
Regression-tree Tuning in a Streaming Setting

Samory Kpotufe*

Toyota Technological Institute at Chicago[†]
firstname@ttic.edu

Francesco Orabona*

Toyota Technological Institute at Chicago
francesco@orabona.com

Abstract

We consider the problem of maintaining the data-structures of a partition-based regression procedure in a setting where the training data arrives sequentially over time. We prove that it is possible to maintain such a structure in time $\mathcal{O}(\log n)$ at any time step n while achieving a nearly-optimal regression rate of $\tilde{\mathcal{O}}(n^{-2/(2+d)})$ in terms of the unknown metric dimension d . Finally we prove a new regression lower-bound which is independent of a given data size, and hence is more appropriate for the streaming setting.

1 Introduction

Traditional nonparametric regression such as kernel or k -NN can be expensive to estimate given modern large training data sizes. It is therefore common to resort to cheaper methods such as tree-based regression which precompute the regression estimates over a partition of the data space [7]. Given a future query x , the estimate $f_n(x)$ simply consists of finding the *closest* cell of the partition by traversing an appropriate tree-structure and returning the precomputed estimate. The partition and precomputed estimates depend on the training data and are usually maintained in batch-mode. We are interested in maintaining such a partition and estimates in a real-world setting where the training data arrives sequentially over time. Our constraints are that of fast-update at every time step, while maintaining a near-minimax regression error-rate at any point in time.

The error-rate of tree-based regression is well known to depend on the *size* of the partition's cells. We will call this size the *binwidth*. The minimax-optimal binwidth ϵ_n is known to be of the form $\mathcal{O}(n^{-1/(2+d)})$, assuming a training data of size n from a metric space of unknown dimension d , and unknown Lipschitz target function f . This setting of ϵ_n would then yield a minimax error rate of $\mathcal{O}(n^{-2/(2+d)})$. Thus, the dimension d is the most important problem variable entering the rate (and the tuning of ϵ_n), while other problem variables such as the Lipschitz properties of f are less crucial in comparison. The main focus of this work is therefore that of adapting to the unknown d while maintaining fast partition estimates in a streaming setting.

A first idea would be to start with an initial dimension estimation phase (where the regression estimates are suboptimal), and using the estimated dimension for subsequent data in a following phase, which leaves only the problem of maintaining partition estimates over time. However, while this sounds reasonable, it is generally unclear when to confidently stop such an initial phase since this would depend on the unknown d and the distribution of the data.

Our solution is to interleave dimension estimation with regression updates as the data arrives sequentially. However the algorithm never relies on the estimated dimensions and views them rather as *guesses* d_i . Even if $d_i \neq d$, it is kept as long as it is not hurtful to regression performance. The guess d_i is discarded once we detect that it hurts the regression, a new d_{i+1} is then estimated and a new phase $i+1$ is started. The decision to discard d_i relies on monitoring quantities that play into the tradeoff between regression variance and bias, more precisely we monitor the size of the partition

*SK and FO contributed equally to this paper.

[†]Other affiliation: Max Planck Institute for Intelligent Systems, Germany

and the partition’s binwidth ϵ_n . We note that the idea can be applied to other forms of regression where other quantities control the regression variance and bias (see longer version of the paper).

1.1 Technical Overview of Results

We assume that training data (x_i, Y_i) is sampled sequentially over time, x_i belongs to a general metric space \mathcal{X} of unknown dimension d , and Y_i is real. The exact setup is given in Section 2.

The algorithm (presented in Section 2.3) maintains regression estimates for all training samples $x^n \triangleq \{x_t\}_{t=1}^n$ arriving over time, while constantly updating a partition of the data and partition binwidth. At any time $t = n$, all updates are provably of order $\log n$ with constants depending on the unknown dimension d of \mathcal{X} .

At time $t = n$, the estimate for a query point x is given by the precomputed estimate for the closest point to x in x^n , which can be found fast using an off-the-shelf similarity search structure, such as those of [2, 10]. We prove that the L_2 error of the algorithm is $\tilde{O}(n^{-2/(2+d)})$, nearly optimal in terms of the unknown dimension d of the metric \mathcal{X} .

Finally, we prove a new lower-bound for regression on a generic metric \mathcal{X} , where the worst-case distribution is the same as n increases. Note that traditional lower-bounds for the offline setting derive a different worst-case distribution for each sample size n . Thus, our lower-bound is more appropriate to the streaming setting where the data arrives over time from the same distribution.

The results are discussed in more technical details in Section 3.

1.2 Related Work

Although various interesting heuristics have been proposed for maintaining tree-based learners in streaming settings (see e.g. [1, 5, 11, 15]), the problem has not received much theoretical attention. This is however an important problem given the growing size of modern datasets, and given that in many modern applications, training data is actually acquired sequentially over time and incremental updates have to be efficient (see e.g. Robotics [12, 16], Finance [8]).

The most closely related theoretical work is that of [6] which treats the problem of tuning a local polynomial regressor where the training data is acquired over time. Their setting however is that of a Euclidean space where d is known (ambient Euclidean dimension). [6] is thus concerned with maintaining a minimax error rate w.r.t. the known dimension d , while efficiently tuning regression bandwidth.

A possible alternative to the method analyzed here is to employ some form of cross-validation or even online solutions based on *mixture of experts* [3], by keeping track of different partitions, each corresponding to some setting of the bandwidth ϵ_n . This is however likely expensive to maintain in practice if good prediction performance is desired.

2 Preliminaries

2.1 Notions of metric dimension

We consider the following notion of dimension which extends traditional notions of dimension (e.g. Euclidean dimension and manifold dimension) to general metric spaces [4]. We assume throughout, w.l.o.g. that the space \mathcal{X} has diameter at most 1 under a metric ρ .

Definition 1. *The metric measure space (\mathcal{X}, μ, ρ) has **metric measure-dimension** d , if there exist $\tilde{C}_\mu, \hat{C}_\mu$ such that for all $\epsilon > 0$, and $x \in \mathcal{X}$, $\tilde{C}_\mu \epsilon^d \leq \mu(B(x, \epsilon)) \leq \hat{C}_\mu \epsilon^d$.*

The assumption of finite metric-measure dimension ensures that the measure μ has mass everywhere on the space ρ . This assumption is a generalization (to a metric space) of common assumptions where the measure has an upper and lower-bounded density on a compact Euclidean space, however is more general in that it does not require the measure μ to have a density (relative to any reference measure). The metric-measure dimension implies the following other notion of metric dimension.

