

# Efficient Regression in Metric Spaces via Approximate Lipschitz Extension

Lee-Ad Gottlieb and Aryeh Kontorovich and Robert Krauthgamer

**Abstract**—We present a framework for performing efficient regression in general metric spaces. Roughly speaking, our regressor predicts the value at a new point by computing an approximate Lipschitz extension — the smoothest function consistent with the observed data — after performing structural risk minimization to avoid overfitting. We obtain finite-sample risk bounds with minimal structural and noise assumptions, and a natural runtime-precision tradeoff. The offline (learning) and online (prediction) stages can be solved by convex programming, but this naive approach has runtime complexity  $O(n^3)$ , which is prohibitive for large datasets. We design instead a regression algorithm whose speed and generalization performance depend on the intrinsic dimension of the data, to which the algorithm adapts. While our main innovation is algorithmic, the statistical results may also be of independent interest.

## I. INTRODUCTION

The classical problem of estimating a continuous-valued function from noisy observations, known as *regression*, is of central importance in statistical theory with a broad range of applications, see e.g. [41], [7], [38], [22], [20]. When no structural assumptions concerning the target function are made, the regression problem is termed *nonparametric*. Informally, the main objective in the study of nonparametric regression is to understand the relationship between the regularity conditions that a function class might satisfy (e.g., Lipschitz or Hölder continuity, or sparsity in some representation) and the minimax risk convergence rates [47], [51]. A further consideration is the computational efficiency of constructing the regression function.

The general (univariate) nonparametric regression problem may be stated as follows. Let  $(\mathcal{X}, \rho)$  be a metric space, and let  $\mathcal{H}$  be a collection of functions (“hypotheses”)  $h : \mathcal{X} \rightarrow [0, 1]$ . (Although in general,

$h$  need not be restricted to a bounded range, typical assumptions on the diameter of  $\mathcal{X}$  and the noise distribution amount to an effective truncation [26], [36].) The space  $\mathcal{X} \times [0, 1]$  is endowed with some fixed, unknown probability distribution  $\mu$ , and the learner observes  $n$  iid draws  $(X_i, Y_i) \sim \mu$ . The learner then seeks to fit the observed data with some hypothesis  $h \in \mathcal{H}$  so as to minimize the *risk*, usually defined as the expected loss  $\mathbb{E} |h(X) - Y|^q$  for  $(X, Y) \sim \mu$  and some  $q \geq 1$ . This is known in machine learning theory as the *agnostic* setting. The agnostic setting is considerably more general than the additive (typically Gaussian) noise model prevalent in statistics (see [42] for a recent paper on agnostic regression).

We consider two kinds of risk:  $L_1$  (mean absolute) and  $L_2$  (mean square). More precisely, for  $q \in \{1, 2\}$  we associate to each hypothesis  $h \in \mathcal{H}$  the empirical  $L_q$ -risk

$$R_n(h, q) = \frac{1}{n} \sum_{i=1}^n |h(X_i) - Y_i|^q, \quad (1)$$

and the (expected)  $L_q$ -risk

$$R(h, q) = \mathbb{E} |h(X) - Y|^q = \int_{\mathcal{X} \times [0, 1]} |h(x) - y|^q d\mu(x, y). \quad (2)$$

It is well-known that  $h(x) = \mathbb{M}[Y | X = x]$  (where  $\mathbb{M}$  is a median) minimizes  $R(\cdot, 1)$  over all integrable  $h \in [0, 1]^{\mathcal{X}}$  and  $h(x) = \mathbb{E}[Y | X = x]$  minimizes  $R(\cdot, 2)$ . However, these characterizations are of little practical use as neither is computable without knowledge of  $\mu$ . Hence, the standard route is to minimize the regularized empirical risk and provide generalization bounds for this procedure. A naive implementation of this approach requires solving a linear (or quadratic) program, which incurs a prohibitive  $O(n^3)$  time complexity.

*Our contribution:* Our approach to the regression problem departs from that of classical statistics in several important ways. Statistics has traditionally been more concerned with establishing minimax risk rates than with the computational efficiency (or even explicit construction) of the regression procedure. In contradistinction,

Ariel University, lee.ad.gottlieb@ariel.ac.il.

Ben-Gurion University, karyeh@cs.bgu.ac.il. This research was partially supported by the Israel Science Foundation (grant #1141/12) and a Yahoo Faculty award.

Weizmann Institute of Science, robert.krauthgamer@weizmann.ac.il. This work was supported in part by the Israel Science Foundation (grants #452/08 and #897/13), and by a Minerva grant.

our framework involves a precision parameter  $\eta$ , which controls the runtime-precision tradeoff. In particular, this means that Bayes-consistency is not achievable for  $\eta > 0$ . Further, our results rely on the structure of the metric space, but only to the extent of assuming that it has a low “intrinsic” dimensionality. Specifically, we consider the doubling dimension of  $\mathcal{X}$ , denoted  $\text{ddim}(\mathcal{X})$ , which was introduced by [19] based on earlier work of [1], [9], and has been since utilized in several algorithmic contexts, including networking, combinatorial optimization, and similarity search, see e.g. [23], [46], [31], [5], [21], [11], [10]. (A formal definition and typical examples appear in Section II.) Following the work of [16] on classification problems, our risk bounds and algorithmic runtime bounds are stated in terms of the doubling dimension of the data space and the Lipschitz constant of the regression hypothesis, although neither of these quantities need be known in advance. Note that any continuous function can be uniformly approximated by Lipschitz functions, with the Lipschitz constant as a measure of regularity — to which our algorithm adapts in a data-dependent fashion.

Our paper’s main contribution is computational. The algorithm in Theorem III.1 computes an  $\eta$ -additive approximation to the Lipschitz-regularized empirical risk minimizer in time  $\eta^{-O(\text{ddim}(\mathcal{X}))} n \ln^3 n$  (recall  $\eta > 0$  is a parameter that controls the desired precision). By Theorem IV.1, this hypothesis can be evaluated on new points in time  $\eta^{-O(\text{ddim}(\mathcal{X}))} \ln n$ . A novel feature of our construction is the use of a spanner to reduce the runtime of a linear program, and the spanner construction in Appendix A is itself of independent interest, having already been invoked in [13], [45]. We also present some statistical risk bounds (culminating in Theorem V.1).

A simple no-free-lunch argument shows that it is impossible to learn functions with arbitrary oscillation, and hence Lipschitzness is a natural and commonly used regularization constraint [47], [51], [43]. In this sense, our work fits into the so-called *luckiness* paradigm [44], of which SVM is a classic instance. Rather than guaranteeing a priori Bayes-consistency or excess risk bounds, luckiness bounds are data-dependent. Thus, in the case of SVM, a *lucky* sample is one that admits a large-margin separator; this in turn allows for optimistic generalization bounds — as opposed to a less lucky sample with a smaller margin and correspondingly more pessimistic bounds. More recently, this data-dependent approach was applied to general metric spaces [16] and was later shown to be Bayes-consistent [28].

Our runtime and generalization bounds explicitly depend on the doubling dimension of  $\mathcal{X}$ , but as we discuss in Remark 2, recent results with data-dependent general-

ization [17] renders our approach adaptive to the *intrinsic* dimension of the samples, offering large savings when the latter is even moderately smaller than the ambient metric dimension.

*Paper outline:* We start by defining the basic concepts in Section II. Our efficient model selection procedure is described in Section III, and the prediction algorithm (for a test point) is described in Section IV. The risk guarantees of our method are provided in Section V.

*Related work:* There are many excellent references for classical Euclidean nonparametric regression assuming iid noise, see for example [12], [20]. For metric regression, a simple risk bound follows from classic VC theory via the pseudo-dimension, see e.g. [40], [48], [39]. However, the pseudo-dimension of many natural function classes, including Lipschitz functions, is infinite — yielding a vacuous bound. An approach to nonparametric regression based on empirical risk minimization, though only for the Euclidean case, may already be found in [34]; see the comprehensive historical overview therein. Indeed, [20, Theorem 5.2] provides a kernel regressor for Lipschitz functions that achieves the minimax rate. Note however that (a) the setting is restricted to Euclidean spaces; and (b) the runtime cost of evaluating the hypothesis at a new point grows linearly with the sample size (while our complexity is roughly logarithmic).

More recently, risk bounds in terms of doubling dimension and Lipschitz constant were given in [29]. These results assumed an additive noise model, and hence are incomparable to ours. Following up, a regression technique based on random partition trees was proposed in [30], based on mappings between Euclidean spaces and also assuming an additive noise model. Another recent advance in nonparametric regression was Rodeo [33], which escapes the curse of dimensionality by adapting to the sparsity of the regression function. In contrast, our results apply to general metric spaces and exploit Lipschitz smoothness rather than sparsity.

Our work was inspired by the paper of von Luxburg and Bousquet [50], who established a connection between Lipschitz classifiers in metric spaces and large-margin hyperplanes in Banach spaces, thereby providing a novel generalization bound for nearest-neighbor classifiers. They developed a powerful statistical framework whose core idea may be summarized as follows: to predict the behavior at new points, find the smoothest function consistent with the training sample, and then extend the function to the new points. Since the regression function is defined implicitly by the labeled sample, the work of [50] raises natural algorithmic issues, such

as efficiently evaluating this function on test points (prediction) and performing model selection (Structural Risk Minimization) to avoid overfitting. Subsequent work (by the current authors) [16] leveraged the doubling dimension for both statistical and computational efficiency, and designed an efficient classifier for doubling metric spaces. Its key feature is an efficient algorithm to optimize the balance between the empirical risk and the penalty term for a given input. The present work extends these techniques from binary classification to real-valued regression, which presents a host of technical challenges.

## II. TECHNICAL BACKGROUND

We use standard notation and definitions throughout. The long-standing custom of ignoring measurability issues in learning-theoretic papers is more than justified in this case: we (effectively) only consider a class of functions computable to fixed precision by a fixed algorithm, and thus no loss of generality is incurred in treating this set of functions as countable. We write  $\ln$  for the natural logarithm and  $\log_b$  to specify a different base  $b$ .

*Metric spaces, Lipschitz constants:* A metric  $\rho$  on a set  $\mathcal{X}$  is a symmetric function that is positive (except for  $\rho(x, x) = 0$ ) and satisfies the triangle inequality  $\rho(x, y) \leq \rho(x, z) + \rho(z, y)$ ; together the two comprise the metric space  $(\mathcal{X}, \rho)$ . The diameter of a set  $A \subseteq \mathcal{X}$  is defined by  $\text{diam}(A) = \sup_{x, y \in A} \rho(x, y)$ . There is no loss of generality in assuming  $\text{diam}(\mathcal{X}) = 1$  since we can always scale the distances (when they are bounded). The *Lipschitz constant* of a function  $f : \mathcal{X} \rightarrow \mathbb{R}$ , denoted  $\|f\|_{\text{Lip}}$  (or  $\|f\|_{\text{Lip}(\rho)}$  if we wish to make the metric explicit) is defined to be the smallest  $L \geq 0$  such that  $|f(x) - f(y)| \leq L\rho(x, y)$  holds for all  $x, y \in \mathcal{X}$ . In addition to the metric  $\rho$  on  $\mathcal{X}$ , we will endow the space of all functions  $f : \mathcal{X} \rightarrow \mathbb{R}$  with the  $L_\infty$  metric:

$$\|f - g\|_\infty = \sup_{x \in \mathcal{X}} |f(x) - g(x)|.$$

A function is called *L-Lipschitz* if  $\|f\|_{\text{Lip}} \leq L$ . We will denote by  $\mathcal{H}_L$  the collection of all  $L$ -Lipschitz functions  $\mathcal{X} \rightarrow [0, 1]$ . It will occasionally be convenient to restrict this class to functions with  $\|f\|_{\text{Lip}} \geq 1$ ; the latter collection will be denoted by  $\mathcal{H}_{L \geq 1}$ . This incurs no loss of generality in our results, as our Structural Risk Minimization procedure in general selects hypotheses whose Lipschitz constant grows with sample size. (See for example the risk bound presented at the beginning of Section III.)

