

Functions with average smoothness: structure, algorithms, and learning

Yair Ashlagi

*Computer Science Department
Ben-Gurion University
Beer Sheva, Israel*

ASHLAGI@POST.BGU.AC.IL

Lee-Ad Gottlieb

*Computer Science Department
Ariel University
Shomron, Israel*

LEEAD@ARIEL.AC.IL

Aryeh Kontorovich

*Computer Science Department
Ben-Gurion University
Beer Sheva, Israel*

KARYEH@CS.BGU.AC.IL

Editor: Po-Ling Loh

Abstract

We initiate a program of average smoothness analysis for efficiently learning real-valued functions on metric spaces. Rather than using the Lipschitz constant as the regularizer, we define a local slope at each point and gauge the function complexity as the average of these values. Since the mean can be dramatically smaller than the maximum, this complexity measure can yield considerably sharper generalization bounds — assuming that these admit a refinement where the Lipschitz constant is replaced by our average of local slopes.

Our first major contribution is to obtain just such distribution-sensitive bounds. This required overcoming a number of technical challenges, perhaps the most formidable of which was bounding the *empirical* covering numbers, which can be much worse-behaved than the ambient ones. Our combinatorial results are accompanied by efficient algorithms for smoothing the labels of the random sample, as well as guarantees that the extension from the sample to the whole space will continue to be, with high probability, smooth on average. Along the way we discover a surprisingly rich combinatorial and analytic structure in the function class we define.

Keywords: metric space, doubling dimension, Lipschitz, smoothness, average case

1. Introduction

Smoothness is a natural measure of complexity commonly used in learning theory and statistics. Perhaps the simplest method of quantifying the smoothness of a function is via the Lipschitz seminorm. This quantifier of regularity has the advantage of being an analytically and algorithmically convenient, broadly applicable complexity measure, requiring only a metric space (as opposed to additional differentiable structure). In particular, the Lipschitz constant yields immediate bounds on the fat-shattering dimension (Gottlieb et al., 2014), covering numbers (Kolmogorov and Tihomirov, 1961), and sample compression (Gottlieb

et al., 2018) of a function class, which in turn directly imply generalization bounds for classification and regression, and also bounds the run-time of associated learning algorithms.

The simplicity of the Lipschitz seminorm, however, has a downside: it is a worst-case measure, insensitive to the underlying distribution. As such, it can be overly pessimistic in that a single point pair can drive the Lipschitz constant of a function arbitrarily high, even if the function is nearly constant everywhere else. Intuitively, we expect the complexity of learning a function that is highly smooth, apart from low-density regions of high fluctuation, to be determined by its average — rather than worst-case — behavior. To this end, we seek a complexity measure that is resilient to local fluctuations in low-density regions. Formalizing this intuition and exploring its analytic and algorithmic ramifications is the main contribution of this paper.

Very roughly speaking, to learn an L -Lipschitz (in the Euclidean metric) function $f : [0, 1]^d \rightarrow [0, 1]$ at fixed precision and confidence requires on the order of L^d examples (Wainwright, 2019), and this continues to hold in more general metric spaces (Kpotufe, 2011; Kpotufe and Dasgupta, 2012; Kpotufe and Garg, 2013; Gottlieb et al., 2017, 2014; Chaudhuri and Dasgupta, 2014). The goal of this paper is to replace the worst-case Lipschitz constant L by an average one \bar{L} , while still obtaining bounds of the general form \bar{L}^d . Further, we seek fully empirical generalization bounds, without making any a priori assumptions on either the target function or the distribution.

1.1 Our contributions

A detailed roadmap of our results is given in Section 2; here we only provide a brief overview. We initiate a program of average-case smoothness analysis for efficiently learning binary and real-valued functions on metric spaces. To any function $f : \Omega \rightarrow \mathbb{R}$ acting on a metric probability space (Ω, ρ, μ) , we associate a complexity measure $\bar{\Lambda}_f = \bar{\Lambda}_f(\Omega, \rho, \mu) \in [0, \infty]$, which corresponds to an average slope. Our measure always satisfies $\bar{\Lambda}_f \leq \|f\|_{\text{Lip}}$ and, as we illustrate below, the gap can be considerable. Having defined our notion of average smoothness, we show that the worst-case Lipschitz constant L can essentially be replaced by its averaged variant $\bar{\Lambda}_f$ in the covering number bounds.

Our results are fully empirical in that we make no a priori assumptions on the target function or the sampling distribution, and only require a finite diameter and doubling dimension of the metric space. A curious and unique feature of our setting — which also presents the bulk of the technical challenges — is the fact that although our hypothesis class is fixed before observing the data, it is defined in terms of the unknown sampling distribution, and hence not explicitly known to the learner. This is in stark contrast with all previous supervised learning settings, where the function classes are fully known a priori. Having observed a sufficiently large sample allows the learner to construct an explicit hypothesis and conclude that, with high probability, it belongs to the average smoothness class (to which our generalization bounds then apply).

The statistical generalization bounds are accompanied by efficient algorithms for performing sample smoothing and a Lipschitz-type extension for label prediction on test points. The function classes we define turn out to exhibit a surprisingly rich structure, making them an object worthy of future study. See Section 2 for a comprehensive overview of our techniques and central results, along with comparisons to the current state-of-art bounds.

1.2 Related work

For the line segment metric $(\Omega, \rho) = ([a, b], |\cdot|)$, the *bounded variation* (BV) of any $f : [a, b] \rightarrow \mathbb{R}$ with integrable derivative is given by $V_a^b(f) = \int_a^b |f'(x)| dx$; this is perhaps the most basic notion of average smoothness. BV does not require differentiability; see Appell et al. (2014) for an encyclopedic reference. Generalization bounds for BV functions may be obtained via covering numbers (Bartlett et al., 1997; Long, 2004) (the latter also gave an efficient algorithm for learning BV functions via linear programming) or the fat-shattering dimension (Anthony and Bartlett, 1999, Theorem 11.12). The aforementioned results correspond to the case of a uniformly distributed μ on $[a, b]$, and thus are not distribution-sensitive. A natural extension of BV to general measures would be to define $V_a^b(f) = \int_a^b |f'(x)| d\mu(x)$, but then the known fat-shattering and covering number estimates break down — especially if μ is not known to the learner.

Generalizing the notion of BV to higher dimensions is not nearly as straightforward. A common approach is via the Hardy-Krause variation (Appell et al., 2014; Kuipers and Niederreiter, 1974; Niederreiter and Talay, 2006). Even the two-dimensional case evades a simple characterization; counter-intuitively, Lipschitz functions $f : [0, 1]^2 \rightarrow \mathbb{R}$ may fail to have finite variation in the Hardy-Krause sense (Basu and Owen, 2016, Lemma 1). Some (rather loose) L_1 covering numbers for BV functions on $[0, 1]^n$ were obtained by Dutta and Nguyen (2018, Theorem 3.1); these are not distribution-sensitive. Generalizations of BV to metric measure spaces beyond the Euclidean are known (Ambrosio and Ghezzi, 2016); we are not aware of any covering number or combinatorial dimension estimates for these.

If one considers *bracketing* (rather than covering) numbers, there are known results for controlling these in terms of various measures of average smoothness. Nickl and Pötscher (2007) bound the bracketing numbers of Besov- and Sobolev-type classes. Malykhin (2010) also gave bracketing number bounds, using a different notion of smoothness: the *averaged modulus of continuity* developed by Sendov and Popov (1988). We note that covering numbers asymptotically always give tighter estimates than bracketing ones (Hanneke et al., 2023). More significantly, to our knowledge, all of the previous results bound the *ambient* rather than the *empirical* covering numbers (see Section 1.3 for definitions, and Section 3 for bounds), when it is precisely the latter that are needed for Uniform Glivenko-Cantelli laws (Giné and Zinn, 1984).

A seminal work on recovering functions with spacially inhomogeneous smoothness from noisy samples is Donoho and Johnstone (1998). Another tangentially related notion is that of *Probabilistic Lipschitzness* (Kpotufe et al., 2015), which seeks to relax a hard Lipschitz condition on the labeling function in binary classification. The authors position it as a “data niceness” condition, analogous to that in Mammen and Tsybakov (1999). These significantly differ from our notion of average slope. Most importantly, PL and the various Tsybakov-type noise conditions are *assumptions* on the data-generating distribution rather than *empirically computable* quantities on a given sample. Our approach is fully empirical in the sense of not making a priori assumptions on the distribution or the target function. Additionally, PL is specifically designed for binary classification with deterministic labels — unlike our notion, which is applicable to any real-valued function and any conditional label distribution.

In this paper, we make systematic use of a Lipschitz-type extension (PMSE, defined in Section B) explicitly tailored to our framework. This extension is closely related to one introduced by Oberman (2008) (also relevant is Shvartsman (2017), p. 385 and p. 416, Remark 5.1). An extended abstract of this work was presented at COLT 2021 (Ashlagi et al., 2021). Very recently, some of the results have been extended to include average Hölder-type smoothness (Hanneke et al., 2023, 2024). The latter work achieved near-optimal rates L_1 risk rates for regression via bracketing numbers.

1.3 Definitions and preliminaries

Metric probability spaces. We assume a basic familiarity with metric measure spaces and refer the reader to a standard reference, such as Heinonen (2001). Standard set-theoretic notation is used throughout; in particular, for $f : \Omega \rightarrow \mathbb{R}$ and $A \subset \Omega$, we denote the restriction of f to A by $f|_A$. The triple (Ω, ρ, μ) is a *metric probability space* if μ is a probability measure supported on the Borel σ -algebra induced by the open sets of ρ . For Ω -valued random variables, the notation $X \sim \mu$ means that $\mathbb{P}(X \in A) = \mu(A)$ for all Borel sets A .

Covers, packings, nets, hierarchies, partitions. The *diameter* of $A \subseteq \Omega$ is the maximal interpoint distance: $\text{diam}(A) = \sup_{x, x' \in A} \rho(x, x')$. For $t > 0$ and $A, B \subseteq \Omega$, we say that A is a *t-cover* of B if

$$\sup_{b \in B} \inf_{a \in A} \rho(a, b) \leq t,$$

and define the *t-covering number* of B to be the minimum cardinality of any *t-cover*, denoted by $\mathcal{N}(t, B, \rho)$. We say that $A \subseteq B \subseteq \Omega$ is a *t-packing* of B if $\rho(a, a') > t$ for all distinct $a, a' \in A$. Finally, A is a *t-net* of B if it is simultaneously a *t-cover* and a *t-packing*. A family of sets $H_{t^0} \subseteq H_{t^{-1}} \subseteq \dots \subseteq H_{t^{-m}}$ is a *hierarchy* for the set $H_{t^{-m}}$ if each $H_{t^{-i}}$ ($i < m$) is a t^{-i} -net of $H_{t^{-(i+1)}}$, where we have assumed that $H_{t^{-m}}$ has diameter 1 and so H_1 contains a single point.

We denote by $B(x, r) = \{x' \in \Omega : \rho(x, x') \leq r\}$ the (closed) r -ball about x . If there is a $D < \infty$ such that every r -ball in Ω is contained in the union of some D $r/2$ -balls, the metric space (Ω, ρ) is said to be *doubling*. Its *doubling dimension* is defined as $\text{ddim}(\Omega) = \text{ddim}(\Omega, \rho) =: \log_2 D^*$, where D^* is the smallest D verifying the doubling property. It is well-known (Krauthgamer and Lee, 2004; Gottlieb et al., 2016) that

$$\mathcal{N}(t, \Omega, \rho) \leq \left(\frac{2 \text{diam}(\Omega)}{t} \right)^{\text{ddim}(\Omega)}, \quad t > 0, \tag{1}$$

which will be referred to as the covering property of doubling spaces. The packing property of doubling spaces asserts an analogous packing number bound, up to constants in the exponent. A hierarchy for any n -point Ω set can be constructed in time $2^{O(\text{ddim}(\Omega))} n \min\{\log n, \log \Delta\}$, where Δ is the *aspect ratio* (ratio between the maximum and minimum interpoint distances) of Ω (Krauthgamer and Lee, 2004; Har-Peled and Mendel, 2006; Cole and Gottlieb, 2006).

To any finite $V \subseteq \Omega$ we associate the map $\varphi_V : \Omega \rightarrow V$ taking each $x \in \Omega$ to its nearest neighbor in V , with ties broken arbitrarily (say, via some fixed ordering on Ω)¹. The

1. A measurable total order always exists (Hanneke et al., 2021).

collection of sets $\{\varphi_V^{-1}(v) : v \in V\}$ is said to comprise the *Voronoi* partition of Ω induced by V . If V happens to be a t -net of Ω , then

$$\rho(x, \text{phi}_V(x)) \leq t, \quad x \in \Omega. \quad (2)$$

Indices, norms. We write $[n] := \{1, \dots, n\}$ and use the shorthand $z_{[n]} := (z_1, \dots, z_n)$ for sequences. For any metric probability space (Ω, ρ, μ) , $p \geq 1$, and any $f : \Omega \rightarrow \mathbb{R}$, we define the norm $\|f\|_{L_p(\Omega, \rho, \mu)}^p = \mathbb{E}_{X \sim \mu}[|f(X)|^p] = \int_{\Omega} |f(x)|^p d\mu(x)$.

This work assumes a fixed metric probability space (Ω, ρ, μ) ; this will be termed the **ambient space**. Several derived metric probability spaces (Ω', ρ', μ') will be considered, which will all be **induced subspaces** of (Ω, ρ) in the sense that $\Omega' \subseteq \Omega$ and $\rho' = \rho|_{\Omega' \times \Omega'}$. To lighten the notation, we will often suppress the common metric ρ and use the shorthand $\|\cdot\|_{L_p(\mu')} := \|\cdot\|_{L_p(\Omega', \rho, \mu')}$. For any $f, g : \Omega \rightarrow \mathbb{R}$ and any induced subspace (Ω', ρ) of Ω with measure μ' , we use the shorthand

$$\|f - g\|_{L_p(\mu')} := \|f|_{\Omega'} - g|_{\Omega'}\|_{L_p(\mu')} = \|f|_{\Omega'} - g|_{\Omega'}\|_{L_p(\Omega', \rho, \mu')}. \quad (3)$$

In particular, sampling $\Omega_n := X_{[n]} = (X_1, \dots, X_n) \sim \mu^n$ induces the **empirical space** (Ω_n, ρ, μ_n) with the norm $\|\cdot\|_{L_p(\mu_n)}$, where μ_n is the empirical measure on Ω_n , formally given by

$$\mu_n(x) = n^{-1} \sum_{i=1}^n \mathbf{1}[X_i = x]. \quad (4)$$

We say that (Ω_n, ρ, μ_n) is the **empirical realization** of (Ω, ρ, μ) . The ℓ_{∞} norm $\|f\|_{\infty} = \sup_{x \in \Omega} |f(x)|$ is measure-independent and dominates all of the measure-induced norms:

$$\|f - g\|_{L_p(\mu')} \leq \|f - g\|_{\infty}, \quad p \geq 1.$$

The Lipschitz seminorm $\|f\|_{\text{Lip}}$ is the smallest $L \in [0, \infty]$ satisfying $|f(x) - f(x')| \leq L\rho(x, x')$ for all $x, x' \in \Omega$.

Strong and weak mean. We define the **weak mean** of a non-negative random variable Z by

$$\mathbb{W}[Z] = \sup_{t>0} t \mathbb{P}(Z \geq t). \quad (5)$$

In contrast, the **strong mean** is just the usual expectation $\mathbb{E}[Z]$. By Markov's inequality, we always have $\mathbb{W}[Z] \leq \mathbb{E}[Z]$; further, the latter might be infinite while the former is finite; see Section 1.4. A partial reverse inequality for finite measure spaces is given in Lemma 25.

Local and average slope. For $f : \Omega \rightarrow \mathbb{R}$, we define the *slope*² of f at $x \in \Omega$ with respect to an $A \subseteq \Omega$ as

$$\Lambda_f(x, A) := \sup_{x' \in A \setminus \{x\}} \frac{|f(x) - f(x')|}{\rho(x, x')}. \quad (6)$$

2. To forestall potential confusion, note that for differentiable $f : \mathbb{R} \rightarrow \mathbb{R}$, it might be that $\Lambda_f(x, \mathbb{R})$ is different from $f'(x)$ and $|f'(x)|$. A more descriptive (though somewhat more cumbersome) term might be “worst-neighbor average slope”, “neighborhood slope”, or “worst-case slope from x ”.

Thus,

$$\|f\|_{\text{Lip}} = \sup_{x \in \Omega} \Lambda_f(x, \Omega). \quad (7)$$

We will define two notions of *average* slope: strong and weak, corresponding, respectively, to the strong and weak L_1 norms of the random variable $\Lambda_f(X)$, where $X \sim \mu$. The two averages are defined, respectively, as

$$\bar{\Lambda}_f(\mu, \Omega) := \mathbb{E}_{X \sim \mu} [\Lambda_f(X, \Omega)] = \|\Lambda_f(\cdot, \Omega)\|_{L_1(\mu)}, \quad (8)$$

$$\tilde{\Lambda}_f(\mu, \Omega) := \mathbb{W}_{X \sim \mu} [\Lambda_f(X, \Omega)] = \sup_{t > 0} t \mu(M_f(t)), \quad (9)$$

where $M_f(t)$ is the t -level set, a central object in this paper:

$$M_f(t) := \{x \in \Omega : \Lambda_f(x, \Omega) \geq t\}. \quad (10)$$

Intuitively speaking, the t -level set is the “rough” region of Ω associated with the function f , where the local slope is large. The strong-weak mean inequality above implies that

$$\tilde{\Lambda}_f(\mu, \Omega) \leq \bar{\Lambda}_f(\mu, \Omega) \leq \|f\|_{\text{Lip}} \quad (11)$$

always holds (the second inequality is obvious); further, $\bar{\Lambda}_f(\mu, \Omega)$ might be infinite while $\tilde{\Lambda}_f(\mu, \Omega)$ is finite (as demonstrated by the step function on $[0, 1]$ with the uniform measure, see Section 1.4). Since the above definitions were stated for *any* metric probability space, $\Lambda_f(x, \Omega_n)$, $\bar{\Lambda}_f(\mu_n, \Omega_n)$, and $\tilde{\Lambda}_f(\mu_n, \Omega_n)$ are well-defined as well. (To appreciate the subtle choice of our definitions, note that some intuitively appealing variants irreparably fail, as discussed in Section 1.4.)

The collection of all $[0, 1]$ -valued L -Lipschitz functions on Ω , as well as its strong and weak mean-slope counterparts are denoted, respectively, by

$$\text{Lip}_L(\Omega, \rho) = \left\{ f \in [0, 1]^\Omega; \|f\|_{\text{Lip}} \leq L \right\}, \quad (12)$$

$$\overline{\text{Lip}}_L(\Omega, \rho, \mu) = \left\{ f \in [0, 1]^\Omega; \bar{\Lambda}_f(\mu, \Omega) \leq L \right\}, \quad (13)$$

$$\widetilde{\text{Lip}}_L(\Omega, \rho, \mu) = \left\{ f \in [0, 1]^\Omega; \tilde{\Lambda}_f(\mu, \Omega) \leq L \right\}. \quad (14)$$

It follows from (11) that $\text{Lip}_L(\Omega, \rho, \mu) \subset \overline{\text{Lip}}_L(\Omega, \rho, \mu) \subset \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$, where all containments are, in general, strict, and

$$\mu(M_f(L/s)) \leq s, \quad s > 0 \quad (15)$$

holds for all $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$. For most of this paper, we shall be interested in the larger latter class, but occasional results for $\overline{\text{Lip}}_L(\Omega, \rho, \mu)$ will be presented, when of independent interest.

Remark: Observe that the classes $\overline{\text{Lip}}$ and $\widetilde{\text{Lip}}$ are defined in terms of the unknown sampling distribution μ . Given full knowledge of a function $f : \Omega \rightarrow \mathbb{R}$, a learner can verify that $f \in \text{Lip}_L$ but, absent full knowledge of μ , it is impossible to know for certain whether $f \in \widetilde{\text{Lip}}_L$ (or $f \in \overline{\text{Lip}}_L$). As increasingly larger samples are observed, the learner will be able to assert the latter inclusions with increasing confidence.

Empirical and true risk. For any probability measure ν on $\Omega \times [0, 1]$, we associate to any measurable $f : \Omega \rightarrow \mathbb{R}$ its *risk* $R(f; \nu) := \mathbb{E}_{(X,Y) \sim \nu} |f(X) - Y|$. In the special case of the empirical measure ν_n induced by a sample $(X_i, Y_i)_{i \in [n]} \sim \nu^n$, $R(f; \nu_n)$ is the empirical risk. For regression with real-valued f , this is the L_1 -risk; for classification with $\{0, 1\}$ -valued f , this is the 0-1 error. (See Mohri et al. (2012) for a standard reference.)

Miscellanea. Additional standard inequalities and notations are deferred to Section A in the Appendix.

1.4 Comparing average and worst-case smoothness

For $\gamma \in (0, 1/2)$, consider the metric space $\Omega = [0, 1/2 - \gamma] \cup [1/2 + \gamma, 1]$ equipped with the standard metric $\rho(x, x') = |x - x'|$, the uniform distribution μ , and $f : \Omega \rightarrow \mathbb{R}$ given by the step function $f(x) = \mathbf{1}[x > 1/2]$. Then $\|f\|_{\text{Lip}} = 1/(2\gamma)$ and

$$\bar{\Lambda}_f(\Omega, \rho, \mu) = \frac{1}{1 - 2\gamma} \int_{[0, 1/2 - \gamma] \cup [1/2 + \gamma, 1]} \frac{1}{|x - 1/2| + \gamma} dx = \frac{2}{1 - 2\gamma} \log \left(\frac{1 + 2\gamma}{4\gamma} \right). \quad (16)$$

For small γ , we have $\|f\|_{\text{Lip}} = \Theta(\gamma^{-1})$ and $\bar{\Lambda}_f = \Theta(\log \gamma^{-1})$, so even the cruder strong average smoothness measure provides an exponential savings over the worst-case one. This example has natural higher-dimensional analogues (i.e., $\Omega = [0, 1]^{d-1} \times ([0, 1/2 - \gamma] \cup [1/2 + \gamma, 1])$), where the phenomenon persists.

