immediate from the definition. We show Part (ii) in the special case of $R = 2$; the general case follows by relatively scaling ε , since relu is *homogeneous*, namely, $\text{relu}(ax) = a \cdot \text{relu}(x)$. Consider any $U \subseteq \mathcal{X}$ such that $\|u\| = 1$ and $\langle u, v \rangle \leq 1 - \varepsilon$ for all $u, v \in U$. It holds that $\underline{\text{Sdim}}_{f^*}(\mathcal{F}^{\text{relu}}, \varepsilon) \geq |U|$ when f^* is the identically zero function, since the function $f_u(x) = \text{relu}(\langle u, x \rangle - (1 - \varepsilon))$ is such that $f_u(v) = 0$ for all $v \in U \setminus \{u\}$, whereas $f_u(u) = \varepsilon$. A standard sphere packing argument shows that such a set U exists with $|U| \geq (1/2\varepsilon)^{d/2}$ for all $\varepsilon < 1/2$. In particular, the δ -packing number of the unit sphere is at least $(1/\delta)^d$ [43, Cor. 4.2.13]. Thus, we can find $(1/\delta)^d$ points such that each pair u, v satisfies $\|u - v\| \geq \delta$, or equivalently $\langle u, v \rangle \leq 1 - \delta^2/2$. Setting $\delta = \sqrt{2\varepsilon}$ proves the claimed lower bound. \square

[Theorem 7](#) was independently shown by Dong et al. [12, Thm. 5.1].

We remark that while we considered the variant of eluder dimension as defined by Foster et al. [19], the lower bound on eluder dimension in [Theorem 7](#) immediately holds for the notion defined by Russo and Van Roy [37]. On the other hand, the upper bound on eluder dimension in [Theorem 6](#) can be shown to hold even with the definition of Russo and Van Roy [37] (by replacing every instance of a^* by a'_i).

	f_1	f_2	f_3	f_4	f_5	f_6
x_1	1	0	0	0	0	0
x_2	*	1	0	0	0	0
x_3	*	*	1	0	0	0
x_4	*	*	*	1	0	0
x_5	*	*	*	*	1	0
x_6	*	*	*	*	*	1

Eluder sequence

	f_1	f_2	f_3	f_4	f_5	f_6
x_1	1	0	0	0	0	0
x_2	0	1	0	0	0	0
x_3	0	0	1	0	0	0
x_4	0	0	0	1	0	0
x_5	0	0	0	0	1	0
x_6	0	0	0	0	0	1

Star sequence

	f_1	f_2	f_3	f_4	f_5	f_6
x_1	1	0	0	0	0	0
x_2	1	1	0	0	0	0
x_3	1	1	1	0	0	0
x_4	1	1	1	1	0	0
x_5	1	1	1	1	1	0
x_6	1	1	1	1	1	1

Threshold sequence

Figure 2: Illustration of witnessing sequences of length 6 for eluder dimension, star number and threshold dimension with respect to $f^* = 0$ (we use the 0/1 representation of functions for clarity). ‘*’ in the eluder witness sequence refers to a free value, either 0 or 1.

5 Relationships between combinatorial measures

In this section, we specialize the eluder dimension to binary-valued function classes and prove several additional characterizations that relate this combinatorial eluder dimension to other learning-theoretic quantities. First, we (re)define these learning-theoretic quantities. The first two definitions (of combinatorial eluder dimension and star number respectively) are the specialization of the scale-sensitive versions (cf. Definition 1 and 2) to the binary-valued function setting. We abuse notation by dropping the argument ε from previous definitions in order to be consistent.

Definition 4. Fix any function class $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \{1, -1\})$ and any $f^* \in \mathcal{F}$.

- The **combinatorial eluder dimension** w.r.t. f^* , denoted $\text{Edim}_{f^*}(\mathcal{F})$, is defined as the largest m such that there exists $(x_1, f_1), \dots, (x_m, f_m) \in \mathcal{X} \times \mathcal{F}$ satisfying for all $i \in [m]$:

$$f_i(x_i) \neq f^*(x_i), \quad \text{and} \quad \text{for all } j < i : \quad f_i(x_j) = f^*(x_j).$$