*Minkowski sums and perturbations:* If  $A, B$  are two families of functions mapping  $\mathcal{X}$  to  $\mathbb{R}$ , then their *Minkowski sum* is  $A \oplus B := \{a + b : a \in A, b \in B\}$ .

For  $\eta > 0$ , define  $\llbracket \eta \rrbracket := [-\eta, \eta]^\mathcal{X}$ . Hence,  $\mathcal{H}_L \oplus \llbracket \eta \rrbracket$  represents the collection of all  $[0, 1]$ -valued  $L$ -Lipschitz functions perturbed pointwise by at most  $\eta$ .

*Doubling dimension:* For a metric space  $(\mathcal{X}, \rho)$ , let  $\lambda > 0$  be the smallest value such that every ball in  $\mathcal{X}$  can be covered by  $\lambda$  balls of half the radius. The *doubling dimension* of  $\mathcal{X}$  is  $\text{ddim}(\mathcal{X}) = \log_2 \lambda$ . A metric space (or family of metrics) is called *doubling* if its doubling dimension is uniformly bounded. Note that while a low Euclidean dimension implies a low doubling dimension (Euclidean metrics of dimension  $d$  have doubling dimension  $O(d)$ ), low doubling dimension is strictly more general than low Euclidean dimension.

Doubling metric spaces occur naturally in many data analysis applications, including for instance the geodesic distance of a low-dimensional manifold residing in a possibly high-dimensional space assuming mild conditions, e.g., on curvature. Some concrete examples for doubling metric spaces include: (i)  $\mathbb{R}^d$  for fixed  $d$  equipped with an arbitrary norm, e.g.  $\ell_p$  or a mix between  $\ell_1$  and  $\ell_2$ ; (ii) the planar earthmover metric between point sets of fixed size  $k$  [16]; (iii) the  $n$ -cycle graph and its continuous version, the quotient  $\mathbb{R}/\mathbb{Z}$ , and similarly bounded-dimensional tori. In addition, various networks that arise in practice, such as peer-to-peer communication networks and online social networks, can be modeled reasonably well by a doubling metric space.

*Graph spanner:* A  $(1 + \delta)$ -stretch spanner for a graph  $G$  (which may have positive edge-lengths) is a subgraph  $H$  that contains all nodes of  $G$  (but not all edges), and  $\rho_H(u, v) \leq (1 + \delta)\rho_G(u, v)$  for all  $u, v \in G$ , where  $\rho_G(u, v)$  denotes the shortest-path distance between  $u$  and  $v$  in  $G$  (and similarly  $\rho_H(u, v)$  for  $H$ ). If a spanner  $H$  achieves this stretch bound even when  $\rho_H$  is evaluated only on paths in  $H$  with at most  $k$  edges, then  $H$  is called a  $(1 + \delta)$ -stretch  $k$ -hop spanner for  $G$ .

A spanner for a finite metric space  $\mathcal{X}$  is defined by viewing the metric space as a complete graph  $G$  on the vertex set  $\mathcal{X}$ , with edge-lengths corresponding to distances in  $\mathcal{X}$ . Doubling metrics are known to admit good spanners [8], [21], [18]. We will use a specific variant described in Appendix A.

## III. REGRESSION ALGORITHM

Let us fix the user-specified parameters  $q \in \{1, 2\}$  (risk type),  $\delta > 0$  (confidence level), and  $\eta > 0$  (precision parameter). Given the training sample  $(X_i, Y_i)_{i \in [n]}$ , our goal is to construct a hypothesis  $\tilde{h} : \mathcal{X} \rightarrow [0, 1]$  with small expected risk  $R(\tilde{h}, q)$ . Since the expected risk cannot be computed exactly (it depends on the

unknown distribution  $\mu$ ), we will instead seek to minimize an upper estimate of the risk. Theorem V.1 shows that with probability at least  $1 - \delta$ , for all  $\tilde{L} \geq 1$ ,  $\tilde{\eta} \in \{\eta, 2\eta, \dots, \eta \lceil 1/\eta \rceil, 1\}$  and hypothesis  $\tilde{h} \in \mathcal{H}_{\tilde{L}} \oplus [\tilde{\eta}]$  (that is,  $\tilde{h}$  is  $\tilde{\eta}$ -close to some  $\tilde{L}$ -Lipschitz function),

$$R(\tilde{h}, q) \leq R_n(\tilde{h}, q) + 4(2q - 1)\tilde{\eta} + (1 + o(1)) \sqrt{\frac{32 \ln \frac{1}{\delta}}{\tilde{\eta}^2}} \frac{1 + \text{ddim}(\mathcal{X})}{\tilde{\eta}} \quad (3)$$

Denote the RHS by  $Q(\tilde{h}, \tilde{L}, \tilde{\eta})$ ; when  $\tilde{L}, \tilde{\eta}$  are clear from the context, it may be convenient to write just  $Q(\tilde{h})$ . In this section, we design an algorithm to find a hypothesis that approximately minimizes  $Q(\tilde{h}, \tilde{L}, \tilde{\eta})$ . (A technique for quickly evaluating this hypothesis on new points is presented in Theorem IV.1.)

Suppose that for some training sample,  $Q(\cdot)$  is *minimized* by some  $(h^*, L^*, \eta^*)$ , where the minimum is taken over  $\eta^* \geq 0$ ,  $L^* \geq 1$ , and hypothesis  $h^* : \mathcal{X} \rightarrow [0, 1]$  that is  $\eta^*$ -close to some  $L^*$ -Lipschitz function.

**Theorem III.1.** *There is an algorithm that, given a precision parameter  $\eta \in (0, \frac{1}{4})$  and a training sample  $(X_i, Y_i) \in \mathcal{X} \times [0, 1]$ ,  $i \in [n]$ , computes  $\tilde{\eta} > 0$ ,  $\tilde{L} \geq 1$  and a hypothesis  $\tilde{h} : \mathcal{X} \rightarrow [0, 1]$ ,  $\tilde{h} \in \mathcal{H}_{\tilde{L}} \oplus [\tilde{\eta}]$  that satisfy*

$$Q(\tilde{h}, \tilde{L}, \tilde{\eta}) \leq Q(h^*, L^*, \eta^*) + \eta, \quad (3)$$

in time  $\eta^{-O(\text{ddim}(\mathcal{X}))} n \ln^3 n$ .

*Remark 1.* The role of the precision parameter  $\eta$  is to facilitate the construction of an approximate Lipschitz hypothesis with much greater efficiency than its exact Lipschitz counterpart. The bound (3) shows that the computed hypothesis  $\tilde{h}$  is competitive not only against any unperturbed Lipschitz hypothesis, but also against any  $\eta^*$ -perturbed hypothesis. Moreover, the pointwise  $\eta^*$ -perturbations might conspire to yield a lower empirical risk than unperturbed hypotheses. Theorem III.1 shows our approximate minimizer  $\tilde{h}$  is competitive even against an “optimally perturbed” hypothesis  $h^*$ .

The rest of this section is devoted to proving Theorem III.1 for  $q = 1$  (Sections III-A and III-B) and for  $q = 2$  (Section III-C). We consider the  $n$  observed samples as fixed values given as input to the algorithm (as opposed to random samples), so we will denote them  $(x_i, y_i)$  instead of  $(X_i, Y_i)$ . We will also restrict our attention to hypotheses for which  $Q(\cdot) < 1$ , since otherwise our bounds are vacuous. Indeed, the minimizer  $h^*$  must satisfy this condition, which holds even for the flat hypothesis mapping all points to  $\frac{1}{2}$  (for sufficiently large  $n$ ).

### A. Motivation and construction

We wish to find an optimal perturbed hypothesis  $h^* \in \mathcal{H}_{L^* \geq 1} \oplus [\eta^*]$  minimizing  $Q(\cdot)$ . Suppose that the Lipschitz and perturbation constants  $L^*, \eta^*$  of a minimizer  $h^*$  were known. Then the problem of computing both  $h^*$  and its empirical risk  $R_n(h^*, q)$  can be described as the following optimization program where the variables  $z_i$  representing the underlying smooth hypothesis of which  $h^*$  is an  $\eta^*$ -perturbation. Note that the optimization program is a Linear Program (LP) when  $q = 1$  and a quadratic program when  $q = 2$ .

$\begin{aligned} &\text{Minimize} && \frac{1}{n} \sum_{i \in [n]} w_i^q \\ &\text{subject to} &&  z_i - z_j  \leq L^* \cdot \rho(x_i, x_j) \quad \forall i, j \in [n] \\ &&& w_i \geq  y_i - z_i  - \eta^* \quad \forall i \in [n] \\ &&& 0 \leq z_i \leq 1 \quad \forall i \in [n] \\ &&& 0 \leq w_i \leq 1 \quad \forall i \in [n] \end{aligned} \quad (4)$
---

After solving the program for variables  $z_i$ , a minimizer  $h^*$  can easily be derived: If solution  $z_i$  is less than  $y_i$  then  $h^*(x_i) = \min\{z_i + \eta^*, y_i\}$ , and otherwise  $h^*(x_i) = \max\{z_i - \eta^*, y_i\}$ . It follows that  $h^*$  could be computed by first obtaining  $L^*$  and  $\eta^*$ , and then solving the above program. However, both computing  $L^*, \eta^*$  and solving the program appear to be expensive computations, which motivates our approximate solution. Note that supplying the LP with only a crude upper-bound on either  $L^*$  or  $\eta^*$  could yield a hypothesis with large Lipschitz constant or perturbation, and potentially poor generalization bounds. We show below how to derive relatively tight estimates for  $L^*, \eta^*$ , and in Section III-B we show how to solve the program quickly.

We first obtain a target perturbation constant  $\bar{\eta}$  that “approximates” the unknown  $\eta^*$ . In particular, we discretize candidate values of  $\bar{\eta}$  to be of the form  $i\eta$  for integral  $i \in [0, \lceil 1/\eta \rceil]$ , and search over all these values. (Recall that  $\eta$  is the input to Theorem III.1.) It follows that there are only  $O(1/\eta)$  candidates for  $\bar{\eta}$ , and that one of these candidates satisfies  $\eta^* \leq \bar{\eta} < \eta^* + \eta$ .