Definition 2. *The metric (\mathcal{X}, ρ) has **metric dimension** d , if there exists \hat{C}_ρ such that, for all $\epsilon > 0$, \mathcal{X} has an ϵ -cover of size at most $\hat{C}_\rho \epsilon^{-d}$.*

The relation between the two notions of dimension is stated in the following lemma of [9], which allows us to use either notion as needed.

Lemma 1 ([9]). *If (\mathcal{X}, μ, ρ) has metric-measure dimension d , then there exists $\tilde{C}_\rho, \hat{C}_\rho$ such that, for all $\epsilon > 0$, any ball $B(x, r)$ centered on (\mathcal{X}, ρ) has an ϵr -cover of size in $[\tilde{C}_\rho \epsilon^{-d}, \hat{C}_\rho \epsilon^{-d}]$.*

2.2 Problem Setup

We receive data pairs $(x_1, Y_1), (x_2, Y_2), \dots$ sequentially, i.i.d. The input x_t belongs to a metric measure space (\mathcal{X}, ρ, μ) of diameter at most 1, and of metric measure dimension d . The output Y_t belongs to a subset of \mathbb{R} of bounded diameter Δ_Y , and satisfies $Y_t = f(x_t) + \eta(x_t)$. The noise $\eta(x_t)$ has 0 mean. The unknown function f is assumed to be λ -Lipschitz w.r.t. ρ for an unknown parameter λ , that is $\forall x, x' \in \mathcal{X}, |f(x) - f(x')| \leq \lambda \rho(x, x')$.

L_2 error: Our main performance result bounds the excess L_2 risk

$$\mathbb{E}_{x^n, Y^n} \|f_n - f\|_{2, \mu}^2 \triangleq \mathbb{E}_{x^n, Y^n} \mathbb{E}_X |f_n(X) - f(X)|^2.$$

We will often also be interested in the average error on the training sample: recall that at any time t , an estimate $f_t(x_s)$ of f is produced for every $x_s \in x^t$. The average error on x^n at $t = n$ is denoted

$$\mathbb{E}_n \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 \triangleq \frac{1}{n} \sum_{t=1}^n |f_n(x_t) - f(x_t)|^2.$$

2.3 Algorithm

The procedure (Algorithm 1) works by partitioning the data into small regions of size roughly $\epsilon_t/2$ at any time t , and computing the regression estimate of the centers of each region. All points falling in the same region (identified by a center point) are assigned the same regression estimate: the average Y values of all points in the region.

The procedure works in phases, where each Phase i corresponds to a guess d_i of the metric dimension d . Where ϵ_t might have been set to $t^{-1/(2+d)}$ if we knew d , we set it to $t_i^{-1/(2+d_i)}$ where t_i is the current time step within Phase i .

We ensure that in each phase our guess d_i does not hurt the variance-bias tradeoff of the estimates: this is done by monitoring the size of the partition ($|\mathbf{X}_i|$ in the algorithm), which controls the variance (see analysis in Section 4), relative to the bandwidth ϵ_t , which controls bias. Whenever $|\mathbf{X}_i|$ is too large relative to ϵ_t , the variance of the procedure is likely too large, so we start a new phase with an new guess of d_i .

Since the algorithm maintains at any time n an estimate $f_n(x_t)$ for all $x_t \in x^n$, for any query point $x \in \mathcal{X}$, we simply return $f_n(x) = f_n(x_t)$ where x_t is the closest point to x in x^n .

Despite having to adaptively tune to the unknown d , the main computation at each time step consists of just a 2-approximate nearest neighbor search for the closest center. These searches can be done fast in time $\mathcal{O}(\log n)$ by employing the off-the-shelf online search-procedure of [10]. This is emphasized in Lemma 2 below.

Finally, the algorithm employs a constant \tilde{C} which is assumed to upper-bound the constant C_ρ in Definition 2. This is a minor assumption since C_ρ is generally taken to be small, e.g. 1, in machine learning literature, and is exactly quantifiable for various metrics [4, 10].

3 Discussion of Results

3.1 Time complexity

The time complexity of updates is emphasized in the following Lemma.

Lemma 2. *Suppose (\mathcal{X}, ρ, μ) has metric dimension d . Then there exist C depending on d such that all computations of the algorithm at any time $t = n$ can be done in time $C \log n$.*

Algorithm 1 Incremental tree-based regressor.

```

1: Initialize:  $i = 1, d_i = 1, t_i = 0, \text{Centers } \mathbf{X}_i = \emptyset$ 
2: for  $t = 1, 2, \dots, T$  do
3:   Receive  $(x_t, y_t)$ 
4:    $t_i \leftarrow t_i + 1$  // counts the time steps within Phase  $i$ 
5:    $\epsilon_t \leftarrow t_i^{-1/(2+d_i)}$ 
6:   Set  $x_s \in \mathbf{X}_i$  to the 2-approximate nearest neighbor of  $x_t$ 
7:   if  $\rho(x_t, x_s) \leq \epsilon_t$  then
8:     Assign  $x_t$  to  $x_s$ 
9:      $f_n(x_s) \leftarrow$  update average  $Y$  for center  $x_s$  with  $y_t$ 
10:    For every  $r \leq t$  assigned to  $x_s, f_n(x_r) = f_n(x_s)$ 
11:   else
12:     if  $|\mathbf{X}_i| + 1 > \hat{C} 4^{d_i} \epsilon_t^{-d_i}$  then
13:       // Start of Phase  $i + 1$ 
14:        $d_{i+1} \leftarrow \left\lceil \log\left(\frac{|\mathbf{X}_i| + 1}{\hat{C}}\right) / \log(4/\epsilon_t) \right\rceil$ 
15:        $i \leftarrow i + 1$ 
16:     end if
17:     Add  $x_t$  as a new center in  $\mathbf{X}_i$ 
18:   end if
19: end for

```

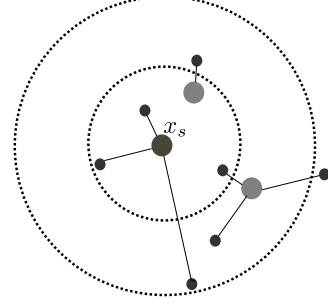


Figure 1: As ϵ_t varies over time, a ball around a center x_s can eventually contain both points assigned to x_s and points non-assigned to it, and even contain other centers. This results in a complex partitioning of the data.

