
Active Learning for Single Neuron Models with Lipschitz Non-Linearities

Aarshvi Gajjar
New York University

Chinmay Hegde
New York University

Christopher Musco
New York University

Abstract

We consider the problem of active learning for single neuron models, also sometimes called “ridge functions”, in the agnostic setting (under adversarial label noise). Such models have been shown to be broadly effective in modeling physical phenomena, and for constructing surrogate data-driven models for partial differential equations.

Surprisingly, we show that for a single neuron model with any Lipschitz non-linearity (such as the ReLU, sigmoid, absolute value, low-degree polynomial, among others), strong provable approximation guarantees can be obtained using a well-known active learning strategy for fitting *linear functions* in the agnostic setting. Namely, we can collect samples via statistical *leverage score sampling*, which has been shown to be near-optimal in other active learning scenarios. We support our theoretical results with empirical simulations showing that our proposed active learning strategy based on leverage score sampling outperforms (ordinary) uniform sampling when fitting single neuron models.

1 INTRODUCTION

This paper considers active learning methods for functions of the form $g(\mathbf{x}) = f(\langle \mathbf{w}, \mathbf{x} \rangle)$, where \mathbf{w} is a weight vector and f is a non-linearity. For a given distribution \mathcal{D} on $\mathbb{R}^d \times \mathbb{R}$, a random vector (\mathbf{x}, y) sampled from \mathcal{D} , and scalar function $f : \mathbb{R} \rightarrow \mathbb{R}$, our goal is to find \mathbf{w} which minimizes the expected squared error:

$$\mathbb{E}_{\mathbf{x}, y \sim \mathcal{D}} (f(\langle \mathbf{w}, \mathbf{x} \rangle) - y)^2.$$

Functions of the form $f(\langle \mathbf{w}, \mathbf{x} \rangle)$ find applications in a variety of settings under various names: they are called “single

neuron” or “single index” models, “ridge functions”, and “plane waves” in different communities [Pinkus, 1997, 2015, Yehudai and Shamir, 2020, Rao et al., 2017, Candès, 2003]. Single neuron models are studied in machine learning theory as tractable examples of simple neural networks [Diakonikolas et al., 2020a, Goel et al., 2017]. Moreover, these models are known to be adept at modeling a variety of physical phenomena [Constantine et al., 2016] and for that reason are effective e.g., in building surrogate models for efficiently solving parametric partial differential equations (PDEs) and for approximating quantity of interest (QoI) surfaces for uncertainty quantification, model-driven design, and data assimilation [O’Leary-Roseberry et al., 2022, Constantine et al., 2017, Cohen et al., 2012, Le Maître and Knio, 2010, Lassila and Rozza, 2010, Binev et al., 2017].

In these applications, single neuron models are used to fit complex functions over \mathbb{R}^d based on queries from those functions. Often, the cost of obtaining a query (\mathbf{x}, y) from the target function dominates the computational cost of fitting the model: each training point collected requires numerically solving the PDE under consideration with parameters given by \mathbf{x} [Adcock et al., 2022a, Cohen and DeVore, 2015]. At the same time, we often have the freedom in exactly *how* the query is obtained; we are not restricted to simply sampling from \mathcal{D} , but rather can specify a target location \mathbf{x} and sample $y \sim \mathcal{D} \mid \mathbf{x}$ (or compute y deterministically, since in most applications it is a deterministic function of \mathbf{x}). Given these considerations, the focus of our work is on developing efficient *active learning* and experimental design methods¹ for fitting single neuron models using as few carefully chosen (\mathbf{x}, y) observations as possible.

We study this active learning problem in the challenging *agnostic learning* or adversarial noise setting. Again, this is motivated by applications of single neuron models in computational science. Typically, while it can be approximated by a single neuron model, the QoI or surrogate function under consideration is not itself of the form $f(\langle \mathbf{w}, \mathbf{x} \rangle)$. For this reason, the agnostic setting has become the standard

¹We use “experimental design” to refer to methods that collect samples in a non-adaptive way. In other words, a set of points $\mathbf{x}_1, \dots, \mathbf{x}_s$ are specified upfront and the corresponding y values are observed all at once. In contrast, in standard active learning methods, the choice of \mathbf{x}_j can depend on the response values of all prior points $\mathbf{x}_1, \dots, \mathbf{x}_{j-1}$.

in work on PDE models involving other common function families, like structured polynomials [Chkifa et al., 2018, Cohen and DeVore, 2015, Adcock et al., 2022b, Hampton and Doostan, 2015b]. In the agnostic setting, for a constant C (or more stringently, $C = 1 + \varepsilon$), our goal is always to return with high probability some $\tilde{\mathbf{w}}$ such that:

$$\begin{aligned} \mathbb{E}_{\mathbf{x}, y \sim \mathcal{D}} (f(\langle \tilde{\mathbf{w}}, \mathbf{x} \rangle) - y)^2 \\ \leq C \cdot \min_{\mathbf{w}} \mathbb{E}_{\mathbf{x}, y \sim \mathcal{D}} (f(\langle \mathbf{w}, \mathbf{x} \rangle) - y)^2. \end{aligned}$$

1.1 Our Contributions

For ease of exposition, we consider the case when \mathcal{D} is a uniform distribution over n points in \mathbb{R}^d . This is without loss of generality, since any continuous distribution can be approximated by the uniform distribution over a sufficient large, but finite, subset of points in \mathbb{R}^d .² In this case, we have the following problem statement.

Problem 1 (Single Neuron Regression). *Given a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ and query access to a target vector $\mathbf{y} \in \mathbb{R}^n$, for a given function $f : \mathbb{R} \rightarrow \mathbb{R}$, find $\mathbf{w} \in \mathbb{R}^d$ to minimize $\|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$ using as few queries from \mathbf{y} as possible.*

When f is an identity function, Problem 1 reduces to active least squares regression, which has received a lot of recent attention in computer science and machine learning. In the agnostic setting, state-of-the-art results can be obtained via “leverage score” sampling, also known as “coherence motivated” or “effective resistance” sampling [Avron et al., 2019, Cohen and Migliorati, 2017, Rauhut and Ward, 2012, Hampton and Doostan, 2015a]. The idea behind leverage scores sampling methods is to collect samples from \mathbf{y} randomly but non-uniformly, using an importance sampling distribution based on the rows of \mathbf{X} . More “unique” rows are selected with higher probability. Formally, rows are selected with probability proportional to their statistical leverage scores:

Definition 1 (Statistical Leverage Score). *The leverage score, $\tau_i(\mathbf{X})$ of the i^{th} row, \mathbf{x}_i of a matrix, $\mathbf{X} \in \mathbb{R}^{n \times d}$ is equal to:*

$$\tau_i(\mathbf{X}) = \mathbf{x}_i^T (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}_i = \max_{\mathbf{w} \in \mathbb{R}^d} \frac{[\mathbf{X}\mathbf{w}]_i^2}{\|\mathbf{X}\mathbf{w}\|_2^2}.$$

Here, $[\mathbf{X}\mathbf{w}]_i$ denotes the i^{th} entry of the vector $\mathbf{X}\mathbf{w}$.

We always have that $0 \leq \tau_i(\mathbf{X}) \leq 1$, and a well-known property of the statistical leverage scores is that $\sum_{i=1}^n \tau_i(\mathbf{X}) = \text{rank}(\mathbf{X}) \leq d$. The leverage score of a row is large (close to 1) if that row has large inner product

²For other function families (e.g. polynomials, or sparse Fourier functions) there has been recent work on active learning algorithms based on leverage score sampling that skip the discrete approximation step by developing algorithms directly tailored to common continuous distributions, like uniform or Gaussian [Erdélyi et al. [2020]]. We believe our techniques should be directly extendable to give comparable results for single neuron models.

with some vector in \mathbb{R}^d , as compared to that vector’s inner product with all other rows in the matrix \mathbf{X} . This means that the particular row enjoys significance in forming the row space of \mathbf{X} . For linear regression, it can be shown that when \mathbf{X} has d columns, leverage score sampling yields a sample complexity of $O(d \log d + d/\varepsilon)$ to find $\hat{\mathbf{w}}$ satisfying $\|\mathbf{X}\hat{\mathbf{w}} - \mathbf{y}\|_2^2 \leq (1 + \varepsilon) \min_{\mathbf{w}} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|_2^2$; moreover, this is optimal up to the log factor [Chen and Price, 2019].

Our main contribution is to establish that, surprisingly, when combined with a novel regularization strategy, leverage score sampling *also* yields theoretical guarantees for the more general case (Problem 1) for a broad class of non-linearities f . In fact, we only require that f is L -Lipschitz for constant L , a property that holds for most non-linearities used in practice (ReLU, sigmoid, absolute value, low-degree polynomials, etc.). We prove the following main result, which shows that $\tilde{O}(d^2/\varepsilon^4)$ samples, collected via leverage score sampling, suffice for provably learning a single neuron model with Lipschitz non-linearity.