An analogous behavior is exhibited by the “margin loss” function $f(x) = \max\{1, \min\{0, 1 - x/\gamma\}\}$ defined on $([0, 1], |\cdot|, \mu)$, where μ is the uniform distribution on $[0, 1]$. In this case, $\|f\|_{\text{Lip}} = 1/\gamma$, while

$$\bar{\Lambda}_f([0, 1], |\cdot|, \mu) = \gamma \cdot \frac{1}{\gamma} + (1 - \gamma) \int_{\gamma}^1 \frac{1}{x} dx = 1 + (1 - \gamma) \log \frac{1}{\gamma} = \Theta \left(\log \frac{1}{\gamma} \right) \quad (17)$$

— again, an exponential savings.

For a more dramatic gap between the two measures, consider the family of functions $f_p(x) = x^p$ for $p \in (0, 1)$, on $\Omega = [0, 1]$ with the uniform distribution μ and the standard metric ρ . These all have $\|f_p\|_{\text{Lip}} = \infty$, while

$$\bar{\Lambda}_f([0, 1], |\cdot|, \mu) = \int_0^1 x^{p-1} dx = \frac{1}{p}. \quad (18)$$

Consider now the case of the step function $f : [0, 1] \rightarrow [0, 1]$ given by $f(x) = \mathbf{1}[x > 0]$. Taking Ω and μ as above, we have $\|f\|_{\text{Lip}} = \bar{\Lambda}_f(\mu, \Omega) = \infty$, so here the strong mean offers no advantage over the worst-case. However, since $\mu(M_f(t)) = 1/t$, we have that $\tilde{\Lambda}_f(\mu, \Omega) = 1$. More generally, Elperin and Kontorovich (2021+) have shown that $BV[0, 1]$, the class of all bounded-variation functions on $[0, 1]$, satisfies $BV[0, 1] \subset \widetilde{\text{Lip}}([0, 1], \rho, \mu)$, and the containment is strict.

Alternative notions of average smoothness. One might consider a natural alternative definition of average smoothness, less stringent than our $\bar{\Lambda}$ and our $\tilde{\Lambda}$:

$$\bar{\Lambda}_f^{\text{alt}}(\mu, \Omega) = \int_{\Omega} \int_{\Omega} \frac{|f(x) - f(y)|}{\rho(x, y)} d\mu(x) d\mu(y).$$

We argue that $\bar{\Lambda}_f^{\text{alt}}(\mu, \Omega)$ fails as an average measure of smoothness for bounding empirical covering numbers and obtaining generalization guarantees. Indeed, consider $\Omega = [0, 1]$ endowed with the standard metric $\rho(x, x') = |x - x'|$ and uniform distribution μ . Let \mathcal{F} be the collection of all $f : \Omega \rightarrow \{0, 1\}$ with finite support. It is well-known that \mathcal{F} is not UGC under μ , and has typical empirical covering numbers exponential in sample size. However, since $|f(x) - f(y)|$ vanishes μ^2 -almost-everywhere on $[0, 1]^2$, we have that $\bar{\Lambda}_f^{\text{alt}}(\mu, \Omega) = 0$ for all $f \in \mathcal{F}$. As a consistency check, note that no uniform bound over all $f \in \mathcal{F}$ is possible for either $\bar{\Lambda}_f$ or $\tilde{\Lambda}_f$. Furthermore, since the VC-dimension of \mathcal{F} is infinite, a classic VC lower bound shows that for every n , there exists an adversarial distribution that forces the expected excess risk of *any* learning algorithm to be at least $c_\infty > 0.169$ (Kontorovich and Pinelis, 2019).

2. Main results and roadmap

This section assumes a familiarity with the terminology and notation defined in Section 1.3.

Combinatorial structure. Our point of departure is Theorem 1, which bounds the $L_2(\mu)$ (*ambient*) covering numbers of the function class $\widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ — and, a fortiori, of $\overline{\text{Lip}}_L(\Omega, \rho, \mu)$ [defined in (13, 14)] — in terms of the average slope L , $\text{diam}(\Omega)$, and $\text{ddim}(\Omega)$. Crucially, there is no dependence on the Lipschitz constant $\|\cdot\|_{\text{Lip}}$. At scale t , Theorem 1 gives a bound of roughly

$$(L/t)^{\tilde{O}(\text{ddim})} \tag{19}$$

instead of the previous state-of-the-art bound of $(\|\cdot\|_{\text{Lip}}/t)^{\tilde{O}(\text{ddim})}$. The improvement can be dramatic, as the worst-case may be significantly (even infinitely) larger than the mean (7).

This simple result appears to be novel and interesting in its own right, but is insufficient to guarantee generalization bounds (via a Uniform Glivenko-Cantelli law), since the uniform convergence requires control over the *empirical* (i.e., $L_2(\mu_n)$) covering numbers. Bounding these proved to be a formidable challenge. The calculation in Theorem 4 reduces this problem to the one of bounding the empirical measure of the level set, $\mu_n(M_f(\ell))$, uniformly over all the functions in our class. We make the perhaps surprising discovery that (i) uniform control over the $\mu_n(M_f(\ell))$ is possible for the sub-class of functions free of certain local defects (Lemma 5) and (ii) any $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ is approximable in ℓ_∞ by a defect-free function (Lemma 6); see the beginning of Section 4 for some discussion and intuition. Together, these enable us to overcome the central challenge of controlling the empirical covering numbers (Theorem 3), yielding a bound comparable to (19). The implied generalization bounds (Section D) enjoy a dependence on $\sqrt{L}/n^{1/O(\text{ddim})}$, while all previously known generalization results for classification and regression feature a dependence on $\|\cdot\|_{\text{Lip}}$ (Tsybakov, 2004; Wainwright, 2019; Gottlieb et al., 2017).

Optimization and learning. From the perspective of supervised learning theory, our statistical bounds imply a non-trivial algorithmic problem: Given a labeled sample, produce a hypothesis whose true risk does not significantly exceed its empirical risk (with high probability). These notions are briefly defined in Section 1.3 and discussed in more detail in Section D.2. In light of the aforementioned generalization bounds, the learning procedure

may be recast as follows: The learner is given a “complexity budget” $L > 0$. Given a labeled sample, $(X_i, Y_i) \in \Omega \times [0, 1]$, $i \in [n]$, the learner seeks to fit to the data some function with average slope not exceeding L , while minimizing the empirical risk. The latter is induced by either the 0-1 loss (classification) or the L_1 loss (regression). Approximation algorithms for this problem are presented in Section 5. Briefly, we cast the regression problem as an optimization problem amenable to the mixed packing-covering framework of Koufogiannakis and Young (2014), and further improve the algorithmic run-time by reducing the number of constraints in the program (Section 5.1). Interestingly, the classification problem admits an efficient bi-criteria approximation when casting the “smoothness budget” in terms of the weak mean, but we were unable to find an efficient solution for the strong mean, and provide some indication that this may in fact be a hard problem (Section 5.2).

Adversarial extension. Having solved the learning problem, we have obtained an approximate minimizer of the empirical risk, but this does not immediately imply a bound on the true risk. To obtain such a bound, we demonstrate that with high probability, average smoothness under the empirical measure μ_n translates to average smoothness under the true sampling measure μ , from which a bound on true risk follows.

To this end, we define the following *adversarial extension* problem: An adversary draws n points $\Omega_n \subset \Omega$ from μ and labels them with $y \in [0, 1]$. This induces an average slope $\bar{\Lambda}_y$ (or $\tilde{\Lambda}_y$) under the empirical measure μ_n . The adversary’s goal is to force *any* extension of y from Ω_n to all of Ω to have a significantly larger average slope under the true measure μ . In the case of regression, we show that if the learner is willing to tolerate a small distortion of y under $L_1(\mu_n)$, it is possible to guarantee an at-most constant factor increase in both $\bar{\Lambda}$ and $\tilde{\Lambda}$ with high probability (Section 6). In the case of classification, we show how to achieve an most $2^{O(\text{ddim})}$ polylog(n) factor increase (with high probability), without incurring any distortion (Section 7).

3. Covering numbers

Our covering-numbers results will be stated for $\widetilde{\text{Lip}}_L$. It follows from (11) that these results hold verbatim for $\bar{\text{Lip}}_L$ as well. (These function classes are defined in (12, 13, 14)). It will be convenient to assume throughout that $\text{diam}(\Omega) = 1$; this incurs no loss of generality, since any bounded metric can be rescaled.

3.1 Ambient covering numbers

Our empirical covering-numbers results build upon the following simpler result for the ambient covering numbers. Results of this type trace back to the seminal work of Kolmogorov and Tihomirov (1961) and express the covering numbers of a well-behaved function class in terms the covering numbers of its metric space domain.

Theorem 1 (Ambient L_2 Covering Numbers) *For $L > 0$ and $0 < t < 1$,*

$$\log \mathcal{N}(t, \widetilde{\text{Lip}}_L(\Omega, \rho, \mu), L_2(\mu)) \leq \mathcal{N}(t^3/(64L), \Omega, \rho) \log(16/t).$$

In particular, for doubling spaces we have

$$\log \mathcal{N}(t, \widetilde{\text{Lip}}_L(\Omega, \rho, \mu), \mathbb{L}_2(\mu)) \leq \left(\frac{128L}{t^3} \right)^{\text{ddim}(\Omega)} \log \frac{16}{t}, \quad 0 < t < 1.$$

The proof of Theorem 1 will be based upon following result:

Lemma 2 (Gottlieb et al. (2017), Lemma 5.2) For $L > 0$ and $0 < t < 1$,

$$\log \mathcal{N}(t, \text{Lip}_L(\Omega, \rho), \ell_\infty) \leq \mathcal{N}(t/(8L), \Omega, \rho) \log(8/t).$$

In particular, for doubling spaces we have

$$\log \mathcal{N}(t, \text{Lip}_L(\Omega, \rho), \ell_\infty) \leq \left(\frac{16L}{t} \right)^{\text{ddim}(\Omega)} \log \frac{8}{t}, \quad 0 < t < 1.$$

Proof [Proof of Theorem 1] Recall the definition of the level set $M_f(\cdot)$ in (10). By (15), we have $\mu(M_f(L/t)) \leq t$ for any $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ and $t > 0$, and by construction, $\Lambda_f(x, \Omega) \leq L/t$ for all $x \in M'_f(L/t) := \Omega \setminus M_f(L/t)$. Thus, for all $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$, we have

$$f|_{M'_f(L/t)} \in \text{Lip}_{L/t}(M'_f(L/t), \rho). \quad (20)$$

Let \hat{F} be a $t/2$ -cover of $\text{Lip}_{4L/t^2}(\Omega, \rho)$ under ℓ_∞ . We claim that \hat{F} is a t -cover of $\widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ under $\mathbb{L}_2(\Omega, \rho, \mu)$. Indeed, choose an $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$. It follows from (20) that

$$f|_{M'_f(4L/t^2)} \in \text{Lip}_{4L/t^2}(M'_f(4L/t^2), \rho). \quad (21)$$

Via the McShane-Whitney Lipschitz extension (McShane, 1934; Whitney, 1934), there is an $\tilde{f} \in \text{Lip}_{4L/t^2}(\Omega, \rho)$ coinciding with f on $M'_f(4L/t^2)$. Since \hat{F} is a $t/2$ -cover, there is an $\hat{f} \in \hat{F}$ such that $\|\tilde{f} - \hat{f}\|_{\mathbb{L}_2(\mu)} \leq \|\tilde{f} - \hat{f}\|_\infty \leq t/2$. Therefore

$$\begin{aligned} \|f - \hat{f}\|_{\mathbb{L}_2(\mu)} &\leq \|f - \tilde{f}\|_{\mathbb{L}_2(\mu)} + \|\tilde{f} - \hat{f}\|_{\mathbb{L}_2(\mu)} \\ &= \left(\int_{\Omega} (f(x) - \tilde{f}(x))^2 d\mu(x) \right)^{1/2} + \|\tilde{f} - \hat{f}\|_{\mathbb{L}_2(\mu)} \\ &= \left(\int_{M'_f(4L/t^2)} (f(x) - \tilde{f}(x))^2 d\mu(x) \right)^{1/2} + \|\tilde{f} - \hat{f}\|_{\mathbb{L}_2(\mu)} \\ &\leq \left(\int_{M'_f(4L/t^2)} 1^2 d\mu(x) \right)^{1/2} + \|\tilde{f} - \hat{f}\|_{\mathbb{L}_2(\mu)} \\ &\leq t/2 + t/2 = t. \end{aligned}$$

The claim follows from Lemma 2, which bounds the size of (a minimal) \hat{F} , when plugging in $t/2$ in place of t and $4L/t^2$ in place of L . \blacksquare

3.2 Empirical covering numbers

The main result of this section is a bound on the empirical covering numbers. To avoid trivialities, we state our asymptotic bounds in n under the assumption that $\text{ddim}(\Omega), L \geq 1$.

Theorem 3 (Empirical L_2 Covering Numbers) *Let (Ω, ρ, μ) be a doubling metric measure space (the ambient space) with $\text{diam}(\Omega) \leq 1$ and (Ω_n, ρ, μ_n) its empirical realization in the sense of (4).*

Then, for constant $\delta > 0$ and $L, \text{ddim}(\Omega) \geq 1$, we have that

$$\log \mathcal{N}((\alpha + 1)\varepsilon_0, \widetilde{\text{Lip}}_L(\Omega, \rho, \mu), L_2(\mu_n)) \leq \left(\frac{L}{\alpha^3 \varepsilon_0^3} \right)^{\text{ddim}(\Omega)} \log \frac{1}{\alpha \varepsilon_0}, \quad \alpha > 0$$

holds simultaneously for all α with probability at least $1 - 2\delta$, where $\varepsilon_0 \leq C_\delta \sqrt{L} n^{-1/8d}$ and $C_\delta > 0$ is a constant depending only on δ .

The proof will be given below, and follows directly from Theorem 1 and the following result:

Theorem 4 (Preserving distances between $L_2(\mu)$ and $L_2(\mu_n)$) *Let (Ω, ρ, μ) be a doubling metric measure space (the ambient space) with $\text{diam}(\Omega) \leq 1$ and (Ω_n, ρ, μ_n) its empirical realization in the sense of (4). Then, with probability at least $1 - 2\delta$, all $f, g \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ satisfy*

$$\begin{aligned} \|f - g\|_{L_2(\mu_n)} &\leq 6\|f - g\|_{L_2(\mu)} \\ &+ 25n^{-1/4d} + 15\sqrt{L}n^{-1/8d} + (6 + 2^{d/4})n^{-1/8} + \left(\frac{162}{n} \log \frac{2}{\delta} \right)^{1/4}, \end{aligned}$$

where $d = \text{ddim}(\Omega)$ and we adhere to the notational convention in (3).

Proof Let $r, t, \eta \geq 0$ be parameters to be chosen later. Our first step is to approximate the function class $\widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ by a smaller, better-behaved class of functions, which have been “ η -smoothed” so as to repair certain defects defined in Section 4. The precise nature of these defects will not be important for the proof; at this point, we only define

$$\widetilde{\text{Lip}}_L^\eta(\Omega, \rho, \mu) := \left\{ f^\eta : f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu) \right\} \subseteq \widetilde{\text{Lip}}_L(\Omega, \rho, \mu),$$

where f^η is the function constructed in Lemma 6, when the latter is invoked with the parameter $\ell = L/t$. Lemma 6 guarantees that for every $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ there is an $f^\eta \in \widetilde{\text{Lip}}_L^\eta(\Omega, \rho, \mu)$ such that $\|f - f^\eta\|_\infty \leq 4\eta$ and $\Lambda_{f^\eta}(x, \Omega) \leq \Lambda_f(x, \Omega)$ for all $x \in \Omega$. (An additional fact about f^η , contained in Corollary 7, will be exploited later in this proof.) Thus, for $f, g \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$,

$$\begin{aligned} \|f - g\|_{L_2(\mu_n)} &\leq \|f - f^\eta\|_\infty + \|g - g^\eta\|_\infty + \|f^\eta - g^\eta\|_{L_2(\mu_n)} \\ &\leq 8\eta + \|f^\eta - g^\eta\|_{L_2(\mu_n)}, \end{aligned} \tag{22}$$

and so it will suffice to bound $\|f^\eta - g^\eta\|_{L_2(\mu_n)}$.

Let $V \subset \Omega$ be an r -net of (Ω, ρ) ; by the doubling property (1),

$$|V| \leq (2/r)^{\text{ddim}(\Omega)}. \quad (23)$$

The net V induces the Voronoi partition $\Omega = \bigcup_{v \in V} W(v)$, such that for each cell $W(v)$ we have $x \in W(v) \implies \rho(x, v) \leq r$, as well as the measure under μ , denoted by $\pi(v) := \mu(W(v))$. Together, these induce the finite metric measure space (V, ρ, π) . The map $\varphi_V : \Omega \rightarrow V$ takes each $x \in \Omega$ with its Voronoi cell; thus, $\varphi_V^{-1}(\varphi_V(x)) = W(v)$ for all $x \in W(v)$.

The proof proceeds in several steps, in which always $f, g \in \widetilde{\text{Lip}}_L^\eta(\Omega, \rho, \mu)$ and the notational convention in (3) is used. Define $M_f(L/t), M_g(L/t) \subset \Omega$ as in (10). As in the proof of Theorem 1, we have that f is L/t -Lipschitz on $M'_f(L/t) := \Omega \setminus M_f(L/t)$, with extension $\tilde{f} \in \text{Lip}_{L/t}(\Omega, \rho)$. Define $\tilde{g} \in \text{Lip}_{L/t}(\Omega, \rho)$ analogously, as extending $g|_{M'_g(L/t)} \in \text{Lip}_{L/t}(M'_g(L/t), \rho)$.

Comparing the norms $\|f - g\|_{\mathbb{L}_2(\mu_n)}^2 = \frac{1}{n} \sum_{i=1}^n (f(X_i) - g(X_i))^2$ and $\|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\pi)}^2 = \sum_{v \in V} (\tilde{f}(v) - \tilde{g}(v))^2 \pi(v)$. We begin by invoking (50):

$$\|f - g\|_{\mathbb{L}_2(\mu_n)}^2 \leq 3\|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\mu_n)}^2 + 3\|f - \tilde{f}\|_{\mathbb{L}_2(\mu_n)}^2 + 3\|g - \tilde{g}\|_{\mathbb{L}_2(\mu_n)}^2. \quad (24)$$

The second and third terms in the bound (24) are bounded identically:

$$\begin{aligned} \|f - \tilde{f}\|_{\mathbb{L}_2(\mu_n)}^2 &= \frac{1}{n} \sum_{i=1}^n (f(X_i) - \tilde{f}(X_i))^2 \\ &\leq \frac{1}{n} \sum_{i: X_i \in M_f(L/t)} 1 = \mu_n(M_f(L/t)). \end{aligned} \quad (25)$$

To estimate the first term in the bound (24), recall that \tilde{f} is L/t -Lipschitz on Ω and $\rho(x, \varphi_V(x)) \leq r$ (and the same holds for \tilde{g}), whence

$$\begin{aligned} |\tilde{f}(X_i) - \tilde{g}(X_i)| &\leq |\tilde{f}(X_i) - \tilde{f}(\varphi_V(X_i))| + |\tilde{g}(X_i) - \tilde{g}(\varphi_V(X_i))| + |\tilde{f}(\varphi_V(X_i)) - \tilde{g}(\varphi_V(X_i))| \\ &\leq |\tilde{f}(\varphi_V(X_i)) - \tilde{g}(\varphi_V(X_i))| + 2Lr/t. \end{aligned}$$

Using $(a + b)^2 \leq 2a^2 + 2b^2$, this yields

$$\begin{aligned} \|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\mu_n)}^2 &= \frac{1}{n} \sum_{i=1}^n (\tilde{f}(X_i) - \tilde{g}(X_i))^2 \\ &\leq \frac{2}{n} \sum_{i=1}^n (\tilde{f}(\varphi_V(X_i)) - \tilde{g}(\varphi_V(X_i)))^2 + 8(Lr/t)^2 \\ &= 2 \sum_{v \in V} (\tilde{f}(v) - \tilde{g}(v))^2 \mu_n(W(v)) + 8(Lr/t)^2 \\ &= 2\|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\pi_n)}^2 + 8(Lr/t)^2, \end{aligned}$$

where π_n is the measure on V given by $\pi_n(v) = \mu_n(W(v))$. Observe that

$$\begin{aligned} \|\|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\pi)}^2 - \|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\pi_n)}^2\| &\leq \sum_{v \in V} (\tilde{f}(v) - \tilde{g}(v))^2 |\pi(v) - \pi_n(v)| \\ &\leq \sum_{v \in V} |\pi(v) - \pi_n(v)| = \|\pi - \pi_n\|_1. \end{aligned}$$

A bound on $\|\pi - \pi_n\|_1$ is provided by (54): with probability at least $1 - \delta$,

$$\|\pi - \pi_n\|_1 \leq \sqrt{\frac{|V|}{n}} + \sqrt{\frac{2}{n} \log \frac{2}{\delta}} \leq \sqrt{\frac{(2/r)^{\text{ddim}(\Omega)}}{n}} + \sqrt{\frac{2}{n} \log \frac{2}{\delta}}.$$

To bound $\mu_n(M_f(L/t))$ in (25), we invoke Corollary 7: with probability at least $1 - \delta$,

$$\sup \left\{ \mu_n(M_f(L/t)) : f \in \widetilde{\text{Lip}}_L^\eta(\Omega, \rho, \mu) \right\} \leq 24t + \frac{1}{2} \sqrt{\frac{(2L/\eta t)^{\text{ddim}(\Omega)}}{n}} + \frac{1}{2} \sqrt{\frac{2}{n} \log \frac{2}{\delta}}.$$

These calculations culminate in the bound

$$\begin{aligned} \|f - g\|_{\mathbb{L}_2(\mu_n)}^2 &\leq 6\|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\pi)}^2 + 24(Lr/t)^2 + 144t \\ &\quad + 6\sqrt{\frac{(2/r)^{\text{ddim}(\Omega)}}{n}} + 9\sqrt{\frac{2}{n} \log \frac{2}{\delta}} + 3\sqrt{\frac{(2L/\eta t)^{\text{ddim}(\Omega)}}{n}}. \end{aligned} \quad (26)$$

Comparing the norms $\|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\pi)}$ **and** $\|f - g\|_{\mathbb{L}_2(\mu)}$. Since \tilde{f} is L/t -Lipschitz and $\text{diam}(W(v)) \leq r$, we have $\tilde{f}(W(v)) \subseteq f(v) \pm Lr/t$, and analogously for \tilde{g} . It follows that

$$|\tilde{f}(v) - \tilde{g}(v)| \leq |\tilde{f}(x) - \tilde{g}(x)| + 2Lr/t, \quad x \in W(v) \quad (27)$$

and hence

$$(\tilde{f}(v) - \tilde{g}(v))^2 \leq 2(\tilde{f}(x) - \tilde{g}(x))^2 + 8(Lr/t)^2, \quad x \in W(v). \quad (28)$$

Integrating,

$$\begin{aligned} \|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\pi)}^2 &= \sum_{v \in V} (\tilde{f}(v) - \tilde{g}(v))^2 \pi(v) = \sum_{v \in V} \int_{W(v)} (\tilde{f}(v) - \tilde{g}(v))^2 d\mu(x) \\ &\leq \sum_{v \in V} \int_{W(v)} [2(\tilde{f}(x) - \tilde{g}(x))^2 + 8(Lr/t)^2] d\mu(x) \\ &= 2\|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\mu)}^2 + 8(Lr/t)^2. \end{aligned}$$

Using (50) and the triangle inequality, we have

$$\begin{aligned} \|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\mu)}^2 &\leq 3\|f - g\|_{\mathbb{L}_2(\mu)}^2 + 3\|f - \tilde{f}\|_{\mathbb{L}_2(\mu)}^2 + 3\|g - \tilde{g}\|_{\mathbb{L}_2(\mu)}^2 \\ &\leq 3\|f - g\|_{\mathbb{L}_2(\mu)}^2 + 6t \end{aligned}$$

(since each of $\|f - \tilde{f}\|_{\mathbb{L}_2(\mu)}$, $\|g - \tilde{g}\|_{\mathbb{L}_2(\mu)}$ is at most \sqrt{t}).