Next, we obtain a target Lipschitz constant  $\bar{L}$  that approximates  $L^*$ . Recall that we have assumed that  $L^* \geq 1$ , and also have that  $L^* < n$ , as otherwise the value of  $Q(h^*, L^*, \eta^*)$  is necessarily greater than 1. We discretize the candidate values of  $\bar{L}$  to be of the form  $\left(1 + \frac{\eta}{\text{ddim}(\mathcal{X})+1}\right)^i$  for integral  $i \geq 0$ , and search over all these values. It follows that there are only  $O\left(\frac{\text{ddim}(\mathcal{X})}{\eta} \ln n\right)$  discretized candidate values for  $\bar{L}$ , and that one of these candidates satisfies  $L^* \leq \bar{L} < \left(1 + \frac{\eta}{\text{ddim}(\mathcal{X})+1}\right) L^*$ . We note that

$$\left(1 + \frac{\eta}{\text{ddim}(\mathcal{X})+1}\right)^{\text{ddim}(\mathcal{X})+1} \leq e^\eta \leq 1 + 2\eta.$$

Now replace  $\eta^*, L^*$  in program (4) with approximations  $\bar{\eta}, \bar{L}$ , and let the hypothesis  $\bar{h}$  be an optimal solution for the modified program; this can only decrease the objective, i.e.,  $R_n(\bar{h}, q) \leq R_n(h^*, q)$ . Recall that  $Q(h^*) \leq 1$ , and so by the definition of  $Q(\cdot)$  and the above bounds on  $\bar{\eta}, \bar{L}$  we have

$$Q(\bar{h}) < Q(h^*) \cdot (1 + 2\eta) + 4\eta = Q(h^*) + 6\eta.$$

It remains to show that for each of the  $O\left(\frac{\text{ddim}(\mathcal{X})}{\eta^2} \ln n\right)$  candidate pairs of  $\bar{L}$  and  $\bar{\eta}$ , the modified linear program may be solved quickly (within fixed precision), which we do in Sections III-B and III-C.

### B. Solving the linear program

We show how to approximately solve the modified linear program, given target Lipschitz constant  $\bar{L}$  and perturbation parameter  $\bar{\eta}$  (recall  $\bar{h}$  is an optimal solution for this modified LP). Our solution will yield a hypothesis  $\tilde{h}$  satisfying

$$Q(\tilde{h}) \leq Q(\bar{h}) + O(\eta).$$

*Reduced constraints:* A central difficulty in obtaining a near-linear runtime for LP (4) is that the number of constraints is  $\Theta(n^2)$ ; in particular, there are  $\Theta(n^2)$  constraints of the form  $|z_i - z_j| \leq \bar{L} \cdot \rho(x_i, x_j)$ . We show how to reduce the number of these constraints (and only these constraints) to near-linear in  $n$ , namely,  $\eta^{-O(\text{ddim}(\mathcal{X}))} n$ . We will further guarantee that each of the  $n$  variables  $z_i$  appears in only  $\eta^{-O(\text{ddim}(\mathcal{X}))}$  constraints. Both these properties will prove useful for solving the program quickly.

Recall that the purpose of the  $\Theta(n^2)$  constraints is to ensure that the underlying hypothesis is smooth in the sense that the target Lipschitz constant is not violated between any pair of points. We show that this property can be approximately maintained with many fewer constraints. To see this, consider a  $1 + \delta$  stretch spanner for the point set, with spanner edge-set  $E$ . We claim that it suffices to enforce the Lipschitz condition  $\bar{L}$  only on pairs that are endpoints in  $E$ : Let  $x_{k_1}, x_{k_j}$  be any pair that are not connecting by a single in  $E$ , and let  $x_{k_2}, \dots, x_{k_{j-1}}$  be the vertices encountered on the minimum stretch path in  $E$  connecting  $x_{k_1}$  and  $x_{k_j}$ . Then by the stretch guarantee of the spanner and the Lipschitz condition on its endpoints we have

$$\frac{|y_{k_1} - y_{k_j}|}{\rho(x_{k_1}, x_{k_j})} \leq \frac{\sum_{i=1}^{j-1} |y_{k_i} - y_{k_{i+1}}|}{\rho(x_{k_1}, x_{k_j})} \leq \frac{\sum_{i=1}^{j-1} \bar{L} \rho(x_{k_i}, x_{k_{i+1}})}{\rho(x_{k_1}, x_{k_j})} \leq \frac{\bar{L}(1 + \delta) \sum_{i=1}^{j-1} \rho(x_{k_i}, x_{k_{i+1}})}{\rho(x_{k_1}, x_{k_j})} \leq \frac{\bar{L}(1 + \delta) \sum_{i=1}^{j-1} \rho(x_{k_i}, x_{k_j})}{\rho(x_{k_1}, x_{k_j})} \leq \frac{\bar{L}(1 + \delta)}{\bar{L}} \cdot \rho(x_i, x_j),$$

More formally, the constraints are reduced as follows: The spanner described in Appendix A has stretch  $1 + \delta$ , degree  $\delta^{-O(\text{ddim}(\mathcal{X}))}$  and hop-diameter  $c' \ln n$  for some

constant  $c' > 0$ , and can be computed quickly. Build this spanner for the observed sample points  $\{x_i : i \in [n]\}$  with stretch  $1 + \frac{\eta^2}{2}$  (i.e., set  $\delta = \frac{\eta^2}{2}$ ) and retain a constraint in LP (4) if and only if its two variables  $z_i, z_j$  correspond to two vertices connected by a spanner edge (that is, edge  $(x_i, x_j)$  is found in spanner's edge set  $E$ ). It follows from the bounded degree of the spanner that each variable appears in  $\eta^{-O(\text{ddim}(\mathcal{X}))}$  constraints, which implies that a total of  $\eta^{-O(\text{ddim}(\mathcal{X}))} n$  constraints are retained. Constructing the spanner (and thus the LP) takes time  $\eta^{-O(\text{ddim}(\mathcal{X}))} n \ln n$ . The complete analysis of the Lipschitz guarantee appears below.

*Fast LP-solver framework:* To solve the modified LP for fixed candidate values  $\bar{L}$  and  $\bar{\eta}$ , we utilize the framework presented by Young [53] for LPs of the following form: Given non-negative matrices  $P, C$ , vectors  $p, c$  and precision  $\beta > 0$ , find a non-negative vector  $x$  such that  $Px \leq p$  and  $Cx \geq c$ . Young shows that if there exists a feasible solution to the input instance, then a solution to a relaxation of the input program — specifically,  $Px \leq (1 + \beta)p$  and  $Cx \geq c$  — can be found in time  $O(md(\ln m)/\beta^2)$ , where  $m$  is the number of constraints in the program and  $d$  is the maximum number of constraints in which a single variable may appear. We may assume that constraints of the form  $0 \leq z_i \leq 1$  and  $0 \leq w_i \leq 1$  can be satisfied exactly: Since  $y_i \leq 1$ , we can always round down a solution variable to 1 without affecting the quality of the solution.

*Modifying the Lipschitz constraints:* In utilizing Young's framework for our problem, we encounter a difficulty that both the input matrices and output vector must be non-negative, while our LP (4) has difference constraints. To bypass this limitation, we first consider the LP variables  $z_i$ , and for each one introduce a new variable  $0 \leq \tilde{z}_i \leq 1$  and two new constraints:

$$\begin{aligned} z_i + \tilde{z}_i &\leq 1, \\ z_i + \tilde{z}_i &\geq 1. \end{aligned}$$

These constraints require that  $\tilde{z}_i = 1 - z_i$ , but by the relaxed guarantees of the LP solver, we have that in the returned solution  $1 - z_i \leq \tilde{z}_i \leq 1 - z_i + \beta$ . This technique allows us to introduce negated variables  $-z_i$  into the linear program, at the loss of additive precision.

Each retained spanner-edge constraint  $|z_i - z_j| \leq \bar{L} \cdot \rho(x_i, x_j)$  is replaced by a pair of constraints

Taken together, the above four constraints require that  $1 + |z_i - z_j| \leq 1 + \bar{L} \cdot \rho(x_i, x_j)$ . The modified program is found in (5).

Minimize	$\frac{1}{n} \sum_{i \in [n]} w_i$	
subject to	$1 \leq z_i + \tilde{z}_i \leq 1$	$\forall i \in [n]$
	$z_i + \tilde{z}_j \leq \bar{L} \cdot \rho(x_i, x_j)$	$\forall (x_i, x_j) \in E$
	$w_i \geq  y_i - z_i  - \bar{\eta}$	$\forall i \in [n]$
	$0 \leq z_i \leq 1$	$\forall i \in [n]$
	$0 \leq \tilde{z}_i \leq 1$	$\forall i \in [n]$
	$0 \leq w_i \leq 1$	$\forall i \in [n]$

Below we will address the objective function and the related constraint  $w_i \geq |y_i - z_i| - \bar{\eta}$ , and show that they can be modified to fit into Young's LP framework. But first, we will show that our modification of the Lipschitz constraints, along with the approximate guarantees of the LP solver, still yield a hypothesis that is close to Lipschitz:

Recall that in the returned solution of the LP solver,  $z_i \leq 1$ , and so necessarily  $|z_i - z_j| \leq 1$ . By the approximate guarantees of the LP solver, we have that in the returned solution to LP (5), each spanner edge constraint will satisfy

$$\begin{aligned} |z_i - z_j| &\leq \min\{1, -1 + (1 + \beta)[1 + \bar{L} \cdot \rho(z_i, z_j)]\} \\ &= \min\{1, \beta + (1 + \beta)\bar{L} \cdot \rho(z_i, z_j)\} \quad (6) \\ &\leq 2\beta + \bar{L} \cdot \rho(z_i, z_j), \end{aligned}$$

where the last inequality follows by splitting into two cases, depending on whether  $\bar{L} \cdot \rho(z_i, z_j) \leq 1$ .

To obtain a similar bound for point pairs not connected by a spanner edge: Let  $x_1, \dots, x_{k+1}$  be a  $(1 + \frac{\eta^2}{2})$ -stretch  $k$ -hop spanner path connecting points  $x_1$  and  $x_{k+1}$ , for  $k \leq c' \ln n$ ; then the stretch guarantee implies that  $\sum_{i=1}^k \rho(x_i, x_{i+1}) \leq (1 + \frac{\eta^2}{2})\rho(x_1, x_{k+1})$ . Using the triangle inequality and (6), and recalling the relaxed guarantees of the LP solver, we have that in the returned solution to LP (5)

$$\begin{aligned} |z_1 - z_{k+1}| &\leq \min\{1, \sum_{i=1}^k |z_i - z_{i+1}|\} \\ &\leq \min\{1, \sum_{i=1}^k [\beta + (1 + \beta)\bar{L} \cdot \rho(x_i, x_{i+1})]\} \\ &\leq \min\{1, \beta k + (1 + \beta)\bar{L} \cdot (1 + \eta^2/2) \rho(x_1, x_{k+1})\} \\ &\leq \min\{1, \beta(k + 1) + (1 + \eta^2/2)\bar{L} \cdot \rho(x_1, x_{k+1})\} \\ &\leq \beta(c' \ln n + 1) + \frac{\eta^2}{2} + \bar{L} \cdot \rho(x_1, x_{k+1}), \end{aligned}$$

where the fourth and fifth inequalities each follow by splitting into two cases.

Choosing  $\beta = \frac{\eta^2}{2c' \ln n + 1}$ , we have that for all point pairs in the returned solution to LP (5)

$$|z_i - z_j| \leq \eta^2 + \bar{L} \cdot \rho(z_i, z_j).$$

Now let  $h_z$  be the hypothesis mapping  $x_i$  to the value of  $z_i$  in the returned solution of the modified LP (5).