- The **star number** w.r.t. f^* , denoted $\text{Sdim}_{f^*}(\mathcal{F})$, is defined as the largest m such that there exists $(x_1, f_1), \dots, (x_m, f_m) \in \mathcal{X} \times \mathcal{F}$ satisfying for all $i \in [m]$:

$$f_i(x_i) \neq f^*(x_i), \quad \text{and} \quad \text{for all } j \neq i : \quad f_i(x_j) = f^*(x_j).$$

- The **threshold dimension** w.r.t. f^* , denoted $\text{Tdim}_{f^*}(\mathcal{F})$, is defined as the largest m such that there exists $(x_1, f_1), \dots, (x_m, f_m) \in \mathcal{X} \times \mathcal{F}$ satisfying for all $i \in [m]$:

$$\text{for all } k \geq i : \quad f_i(x_k) \neq f^*(x_k), \quad \text{and} \quad \text{for all } j < i : \quad f_i(x_j) = f^*(x_j).$$

As before, we define the combinatorial eluder dimension (resp. star number and threshold dimension) to be $\text{Edim}(\mathcal{F}) := \sup_{f \in \mathcal{F}} \text{Edim}_f(\mathcal{F})$ ($\text{Sdim}(\mathcal{F})$ and $\text{Tdim}(\mathcal{F})$ are defined similarly).

Let us pause to unpack these definitions and give some background. In fact, the star number definition stated above is the original definition of Hanneke and Yang [21], who give tight upper and lower bounds on the label complexity of *pool-based active learning* via the star number $\text{Sdim}(\mathcal{F})$ and show that almost every previously proposed complexity measure for active learning takes a worst case value equal to the star number. Roughly speaking, the star number corresponds to the number of ‘‘singletons’’ one can embed in a function class; that is, the maximum number of functions that differ from a base function f^* at exactly one point among a subset of the domain $\{x_1, \dots, x_m\} \subseteq \mathcal{X}$.

The threshold dimension has recently gained attention due to its role in proving an equivalence relationship between private PAC learning and online learning [see, e.g., 3, 9]. We slightly generalize the definition of Alon et al. [3] to allow for any base function f^* , in the spirit of the other two definitions. A classical result in model theory provides a link between the threshold dimension and Littlestone dimension (which we denote Ldim), a quantity which is both necessary and sufficient for online learnability [8, 4]. In particular, results by Shelah [40] and Hodges et al. [22] show that for any binary-valued \mathcal{F} and any $f^* \in \mathcal{F}$:

$$\lfloor \log \text{Tdim}_{f^*}(\mathcal{F}) \rfloor \leq \text{Ldim} \leq 2^{\text{Tdim}_{f^*}(\mathcal{F})}.$$

A combinatorial proof of this fact can be found in Thm. 3 of Alon et al. [3].⁶ Thus, finiteness of threshold dimension is necessary and sufficient for online learnability (albeit in a much weaker, “qualitative” sense).

5.1 A qualitative equivalence

Now, we prove a rather surprising characterization that closely ties all three quantities in [Definition 4](#) together. Our result implies that for any binary-valued function class, finiteness of the combinatorial eluder dimension is *equivalent* to finiteness of both star number and threshold dimension.

Theorem 8. *For any function class $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \{1, -1\})$ and any $f^* \in \mathcal{F}$, the following holds:*

$$\max\{\text{Sdim}_{f^*}(\mathcal{F}), \text{Tdim}_{f^*}(\mathcal{F})\} \leq \text{Edim}_{f^*}(\mathcal{F}) \leq 4^{\max\{\text{Sdim}_{f^*}(\mathcal{F}), \text{Tdim}_{f^*}(\mathcal{F})\}}.$$

The proof of [Theorem 8](#) can be found in [Appendix D](#). The lower bound is trivial by examining [Definition 4](#). To prove the upper bound, we rely on a novel connection to Ramsey theory. In particular, we show that sequences $(x_1, f_1), \dots, (x_m, f_m)$ which witness $\text{Edim}_{f^*} = m$ form a bijection with red-blue colorings of the complete graph K_m , while subsequences of the witnessing eluder sequence that are valid star number witnesses or threshold dimension witnesses can be interpreted as monochromatic colorings of subgraphs of K_m (see [Figure 2](#)). Thus, applying classical bounds from Ramsey theory implies the result.