Combining these yields

$$\|\tilde{f} - \tilde{g}\|_{\mathbb{L}_2(\pi)}^2 \leq 6\|f - g\|_{\mathbb{L}_2(\mu)}^2 + 12t + 8(Lr/t)^2. \quad (29)$$

Finishing up. Combining (26) and (29) yields

$$\begin{aligned} \|f - g\|_{\mathbb{L}_2(\mu_n)}^2 &\leq 36\|f - g\|_{\mathbb{L}_2(\mu)}^2 + 72(Lr/t)^2 + 216t \\ &\quad + 6\sqrt{\frac{(2/r)^{\text{ddim}(\Omega)}}{n}} + 9\sqrt{\frac{2}{n} \log \frac{2}{\delta}} + 3\sqrt{\frac{(2L/\eta t)^{\text{ddim}(\Omega)}}{n}}. \end{aligned} \quad (30)$$

Choosing r, t, η . Putting $d = \text{ddim}(\Omega)$, we choose $r = n^{-1/2d}$, $t = Ln^{-1/4d}$, and $\eta = 2n^{-1/4d}$. For this choice,

$$\begin{aligned} \|f - g\|_{\mathbf{L}_2(\mu_n)}^2 &\leq 36\|f - g\|_{\mathbf{L}_2(\mu)}^2 + 72n^{-1/2d} + 216Ln^{-1/4d} \\ &\quad + 2^{d/2} \cdot 6n^{-1/4} + 9\sqrt{\frac{2}{n} \log \frac{2}{\delta}} + 3 \left(n^{1/4d}\right)^{d/2} n^{-3/8}. \end{aligned}$$

Applying the inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ to (30) proves the claim. \blacksquare

Proof [Proof of Theorem 3] Let ε_+ be the additive term in the bound in Theorem 4:

$$\varepsilon_+ = 25n^{-1/4d} + 15\sqrt{Ln}^{-1/8d} + (6 + 2^{d/4})n^{-1/8} + \left(\frac{162}{n} \log \frac{2}{\delta}\right)^{1/4}. \quad (31)$$

Since for fixed $\delta > 0$ and $d, L \geq 1$ the dominant term in (32) is $15\sqrt{Ln}^{-1/8d}$, and so there is a $C = C_\delta < 0$ such that

$$\varepsilon_+ \leq C_\delta \sqrt{Ln}^{-1/8d} =: \varepsilon_0. \quad (32)$$

Now Theorem 4 implies that any ε -cover of $\widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ under $\mathbf{L}_2(\mu)$ also provides a $(6\varepsilon + \varepsilon_0)$ -cover under $\mathbf{L}_2(\mu_n)$. Equivalently, an $\alpha\varepsilon_0/6$ -cover of the former yields an $(\alpha + 1)\varepsilon_0$ -cover of the latter, for $\alpha > 0$. Hence,

$$\begin{aligned} \log \mathcal{N}((\alpha + 1)\varepsilon_0, \widetilde{\text{Lip}}_L(\Omega, \rho, \mu), \mathbf{L}_2(\mu_n)) &\leq \log \mathcal{N}(\alpha\varepsilon_0/6, \widetilde{\text{Lip}}_L(\Omega, \rho, \mu), \mathbf{L}_2(\mu)) \\ &\leq \left(\frac{cL}{\alpha^3\varepsilon_0^3}\right)^{\text{ddim}(\Omega)} \log \frac{1}{\alpha\varepsilon_0}, \end{aligned}$$

where $c > 0$ is a universal constant. \blacksquare

4. Defect free functions

This section presents results that were invoked in the proofs in Section 3. It constitutes the core of the analytic and combinatorial structure we discovered in the very general setting of real-valued functions on metric spaces. Such a function may fail to be locally smooth for two “moral” reasons: due to “large jumps” or “small jumps”. The former is witnessed by two nearby points x, x' for which $|f(x) - f(x')|$ is large — say, 1. The latter is witnessed by two nearby (say, ε -close) points x, x' for which $\varepsilon \ll |f(x) - f(x')| \leq T(\varepsilon) \ll 1$ — say, $T(\varepsilon) = \sqrt{\varepsilon}$. It turns out that the large jumps do not present a problem for the combinatorial structure we seek in Lemma 5, which forms the basis for Corollary 7, the latter a crucial component in the empirical covering number bound, Theorem 4. Rather, it is the small jumps — which we formalize as *defects* below — that present an obstruction. Fortunately, as we show in Lemma 6, *any* bounded real-valued function on a doubling metric space admits a defect-free approximation under ℓ_∞ .

4.1 Definition and structure

For a given $f : \Omega \rightarrow \mathbb{R}$, $\ell > 0$, and $x, y \in \Omega$, we say that y is an ℓ -**slope witness** for x (w.r.t. f) if $|f(x) - f(y)|/\rho(x, y) \geq \ell$. For $\eta, \ell > 0$ and $c \geq 1$, we say that an $x \in \Omega$ is an (η, ℓ, c) -**defect** of f if:

- (a) $\Lambda_f(x) \geq \ell$
- (b) Every ℓ/c -slope witness y of x satisfies $|f(x) - f(y)| \leq \eta$.

Define $\Xi_f(\eta, \ell, c) \subseteq \Omega$ to be the set of all (η, ℓ, c) -defects of f . Note that $\Xi_f(\eta, \ell, c) \subseteq \Xi_f(\eta, \ell, c')$ whenever $c' \leq c$. For $\eta, \ell > 0$, $c \geq 1$ define $\mathcal{G}(\eta, \ell, c)$ to be the collection of all $f : \Omega \rightarrow [0, 1]$ such that f does not have any (η, ℓ, c) -defects.

Lemma 5 (Combinatorial structure of defect-free functions) *For every $0 < \eta < \ell$, $c \geq 1$ there is a partition $\Pi = \{B_1, \dots, B_N\}$ of Ω of size $N \leq (2\ell/\eta)^{\text{ddim}(\Omega)}$ such that for each $f \in \mathcal{G}(\eta, \ell, c)$, we have*

$$M_f(\ell) \subseteq U_f \subseteq M_f(\ell/4c), \quad (33)$$

where $M_f(\cdot)$ is the level set defined in (10) and

$$U_f := \bigcup \{B \in \Pi : B \cap M_f(\ell) \neq \emptyset\}. \quad (34)$$

Proof Let Π be the Voronoi partition induced by an η/ℓ -net of Ω . Then the claimed bound on $|\Pi|$ holds by (1) and the first inclusion in (33) is obvious by construction; it only remains to show that $U_f \subseteq M_f(\ell/4c)$.

Choose any $u \in U_f$. Since Π is induced by a net, there is some $x \in M_f(\ell)$ for which $\rho(x, u) \leq \eta/\ell$. Since f has no (η, ℓ, c) -defects and $\Lambda_f(x) \geq \ell$, there must be some ℓ/c -slope witness $y \in \Omega$ of x for which $|f(x) - f(y)| > \eta$. Invoking (51), we have $|f(u) - f(x)| \vee |f(u) - f(y)| \geq |f(x) - f(y)|/2 > \eta/2$. We consider the two cases:

- (i) $\rho(x, y) \leq \eta/\ell$
- (ii) $\rho(x, y) > \eta/\ell$.

In the first case, the triangle inequality implies that $\rho(u, x) \vee \rho(u, y) \leq 2\eta/\ell$, and hence

$$\Lambda_f(u) \geq \frac{|f(u) - f(x)|}{\rho(u, x)} \vee \frac{|f(u) - f(y)|}{\rho(u, y)} \geq \frac{\eta/2}{2\eta/\ell} = \frac{\ell}{4} \geq \frac{\ell}{4c} \implies u \in M_f(\ell/4c).$$

For the second case, if $|f(u) - f(x)| \geq \eta/2$, the proof is the same as in the first case. Otherwise, $|f(u) - f(y)| \geq |f(x) - f(y)|/2$ and so

$$\frac{|f(u) - f(y)|}{\rho(u, y)} \geq \frac{|f(x) - f(y)|}{2\rho(u, y)} \geq \frac{|f(x) - f(y)|}{2 \cdot 2\rho(x, y)} \geq \frac{\ell}{4c} \implies u \in M_f(\ell/4c),$$

where the second inequality is a result of applying the triangle inequality to the fact that $\rho(u, x) < \rho(x, y)$. ■

4.2 Repairing defects

The main result of this section is that the problematic “small jumps” alluded to in the beginning of Section 4 can be smoothed out via an ℓ_∞ approximation.

Lemma 6 (Defect repair) *For each $\eta, \ell > 0$ and $f : \Omega \rightarrow [0, 1]$, there is an $\bar{f} \in \mathcal{G}(\eta, \ell, c = 6)$ such that $\|f - \bar{f}\|_\infty \leq 4\eta$ and*

$$\Lambda_{\bar{f}}(x, \Omega) \leq \Lambda_f(x, \Omega), \quad x \in \Omega. \quad (35)$$

Proof We will prove the equivalent claim that $\bar{f} \in \mathcal{G}(\eta/2, \ell, c = 6)$ and $\|f - \bar{f}\|_\infty \leq 2\eta$. We begin by constructing \bar{f} . Let $M_f(\ell)$ be as defined in (10) and V be a η/ℓ -net of this set. Partition $V = V_0 \cup V_1$, where V_0 is “smooth,”

$$V_0 := \{v \in V : B(v, \eta/\ell) \not\subseteq \Xi_f(\eta, \ell, 1)\},$$

and V_1 is “rough,”

$$V_1 := \{v \in V : B(v, \eta/\ell) \subseteq \Xi_f(\eta, \ell, 1)\}.$$

Define

$$A_f := \bigcup_{v_1 \in V_1} B(v_1, \eta/\ell) \setminus \left(\bigcup_{v_0 \in V_0} B(v_0, \eta/\ell) \cup V \right).$$

In words, A_f consists of the entirely defective (or “rough”) balls without their center-points or their intersections with smooth balls. Define \bar{f} as the PMSE extension of f from $\Omega \setminus A_f$ to Ω , as in Definition 18. Having constructed the \bar{f} , we proceed to verify its properties.

Proof that (35) holds. This is an immediate consequence of Theorem 23.

Proof that $\|f - \bar{f}\|_\infty \leq 2\eta$. Since PMSE is an extension, we need only establish $|f(x) - \bar{f}(x)| \leq 2\eta$ for $x \in A_f$. For any such x , the definition of A_f implies the existence of some $v_1 \in V_1$ for which $\rho(x, v_1) \leq \eta/\ell$. Since $|f(x) - \bar{f}(x)| \leq |f(v_1) - f(x)| + |\bar{f}(x) - f(v_1)|$, it is sufficient to bound each term separately by η .

To bound the first term, assume, for a contradiction, that $|f(v_1) - f(x)| > \eta$. Then $\frac{|f(v_1) - f(x)|}{\rho(v_1, x)} > \frac{\eta}{\eta/\ell} = \ell$, contradicting the defectiveness of v_1 .

To bound the second term, again assume for a contradiction that $\bar{f}(x) - f(v_1) > \eta$; the case $f(v_1) - \bar{f}(x) > \eta$ is handled analogously. Our assumption implies $\Lambda_{\bar{f}}(x) \geq \frac{\bar{f}(x) - \bar{f}(v_1)}{\rho(x, v_1)} \geq \ell$. By properties (i) and (v) of Corollary 24, there is an $x' \in \Omega \setminus A_f$ for which $\Lambda_{\bar{f}}(x) = \frac{\bar{f}(x') - \bar{f}(x)}{\rho(x', x)} \geq \ell$ and $\bar{f}(x') \geq \bar{f}(x)$. Invoking (53), we have $\frac{f(x') - f(v_1)}{\rho(x', v_1)} \geq \ell$. Additionally, we have $\bar{f}(x') - \bar{f}(v_1) > \bar{f}(x) - \bar{f}(v_1) > \eta$, again contradicting the defectiveness of v_1 .

Proof that $\bar{f} \in \mathcal{G}(\eta/2, \ell, 6)$. A statement equivalent to $\bar{f} \in \mathcal{G}(\eta/2, \ell, 6)$ is that $\Xi_{\bar{f}}(\eta/2, \ell, 6) = \emptyset$. Let us define the sets

$$\begin{aligned} E_1 &:= \bigcup_{v_0 \in V_0} B(v, \eta/\ell), \\ E_2 &:= A_f \cup V_1, \\ E_3 &:= \Omega \setminus M_f(\ell), \end{aligned}$$

which are, by construction, a (not necessarily disjoint) cover of Ω . Hence, it suffices to show that $\Xi_{\bar{f}}(\eta/2, \ell, 6) \cap E_i = \emptyset$ for $i \in [3]$.

Let $x \in E_3$. By Theorem 23, $\Lambda_{\bar{f}}(x, \Omega) \leq \Lambda_f(x, \Omega)$. Since $x \notin M_f(\ell)$, we have that $\Lambda_{\bar{f}}(x) < \ell$, implying that x is not $(\eta/2, \ell, c)$ -defect for any $c \geq 1$ with respect to \bar{f} .

Let $x \in E_1$. Then there is a $v_0 \in V_0$ such that $\rho(x, v_0) \leq \eta/\ell$. Being in V_0 implies that v_0 has some ℓ -slope witness $v'_0 \in E_1$ such that $|f(v_0) - f(v'_0)| > \eta$. This implies by (51) that $|\bar{f}(x) - \bar{f}(v_0)| \vee |\bar{f}(x) - \bar{f}(v'_0)| > |f(v_0) - f(v'_0)|/2 > \eta/2$, since f and \bar{f} must agree on v_0 and v'_0 . The triangle inequality yields a slope of at least $\ell/4$ witnessed by at least one of $\{v_0, v'_0\}$.

Let $x \in E_2$. Then there is a $v_1 \in V_1 \subseteq \Xi_f(\eta, \ell, 1)$ for which $\rho(x, v_1) \leq \eta/\ell$. By Remark 21, the maximal slope at x is achieved at the two distinct points u^* and v^* by which it is determined. Suppose $\rho(x, u^*) \vee \rho(x, v^*) > \eta/2\ell$. If $\Lambda_f(x) \geq \ell$ and one of $\{u^*, v^*\}$ — say, u^* — satisfies the inequality. Then

$$|f(x) - f(u^*)| \geq \ell \rho(x, u^*) > \ell \cdot \frac{\eta}{2\ell} = \frac{\eta}{2},$$

which contradicts the second condition for an $(\eta/2, \ell, c)$ -defect. Otherwise, $\rho(x, u^*) \vee \rho(x, v^*) \leq \eta/2\ell$. Since V is an η/ℓ -net, it is not possible that both $u^*, v^* \in V$. Therefore (without loss of generality) $u^* \in E_1 \cup E_3$. If $u^* \in E_3$, it follows from property (ii) in Corollary 24 that

$$\Lambda_{\bar{f}}(x, \Omega) \leq \Lambda_{\bar{f}}(u^*, \Omega) \leq \Lambda_f(u^*, \Omega) < \ell,$$

which implies that $x \notin \Xi_{\bar{f}}(\eta/2, \ell, c)$ for any $c \geq 1$. If $u^* \in E_1$ then there is some $v_0 \in V_0$ for which $\rho(u^*, v_0) \leq \eta/\ell$. Since $\Lambda_f(v_0) > \ell$ and $v_0 \notin \Xi_f(\eta, \ell, 1)$, it must have some witness v'_0 such that $|f(v_0) - f(v'_0)| > \eta$ and the slope between them is at least ℓ . Similarly to previous arguments, $|\bar{f}(x) - \bar{f}(v_0)| \vee |\bar{f}(x) - \bar{f}(v'_0)| > |f(v_0) - f(v'_0)|/2 \geq \eta/2$. In either case, applying the triangle inequality yields a slope of at least $\ell/6$ witnessed by at least one of $\{v_0, v'_0\}$, which shows that $x \notin \Xi_{\bar{f}}(\eta/2, \ell, 6)$. ■

The culmination of this section is the following crucial uniform convergence result invoked in the course of proving Theorem 4:

Corollary 7 *Let f^η be the function constructed from f as in Lemma 6, when the latter is invoked with the parameter $\ell = L/t$, and let*

$$\widetilde{\text{Lip}}_L^\eta(\Omega, \rho, \mu) := \left\{ f^\eta : f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu) \right\} \subseteq \widetilde{\text{Lip}}_L(\Omega, \rho, \mu).$$

Then, with probability at least $1 - \delta$, we have

$$\sup \left\{ \mu_n(M_f(L/t)) : f \in \widetilde{\text{Lip}}_L^\eta(\Omega, \rho, \mu) \right\} \leq 24t + \frac{1}{2} \sqrt{\frac{(2L/\eta t)^{\text{ddim}(\Omega)}}{n}} + \frac{1}{2} \sqrt{\frac{2}{n} \log \frac{2}{\delta}},$$

where μ_n is the empirical measure induced by μ .

Proof Let Π be as in Lemma 5. For each $f \in \widetilde{\text{Lip}}_L^\eta(\Omega, \rho, \mu)$, let U_f be as defined in (34). Then, invoking the inclusion in (33) and recalling that $\mu(M_f(L/t)) \leq t$ for all $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ and $t > 0$ (and that Lemma 6 sets $c = 6$),

$$\begin{aligned}
 \sup_f \mu_n(M_f(L/t)) &\leq \sup_f \mu_n(U_f) \\
 &= \sup_f [\mu_n(U_f) - \mu(U_f)] + \mu(U_f) \\
 &\leq \sup_f (\mu_n(U_f) - \mu(U_f)) + \sup_f \mu(U_f) \\
 &\leq \sup_{U \subseteq \Pi} (\mu_n(U) - \mu(U)) + \sup_f \mu(M_f(L/4ct)) \\
 &\leq \sup_{U \subseteq \Pi} (\mu_n(U) - \mu(U)) + 24t \\
 &= \frac{1}{2} \sum_{B \in \Pi} |\mu(B) - \mu_n(B)| + 24t,
 \end{aligned}$$

where the last step used the variational characterization of the total variation distance. The latter is bounded as in (54), completing the proof. ■

5. Learning algorithms: training

We consider two learning problems — classification and regression — in a unified *agnostic* setting (Mohri et al., 2012). In each case, the learner receives a labeled sample, $(X_i, Y_i)_{i \in [n]}$, where $X_i \in \Omega$ and $Y_i \in \{0, 1\}$ for classification or $Y_i \in [0, 1]$ for regression. The learner then selects a hypothesis $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$, where L is fixed a priori.³ Finally, given a test point $x' \in \Omega$, the learner’s predicted label is either $f(x') \in [0, 1]$ (regression) or $\mathbf{1}[f(x') > 1/2] \in \{0, 1\}$ (classification); this is elaborated in greater detail in Section D.2. Computational considerations, as well as the learner’s inherent uncertainty regarding whether $f \in \widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ (see below), will lead us to consider relaxed versions of the learning problem, where the “complexity budget” will increase from L to $O(L)$ for regression and $O(L \cdot \text{polylog } n)$ for classification.