In the original LP (4), the variables  $z_i$  represented the  $L^*$ -Lipschitz underlying function. In the solver solution for LP (5), the variables  $z_i$  are a  $4\eta$ -perturbation of an  $\bar{L}$ -Lipschitz function:

**Lemma III.2.** *With  $h_z$  defined as above,  $h_z \in \mathcal{H}_{\bar{L}} \oplus \llbracket 4\eta \rrbracket$ .*

**Proof:** Let us construct a function  $\tilde{h}_z$  as follows: Let  $S$  be the sample points  $\{x_i\}_{i \in [n]}$ , and extract from  $S$  an  $\eta/\bar{L}$ -net  $N$ .<sup>1</sup> For every net-point  $v \in N$  set  $\tilde{h}_z(v) = \frac{h_z(v)}{1+\eta}$ . Then extend hypothesis  $\tilde{h}_z$  from  $N$  to all of the sample  $S$  without increasing Lipschitz constant by using the McShane-Whitney extension theorem [35], [52] for real-valued functions.<sup>2</sup> This completes the description of  $\tilde{h}_z$ .

We first show that  $\|\tilde{h}_z\|_{\text{Lip}} \leq \bar{L}$ . Indeed, for every two net-points  $v \neq v' \in N$  we have  $\rho(v, v') \geq \eta/\bar{L}$  and so

$$\begin{aligned} |\tilde{h}_z(v) - \tilde{h}_z(v')| &= \frac{|\tilde{h}_z(v) - \tilde{h}_z(v')|}{1+\eta} \\ &\leq \frac{\eta^2 + \bar{L} \cdot \rho(v, v')}{1+\eta} \\ &\leq \bar{L} \cdot \rho(v, v'). \end{aligned}$$

It follows that  $\tilde{h}_z$  indeed satisfies the  $\bar{L}$ -Lipschitz condition on the net-points. By the extension theorem,  $\tilde{h}_z$  achieves Lipschitz constant  $\bar{L}$  on all points of  $S$ .

It remains to show that  $\|h_z - \tilde{h}_z\|_{\infty} \leq 4\eta$ : Consider any point  $x \in S$  and its closest net-point  $v \in N$ ; then  $\rho(x, v) < \eta/\bar{L}$  and we have

$$\begin{aligned} |h_z(x) - \tilde{h}_z(x)| &\leq |h_z(x) - h_z(v)| + |h_z(v) - \tilde{h}_z(v)| + |\tilde{h}_z(v) - \tilde{h}_z(x)| \\ &< [\eta^2 + \bar{L} \cdot \frac{\eta}{\bar{L}}] + [1 - \frac{1}{1+\eta}] + [\eta^2 + \bar{L} \cdot \frac{\eta}{\bar{L}}] \cdot \frac{1}{1+\eta} \\ &< 4\eta, \end{aligned}$$

and we conclude that  $h_z$  is a  $4\eta$ -perturbation of  $\tilde{h}_z$ .  $\square$

*Modifying the objective function:* We now turn to the constraints  $w_i \geq |y_i - z_i| - \bar{\eta}$  and the objective function  $\frac{1}{n} \sum_{i \in [n]} w_i$ . Each LP constraint is replaced by a constraint pair

$$\begin{aligned} w_i + z_i &\geq y_i - \bar{\eta}, \\ w_i + \tilde{z}_i &\geq 1 - y_i - \bar{\eta}, \end{aligned}$$

and together these require that  $w_i \geq |y_i - z_i| - \bar{\eta}$ . Note however that in the returned solution we are guaranteed only that  $w_i \geq |y_i - z_i| - \bar{\eta} - \beta$ . Hence, the empirical

<sup>1</sup>The notion of a net referred to here means that (i) the distance between every two points in  $N$  is at least  $\eta/\bar{L}$ ; and (ii) every point in  $S$  is within distance  $\eta/\bar{L}$  from at least one point in  $N$ . It can be easily constructed by a greedy process.

<sup>2</sup>The McShane-Whitney extension theorem says that for every metric space  $M$  and subset  $N \subset M$ , every  $L$ -Lipschitz  $f : N \rightarrow \mathbb{R}$  can be extended to all of  $M$  while preserving the  $L$ -Lipschitz condition.

error of the hypothesis is bounded by  $\beta + \frac{1}{n} \sum_{i \in [n]} w_i$  instead of  $\frac{1}{n} \sum_{i \in [n]} w_i$ .

The objective function is replaced by the constraint

$$\frac{1}{n} \sum_{i \in [n]} w_i \leq r,$$

where  $r$  itself is guessed by discretizing into multiples of  $\eta$  — that is  $\bar{r} = i\eta^2$  for integral  $i \in [1, \lceil 1/\eta^2 \rceil]$  — which gives  $O(1/\eta^2)$  candidate values for  $r$ . By the discretization of  $r$ , the relaxed guarantees of the LP solver, and the above bound on the empirical error, the empirical error of the solution hypothesis  $\tilde{h}$  is within an additive term  $\eta^2 + \beta + \beta < 2\eta^2$  of optimal. The final program is found in (7).

Find	$0 \leq z_i \leq 1$	$\forall i \in [n]$
	$0 \leq \tilde{z}_i \leq 1$	$\forall i \in [n]$
subject to	$0 \leq w_i \leq 1$	$\forall i \in [n]$
	$1 \leq z_i + \tilde{z}_i \leq 1$	$\forall i \in [n]$
	$z_i + \tilde{z}_j \leq \bar{L} \cdot \rho(x_i, x_j)$	$\forall (x_i, x_j) \in E$
	$\sum_i w_i \leq \bar{r}$	$\forall i \in [n]$
	$w_i + z_i \geq y_i - \bar{\eta}$	$\forall i \in [n]$
	$w_i + \tilde{z}_i \geq 1 - y_i - \bar{\eta}$	$\forall i \in [n]$

*Correctness and runtime analysis:* Consider the choice of  $\bar{L}, \bar{\eta}$ , closest to the values  $L^*, \eta^*$ , and recall that for these values there exists a hypothesis  $\bar{h} \in \mathcal{H}_{\bar{L} \geq 1} \oplus \llbracket \bar{\eta} \rrbracket$  satisfying

$$Q(\bar{h}, \bar{L}, \bar{\eta}) < Q(h^*, L^*, \eta^*) + 6\eta.$$

As shown above, running program (7) on this  $\bar{L}, \bar{\eta}$ , we obtain a hypothesis  $\tilde{h} \in \mathcal{H}_{\bar{L}} \oplus \llbracket 4\eta + \bar{\eta} \rrbracket$  whose empirical risk is within an additive term  $2\eta^2$  of the empirical risk of the optimal  $h^*$ . It follows that

$$Q(\tilde{h}, \bar{L}, 4\eta + \bar{\eta}) \leq Q(\bar{h}, \bar{L}, \bar{\eta}) + 2\eta^2 + 4\eta \leq Q(h^*, L^*, \eta^*) + 11\eta.$$

The result claimed in Theorem III.1 is achieved, up to scaling  $\eta$ , i.e., applying the above for  $\eta = \eta_1/11$ , by exhaustively trying all pairs of candidates  $\bar{L}, \bar{\eta}$  and picking the pair that minimizes  $Q(\cdot)$ .

We turn to analyze the algorithmic runtime. Recall that the spanner can be constructed in time  $O(\eta^{-O(\text{ddim}(\mathcal{X}))} n \ln n)$ . Young's LP solver [53] is invoked on  $O(\eta^{-2 \text{ddim}(\mathcal{X})} \ln n)$  pairs of  $\bar{L}, \bar{\eta}$  and  $O(1/\eta^2)$  candidate values of  $\bar{r}$ , for a total of  $O(\eta^{-4 \text{ddim}(\mathcal{X})} \ln n)$  times. To determine the runtime per invocation, recall that each variable of the program appears in  $d = \eta^{-O(\text{ddim}(\mathcal{X}))}$  constraints, implying that there are in total  $m = \eta^{-O(\text{ddim}(\mathcal{X}))} n$  constraints. Since we set  $\beta = O(\eta^2/\ln n)$ , we have that each call to the solver takes time

$O(m d (\ln m) / \beta^2) \leq \eta^{-O(\text{ddim}(\mathcal{X}))} n \ln^2 n$ , and the total runtime is  $\eta^{-O(\text{ddim}(\mathcal{X}))} n \ln^3 n$ . This completes the proof of Theorem III.1 for  $q = 1$ .

### C. Solving the quadratic program

We proceed to the case of a quadratic loss function, i.e.,  $q = 2$  in our original program (4). A recent line of work on fast solvers for Laplacian systems and for electrical flows, see e.g. [49, Sections 3 and 11], provides powerful algorithms that can speed up Laplacian-based machine-learning tasks [54]. However, these algorithms are not directly applicable here, because our quadratic program (4) contains hard non-quadratic constraints to enforce a Lipschitz-constant bound  $L^*$ . In fact, our program can be viewed as minimizing simultaneously the  $\ell_\infty$ -Laplacian on the graph edges and some  $\ell_2$ -Laplacian related to the point values. See also [32] for a discussion of Lipschitz extension on graphs and additional references.

Our approach is to modify the methodology we developed above for linear loss, to cover the case of a quadratic loss function  $\frac{1}{n} \sum_i w_i^2$ . Specifically, we introduce variables  $v_i \geq w_i^2$ , and replace the objective function with  $\frac{1}{n} \sum_i v_i$ . It remains to show how to model the constraints  $v_i \geq w_i^2$ .

First consider a parabola  $y = x^2$ , and note that a line  $y = (2a)x - a^2$  is tangent to the parabola, intersecting it at  $x = a$ . Hence, the constraint  $v_i \geq w_i^2$  can be approximated by a constraint set  $v_i \geq (2a)w_i - a^2$  for  $a = i\eta$  and integral  $i \in [0, \lceil 1/\eta \rceil]$ . These lines have slope in the range  $[0, 2]$ , and so the approximation may cause the value of  $v_i$  to be underestimated by  $2\eta$ . This is in addition to the previous underestimate of  $w_i$ , and by the above scaling of  $\eta$  this maintains the asymptotic error guarantee of the theorem. Turning to the runtime analysis, the replacement of a single constraint by  $O(1/\eta)$  new constraints does not change the asymptotic runtime.

## IV. APPROXIMATE LIPSCHITZ EXTENSION

In this section, we show how to evaluate our hypothesis on a new point. We take the underlying smooth hypothesis on set  $S$  implicit in Lemma III.2 — call it  $\tilde{h}_z(\cdot)$  — and we wish to evaluate a minimum Lipschitz extension of  $\tilde{h}_z$  on a new point  $x \notin S$ . That is, denoting  $S = \{x_1, \dots, x_n\}$ , we wish to return a value  $y = \tilde{h}_z(x)$  that minimizes  $\max_{i \in [n]} \frac{|y - \tilde{h}_z(x_i)|}{\rho(x, x_i)}$ . By the McShane-Whitney extension theorem, the extension of  $\tilde{h}_z$  to the new point does not increase the Lipschitz

constant of  $\tilde{h}_z$ , and so the risk bound in Theorem V.1 applies.<sup>3</sup>

First note that the Lipschitz extension label  $y$  of  $x \notin S$  will be determined by a pair of points of  $S$ : There exist points  $x_i, x_j \in S$ , one with label greater than  $y$  and one with a label less than  $y$ , such that the Lipschitz constant of  $x$  relative to each of these points (that is,  $L = \frac{\tilde{h}_z(x_i) - y}{\rho(x, x_i)} = \frac{y - \tilde{h}_z(x_j)}{\rho(x, x_j)}$ ) is maximum over the Lipschitz constant of  $x$  relative to any point in  $S$ . Hence,  $y$  cannot be increased or decreased without increasing the Lipschitz constant with respect to one of these points. Hence, an exact Lipschitz extension may be computed in  $\Theta(n^2)$  time in brute-force fashion, by enumerating all point pairs in  $S$ , calculating the optimal Lipschitz extension for  $x$  with respect to each pair alone, and then choosing the candidate value for  $y$  with the highest Lipschitz constant. However, we demonstrate that an approximate solution to the Lipschitz extension problem can be obtained more efficiently.