Proof. The main computation at time n consists of finding the 2-approximate nearest neighbor x_n in \mathbf{X}_i and update the data structure for the nearest neighbor search. These centers are all at least $\frac{\epsilon_n}{2}$ far-apart. Using the results of [10], this can be done online in time $\mathcal{O}(\log(1/\epsilon_n) + \log \log(1/\epsilon_n))$. \square

3.2 Convergence rates

The main theorem below bounds the L_2 error of the algorithm at any given point in time. The main difficulty is in the fact that the data is partitioned in a complicated way due to the ever-changing bandwidth ϵ_t : every ball around a center can eventually contain both points assigned to the center and points not assigned to the center, in fact can contain other centers (see Figure 1). This makes it hard to get a handle on the number of points assigned to a single center x_t (contributing to the variance of $f_n(x_t)$) and the distance between points assigned to the same center (contributing to the bias). This is not the case in classical analyses of tree-based regression since the data partitioning is usually clearly defined.

The problem is handled by first looking at the average error over points in x^n , which is less difficult.

Theorem 1. *Suppose the space (\mathcal{X}, μ, ρ) has metric-measure dimension d .*

For any $x \in \mathcal{X}$, define $f_n(x) = f_n(x_t)$ where x_t is the closest point to x in x^n . Then at any time $t = n$, we have for some C independent of n ,

$$\mathbb{E}_{x^n, Y^n} \|f_n - f\|_{2, \mu}^2 \leq C(d \log n) \sup_{x^n} \mathbb{E}_n \mathbb{E}_{Y^n} \|f_n(X) - f(X)\|^2 + C\lambda^2 \left(\frac{d \log n}{n} \right)^{2/d} + \frac{\Delta_Y^2}{n}.$$

If the algorithm parameter $\hat{C} \geq \hat{C}_\rho$, then for some constant C' independent of n , we have at any time n that

$$\sup_{x^n} \mathbb{E}_n \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 \leq C' (\Delta_Y^2 + \lambda^2) n^{-2/(2+d)}.$$

The convergence rate is therefore $\tilde{\mathcal{O}}(n^{-2/(2+d)})$, nearly optimal in terms of the unknown d (up to a $\log n$ factor). In the simulation of Figure 2(Left) we compare our procedure to tree-based regressors with a fixed setting of d and of $\epsilon_t = t^{-1/(2+d)}$. We use the classic rotating-Teapot dataset, where the target output values are the cosine of the rotation angles. Our method attains the same performance as the one with the right fixed setting of d .

As alluded to above, the proof of Theorem 1 proceeds by first bounding the average error $\mathbb{E}_n \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2$ on the sample x^n . Interestingly, the analysis of the average error is




Figure 2: Simulation results on Teapot (Left) and Synthetic (Right) datasets. \hat{C} set to 1, size of the test sets 1800 and 12500, respectively.

of a worst-case nature where the data x_1, x_2, \dots is allowed to arrive adversarially (see analysis of Section 4.1). This shows a sense in which the algorithm is robust to bad dimension estimates: the average error is of the optimal form in terms of d , even though the data could trick us into picking a bad guess d_i of d . Thus the insights behind the algorithm are perhaps of wider applicability to problems of a more adversarial nature. This is shown empirically in Figure 2(Right), where we created a synthetic datasets with $d = 5$, while the first 1000 samples are from a line in \mathcal{X} . An algorithm that estimates the dimension in a first phase would end up using the suboptimal setting $d = 1$, while our algorithm robustly updates its estimate over time.

As mentioned in the introduction, the same insights can be applied to other forms of regression in a streaming setting. We show in the longer version of the paper a procedure more akin to kernel regression, which employs other quantities (appropriate to the method) to control the bias-variance tradeoff while deciding on keeping or rejecting the guess d_i .

3.3 Lower-bounds

We have to produce a distribution for which the problem is hard, and which matches our streaming setting as well as possible. With this in mind, our lower-bound result differs from existing non-parametric lower-bounds by combining two important aspects. First, the lower-bound holds for any given metric measure space (\mathcal{X}, ρ, μ) with finite measure-dimension: we constrain the worst-case distribution to have the marginal μ that nature happens to choose. In contrast, lower-bounds in literature would commonly pick a suitable marginal on the space \mathcal{X} [13, 14]. Second, the worst-case distribution does not depend on the sample size as is common in literature. Instead, we show that the rate of $n^{-2/(2+d)}$ holds for infinitely many n for a distribution fixed beforehand. This is more appropriate for the online setting where the data is generated over time from a fixed distribution.

The lower-bound result of [9] also holds for a given measure space (\mathcal{X}, μ, ρ) , but the worst-case distribution depends on sample size. A lower-bound of [7] holds for infinitely many n , but is restricted to distributions on a Euclidean cube, and thus benefits from the regularity of the cube. Our result combines some technical intuition from these two results in a way described in Section 4.3.

We need the following definition.

Definition 3. Given a metric-measure space (\mathcal{X}, μ, ρ) , we let $\mathcal{D}_{\mu, \lambda}$ denote the set of distributions on $X, Y, X \in \mathcal{X}, Y \in \mathbb{R}$, where the marginal on \mathcal{X} is μ , and where the function $f(x) = \mathbb{E}[Y|X = x]$ is λ -Lipschitz w.r.t. ρ .

Theorem 2. Let (\mathcal{X}, μ, ρ) be a metric space of diameter 1, and metric-measure dimension d . For any $n \in \mathbb{N}$, define $r_n^2 = (\lambda^2 n)^{-2/(2+d)}$. Pick any positive sequence $\{\beta_n\}_{n \in \mathbb{N}}$, $\beta_n = o(\lambda^2 r_n^2)$. There exists an indexing subsequence $\{n_t\}_{t \in \mathbb{N}}$, $n_{t+1} > n_t$, such that

$$\inf_{\{f_n\}} \sup_{\mathcal{D}_{\mu, \lambda}} \lim_{t \rightarrow \infty} \frac{\mathbb{E}_{X^{n_t}, Y^{n_t}} \|f_{n_t} - f\|_{2, \mu}^2}{\beta_{n_t}} = \infty,$$

where the infimum is taken over all sequences $\{f_n\}$ of estimators $f_n : X^n, Y^n \mapsto L_{2, \mu}$.

By the statement of the theorem, if we pick any rate β_n faster than $n^{-2/(2+d)}$, then there exists a distribution with marginal μ for which $\mathbb{E} \|f_n - f\|^2 / \beta_n$ either diverges or tends to ∞ .

4 Analysis

We first analyze the average error of the algorithm over the data x^n in Section 4.1. The proof of the main theorem follows in Section 4.2.

4.1 Bounds on Average Error

We start by bounding the average error on the sample x^n at time n , that is we upper-bound $\mathbb{E}_n \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2$.