[Theorem 8](#) has an exponential gap between the upper and lower bounds. Can we improve either of the inequalities? The lower bound cannot be improved, by considering the simple examples over $\mathcal{X} = [n]$ of the singleton class $\mathcal{F}_n^{\text{sing}} := \{x \mapsto \mathbb{1}\{x = i\} \mid i \in [n+1]\}$ and the threshold class $\mathcal{F}_n^{\text{th}} := \{x \mapsto \mathbb{1}\{x \geq i\} \mid i \in [n+1]\}$. We also show that the upper bound cannot be improved, for example, to $\text{Edim}(\mathcal{F}) \leq \text{poly}(\text{Sdim}(\mathcal{F}), \text{Tdim}(\mathcal{F}))$ in general.

Theorem 9. *For every $N > 0$, there exists a function class \mathcal{F}_N such that $\text{Edim}(\mathcal{F}_N) = N$ and $\max\{\text{Sdim}(\mathcal{F}), \text{Tdim}(\mathcal{F})\} < c \cdot \log_2 N$, where $c > 1/2$ is some absolute numerical constant.*

The proof of [Theorem 9](#) can be found in [Appendix E](#). It relies on the probabilistic method to show the existence of a randomly constructed \mathcal{F} which satisfies the desired properties.

5.2 Comparisons with sign-rk

In this section, we investigate the comparison of these combinatorial quantities (eluder, star, threshold) with sign-rk, and examine whether we can prove stronger separations. One direction is clear; for the class of linear classifiers in \mathbb{R}^d , the combinatorial eluder dimension, star number, and threshold dimension are all infinite. However, we ask if the other separation is also possible: can we construct \mathcal{F} where eluder/star/threshold are finite but sign-rk = ∞ ? We have already provided an explicit exponential separation: [Theorem 6](#) shows that for the function class \mathcal{F}^\oplus , we had $\text{Edim}(\mathcal{F}^\oplus) = \text{Sdim}(\mathcal{F}^\oplus) = d$ but $\text{sign-rk}(\mathcal{F}^\oplus) \geq 2^{d/2}$. (One can also show that $\text{Tdim}(\mathcal{F}^\oplus) = d$.) We are able to show stronger (but nonconstructive) separations by extending the probabilistic techniques of Alon et al. [2], who recently provided similar separations for VC dimension versus sign-rk. In view of [Proposition 3](#), this result also provides a separation for the scale-sensitive [Definition 1](#).

Theorem 10. *For every $N > 0$, there exists a function class $\mathcal{F}_N \subseteq ([N] \rightarrow \{1, -1\})$ such that $\text{Edim}_1(\mathcal{F}_N) = 4$ and $\text{sign-rk}(\mathcal{F}_N) \geq \Omega(N^{1/9} / \log N)$, where 1 is shorthand for the all 1s function.*

The proof of [Theorem 10](#) can be found in [Appendix F](#). It is straightforward to replace the reference function $f^*(x) = 1$ with any fixed reference function $f^* : [N] \rightarrow \{1, -1\}$. First, we use Lemma 22 of Alon et al. [2] which bounds the number of distinct matrices with sign-rk = r ; then using a probabilistic argument we show that there must be many (more) matrices with $\text{Edim}_1 = 4$, so at least one of them must have large sign-rk.

The careful reader might notice that we do not prove the existence of a function class where $\text{Edim}(\mathcal{F})$ is constant and the sign-rk is infinite; instead we prove the weaker statement that a function class exists with Edim w.r.t. any *fixed* function f^* is bounded. We conjecture that the stronger statement holds; see [Appendix F.1](#) for more details.

⁶Alon et al. [3] prove the result for $f^*(x) = -1$, but it is easy to extend their proof to hold for any f^* .

5.3 Back to scale-sensitive?

It is natural to ask if our results can be extended back to the scale-sensitive definitions. First, we require a scale-sensitive version of threshold dimension. One proposal is the following:

Definition 5. For any function class $\mathcal{F} \subseteq (\mathcal{X} \rightarrow \mathbb{R})$, $f^* \in \mathcal{F}$, and scale $\varepsilon \geq 0$, the **exact threshold dimension** $\underline{\text{Tdim}}_{f^*}(\mathcal{F}, \varepsilon)$ is the largest m such that there exists $(x_1, f_1), \dots, (x_m, f_m) \in \mathcal{X} \times \mathcal{F}$ satisfying for all $i \in [m]$:

$$\forall j \geq i : |f_i(x_j) - f^*(x_j)| > \varepsilon, \quad \text{and} \quad \sum_{j < i} (f_i(x_j) - f^*(x_j))^2 \leq \varepsilon^2.$$