**Theorem IV.1.** *An  $\eta$ -additive approximation to the Lipschitz extension problem on a function  $f : S \rightarrow [0, 1]$  can be computed in time  $\left(\frac{1}{\eta}\right)^{-O(\text{ddim}(\mathcal{X}))} \ln n$ .*

**Proof:** The algorithm is as follows. Round up all labels  $f(x_i)$  to the nearest multiple of  $j\eta/2$  (for any integer  $0 \leq j \leq 2/\eta$ ), and call the new label function  $\tilde{f}$ . We seek the value of  $\tilde{f}(x)$ , the value at point  $x$  of the optimal Lipschitz extension function  $\tilde{f}$ . Trivially,  $f(x) \leq \tilde{f}(x) \leq f(x) + \eta/2$ . Now, if we were given for each  $j$  the point with label  $j\eta/2$  that is the nearest neighbor of  $x$  (among all points with this label), then we could run the brute-force algorithm described above on these  $2/\eta$  points in time  $O(1/\eta^2)$  and compute  $\tilde{f}(x)$ . However, exact metric nearest neighbor search is potentially expensive, and so we cannot find these points efficiently. We instead find for each  $j$  a point  $x' \in S$  with label  $\tilde{f}(x') = j\eta/2$  that is a  $(1 + \frac{\eta}{2})$ -approximate nearest neighbor of  $x$  among points with this label. (This can be done by presorting the points of  $S$  into  $2/\eta$  buckets based on their  $\tilde{f}$  label, and once  $x$  is received, running on each bucket a  $(1 + \frac{\eta}{2})$ -approximate nearest neighbor search algorithm due to [11] that takes  $(1/\eta)^{O(\text{ddim}(\mathcal{X}))} \ln n$  time.) We then run the brute force algorithm on these  $2/\eta$  points in time  $O(1/\eta^2)$ . The nearest neighbor search achieves approximation factor

<sup>3</sup>Theorems IV.1 and V.1 are “local” in the following sense. At a test point  $x$ , Theorem IV.1 returns the value  $h(x)$ , where  $h : \mathcal{X} \rightarrow [0, 1]$  is an  $\eta$ -perturbed  $L$ -Lipschitz function. At a different test point  $x'$ , a different  $h' : \mathcal{X} \rightarrow [0, 1]$  is evaluated. There is no consistency requirement between  $h$  and  $h'$  — there need not exist any  $\eta$ -perturbed  $L$ -Lipschitz function  $h''$  such that  $h''(x) = h(x)$  and  $h''(x') = h'(x')$ .

$1 + \frac{\eta}{2}$ , implying a similar multiplicative approximation to  $L$ , and thus also to  $|y - f(x')| \leq 1$ , which means at most  $\eta/2$  additive error in the value  $y$ . We conclude that the algorithm’s output solves the Lipschitz extension problem within additive approximation  $\eta$ .  $\square$

## V. RISK BOUNDS

The algorithm in Section III produces a hypothesis  $h : \mathcal{X} \rightarrow [0, 1]$ ,<sup>4</sup> which is an  $\bar{\eta}$ -perturbation of some hypothesis in  $H_L$  (the notation there was  $\tilde{h}$  and  $\tilde{L}$ ). Recalling the definitions of empirical risk and expected risk in (1) and (2), this section is devoted to proving that with high probability,  $R(h, q)$  is not much greater than  $R_n(h, q)$ .

**Theorem V.1.** *Fix  $q \in \{1, 2\}$ ,  $\eta \in (0, 1]$ , and  $\bar{\eta} \in \{\eta, 2\eta, \dots, \eta \lfloor 1/\eta \rfloor, 1\}$ . Then for all  $\delta > 0$ , with probability at least  $1 - \delta$ , the following holds uniformly for all  $L \geq 1$  and all  $h \in \mathcal{H}_L \oplus \llbracket \bar{\eta} \rrbracket$ :*

$$R(h, q) \leq R_n(h, q) + 4(2q - 1)\bar{\eta} + \sqrt{\frac{32 \ln \frac{8}{(2q-1)\bar{\eta}}}{n}} \left( \frac{16q^{3/2}L}{(2q-1)\bar{\eta}} \right)^{1+d} + \sqrt{\frac{\ln \log_2(2L^{1+\text{ddim}(\mathcal{X}))}}{n}} + 3\sqrt{\frac{\ln \frac{4}{\delta\eta}}{2n}}.$$

The proof of Theorem V.1 proceeds in two conceptual steps. We first bound the covering numbers for classes of Lipschitz functions (in Section V-A) and then use those to estimate Rademacher complexities (in Section V-B).

### A. Covering numbers for Lipschitz function classes

We begin by obtaining complexity estimates for Lipschitz functions in doubling spaces. In the conference version [15] this was done in terms of the fat-shattering dimension, but here we obtain considerably simpler and tighter bounds by direct control over the covering numbers.

The following variant of the classic “covering numbers by covering numbers” estimate [24] was proved together with Roi Weiss (cf. [27, Lemma 2]):

**Lemma V.2.** *Let  $\mathcal{F}_L$  be the collection of  $L$ -Lipschitz functions mapping the metric space  $(\mathcal{X}, \rho)$  to  $[0, 1]$ . Then the covering numbers of  $\mathcal{F}_L$  may be estimated in terms of the covering numbers of  $\mathcal{X}$ :*

$$\mathcal{N}(\varepsilon, \mathcal{F}_L, \|\cdot\|_\infty) \leq \left( \frac{8}{\varepsilon} \right)^{\mathcal{N}(\varepsilon/8L, \mathcal{X}, \rho)}.$$

<sup>4</sup> Since  $Y_i \in [0, 1]$ , there is no loss in assuming that the hypothesis also has this range; this is trivially ensured by a truncation, which preserves the Lipschitz constant.



Hence, for doubling spaces with  $\text{diam}(\mathcal{X}) = 1$ ,

$$\ln \mathcal{N}(\varepsilon, \mathcal{F}_L, \|\cdot\|_\infty) \leq \left(\frac{16L}{\varepsilon}\right)^{\text{ddim}(\mathcal{X})} \ln\left(\frac{8}{\varepsilon}\right).$$

**Proof:** Fix a covering of  $\mathcal{X}$  consisting of  $|N| = \mathcal{N}(\varepsilon/8L, \mathcal{X}, \rho)$  balls  $\{U_1, \dots, U_{|N|}\}$  of radius  $\varepsilon' = \varepsilon/8L$  and choose  $|N|$  points  $N = \{x_i \in U_i\}_{i=1}^{|N|}$ . We will construct an  $\varepsilon$ -cover  $\hat{F} = \{\hat{f}_1, \dots, \hat{f}_{|\hat{F}|}\}$  as follows.

At every point  $x_i \in N$ , we choose  $\hat{f}(x_i)$  to be some multiple of  $2L\varepsilon' = \varepsilon/4$ , while maintaining  $\|\hat{f}\|_{\text{Lip}} \leq 2L$ . Construct a  $2L$ -Lipschitz extension for  $\hat{f}$  from  $N$  to all over  $\mathcal{X}$  (such an extension always exists, [35], [52]).

We claim that every  $f \in \mathcal{F}_L$  is close to some  $\hat{f} \in \hat{F}$ , in the sense that  $\|f - \hat{f}\|_\infty \leq \varepsilon$ . Indeed, every point  $x \in \mathcal{X}$  is  $\varepsilon'$ -close to some point  $x_N \in N$ , and since  $f$  is  $L$ -Lipschitz and  $\hat{f}$  is  $2L$ -Lipschitz,

$$\begin{aligned} |f(x) - \hat{f}(x)| &\leq |f(x) - f(x_N)| + |f(x_N) - \hat{f}(x_N)| + \left| \frac{\hat{f}(x_N) - \hat{f}(x)}{d_q((x, y), (x', y'))} \right| = (\rho(x, x_N)^q + |y - y'|^q)^{1/q}. \\ &\leq L \cdot \rho(x, x_N) + \varepsilon/4 + 2L \cdot \rho(x, x_N) = \varepsilon. \end{aligned} \quad (10)$$

It is easy to verify that  $|\hat{F}| \leq (8/\varepsilon)^{|N|}$ , since by construction, the functions  $\hat{f}$  are determined by their values on  $N$ . This provides a covering of  $\mathcal{F}_L$  using  $|\hat{F}|$  balls of radius  $\varepsilon$ .

The bound for doubling spaces follows immediately by applying the so-called doubling property (see for example [31]) and the diameter bound, to obtain

$$\mathcal{N}(\varepsilon, \mathcal{X}, \rho) \leq \left(\frac{2}{\varepsilon}\right)^{\text{ddim}(\mathcal{X})}.$$

□

Let us consider two additional properties that a metric space  $(\mathcal{X}, \rho)$  might possess:

- 1)  $(\mathcal{X}, \rho)$  is *connected* if for all  $x, x' \in \mathcal{X}$  and all  $\varepsilon > 0$ , there is a finite sequence of points  $x = x_1, x_2, \dots, x_m = x'$  such that  $\rho(x_i, x_{i+1}) < \varepsilon$  for all  $1 \leq i < m$ .
- 2)  $(\mathcal{X}, \rho)$  is *centered* if for all  $r > 0$  and all  $A \subset \mathcal{X}$  with  $\text{diam}(A) \leq 2r$ , there exists a point  $x \in \mathcal{X}$  such that  $\rho(x, a) \leq r$  for all  $a \in A$ .

The estimate in Lemma V.2 may be improved for doubling spaces that are additionally connected and centered, as follows.

**Lemma V.3** ([24]). *If  $(\mathcal{X}, \rho)$  is connected and centered, then, for constant  $\text{ddim}(\mathcal{X})$ ,*

$$\ln \mathcal{N}(\varepsilon, \mathcal{F}_L, \|\cdot\|_\infty) = O\left(\left(\frac{L}{\varepsilon}\right)^{\text{ddim}(\mathcal{X})} + \ln\left(\frac{1}{\varepsilon}\right)\right).$$

## B. Rademacher complexities

The (empirical) *Rademacher complexity* [4], [25] of a collection of functions  $\mathcal{F}$  mapping some set  $\mathcal{Z}$  to  $\mathbb{R}$  is defined, with respect to a sequence  $Z = (Z_i)_{i \in [n]} \in \mathcal{Z}^n$ , by

$$\widehat{\mathcal{R}}_n(\mathcal{F}; Z) = \mathbb{E} \left[ \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \sigma_i f(Z_i) \right], \quad (8)$$

where the expectation is over the  $\sigma_i$ , which are iid with  $\mathbb{P}(\sigma_i = +1) = \mathbb{P}(\sigma_i = -1) = 1/2$ .