The proof idea of the upper bound is the following. We bound the error in a given phase (Lemma 4), then combine these errors over all phases to obtain the final bounds (Corollary 1). To bound the error in a phase, we decompose the error in terms of squared bias and variance. The main technical difficulty is that the bandwidth ϵ_t varies over time and thus points at varying distances are included in each estimate. Nevertheless, if n_i is the number of steps in Phase i , we will see that both average squared bias and variance can be bounded by roughly $n_i^{-2/(2+d)}$.

Finally, the algorithm ensures that the guess d_i is always an under-estimate of the unknown dimension d , as proven in Lemma 3 (proof in the supplemental appendix), so integrating over all phases yields an adaptive bound on the average error. We assume throughout this section that the space (\mathcal{X}, ρ) has dimension d for some \hat{C}_ρ (see Def. 2).

Lemma 3. *Suppose the algorithm parameter $\hat{C} \geq \hat{C}_\rho$. The following invariants hold throughout the procedure for all phases $i \geq 1$ of Algorithm 1:*

- $i \leq d_i \leq d$.
- For any $t \in \text{Phase } i$ we have $|\mathbf{X}_i| \leq \hat{C} 4^{d_i} \epsilon_t^{-d_i}$.

Lemma 4 (Bound on single phase). *Suppose the algorithm parameter $\hat{C} \geq \hat{C}_\rho$. Consider Phase $i \geq 1$ and suppose this phase lasts n_i steps. Let \mathbb{E}_{n_i} denote expectation relative to the uniform choice of X out of $\{x_t : t \in \text{Phase } i\}$. We have the following bound:*

$$\mathbb{E}_{n_i} \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 \leq \left(\hat{C} 4^d \Delta_Y^2 + 12\lambda^2 \right) n_i^{-\frac{2}{2+d}}.$$

Proof. Let $\mathbf{X}_i(X)$ denote the center closest to X in \mathbf{X}_i . Suppose $\mathbf{X}_i(X) = x_s, s \in [n]$, we let n_{x_s} denote the number of points assigned to the center x_s . We use the notation $x_t \rightarrow x_s$ to say that x_t is assigned to center x_s .

First fix $X \in \{x_t : t \in \text{Phase } i\}$ and let $x_s = \mathbf{X}_i(x_t)$. Define $\tilde{f}_n(X) \equiv \mathbb{E}_{Y^n} f_n(X) = \frac{1}{n_{x_s}} \sum_{x_t \rightarrow x_s} f(x_t)$. We proceed with the following standard bias-variance decomposition

$$\mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 = \mathbb{E}_{Y^n} \left| f_n(X) - \tilde{f}_n(X) \right|^2 + \left| \tilde{f}_n(X) - f(X) \right|^2. \quad (1)$$

Let $X = x_r, r \geq s$. We first bound the bias term. Using the Lipschitz property of f , and Jensen's inequality, we have

$$\begin{aligned} \left| \tilde{f}_n(X) - f(X) \right|^2 &\leq \left(\frac{1}{n_{x_s}} \sum_{x_t \rightarrow x_s} \lambda \rho(x_r, x_t) \right)^2 \leq \frac{1}{n_{x_s}} \sum_{x_t \rightarrow x_s} \lambda^2 \rho(x_r, x_t)^2 \\ &\leq \frac{2\lambda^2}{n_{x_s}} \sum_{x_t \rightarrow x_s} \left(\rho(x_r, x_s)^2 + \rho(x_s, x_t)^2 \right) \leq \frac{2\lambda^2}{n_{x_s}} \sum_{x_t \rightarrow x_s} (\epsilon_r^2 + \epsilon_t^2). \end{aligned}$$

The variance term is easily bounded as follows:

$$\mathbb{E}_{Y^n} \left| f_n(X) - \tilde{f}_n(X) \right|^2 = \sum_{x_t \rightarrow x_s} \frac{\mathbb{E}_{Y^n} |Y_t - f(x_t)|^2}{n_{x_s}^2} \leq \frac{\Delta_Y^2}{n_{x_s}}.$$

Now take the expectation over $X \sim \mathcal{U} \{x_t : t \in \text{Phase } i\}$. We have:

$$\begin{aligned}
\mathbb{E}_{n_i} \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 &= \sum_{x_s \in \mathbf{X}_i} \mathbb{E}_n \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 \cdot \mathbf{1}\{X \rightarrow x_s\} \\
&\leq \frac{1}{n_i} \sum_{x_s \in \mathbf{X}_i} \sum_{x_r \rightarrow x_s} \left(\frac{\Delta_Y^2}{n_{x_s}} + \frac{2\lambda^2}{n_{x_s}} \sum_{x_t \rightarrow x_s} (\epsilon_r^2 + \epsilon_t^2) \right) \\
&= \frac{1}{n_i} \sum_{x_s \in \mathbf{X}_i} \Delta_Y^2 + \frac{2\lambda^2}{n_i} \sum_{x_s \in \mathbf{X}_i} \frac{1}{n_{x_s}} \sum_{x_r \rightarrow x_s} \sum_{x_t \rightarrow x_s} (\epsilon_r^2 + \epsilon_t^2) \\
&= \frac{\Delta_Y^2 \cdot |\mathbf{X}_i|}{n_i} + \frac{4\lambda^2}{n_i} \sum_{x_s \in \mathbf{X}_i} \sum_{x_t \rightarrow x_s} \epsilon_t^2 = \frac{\Delta_Y^2 \cdot |\mathbf{X}_i|}{n_i} + \frac{4\lambda^2}{n_i} \sum_{t \in \text{Phase } i} \epsilon_t^2.
\end{aligned}$$

To bound the last term, we have

$$\sum_{t \in \text{Phase } i} \epsilon_t^2 = \sum_{t_i \in [n_i]} t^{-\frac{2}{2+d_i}} \leq \int_0^{n_i} \tau^{-\frac{2}{2+d_i}} d\tau \leq 3n_i^{1-\frac{2}{2+d_i}}.$$

Combine with the previous derivation and with both statements of Lemma 3 to get

$$\mathbb{E}_{n_i} \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 \leq \frac{\Delta_Y^2 \cdot |\mathbf{X}_i|}{n_i} + 12\lambda^2 n_i^{-\frac{2}{2+d_i}} \leq \left(\hat{C} 4^d \Delta_Y^2 + 12\lambda^2 \right) n_i^{-\frac{2}{2+d}}.$$

□

Corollary 1 (Combined phases). *Suppose the algorithm parameter $\hat{C} \geq \hat{C}_\rho$, then we have*

$$\mathbb{E}_n \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 \leq 2 \left(\hat{C} 4^d \Delta_Y^2 + 12\lambda^2 \right) n^{-\frac{2}{2+d}}.$$