Theorem 1 (Main Result). *Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ be a data matrix and $\mathbf{y} \in \mathbb{R}^n$ be a target vector. Let f be an L -Lipschitz non-linearity with $f(0) = 0$ and let $OPT = \min_{\mathbf{w}} \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$. There is an algorithm (Algorithm 1) that, based on the leverage scores of \mathbf{X} , observes $m = O\left(\frac{d^2 \log(1/\varepsilon)}{\varepsilon^4}\right)$ random entries from \mathbf{y} and returns with probability $> 9/10$ a vector $\hat{\mathbf{w}}$ satisfying:*

$$\|f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{y}\|_2^2 \leq C \cdot (OPT + \varepsilon L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2).$$

The assumption $f(0) = 0$ in Theorem 1 is without loss of generality. If $f(0)$ is non-zero, we can simply solve a transformed problem with $\mathbf{y}' = \mathbf{y} - f(0)$ and $f'(x) = f(x) - f(0)$. The theorem mirrors previous results in the linear setting, and in contrast to prior work on agnostically learning single neuron models, does not require any assumptions on \mathbf{X} [Diakonikolas et al., 2022b, Tyagi and Cevher, 2012]. In addition to multiplicative error C , the theorem has an additive error term of $C\varepsilon L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2$. An additive error term is necessary; as we will show in Section 5, it is provably impossible to achieve purely relative error with a number of samples polynomial in d . Similar additive error terms arise in related work on leverage score sampling for problems like logistic regression [Munteanu et al., 2018, Mai et al., 2021]. On the other hand, we believe the d^2 dependence in our bound is not necessary, and should be improvable to linear in d . The dependence on ε is also likely improvable.

In Section 4 we support our main theoretical result with an empirical evaluation on both synthetic data and several test problems that require approximating PDE quantity of interest surfaces. In all settings, leverage score sampling outperforms ordinary uniform sampling, often improving on the error obtained using a fixed number of samples by an order of magnitude or more.

1.2 Related Work

Single neuron models have been widely studied in a number of communities, including machine learning, computational science, and approximation theory. These models can be generalized to the “multi-index” case, where $g(\mathbf{x}) = f_1(\langle \mathbf{x}, \mathbf{w}_1 \rangle) + \dots + f_q(\langle \mathbf{x}, \mathbf{w}_1 \rangle)$ [Klusowski and Barron, 2018, Tyagi and Cevher, 2012, Candès, 2003] or to the case when f is not known in advance (but might be from a parameterized function family, such as low-degree polynomials) [Cohen et al., 2012, Hokanson and Constantine, 2018, Constantine et al., 2017]. While we do not address these generalizations in this paper, we hope that our work can provide a foundation for further work in this direction.

Beyond sample complexity, there has also been a lot of interest in understanding the computational complexity of fitting single neuron models, including in the agnostic setting [Yehudai and Shamir, 2020]. There are both known hardness results for general data distributions [Diakonikolas et al., 2020b, Goel et al., 2019, Diakonikolas et al., 2022a], as well as positive results on efficient algorithms under additional assumptions [Goel et al., 2017, Diakonikolas et al., 2022b, 2020a]. Our work differs from this setting in two ways: 1) we focus on sample complexity; 2) we make no assumptions on the data distribution \mathcal{D} (i.e., no assumptions on the data matrix \mathbf{X}); 3) we allow for active sampling strategies. Obtaining results under i.i.d. sampling, even for well behaved distributions, inherently requires additional assumptions, like $\mathbf{w}^* = \min_{\mathbf{w}} \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2$ being bounded in norm or \mathbf{X} having bounded condition number.

In computational science, there has been significant work on active learning and experimental design for fitting other classes of functions adept at modeling high dimensional physical phenomena [Chkifa et al., 2018, Cohen and DeVore, 2015, Adcock et al., 2022b, Hampton and Doostan, 2015b]. Most such results focus on minimizing squared prediction error. For situations involving model mismatch, the agnostic (adversarial noise) setting is more appropriate than assuming i.i.d. zero mean noise, which is more typical in classical statistical results on experimental design [Pukelsheim, 2006]. To obtain results in the agnostic setting for *linear* function families, much of the prior work in computational science uses “coherence motivated” sampling techniques [Cohen and Migliorati, 2017, Rauhut and Ward, 2012, Hampton and Doostan, 2015a]. Such methods are equivalent to leverage score sampling [Avron et al., 2019].

Leverage score sampling and related methods have found widespread applications in the design of efficient algorithms, including in the construction of graph sparsifiers, coresets, randomized low-rank approximations, and in solving over-constrained linear regression problems [Spielman and Srivastava, 2011, Feldman and Langberg, 2011, Drineas et al., 2008, Cohen et al., 2017, Musco and Musco, 2017, Dasgupta et al., 2008, Drineas et al., 2006, Cohen et al., 2015, Kacham and Woodruff, 2020].

Recently, there has been renewed attention on using leverage score sampling for solving the active linear regression problem for various loss functions [Erdélyi et al., 2020, Chen and Derezhinski, 2021, Meyer et al., 2023, Musco et al., 2022]. While all of these results are heavily tailored to linear models, a small body of works addresses the problem of active learning for non-linear problems. This includes problems of the form $\|f(\mathbf{X}\mathbf{w} - \mathbf{y})\|_2^2$, where f is Lipschitz [Munteanu et al., 2018, Mai et al., 2021]. While not equivalent to our problem, this formulation captures important tasks like logistic regression. Our work can be viewed as broadening the range of non-linear problems for which leverage score sampling yields natural worst-case guarantees.

Finally, we mention a few papers that directly address the active learning problem for functions of form $f(\mathbf{X}\mathbf{w})$. Cohen et al. [2012] studies adaptive query algorithms, but in a different setting than our paper. Specifically, they address the easier (noiseless) realizable setting, where samples are obtained from a function of the form $f(\mathbf{X}\mathbf{w}^*)$ for a ground truth \mathbf{w}^* . They also make stronger smoothness assumptions on f , although their algorithm can handle the case when f is not known in advance. Follow-up work also address the multi-index problem in the same setting [Fornasier et al., 2012]. Also motivated by applications in efficient PDE surrogate modeling, Tyagi and Cevher [2012] study the multi-index problem, but again in the realizable setting. Their techniques can handle mean centered i.i.d. noise, but not adversarial noise or model mismatch.

2 PRELIMINARIES

Notation. Throughout the paper, we use bold lower-case letters for vectors and bold upper-case letters for matrices. We let \mathbf{e}_i denotes the i^{th} standard basis vector (all zeros, but with a 1 in position i). The dimension of \mathbf{e}_i will be clear from context. For a vector $\mathbf{x} \in \mathbb{R}^n$ with entries x_1, \dots, x_n , $\|\mathbf{x}\|_2 = (\sum_{i=1}^n x_i^2)^{1/2}$ denotes the Euclidean norm. $\mathcal{B}^d(r)$ denotes a ball of radius r centered at 0, i.e. $\mathcal{B}^d(r) = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq r\}$. For a scalar function $f : \mathbb{R} \rightarrow \mathbb{R}$ and vector \mathbf{y} , we use $f(\mathbf{y})$ to denote the vector obtained by applying f to each entry of \mathbf{y} . For a fixed matrix \mathbf{X} , unobserved target vector \mathbf{y} , and non-linearity f , we denote $\|f(\mathbf{X}\mathbf{w}^*) - \mathbf{y}\|_2^2$ by OPT where $\mathbf{w}^* \in \arg \min_{\mathbf{w}} \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$.

Importance Sampling. As discussed, our approach is based on importance sampling according to statistical leverage scores: we fix a set of probabilities p_1, \dots, p_n , use them to sample m rows from the regression problem $\|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$, and solve a reweighted least squares problem to find a near optimal choice for \mathbf{w} . Formally, this can be implemented by defining a sampling matrices of the following form:

Definition 2 (Importance Sampling Matrix). *Let*

$\{p_1, \dots, p_n\} \in (0, 1]^n$ be a given set of probabilities (so that $\sum_i p_i = 1$). A matrix \mathbf{S} is an $m \times n$ importance sampling matrix if each of its rows is chosen to equal $\frac{1}{\sqrt{m \cdot p_i}} \cdot \mathbf{e}_i$ with probability proportional to p_i .

To compute an approximate to $\min \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$ we will solve an optimization problem involving the sub-sampled objective $\|\mathbf{S}f(\mathbf{X}\mathbf{w}) - \mathbf{S}\mathbf{y}\|_2^2$. It is easily verified that for any choice of p_1, \dots, p_n and any vector \mathbf{z} , $\mathbb{E} [\|\mathbf{S}\mathbf{z}\|_2^2] = \|\mathbf{z}\|_2^2$. We will use this fact repeatedly.