To any collection  $\mathcal{G}$  of hypotheses mapping  $\mathcal{X}$  to  $\mathbb{R}$ , we associate the *q-loss class*, whose members map  $\mathcal{X} \times \mathbb{R}$  to  $\mathbb{R}$ . The latter is denoted by  $q \circ \mathcal{G}$  and defined to be

$$q \circ \mathcal{G} = \{f : (x, y) \mapsto |g(x) - y|^q; g \in \mathcal{G}\}. \quad (9)$$

It will also be convenient to define the auxiliary metric space  $(\mathcal{Z}, d_q)$ , where  $\mathcal{Z} = \mathcal{X} \times [0, 1]$  and

Let us recall the relevance of Rademacher complexities to risk estimates [37, Theorem 3.1]: for every  $\delta > 0$ , we have that, with probability at least  $1 - \delta$ ,

$$R(g, q) \leq R_n(g, q) + 2\widehat{\mathcal{R}}_n(q \circ \mathcal{G}; Z) + 3\sqrt{\frac{\ln(2/\delta)}{2n}}, \quad (11)$$

holds uniformly over all  $g \in \mathcal{G}$ , where  $Z = (X_i, Y_i)_{i \in [n]}$  is the training sample.

The following simple and well-known estimate of Rademacher complexity is obtained via covering numbers; see, e.g., [3, Theorem 1.1] for the proof of a closely related fact.

**Lemma V.4.** *For all function classes  $\mathcal{F} \subset [0, 1]^{\mathcal{Z}}$ , all  $Z \in \mathcal{Z}^n$ , and all  $\varepsilon > 0$ ,*

$$\widehat{\mathcal{R}}_n(\mathcal{F}; Z) \leq \varepsilon + \sqrt{\frac{2 \ln \mathcal{N}(\varepsilon, \mathcal{F}, \|\cdot\|_\infty)}{n}}.$$

Having reduced the problem to one of estimating covering numbers, we would like to invoke results from Section V-A, such as Lemma V.2. The following result sheds light on the relation between  $\mathcal{H}_L$  and its loss class. Its proof appears in Appendix B.

**Lemma V.5.** *Let  $(\mathcal{Z}, d_q)$  be as defined in (10) and  $q \in \{1, 2\}$ . The following relations hold:*

- (i) *if  $f \in q \circ \mathcal{H}_{L \geq 1}$  with witness  $h \in \mathcal{H}_{L \geq 1}$ , then  $\|f\|_{\text{Lip}(d_q)} \leq q^{3/2} \|h\|_{\text{Lip}(\rho)}$ ,*
- (ii)  *$\text{ddim}(\mathcal{Z}, d_q) \leq 2 + 2\text{ddim}(\mathcal{X}, \rho)$ .*

We are ready to prove an ‘‘unperturbed’’ version of Theorem V.1, as follows.

**Theorem V.6.** For  $q \in \{1, 2\}$ ,  $L \geq 1$ , and  $0 < \delta < 1$ , with probability at least  $1 - \delta$ , the following holds uniformly over all  $h \in \mathcal{H}_L$ :

$$R(h, q) - R_n(h, q) \leq 3\sqrt{\frac{\ln \frac{2}{\delta}}{2n}} + 2 \inf_{\varepsilon > 0} \left[ \varepsilon + \sqrt{\frac{2 \ln \frac{8}{\varepsilon}}{n}} \left( \frac{16q^{3/2}L}{\varepsilon} \right)^{1+\text{ddim}(\mathcal{X})} \right]$$

**Proof:** Let  $Z = (X_i, Y_i)_{i \in [n]}$  be the training sample and fix some  $L \geq 1$  and  $\varepsilon > 0$ . We begin by applying (11) to  $\mathcal{G} = \mathcal{H}_L$ , and get that with probability at least  $1 - \delta$ , uniformly for all hypotheses  $h \in \mathcal{H}_L$ ,

$$R(h, q) \leq R_n(h, q) + 2\widehat{\mathcal{R}}_n(q \circ (\mathcal{H}_L); Z) + 3\sqrt{\frac{\ln(2/\delta)}{2n}}.$$

Further,

$$\begin{aligned} \widehat{\mathcal{R}}_n(q \circ (\mathcal{H}_L); Z) &\leq \varepsilon + \sqrt{\frac{2 \ln \mathcal{N}(\varepsilon, q \circ \mathcal{H}_L, \|\cdot\|_\infty)}{n}} \\ &\leq \varepsilon + \sqrt{\frac{2 \left( \frac{16q^{3/2}L}{\varepsilon} \right)^{2+2\text{ddim}(\mathcal{X})} \ln \frac{8}{\varepsilon}}{n}} \\ &= \varepsilon + \sqrt{\frac{2 \ln \frac{8}{\varepsilon}}{n} \left( \frac{16q^{3/2}L}{\varepsilon} \right)^{1+\text{ddim}(\mathcal{X})}} \end{aligned}$$

where the first inequality follows from Lemma V.4 and the second one by applying the covering number estimate in Lemma V.2 to  $q \circ \mathcal{H}_L$ , after the appropriate ‘‘conversion’’ of Lipschitz constants and doubling dimensions furnished by Lemma V.5.  $\square$

For completeness, we relate the empirical risk to the optimal risk.

**Corollary V.7.** Fix  $q \in \{1, 2\}$ ,  $L \geq 1$ , and  $0 < \delta < 1$ , and define

$$\begin{aligned} \widehat{R}_n(q) &:= \inf_{h \in \mathcal{H}_L} R_n(h, q), \\ R^*(q) &:= \inf_{h \in \mathcal{H}_L} R(h, q). \end{aligned}$$

Then

$$\widehat{R}_n(q) - R^*(q) \leq 3\sqrt{\frac{\ln \frac{2}{\delta}}{2n}} + 2 \inf_{\varepsilon > 0} \left[ \varepsilon + \sqrt{\frac{2 \ln \frac{8}{\varepsilon}}{n}} \left( \frac{16q^{3/2}L}{\varepsilon} \right)^{1+\text{ddim}(\mathcal{X})} \right]$$

holds with probability at least  $1 - \delta$ .

**Proof:** It will be convenient to denote the right-hand side of (12) by  $\Delta(\delta)$  and to assume, without loss of generality, the existence of minimizers  $\hat{h}_n$  and  $h^*$  of  $R_n(\cdot, q)$  and  $R(\cdot, q)$ , respectively, over  $\mathcal{H}_L$ ; this is justified via a standard approximation argument. A standard symmetrization argument (e.g., swapping  $\Phi(S)$

and  $\Phi(S')$  in [37, Eq. (3.6)]) shows that the estimate of Theorem V.6 holds in the other direction as well:

$$\begin{aligned} \mathbb{P} \left( \sup_{h \in \mathcal{H}_L} R_n(h, q) - R(h, q) > \Delta(\delta) \right) &\leq \delta. \\ \widehat{R}_n(q) - R^*(q) &= R_n(\hat{h}_n, q) - R(h^*, q) \\ &\leq R_n(h^*, q) - R(h^*, q), \end{aligned}$$

whence

$$\mathbb{P} \left( \widehat{R}_n(q) - R^*(q) \leq \Delta(\delta) \right) \geq 1 - \delta.$$

$\square$

To extend Theorem V.6 to perturbed hypotheses, we will need the following decomposition, whose proof appears in Appendix B.

**Lemma V.8.** If  $\eta > 0$  and  $\mathcal{H}$  is any collection of functions mapping  $\mathcal{X}$  to  $[0, 1]$ , then

$$q \circ (\mathcal{H} \oplus [\eta]) \subseteq (q \circ \mathcal{H}) \oplus [(2q - 1)\eta].$$

**Corollary V.9.** For  $q \in \{1, 2\}$ ,  $L \geq 1$ ,  $\eta > 0$ , and  $0 < \delta < 1$ , with probability at least  $1 - \delta$ , the following holds uniformly over all  $h \in \mathcal{H}_L \oplus [\eta]$ :

$$R(h, q) \leq 4(2q - 1)\eta + R_n(h, q) + \sqrt{\frac{8 \ln \frac{8}{(2q-1)\eta}}{n}} \left( \frac{16q^{3/2}L}{(2q-1)\eta} \right)^{1+\text{ddim}(\mathcal{X})}$$

**Proof:** For any sequence  $Z = (X_i, Y_i)_{i \in [n]}$ , we have

$$\begin{aligned} \widehat{\mathcal{R}}_n(q \circ (\mathcal{H}_L \oplus [\eta]); Z) &\leq \widehat{\mathcal{R}}_n((q \circ \mathcal{H}_L) \oplus [(2q - 1)\eta]; Z) \\ &\leq \widehat{\mathcal{R}}_n(q \circ \mathcal{H}_L; Z) + (2q - 1)\eta \\ &\leq 2(2q - 1)\eta + \sqrt{\frac{2 \ln \frac{8}{(2q-1)\eta}}{n}} \left( \frac{16q^{3/2}L}{(2q-1)\eta} \right)^{1+\text{ddim}(\mathcal{X})} \end{aligned}$$

where the first inequality follows from Lemma V.8, the second from the sub-additivity of Rademacher complexities ([6, Theorem 3.3]), and the third from (13) (with  $\varepsilon = (2q - 1)\eta$ ). Invoking (11) to bound the risk in terms of  $\widehat{\mathcal{R}}_n$  completes the proof.  $\square$

**Proof:** [Proof of Theorem V.1] In light of Corollary V.9, it only remains to extend the risk bound from a fixed  $(L, \bar{\eta})$  to hold uniformly over all  $L \geq 1$  and  $\bar{\eta} \in \{\eta, 2\eta, \dots, \eta \lfloor 1/\eta \rfloor, 1\}$ . This is carried out via a standard stratification argument, such as the one given in [37, Theorem 4.5]. To stratify over  $L$ , take  $\rho^{-1} = L^{1+\text{ddim}(\mathcal{X})}$  in (4.42) *ibid.*, we have that with probability at least  $1 - \delta$ ,

$$R(h, q) \leq R_n(h, q) + \frac{4}{\rho} \widehat{\mathcal{R}}_n(q \circ \mathcal{H}_1; Z) + \sqrt{\frac{\ln \log_2 \frac{2}{\rho}}{n}} + 3\sqrt{\frac{\ln \frac{2}{\delta}}{2n}}$$

holds uniformly over all  $h \in \bigcup_{L \geq 1} \mathcal{H}_L$ . As in the proof of Corollary V.9, the cumulative effect of  $\bar{\eta}$ -perturbation is an additive error term of  $4(2q-1)\bar{\eta}$ . To stratify over  $\bar{\eta}$ , notice that  $\bar{\eta}$  is chosen from an a-priori fixed set of size  $\lceil 1/\eta \rceil \leq 2/\eta$  — and so taking a union bound amounts to replacing  $\delta$  by  $\delta\eta/2$ .  $\square$

*Remark 2.* The runtime guarantees of Theorems III.1 and IV.1, as well as the risk bound of Theorem V.1, all depend exponentially on the doubling dimension of the metric space  $\mathcal{X}$ , hence even a modest dimensionality reduction yields dramatic savings in algorithmic and sample complexities. This was exploited in [17], which develops a technique that may roughly be described as a metric analogue of PCA. A set  $X = \{x_1, \dots, x_n\} \subset \mathcal{X}$  inherits the metric  $\rho$  of  $\mathcal{X}$  and hence  $\text{ddim}(X) \leq 2\text{ddim}(\mathcal{X})$  is well-defined [14, Lemma 6.6]. Let us say that  $\tilde{X} = \{\tilde{x}_1, \dots, \tilde{x}_n\} \subset \mathcal{X}$  is an  $(\alpha, \beta)$ -perturbation of  $X$  if  $\frac{1}{n} \sum_{i=1}^n \rho(x_i, \tilde{x}_i) \leq \alpha$  and  $\text{ddim}(\tilde{X}) \leq \beta$ . Intuitively, the data is “essentially” low-dimensional if it admits an  $(\alpha, \beta)$ -perturbation with small  $\alpha, \beta$ , which leads to improved Rademacher estimates.