Then for all $\varepsilon > 0$:

- the **threshold dimension** is $\text{Tdim}_{f^*}(\mathcal{F}, \varepsilon) = \sup_{\varepsilon' \geq \varepsilon} \underline{\text{Tdim}}_{f^*}(\mathcal{F}, \varepsilon')$.
- $\underline{\text{Tdim}}(\mathcal{F}, \varepsilon) := \sup_{f^* \in \mathcal{F}} \underline{\text{Tdim}}_{f^*}(\mathcal{F}, \varepsilon)$ and $\text{Tdim}(\mathcal{F}, \varepsilon) := \sup_{f^* \in \mathcal{F}} \text{Tdim}_{f^*}(\mathcal{F}, \varepsilon)$.

[Definition 5](#) mirrors the scale-sensitive definitions for eluder and star; it also recovers the combinatorial definition when \mathcal{F} is binary-valued. By definition, the relationship that $\text{Edim}_{f^*}(\mathcal{F}, \varepsilon) \geq \max\{\text{Sdim}_{f^*}(\mathcal{F}, \varepsilon), \text{Tdim}_{f^*}(\mathcal{F}, \varepsilon)\}$ for every $\mathcal{F}, f^* \in \mathcal{F}, \varepsilon > 0$ is trivial. However, one cannot hope to prove the corresponding upper bound under this definition. For example, for any $\varepsilon > 0$, take the function class which is represented by the $N \times N$ matrix:

$$f_j(x_i) = \begin{cases} 0 & i < j \\ \varepsilon & i = j \\ 0.99\varepsilon & i > j. \end{cases}$$

It is easy to see that $\text{Edim}_0(\mathcal{F}, \varepsilon) = N$, while $\text{Sdim}_0(\mathcal{F}, \varepsilon) = 2$ and $\text{Tdim}_0(\mathcal{F}, \varepsilon) = 1$. Notice that this class is still “threshold-like”, but [Definition 5](#) does not capture this for said value of ε . Generally speaking, it is unclear how to carry over the intuition from Ramsey theory that applies in the combinatorial case to the scale-sensitive case; we leave this to future work.

Acknowledgments and Disclosure of Funding

We thank Gaurav Mahajan for allowing us to include the proof of [Proposition 5](#) [33]. We thank Akshay Krishnamurthy, Tengyu Ma, and Ruosong Wang for helpful discussions. GL was partially supported by NSF award IIS-1764032. PK was partially supported by NSF BIGDATA award 1546500. Part of this work was done while GL, PK, and DF were participating in the Simons program on the Theoretical Foundations of Reinforcement Learning.