Properties of Leverage Scores. Our importance sampling mechanism is based on sampling by the leverage scores $\tau_1(\mathbf{X}), \dots, \tau_n(\mathbf{X})$ of the design matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$. For any full-rank $d \times d$ matrix \mathbf{R} , we have that $\tau_i(\mathbf{X}\mathbf{R}) = \tau_i(\mathbf{X})$. This is clear from Definition 1 and implies that $\tau_i(\mathbf{X})$ only depends on the column span of \mathbf{X} . In our proofs, this property will allow us to easily reduce to the setting where \mathbf{X} is assumed to be a matrix with orthonormal columns.

We will also use the following well-known fact about using leverage score sampling to construct a “subspace embedding” for a matrix \mathbf{X} .

Lemma 1 (Subspace Embedding (see e.g. Theorem 17 in Woodruff [2014]). *Given $\mathbf{X} \in \mathbb{R}^{n \times d}$ with leverage scores $\tau_1(\mathbf{X}), \dots, \tau_n(\mathbf{X})$, let $p_i = \tau_i(\mathbf{X}) / \sum_i \tau_i(\mathbf{X})$. Let $\mathbf{S} \in \mathbb{R}^{m \times n}$ be a sampling matrix constructed as in Definition 2 using the probabilities p_1, \dots, p_n . For any $0 < \gamma < 1$, as long as $m \geq c \cdot d \log(d/\delta)/\gamma^2$ for some fixed constant c , then with probability $1 - \delta$ we have that for all $\mathbf{w} \in \mathbb{R}^d$,*

$$(1 - \gamma)\|\mathbf{X}\mathbf{w}\|_2^2 \leq \|\mathbf{S}\mathbf{X}\mathbf{w}\|_2^2 \leq (1 + \gamma)\|\mathbf{X}\mathbf{w}\|_2^2.$$

Lemma 1 establishes that, with high probability, leverage score sampling preserves the norm of any vector $\mathbf{X}\mathbf{w}$ in the column span of \mathbf{X} . This guarantee can be proven using an argument that reduces to a matrix Chernoff bound [Spielman and Srivastava, 2011]. This is a critical component for previously known active learning guarantees for fitting linear functions using leverage score sampling [Sarlos, 2006].

3 MAIN RESULT

With preliminaries in place, we are ready to prove Theorem 1. We begin with pseudocode for Algorithm 1, which obtains the guarantee of the theorem via leverage score sampling combined with a novel regularization strategy.

The core of Algorithm 1 is the optimization problem:

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}: \|\mathbf{S}\mathbf{X}\mathbf{w}\|_2^2 \leq \frac{1}{\varepsilon \cdot L^2} \|\mathbf{S}(\mathbf{y})\|_2^2} \|\mathbf{S}f(\mathbf{X}\mathbf{w}) - \mathbf{S}\mathbf{y}\|_2^2. \quad (1)$$

Since this problem involves a Euclidean constraint on $\mathbf{S}\mathbf{X}\mathbf{w}$, it is notably different from the more standard (weighted) empirical risk minimization problem: $\min_{\mathbf{w}} \|\mathbf{S}f(\mathbf{X}\mathbf{w}) - \mathbf{S}\mathbf{y}\|_2^2$.

Algorithm 1 Leverage Score Based Active Learning for Single Neuron Models

input: Matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, L -Lipschitz non-linearity $f: \mathbb{R} \rightarrow \mathbb{R}$ with $f(0) = 0$, query access to target vector $\mathbf{y} \in \mathbb{R}^n$, number of samples m .

output: Approximate solution to $\min_{\mathbf{w}} \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$.

- 1: Compute $\tau_i(\mathbf{X})$ for all $i = 1, \dots, n$.
- 2: Set $p_i = \tau_i(\mathbf{X}) / \sum_{j=1}^n \tau_j(\mathbf{X})$ for all $i = 1, \dots, n$.
- 3: Construct $\mathbf{S} \in \mathbb{R}^{m \times n}$ according to Definition 2 with probabilities p_1, \dots, p_n .
- 4: Query \mathbf{y} at m locations to obtain $\mathbf{S}\mathbf{y} \in \mathbb{R}^m$.
- 5: Solve the minimization problem:

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}: \|\mathbf{S}\mathbf{X}\mathbf{w}\|_2^2 \leq \frac{1}{\varepsilon \cdot L^2} \|\mathbf{S}(\mathbf{y})\|_2^2} \|\mathbf{S}f(\mathbf{X}\mathbf{w}) - \mathbf{S}\mathbf{y}\|_2^2$$

- 6: **return** $\hat{\mathbf{w}}$

We believe that the norm constraint is necessary for getting acceptable upper bounds in the agnostic setting and cannot be eliminated. However, in our experiments (Section 4) we were able to safely ignore the constraint without hurting empirical performance. In any case, with or without constraint, minimizing (1) is a non-convex neuron fitting problem, and we do not attempt to theoretically analyze its computational complexity in this paper; however, it can be solved easily in practice using standard first-order optimization methods (such as gradient descent or its projected version).

As an first step to proving Theorem 1, we link the quality of the solution to (1) to that of the optimum regressor, $\mathbf{w}^* \in \arg \min_{\mathbf{w}} \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2$, as follows.

Claim 1. *Let $\hat{\mathbf{w}}$ be the vector returned by Algorithm 1 and let $OPT = \|f(\mathbf{X}\mathbf{w}^*) - \mathbf{y}\|_2^2$. With probability 49/50, for a fixed constant $C > 0$, we have*

$$\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}\mathbf{y}\|_2^2 \leq C \cdot (OPT + \varepsilon L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2).$$

Proof. Consider the case when $\|\mathbf{S}\mathbf{X}\mathbf{w}^*\|_2^2 \leq \frac{1}{\varepsilon L^2} \|\mathbf{S}\mathbf{y}\|_2^2$. Then \mathbf{w}^* satisfies the constraint of the above optimization problem so we have that

$$\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}\mathbf{y}\|_2^2 \leq \|\mathbf{S}f(\mathbf{X}\mathbf{w}^*) - \mathbf{S}\mathbf{y}\|_2^2 \leq C \cdot OPT.$$

The last inequality follows with probability 49/50 via Markov’s inequality since

$$\mathbb{E} [\|\mathbf{S}f(\mathbf{X}\mathbf{w}^*) - \mathbf{S}\mathbf{y}\|_2^2] = \|f(\mathbf{X}\mathbf{w}^*) - \mathbf{y}\|_2^2 = OPT.$$

On the other hand, consider the case where $\|\mathbf{S}\mathbf{X}\mathbf{w}^*\|_2^2 \geq \frac{1}{\varepsilon L^2} \|\mathbf{S}\mathbf{y}\|_2^2$. Then we have that $\|\mathbf{S}\mathbf{y}\|_2^2 \leq \varepsilon L^2 \cdot \|\mathbf{S}\mathbf{X}\mathbf{w}^*\|_2^2$. In this case, we can plug in the zero vector to the above minimization problem (since zero clearly satisfies the constraint) and conclude again that:

$$\begin{aligned} \|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}\mathbf{y}\|_2^2 &\leq \|\mathbf{S}f(\mathbf{X}\mathbf{0}) - \mathbf{S}\mathbf{y}\|_2^2 = \|\mathbf{S}\mathbf{y}\|_2^2 \\ &\leq \varepsilon L^2 \|\mathbf{S}\mathbf{X}\mathbf{w}^*\|_2^2 \leq 2\varepsilon L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2. \end{aligned}$$

The last inequality follows from the subspace embedding inequality from Lemma 1. Note that the constraint of $f(\mathbf{0}) = 0$ is used above as $f(\mathbf{X}\mathbf{0}) = f(\mathbf{0}) = \mathbf{0}$. \square

3.1 Concentration Bounds

Claim 1 upper bounds the error of $\hat{\mathbf{w}}$ in solving the *subsampled* regression problem $\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}\mathbf{y}\|_2^2$. To show that it also provides a good solution for the original problem $\min_{\mathbf{w}} \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$ we require several concentration results that are a consequence of leverage score sampling. These results are similar to Lemma 1, except that they show that sampling with \mathbf{S} also preserves the norm of vectors obtained via *non-linear* transformations of the form $f(\mathbf{X}\mathbf{w})$.

The first bound gives a guarantee on preserving the distance between two fixed vectors, $f(\mathbf{X}\mathbf{w}_1) - f(\mathbf{X}\mathbf{w}_2)$. In contrast to the relative error subspace embedding of Lemma 1, the bound involved an additive error term; this extra term is likely unavoidable in the most general case of Lipschitz f .