The data-dependent nature of  $\hat{\mathcal{R}}_n$  was used in [17] to develop generalization bounds that can exploit data that is essentially low-dimensional in the above sense. That paper dealt with the binary classification setting, and the technique was applied to the multiclass case by [27]. The same dimensionality reduction technique applies just as directly in our context of regression (the proof is deferred to Appendix B).

*Theorem V.10.* Let  $Z = (X, Y) \in \mathcal{X}^n \times [0, 1]^n$  be the training sample and suppose that  $X$  admits an  $(\alpha, \beta)$ -perturbation  $\tilde{X}$ . Then, for  $L \geq 1$ ,

$$\hat{\mathcal{R}}_n(q \circ (\mathcal{H}_L \oplus \llbracket \eta \rrbracket); Z) \leq 2(2q-1)\eta + q^{3/2}L\alpha + \sqrt{\frac{2 \ln \frac{1}{2q}}{n}}$$

A key feature of the bound above is that it does not explicitly depend on  $\text{ddim}(\mathcal{X})$  (the dimension of the ambient space) or even on  $\text{ddim}(X)$  (the dimension of the data).

### Acknowledgements

We thank Larry Wasserman for helpful comments on the manuscript, and Robert Schapire for useful discussion and feedback.

### REFERENCES

[1] P. Assouad. Plongements lipschitziens dans  $\mathbf{R}^n$ . *Bull. Soc. Math. France*, 111(4):429–448, 1983.  
[2] Amitabha Bagchi, Adam L. Buchsbaum, and Michael T. Goodrich. Biased skip lists. *Algorithmica*, 42:31–48, 2005.

[3] Peter Bartlett. Covering numbers, chaining & dudley’s entropy integral (lecture 24). Class Notes on Statistical Learning Theory, <https://people.eecs.berkeley.edu/~bartlett/courses/281b-sp06/lecture24.ps>, 2006.  
[4] Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.  
[5] Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *ICML ’06: Proceedings of the 23rd international conference on Machine learning*, pages 97–104, New York, NY, USA, 2006. ACM.  
[6] Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: a survey of some recent advances. *ESAIM: Probability and Statistics*, 9:323–375, 2005.  
[7] Leo Breiman, Jerome H. Friedman, Richard A. Olshen, and Charles J. Stone. *Classification and regression trees*. Wadsworth Statistics/Probability Series. Wadsworth Advanced Books and Software, Belmont, CA, 1984.  
[8] Hubert T-H. Chan, Anupam Gupta, Bruce M. Maggs, and Shuheng Zhou. On hierarchical routing in doubling metrics. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’05, pages 762–771. SIAM, 2005.  
[9] Kenneth L. Clarkson. Nearest neighbor queries in metric spaces. *Discrete Comput. Geom.*, 22(1):63–93, 1999.  
[10] Kenneth L. Clarkson. Nearest-neighbor searching and metric space dimensions. In G. Shakhnarovich, T. Darrell, and P. Indyk, editors, *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*, pages 15–59. MIT Press, 2006.  
[11] Richard Cole and Lee-Ad Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *STOC*, pages 574–583, 2006.  
[12] Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 1996.  
[13] Michael Elkin and Shay Solomon. Optimal Euclidean spanners: really short, thin and lanky. In *STOC*, pages 645–654, 2013.  
[14] Andreas Emil Feldmann, Wai Shing Fung, Jochen Könemann, and Ian Post. A  $(1 + \epsilon)$ -embedding of low highway dimension graphs into bounded treewidth graphs. *CoRR*, abs/1502.04588, 2015.  
[15] Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Efficient regression in metric spaces via approximate Lipschitz extension. In *SIMBAD*, pages 43–58, 2013.  
[16] Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Efficient classification for metric data. *IEEE Transactions on Information Theory*, 60(9):5750–5759, 2014.  
[17] Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Adaptive metric dimensionality reduction. *Theoretical Computer Science*, pages 105–118, 2016. Preliminary version in ALT 2013.  
[18] Lee-Ad Gottlieb and Liam Roditty. An optimal dynamic spanner for doubling metric spaces. In *ESA ’08: Proceedings of the 16th annual European symposium on Algorithms*, pages 478–489, 2008.  
[19] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS*, pages 534–543, 2003.  
[20] László Györfi, Michael Kohler, Adam Krzyżak, and Harro Walk. *A distribution-free theory of nonparametric regression*. Springer Series in Statistics. Springer-Verlag, New York, 2002.  
[21] Sarel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006.  
[22] Wolfgang Härdle. *Applied nonparametric regression*, volume 19 of *Econometric Society Monographs*. Cambridge University Press, Cambridge, 1990.  
[23] Jon Kleinberg, Aleksandrs Slivkins, and Tom Wexler. Triangulation and embedding using small sets of beacons. *J. ACM*, 56:32:1–32:37, September 2009.

- [24] Andreĭ N. Kolmogorov and Vladimir M. Tihomirov.  $\varepsilon$ -entropy and  $\varepsilon$ -capacity of sets in functional space. *Amer. Math. Soc. Transl. (2)*, 17:277–364, 1961.
- [25] Vladimir Koltchinskii and Dmitriy Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *Ann. Statist.*, 30(1):1–50, 2002.
- [26] Aryeh Kontorovich. Concentration in unbounded metric spaces and algorithmic stability. In *ICML (2)*, 2014.
- [27] Aryeh Kontorovich and Roi Weiss. Maximum margin multiclass nearest neighbors. In *ICML*, 2014.
- [28] Aryeh Kontorovich and Roi Weiss. A bayes consistent 1-NN classifier. In *18th International Conference on Artificial Intelligence and Statistics, AISTATS 2015*, 2015.
- [29] Samory Kpotufe. Fast, smooth and adaptive regression in metric spaces. In *Advances in Neural Information Processing Systems 22*, pages 1024–1032, 2009.
- [30] Samory Kpotufe and Sanjoy Dasgupta. A tree-based regressor that adapts to intrinsic dimension. *J. Comput. Syst. Sci.*, 78(5):1496–1515, September 2012.
- [31] Robert Krauthgamer and James R. Lee. Navigating nets: Simple algorithms for proximity search. In *15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 791–801, January 2004.
- [32] Rasmus Kyng, Anup Rao, Sushant Sachdeva, and Daniel A. Spielman. Algorithms for Lipschitz learning on graphs. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015*, pages 1190–1223, 2015.
- [33] John Lafferty and Larry Wasserman. Rodeo: Sparse, greedy nonparametric regression. *Ann. Stat.*, 36(1):28–63, 2008.
- [34] Gábor Lugosi and Kenneth Zeger. Nonparametric estimation via empirical risk minimization. *IEEE Transactions on Information Theory*, 41(3):677–687, 1995.
- [35] Edward James McShane. Extension of range of functions. *Bull. Amer. Math. Soc.*, 40(12):837–842, 1934.
- [36] Ron Meir and Tong Zhang. Generalization error bounds for bayesian mixture algorithms. *J. Mach. Learn. Res.*, 4:839–860, December 2003.
- [37] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations Of Machine Learning*. The MIT Press, 2012.
- [38] É. A. Nadaraya. *Nonparametric estimation of probability densities and regression curves*, volume 20 of *Mathematics and its Applications (Soviet Series)*. Kluwer Academic Publishers Group, Dordrecht, 1989. Translated from the Russian by Samuel Kotz.
- [39] Tyler Neylon. *Sparse solutions for linear prediction problems*. PhD thesis, New York University, 2006.
- [40] David Pollard. *Convergence of Stochastic Processes*. Springer-Verlag, 1984.
- [41] B. L. S. Prakasa Rao. *Nonparametric functional estimation*. Probability and Mathematical Statistics. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], New York, 1983.
- [42] Alexander Rakhlin, Karthik Sridharan, and Alexandre B. Tsybakov. Empirical entropy, minimax regret and minimax risk. *CoRR*, abs/1308.1147, 2013.
- [43] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- [44] John Shawe-Taylor, Peter L. Bartlett, Robert C. Williamson, and Martin Anthony. Structural risk minimization over data-dependent hierarchies. *IEEE Transactions on Information Theory*, 44(5):1926–1940, 1998.
- [45] Shay Solomon. From hierarchical partitions to hierarchical covers: optimal fault-tolerant spanners for doubling metrics. In *STOC*, pages 363–372, 2014.
- [46] Kunal Talwar. Bypassing the embedding: Algorithms for low dimensional metrics. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing, STOC '04*, pages 281–290, New York, NY, USA, 2004. ACM.
- [47] Alexandre B. Tsybakov. *Introduction à l’estimation non-paramétrique*, volume 41 of *Mathématiques & Applications (Berlin) [Mathematics & Applications]*. Springer-Verlag, Berlin, 2004.
- [48] Vladimir N. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, New York, 1995.
- [49] Nisheeth K. Vishnoi.  $L_x = b$ . *Foundations and Trends in Theoretical Computer Science*, 8(12):1–141, 2013.
- [50] Ulrike von Luxburg and Olivier Bousquet. Distance-based classification with Lipschitz functions. *Journal of Machine Learning Research*, 5:669–695, 2004.
- [51] Larry Wasserman. *All of nonparametric statistics*. Springer Texts in Statistics. Springer, New York, 2006.
- [52] Hassler Whitney. Analytic extensions of differentiable functions defined in closed sets. *Transactions of the American Mathematical Society*, 36(1):63–89, 1934.
- [53] Neal E. Young. Sequential and parallel algorithms for mixed packing and covering. In *42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 538–546, 2001.
- [54] Xiaojin Zhu, Zoubin Ghahramani, and John D. Lafferty. Semi-supervised learning using gaussian fields and harmonic functions. In *Machine Learning, Proceedings of the 20th International Conference, ICML 2003*, pages 912–919, 2003.

## APPENDIX

### A. A small-hop spanner

In this section, we prove the following theorem. See Section II for the definition of a spanner.

**Theorem A.1.** *Every finite metric space  $\mathcal{X}$  on  $n$  points admits a  $(1 + \delta)$ -stretch spanner with degree  $\delta^{-O(\text{ddim}(\mathcal{X}))}$  (for  $0 < \delta \leq \frac{1}{2}$ ) and hop-diameter  $O(\ln n)$ , that can be constructed in time  $\delta^{-O(\text{ddim}(\mathcal{X}))} n \ln n$ .*

Gottlieb and Roditty [18] presented for general metrics a  $(1 + \delta)$ -stretch spanner with degree  $\delta^{-O(\text{ddim}(\mathcal{X}))}$  and construction time  $\delta^{-O(\text{ddim}(\mathcal{X}))} n \ln n$ , but this spanner has potentially large hop-diameter. Our goal is to modify this spanner to have low hop-diameter, without significantly increasing the spanner degree. Now, as described in [18], the points of  $\mathcal{X}$  are arranged in a tree of degree  $\delta^{-O(\text{ddim}(\mathcal{X}))}$ , and a spanner path is composed of three consecutive parts: (a) a path ascending the edges of the tree; (b) a single edge; and (c) a path descending the edges of the tree. We will show how to decrease the number of hops in parts (a) and (c). Below we will prove the following lemma.