Proof. Let I denote the number of phases up to time n . We will decompose the expectation \mathbb{E}_n in terms of the various phases $i \in [I]$ and apply Lemma 4. Let $B \triangleq \hat{C} 4^d \Delta_Y^2 + 12\lambda^2$. We have:

$$\begin{aligned}
\mathbb{E}_n \mathbb{E}_{Y^n} |f_n(X) - f(X)|^2 &\leq B \sum_{i=1}^I \frac{n_i}{n} n_i^{-\frac{2}{2+d}} = B \frac{I}{n} \sum_{i=1}^I \frac{1}{I} n_i^{\frac{d}{2+d}} \leq B \frac{I}{n} \left(\sum_{i=1}^I \frac{n_i}{I} \right)^{\frac{d}{2+d}} \\
&= B \frac{I}{n} \left(\frac{n}{I} \right)^{\frac{d}{2+d}} = B \cdot I^{\frac{2}{2+d}} n^{-\frac{2}{2+d}} \leq B \cdot d^{\frac{2}{2+d}} n^{-\frac{2}{2+d}},
\end{aligned}$$

where in the second inequality we use Jensen's inequality, and in the last inequality Lemma 3. □

4.2 Bound on L_2 Error

We need the following lemma, whose proof is in the supplemental appendix, which bounds the probability that a ρ -ball of a given radius contains a sample from x^n . This will then allow us to bound the bias induced by transforming a solution for the adversarial setting to a solution for the stochastic setting.

Lemma 5. *Suppose (\mathcal{X}, ρ, μ) has metric measure dimension d . Let μ be a distribution on \mathcal{X} and let μ_n denote an empirical distribution on an i.i.d sample x^n from μ . For $\epsilon > 1/n$, let \mathcal{B}_ϵ denote the class of ρ -balls centered on \mathcal{X} of radius ϵ . There exists C depending on d such that the following holds. Let $0 < \delta < 1$, Define $\alpha_{n,\delta} = C(d \log n + \log(1/\delta))$. Then, with probability at least $1 - \delta$, for all $B \in \mathcal{B}_\epsilon$ satisfying $\mu(B) \geq \alpha_{n,\delta}/n$ we have $\mu_n(B) > 1/n$.*

We are now ready to prove Theorem 1.

Proof of Theorem 1. Fix $\delta = 1/n$ and define $\alpha_{n,\delta}$ as in Lemma 5. Pick $\epsilon = (\alpha_{n,\delta}/C_1 n)^{1/d} \geq 1/n$, where C_1 is such that every $B \in \mathcal{B}_\epsilon$ has mass at least $C_1 \epsilon^d$. Since for every $B \in \mathcal{B}_\epsilon$, $\mu(B) \geq C_1 \epsilon^{-d} \geq \alpha_{n,\delta}/n$, we have by Lemma 5, that with probability at least $1 - \delta$, all $B \in \mathcal{B}_\epsilon$ contain a point from x^n . In other words, the event \mathcal{E} that x^n forms an ϵ -cover of \mathcal{X} is $(1 - \delta)$ -likely.

Suppose x_t is the closest point in x^n to $x \in \mathcal{X}$. We write $x \rightarrow x_t$. Then, under \mathcal{E} , we have, $\|f(x) - f(x_t)\| \leq \lambda\epsilon$. We therefore have by Fubini's theorem

$$\begin{aligned} \mathbb{E}_{x^n, Y^n} \|f_n - f\|_{2, \mu}^2 &= \mathbb{E}_{x^n} \mathbb{E}_X \mathbb{E}_{Y^n | x^n} |f_n(X) - f(X)|^2 \cdot (\mathbf{1}\{\mathcal{E}\} + \mathbf{1}\{\bar{\mathcal{E}}\}) \\ &\leq \mathbb{E}_{x^n} \sum_{t=1}^n 2\mu(x : x \rightarrow x_t) \mathbb{E}_{Y^n | x^n} |f_n(x_t) - f(x_t)|^2 + 2\lambda^2\epsilon^2 + \delta\Delta_Y^2 \\ &\leq \mathbb{E}_{x^n} \sum_{t=1}^n 2C_2\epsilon^d \mathbb{E}_{Y^n | x^n} |f_n(x_t) - f(x_t)|^2 + 2\lambda^2\epsilon^2 + \delta\Delta_Y^2 \\ &\leq \frac{2C_2\alpha_{n, \delta}}{C_1} \sup_{x^n} \mathbb{E}_n \mathbb{E}_{Y^n} |f_n(x_t) - f(x_t)|^2 + 2\lambda^2\epsilon^2 + \delta\Delta_Y^2, \end{aligned}$$

where in the first inequality we break the integration over the Voronoi partition of \mathcal{X} defined by the points in x^n , and introduce $f(x_t)$; the second inequality uses $\{x : x \rightarrow x_t\} \subset B(x_t, \epsilon)$ under \mathcal{E} . \square

4.3 Lower-bound

Let's consider first the case of a fixed n . The idea behind the proof is as follows: for μ fixed, we have to come up with a class \mathcal{F} of functions which vary considerably on the space \mathcal{X} . To this end we discretize \mathcal{X} into as many cells as possible, and let any $f \in \mathcal{F}$ potentially change sign from one cell to the other. The larger the dimension d the more we can discretize the space and the more complex \mathcal{F} , subject to a Lipschitz constraint. The problem of picking the right f can thus be reduced to that of classification, since the learner has to discover the sign of f on sufficiently many cells.

In order to handle many data sizes n simultaneously, we borrow from the idea above. Say we want to show that the lower-bound holds for a subsequence $\{n_i\}$ simultaneously. Then we reserve a subset of the space \mathcal{X} for each n_1, n_2, \dots , and discretize each subset according to n_i . The functions in \mathcal{F} have to then vary sufficiently in each subset of the space \mathcal{X} according to the corresponding n_i . This is illustrated in Figure 3.




Figure 3: Lower bound proof idea.

We can then apply the same idea of reduction to classification for each n_t separately. This sort of idea appears in [7] where μ is uniform on the Euclidean cube, where they use the regularity of the cube to set up the right sequence of discretizations over subsets of the cube. The main technicality in our result is that we work with a general space without much regularity. The lack of regularity makes it unclear a priori how to *divide* such a space into subsets of the proper *size* for each n_i .

Last, we have to ensure that the functions $f \in \mathcal{F}$ resulting from our discretization of a general metric space \mathcal{X} are in fact Lipschitz. For this, we extend some of the ideas from [9] which handles the case of a fixed n . For lack of space, the complete proof is in the extended version of the paper.