As described in Section 1.3, the sample is drawn from the joint measure ν over $\Omega \times [0, 1]$ — whose first marginal, by definition, necessarily coincides with μ — and, once drawn, induces the empirical measure ν_n . The *empirical* (respectively, *true*) risk of $f : \Omega \rightarrow \mathbb{R}$ is the expected value of $|f(X) - Y|$ under ν (respectively, ν_n); these are denoted by $R(f; \nu)$ and $R(f; \nu_n)$. The learner seeks to minimize $R(f; \nu)$ but can only directly access $R(f; \nu_n)$; hence, an *optimization* algorithm will seek to minimize the latter, while a *generalization* bound will provide a high-confidence bound on the former.

Our learning problem presents a novel challenge, not typically encountered in the classic supervised learning setting. Namely, ensuring that the learner’s hypothesis belongs to $\widetilde{\text{Lip}}_L(\Omega, \rho, \mu)$ (or $\overline{\text{Lip}}_L(\Omega, \rho, \mu)$) is non-trivial, and is certainly not guaranteed “by construction”. The core of the challenge is that although the hypothesis class of on-average smooth

3. Assuming L fixed and known incurs no loss of generality, as discussed at the beginning of Section D.

functions is fixed in advance of seeing the data, it is distribution-dependent and not explicitly known to the learner, insofar as the distribution itself is unknown. Indeed, let us break down the learning process into its basic stages: training and prediction. The training stage, which may be called *smoothing* or *denoising* (or yet *regularization*), involves solving the following optimization problem: Choose a hypothesis f that stays within the “smoothness budget” and achieves a low empirical risk $R(f; \nu_n)$. Algorithmically, this is done by computing an $\hat{f} : \Omega_n \rightarrow [0, 1]$, where $\Omega_n = (X_i)_{i \in [n]}$ and $\hat{f}(X_i)$ is a “smoothed” version of the Y_i , achieving a desired average empirical slope $\hat{\Lambda}_{\hat{f}}(\mu_n, \Omega_n)$ (or $\bar{\Lambda}_{\hat{f}}(\mu_n, \Omega_n)$). The function \hat{f} is then extended via (a variant of) PMSE from Ω_n to all of Ω . The novel challenge arises at the prediction stage, and entails ensuring that $\hat{\Lambda}_{\hat{f}}(\mu, \Omega)$ (respectively, $\bar{\Lambda}_{\hat{f}}(\mu, \Omega)$) does not much exceed its empirical version. We term this problem *adversarial extension* and address it in Sections 6 and 7.

The results for regression are conceptually simpler and are presented first; those for classification follow. Throughout this section, we assume $\text{diam}(\Omega) \leq 1$ and $d := \text{ddim}(\Omega) < \infty$.

5.1 Regression

Theorem 8 (Training and generalization for regression, strong mean.) *Let ν be some distribution on $\Omega \times [0, 1]$, and $S_n = (X_i, Y_i)_{i \in [n]} \sim \nu^n$ be a set sampled i.i.d. from ν . Denote by $\hat{f} \in \overline{\text{Lip}}_L(\Omega_n, \rho, \mu_n)$ the minimizer of $R(\cdot; \nu_n)$. Then there is an efficient learning algorithm \mathcal{A} that constructs a hypothesis $f = \mathcal{A}(S_n)$ such that for any given $L > 0$, $0 < \varepsilon, \delta < 1$ and $0 < c < 1$:*

(a) *With probability at least $1 - \exp(-n(\varepsilon/8)^{d+1} + d \ln(8/\varepsilon)) - 3\delta$,*

$$R(f; \nu) \leq (1 + c)R(\hat{f}; \nu_n) + O\left(\varepsilon L + \frac{C_\delta \sqrt{L}}{n^{1/8d}}\right) + \frac{C_\delta^{-d/2} \sqrt{2}}{n^{5/16}} + 3\sqrt{\frac{\log(2/\delta)}{2n}},$$

(b) *$f(x)$ can be evaluated at each $x \in \Omega$ in time $O(n^2)$ after a one-time “smoothing” computation of $\min\{2^{O(d)}(n/c^2) \log \Delta, O((n/c)^2 \log n)\}$ where $\Delta = 1/\min_{x \neq x' \in x_{[n]}} \rho(x, x')$*

where C_δ is a constant depending only on δ .

Remark. The multiplicative empirical risk parameter c in (a) can be made arbitrarily small, at the expense of an increased runtime, as in (b). As c becomes arbitrarily small and n grows, the risk of $f = \mathcal{A}(S_n)$ approaches that of $\inf_{h \in \overline{\text{Lip}}_L(\Omega, \rho, \mu)} R(h, \nu)$ via a standard uniform convergence argument (67). The rate of risk decay in (a) is far from optimal. In the realizable (i.e., $\inf_{h \in \overline{\text{Lip}}_L(\Omega, \rho, \mu)} R(h, \nu) = 0$) case, near-optimal rates have been obtained in a follow-up work to this one (Hanneke et al., 2023). Computing and achieving minimax rates for the agnostic case is currently an open problem.

Proof The smoothing algorithm described in Lemma 9 constructs an approximate minimizer $\tilde{f} \in \overline{\text{Lip}}_{O(1)L}(\Omega_n, \rho, \mu_n)$ of $R(\cdot; \nu_n)$ and the “adversarial extension” algorithm in Lemma 13 provides an extension f of \tilde{f} from Ω_n to Ω that, with high probability, belongs to $\overline{\text{Lip}}_{O(1)L}(\Omega, \rho, \mu)$ and increases the empirical risk by at most an additive $O(\varepsilon L)$. The bound in (a) is then a direct application of (67).

■

Lemma 9 (Smoothing for regression, strong mean) *Let (Ω, ρ) be a metric space with $\text{diam}(\Omega) \leq 1$ and $\text{ddim}(\Omega) < \infty$. Suppose that $(x, y) = (x_{[n]}, y_{[n]}) \in (\Omega^n, \mathbb{R}^n)$ and $L > 0$ are given, and denote*

$$A(f; x, y, L) := \{\|f - y\|_{L_1(\mu_n)} : f \in \overline{\text{Lip}}_L(x_{[n]}, \rho, \mu_n)\}, \quad (36)$$

where μ_n is the counting measure on $x_{[n]}$.

Then a $(1 + c)$ -approximate minimizer $\hat{f} \in \overline{\text{Lip}}_L(x_{[n]}, \rho, \mu_n)$ of $A(\cdot; x, y, L)$ can be computed in time

$$\min\{2^{O(\text{ddim}(\Omega))}(n/c^2) \log \Delta, O((n/c)^2 \log n)\}.$$

Proof We cast the optimization problem as a linear program over the variables L_i, w_i, z_i :

$$\begin{aligned} \text{Minimize} \quad & W = \sum_{i \in [n]} w_i \\ \text{subject to} \quad & \sum_{i \in [n]} L_i \leq L \\ & w_i \geq |z_i - y_i| \quad \forall i \in [n] \\ & |z_i - z_j| \leq L_i \rho(x_i, x_j) \quad \forall i, j \in [n] \\ & 0 \leq w_i, z_i \leq 1 \quad \forall i \in [n]. \end{aligned}$$

A linear program in $O(n)$ variables and constraints can be solved in time $\tilde{O}(n^\omega)$ (Cohen et al., 2019), where ω is the best exponent for matrix inversion, currently $\omega \approx 2.37$.

First runtime improvement. To improve on the runtime, we will utilize the packing-covering framework of Koufogiannakis and Young (2014). For a constraint matrix of at most m rows and columns with all non-negative entries and at most ζ non-zero entries, the algorithm computes in time $O((m/c^2) \log \zeta + \zeta)$ a $(1 + c)$ -approximate solution satisfying all constraints. A difficulty in utilizing this framework is that our constraint matrix has negative entries; in particular, each constraint of the form

$$|z_i - z_j| \leq L_i \cdot \rho(x_i, x_j)$$

reduces to solving two constraints of the form

$$\begin{aligned} z_i - z_j &\leq L_i \cdot \rho(x_i, x_j) \\ z_j - z_i &\leq L_i \cdot \rho(x_i, x_j). \end{aligned}$$

To address this, we introduce dummy variables \tilde{z}_i satisfying $z_i + \tilde{z}_i = 1$. Then the above constraints become:

$$\begin{aligned} L_i \cdot \rho(x_i, x_j) + \tilde{z}_i + z_j &\geq 1 \\ L_i \cdot \rho(x_i, x_j) + z_i + \tilde{z}_j &\geq 1. \end{aligned}$$

Similarly, the constraint

$$w_i \geq |z_i - y_i|$$

is replaced by two constraints

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

For the runtime, we have that both terms m, ζ are bounded by $O(n^2)$, for a total runtime of $O(n^2 \log n)$.

Second runtime improvement. The main obstacle to improving the above runtime lies in the quadratic number of constraints necessary to compute the average slope. Here we show that we can reduce these to only $2^{O(\text{ddim})} n \log \Delta$ constraints, each with a constant number of variables, and so the linear program of Koufogiannakis and Young (2014) will run in time $2^{O(\text{ddim})} \tilde{O}(n \log \Delta)$. However, this comes at a cost of increasing the average slope by a constant factor.

We first extract from $x_{[n]} = (x_1, \dots, x_n)$ a point hierarchy $\{H_{2^{-k}}\}_{k=0}^{\lceil \log \Delta \rceil}$. Let $P(x, k)$ be the nearest neighbor of $x \in x_{[n]}$ in level $H_{2^{-k}}$, and for each point $x' \in x_{[n]}$, let neighborhood $N(x', k)$ include all points x for which $P(x, k) = x'$. (Of course, $N(x', q)$ can be non-empty only if $x' \in H_{2^{-q}}$.) Now let representative set $C(x, k)$ include all net points in H_k satisfying $2 \cdot 2^k \leq \rho(x, y) < 4 \cdot 2^k$.

Instead of computing the mean slope averaged over all points, we will record for each hierarchical point $x_j \in H_k$ the maximum and minimum labels of points in its neighborhood ($z'_{j,k}, z''_{j,k}$, respectively), and for each point $x_i \in S$, compare its label to the maximum and minimum among the neighborhoods of the points of representative set $C(x, k)$ for all k . For any point pair $x, x' \in S$, the triangle inequality implies that for level k satisfying $2 \cdot 2^k \leq \rho(x, x') < 4 \cdot 2^k$ we have $2^k \leq \rho(x, C(x', k)) < 5 \cdot 2^k$, and so the average slope is preserved up to constant factors.

$$\begin{aligned} \text{Minimize} \quad & W = \sum_{i \in [n]} w_i \\ \text{subject to} \quad & \frac{1}{n} \sum_{i \in [n]} L_i \leq L \\ & w_i \geq |z_i - y_i| \quad \forall i \in [n] \\ & \max\{|z_i - z'_{j,k}|, |z_i - z''_{j,k}|\} \leq L_i \cdot \rho(x_i, x_j) \quad \forall i \in [n], k \in [\lceil \log \Delta \rceil], x_j \in C(x_i, k) \\ & z''_{i,k} \leq z_j \leq z'_{i,k} \quad \forall i \in [n], k \in [\lceil \log \Delta \rceil], x_j \in N(x_i, k) \\ & 0 \leq z_i, z'_i, z''_i \leq 1 \quad \forall i \in [n]. \end{aligned}$$

By the packing property (1), each point of S can be found in at most $2^{O(\text{ddim})}$ neighborhoods of each level, so that the sum of sizes all all neighborhoods is $2^{O(\text{ddim})} n \log \Delta$. Similarly, $|C(x)| = 2^{O(\text{ddim})} \log \Delta$, and so the sum of sizes of all representative sets is $2^{O(\text{ddim})} n \log \Delta$. It follows that the program has $2^{O(\text{ddim})} n \log \Delta$ constraints, each with only a constant number of non-zero variables. As before, the program can be adapted to the framework of Koufogiannakis and Young (2014) by separating the max term into two separate constraints, and introducing dummy variables $\tilde{z}_i, \tilde{z}'_i, \tilde{z}''_i$ respectively satisfying $z_i + \tilde{z}_i = 1, z'_i + \tilde{z}'_i = 1$ and $z''_i + \tilde{z}''_i = 1$. The claimed runtime follows. \blacksquare

Extension to the weak mean. In light of Corollary 26, relaxing the constraint in (36) from $f \in \overline{\text{Lip}}_L(x_{[n]}, \rho, \mu_n)$ to $f \in \widetilde{\text{Lip}}_L(x_{[n]}, \rho, \mu_n)$ will yield an improvement in the objective function that can also be achieved via the relaxation $f \in \overline{\text{Lip}}_{2L \log n}(x_{[n]}, \rho, \mu_n)$, and hence we forgo designing a specialized algorithm for this case.

5.2 Classification

We show below that the sample smoothing problem for classification under average slope constraints in the strong-mean sense admits an algorithmic solution, but this solution reduces to solving an NP-hard problem. (This does not necessarily imply however that the smoothing problem in the strong-mean sense is NP-hard.) Fortunately, we are able to produce an efficient bi-criteria approximation algorithm for the sample smoothing problem under average slope constraints in the *weak*-mean sense. Given our current state of knowledge, the weak mean provides us an unexpected computational advantage over the strong mean, in addition to its being a more refined indicator of average smoothness.

Smoothing under the strong mean. Let ν be some distribution on $\Omega \times \{0, 1\}$, and $S = (X_i, Y_i)_{i \in [n]} \sim \nu^n$ be a set sampled i.i.d. from ν . At constant confidence level δ , the generalization bound (70) implies that any $f : \Omega \rightarrow [0, 1]$ with $\overline{\Lambda}_f(\mu, \Omega) \leq L$ that makes $k = \sum_{i=1}^n \mathbf{1}[f(X_i) \neq Y_i]$ or fewer mistakes on the sample will achieve, with high probability, a generalization error

$$\mathbb{P}_{(X, Y) \sim \nu} (\mathbf{1}[f(X) > 1/2] \neq Y) \leq \frac{k}{n} + G(L, n) =: Q(f, L), \quad (37)$$

where $G(\cdot, \cdot)$ is the bound in the right-hand side of (70).

We wish to find a hypothesis approximately minimizing the bound $Q(\cdot, \cdot)$ in (37). An intuitive approach might involve solving the following problem, which we call the Minimum Removal Average Slope Problem: Given an average slope target value L , remove the smallest number points from S so that the resulting point set attains average slope at most L . Clearly, an algorithm solving or approximating the Minimum Removal Average Slope Problem can be leveraged to find a minimizer for (37). However, we can show that such an approach is algorithmically infeasible:

Claim 10 *The Minimum Removal Average Slope Problem is NP-hard. Assuming the Exponential Time Hypothesis (ETH), it is hard to approximation within a factor $n^{1/\log^r \log n}$ for some universal constant r .*

Proof The hardness follows via a reduction from the Minimum k -Union Problem. In this problem we are given a collection C of n sets and a parameter k , and must find a subset $C' \subset C$ of size $|C'| = k$ so that the union of all sets in C' is minimized. The Minimum k -Union Problem is known to be NP-hard (Chlamtáč et al., 2016), and under the ETH, it is hard to approximate the minimum union within a factor of $n^{1/\log^{r'} \log n}$ for some universal constant r' . (The hardness of approximation follows directly from the Densest k -Subgraph Problem, which can be viewed as a special case of Minimum k -Union Problem (Manurangsi, 2017; Chlamtáč, 2020).)

The reduction is as follows: Given an instance (C, k) of Minimum k -Union, we create an instance of Minimum Removal Average Slope. Create bipartite point set $S = S_e \cup S_s$ thus: Set $S_e \in S$ has a point corresponding to each element in the element-universe of C . Set $S_s \in S$ has a point of weight $m = |S_e| + 1$ corresponding to each set in C . (A point can be assigned weight m by placing m copies of the same point in S_e .) For each point pair $s \in S_s, e \in S_e$ we set $\rho(s, e)$ equal to 2 if $e \in s$, and 1 otherwise. Now let the target average slope be $\frac{2|S|-km}{|S|}$. Clearly, this can only be attained by deleting the minimum number of points in S_e so that at least k points of S_s are not within distance 1 of any point of S_e . This is equivalent to finding k sets of C of minimum union. The reduction preserves hardness-of-approximation as well. \blacksquare

One attempt around the hardness result would be to mimic the approach taken for regression: Identify a target error term $\varepsilon \in [\frac{1}{n}, 1]$ (via binary search), and remove from S the “worst” $2\varepsilon n$ points in order to minimize the *maximum* slope of the remaining points. This in turn may be approximated using the algorithm of Gottlieb et al. (2014), which runs in time $2^{O(\text{ddim})} n \log n + \text{ddim}^{O(\text{ddim})} n$. Such an approach would yield a classifier achieving a value of $Q(\cdot, \cdot)$ within a factor of $(1/\varepsilon)^{O(\text{ddim})}$ of the optimal one. A better approximation factor of $2^{O(\text{ddim})} \log n$ is feasible, however, as we shall see below.

Theorem 11 (Training and generalization for classification, weak mean) *Let ν be some distribution on $\Omega \times \{0, 1\}$, and $S_n = (X_i, Y_i)_{i \in [n]} \sim \nu^n$ be a set sampled i.i.d. from ν . Denote by $\hat{f} \in \widetilde{\text{Lip}}_L(\Omega_n, \rho, \mu_n)$ the minimizer of $R(\cdot; \nu_n)$. Then there is an efficient learning algorithm \mathcal{A} , which constructs a classifier $f = \mathcal{A}(S_n)$ such that for any given $L > 0$ and $0 < \delta < 1$:*

(a) *With probability at least $1 - 2^{O(\text{ddim})} \log^3(n)/n - 3\delta$,*

$$R(f; \nu) \leq 2^{O(\text{ddim})} \log(n) R(\hat{f}; \nu_n) + \frac{C_\delta \sqrt{2^{O(\text{ddim})} \log^3(n) L}}{n^{1/8d}} + \frac{C_\delta^{-d/2} \sqrt{2}}{n^{5/16}} + 3 \sqrt{\frac{\log(2/\delta)}{2n}},$$

(b) *$f(x)$ can be evaluated at each $x \in \Omega$ in time $O(n^2)$ after a one-time “smoothing” computation of $O(n^2) + 2^{O(\text{ddim})} n \log^2 n$.*

where C_δ is a constant depending only on δ .

Proof The bi-criteria approximation algorithm in Lemma 12 yields an $\tilde{f} \in \widetilde{\text{Lip}}_{O(1)L}(\Omega_n, \rho, \mu_n)$ whose empirical risk is within a $2^{O(\text{ddim})} \log(n)$ factor of the optimal $R(\hat{f}; \nu_n)$ and the adversarial extension procedure in Lemma 17 for classification guarantees that the PMSE extension of \tilde{f} from Ω_n to Ω satisfies $f \in \widetilde{\text{Lip}}_{2^{O(\text{ddim})} \log^3(n)L}(\Omega, \rho, \mu)$ with high probability. The generalization bound in (70) then applies directly to yield (a). The runtimes claimed in (b) are demonstrated in Remark 21 (which argues that PMSE can be evaluated in time $O(n^2)$) and the proof of Lemma 12. \blacksquare

Bi-criteria approximation for smoothing under weak mean. We wish to perform smoothing of $\tilde{\Lambda}_f(\mu_n, \Omega_n)$. For this, we define the *continuous local slope removal problem* (CLSRP) as follows: Let (Ω, ρ) be a metric space with $\text{diam}(\Omega) \leq 1$ and $\text{ddim}(\Omega) < \infty$. Given $S = (x_{[n]}, y_{[n]}) \in (\Omega^n, \{0, 1\}^n)$ and $L > 0$, relabel the minimal amount of points in S with any real label in $[0, 1]$, so that for the resulting label-set the number of points with local slope tL or greater is at most n/t for all real $t \in [1, n+1]$. Notice that solving CLSRP for a given L implies that $\tilde{\Lambda}_f(\mu_n, \Omega_n) \leq L$. By definition of CLSRP, $k\mu_n(M_f(k)) \leq L$ for $L \leq k \leq (n+1)L$, while this extends trivially for all $k \leq L$ and $k \geq (n+1)L$.

Suppose that the solution I of CLSRP consists of relabeling $k > 0$ points in S . Then an (a, b) -bicriteria approximation $(a, b \geq 1)$ to the solution of CLSRP on I is one in which at most ak points are relabeled, while the number of points with local slope btL or greater is at most n/t for all real $t \in [1, n+1]$. We can show the following:

Lemma 12 *CLSRP admits a $(2^{O(\text{ddim})} \log n, O(1))$ -bicriteria approximation in time $O(n^2) + 2^{O(\text{ddim})} n \log^2 n$.*

Proof The construction is as follows. Let $t_i = 2^i$ for integer $i \in [0, \lceil \log n \rceil]$. For each t_i we will construct a points set P_i of points to be relabelled, and the final solution will be $P = \cup_i P_i$:

For each i , construct a $\frac{1}{2t_i L}$ -net of S called $T_i \subset S$. Associate every point in S with its nearest neighbor in T_i , and let the neighborhood of $p \in T_i$ ($N(p)$) include all points of S associated with p . Now, if not all points in $N(p)$ have the same sign, then create a new point $p' \in T_i$ which is a copy of p but with label $1 - l(p)$. Remove from $N(p)$ all points with label $1 - l(p)$, and place them in $N(p')$ instead. (Note that p is found in $N(p)$, but p' is not found in $N(p')$.) This can all be done in time $2^{O(\text{ddim})} n \log n$ using a standard point hierarchy.