Lemma 2. *Let f and \mathbf{X} be as in Theorem 1. Let $\mathbf{S} \in \mathbb{R}^{m \times n}$ be an importance sampling matrix chosen with probabilities $\{p_1, \dots, p_n\}$, where $p_i = \tau_i(\mathbf{X}) / \sum_{i=1}^n \tau_i(\mathbf{X})$. As long as $m \geq \frac{3d \log(2/\delta)}{\varepsilon^2}$, then with probability $\geq 1 - \delta$, for any fixed pair of vectors $\mathbf{w}_1, \mathbf{w}_2 \in \mathbb{R}^d$, we have:*

$$\begin{aligned} & \|f(\mathbf{X}\mathbf{w}_1) - f(\mathbf{X}\mathbf{w}_2)\|_2^2 - \varepsilon L^2 \|\mathbf{X}\mathbf{w}_1 - \mathbf{X}\mathbf{w}_2\|_2^2 \\ & \leq \|\mathbf{S}f(\mathbf{X}\mathbf{w}_1) - \mathbf{S}f(\mathbf{X}\mathbf{w}_2)\|_2^2 \\ & \leq \|f(\mathbf{X}\mathbf{w}_1) - f(\mathbf{X}\mathbf{w}_2)\|_2^2 + \varepsilon L^2 \|\mathbf{X}\mathbf{w}_1 - \mathbf{X}\mathbf{w}_2\|_2^2. \end{aligned}$$

By combining Lemma 2 with an ε -net argument, we can extend the bound to obtain a one-sided guarantee that involves the distance between $f(\mathbf{X}\mathbf{w}^*)$ and $f(\mathbf{X}\mathbf{w})$ for *all* \mathbf{w} within a ball of radius R .

Lemma 3. *Let f , \mathbf{X} , and \mathbf{y} be as in Theorem 1. Let $\mathbf{w}^* = \arg \min_{\mathbf{w}} \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$ and let R be a fixed radius. Let $\mathbf{S} \in \mathbb{R}^{m \times n}$ be an importance sampling matrix as in Lemma 2. As long as $m \geq c \frac{d^2 \log(1/\varepsilon)}{\varepsilon^2}$ for $\varepsilon < 1$ and fixed constant c , then with probability $49/50$, for all $\hat{\mathbf{w}} \in \mathcal{B}^d(R)$,*

$$\begin{aligned} \|f(\mathbf{X}\hat{\mathbf{w}}) - f(\mathbf{X}\mathbf{w}^*)\|_2^2 & \leq 4 \cdot \|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}f(\mathbf{X}\mathbf{w}^*)\|_2^2 \\ & \quad + \varepsilon^2 L^2 R^2 + \varepsilon^2 L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2. \end{aligned}$$

Proofs of Lemma 1 and 3 are deferred to Appendix 5.

3.2 Proof of Main Result

With Claim 1 and Lemma 3 in place, we are now ready to prove our main result.

Proof of Theorem 1. First note that, without loss of generality, we can assume that \mathbf{X} has orthonormal columns. In particular, if \mathbf{X} is not orthonormal, we can write it as $\mathbf{X} = \mathbf{Q}\mathbf{R}$ where $\mathbf{Q} \in \mathbb{R}^{n \times \text{rank}(\mathbf{X})}$ has orthonormal columns and \mathbf{R}

is a square full-rank matrix. The leverage scores of \mathbf{Q} are equal to those of \mathbf{X} . Moreover, any solution $\hat{\mathbf{w}}$ to (1) has a corresponding solution $\mathbf{R}\hat{\mathbf{w}}$ to the minimization problem if \mathbf{X} were replaced by \mathbf{Q} . So solving the above problem is equivalent to first explicitly orthogonalizing \mathbf{X} and solving the same problem.

Next, we use the elementary fact that for any vectors \mathbf{a} and \mathbf{b} , $\|\mathbf{a} + \mathbf{b}\|_2^2 \leq 2\|\mathbf{a}\|_2^2 + 2\|\mathbf{b}\|_2^2$. This give the bound:

$$\begin{aligned} & \|f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{y}\|_2^2 \\ & \leq 2\|f(\mathbf{X}\hat{\mathbf{w}}) - f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 2\|f(\mathbf{X}\mathbf{w}^*) - \mathbf{y}\|_2^2 \\ & \leq 2\|f(\mathbf{X}\hat{\mathbf{w}}) - f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 2 \cdot \text{OPT}. \end{aligned} \quad (2)$$

We need to bound the first term. To do so, we first observe that, thanks to the constraint imposed in (1), the norm of $\hat{\mathbf{w}}$ can be bounded, which allows us to apply Lemma 3. In particular, we claim that with probability $49/50$,

$$\|\hat{\mathbf{w}}\|_2^2 \leq \frac{100}{\varepsilon L^2} \cdot \|\mathbf{y}\|_2^2. \quad (3)$$

To see that this is the case, note that under our assumption that \mathbf{X} is orthogonal, we have $\|\hat{\mathbf{w}}\|_2^2 = \|\mathbf{X}\hat{\mathbf{w}}\|_2^2$. We can bound $\|\mathbf{X}\hat{\mathbf{w}}\|_2^2$ as follows:

$$\begin{aligned} \|\mathbf{X}\hat{\mathbf{w}}\|_2^2 & \leq 2\|\mathbf{S}\mathbf{X}\hat{\mathbf{w}}\|_2^2 \quad (\text{Lemma 1}) \\ & \leq 2 \frac{1}{\varepsilon \cdot L^2} \|\mathbf{S}\mathbf{y}\|_2^2 \quad (\text{From the constraint in (1)}) \\ & \leq \frac{100}{\varepsilon \cdot L^2} \|\mathbf{y}\|_2^2 \quad (\text{Markov's inequality}). \end{aligned}$$

In the last inequality, we used that $\mathbb{E}[\|\mathbf{S}\mathbf{y}\|_2^2] = \|\mathbf{y}\|_2^2$.

Since $\hat{\mathbf{w}}$ lies in $\mathcal{B}^d(R)$, where $R^2 = \frac{100}{\varepsilon L^2} \cdot \|\mathbf{y}\|_2^2$, we can apply Lemma 3 and conclude that, as long as $m \geq c \frac{d^2 \log(1/\varepsilon)}{\varepsilon^2}$:

$$\begin{aligned} & \|f(\mathbf{X}\hat{\mathbf{w}}) - f(\mathbf{X}\mathbf{w}^*)\|_2^2 \\ & \leq 4\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 100\varepsilon\|\mathbf{y}\|_2^2 \\ & \quad + \varepsilon^2 L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2 \\ & \leq 8\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}\mathbf{y}\|_2^2 + 8\|\mathbf{S}f(\mathbf{X}\mathbf{w}^*) - \mathbf{S}\mathbf{y}\|_2^2 \\ & \quad + 100\varepsilon\|\mathbf{y}\|_2^2 + \varepsilon^2 L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2 \\ & \leq 8\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}\mathbf{y}\|_2^2 + C \cdot \text{OPT} \\ & \quad + 100\varepsilon\|\mathbf{y}\|_2^2 + \varepsilon^2 L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2. \end{aligned}$$

As in the proof of Claim 1, the last inequality follows with probability $49/50$ via Markov's inequality since $\mathbb{E}[\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}\mathbf{y}\|_2^2] = \|f(\mathbf{X}\mathbf{w}^*) - \mathbf{y}\|_2^2 = \text{OPT}$.

Next, we apply Claim 1 to bound $\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}\mathbf{y}\|_2^2 \leq O(\text{OPT} + \varepsilon L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2)$. So overall, we conclude that for a constant C ,

$$\|f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{y}\|_2^2 \leq C \cdot (\text{OPT} + \varepsilon L^2 \|\mathbf{X}\mathbf{w}^*\|_2^2 + \varepsilon \|\mathbf{y}\|_2^2). \quad (4)$$

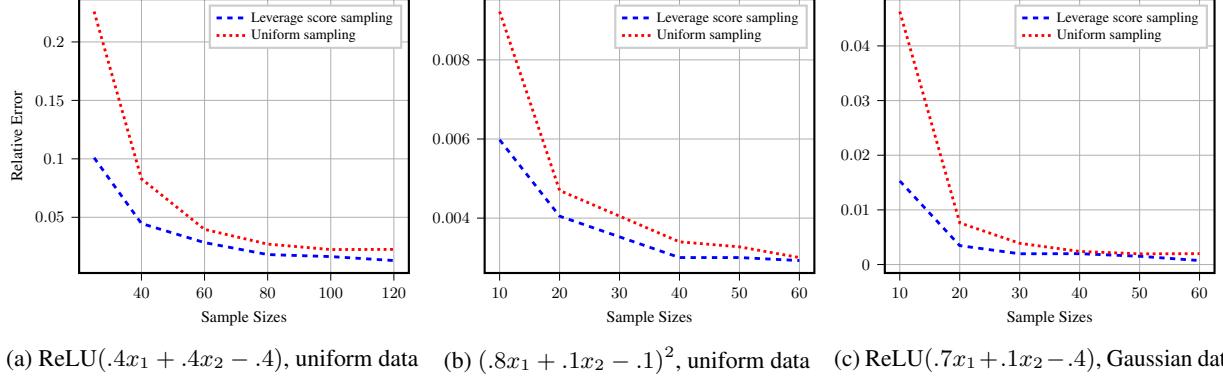


Figure 1: Median relative error for learning two-dimensional single neuron models involving a ReLU non-linearity for a synthetic \mathbf{X} with rows selected from a uniform or Gaussian distribution. The target vector \mathbf{y} was obtained by corrupting the ground truth with Gaussian noise with variance .05. In all cases our active leverage score sampling method outperforms naive uniform sampling. As expected, the improvement is more significant when a small number of samples are taken.