5 Conclusions

We presented an efficient and nearly minimax optimal approach to nonparametric regression in a streaming setting. The streaming setting is gaining more attention as modern data sizes are getting larger, and as data is being acquired online in many applications.

The main insights behind the approach presented extend to other nonparametric methods, and are likely to extend to settings of a more adversarial nature. We left open the question of optimal adaptation to the smoothness of the unknown function, while we efficiently solve the equally or more important question of adapting to the unknown dimension of the data, which generally has a stronger effect on the convergence rate.

References

- [1] Y. Ben-Haim and E. Tom-Tov. A streaming parallel decision tree algorithm. *Journal of Machine Learning Research*, 11:849–872, 2010.
- [2] A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbors. *ICML*, 2006.
- [3] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.
- [4] K. Clarkson. Nearest-neighbor searching and metric space dimensions. *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*, 2005.
- [5] P. Domingos and G. Hulten. Mining high-speed data streams. In *Proceedings of the 6th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 71–80, 2000.
- [6] H. Gu and J. Lafferty. Sequential nonparametric regression. *ICML*, 2012.
- [7] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution Free Theory of Nonparametric Regression*. Springer, New York, NY, 2002.
- [8] A. Kalai and S. Vempala. Efficient algorithms for universal portfolios. *Journal of Machine Learning Research*, 3:423–440, 2002.
- [9] S. Kpotufe. k-NN Regression Adapts to Local Intrinsic Dimension. *NIPS*, 2011.
- [10] R. Krauthgamer and J. R. Lee. Navigating nets: simple algorithms for proximity search. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '04, pages 798–807, Philadelphia, PA, USA, 2004. Society for Industrial and Applied Mathematics.
- [11] B. Pfahringer, G. Holmes, and R. Kirkby. Handling numeric attributes in hoeffding trees. In *Advances in Knowledge Discovery and Data Mining: Proceedings of the 12th Pacific-Asia Conference (PAKDD)*, volume 5012, pages 296–307. Springer, 2008.
- [12] S. Schaal and C. Atkeson. Robot Juggling: An Implementation of Memory-based Learning. *Control Systems Magazine, IEEE*, 1994.
- [13] C. J. Stone. Optimal rates of convergence for non-parametric estimators. *Ann. Statist.*, 8:1348–1360, 1980.
- [14] C. J. Stone. Optimal global rates of convergence for non-parametric estimators. *Ann. Statist.*, 10:1340–1353, 1982.
- [15] M. A. Taddy, R. B. Gramacy, and N. G. Polson. Dynamic trees for learning and design. *Journal of the American Statistical Association*, 106(493), 2011.
- [16] S. Vijayakumar and S. Schaal. Locally weighted projection regression: An $O(n)$ algorithm for incremental real time learning in high dimensional space. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, pages 1079–1086, 2000.

6 Appendix

6.1 Proof of Lemma 3

We require the following lemma that relates the size of a packing to that of a cover.

Lemma 6 ([4]). *If (\mathcal{X}, ρ) has metric dimension d for some C_ρ (see Definition 2), then any ϵ -packing Z of \mathcal{X} (i.e. $\rho(z, z') > \epsilon$ for pairs $z, z' \in Z$) has size at most that $2^d C_\rho \epsilon^{-d}$.*

Proof. By the design of the algorithm, and the fact that we use a 2-approximate nearest neighbor search, we know that \mathbf{X}_i is a $\frac{\epsilon_t}{2}$ packing, hence we can apply Lemma 6 to have $|\mathbf{X}_i| \leq 2^d C_\rho \frac{\epsilon_t}{2}^{-d}$, hence, by definition of d_i , we have $d_i \leq d$ whenever $\hat{C} \geq \hat{C}_\rho$. The lower-bound on d_i is obtained by induction as follows. We have $d_1 = 1$. For $i > 1$, we have $d_{i+1} > d_i$ since it held that $|\mathbf{X}_i| + 1 > \hat{C} 4^{d_i} \epsilon_t^{-d_i}$. Thus $d_{i+1} \geq d_i + 1 \geq i + 1$.

For the second part, the bound holds by the condition on ending a phase and because ϵ_t decreases over time within a phase. \square

6.2 Proof of Lemma 5

Proof. By Lemma 1, there exists an $\epsilon/4$ -cover Z of \mathcal{X} of size at most $(\epsilon/2)^{-d} \leq \hat{C}_0 (4n)^d$ for some C_0 . Apply Bernstein's inequality for every $z \in Z$, followed by a union bound over $z \in Z$ to obtain that, with probability at least $1 - \delta$, $\mu_n(B(z, \epsilon/2)) > 1/n$ whenever $\mu(B(z, \epsilon/2)) \geq \alpha'_{n,\delta}/n$, where $\alpha'_{n,\delta} = C_1 (d \log n + \log(1/\delta))$ an appropriate setting of C_1 . Now remark that every ball $B(x, \epsilon)$ in \mathcal{B}_ϵ contains a ball $B(z, \epsilon/2)$ which in turn contains $B(x, \epsilon/4)$. By assumption, there exist C_2 such that for all $x \in \mathcal{X}$, $C_2 \mu(B(x, \epsilon/4)) \geq \mu(B(x, \epsilon))$. Thus for $\mu(B(x, \epsilon)) \geq C_2 \alpha'_{n,\delta}/n$, we have $\mu(B(z, \epsilon/2)) \geq \mu(B(x, \epsilon/4)) \geq \alpha'_{n,\delta}/n$ and therefore $\mu_n(B(x, \epsilon)) \geq \mu_n(B(z, \epsilon/2)) \geq 1/n$. \square

6.3 Lower Bound Proof

We start with the following definition.

Definition 4 (ϵ -net). *An ϵ -net is both an ϵ -cover and an ϵ -packing.*

Theorem 3. *Suppose (\mathcal{X}, μ, ρ) has diameter 1, and measure-dimension d . Suppose (\mathcal{X}, ρ) admits a metric measure μ . For any $n \in \mathbb{N}$, define $r_n^2 = (\lambda^2 n)^{-\frac{2}{2+d}}$. Pick any positive sequence $\{\alpha_n\}_{n \in \mathbb{N}}$, $\alpha_n \xrightarrow{n \rightarrow \infty} 0$. There exists $C > 0$, and an indexing subsequence $\{n_t\}_{t \in \mathbb{N}}$, $n_{t+1} > n_t$, such that*

$$\inf_{\{f_n\}} \sup_{\mathcal{D}_\lambda} \overline{\lim}_{t \rightarrow \infty} \frac{\mathbb{E}_{X^{n_t}, Y^{n_t}} \|f_{n_t} - f\|_{2,\mu}^2}{\alpha_{n_t} \lambda^2 r_{n_t}^2} > C,$$

where the infimum is taken over all sequences $\{f_n\}$ of estimators $f_n : X^n, Y^n \mapsto L_{2,\mu}$.