Now create a new subset $T'_i \subset T_i$ thus: $p \in T_i$ is added to T'_i only if there is some point $q \in T_i$ with $l(p) \neq l(q)$ satisfying $d(p, q) \leq \frac{2}{t_i L}$. We will show below that (roughly speaking) for $p \in T'_i$ all points in $N(p)$ have high local slope constant, while for $p \in T_i \setminus T'_i$ all points in $N(p)$ have low local slope constant. We will associate a weight with each $p \in T'_i$ thus: For all $p \in T'_i$, let $T_0(p), T_1(p)$ consist of all points of T'_i within distance $\frac{2}{t_i L}$ of p , and with respective labels 0,1. Define $S_0(p) = \cup_{q \in T_0(p)} N(q)$ and $S_1(p) = \cup_{q \in T_1(p)} N(q)$. With each point $p \in T$ we associate weight

$$w(p) = \min\{|S_0(p)|, |S_1(p)|\}.$$

Intuitively, this weight reflects the cost of reducing the local slope constant of all points in $N(p)$; this requires relabeling all points in either $S_0(p)$ or $S_1(p)$.

Let $m = \sum_{p \in T'_i} |N(p)|$. Let $m' = m - \frac{6n}{t_i L}$ (a value which we will motivate below), and we wish to find a minimal weight subset $C_i^* \subset T'_i$ satisfying that $\sum_{p \in C_i^*} |N(p)| \geq m'$. This is a version of the NP-hard Minimum Knapsack Problem, but we can find in time $O(n \log n)$ a subset $C_i \subset T'_i$ satisfying that $\sum_{p \in C_i} |N(p)| \geq m'$, with $w(C_i) \leq 2w(C_i^*)$ (Csirik et al., 1991). Then the set of points to be relabeled is $P_i = \cup_{p \in C_i} N(p)$, and as above the final solution is $P = \cup_i P_i$. This completes the construction.

To prove correctness, first fix some t_i . Consider a pair $p, q \in S$ which are found in the neighborhoods of respective points $p', q' \in T$. If $d(p, q) \leq \frac{1}{t_i L}$, then by the triangle inequality

$$\rho(p', q') \leq \rho(p, q) + \rho(p, p') + \rho(q, q') \leq \frac{1}{t_i L} + \frac{1}{2t_i L} + \frac{1}{2t_i L} = \frac{2}{t_i L}.$$

It follows that if $\rho(p', q') > \frac{2}{t_i L}$ then $\rho(p, q) > \frac{1}{t_i L}$. Also, if $\rho(p', q') \leq \frac{2}{t_i L}$ then

$$\rho(p, q) \leq \rho(p', q') + \rho(p, p') + \rho(q, q') \leq \frac{2}{t_i L} + \frac{1}{2t_i L} + \frac{1}{2t_i L} = \frac{3}{t_i L}.$$

For the approximation bound on the number of relabelled points: Consider some $p \in T'_i$. By the above calculation, if at least one point in each of $S_0(p)$ and $S_1(p)$ is not chosen for relabelling, then the above bound along with Corollary 24 imply that no point of $N(p)$ can attain local slope constant less than $\frac{1}{\frac{3}{t_i L} + \frac{3}{t_i L}} = \frac{t_i L}{6}$. Now, as $N(p) \subset S_0(p) \cup S_1(p)$, and since the exact solution to CLSRP has slope constant $\frac{t_i L}{6}$ or greater on at most $\frac{6n}{t_i L}$ points, the exact solution must relabel all but at most $\frac{6n}{t_i L}$ points of $\cup_{p \in T'_i} N(p)$. And further, for any point p relabelled by the exact solution to achieve local slope constant $\frac{t_i L}{6}$ or less, the exact solution must also relabel all of either $S_0(p)$ or $S_1(p)$. By the packing property (1), for any t_i , any point $q \in S$ appears in $2^{O(\text{ddim})}$ sets of the form $S_0(p), S_1(p)$, and so it follows that the number of points relabelled by the above construction for t_i is within a factor $2^{O(\text{ddim})}$ of the number of points relabelled by the exact solution. Summing over $O(\log n)$ values of i , we have that the number of points relabelled by the approximation algorithm is within a factor $2^{O(\text{ddim})} \log n$ of the exact solution.

For the bound on the local slope constant: Fix some t_i . For any $p \in T_i$, if all points in $S_0(p)$ or $S_1(p)$ are relabelled according to PMSE, then as shown above the distance from any point $p' \in N(p)$ to a point of label 0 or 1 (respectively) is greater than $\frac{2}{t_i L}$. Then by Corollary 24(ii), the local slope constant of p will be at most $\frac{t_i L}{2}$. By construction, at most $\frac{6n}{t_i L}$ points remain with this slope constant, and the result follows. \blacksquare

6. Adversarial extension: regression

As we discussed in Section 2 (and, in greater detail, at the beginning of Section 5), ensuring that a function with on-average smooth behavior on the sample also possesses this property on the whole space is non-trivial. To this end, we introduce the following **adversarial extension** game. First, $\Omega_n = X_{[n]} \sim \mu^n$ is drawn, which induces the usual empirical measure μ_n . Next, the adversary picks an $\varepsilon > 0$ and $y : \Omega_n \rightarrow [0, 1]$ arbitrarily. Finally, the learner is challenged to construct a function $f : \Omega \rightarrow [0, 1]$ satisfying the following criteria:

- (a) f is close to y on the sample
 - (i) (w.r.t. strong average slope): $\|f - y\|_{L_1(\mu_n)} \leq O(\varepsilon)(1 \vee \bar{\Lambda}_y(\mu_n, \Omega_n))$,
 - (ii) (w.r.t. weak average slope): $\|f - y\|_{L_1(\mu_n)} \leq \tilde{O}(\varepsilon)(1 \vee \tilde{\Lambda}_y(\mu_n, \Omega_n))$;
- (b) f 's average slope does not much exceed the sample one

- (i) (w.r.t. strong average): $\bar{\Lambda}_f(\mu, \Omega) \leq O(1)\bar{\Lambda}_y(\mu_n, \Omega_n)$,
- (ii) (w.r.t. weak average): made precise in Lemma 31.

Two immediate observations are in order. First, we notice the tension between the criteria (a) and (b). Each one is trivial to satisfy individually — (b) by a constant function and (a) by *any* proper extension, including PMSE — but it is not obvious that both can be satisfied simultaneously. Second, (b) can at best hold with high probability. Indeed, let μ be a distribution over $\Omega = [0, 1]$ with high density near $1/2$ and low density at the endpoints. Suppose further that the sample Ω_n has turned out rather unrepresentative: many points near the endpoints and only two near $1/2$. In such a setting, the adversary can force, say, $\bar{\Lambda}_f(\mu, \Omega)/\bar{\Lambda}_y(\mu_n, \Omega_n)$ to be large for *any* extension f of y .

Throughout this section, we assume $\text{diam}(\Omega) \leq 1$ and $d := \text{ddim}(\Omega) < \infty$.

6.1 Proving (a.i) and (b.i), strong mean

We begin by handling the technically simpler case of addressing the adversarial extension problem for the strong mean.

Lemma 13 (Adversarial extension for strong mean) *In the adversarial extension game, there is an efficient algorithm for satisfying conditions (a.i) and (b.i), the latter with probability at least*

$$1 - \exp\left(-n(\varepsilon/8)^{\text{ddim}+1} + d \log((8/\varepsilon))\right),$$

where $\varepsilon < 1$ and $n \geq (8/\varepsilon)^{\text{ddim}+2}$. (For concrete constants, $O(\varepsilon)$ in (a.i) may be replaced by 3ε and $O(1)$ in (b.i) by 5.)

Proof We begin by constructing the extension f .

1. Sort the $\Lambda_y(X_i, \Omega_n)$ in decreasing order, and let $\Omega_n(\varepsilon) \subset \Omega_n$ consist of the $\lfloor \varepsilon n \rfloor$ points with the largest values (breaking ties arbitrarily).
2. Put $\Omega'_n(\varepsilon) := \Omega_n \setminus \Omega_n(\varepsilon)$.
3. Let V be an ε -net of $\Omega'_n(\varepsilon)$.
4. Define $f : \Omega \rightarrow \mathbb{R}$ as the PMSE extension of y from V to Ω , as defined in Definition 18.

Since $|\Omega'_n(\varepsilon)| < n$, the value of $f(x)$ can be computed in time $O(n^2)$ at any given $x \in \Omega$. The computational runtime for net construction is $\min\{2^{O(\text{ddim})}n \log(n) \log \Delta, O(n^2)\}$ (Krauthgamer and Lee, 2004).

Proof of (a.i). Recalling that $f \equiv y$ on V , we have

$$\begin{aligned} \|f - y\|_{L_1(\mu_n)} &= \frac{1}{n} \sum_{x \in \Omega_n \setminus V} |f(x) - y(x)| \\ &= \frac{1}{n} \sum_{x \in \Omega_n(\varepsilon) \setminus V} |f(x) - y(x)| + \frac{1}{n} \sum_{x \in \Omega'_n(\varepsilon) \setminus V} |f(x) - y(x)|. \end{aligned} \quad (38)$$

Since $0 \leq f, y \leq 1$, the first term is trivially bounded by $|\Omega_n(\varepsilon)|/n \leq \varepsilon$. To bound the second term, we recall the map $\varphi_V : \Omega_n(\varepsilon) \rightarrow V$ defined in Section 1.3 — and in particular, that $\rho(x, \varphi_V(x)) \leq \varepsilon$ — and compute

$$\begin{aligned}
 \frac{1}{n} \sum_{x \in \Omega'_n(\varepsilon) \setminus V} |f(x) - y(x)| &\leq \frac{1}{n} \sum_{x \in \Omega'_n(\varepsilon) \setminus V} \frac{\varepsilon}{\rho(x, \varphi_V(x))} |f(x) - y(x)| \\
 &\leq \frac{\varepsilon}{n} \sum_{x \in \Omega'_n(\varepsilon) \setminus V} \frac{|y(\varphi_V(x)) - y(x)| + |f(x) - y(\varphi_V(x))|}{\rho(x, \varphi_V(x))} \\
 &\leq \frac{\varepsilon}{n} \cdot 2 \sum_{x \in \Omega'_n(\varepsilon) \setminus V} \Lambda_y(x, \Omega_n) \\
 &\leq 2\varepsilon \bar{\Lambda}_y(\mu_n, \Omega_n),
 \end{aligned} \tag{39}$$

where the third inequality follows from Theorem 23:

$$\frac{|f(x) - y(\varphi_V(x))|}{\rho(x, \varphi_V(x))} = \frac{|f(x) - f(\varphi_V(x))|}{\rho(x, \varphi_V(x))} \leq \Lambda_f(x, V) \leq \Lambda_y(x, V) \leq \Lambda_y(x, \Omega_n).$$

Hence f satisfies (a.i) with 3ε .

Proof of (b.i). Let $q > 0$ be a parameter to be determined later. Let $U \subseteq \Omega$ be an $\varepsilon/4$ -net of Ω , with induced Voronoi partition $\Pi = \{\varphi_U^{-1}(u) : u \in U\}$. Put $m = |\Pi| \leq (8/\varepsilon)^{\text{ddim}}$ (see (1)) and segregate the elements of Π into “light,” Π_0 , and “heavy,” Π_1 :

$$\Pi_0 := \{B \in \Pi : \mu_n(B) < nq/m\}, \quad \Pi_1 := \{B \in \Pi : \mu_n(B) \geq nq/m\}.$$

We will need three auxiliary lemmata (whose proof is deferred to the Appendix).

Lemma 14 (Local slope smoothness of the PMSE) *Suppose that $A \subseteq \Omega$ and f is the PMSE of some function from A to Ω . Suppose further that $E \subseteq \Omega$ satisfies*

$$\text{diam}(E) \leq \frac{1}{2} \min_{x \neq x' \in A} \rho(x, x'). \tag{40}$$

Then

$$\sup_{x, x' \in E} \frac{\Lambda_f(x, \Omega)}{\Lambda_f(x', \Omega)} \leq 2.$$

Lemma 15 (Accuracy of empirical measure) *The individual heavy cells have empirical measure within a constant factor of their true measure, with high probability:*

$$\mathbb{P} \left(\min_{B \in \Pi_1} \frac{\mu_n(B)}{\mu(B)} \leq \frac{1}{2} \right) \leq m \exp(-nq/4m), \tag{41}$$

$$\mathbb{P} \left(\max_{B \in \Pi_1} \frac{\mu_n(B)}{\mu(B)} \geq 2 \right) \leq m \exp(-nq/3m). \tag{42}$$

Additionally, the combined μ -mass of the light cells is not too large:

$$\mathbb{P} \left(\sum_{B \in \Pi_0} \mu(B) \geq 2q \right) \leq \exp \left[-(m + nq^2)/2 + q\sqrt{mn} \right], \quad nq^2 \geq m. \tag{43}$$

Finally, we bound the Lipschitz constant of the PMSE f of y :

Lemma 16 (Lipschitz constant of f)

$$\|f\|_{\text{Lip}} \leq 2\varepsilon^{-1}\tilde{\Lambda}_y(\mu_n, \Omega_n) \leq 2\varepsilon^{-1}\bar{\Lambda}_y(\mu_n, \Omega_n).$$

Armed with these results, we are in a position to prove (b.i). We choose $q := \varepsilon/8$ and calculate

$$\begin{aligned} \bar{\Lambda}_f(\mu, \Omega) &= \int_{\Omega} \Lambda_f(x, \Omega) d\mu = \sum_{B \in \Pi} \int_B \Lambda_f(x, \Omega) d\mu \\ &= \sum_{B \in \Pi_0} \int_B \Lambda_f(x, \Omega) d\mu + \sum_{B \in \Pi_1} \int_B \Lambda_f(x, \Omega) d\mu. \end{aligned} \quad (44)$$

The first term in (44) is bounded using (43) in Lemma 15 and Lemma 16

$$\begin{aligned} \sum_{B \in \Pi_0} \int_B \Lambda_f(x, \Omega) d\mu &\leq \sum_{B \in \Pi_0} \int_B \frac{2\bar{\Lambda}_y(\mu_n, \Omega_n)}{\varepsilon} d\mu \\ &= \frac{2\bar{\Lambda}_y(\mu_n, \Omega_n)}{\varepsilon} \sum_{B \in \Pi_0} \mu(B) \\ &\leq_p \frac{2\bar{\Lambda}_y(\mu_n, \Omega_n)}{\varepsilon} \left(2 \cdot \frac{\varepsilon}{8}\right) \leq \bar{\Lambda}_y(\mu_n, \Omega_n), \end{aligned}$$

where the inequality indicated by \leq_p holds with high probability, as in (43).

To bound the second term in (44), we first observe that for all $B \in \Pi$ and $x' \in B$,

$$\Lambda_f(x', \Omega) \leq 2 \min_{x \in \Omega_n \cap B} \Lambda_f(x, \Omega).$$

Indeed, this follows from Lemma 14, invoked with $A = V$ and $E = \Omega_n \cap B$. Proceeding,

$$\begin{aligned} \sum_{B \in \Pi_1} \int_B \Lambda_f(x, \Omega) d\mu &\leq \sum_{B \in \Pi_1} \int_B 2 \min_{x \in \Omega_n \cap B} \Lambda_f(x, \Omega) d\mu \\ &= \sum_{B \in \Pi_1} 2 \min_{x \in \Omega_n \cap B} \Lambda_f(x, \Omega) \mu(B) \\ &\leq_p 4 \sum_{B \in \Pi_1} \min_{x \in \Omega_n \cap B} \Lambda_f(x, \Omega) \mu_n(B) \\ &= \frac{4}{n} \sum_{B \in \Pi_1} \sum_{x' \in \Omega_n \cap B} \min_{x \in \Omega_n \cap B} \Lambda_f(x, \Omega) \\ &\leq \frac{4}{n} \sum_{B \in \Pi_1} \sum_{x' \in \Omega_n \cap B} \Lambda_f(x', \Omega) \\ &\leq \frac{4}{n} \sum_{x' \in \Omega_n} \Lambda_f(x', \Omega) \leq 4\bar{\Lambda}_y(\mu_n, \Omega_n), \end{aligned}$$

where the inequality indicated by \leq_p holds with high probability, as in Lemma 15, and the last inequality follows from Corollary 24(i), since $\Lambda_f(x', \Omega)$ is determined by f 's values on V :

$$\Lambda_f(x', \Omega) = \Lambda_f(x', V) = \Lambda_f(x', \Omega_n) \leq \Lambda_y(x', \Omega_n).$$

Plugging these estimates back into (44) yields (b.i) with 5 in place of $O(1)$. Applying the union bound to the probabilistic inequalities marked by \leq_p above yields the claim with probability at least $1 - \delta$, where

$$\delta = \left[\exp\left(-\frac{m + nq^2}{2} - q\sqrt{mn}\right) + m \exp\left(-\frac{nq}{4m}\right) \right] \leq \exp\left(-n(\varepsilon/8)^{\text{ddim}+1} + d \log((8/\varepsilon))\right).$$

■

6.2 Proving (a.ii) and (b.ii), weak mean

As discussed at the end of Section 5.1, our training procedure for regression obtains comparable results for both strong- and weak-mean regularization. Hence, only the former is fleshed out, and the latter, corresponding to claims (a.ii) and (b.ii) above, is not directly invoked in this paper. We find the proof of (a.ii) and (b.ii) to be of independent interest, and present it in Section E.

7. Adversarial extension: classification

The adversarial extension for classification differs in several aspects from its regression analogue in Section 6. Conceptually, there is a “type mismatch” between a $[0, 1]$ -valued function and the $\{0, 1\}$ labels it is supposed to predict. The actual prediction is performed by rounding via $f(x) \mapsto \mathbf{1}[f(x) > 1/2]$, but the sample risk charges a unit loss for every $f(x_i) \neq y_i$, regardless of how close the two might be.⁴ Thus, for adversarial extension, no distortion of the adversary’s labels y is allowed — the only changes the learner makes to the sample labels occur during the smoothing procedure in Section 5.2 — and hence there is no ε parameter. The strict adherence to y incurs the cost of a $2^{O(\text{ddim})}$ polylog n increase in the average slope of the extension (unlike in regression, where a small distortion of y afforded an at most constant increase). Finally, note that though intermediate results for the strong mean are obtained, only those for the weak mean are algorithmically useful, in light of the hardness result in Section 5.2.

Lemma 17 (Adversarial extension for classification) *Suppose that $\Omega_n = X_{[n]} \sim \mu^n$ and $y : \Omega_n \rightarrow \{0, 1\}$ satisfies $\tilde{\Lambda}_y(\mu_n, \Omega_n) \leq n$, but is otherwise arbitrary. Then there is an efficient algorithm for computing a function $f : \Omega \rightarrow [0, 1]$ that coincides with y on Ω_n and*

4. A simple no-free-lunch argument shows that one could not hope to obtain a generalization bound with sample risk based on the $\mathbf{1}[f(x_i) > 1/2] \neq y_i$ loss. Indeed, the function $f(x_i) = 1/2 + \varepsilon y(x_i)$ would achieve zero empirical risk while also having an arbitrarily small Lipschitz constant.

satisfies

$$\begin{aligned}
 \tilde{\Lambda}_f(\mu, \Omega) &\leq \bar{\Lambda}_f(\mu, \Omega) \\
 &\leq 2^{O(\text{ddim})} \log^2(n) \bar{\Lambda}_y(\mu_n, \Omega_n) \\
 &\leq 2^{O(\text{ddim})} \log^3(n) \tilde{\Lambda}_y(\mu_n, \Omega_n).
 \end{aligned} \tag{45}$$

with probability at least

$$1 - 2^{O(\text{ddim})} \log^3(n)/n.$$

Remark. The assumption $\tilde{\Lambda}_y(\mu_n, \Omega_n) \leq n$ incurs no loss of generality because $\tilde{\Lambda}_f(\mu, \Omega) > n$ yields vacuous generalization bounds.

Proof Only the estimate in (45) requires proof; the rest hold everywhere (and in particular, with probability 1) by (11) and Corollary 26, respectively. Our algorithm computes f as the PMSE extension of y from Ω_n to all of Ω . It remains to show that (45) holds with the claimed probability. Throughout the proof, $d := \text{ddim}(\Omega)$.

Let $\{H_{2^{-i}}\}_{i=0}^{\lceil 2 \log 2n \rceil}$ be a point hierarchy for Ω and to each net-point $p \in H_r$ associate a ball $B(p, r)$. Note that the balls associated with points in the lowest level of the hierarchy have a radius of $r \leq \frac{1}{4n^2}$. Hence, any ball B associated to some point in the lowest level in the hierarchy contains y -homogeneously labeled points of Ω_n . That is, for any $x, x' \in B$ we have $y(x) = y(x')$, otherwise contradicting the assumption that $\tilde{\Lambda}_y(\mu_n, \Omega_n) \leq n$. Furthermore, the nearest opposite-label point is at least at a distance of $1/n^2$ from any point in B ; we will refer to this property as the *extended monochromatic property*. Denote by \mathcal{B}_j the set of balls corresponding to points of H_j (note that $j = 2^{-i}$ is typically not an integer).