By triangle inequality, we have that

$$\begin{aligned}\|\mathbf{y}\|_2^2 &\leq 2OPT + 2\|f(\mathbf{X}\mathbf{w}^*)\|_2^2 \\ &\leq 2OPT + 2L^2\|\mathbf{X}\mathbf{w}^*\|_2^2.\end{aligned}$$

Using this fact, plugging (4) into (2), and rearranging terms yields the stated main result with probability. Union bounding overall all events assumed to hold in the proof, the result holds with probability $> 45/50 = 9/10$. \square

4 EXPERIMENTAL RESULTS

To complement our theoretical analysis, we also provide experimental results showing the promise of leverage score sampling for actively learning single neuron models. We consider both synthetic data problems, as well as several tests derived from differential equation approximation problems. We focus on sample efficiency – i.e., how many samples from \mathbf{y} are required to obtain a good approximation to $\min_{\mathbf{w}} \|f(\mathbf{X}\mathbf{w}) - \mathbf{y}\|_2^2$. Computational efficiency is not a major concern: as discussed, in typical applications of single-neuron learning in computational science, collecting samples requires numerically solving a differential equation, which dominates any runtime cost of the actual fitting procedure [Adcock et al., 2022a]. Moreover, leverage score sampling has already proven an efficient active learning tool for linear function classes [Cohen and DeVore, 2015].

Overall, for all problems tested, our experiments show that leverage score sampling obtains a better sample/accuracy trade-off than the standard approach of choosing sample points uniformly at random from \mathbf{X} .

Synthetic Data. For the synthetic data problems, we set \mathbf{X} to contain 10^5 random vectors drawn from either a two dimensional Gaussian distribution (“Gaussian data”), or uniformly from the two-dimensional box $[-1, 1]^2$ (“uniform data”). We also add a column of all 1’s to \mathbf{X} , which corresponds to including a bias term in the single-neuron model.

We select a ground truth \mathbf{w}^* , and create $\mathbf{y} = f(\mathbf{X}\mathbf{w}^*) + \mathbf{g}$, where \mathbf{g} is a vector of mean-centered Gaussian noise and f is the ReLU non-linearity. We then compute $\hat{\mathbf{w}}$ by subsampling data via leverage scores (as in Algorithm 1) and minimizing $\|\mathbf{S}f(\mathbf{X}\mathbf{w}) - \mathbf{S}\mathbf{y}\|_2^2$ over our subsampled data.

In our experiments we found that the constraint in Eq. 1 could be dropped without hurting the performance of leverage score sampling. For these low-dimensional synthetic problems, we simply used brute force search to optimize weights to ensure that a true minimum was found. We then run 100 trials each for various subsample sizes, and report median relative error: $\|f(\mathbf{X}\mathbf{w}^*) - \mathbf{y}\|_2^2/\|\mathbf{y}\|_2^2$. As shown in Figure 1, leverage scores sampling outperforms uniform sampling, especially for a relatively small number of samples. As expected, for a large number of samples, both methods eventually perform comparably, as both will obtain a $\hat{\mathbf{w}}$ very close to the optimal \mathbf{w}^* .

Test Problems. We consider three test problems involving the approximation of various Quantities of Interest (QoI’s) for three parametric differential equations: a damped harmonic oscillator, the heat equation, and the steady viscous Burger’s equation.

Test 1. We first consider a second order ODE modeling a damped harmonic oscillator with a sinusoidal force applied, which corresponds to the parametric differential equation:

$$\begin{aligned}\frac{d^2x}{dt^2}(t) + c \cdot \frac{dx}{dt}(t) + k \cdot x(t) &= f \cdot \cos(\omega t); \\ x(0) = x_0, \quad \frac{dy}{dt}(0) &= x_1.\end{aligned}$$

Here, (x, t) is the oscillator’s space and time coordinates, and c, k, f, ω are parameters. The choice of parameters significantly impact the final solution; for example, if the frequency term ω is close to the resonant frequency of the oscillator, we expect the driving force to lead to large oscillations. We consider as our QoI the maximum oscillator

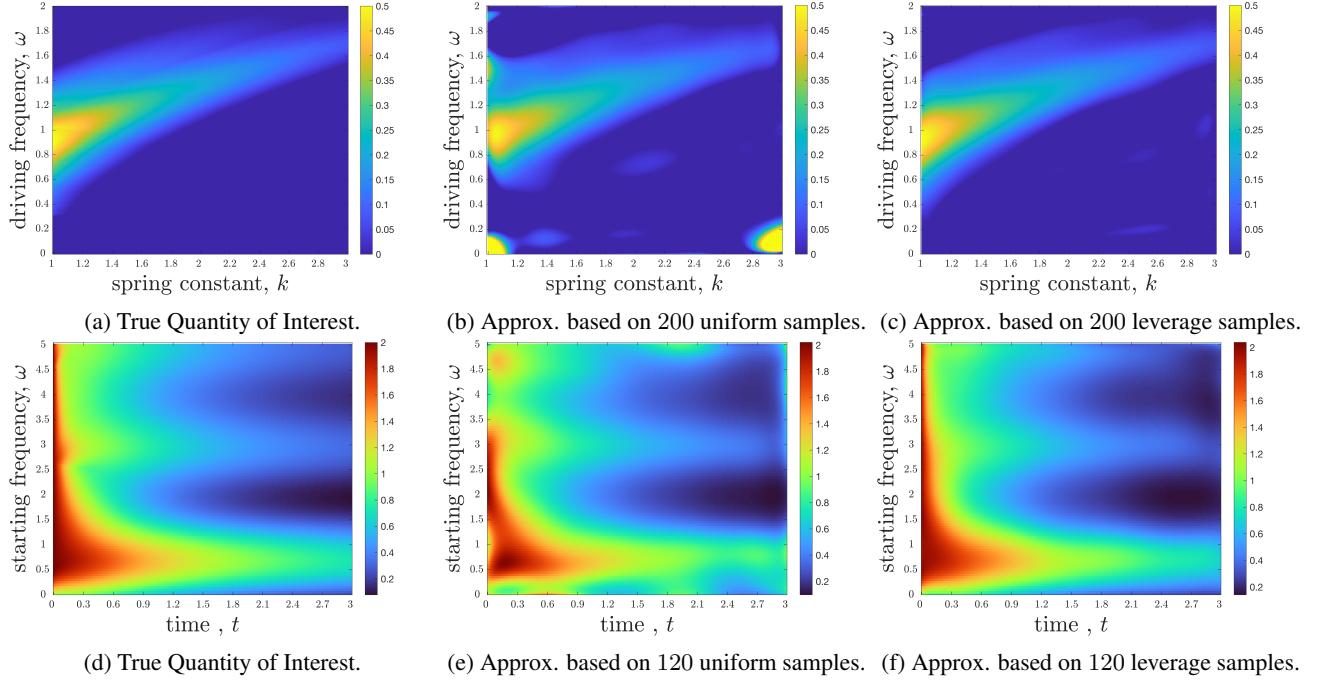


Figure 2: The top 3 images represent a plot of single neuron model fit to the maximum displacement QoI for a parametric ODE modeling a driven harmonic oscillator; and, the bottom 3 images represent the fit of the maximum temperature QoI for the heat equation PDE with sinusoidal initial condition. Evidently, leverage score sampling provides a better fit.

displacement after 20 seconds, and the goal is to estimate this value for all k and ω in the rectangle $\mathcal{U} = [1, 3] \times [0, 2]$.

We choose to approximate the QoI (which is always positive) with a function of the form $\text{ReLU}(p(k, \omega))$, where p is bivariate polynomial with total degree $q = 9$. This is accomplished by setting \mathbf{X} to be a Vandermonde matrix of Legendre polynomials evaluated at a grid of values on $[1, 3] \times [0, 2]$. \mathbf{X} has $55 = (q+1)(q+2)/2$ columns, which is the total number of terms in a degree q bivariate polynomial. Each row in \mathbf{X} corresponds to a different choice of parameters k, ω . We fit our single neuron model to the QoI using gradient descent with a standard adaptive step-size, again dropping the constraint in (1). As shown in the top three images of Figure 2, for a fixed number of samples, leverage score sampling leads to a visually better fit than uniform sampling. Quantitatively, we see in Figure 4 that leverage score sampling gives almost an order of magnitude lower error across a wide range of sampling numbers.