Proof. For $t \in \mathbb{N}$, let $n_0 = 0$ and define n_t recursively as the smallest value of $n > n_{t-1}$ such that $2^{-t} > \max\{8r_n, \alpha_n\}$.

Given $z \in \mathcal{X}$, and $t \in \mathbb{N}$, define the following z -centered function over \mathcal{X} :

$$g_{z,t}(x) \triangleq \frac{\lambda}{5} (\tau r_{n_t} - \rho(x, z))_+,$$

for some fixed $\tau \leq 1/2$ which is further specified below. In other words, $g_{z,n}(x) = 0$ whenever $\rho(x, z) \geq \tau r_{n_t}$. It is not hard to see that, by triangle-inequality, $g_{z,t}(x)$ is $\lambda/5$ -Lipschitz.

We will choose centers z and indices t to form a collection $\mathcal{G} = \{g_{z,t}\}$ which generates a class of functions

$$\mathcal{F} = \left\{ f_\sigma(\cdot) = \sum_{g_{z,t} \in \mathcal{G}} \sigma_{z,t} \cdot g_{z,t}(\cdot), \sigma_{z,t} \in \{-1, 1\} \right\}.$$

The collection \mathcal{G} is defined recursively as follows.

For $t \in \mathbb{N}$, let Z_t denote a 2^{-t} -net of \mathcal{X} . We will pick points in $Z'_t \subset Z_t$, and pick centers for functions $g_{z,t}$ near these points in Z'_t .

Let $Z'_0 = \emptyset$. For $t \geq 1$, let \mathcal{B}_{t-1} denote $\cup_{s=0}^{t-1} \{B(z, 2^{-s-1}) : z \in Z'_s\}$. Let $c_t = C_0 2^{t(d-1)}$ for $C_0 < C_1/(C_2+1)$. Pick c_t points $z \in Z_t \setminus \mathcal{B}_{t-1}$. These are the points in Z'_t . This is always possible since

$$\begin{aligned} |Z_t \setminus \mathcal{B}_{t-1}| &\geq |Z_t| - \sum_{s=0}^{t-1} |\{B(z, 2^{-s-1}) \cap Z_t : z \in Z'_s\}| \\ &\geq |Z_t| - \sum_{s=0}^{t-1} c_s C_2 2^{(t-s-1)d} = |Z_t| C_0 C_2 \sum_{s=0}^{t-1} 2^{(t-1)d-s} \\ &\geq C_1 2^{td} - C_0 C_2 2^{(t-1)d} \geq C_0 2^{t(d-1)} = c_t. \end{aligned}$$

Now, for every $z' \in Z'_t$, let $Z_{n_t}(z')$ be an r_{n_t} packing of $B(z', 2^{-t}/8)$. We let $Z''_t = \cup_{z' \in Z'_t} Z_{n_t}(z')$. We can now define $\mathcal{G} \equiv \{g_{z,t} : z \in Z''_t, t \geq 1\}$. The functions in \mathcal{G} are mutually orthogonal in $L_{2,\mu}$. We show this next.

Pick any $s \geq t \geq 1$, different centers $z_t \in Z''_t$, $z_s \in Z''_s$, and suppose $z_t \in Z_{n_t}(z'_t)$ and $z_s \in Z_{n_s}(z'_s)$ for some $z'_t \in Z'_t$ and $z'_s \in Z'_s$. If $z'_t = z'_s$, then $t = s$, and $\rho(z_t, z_s) > r_{n_t}$ by construction. Otherwise, if $z'_t \neq z'_s$, then

$$\begin{aligned} \rho(z_t, z_s) &\geq \rho(z'_t, z'_s) - \rho(z'_t, z_t) - \rho(z'_s, z_s) \\ &\geq 2^{-t}/2 - 2^{-t}/4 > r_{n_t}. \end{aligned}$$

In either case, g_{z_t, n_t} and g_{z_s, n_s} are respectively nonzero only on balls $B(z_t, \tau r_{n_t})$ and $B(z_s, \tau r_{n_s})$, and these balls do not intersect since $\rho(z_t, z_s) > r_{n_t} \geq 2\tau r_{n_t}$. Thus $\langle g_{z_t, n_t}, g_{z_s, n_s} \rangle = \mathbb{E}_{X \sim \mu} |g_{z_t, n_t}(X) \cdot g_{z_s, n_s}(X)| = 0$.

One can show that for τ sufficiently small (relative to the growth of μ), the functions $f_\sigma \in \mathcal{F}$ are λ -Lipschitz. The argument, which can be found in Theorem 2 of [], relies on the fact that the functions $g_{z,t} \in \mathcal{G}$ are $\lambda/5$ -Lipschitz, and that μ is doubling.

We will consider the class of distributions $\mathcal{D}_{\mathcal{F}}$ of X, Y , where $X \sim \mu$, and $Y = f_\sigma(X) + \mathcal{N}(0, 1)$, $f_\sigma \in \mathcal{F}$. By construction $\mathcal{D}_{\mathcal{F}} \subset \mathcal{D}_\lambda$ so we only need to bound the hardness of learning from $\mathcal{D}_{\mathcal{F}}$.

Now given the above construction, the rest of the proof relies on standard techniques (see e.g. [7]), where the regression problem is reduced to some classification problems.

We will need the following fact: there exists C_3, C_4 such that, for any $g_{z,t} \in \mathcal{G}$, we have

$$C_3 \lambda^2 r_{n_t}^{2+d} \leq \|g_{z,t}\|^2 \leq C_4 \lambda^2 r_{n_t}^{2+d}.$$

To see this, remark that $g_{z,t}$ is at most $\lambda \tau r_{n_t}$ on $B(z, \tau r_{n_t})$ and at least $\lambda \tau r_{n_t}/10$ on $B(z, \tau r_{n_t}/2)$; furthermore, by the doubling assumption on μ , these balls have mass proportional to $r_{n_t}^d$.