For any $x \in \Omega$, denote by $p_x \in \underset{p \in \Omega_n}{\text{argmin}} \rho(p, x)$ a nearest neighbor of x within the set Ω_n . By property (ii) of corollary 24, we know that

$$\Lambda_f(x, \Omega) = \frac{|f(x) - f(p_x)|}{\rho(x, p_x)} \leq \Lambda_f(p_x, \Omega_n). \tag{46}$$

Let $i(x)$ be the minimum between the following: the smallest power of $1/2$ such that $(1/2)^{i(x)} \leq \rho(x, p_x)/4$ and $\lceil 2 \log 2n \rceil$. Notice that in either case, the inequality

$$\rho(x, p_x)/8 \leq (1/2)^{i(x)} \tag{47}$$

holds. Let $B(x) \in \mathcal{B}_j$ be the ball of radius $(1/2)^{i(x)}$ covering x . By the triangle inequality and the extended monochromatic property of the lowest level balls, for any $x' \in B(x)$ we have that

$$\frac{1}{2} \Lambda_f(x, \Omega) \leq \Lambda_f(x', \Omega) \leq 2 \Lambda_f(x, \Omega). \tag{48}$$

Now let $\Omega(i) \subseteq \Omega$ consist of all points $x \in \Omega$ for which $B(x)$ has radius $(1/2)^i$, and let

$$\tilde{\mathcal{B}}_{(1/2)^i} = \{B(x) \in \mathcal{B}_{(1/2)^i} : x \in \Omega(i)\}$$

and $\tilde{\mathcal{B}} = \cup_j \tilde{\mathcal{B}}_j$. We claim that $|\tilde{\mathcal{B}}| \leq 2^{O(\text{ddim})} n \log n$. Indeed, for any $p \in \Omega_n$, let $N(p, i) \subseteq \Omega(i)$ include every point $x \in \Omega(i)$ for which $p_x = p$, and define

$$\tilde{\mathcal{B}}_{(1/2)^i}(p) = \{B(x) \in \tilde{\mathcal{B}}_{(1/2)^i} : x \in N(p, i)\}$$

and $\tilde{\mathcal{B}}(p) = \cup_j \tilde{\mathcal{B}}_j(p)$. Since balls in $\tilde{\mathcal{B}}_{(1/2)^i}(p)$ have radius $(1/2)^i$ and their centers are within distance $8(1/2)^i$ of p , the packing property (1) gives that $|\tilde{\mathcal{B}}_{(1/2)^i}(p)| = 2^{O(\text{ddim})}$ and so $|\tilde{\mathcal{B}}(p)| = 2^{O(\text{ddim})} \log n$. It follows that $|\tilde{\mathcal{B}}| \leq \sum_{p \in \Omega_n} |\tilde{\mathcal{B}}_j(p)| = 2^{O(\text{ddim})} n \log n$.

We would like to claim that the empirical measure of each $B \in \tilde{\mathcal{B}}$ is close to its true measure. Some care must be taken here, since $\tilde{\mathcal{B}}$ is itself a random set, determined by the same sample that determines the empirical measure. However, conditional on any given $p \in \Omega_n$ — which determines the set $\tilde{\mathcal{B}}(p)$ — we can use the remaining $n - 1$ sample points to estimate the mass of each $B \in \tilde{\mathcal{B}}(p)$. To avoid the notational nuisance of distinguishing fractions involving n and $n - 1$, we will write $a \approx b$ to mean $a = (1 + \Theta(1/n))b$ and $a \lesssim b$ to mean $a = (1 + O(1/n))b$ for the remainder of the proof.

For any $B \in \tilde{\mathcal{B}}$, let \mathcal{E}_B be the event that

$$\mu(B) \gtrsim \frac{2 \log n}{n} \implies \mu(B) \lesssim \mu_n(B) \log n$$

(in words: if B is sufficiently “heavy” then it is not under-sampled). It follows from Theorem 32 that $\mathbb{P}(\mathcal{E}_B) \geq 1 - O\left(\frac{\log^2 n}{n^2}\right)$ holds for each $B \in \tilde{\mathcal{B}}$. We argued above that $|\tilde{\mathcal{B}}(p)| = 2^{O(\text{ddim})} \log n$, whence the event $\mathcal{E}(p) := \cap_{B \in \tilde{\mathcal{B}}(p)} \mathcal{E}_B$, conditional on the given $p \in \Omega_n$, occurs with probability at least $1 - \frac{2^{O(\text{ddim})} \log^3 n}{n^2}$. Finally, taking a union bound over the n draws of Ω_n , we have that

$$\mathcal{E} := \bigcap_{p \in \Omega_n} \mathcal{E}(p)$$

holds with probability at least $1 - \frac{2^{O(\text{ddim})} \log^3 n}{n}$. The remainder of the proof proceeds conditionally on the high-probability event \mathcal{E} .

Let $\tilde{\mathcal{C}}_0$ be the set of $B \in \tilde{\mathcal{B}}$ such that $\mu_n(B) = 0$ and $\tilde{\mathcal{C}}_1 = \tilde{\mathcal{B}} \setminus \tilde{\mathcal{C}}_0$. Then

$$\int_{\Omega} \Lambda_f(x, \Omega) d\mu \leq \sum_{B \in \tilde{\mathcal{C}}_0} \int_B \Lambda_f(x, \Omega) d\mu + \sum_{B \in \tilde{\mathcal{C}}_1} \int_B \Lambda_f(x, \Omega) d\mu. \quad (49)$$

We begin by bounding the first term in (49). Since $\mu_n(B) = 0$ for each $B \in \tilde{\mathcal{C}}_0$ and we are assuming event \mathcal{E} , it follows that each of these balls must verify $\mu(B) \lesssim 2 \log(n)/n$.

Recalling the definition of $\tilde{\mathcal{B}}(p)$ for $p \in \Omega_n = \{x_1, \dots, x_n\}$,

$$\begin{aligned}
 \sum_{B \in \tilde{\mathcal{C}}_0} \int_B \Lambda_f(x, \Omega) d\mu &= \sum_{i \in [n]} \sum_{B \in \tilde{\mathcal{B}}(x_i) \cap \tilde{\mathcal{C}}_0} \int_B \Lambda_f(x, \Omega) d\mu \\
 &\leq \sum_{i \in [n]} \sum_{B \in \tilde{\mathcal{B}}(x_i) \cap \tilde{\mathcal{C}}_0} \int_B \Lambda_f(x_i, \Omega_n) d\mu \quad (\text{by (46)}) \\
 &\leq \sum_{i \in [n]} \sum_{B \in \tilde{\mathcal{B}}(x_i) \cap \tilde{\mathcal{C}}_0} \int_B \Lambda_y(x_i, \Omega_n) d\mu \quad (\text{pointwise optimality of PMSE}) \\
 &\lesssim \frac{2 \log n}{n} \sum_{i \in [n]} \sum_{B \in \tilde{\mathcal{B}}(x_i)} \Lambda_y(x_i, \Omega_n) \quad (\mu(B) \lesssim 2 \log(n)/n) \\
 &\leq 2^{O(\text{ddim})} \log^2 n \cdot \frac{1}{n} \sum_{i \in [n]} \Lambda_y(x_i, \Omega_n) \quad (|\tilde{\mathcal{B}}(x_i)| = 2^{O(\text{ddim})} \log n) \\
 &= 2^{O(\text{ddim})} \log^2(n) \bar{\Lambda}_y(\mu_n, \Omega_n).
 \end{aligned}$$

We now proceed to bound the second term in (49). To this end, we analyze two possibilities: either a $B \in \tilde{\mathcal{C}}_1$ is “light,” meaning that $\mu(B) \lesssim 2 \log(n)/n$, or else “heavy,” meaning that $\mu(B) \gtrsim 2 \log(n)/n$. Now $\mu_n(B) > 0 \implies \mu_n(B) \geq 1/n$, and so for any light ball we have, by construction, $\mu(B) \lesssim 2 \log(n) \mu_n(B)$. On the other hand, conditional on event \mathcal{E} , a heavy ball satisfies $\mu(B) \lesssim \log(n) \mu_n(B)$.

$$\begin{aligned}
 \sum_{B \in \tilde{\mathcal{C}}_1} \int_B \Lambda_f(x, \Omega) d\mu &\leq \sum_{B \in \tilde{\mathcal{C}}_1} \int_B 2 \min_{x' \in B \cap \Omega_n} \Lambda_f(x', \Omega) d\mu \quad (\text{by (46) and (48)}) \\
 &\lesssim 2 \log n \sum_{B \in \tilde{\mathcal{C}}_1} \mu_n(B) \min_{x' \in B \cap \Omega_n} \Lambda_f(x', \Omega_n) \\
 &= 2 \log n \sum_{B \in \tilde{\mathcal{C}}_1} \frac{1}{n} \sum_{x \in B \cap \Omega_n} \min_{x' \in B \cap \Omega_n} \Lambda_f(x', \Omega_n) \\
 &\leq 2 \log n \sum_{B \in \tilde{\mathcal{B}}} \frac{1}{n} \sum_{x \in B \cap \Omega_n} \Lambda_f(x, \Omega_n) \\
 &\leq 2^{O(\text{ddim})} \log^2(n) \frac{2}{n} \sum_{x \in \Omega_n} \Lambda_f(x, \Omega_n) \quad (\text{because } |\tilde{\mathcal{B}}| \leq 2^{O(\text{ddim})} \log n) \\
 &\leq 2^{O(\text{ddim})} \log^2(n) \frac{2}{n} \sum_{x \in \Omega_n} \Lambda_y(x, \Omega_n) \\
 &= 2^{O(\text{ddim})} \log^2(n) \bar{\Lambda}_y(\mu_n, \Omega_n).
 \end{aligned}$$

■

Appendix A. Miscellaneous inequalities and notations

Numerical inequalities, \vee, \wedge . We will use the following elementary facts: for all $a, b, c \in \mathbb{R}$, we have

$$(a + b + c)^2 \leq 3a^2 + 3b^2 + 3c^2 \quad (50)$$

and

$$|a - b| \vee |a - c| \geq \frac{1}{2}|b - c|; \quad (51)$$

for all $a, b, c, d \in \mathbb{R}_+$ such that $\frac{a}{c} \neq \frac{b}{d}$, we have

$$\frac{a + b}{c + d} > \frac{a}{c} \wedge \frac{b}{d}, \quad (52)$$

and for or all $f : \Omega \rightarrow \mathbb{R}$ and $x, y, z \in \Omega$ such that $f(x) \leq f(y) \leq f(z)$, we have

$$\frac{f(z) - f(x)}{\rho(z, x)} \geq \frac{f(y) - f(x)}{\rho(y, x)} \wedge \frac{f(z) - f(y)}{\rho(z, y)}, \quad (53)$$

where $s \vee t := \max\{s, t\}$ and $s \wedge t := \min\{s, t\}$. The floor $\lfloor \cdot \rfloor$ and ceiling $\lceil \cdot \rceil$ functions map a real number t to its closest integers below and above, respectively.

Bound on $\|\mu - \mu_n\|_1$. If μ is a probability measure with support size m and μ_n is its empirical realization (in the sense of (4)), then the following bound is well-known (see, e.g., Berend and Kontorovich (2013, Eqs. (5) and (17))):

$$\|\mu - \mu_n\|_1 \leq \sqrt{\frac{m}{n}} + \sqrt{\frac{2}{n} \log \frac{2}{\delta}}, \quad 0 < \delta < 1 \quad (54)$$

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

Order of magnitude. We use standard order-of-magnitude notation $f = O(g)$ to mean that $0 \leq f(\cdot) \leq cg(\cdot)$ for some universal $c > 0$. We write $f = \Theta(g)$ to indicate that both $f = O(g)$ and $g = O(f)$ hold. In the tilde notation ($\tilde{O}(\cdot)$, $\tilde{\Theta}(\cdot)$), logarithmic factors are ignored.

Appendix B. Pointwise Minimum Slope Extension

As mentioned in Related work, the material in this section turns out to have been largely anticipated by Oberman (2008) and is included here for self-containment and uniformity of notation and terminology. The term PMSE is ours, and the pointwise minimality property of this extension was not explicitly mentioned in Oberman (2008) (though is easily derivable from the results presented therein).

Let $f : \Omega \rightarrow \mathbb{R}$, $x \in \Omega$ and $\emptyset \neq A \subseteq \Omega$ be fixed (and hence frequently suppressed in the notation for readability). We assume for now that $|A| \geq 2$; the degenerate case $|A| = 1$ will

be handled below. For $u, v \in A$, define

$$\begin{aligned} R_x(u, v) &:= \frac{f(v) - f(u)}{\rho(x, v) + \rho(x, u)}, \\ F_x(u, v) &:= f(u) + R_x(u, v)\rho(x, u), \\ R_x^* &:= \sup_{u, v \in A} R_x(u, v), \\ W_x(\varepsilon) &:= \{(u, v) \in A^2 : R_x(u, v) > R_x^* - \varepsilon\}, \quad 0 < \varepsilon < R_x^* \\ \varphi_x(\varepsilon) &:= \{F_x(u, v) : (u, v) \in W_x(\varepsilon)\}. \end{aligned}$$

The assumption $|A| \geq 2$ implies that $R_x^* > 0$. Further, $\sup_{x \in \Omega} R_x^* < \infty$ if and only if $\|f|_A\|_{\text{Lip}} < \infty$.

Definition 18 (PMSE) For any metric space (Ω, ρ) , any $f : \Omega \rightarrow \mathbb{R}$ and $\emptyset \neq A \subseteq \Omega$ with $\|f|_A\|_{\text{Lip}} < \infty$ and $\text{diam}(A) \wedge \|f|_A\|_{\infty} < \infty$, define the **Pointwise Minimum Slope Extension (PMSE)** of f from A to Ω , denoted $f_A : \Omega \rightarrow \mathbb{R}$, by

$$f_A(x) := \limsup_{\varepsilon \rightarrow 0} (\Phi_x(\varepsilon)) = \liminf_{\varepsilon \rightarrow 0} (\Phi_x(\varepsilon)), \quad x \in \Omega. \quad (55)$$

In the degenerate case $A = \{a\}$, define $f_A(x) := f(a)$.

The first order of business is to verify that PMSE is well-defined (i.e., that the limit in (55) indeed exists):

Lemma 19 Assume $A \subseteq \Omega$, $|A| \geq 2$, and $\|f|_A\|_{\text{Lip}} < \infty$. Then, for (u, v) and (u', v') in $W_x(\varepsilon)$, $\varepsilon < R_x^*/2$, we have

$$|F_x(u, v) - F_x(u', v')| \leq \varepsilon \min\{4 \text{diam}(A), 16\|f|_A\|_{\infty}/R_x^*\}.$$

Proof Since x is fixed, we omit it from the subscripts for readability. Observe that $R(u, v) > 0$ for $(u, v) \in W(\varepsilon)$ and that $F(u, v) = f(v) - R(u, v)\rho(x, v)$. Thus,

$$f(u) \leq F(u, v) \leq f(v), \quad (56)$$

and the same holds for $(u', v') \in W(\varepsilon)$. There is no loss of generality in assuming $F(u, v) \leq F(u', v')$. In this case, (56) implies that $f(u) \leq f(v')$ and hence

$$\begin{aligned} R^* &\geq \frac{f(v') - f(u)}{\rho(x, v') + \rho(x, u)} = \frac{f(v') - F(u', v') + F(u', v') - F(u, v) + F(u, v) - f(u)}{\rho(x, v') + \rho(x, u)} \\ &= \frac{f(v') - F(u', v') + F(u, v) - f(u)}{\rho(x, v') + \rho(x, u)} + \frac{F(u', v') - F(u, v)}{\rho(x, v') + \rho(x, u)} \\ &= \frac{R(u', v')\rho(x, v') + R(u, v)\rho(x, u)}{\rho(x, v') + \rho(x, u)} + \frac{F(u', v') - F(u, v)}{\rho(x, v') + \rho(x, u)} \\ &\geq \frac{R(u', v')\rho(x, v') + R(u, v)\rho(x, u)}{\rho(x, v') + \rho(x, u)} + \frac{F(u', v') - F(u, v)}{2 \text{diam}(A)} \\ &\geq \frac{R(u', v')\rho(x, v') + (R(u', v') - \varepsilon)\rho(x, u)}{\rho(x, v') + \rho(x, u)} + \frac{F(u', v') - F(u, v)}{2 \text{diam}(A)} \\ &= R(u', v') - \varepsilon + \frac{F(u', v') - F(u, v)}{2 \text{diam}(A)}. \end{aligned}$$

This proves

$$|F_x(u, v) - F_x(u', v')| \leq 4\varepsilon \operatorname{diam}(A). \quad (57)$$

To prove the remaining claim, we argue that for $\varepsilon < R^*/2$, all $(u, v) \in W(\varepsilon)$ satisfy $u, v \in B(x, 2\|f|_A\|_\infty/R^*)$. Indeed, assume for a contradiction that some $(u, v) \in W(\varepsilon)$ violates this assumption, with $\rho(x, v) \vee \rho(x, u) > 2\|f|_A\|_\infty/R^*$. Then,

$$R(u, v) = \frac{f(v) - f(u)}{\rho(x, v) + \rho(x, u)} \leq \frac{\|f|_A\|_\infty}{\rho(x, v) + \rho(x, u)} \leq \frac{R^*}{2} < R^* - \varepsilon,$$

implying that $(u, v) \notin W(\varepsilon)$, a contradiction. Since the diameter of the point pairs in $W(\varepsilon)$ is at most $4\|f|_A\|_\infty/R^*$, we can repeat the calculation leading to (57) to complete the proof. \blacksquare

Corollary 20 *Suppose that $f : \Omega \rightarrow \mathbb{R}$, $x \in \Omega$ and $A \subseteq \Omega$, $|A| \geq 2$ are given. If $\|f|_A\|_{\text{Lip}} < \infty$, then (55) is well-defined.*

Proof Lemma 19 shows that $\|f|_A\|_{\text{Lip}} \vee \|f|_A\|_\infty < \infty$ implies that approximate maximizers of $R_x(\cdot, \cdot)$ all yield approximately the same value of $F_x(\cdot, \cdot)$. In particular,

$$|F_x(u, v) - F_x(u', v')| \leq \varepsilon 16\|f|_A\|_\infty/R_x^*, \quad \varepsilon < R_x^*/2. \quad (58)$$

Now we drop the assumption that $\|f|_A\|_\infty < \infty$. For $n \in \mathbb{N}$, truncate f via $f_n := n \wedge (f \vee -n)$ and define $F_x^{(n)}$, $R_x^{*(n)}$ as above with f_n in place of f . A simple continuity argument shows that $R_x^{*(n)} \rightarrow R_x^*$ and $F_x^{(n)} \rightarrow F_x$ (the latter, pointwise) as $n \rightarrow \infty$. (Note that $R_x^* > 0$ and $\|f|_A\|_{\text{Lip}} < \infty \implies \sup_{x \in \Omega} R_x^* < \infty$.) Taking $\varepsilon = n^{-2}$ in (58), we have

$$|F_x^{(n)}(u, v) - F_x^{(n)}(u', v')| \leq \frac{32}{nR_x^*}$$

for all n sufficiently large that $R_x^{*(n)} \geq R_x^*/2$. This proves the claim. \blacksquare

Remark 21 *When $R_x(\cdot, \cdot)$ has a unique maximizer $(u^*, v^*) \in A^2$, the definition of f_A simplifies to*

$$f_A(x) = f(u^*) + \frac{\rho(u^*, x)}{\rho(u^*, x) + \rho(v^*, x)}(f(v^*) - f(u^*)). \quad (59)$$

In light of Corollary 20, whenever $\|f|_A\|_\infty < \infty$, there is no loss of generality in assuming that for each $x \in \Omega$, there is a unique maximizer $(u^, v^*) = (u^*(x), v^*(x))$. In particular, (59) shows that for finite A , one can compute $f_A(x)$ at any given x in time $O(|A|^2)$.*

Remark 22 *Another consequence of Corollary 20 is that, when $\|f|_A\|_\infty < \infty$, the PMSE extension admits the conceptually simpler form*

$$f_A(x) = \inf_{y \in \mathbb{R}} \sup_{a \in A} \frac{|f(a) - y|}{\rho(a, x)}.$$

This will always be the case for finite A .

For the remainder of this section, the conditions $|A| \geq 2$ and $\|f|_A\|_\infty < \infty$ are assumed. It is readily verified that for $x \in A$, we have $f(x) = f_A(x)$; thus PMSE is indeed an extension.

Theorem 23 (Pointwise minimality of the PMSE) *For $A \subseteq \Omega$ and $f : \Omega \rightarrow \mathbb{R}$, let f_A be the extension of f from A to Ω . Then*

$$\Lambda_{f_A}(x, \Omega) \leq \Lambda_f(x, \Omega), \quad x \in \Omega.$$

Proof We break down the proof into three shorter claims. As argued in Remark 21, there is no loss of generality in assuming, for any $x \in \Omega$, a unique maximizer $(u^*, v^*) = (u^*(x), v^*(x))$ of $R_x(u, v)$ over A^2 .

Claim I. For any $x \in \Omega \setminus A$, PMSE achieves the minimum local slope on A among all functions that agree with f_A on A :

$$\Lambda_{f_A}(x, A) \leq \Lambda_f(x, A), \quad x \in \Omega \setminus A.$$

We first show that f_A achieves the optimal slope at x with respect to (u^*, v^*) , which define $f_A(x)$ as in (59). It is enough to show that $\frac{|f_A(u^*) - f_A(x)|}{\rho(u^*, x)} = \frac{|f_A(v^*) - f_A(x)|}{\rho(v^*, x)}$, since any other value of $f_A(x)$, for which the equality does not hold, will result in a larger slope between x and either u^* or v^* . In light of (56), there is no loss of generality in assuming $f_A(u^*) \leq f_A(x) \leq f_A(v^*)$. Then:

$$\begin{aligned} \frac{|f_A(u^*) - f_A(x)|}{\rho(u^*, x)} &= \frac{\left[f(u^*) + \frac{\rho(u^*, x)}{\rho(v^*, x) + \rho(u^*, x)} (f(v^*) - f(u^*)) \right] - f(u^*)}{\rho(u^*, x)} \\ &= \frac{f(v^*) - f(u^*)}{\rho(v^*, x) + \rho(u^*, x)} \\ &= \frac{f(v^*) - \left[f(v^*) - \frac{\rho(v^*, x)}{\rho(v^*, x) + \rho(u^*, x)} (f(v^*) - f(u^*)) \right]}{\rho(v^*, x)} \\ &= \frac{|f_A(v^*) - f_A(x)|}{\rho(v^*, x)}. \end{aligned} \tag{60}$$

Let $\ell_1 := \frac{|f_A(u^*) - f_A(x)|}{\rho(u^*, x)} = \frac{f(v^*) - f(u^*)}{\rho(v^*, x) + \rho(u^*, x)}$. It remains to show that $\frac{|f_A(x') - f_A(x)|}{\rho(x', x)} \leq \ell_1$ for all $x' \in A$. Assume, to the contrary, the existence of an $x' \in A$ such that $\frac{|f_A(x') - f_A(x)|}{\rho(x', x)} > \ell_1$. Then by (52),

$$\frac{f(x') - f(u^*)}{\rho(x', x) + \rho(x, u^*)} = \frac{f_A(x') - f_A(x) + f_A(x) - f_A(u^*)}{\rho(x', x) + \rho(x, u^*)} > \ell_1,$$

which contradicts the definition of (u^*, v^*) in (59) as maximizers of R_x . This proves Claim I.