References

- [1] N. Alon, P. Frankl, and V. Rödl. Geometrical realization of set systems and probabilistic communication complexity. In *26th Annual Symposium on Foundations of Computer Science*, pages 277–280. IEEE Computer Society, 1985. URL <https://doi.org/10.1109/SFCS.1985.30>.
- [2] N. Alon, S. Moran, and A. Yehudayoff. Sign rank versus VC dimension. In *29th Annual Conference on Learning Theory*, volume 49 of *PMLR*, pages 47–80, 2016. URL <http://proceedings.mlr.press/v49/alon16.html>.
- [3] N. Alon, R. Livni, M. Malliaris, and S. Moran. Private pac learning implies finite littlestone dimension. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 852–860, New York, NY, USA, 2019. Association for Computing Machinery. URL <https://doi.org/10.1145/3313276.3316312>.
- [4] N. Alon, O. Ben-Eliezer, Y. Dagan, S. Moran, M. Naor, and E. Yoge. Adversarial laws of large numbers and optimal regret in online classification. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 447–455, 2021.
- [5] R. I. Arriaga and S. S. Vempala. An algorithmic theory of learning: Robust concepts and random projection. *Mach. Learn.*, 63(2):161–182, 2006. URL <https://doi.org/10.1007/s10994-006-6265-7>.
- [6] A. Ayoub, Z. Jia, C. Szepesvari, M. Wang, and L. Yang. Model-based reinforcement learning with value-targeted regression. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *PMLR*, pages 463–474, 2020. URL <http://proceedings.mlr.press/v119/ayoub20a.html>.
- [7] S. Ben-David, N. Eiron, and H. U. Simon. Limitations of learning via embeddings in euclidean half spaces. *Journal of Machine Learning Research*, 3(Nov):441–461, 2002.
- [8] S. Ben-David, D. Pál, and S. Shalev-Shwartz. Agnostic online learning. In *COLT*, volume 3, page 1, 2009.
- [9] M. Bun, R. Livni, and S. Moran. An equivalence between private classification and online prediction. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 389–402. IEEE, 2020.
- [10] Y. Chen, H. Luo, T. Ma, and C. Zhang. Active online learning with hidden shifting domains. In *International Conference on Artificial Intelligence and Statistics*, pages 2053–2061. PMLR, 2021.
- [11] K. Dong, J. Peng, Y. Wang, and Y. Zhou. Root-n-regret for learning in markov decision processes with function approximation and low bellman rank. In *Conference on Learning Theory*, pages 1554–1557. PMLR, 2020.
- [12] K. Dong, J. Yang, and T. Ma. Provable model-based nonlinear bandit and reinforcement learning: Shelve optimism, embrace virtual curvature. *arXiv*, 2102.04168, 2021. URL <https://arxiv.org/abs/2102.04168>.
- [13] S. Du, S. Kakade, J. Lee, S. Lovett, G. Mahajan, W. Sun, and R. Wang. Bilinear classes: A structural framework for provable generalization in rl. In *International Conference on Machine Learning*, pages 2826–2836. PMLR, 2021.
- [14] S. S. Du, J. D. Lee, G. Mahajan, and R. Wang. Agnostic Q-learning with function approximation in deterministic systems: Near-optimal bounds on approximation error and sample complexity. In *Advances in Neural Information Processing Systems*, volume 33, pages 22327–22337. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/fd5c905bcd8c3348ad1b35d7231ee2b1-Paper.pdf>.
- [15] Y. Fei, Z. Yang, and Z. Wang. Risk-sensitive reinforcement learning with function approximation: A debiasing approach. In *International Conference on Machine Learning*, pages 3198–3207. PMLR, 2021.

[16] F. Feng, W. Yin, A. Agarwal, and L. Yang. Provably correct optimization and exploration with non-linear policies. In *International Conference on Machine Learning*, pages 3263–3273. PMLR, 2021.

[17] S. Filippi, O. Cappe, A. Garivier, and C. Szepesvári. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc., 2010. URL <https://proceedings.neurips.cc/paper/2010/file/c2626d850c80ea07e7511bbae4c76f4b-Paper.pdf>.

[18] J. Forster. A linear lower bound on the unbounded error probabilistic communication complexity. *J. Comput. Syst. Sci.*, 65(4):612–625, 2002. URL [https://doi.org/10.1016/S0022-0000\(02\)00019-3](https://doi.org/10.1016/S0022-0000(02)00019-3).

[19] D. J. Foster, A. Rakhlin, D. Simchi-Levi, and Y. Xu. Instance-dependent complexity of contextual bandits and reinforcement learning: A disagreement-based perspective. *arXiv*, 2010.03104, 2020. URL <https://arxiv.org/abs/2010.03104>.

[20] D. J. Foster, S. M. Kakade, J. Qian, and A. Rakhlin. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.

[21] S. Hanneke and L. Yang. Minimax analysis of active learning. *Journal of Machine Learning Research*, 16(109):3487–3602, 2015. URL <http://jmlr.org/papers/v16/hanneke15a.html>.

[22] W. Hodges et al. *A shorter model theory*. Cambridge university press, 1997.

[23] B. Huang, K. Huang, S. Kakade, J. D. Lee, Q. Lei, R. Wang, and J. Yang. Going beyond linear rl: Sample efficient neural function approximation. *Advances in Neural Information Processing Systems*, 34, 2021.

[24] B. Huang, J. D. Lee, Z. Wang, and Z. Yang. Towards general function approximation in zero-sum markov games. *arXiv preprint arXiv:2107.14702*, 2021.

[25] K. Huang, S. M. Kakade, J. D. Lee, and Q. Lei. A short note on the relationship of information gain and eluder dimension. *arXiv preprint arXiv:2107.02377*, 2021.