Now, for every estimator $\{f_n\}$, we let $f_{n,\mathcal{G}} = \sum_{g_{z,t} \in \mathcal{G}} w_{z,t} g_{z,t}$ denote the projection of $f_n(X^n, Y^n)$ onto the orthonormal system induced by \mathcal{G} . We have for every n_t that

$$\begin{aligned} \|f_{n_t} - f\|_{2,\mu}^2 &\geq \|f_{n_t,\mathcal{G}} - f\|_{2,\mu}^2 \\ &\geq \sum_{z \in Z''_t} (w_{z,t} - \sigma_{z,t})^2 \|g_{z,t}\|^2 \\ &\geq \sum_{z \in Z''_t} \mathbf{1}\{w_{z,t} \cdot \sigma_{z,t} < 0\} \|g_{z,t}\|^2 \\ &\geq C_3 \lambda^2 r_{n_t}^{2+d} \sum_{z \in Z''_t} \mathbf{1}\{w_{z,t} \cdot \sigma_{z,t} < 0\}. \end{aligned}$$

Let \mathbb{E}_σ denote expectation over the random choice of $\sigma = \{\sigma_{z,t}\}$ where each $\sigma_{z,t} = 1$ w.p. $1/2$. Let \mathbb{E} denote $\mathbb{E}_{X^{n_t}, Y^{n_t}}$ for short. We have

$$\begin{aligned}
& \inf_{\{f_n\}} \sup_{\mathcal{D}_{\mathcal{F}}} \overline{\lim}_{t \rightarrow \infty} \frac{\mathbb{E} \|f_{n_t} - f\|_{2,\mu}^2}{\alpha_{n_t} \lambda^2 r_{n_t}^2} \\
& \geq \inf_{\{f_n\}} \mathbb{E} \overline{\lim}_{\sigma t \rightarrow \infty} \frac{\mathbb{E} \|f_{n_t} - f\|_{2,\mu}^2}{\alpha_{n_t} \lambda^2 r_{n_t}^2} \\
& \geq \inf_{\{f_n\}} \mathbb{E} \overline{\lim}_{\sigma t \rightarrow \infty} \frac{C_3 r_{n_t}^d \mathbb{E} \sum_{z \in Z_t''} \mathbf{1}\{w_{z,t} \cdot \sigma_{z,t} < 0\}}{\alpha_{n_t}} \\
& \geq \inf_{\{f_n\}} \mathbb{E} \overline{\lim}_{\sigma t \rightarrow \infty} \frac{C_3 r_{n_t}^d \sum_{z \in Z_t''} \mathbb{E} \mathbf{1}\{w_{z,t} \cdot \sigma_{z,t} < 0\}}{2^{-t}} \\
& \geq \inf_{\{f_n\}} \overline{\lim}_{t \rightarrow \infty} \frac{C_3 r_{n_t}^d \sum_{z \in Z_t''} \mathbb{E} \mathbb{E}_\sigma \mathbf{1}\{w_{z,t} \cdot \sigma_{z,t} < 0\}}{2^{-t}},
\end{aligned}$$

where in the last inequality above we applied, first, the dominated convergence theorem (DCT), then Fubini's theorem. The application of DCT is valid since the term in the limit can be upper-bounded by noticing that

$$\begin{aligned}
& r_{n_t}^d \sum_{z \in Z_t''} \mathbb{E} \mathbf{1}\{w_{z,t} \cdot \sigma_{z,t} < 0\} \\
& \leq r_{n_t}^d |Z_t''| \leq r_{n_t}^d c_t C_2 2^{-(t+3)d} r_{n_t}^{-d} = C_0 C_2 2^{-t+3d}.
\end{aligned}$$

For $z \in Z_t'$, and fix X^{n_t} . Then $\mathbb{E} \mathbb{E}_\sigma \mathbf{1}\{w_{z,t} \cdot \sigma_{z,t} < 0\}$ is the probability of error of a classifier (which outputs $\text{sign}(w_{n,z})$) for the following prediction task. Let $x_{(1)}, x_{(2)}, \dots, x_{(m)}$ denote the values of X falling in $B(z, \tau r_{n_t})$ where $g_{z,t}$ is non zero. Then

$$(Y_{(1)}, \dots, Y_{(m)}) = \varsigma_z(g_{z,t}(x_{(1)}), \dots, g_{z,t}(x_{(m)})) + \mathcal{N}(0, I_m)$$

is a random vector sampled from the equal-weight mixture of two spherical Gaussians in \mathbb{R}^m centered at $u \doteq (g_{z,t}(x_{(1)}), \dots, g_{z,t}(x_{(m)}))$ and $-u$. The prediction task is that of identifying the right mixture component from the single sample $(Y_{(1)}, \dots, Y_{(m)})$. The smallest possible error for this task is that of the Bayes classifier and is well known to be $\Phi(-\|u\|) \geq \Phi\left(-\sqrt{\sum_{i=1}^{n_t} g_{z,t}^2(X_i)}\right)$. We can now take the expectation over X^{n_t} , and since $\Phi(-\sqrt{\cdot})$ is convex, we have by Jensen's inequality that

$$\begin{aligned}
\mathbb{E} \mathbb{E}_\sigma \mathbf{1}\{w_{z,t} \cdot \sigma_{z,t} < 0\} & \geq \mathbb{E}_{X^{n_t}} \Phi\left(-\sqrt{\sum_{i=1}^n g_{z,t}^2(X_i)}\right) \\
& \geq \Phi\left(-\sqrt{\sum_{i=1}^n \mathbb{E}_{X_i} g_{z,t}^2(X_i)}\right) = \Phi\left(-\sqrt{n \|g_{z,t}\|^2}\right) \\
& \geq \Phi\left(\sqrt{C_4}\right).
\end{aligned}$$

Finally, since $|Z_t''| \geq c_t C_1 2^{-(t+3)d} r_{n_t}^{-d} = C_0 C_1 2^{-t+3d} r_{n_t}^{-d}$, we have

$$\inf_{\{f_n\}} \sup_{\mathcal{D}_{\mathcal{F}}} \overline{\lim}_{t \rightarrow \infty} \frac{\mathbb{E} \|f_{n_t} - f\|_{2,\mu}^2}{\alpha_{n_t} \lambda^2 r_{n_t}^2} \geq C_3 C_0 C_1 2^{3d} \Phi\left(\sqrt{C_4}\right).$$

□

The main lower bound statement of Theorem 2 is obtained as follows.

Proof of Theorem 2. Fix any f_n . Let $\alpha_n^2 = \beta_n / (\lambda^2 r_n^2)$. From Theorem 3 there exists a distribution in \mathcal{D}_λ such that

$$\overline{\lim}_{n \rightarrow \infty} \frac{\mathbb{E}_{X^n, Y^n} \|f_{n_t} - f\|_{2,\mu}^2}{\alpha_n^2 \lambda^2 r_n^2} = \infty.$$

Pick a subsequence whose limit is the lim sup. □