Claim II.

$$\Lambda_{f_A}(x, \Omega \setminus A) \leq \Lambda_{f_A}(x, A), \quad x \in \Omega \setminus A.$$

Let us define the *slope* operator $S(u, v) := |f_A(u) - f_A(v)|/\rho(u, v)$. It suffices to show that

$$S(x, y) \leq \Lambda_{f_A}(x, A) \wedge \Lambda_{f_A}(y, A), \quad x, y \in \Omega \setminus A.$$

Let $(u^*(x), v^*(x))$ and $(u^*(y), v^*(y))$ be as defined in (59). It follows from the proof of Claim I that $S(x, u^*(y)) \vee S(x, v^*(y)) \leq \Lambda_{f_A}(x, A)$.

Assume for concreteness that $\Lambda_{f_A}(x, A) \leq \Lambda_{f_A}(y, A)$. As in the proof of Claim I, $f_A(u^*(x)) \leq f_A(x) \leq f_A(v^*(x))$ and $f_A(u^*(y)) \leq f_A(y) \leq f_A(v^*(y))$. Suppose, for a contradiction, that $S(x, y) > \Lambda_{f_A}(x, A)$.

If $f_A(x) \leq f_A(y)$, then

$$f_A(v^*(y)) = f_A(x) + \rho(x, y)S(x, y) + \rho(y, v^*(y))\Lambda_{f_A}(y) > f_A(x) + \rho(x, v^*(y))\Lambda_{f_A}(x),$$

implying that $S(x, v^*(y)) > \Lambda_{f_A}(x)$ — a contradiction.

If $f_A(x) > f_A(y)$, then

$$f_A(x) = f_A(u^*(y)) + \rho(u^*(y), y)\Lambda_{f_A}(y) + \rho(y, x)S(x, y) > f_A(u^*(y)) + \rho(u^*(y), x)\Lambda_{f_A}(x),$$

implying that $S(x, u^*(y)) > \Lambda_{f_A}(x)$ — a contradiction, which proves Claim II.

Claim III. ⁵

$$\Lambda_{f_A}(x, \Omega) = \Lambda_{f_A}(x, A) \leq \Lambda_f(x, \Omega), \quad x \in A.$$

Assume for a contradiction that for some $y \notin A$ we have

$$\Lambda_{f_A}(x, \Omega) \geq S(x, y) = \frac{|f_A(x) - f_A(y)|}{\rho(x, y)} > \Lambda_{f_A}(x, A).$$

We denote by $(u_y^*, v_y^*) \in A \times A$ the maximizer of $R_y(\cdot, \cdot)$. Recall that since $x \in A$, in the proof of Claim I we showed that $S(x, y) \leq S(y, u_y^*) = S(y, v_y^*)$. If $F_A(x) \leq F_A(y) \leq F_A(v_y^*)$ then

$$\begin{aligned} S(x, v_y^*) &= \frac{F_A(v_y^*) - F_A(x)}{\rho(v_y^*, x)^\beta} \geq \frac{F_A(v_y^*) - F_A(y) + F_A(y) - F_A(x)}{\rho(v_y^*, y)^\beta + \rho(x, y)^\beta} \\ &\geq \min\{S(y, v_y^*), S(x, y)\} = S(x, y) > \Lambda_{F_A}^\beta(x, A), \end{aligned}$$

while on the other hand if $F_A(x) > F_A(y) \geq F_A(u_y^*)$ then

$$\begin{aligned} S(x, u_y^*) &= \frac{F_A(x) - F_A(u_y^*)}{\rho(x, u_y^*)^\beta} \geq \frac{F_A(x) - F_A(y) + F_A(y) - F_A(u_y^*)}{\rho(x, y)^\beta + \rho(u_y^*, y)^\beta} \\ &\geq \min\{S(x, y), S(y, u_y^*)\} = S(x, y) > \Lambda_{F_A}(x, A), \end{aligned}$$

where in both calculations we used the mediant inequality. Both inequalities above contradict the definition of $\Lambda_{F_A}(x, A)$, thus proving Claim III.

⁵. Our original proof of this claim had a slight inaccuracy, pointed out by Guy Kornowski. It has since been corrected (and generalized to the Hölder case) in Hanneke et al. (2023, Theorem 16).

Putting it together. Combining Claims II and III yields that the local slope of x with respect to f_A is determined by a point in A :

$$\Lambda_{f_A}(x, \Omega) = \Lambda_{f_A}(x, A), \quad x \in \Omega.$$

Therefore:

$$\Lambda_{f_A}(x, \Omega) = \Lambda_{f_A}(x, A) \leq \Lambda_f(x, A) \leq \Lambda_f(x, \Omega) \quad (61)$$

where the first inequality stems from Claim I and the fact that f and f_A agree on A . \blacksquare

Corollary 24 (Properties of the PMSE) *Suppose that $A \subseteq \Omega$ and $f : \Omega \rightarrow \mathbb{R}$ satisfy $\|f|_A\|_\infty < \infty$. Then for any $x \in \Omega$, $B \subseteq A$, and $(u^*, v^*) = (u^*(x), v^*(x))$ as defined in Remark 21, the following hold:*

(i) *Local slope value:*

$$\Lambda_{f_A}(x, \Omega) = \frac{|f_A(x) - f_A(u^*)|}{\rho(x, u^*)} = \frac{|f_A(x) - f_A(v^*)|}{\rho(x, v^*)} = \frac{|f_A(v^*) - f_A(u^*)|}{\rho(v^*, x) + \rho(x, u^*)};$$

(ii) *Local slope bounds:*

$$\Lambda_{f_A}(x, \Omega) \leq \Lambda_{f_A}(u^*, A) \wedge \Lambda_{f_A}(v^*, A) \leq \Lambda_{f_A}(u^*, \Omega) \wedge \Lambda_{f_A}(v^*, \Omega);$$

(iii) *Lipschitz:* $\|f_A\|_{\text{Lip}} = \|f|_A\|_{\text{Lip}}$;

(iv) *Local slope monotonicity:* $\Lambda_{f_B}(x, \Omega) \leq \Lambda_{f_A}(x, \Omega)$;

(v) *Extension sandwich:* $f_A(u^*) \leq f_A(x) \leq f_A(v^*)$.

Proof All of the claims follow from Theorem 23 and its proof:

(i) Follows from Claim II and the proof of Claim I.

(ii) Follows from (i): $\Lambda_{f_A}(x, \Omega) = \frac{|f_A(v^*) - f_A(u^*)|}{\rho(v^*, x) + \rho(x, u^*)} \leq \frac{|f_A(v^*) - f_A(u^*)|}{\rho(v^*, u^*)} \leq \Lambda_{f_A}(u^*, A)$.

(iii) Follows from (ii): $\Lambda_{f_A}(x, \Omega) \leq \Lambda_{f_A}(u^*, A) \leq \|f|_A\|_{\text{Lip}}$.

(iv) Direct result of Theorem 23 where f_A assumes the role of f .

(v) Was proved in the course of proving Claim I. \blacksquare

Comparison between PMSE and AMLE. At first glance, PMSE might appear similar to the Absolutely Minimal Lipschitz Extension (AMLE) (Juutinen, 2002; Peres et al., 2009). It was already observed by Oberman (2008) that the two are distinct. Note first that AMLE requires a length space in order to be well-defined, while PMSE of f from $A \subset \Omega$ to Ω is well-defined as long as $\|f|_A\|_\infty < \infty$.

A visual comparison of the behaviors of AMLE and PMSE on $\Omega = [0, 1]$ is instructive. We evaluate the step function $f(x) = \mathbf{1}[x > 1/2]$ at 10 uniformly spaced “anchor” points on $[0, 1]$. The AMLE is just the linear interpolation, illustrated by the piecewise-linear (blue) curve in Figure 1. The behavior of PMSE is more interesting: at first, the sawtooth (red) shape looks somewhat odd. Recall, however, that the PMSE is minimizing the *local* slope at each point, which is affected by the values at all of the anchor points. Thus, each bottom spike reflects the tension between the 0-value at the nearby anchor points and the 1-value at the farther anchor points (and similarly for the top spikes). The two curves coincide at the line segment representing the steep rise between the 5th and 6th anchor points.




Figure 1: Comparing PMSE to AMLE on the step function.

Appendix C. Auxiliary results

Lemma 25 *Suppose that X is a random variable with finite support $R \subset \mathbb{R}_+$. Then*

$$\mathbb{E}[X] \leq \mathbb{W}[X] + \mathbb{W}[X] \log \frac{1}{p_*},$$

where $p_* := \min \{\mathbb{P}(X = r) : r \in R\}$ and $\mathbb{W}[X] := \sup_{t>0} t \mathbb{P}(X \geq t)$ is the “weak mean” of X . If, additionally, $p_* \leq 1/3$, then $\mathbb{E}[X] \leq 2 \mathbb{W}[X] \log \frac{1}{p_*}$.

Proof Put $r^* = \max R$. Then

$$\begin{aligned} \mathbb{E}[X] &= \int_0^{r^*} \mathbb{P}(X \geq t) dt \\ &\leq a + \mathbb{W}[X] \int_a^{r^*} \frac{dt}{t}, \quad 0 < a \leq r^*. \end{aligned}$$

The integral evaluates to $\log(r^*/a)$, and our choice $a := p_* r^*$ yields the bound

$$\begin{aligned} \mathbb{E}[X] &\leq p_* r^* + \mathbb{W}[X] \log \frac{1}{p_*} \\ &\leq \mathbb{W}[X] + \mathbb{W}[X] \log \frac{1}{p_*}. \end{aligned}$$

■

Corollary 26 *Let (Ω, ρ) be a finite metric space endowed with the uniform distribution μ . Then, for any $f : \Omega \rightarrow \mathbb{R}$, we have*

$$\bar{\Lambda}_f(\mu, \Omega) \leq \tilde{\Lambda}_f(\mu, \Omega) + \log(|\Omega|) \tilde{\Lambda}_f(\mu, \Omega).$$

When $|\Omega| \geq 3$, we have $\bar{\Lambda}_f(\mu, \Omega) \leq 2 \log(|\Omega|) \tilde{\Lambda}_f(\mu, \Omega)$.

C.1 Adversarial extension: deferred proofs

Proof [Proof of Lemma 14] By Corollary 24(i), we have that the local slope of f at x is determined by a pair of points $u^*, v^* \in \Omega_n$:

$$\Lambda_f(x) = \frac{f(v^*) - f(u^*)}{\rho(v^*, x) + \rho(x, u^*)}.$$

From (40), we have $\rho(v^*, x) + \rho(x, u^*) \geq \rho(v^*, u^*) \geq 2 \operatorname{diam}(E)$, and hence $\rho(v^*, x) + \rho(x, u^*) + 2 \operatorname{diam}(E) \leq 2(\rho(v^*, x) + \rho(x, u^*))$. Thus,

$$\begin{aligned} \Lambda_f(x') &\geq \frac{f(v^*) - f(u^*)}{\rho(v^*, x') + \rho(x', u^*)} \\ &\geq \frac{f(v^*) - f(u^*)}{\rho(v^*, x) + \operatorname{diam}(E) + \rho(x, u^*) + \operatorname{diam}(E)} \\ &\geq \frac{f(v^*) - f(u^*)}{2(\rho(v^*, x) + \rho(x, u^*))} = \frac{\Lambda_f(x)}{2}. \end{aligned}$$

■

Proof [Proof of Lemma 15] The claims in (41) and (42) are standard applications of multiplicative Chernoff bounds (Dubhashi and Panconesi, 2009, Theorem 1.1). To prove (43), observe that

$$\begin{aligned} \mathbb{P}\left(\sum_{B \in \Pi_0} \mu(B) \geq 2q\right) &\leq \mathbb{P}\left(\sum_{B \in \Pi_0} (\mu(B) - \mu_n(B)) \geq q\right) \\ &\leq \mathbb{P}\left(\sum_{B \in \Pi} |\mu(B) - \mu_n(B)| \geq q\right). \end{aligned}$$

To bound the latter, define the random variable $J_n := \sum_{B \in \Pi} |\mu(B) - \mu_n(B)|$. It follows from Berend and Kontorovich (2013, Eqs. (5) and (17)) that $\mathbb{E}[J_n] \leq \sqrt{m/n}$ and

$$\begin{aligned} \mathbb{P}(J_n \geq q) &\leq \mathbb{P}(J_n \geq \mathbb{E}[J_n] + (q - \sqrt{m/n})) \\ &\leq \exp[-n(q - \sqrt{m/n})^2/2], \quad nq^2 \geq m, \end{aligned}$$

which completes the proof. ■

Proof [Proof of Lemma 16] We only prove the first inequality, since the second one is immediate from (11). It follows directly from the definition of $\tilde{\Lambda}_y(\mu_n, \Omega_n)$ that

$$\mu_n(M_y(L)) \geq \alpha \implies L \leq \alpha^{-1} \tilde{\Lambda}_y(\mu_n, \Omega_n).$$

The algorithm removes the $\lfloor \varepsilon n \rfloor$ points with the largest $\Lambda_y(x, \Omega_n)$ from the set Ω_n and it is a basic fact that

$$\varepsilon n \leq 2 \lfloor \varepsilon n \rfloor, \quad \varepsilon \geq n^{-1}.$$

This corresponds to removing a mass of $\alpha \geq \varepsilon/2$ points, and so for $n \geq \varepsilon^{-1}$,

$$\max_{x \in \Omega'_n(\varepsilon)} \Lambda_y(x, \Omega'_n(\varepsilon)) \leq \frac{2\tilde{\Lambda}_y(\mu_n, \Omega_n)}{\varepsilon}. \tag{62}$$

Finally, Corollary 24(iii,iv) implies

$$\|f\|_{\text{Lip}} = \|f|_V\|_{\text{Lip}} \leq \left\| f|_{\Omega'_n(\varepsilon)} \right\|_{\text{Lip}} \leq \frac{2\tilde{\Lambda}_y(\mu_n, \Omega_n)}{\varepsilon}. \tag{63}$$

■

Appendix D. Generalization

Generalization guarantees refer to claims bounding the true risk in terms of the empirical risk, plus confidence and hypothesis complexity terms. Throughout this paper, we assume that the learner has a fixed maximal allowable average slope L . This assumption incurs no loss of generality, since a standard technique, known as Structural Risk Minimization (SRM), (Shawe-Taylor et al., 1998), creates a nested family of function classes with increasing L , allowing the learner to select one based on the sample, so as to optimize the underfit-overfit tradeoff.

D.1 Uniform Glivenko-Cantelli

We begin by dispensing with some measure-theoretic technicalities. Our learner constructs a hypothesis $f : \Omega \rightarrow [0, 1]$ via the PMSE extension, which, by Corollary 24(iii), is a Lipschitz function. Thus, operationally, the learner's function class is

$$\mathcal{H}_L = \widetilde{\text{Lip}}_L(\Omega, \rho, \mu) \cap \left\{ f : \Omega \rightarrow \mathbb{R}; \|f\|_{\text{Lip}} < \infty \right\}. \quad (63)$$

Our assumptions that $\text{ddim}(\Omega), \text{diam}(\Omega) < \infty$ imply that (Ω, ρ) has compact closure; this in turn implies a countable $\mathcal{F} \subset \mathcal{H}_L$ such that every member of \mathcal{H}_L is a pointwise limit of a sequence in \mathcal{F} . This suffices (Dudley, 1999) to ensure the measurability of the empirical process $\sup_{f \in \mathcal{H}_L} \left(\int_{\Omega} f d\mu - \int_{\Omega} f d\mu_n \right)$.

A standard method for bounding this empirical process is via the Rademacher complexity (see, e.g., Mohri et al. (2012, Theorem 3.1)): with probability at least $1 - \delta$,

$$\sup_{f \in \mathcal{H}_L} \left(\int_{\Omega} f d\mu - \int_{\Omega} f d\mu_n \right) \leq 2\mathfrak{R}_n(\mathcal{H}_L | X_{[n]}) + 3\sqrt{\frac{\log(2/\delta)}{2n}}, \quad (64)$$

where $X_{[n]} = (X_1, \dots, X_n) \sim \mu^n$ and

$$\mathfrak{R}_n(\mathcal{H}_L | X_{[n]}) := \mathbb{E}_{\sigma \sim \text{Uniform}(\{-1, 1\}^n)} \sup_{f \in \mathcal{H}_L} \frac{1}{n} \sum_{i=1}^n \sigma_i f(X_i).$$

Finally, an elementary estimate on Rademacher complexity is in terms of the empirical L_2 covering numbers (see, e.g., Bartlett (2006)):

$$\mathfrak{R}_n(\mathcal{H}_L | X_{[n]}) \leq \inf_{\varepsilon > 0} \varepsilon + \sqrt{\frac{2 \log \mathcal{N}(\varepsilon, \mathcal{H}_L, L_2(\mu_n))}{n}}. \quad (65)$$

Invoking the estimate in Theorem 3 with $\alpha = L^{-1/6}$ and combining it with (64, 65) yields that

$$\sup_{f \in \mathcal{H}_L} \left(\int_{\Omega} f d\mu - \int_{\Omega} f d\mu_n \right) \leq \frac{C_\delta \sqrt{L}}{n^{1/8d}} + \frac{C_\delta^{-d/2} \sqrt{2}}{n^{5/16}} + 3\sqrt{\frac{\log(2/\delta)}{2n}} \quad (66)$$

holds with probability at least $1 - 3\delta$, where $d = \text{ddim}(\Omega)$ and C_δ is a constant depending only on δ (assuming, to avoid trivialities, that $L, d \geq 1$). This establishes the so-called Uniform Glivenko-Cantelli (UGC) property of \mathcal{H}_L . We have not attempted to optimize any of the constants.

D.2 Risk bounds

Recall our learning setup: ν is an unknown distribution on $\Omega \times [0, 1]$, from which the learner receives a labeled sample $S_n = (X_i, Y_i)_{i \in [n]} \sim \nu^n$. Based on S_n , the learner constructs a hypothesis $f : \Omega \rightarrow [0, 1]$, to which we associate the (true) risk $R(f; \nu) := \mathbb{E}_{(X, Y) \sim \nu} |f(X) - Y|$ as well as the empirical risk $R(f; \nu_n)$.

We discuss regression and classification separately.

Regression. The learner selects a hypothesis $f \in \mathcal{H}_L$ (seeking to minimize $R(f; \nu_n)$, but this will not be used in our analysis). Then

$$\sup_{f \in \mathcal{H}_L} (R(f; \nu_n) - R(f; \nu)) \leq 2\mathfrak{R}(\mathcal{G}|(X, Y)_{[n]}) + 3\sqrt{\frac{\log(2/\delta)}{2n}}$$

holds with probability at least $1 - \delta$, where \mathcal{G} is the *loss class* consisting of

$$\mathcal{G} = \{g : (x, y) \mapsto |f(x) - y|; x \in \Omega, y \in [0, 1], f \in \mathcal{H}_L\} \subseteq [0, 1]^{\Omega \times [0, 1]}.$$

To bound the Rademacher complexity of \mathcal{G} , we notice that each $g \in \mathcal{G}$ is of the form $g = \varphi \circ h$, where $\varphi(\cdot) = |\cdot|$ and $h : (x, y) \mapsto f(x) - y$. Since $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is 1-Lipschitz, Talagrand’s contraction principle (Ledoux and Talagrand, 1991, Corollary 3.17) implies that

$$\begin{aligned} \mathfrak{R}(\mathcal{G}|(X, Y)_{[n]}) &\leq \mathbb{E} \sup_{\sigma} \sup_{f \in \mathcal{H}_L} \frac{1}{n} \sum_{i=1}^n \sigma_i (f(X_i) - Y_i) \\ &= \mathbb{E} \sup_{\sigma} \sup_{f \in \mathcal{H}_L} \frac{1}{n} \sum_{i=1}^n \sigma_i f(X_i) = \mathfrak{R}(\mathcal{H}_L|X_{[n]}). \end{aligned}$$

Combining this with (66) yields

$$\sup_{f \in \mathcal{H}_L} (R(f; \nu_n) - R(f; \nu)) \leq \frac{C_\delta \sqrt{L}}{n^{1/8d}} + \frac{C_\delta^{-d/2} \sqrt{2}}{n^{5/16}} + 3\sqrt{\frac{\log(2/\delta)}{2n}} \quad (67)$$

with probability at least $1 - 3\delta$.

Classification. For classification, the learning setup is asymmetric with respect to training and prediction (see the discussion at the beginning of Section 7). The distribution ν over $\Omega \times \{0, 1\}$, and again, $S = (X_i, Y_i)_{i \in [n]} \sim \nu^n$ is presented to the learner. The latter produces a hypothesis $f : \Omega \rightarrow [0, 1]$, to which we associate the *sample error*,

$$\widehat{\text{err}}(f) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}[f(X_i) \neq Y_i], \quad (68)$$

and the *generalization error*,

$$\text{err}(f) = \mathbb{P}_{(X, Y) \sim \nu} (\mathbf{1}[f(X) > 1/2] \neq Y). \quad (69)$$

Notice the asymmetry between $\widehat{\text{err}}$ and err : a hypothesis is penalized for every sample point it fails to label correctly, but is only required to “be closer to the correct label” at test time.