[26] H. Ishfaq, Q. Cui, V. Nguyen, A. Ayoub, Z. Yang, Z. Wang, D. Precup, and L. Yang. Randomized exploration in reinforcement learning with general value function approximation. In *International Conference on Machine Learning*, pages 4607–4616. PMLR, 2021.

[27] N. Jiang, A. Krishnamurthy, A. Agarwal, J. Langford, and R. E. Schapire. Contextual decision processes with low Bellman rank are PAC-learnable. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1704–1713, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/jiang17c.html>.

[28] C. Jin, Q. Liu, and S. Miryoosefi. Bellman eluder dimension: New rich classes of RL problems, and sample-efficient algorithms. *arXiv*, 2102.00815, 2021. URL <https://arxiv.org/abs/2102.00815>.

[29] C. Jin, Q. Liu, and T. Yu. The power of exploiter: Provable multi-agent rl in large state spaces. *arXiv preprint arXiv:2106.03352*, 2021.

[30] P. Kamath, O. Montasser, and N. Srebro. Approximate is good enough: Probabilistic variants of dimensional and margin complexity. In *Conference on Learning Theory*, pages 2236–2262. PMLR, 2020.

[31] B. Kveton, M. Zaheer, C. Szepesvari, L. Li, M. Ghavamzadeh, and C. Boutilier. Randomized exploration in generalized linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 2066–2076. PMLR, 2020.

[32] L. Li, Y. Lu, and D. Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2071–2080. PMLR, 06–11 Aug 2017. URL <http://proceedings.mlr.press/v70/li17c.html>.

- [33] G. Mahajan and S. Lovett. Personal communication. 2021.
- [34] W. Mou, Z. Wen, and X. Chen. On the sample complexity of reinforcement learning with policy space generalization. *arXiv preprint arXiv:2008.07353*, 2020.
- [35] D. Mubayi and A. Suk. A survey of quantitative bounds for hypergraph ramsey problems. *arxiv e-prints* (july 2017). *arXiv preprint math.CO/1707.04229*, 2017.
- [36] I. Osband and B. Van Roy. Model-based reinforcement learning and the eluder dimension. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/1141938ba2c2b13f5505d7c424ebae5f-Paper.pdf>.
- [37] D. Russo and B. Van Roy. Eluder dimension and the sample complexity of optimistic exploration. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL <https://proceedings.neurips.cc/paper/2013/file/41bfd20a38bb1b0bec75acf0845530a7-Paper.pdf>.
- [38] D. Russo and B. Van Roy. An information-theoretic analysis of thompson sampling. *The Journal of Machine Learning Research*, 17(1):2442–2471, 2016.
- [39] S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014. URL <https://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/>.
- [40] S. Shelah. *Classification theory: and the number of non-isomorphic models*. Elsevier, 1990.
- [41] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- [42] W. Sun, N. Jiang, A. Krishnamurthy, A. Agarwal, and J. Langford. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Conference on learning theory*, pages 2898–2933. PMLR, 2019.
- [43] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018. URL <https://www.math.uci.edu/~rvershyn/papers/HDP-book/HDP-book.html>.
- [44] R. Wang, R. R. Salakhutdinov, and L. Yang. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. In *Advances in Neural Information Processing Systems*, volume 33, pages 6123–6135. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/440924c5948e05070663f88e69e8242b-Paper.pdf>.
- [45] Y. Wang, R. Wang, S. S. Du, and A. Krishnamurthy. Optimism in reinforcement learning with generalized linear function approximation. *arXiv preprint arXiv:1912.04136*, 2019.
- [46] Z. Wen and B. Van Roy. Efficient exploration and value function generalization in deterministic systems. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL <https://proceedings.neurips.cc/paper/2013/file/bad5f33780c42f2588878a9d07405083-Paper.pdf>.
- [47] Z. Xu and A. Tewari. Representation learning beyond linear prediction functions. *Advances in Neural Information Processing Systems*, 34, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]**
 - (b) Did you describe the limitations of your work? **[Yes]**
 - (c) Did you discuss any potential negative societal impacts of your work? **[N/A]** *Since this is a theory paper regarding several notions of complexity for learning problems, we do not foresee any immediate potential negative impacts.*
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]**
 - (b) Did you include complete proofs of all theoretical results? **[Yes]**
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[N/A]**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[N/A]**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[N/A]**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[N/A]**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[N/A]**
 - (b) Did you mention the license of the assets? **[N/A]**
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[N/A]**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**