**Lemma A.2.** *Let  $T$  be a tree containing directed child-parent edges ( $n = |T|$ ), and let  $p$  be the degree of  $T$ . Then  $T$  may be augmented with directed descendant-ancestor edges to create a DAG  $G$  with the following properties: (i)  $G$  has degree  $p + 3$ ; and (ii) The hop-distance in  $G$  from any node to each of its ancestors is  $O(\ln n)$ .*

Note that Theorem A.1 is an immediate consequence of Lemma A.2 applied to the spanner of [18]. It remains

only to prove Lemma A.2, for which we will need the following preliminary lemma.

**Lemma A.3.** *Consider an ordered path on nodes  $x_1, \dots, x_n$ . Let these nodes be assigned positive weights  $w_i = w(x_i)$ , and let the weight of the path be  $W = \sum_{i=1}^n w(x_i)$ . there exists a DAG  $G$  on these nodes with the following properties:*

- 1) *Edges in  $G$  always point to the antecedent node in the ordering.*
- 2) *The hop-distance from any node  $x_i$  to the root node  $x_1$  is not more than  $O(\ln \frac{W}{w_i})$ .*
- 3) *The hop-distance from any node  $x_i$  to an antecedent  $x_j$  is not more than  $O(\ln \frac{W}{w_i} + \ln \frac{W}{w_j})$ .*
- 4)  *$G$  has degree 3.*

**Proof:** [Proof of Lemma A.3] The construction is essentially the same as in the biased skip-lists of Bagchi et al. [2]. Let  $x_1$  and  $x_n$  be the left and right end nodes of the path, and let the other nodes be the middle nodes. Partition the middle nodes into two child subpaths  $\{x_2, \dots, x_i\}$  (the left child path) and  $\{x_{i+1}, \dots, x_{n-1}\}$  (the right child path), where  $x_i$  is chosen so that the weight of the middle nodes of each child path is not more than half the weight of the middle nodes of the parent path. (If the parent path has three middle nodes or fewer, then there will be a single child path.) The child paths are then recursively partitioned, until the recursion reaches paths with no middle nodes.

The edges are assigned as follows. A right end node of a path has two edges leaving it. One points to the left end node of the path (unless the path has only one node). The other edge points to the right end node of the right (or single) child path. A left end node of a path has one edge leaving it: If this path is a right child path, the edge points to the left sibling path's right end node. If this path is a left or single child path, then the edge points to the parent's left end node. The lemma follows via standard analysis.  $\square$

We are now ready to prove Lemma A.2, which would conclude the proof of Theorem A.1.

**Proof:** [Proof of Lemma A.2] Given tree  $T$ , decompose  $T$  into *heavy paths*: A heavy path is one that begins at the root and continues with the heaviest child, the child with the most descendants. In a heavy path decomposition, all off-path subtrees are recursively decomposed. For each heavy path, let the weight of each node in the path be the number of descendant nodes in its off-path subtrees. For each heavy path, we build the weighted construction of Lemma A.3.

Now, a path from node  $u \in T$  to  $v \in T$  traverses a set of at most  $\lceil \ln n \rceil$  heavy paths, say paths  $P_1, \dots, P_j$ . The

number of hops from  $u$  to  $v$  is bounded by  $O(\ln \frac{w(P_1)}{w(u)} + (\sum_{i=1}^j \ln \frac{w(P_{i-1})}{w(P_i)}) + \ln \frac{n}{w(v)}) = O(\ln n)$ , and the degree of  $G$  is at most  $p + 3$ .  $\square$

## B. Rademacher-complexity proofs

**Proof:** [Proof of Lemma V.5] Suppose that  $h : \mathcal{X} \rightarrow [0, 1]$  with  $\|h\|_{\text{Lip}} = L$ ,  $f(x, y) = |h(x) - y|^q$ , and  $(\mathcal{Z}, d_q)$  is the metric space defined in (10). To prove (i), we consider the cases  $q = 1, 2$  separately. For  $q = 1$ ,

$$\begin{aligned} |f(x, y) - f(x', y')| &= ||h(x) - y| - |h(x') - y'|| \\ &\leq |(h(x) - y) - (h(x') - y')| \\ &\leq |h(x) - h(x')| + |y - y'| \\ &\leq L\rho(x, x') + |y - y'| \\ &\leq \max\{1, L\}(\rho(x, x') + |y - y'|) \\ &= \max\{1, L\}d_1((x, y), (x', y')), \end{aligned}$$

which proves the claim for this case. Now consider the case  $q = 2$  and recall the following basic fact: if  $\varphi$  maps  $E \subset \mathbb{R}^k$  to  $\mathbb{R}$ , then

$$\sup_{x \neq x' \in E} \frac{|\varphi(x) - \varphi(x')|}{\|x - x'\|_2} \leq \sup_{z \in E} \|\nabla \varphi(z)\|_2.$$

Let us take  $\varphi[0, 1]^2 \rightarrow \mathbb{R}$  to be  $\varphi(h, y) = (h - y)^2$ , which satisfies

$$\max_{(h, y) \in [0, 1]^2} \|\nabla \varphi(h, y)\|_2 = 2^{3/2}.$$

It follows that

$$\begin{aligned} |f(x, y) - f(x', y')| &= |(h(x) - y)^2 - (h(x') - y')^2| \\ &\leq 2^{3/2} ((h(x) - h(x'))^2 + (y - y')^2)^{1/2} \\ &\leq 2^{3/2} ((L\rho(x, x'))^2 + (y - y')^2)^{1/2} \\ &\leq 2^{3/2} \max\{1, L\}d_2((x, y), (x', y')), \end{aligned}$$

which completes the proof of (i).

To prove (ii), we will show that

$$\lambda(\mathcal{Z}, d_q) \leq 4\lambda(\mathcal{X}, \rho)^2, \quad (15)$$

where  $\lambda(\cdot)$  is the doubling constant of a given metric space. Consider the case  $q = 1$ , put  $a = \lambda(\mathcal{X}, \rho)$ , and fix any  $d_1$ -ball  $B \subset \mathcal{Z}$  with diameter  $r$ . Define the coordinate projections  $\pi_1 : \mathcal{Z} \rightarrow \mathcal{X}$  and  $\pi_2 : \mathcal{Z} \rightarrow [0, 1]$  in the obvious way and assume without loss of generality that  $\pi_2(B) \subset [b, b + r]$ . Now partition  $B$  into 4 subsets based on the second coordinate:

$$B_i = \left\{ z \in B : \pi_2(z) \in \left[ b + \frac{i}{4}, b + \frac{i+1}{4} \right) \right\}$$

for  $i = 0, 1, 2, 3$ .

By definition of the doubling constant, each  $\pi_1(B_i) \subset \mathcal{X}$  can be covered by  $a^2$  balls  $V \subset \mathcal{X}$  of diameter at most  $r/4$  under the metric  $\rho$ . It follows by construction that each  $B_i$  can be covered by  $a^2$  sets of the form

$$V \times [b + i/4, b + (i + 1)/4),$$

each of  $d_1$ -diameter at most  $r/2$ . Hence, any ball in  $\mathcal{Z}$  can be covered by  $4a^2$  balls of half the diameter, and so the claim is proved for  $q = 1$ . To handle the case  $q = 2$ , observe that

$$d_2((x, y), (x', y')) \leq d_1((x, y), (x', y'))$$

for all  $(x, y), (x', y') \in \mathcal{Z}$ . This proves (ii).  $\square$

**Proof:** [Proof of Lemma V.8] Let  $\tilde{h}(x) = h(x) + \delta(x)$ , with  $\|\delta\|_\infty \leq \eta$  be an  $\eta$ -perturbed version of  $h$ , with the corresponding  $\tilde{f}(x, y) = |h(x) - y|^q$ . Consider the case  $q = 1$ . Then

$$\begin{aligned} |f(x, y) - \tilde{f}(x, y)| &= \left| |h(x) - y| - |\tilde{h}(x) - y| \right| \\ &\leq \left| (h(x) - y) - (\tilde{h}(x) - y) \right| \\ &= |h(x) - \tilde{h}(x)| = |\delta(x)| \leq \eta, \end{aligned}$$

which proves this case. For  $q = 2$ , we have

$$\begin{aligned} |f(x, y) - \tilde{f}(x, y)| &= \left| (h(x) - y)^2 - (\tilde{h}(x) - y)^2 \right| \\ &= \left| [h(x) + \delta(x) - y]^2 - [h(x) - y]^2 \right| \\ &= \delta(x) |2h(x) + \delta(x) - 2y| \leq 3\eta, \end{aligned}$$

since  $0 \leq h, y, \delta \leq 1$ .  $\square$

**Proof:** [Proof of Theorem V.10] Put  $\tilde{Z} = (\tilde{X}, Y)$ . For  $X_i \in X$ ,  $\tilde{X}_i \in \tilde{X}$ , and  $f \in q \circ \mathcal{H}_L$ , define  $\delta_i(f) = f(X_i, Y_i) - f(\tilde{X}_i, Y_i)$ . As in the proof of Corollary V.9, we have

$$\widehat{\mathcal{R}}_n(q \circ (\mathcal{H}_L \oplus [\eta])); Z \leq \widehat{\mathcal{R}}_n(q \circ \mathcal{H}_L; Z) + (2q - 1)\eta.$$

Further,

$$\begin{aligned} \widehat{\mathcal{R}}_n(q \circ \mathcal{H}_L; Z) &= \mathbb{E} \left[ \sup_{f \in q \circ \mathcal{H}_L} \frac{1}{n} \sum_{i=1}^n \sigma_i f(X_i, Y_i) \right] \\ &= \mathbb{E} \left[ \sup_{f \in q \circ \mathcal{H}_L} \frac{1}{n} \sum_{i=1}^n \sigma_i \left( f(\tilde{X}_i, Y_i) - \delta_i(f) \right) \right] \\ &\leq \widehat{\mathcal{R}}_n(q \circ \mathcal{H}_L; \tilde{Z}) + \mathbb{E} \left[ \sup_{f \in q \circ \mathcal{H}_L} \frac{1}{n} \sum_{i=1}^n \sigma_i \delta_i(f) \right]. \end{aligned}$$

The first term is estimated by the same calculation as in the proof of Theorem V.6:

$$\begin{aligned} \widehat{\mathcal{R}}_n(q \circ (\mathcal{H}_L); \tilde{Z}) &\leq (2q - 1)\eta + \sqrt{\frac{2 \ln \mathcal{N}((2q - 1)\eta, q \circ \mathcal{H}_L, \|\cdot\|_\infty)}{n}} \\ &\leq (2q - 1)\eta + \sqrt{\frac{2}{n}} \left( \frac{16q^{3/2}L}{(2q - 1)\eta} \right)^{1+\beta} \left( \ln \frac{8}{(2q - 1)\eta} \right)^{1/2}. \end{aligned}$$

To bound the second term, invoke Lemma V.5(i) to conclude that

$$|\delta_i| = \left| f(X_i, Y_i) - f(\tilde{X}_i, Y_i) \right| \leq q^{3/2}L\rho(X_i, \tilde{X}_i).$$

Hence,

$$\begin{aligned} \mathbb{E} \left[ \sup_{f \in q \circ \mathcal{H}_L} \frac{1}{n} \sum_{i=1}^n \sigma_i \delta_i(f) \right] &\leq \sup_{f \in q \circ \mathcal{H}_L} \frac{1}{n} \sum_{i=1}^n \left| f'(X_i, Y_i) - f'(\tilde{X}_i, Y_i) \right| \\ &\leq n^{-1}q^{3/2}L \sum_{i=1}^n \rho(X_i, \tilde{X}_i) \leq q^{3/2}L\alpha. \end{aligned}$$

$\square$