In Figure 3, we visualize how, for this problem, uniform samples differ from those collected using leverage scores. The Vandermonde matrix \mathbf{X} has higher leverage score for rows corresponding to points near the boundary of $[1, 3] \times [0, 2]$, so more samples are taken for (k, ω) values near the boundary. The benefits of sampling near the boundary are well-known for fitting simple polynomials [Cohen and Migliorati, 2017]. It is interesting that these benefits remain when the polynomial is combined with a non-linearity.

Test 2. We consider the 1-dimensional heat equation for values of $x \in [0, 1]$ with a time-dependent boundary equation and sinusoidal initial condition. This is modeled by the partial differential equation:

$$\begin{aligned} \pi \frac{\partial u}{\partial t} &= \frac{\partial^2 u}{\partial x^2}, \quad \mu(0, t) = 0, \quad \mu(x, 0) = \sin(\omega \pi x) \\ \pi e^{-t} + \frac{\partial u(1, t)}{\partial t} &= 0 \end{aligned}$$

As our QoI we consider the maximum temperate over all values of x for times $t \in [0, 3]$ and frequencies $\omega \in [0, 5]$. This leads to a highly varied QoI surface, which we again choose to fit with a model of the form $f(p(t, \omega))$. We let p be

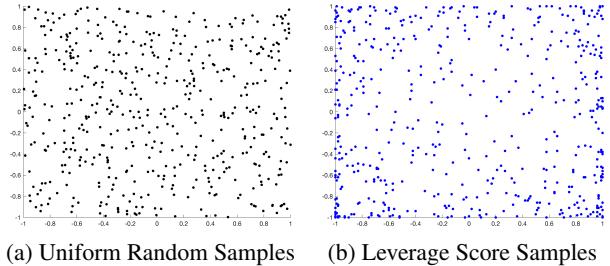


Figure 3: The plots visualize uniform vs. leverage score sampling for selecting example parameter values from the box $[1, 3] \times [0, 2]$ for fitting the QoI for Test Problem 1. Our leverage score method tends to sample more heavily near the perimeter of the box to fit the single neuron model.

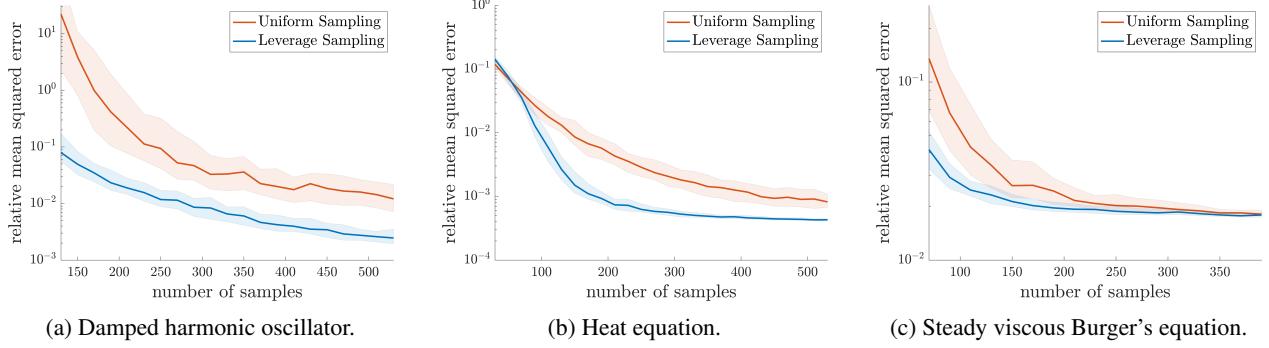


Figure 4: Sample complexity vs. relative error $\|f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{y}\|_2^2 / \|\mathbf{y}\|_2^2$ for fitting the QoI's specified for our three test problems. All experiments were run for 100 trials per sample value and we plot the median error and interquartile range.

a degree $q = 11$ bivariate polynomial, so \mathbf{X} has 78 columns. For this problem we choose $f(a) = e^a$ to be the exponential function; otherwise the experimental setup is identical to Test Problem 1. Despite the fact that this non-linearity is *not Lipschitz*, we again see visually better performances of leverage score sampling for a fixed number of samples in the bottom three plots Figure 2, and quantitatively better error in Figure 4. This result suggests our leverage score based active learning method may be robust to non-linearities that are just “locally” instead of globally Lipschitz.

Test 3. Finally, we consider steady state viscous Burger’s equation given by the following PDE:

$$u \cdot \frac{du}{dx} = \nu \cdot \frac{d^2u}{dx^2}, \quad u(a) = \alpha, u(b) = \beta,$$

where $u(x)$ is defined over the interval $x \in [a, b]$, $\nu > 0$ is the viscosity parameter, and α and β are boundary parameters. We consider the point at which the solution changes its sign as the quantity of interest. It is experimentally known that this QoI is particularly sensitive to the choice of α and β and not to the viscosity. Therefore, we fix $\nu = 0.1$, $[a, b] = [-1, 1]$, and vary $\alpha, \beta \in [0.8, 1.2] \times [-1.2, -0.8]$. We subtract the QoI obtained for these parameters by the minimum value to ensure that the function is always positive and again fit with a single neuron model of the form $\text{ReLU}(p(\alpha, \beta))$, where p has total degree 7. Results in Figure 4 align with the previous test problems: leverage score sampling shows a clear improvement over uniform sampling. For this problem, the improvement was less significant for a larger number of samples, suggesting that the model was simple enough that both the uniform and leverage score methods were able to eventually obtain a near-optimal fit.

5 DISCUSSION AND FUTURE WORK

We believe our main theoretical result can be improved in a number of ways. Most importantly, an ideal result would obtain a near-linear dependence on d instead of a dependence on d^2 , mirroring the $O(d \log d)$ sample complexity obtained

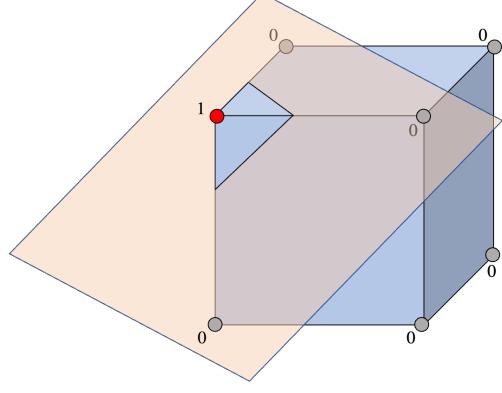


Figure 5: Hard instance for obtaining relative error.

by leverage score sampling for the active linear regression problem. The d^2 dependence is an inherent artifact of our ε -net analysis; possible approaches to improve this include appealing to a more careful net construction, as in Musco et al. [2022], or more directly reducing to matrix concentration, as was done in recent work to obtain a near-linear dependence for a related problem involving ℓ_1 embeddings of vectors transformed by Lipschitz non-linearities [Mai et al., 2021].

One might also hope to improve the error bound of Theorem 1. For example, it would be ideal to obtain a pure relative error bound of the form $\|f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{y}\|_2^2 \leq C \cdot OPT$. Unfortunately, we can argue that this is not possible without taking a number of samples exponential in d . Consider a $(d+1)$ data matrix \mathbf{X} whose first d columns contain all of the 2^d vertices of the d dimensional hypercube (i.e., there is a row containing every binary vector of length d). Let the last column of \mathbf{X} be the all-ones vector. Consider the 1-Lipschitz non-linearity $f(a) = \text{ReLU}(a)$ and let $\mathbf{y} = f(\mathbf{X}\mathbf{w})$ for some ground truth \mathbf{w} , in which case a pure relative error guarantee requires exactly recovering \mathbf{y} (since $OPT = 0$). As visualized in Figure 5, since there is a hyperplane separating any vertex in the hypercube from all other vertices, for any i it is possible to find a \mathbf{w}_i such that $\text{ReLU}(\mathbf{X}\mathbf{w}_i)$

evaluates to 1 in its i^{th} coordinate, and 0 everywhere else. Without observing at least $\Omega(2^d)$ entries from \mathbf{y} , we cannot distinguish between the case when $\mathbf{y} = \text{ReLU}(\mathbf{X}\mathbf{w}_i)$ for a randomly chosen i or $\mathbf{y} = \text{ReLU}(\mathbf{X}\mathbf{0})$.

Finally, we note that a major open direction for future research is to obtain provable active learning methods in the agnostic setting for the more challenging multi-index model, or in the case when f is not known in advance (and must be learned as part of the training process). We have some preliminary progress for the case where f is unknown, but defer a full discussion to future work.

Acknowledgements

This work was supported by the United States Department of Energy grant DE-SC0022266.

References

Ben Adcock, Simone Brugiapaglia, Nick Dexter, and Sebastian Morage. Deep neural networks are effective at learning high-dimensional hilbert-valued functions from limited data. In *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, volume 145 of *Proceedings of Machine Learning Research*, pages 1–36, 2022a.

Ben Adcock, Juan M. Cardenas, Nick Dexter, and Sebastian Morage. *Towards Optimal Sampling for Learning Sparse Approximations in High Dimensions*, pages 9–77. Springer International Publishing, 2022b.

Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir Zandieh. A universal sampling method for reconstructing signals with simple fourier transforms. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC)*, 2019.

Peter Binev, Albert Cohen, Wolfgang Dahmen, Ronald DeVore, Guergana Petrova, and Przemyslaw Wojtaszczyk. Data assimilation in reduced modeling. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):1–29, 2017.

Emmanuel J. Candès. Ridgelets: estimating with ridge functions. *The Annals of Statistics*, 31(5):1561–1599, 2003.

Xue Chen and Michal Derezhinski. Query complexity of least absolute deviation regression via robust uniform convergence. In *Proceedings of the 34th Annual Conference on Computational Learning Theory (COLT)*, volume 134, pages 1144–1179, 2021.

Xue Chen and Eric Price. Active regression via linear-sample sparsification active regression via linear-sample sparsification. In *Proceedings of the 32nd Annual Conference on Computational Learning Theory (COLT)*, 2019.

Abdellah Chkifa, Nick Dexter, Hoang Tran, and Clayton G. Webster. Polynomial approximation via compressed sensing of high-dimensional functions on lower sets. *Math. Comp.*, 87(311):1415–1450, 2018.

Albert Cohen and Ronald DeVore. Approximation of high-dimensional parametric PDEs. *Acta Numerica*, 24:1, 2015.

Albert Cohen and Giovanni Migliorati. Optimal weighted least-squares methods. *SMAI Journal of Computational Mathematics*, 3:181–203, 2017.

Albert Cohen, Ingrid Daubechies, Ronald DeVore, Gerard Kerkyacharian, and Dominique Picard. Capturing ridge functions in high dimensions from point queries. *Constructive Approximation*, 35(2):225–243, 2012.

Michael B. Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. Uniform sampling for matrix approximation. In *Proceedings of the 6th Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 181–190, 2015.

Michael B. Cohen, Cameron Musco, and Christopher Musco. Input sparsity time low-rank approximation via ridge leverage score sampling. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1758–1777, 2017.

Paul G. Constantine, Zachary del Rosario, and Gianluca Iaccarino. Many physical laws are ridge functions. *arXiv:1605.07974*, 2016.

Paul G. Constantine, Armin Eftekhari, Jeffrey Hokanson, and Rachel A. Ward. A near-stationary subspace for ridge approximation. *Computer Methods in Applied Mechanics and Engineering*, 326:402–421, 2017.

Anirban Dasgupta, Petros Drineas, Boulos Harb, Ravi Kumar, and Michael W. Mahoney. Sampling algorithms and coresets for l_p regression. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 932–941, 2008.

Ilias Diakonikolas, Surbhi Goel, Sushrut Karmalkar, Adam R. Klivans, and Mahdi Soltanolkotabi. Approximation schemes for relu regression. In *Proceedings of the 33rd Annual Conference on Computational Learning Theory (COLT)*, volume 125, pages 1452–1485, 2020a.

Ilias Diakonikolas, Daniel Kane, and Nikos Zarifis. Near-optimal sq lower bounds for agnostically learning halfspaces and relus under gaussian marginals. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 13586–13596, 2020b.

Ilias Diakonikolas, Daniel Kane, Pasin Manurangsi, and Lisheng Ren. Hardness of learning a single neuron with adversarial label noise. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 151, pages 8199–8213, 2022a.

Ilias Diakonikolas, Vasilis Kontonis, Christos Tzamos, and Nikos Zarifis. Learning a single neuron with adversarial label noise via gradient descent. In *Proceedings of the 35th Annual Conference on Computational Learning Theory (COLT)*, volume 178, pages 4313–4361, 2022b.

Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Sampling algorithms for ℓ_2 regression and applications. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1127–1136, 2006.

Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relative-error CUR matrix decompositions. *SIAM J. Matrix Anal. Appl.*, 30(2):844–881, 2008.

Tamás Erdélyi, Cameron Musco, and Christopher Musco. Fourier sparse leverage scores and approximate kernel learning. *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 2020.

Dan Feldman and Michael Langberg. A unified framework for approximating and clustering data. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 569–578, 2011.

Massimo Fornasier, Karin Schnass, and Jan Vybral. Learning functions of few arbitrary linear parameters in high dimensions. *Foundations of Computational Mathematics*, 12(2):229–262, 2012.

Surbhi Goel, Varun Kanade, Adam Klivans, and Justin Thaler. Reliably learning the relu in polynomial time. In *Proceedings of the 30th Annual Conference on Computational Learning Theory (COLT)*, volume 65, pages 1004–1042, 2017.

Surbhi Goel, Sushrut Karmalkar, and Adam Klivans. Time/accuracy tradeoffs for learning a relu with respect to gaussian marginals. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, 2019.

Jerrad Hampton and Alireza Doostan. Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression. *Comput. Method. Appl. M.*, 290: 73–97, 2015a.

Jerrad Hampton and Alireza Doostan. Compressive sampling of polynomial chaos expansions: Convergence analysis and sampling strategies. *Journal of Computational Physics*, 280:363–386, 2015b.

Jeffrey M. Hokanson and Paul G. Constantine. Data-driven polynomial ridge approximation using variable projection. *SIAM Journal on Scientific Computing*, 40(3), 2018.

Praneeth Kacham and David Woodruff. Optimal deterministic coresets for ridge regression. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 4141–4150. PMLR, 26–28 Aug 2020. URL <https://proceedings.mlr.press/v108/kacham20a.html>.

Jason M. Klusowski and Andrew R. Barron. Approximation by combinations of relu and squared relu ridge functions with ℓ^1 and ℓ^0 controls. *IEEE Transactions on Information Theory*, 64(12):7649–7656, 2018.

Toni Lassila and Gianluigi Rozza. Parametric free-form shape design with PDE models and reduced basis method. *Computer Methods in Applied Mechanics and Engineering*, 199(23):1583–1592, 2010.

Olivier P. Le Maître and Omar M. Knio. *Spectral methods for uncertainty quantification : with applications to computational fluid dynamics*. Scientific computation. Springer Netherlands, Dordrecht, New York, 2010.

Tung Mai, Anup B. Rao, and Cameron Musco. Coresets for classification – simplified and strengthened. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, 2021.

Raphael Meyer, Cameron Musco, Christopher Musco, and Samson Zhou David P. Woodruff. Near-linear sample complexity for l_p polynomial regression. In *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2023.

Alexander Munteanu, Chris Schwiegelshohn, Christian Sohler, and David Woodruff. On coresets for logistic regression. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, volume 31, 2018.

Cameron Musco and Christopher Musco. Recursive sampling for the Nyström method. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, pages 3833–3845, 2017.

Cameron Musco, Christopher Musco, David P. Woodruff, and Taisuke Yasuda. Active linear regression for ℓ_p norms and beyond. In *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2022.

Thomas O’Leary-Roseberry, Umberto Villa, Peng Chen, and Omar Ghattas. Derivative-informed projected neural networks for high-dimensional parametric maps governed by pdes. *Computer Methods in Applied Mechanics and Engineering*, 388, 2022.

Allan Pinkus. Approximating by ridge functions. *Surface fitting and multiresolution methods*, pages 279–292, 1997.

Allan Pinkus. *Ridge functions*, volume 205. Cambridge University Press, 2015.

Friedrich Pukelsheim. *Optimal Design of Experiments*. Society for Industrial and Applied Mathematics, 2006.

Nikhil Rao, Ravi Ganti, Laura Balzano, Rebecca Willett, and Robert Nowak. On learning high-dimensional structured single index models. In *Proceedings of the AAAI Conference on Artificial (AAAI)*, 2017.

Holger Rauhut and Rachel Ward. Sparse Legendre expansions via ℓ_1 -minimization. *Journal of Approximation Theory*, 164(5):517 – 533, 2012.

Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 143–152, 2006.

Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011. Preliminary version in the 40th Annual ACM Symposium on Theory of Computing (STOC).

Hemant Tyagi and Volkan Cevher. Active learning of multi-index function models. In *Advances in Neural Information Processing Systems 25 (NeurIPS)*, pages 1466–1474, 2012.

Roman Vershynin. *Introduction to the non-asymptotic analysis of random matrices*. Cambridge University Press, 2012.

David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1–2):1–157, 2014.

Gilad Yehudai and Ohad Shamir. Learning a single neuron with gradient methods. In *Proceedings of the 33rd Annual Conference on Computational Learning Theory (COLT)*, volume 125, pages 3756–3786, 2020.

In this section we present the detailed proof of Lemma 2 and Lemma 3.

Proof of Lemma 2. Let \mathbf{x}_i denote the i^{th} row of \mathbf{X} and let $\mathbf{u} = f(\mathbf{X}\mathbf{w}_1) - f(\mathbf{X}\mathbf{w}_2)$ and $\mathbf{v} = \mathbf{X}\mathbf{w}_1 - \mathbf{X}\mathbf{w}_2$. Since f is L -Lipschitz, for every $i \in [n]$, we have that:

$$\begin{aligned} u_i &= |f(\langle \mathbf{x}_i, \mathbf{w}_1 \rangle) - f(\langle \mathbf{x}_i, \mathbf{w}_2 \rangle)|_i \\ &\leq L \cdot |\langle \mathbf{x}_i, \mathbf{w}_1 \rangle - \langle \mathbf{x}_i, \mathbf{w}_2 \rangle|_i \leq L\mathbf{v}_i. \end{aligned} \quad (5)$$

Let $j_i \in [n]$ be the index of the row from \mathbf{X} selected by the i^{th} row in \mathbf{S} . We have that $\|\mathbf{S}\mathbf{u}\|_2^2 = \sum_{i=1}^m \frac{u_{j_i}^2}{m \cdot p_{j_i}}$, where $p_{j_i} = \tau_{j_i}(\mathbf{X}) / \text{rank}(\mathbf{X})$. We thus have that $\mathbb{E} \|\mathbf{S}\mathbf{u}\|_2^2 = \|\mathbf{u}\|_2^2$. Moreover, we can bound the variance in each term of the sum. In particular, we have that:

$$\text{Var} \left[\frac{u_{j_i}^2}{p_{j_i}} \right] \leq \mathbb{E} \left[\left(\frac{u_{j_i}^2}{p_{j_i}} \right)^2 \right] = \sum_{k=1}^n \frac{u_k^4}{p_k^2} \cdot p_k = \sum_{k=1}^n \frac{L^4 v_k^4 \text{rank}(\mathbf{X})}{\tau_k(\mathbf{X})}.$$

In the last step we have used the upper bound from (5), and the fact that $p_k = \tau_k(\mathbf{X}) / \text{rank}(\mathbf{X})$. From the definition of leverage scores (Definition 1), and the fact that \mathbf{v} lies in the span of \mathbf{X} , we have that $\tau_k(\mathbf{X}) \geq \frac{v_k^2}{\|\mathbf{v}\|_2^2}$. So we can further upper bound the variance as follows:

$$\begin{aligned} \text{Var} \left[\frac{u_{j_i}^2}{p_{j_i}} \right] &\leq L^4 \cdot \sum_{k=1}^n v_k^2 \|\mathbf{v}\|_2^2 \text{rank}(\mathbf{X}) \\ &= L^4 \cdot \|\mathbf{v}\|_2^4 \cdot \text{rank}(\mathbf{X}) \\ &\leq L^4 \cdot d \|\mathbf{v}\|_2^4. \end{aligned}$$

Moreover, we have that with probability 1, $\frac{u_{j_i}^2}{p_{j_i}} \leq \max_k L^2 \cdot \frac{v_k^2 \text{rank}(\mathbf{X})}{\tau_k(\mathbf{X})} \leq L^2 \cdot d \|\mathbf{v}\|_2^2$.

Finally, applying Bernstein's to the sum $\|\mathbf{S}\mathbf{u}\|_2^2 = \frac{1}{m} \sum_{i=1}^m \frac{u_{j_i}^2}{p_{j_i}}$, we have that:

$$\begin{aligned} \Pr [|\|\mathbf{S}\mathbf{u}\|_2^2 - \|\mathbf{u}\|_2^2| \geq t/m] \\ \leq 2 \exp \left(-\frac{t^2/2}{m \cdot L^4 \cdot d \|\mathbf{v}\|_2^4 + t \cdot L^2 \cdot d \|\mathbf{v}\|_2^2 / 3} \right). \end{aligned}$$

Setting $m = \frac{3d \log(2/\delta)}{\varepsilon^2}$ and $t = m \cdot \varepsilon \|\mathbf{v}\|_2^2 \cdot L^2$ and plugging in we have:

$$\begin{aligned} \Pr [|\|\mathbf{S}\mathbf{u}\|_2^2 - \|\mathbf{u}\|_2^2| \geq \varepsilon L^2 \|\mathbf{v}\|_2^2] \\ \leq 2 \exp \left(-\frac{\frac{1}{2} m^2 \varepsilon^2 \|\mathbf{v}\|_2^4 L^4}{m \cdot L^4 \cdot d \|\mathbf{v}\|_2^4 + m \varepsilon L^4 \cdot d \|\mathbf{v}\|_2^4 / 3} \right) \leq \delta. \end{aligned}$$

This completes the bound. \square

Proof of Lemma 3. Let N be an (εR) -net in the Euclidean norm on $\mathcal{B}^d(R)$. I.e. for every $\mathbf{v} \in \mathcal{B}^d(R)$, there should be some point $\mathbf{z} \in N$ such that $\|\mathbf{z} - \mathbf{v}\|_2 \leq \varepsilon R$. It is well known that such an N exists with cardinality $|N| \leq (1 + \frac{2}{\varepsilon})^d$ (see e.g. Lemma 5.2 in Vershynin [2012]). Applying Lemma 2 with $\delta = \frac{1}{50|N|}$ and combining with a union bound, we conclude that as long as $m \geq c \frac{d^2 \log(1/\varepsilon)}{\varepsilon^4}$ for a fixed constant c , then with probability 49/50, for all $\mathbf{z} \in N$,

$$\|f(\mathbf{X}\mathbf{z}) - f(\mathbf{X}\mathbf{w}^*)\|_2^2 \in \left[\|\mathbf{S}f(\mathbf{X}\mathbf{z}) - \mathbf{S}f(\mathbf{X}\mathbf{w}^*)\|_2^2 \pm \varepsilon^2 L^2 \|\mathbf{X}\mathbf{z} - \mathbf{X}\mathbf{w}^*\|_2^2 \right]. \quad (6)$$

Now, let \mathbf{z}^* be the closest point to $\hat{\mathbf{w}}$ in N . I.e., $\mathbf{z}^* = \arg \min_{\mathbf{z} \in N} \|\mathbf{z} - \hat{\mathbf{w}}\|_2$. Applying (6) and the fact that for any two

vectors \mathbf{a}, \mathbf{b} , $\|\mathbf{a} + \mathbf{b}\|_2^2 \leq 2\|\mathbf{a}\|_2 + 2\|\mathbf{b}\|_2^2$, we have:

$$\begin{aligned}
 \|f(\mathbf{X}\hat{\mathbf{w}}) - f(\mathbf{X}\mathbf{w}^*)\|_2^2 &\leq 2\|f(\mathbf{X}\mathbf{z}^*) - f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 2\|f(\mathbf{X}\hat{\mathbf{w}}) - f(\mathbf{X}\mathbf{z}^*)\|_2^2 \\
 &\leq 2\|\mathbf{S}f(\mathbf{X}\mathbf{z}^*) - \mathbf{S}f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 2\varepsilon^2 L^2 \|\mathbf{X}\mathbf{z}^* - \mathbf{X}\mathbf{w}^*\|_2^2 + 2\|f(\mathbf{X}\hat{\mathbf{w}}) - f(\mathbf{X}\mathbf{z}^*)\|_2^2 \\
 &\leq 4\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 4\|\mathbf{S}f(\mathbf{X}\mathbf{z}^*) - \mathbf{S}f(\mathbf{X}\hat{\mathbf{w}})\|_2^2 + 2\varepsilon^2 L^2 \|\mathbf{X}\mathbf{z}^* - \mathbf{X}\mathbf{w}^*\|_2^2 + 2\|f(\mathbf{X}\hat{\mathbf{w}}) - f(\mathbf{X}\mathbf{z}^*)\|_2^2 \\
 &\leq 4\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 4\|f(\mathbf{X}\mathbf{z}^*) - f(\mathbf{X}\hat{\mathbf{w}})\|_2^2 + 6\varepsilon^2 L^2 \|\mathbf{X}\mathbf{z}^* - \mathbf{X}\mathbf{w}^*\|_2^2 + 2L^2 \|\mathbf{X}\hat{\mathbf{w}} - \mathbf{X}\mathbf{z}^*\|_2^2 \\
 &\leq 4\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 4L^2 \|\mathbf{X}\mathbf{z}^* - \mathbf{X}\hat{\mathbf{w}}\|_2^2 + 6\varepsilon^2 L^2 (R + \|\mathbf{X}\mathbf{w}^*\|_2)^2 + 2L^2 \|\mathbf{X}\hat{\mathbf{w}} - \mathbf{X}\mathbf{z}^*\|_2^2 \\
 &\leq 4\|\mathbf{S}f(\mathbf{X}\hat{\mathbf{w}}) - \mathbf{S}f(\mathbf{X}\mathbf{w}^*)\|_2^2 + 4\varepsilon^2 L^2 R^2 + 12\varepsilon^2 L^2 R^2 + 12\|\mathbf{X}\mathbf{w}^*\|_2^2 + 2\varepsilon^2 L^2 R^2.
 \end{aligned}$$

Combining terms and adjusting constants on ε yields the bound. \square