Given these definitions, a standard argument (via the *margin function* and Talagrand’s contraction, see Mohri et al. (2012, Theorem 4.4)), yields

$$\sup_{f \in \mathcal{H}_L} (\text{err}(f) - \widehat{\text{err}}(f)) \leq \frac{C_\delta \sqrt{L}}{n^{1/8d}} + \frac{C_\delta^{-d/2} \sqrt{2}}{n^{5/16}} + 3\sqrt{\frac{\log(2/\delta)}{2n}} \quad (70)$$

with probability at least $1 - 3\delta$.

Relation to classic results. Take $\Omega = \{1, 2, \dots\}$ with any probability measure μ and metric $\rho(x, x') = |x - x'|$. Consider the function class $\mathcal{F} = [0, 1]^\Omega$. It is well-known that \mathcal{F} is Uniform Glivenko-Cantelli (UGC) with respect to μ ; this follows, for example, from missing mass arguments (see, e.g., (Efremenko et al., 2020, Theorem 2)). We note in passing that under a fixed distribution, UGC is a strictly stronger property than learnability (Benedek and Itai, 1991). Let us specialize our general techniques to this toy setting.

It is easily seen that $\text{ddim}(\Omega) = 1$, but we must address the technical issue that $\text{diam}(\Omega) = \infty$. A cursory glance at the proof of Theorem 3 shows that in fact only $\text{diam}(\Omega_n) < \infty$ is needed. In fact, even further savings is possible: we can relabel the elements of Ω_n so that

$$\text{diam}(\Omega_n) = \|\mu_n\|_0 \equiv |\text{supp}(\mu_n)| =: |\{x \in \Omega : \mu_n(x) > 0\}|.$$

Renormalizing the empirical diameter to 1 by shrinking the distances by $\|\mu_n\|_0$, we have that $\|f|_{\Omega_n}\|_{\text{Lip}} \leq \|\mu_n\|_0$ for all $f \in \mathcal{F}$. To this we may apply Lemma 2, obtaining a t -covering number bound (under ℓ_∞) of order $O(\frac{\|\mu_n\|_0}{t} \log \frac{1}{t})$. Then the Rademacher bound in (65) yields a rate of $O((\|\mu_n\|_0/n)^{1/3})$, although Dudley’s chaining integral (Vershynin, 2018, Theorem 8.1.3) yields the sharper estimate $O((\|\mu_n\|_0/n)^{1/2})$. A more careful analysis (Cohen et al., 2020) shows that essentially the optimal rate is $O((\|\mu_n\|_{1/2}/n)^{1/2})$. The estimate in (66) loses out considerably, due to the additive error in the covering number bound in Theorem 3; however, it does suffice to conclude that \mathcal{F} is UGC. Note that our techniques establish finite empirical L_2 covering numbers for this class — unlike, say, the covering number estimate of Mendelson and Vershynin (2003), which requires bounded fat-shattering dimension.

Appendix E. Adversarial extension for regression, weak mean

In this section, we prove the weak-mean counterparts (a.ii), (b.ii) of the adversarial extension game for regression, defined in Section 6. These results are not used in the paper and are included for completeness and independent interest. The following notation will be used throughout:

$$\begin{aligned} \ell_{\min} &:= \min_{x \in \Omega_n} \Lambda_f(x, \Omega_n), \\ \ell_{\max} &:= \max_{x \in \Omega_n} \Lambda_f(x, \Omega_n). \end{aligned}$$

The function $f : \Omega \rightarrow [0, 1]$ will refer exclusively to the one constructed by the “adversarial extension” algorithm in the proof of Lemma 13. We proceed to make some observations regarding f .

Lemma 27 (f is defect-free) *For all $\ell > 0$, the function f is $(\ell\varepsilon/2, \ell, 1)$ -defect-free, as defined in Section 4.1.*

Proof By construction, f is the PMSE of an ε -net V . Let $u^*(x)$ and $v^*(x)$ be as in Remark 21. Suppose that $\Lambda_f(x, \Omega) \geq \ell$. Then

$$\begin{aligned} &\max\{|f(x) - f(u^*(x))|, |f(x) - f(v^*(x))|\} \\ &\geq \ell \cdot \max\{\rho(x, u^*(x)), \rho(x, v^*(x))\} \geq \ell \cdot \frac{\varepsilon}{2} = \ell\varepsilon/2. \end{aligned}$$

■

Lemma 28 (Combinatorial structure) For any $t > 0$ and $0 < q < 1$, we have

$$\mu(M_f(t)) \leq 2\mu_n(M_f(t/4)) + \frac{8m}{n^{1/3}} \log(m/\delta), \quad (71)$$

with probability at least $1 - 2\delta$, where $m \leq (8/\varepsilon)^d$.

Lemma 29 (Bounded local slope ratio)

$$\max_{x \neq x' \in \Omega} \frac{\Lambda_f(x, \Omega)}{\Lambda_f(x', \Omega)} \leq \frac{2 \operatorname{diam}(\Omega)}{\varepsilon}.$$

Proofs of Lemmas 28 and 29 are deferred to the end of this section.

Lemma 30 (f is close to y in weak mean) The “adversarial extension” function f satisfies (a.ii) for $0 < \varepsilon < 1$.

Proof We begin with the same decomposition as in (38):

$$\|f - y\|_{L_1(\mu_n)} = \frac{1}{n} \sum_{x \in \Omega_n(\varepsilon) \setminus V} |f(x) - y(x)| + \frac{1}{n} \sum_{x \in \Omega'_n(\varepsilon) \setminus V} |f(x) - y(x)|$$

and, as above, bound the first term by ε and the second term as in (39):

$$\begin{aligned} \frac{1}{n} \sum_{x \in \Omega'_n(\varepsilon) \setminus V} |f(x) - y(x)| &\leq \frac{2\varepsilon}{n} \sum_{x \in \Omega'_n(\varepsilon) \setminus V} \Lambda_y(x, \Omega_n) \\ &= 2\varepsilon \frac{|\Omega'_n(\varepsilon) \setminus V|}{n} \int_{\Omega'_n(\varepsilon) \setminus V} \Lambda_y(x, \Omega_n) d\bar{\mu}(x), \end{aligned}$$

where $\bar{\mu}$ is the uniform measure on $\Omega'_n(\varepsilon) \setminus V$, given by $\bar{\mu}(x) = n|\Omega'_n(\varepsilon) \setminus V|^{-1}\mu_n(x)$. Now

$$\begin{aligned} \frac{|\Omega'_n(\varepsilon) \setminus V|}{n} \int_{\Omega'_n(\varepsilon) \setminus V} \Lambda_y(x, \Omega_n) d\bar{\mu}(x) &= \frac{|\Omega'_n(\varepsilon) \setminus V|}{n} \int_0^\infty \bar{\mu}(\{x \in \Omega'_n(\varepsilon) \setminus V : \Lambda_y(x, \Omega_n) > t\}) dt \\ &= \int_0^\infty \mu_n(\{x \in \Omega'_n(\varepsilon) \setminus V : \Lambda_y(x, \Omega_n) > t\}) dt \\ &= \int_0^\beta \mu_n(\{x \in \Omega'_n(\varepsilon) \setminus V : \Lambda_y(x, \Omega_n) > t\}) dt \\ &\leq \alpha + \tilde{\Lambda}_y(\mu_n, \Omega_n) \int_\alpha^\beta \frac{dt}{t}, \end{aligned}$$

where $\alpha > 0$ is arbitrary and $\beta = \frac{2\tilde{\Lambda}_y(\mu_n, \Omega_n)}{\varepsilon}$, in light of (62). The integral evaluates to $\log(\beta/\alpha)$ and our choice $\alpha := \tilde{\Lambda}_y(\mu_n, \Omega_n)$ yields the estimate

$$\frac{1}{n} \sum_{x \in \Omega'_n(\varepsilon) \setminus V} |f(x) - y(x)| \leq 2\varepsilon(\alpha + \alpha \log \frac{2}{\varepsilon}) = \tilde{O}(\varepsilon) \tilde{\Lambda}_y(\mu_n, \Omega_n).$$

■

Lemma 31 (Satisfying (b.ii)) For $0 < \varepsilon, \delta < 1$, the adversarial extension function f satisfies

$$\tilde{\Lambda}_f(\mu, \Omega) \leq 16\tilde{\Lambda}_f(\mu_n, \Omega_n) + 64\tilde{\Lambda}_y(\mu_n, \Omega_n) \frac{m}{\varepsilon n^{1/3}} \log \left(\frac{m \log(2/\varepsilon)}{\delta} \right)$$

with probability at least $1 - 2\delta$, where $m \leq (8/\varepsilon)^{\text{ddim}}$.

Proof Fix a $\delta > 0$ and put $\delta_{ij} := \delta 2^{-i-j}$; then $\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \delta_{ij} = \delta$. Define also $\tau_{ij} = 2^{i-j}$ for $i, j \in \{1, 2, \dots\}$. Invoking Lemma 28 with the union bound, we have:

$$\mathbb{P} \left(\exists \tau_{ij} : \mu(M_f(\tau_{ij})) > 2\mu_n(M_f(\tau_{ij}/4)) + \frac{8m}{n^{1/3}} \log \frac{m 2^{i+j}}{\delta} \right) \leq 2\delta. \quad (72)$$

For $t < \ell_{\min}$, we have

$$t\mu(M_f(t)) \leq t = t\mu_n(M_f(t)) \quad (73)$$

and for $t > \ell_{\max}$, recalling from Lemma 16 that $\|f\|_{\text{Lip}} = \ell_{\max} \leq 2\varepsilon^{-1}\tilde{\Lambda}_y(\mu_n, \Omega_n)$, we have

$$t\mu(M_f(t)) = 0 = t\mu_n(M_f(t)). \quad (74)$$

For any other $t > 0$, define $\tau(t)$ to be the largest $\tau_{ij} \leq t$. Then, for $t \in [\ell_{\min}, \ell_{\max}]$,

$$\begin{aligned} t\mu(M_f(t)) &\leq 2\tau(t)\mu(M_f(\tau(t))) \\ &\leq_p 4\tau(t)\mu_n(M_f(\tau(t)/4)) + 16\tau(t) \frac{m}{n^{1/3}} \log \frac{m \log(\ell_{\max}/\ell_{\min})}{\delta} \\ &\leq 16(\tau(t)/4)\mu_n(M_f(\tau(t)/4)) + 32\ell_{\max} \frac{m}{n^{1/3}} \log \frac{m \log(\ell_{\max}/\ell_{\min})}{\delta} \\ &\leq 16\tilde{\Lambda}_f(\mu_n, \Omega_n) + 64\varepsilon^{-1}\tilde{\Lambda}_y(\mu_n, \Omega_n) \frac{m}{n^{1/3}} \log \left(\frac{m \log(2/\varepsilon)}{\delta} \right), \end{aligned}$$

where \leq_p holds with probability at least $1 - \delta$ and Lemma 29 was invoked in the last inequality. The claim follows. ■

Deferred proofs. Proof [Lemma 28] Let Π be the partition defined in Lemma 5 and let $\Pi = \Pi_0 \cup \Pi_1$ be the dichotomy of Π into light and heavy cells as in the proof of Lemma 13. Put $U_f = U_0 \cup U_1$, where:

$$\begin{aligned} U_0 &= \{B \in \Pi_0 : B \cap M_f(t) \neq \emptyset\}, \\ U_1 &= \{B \in \Pi_1 : B \cap M_f(t) \neq \emptyset\}. \end{aligned}$$

Then

$$\begin{aligned}
 \mu(M_f(t)) &= \sum_{B \in \Pi_0} \mu(B \cap M_f(t)) + \sum_{B \in \Pi_1} \mu(B \cap M_f(t)) \\
 &\leq \sum_{B \in U_0} \mu(B) + \sum_{B \in U_1} \mu(B) \\
 &\leq_p (2\mu_n(\cup U_0) + 2q) + 2 \sum_{B \in U_1} \mu_n(B) \\
 &\leq 2\mu_n(\cup U_0) + 2\mu_n(\cup U_1) + 2q \\
 &= 2\mu_n(U_f) + 2q \\
 &\leq 2\mu_n(M_f(t/4)) + 2q,
 \end{aligned}$$

where \leq_p follows from (41, 43) and the final inequality from Lemma 5. Hence the bound holds with probability at least $1 - \left[m \exp\left(-\frac{nq}{4m}\right) + \exp\left(-\frac{m+nq^2}{2} + q\sqrt{mn}\right) \right]$. Choosing $q = (4m/n^{1/3}) \log(m/\delta)$, we have that (71) holds with probability at least

$$\begin{aligned}
 &1 - m \left(\frac{\delta}{m}\right)^{n^{2/3}} - \exp\left(-\frac{m + n(4m/n^{1/3})^2 \log(m/\delta)^2}{2} + (4m/n^{1/3}) \log(m/\delta) \sqrt{mn}\right) \\
 &\geq 1 - \delta - \exp(-m/2) \cdot \left(\frac{\delta}{m}\right)^{8m^2 \log(m/\delta) n^{1/3} - 4m^{3/2} \log(m/\delta) n^{1/6}} \\
 &= 1 - \delta - \exp(-m/2) \cdot \left(\frac{\delta}{m}\right)^{\log(m/\delta) n^{1/6} (8m^2 n^{1/6} - 4m^{3/2})} \\
 &\geq 1 - 2\delta.
 \end{aligned}$$

■

Proof [Lemma 29] By Corollary 24(iii), we may assume that for some $u, v \in V$,

$$\|f\|_{\text{Lip}} = \frac{|f(u) - f(v)|}{\rho(u, v)}.$$

Since $\rho(u, v) \geq \varepsilon$, we have, for any $x \in \Omega$,

$$\begin{aligned}
 \Lambda_f(x, \Omega) &\geq \max \left\{ \frac{|f(x) - f(u)|}{\rho(x, u)}, \frac{|f(x) - f(v)|}{\rho(x, v)} \right\} \\
 &\geq \max \left\{ \frac{|f(x) - f(u)|}{\text{diam}(\Omega)}, \frac{|f(x) - f(v)|}{\text{diam}(\Omega)} \right\} \\
 &\geq \frac{|f(u) - f(v)|}{2 \text{diam}(\Omega)} \\
 &= \|f\|_{\text{Lip}} \cdot \frac{\rho(u, v)}{2 \text{diam}(\Omega)} \geq \|f\|_{\text{Lip}} \cdot \frac{\varepsilon}{2 \text{diam}(\Omega)}.
 \end{aligned}$$

■

Appendix F. Chernoff-type bound

Theorem 32 For $X \sim \text{Binomial}(n, p)$ and $p = p(n) \geq 2 \log(n)/n$, $q = q(n) \leq p(n)/\log(n)$,

$$\mathbb{P}(X/n \leq q) \leq \left(\frac{e \log n}{n} \right)^2, \quad n \geq 3. \quad (75)$$

Proof

For all $0 < q < p < 1$ and $X \sim \text{Binomial}(n, p)$, (Dubhashi and Panconesi, 2009, page 4)

$$\mathbb{P}(X/n \leq q) \leq \left(\left(\frac{p}{q} \right)^q \left(\frac{1-p}{1-q} \right)^{1-q} \right)^n.$$

We first consider the case where $p(n) = 2 \log(n)/n$ and $q(n) = p(n)/\log(n) = 2/n$. In this case,

$$\begin{aligned} \mathbb{P}(X/n \leq q) \cdot \left(\frac{n}{\log n} \right)^2 &\leq \left(\left(\frac{p}{q} \right)^q \left(\frac{1-p}{1-q} \right)^{1-q} \right)^n \cdot \left(\frac{n}{\log n} \right)^2 \\ &= (\log n)^2 \cdot \left(\frac{n-2 \log n}{n-2} \right)^{n-2} \cdot \left(\frac{n}{\log n} \right)^2 \\ &= n^2 \left(\frac{n-2 \log n}{n-2} \right)^{n-2} =: a(n) \\ &\xrightarrow{n \rightarrow \infty} e^2. \end{aligned}$$

Furthermore, the sequence $a(n)$ is monotonically increasing for $n \geq 12$, and it is easily verified that $a(n) \leq e^2$ for all $n \geq 3$. This proves (75) for $p(n) = 2 \log(n)/n$, $q(n) = 2/n$.

To prove the full claim, consider the function

$$f(p, k) = k^{p/k} \left(\frac{1-p}{1-p/k} \right)^{1-p/k}, \quad p \in [0, 1], k \in [1, \infty).$$

We claim that (i) f is monotonically decreasing in each argument and (ii) this suffices to establish (75) in its full generality. To see how monotonicity implies the full claim, note that $\mathbb{P}(X/n \leq p/k) \leq f(p, k)^n$, the latter being maximized by the smallest feasible values of p and k . The conditions of the Theorem constrain these at $p \geq 2 \log(n)/n$ and $k \geq \log n \geq \log 3 > 1$, which reduces the problem to the case analyzed above.

To prove monotonicity in p , compute

$$\frac{\partial}{\partial p} \log f(p, k) = \frac{1-k + (1-p) \log k - (1-p) \log[(1-p)/(1-p/k)]}{k(1-p)}.$$

Since the denominator is clearly positive, it suffices to prove that the numerator is negative. Now $1-k + (1-p) \log k \leq 1-k + \log k < 0$ for $k > 1$, and $(1-p)/(1-p/k) > 1$, so the contribution of the remaining term is negative as well.

To prove monotonicity in k , compute

$$\frac{\partial}{\partial k} \log f(p, k) = \frac{p[-\log k + \log((1-p)/(1-p/k))]}{k^2},$$

whose negativity is equivalent to $k > (1 - p)/(1 - p/k)$, the latter a consequence of $k > 1$. ■

Acknowledgments

We thank Luigi Ambrosio and Ariel Elperin for very helpful feedback on earlier attempts to define a notion of average smoothness, and to Sasha Rakhlin for useful discussions. Pavel Shvartsman and Adam Oberman were very helpful in placing PMSE in proper historical context.

This research was partially supported by the Israel Science Foundation (grant No. 1602/19) and an Amazon Research Award. Additional support was provided by the Ariel Cyber Innovation Center in conjunction with the Israel Cyber directorate in the Prime Minister’s Office.

References

- Luigi Ambrosio and Roberta Ghezzi. Sobolev and bounded variation functions on metric measure spaces. In *Geometry, analysis and dynamics on sub-Riemannian manifolds. Vol. II*, EMS Ser. Lect. Math., pages 211–273. Eur. Math. Soc., Zürich, 2016.
- Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, Cambridge, 1999. ISBN 0-521-57353-X. doi: 10.1017/CBO9780511624216. URL <http://dx.doi.org/10.1017/CBO9780511624216>.
- Jürgen Appell, Józef Banaś, and Nelson Merentes. *Bounded variation and around*, volume 17 of *De Gruyter Series in Nonlinear Analysis and Applications*. De Gruyter, Berlin, 2014. ISBN 978-3-11-026507-1; 978-3-11-026511-8.
- Yair Ashlagi, Lee-Ad Gottlieb, and Aryeh Kontorovich. Functions with average smoothness: structure, algorithms, and learning. In Mikhail Belkin and Samory Kpotufe, editors, *Conference on Learning Theory, COLT 2021, 15-19 August 2021, Boulder, Colorado, USA*, volume 134 of *Proceedings of Machine Learning Research*, pages 186–236. PMLR, 2021. URL <http://proceedings.mlr.press/v134/ashlagi21a.html>.
- Peter Bartlett. Covering numbers and metric entropy (lecture 23). Class Notes on Statistical Learning Theory, <https://people.eecs.berkeley.edu/~bartlett/courses/281b-sp06/lecture23.ps>, 2006.
- Peter L. Bartlett, Sanjeev R. Kulkarni, and S. E. Posner. Covering numbers for real-valued function classes. *IEEE Trans. Information Theory*, 43(5):1721–1724, 1997. doi: 10.1109/18.623181. URL <https://doi.org/10.1109/18.623181>.
- Kinjal Basu and Art B. Owen. Transformations and Hardy-Krause variation. *SIAM J. Numer. Anal.*, 54(3):1946–1966, 2016. ISSN 0036-1429. doi: 10.1137/15M1052184. URL <https://doi.org/10.1137/15M1052184>.

- Gyora M. Benedek and Alon Itai. Learnability with respect to fixed distributions. *Theoretical Computer Science*, 86(2):377 – 389, 1991. ISSN 0304-3975. doi: 10.1016/0304-3975(91)90026-X. URL <http://www.sciencedirect.com/science/article/pii/030439759190026X>.
- Daniel Berend and Aryeh Kontorovich. A sharp estimate of the binomial mean absolute deviation with applications. *Statistics & Probability Letters*, 83(4):1254–1259, 2013.
- Kamalika Chaudhuri and Sanjoy Dasgupta. Rates of convergence for nearest neighbor classification. In *NIPS*, 2014.
- Eden Chlamtác. Private Communication, 2020.
- Eden Chlamtác, Michael Dinitz, Christian Konrad, Guy Kortsarz, and George Rabanca. The densest k-subhypergraph problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2016)*, 2016.
- Doron Cohen, Aryeh Kontorovich, and Geoffrey Wolfer. Learning discrete distributions with infinite support. In *Neural Information Processing Systems (NIPS)*, 2020.
- Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. In *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*, pages 938–942, 2019.
- Richard Cole and Lee-Ad Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *STOC*, pages 574–583, 2006.
- J Csirik, JBG Frenk, M Labbe, and S Zhang. Heuristics for the 0-1 min-knapsack problem. *Acta Cybernetica*, 10(1-2):15–20, 1991.
- David L. Donoho and Iain M. Johnstone. Minimax estimation via wavelet shrinkage. *Ann. Statist.*, 26(3):879–921, 1998. ISSN 0090-5364. doi: 10.1214/aos/1024691081. URL <https://doi.org/10.1214/aos/1024691081>.
- Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009. ISBN 978-0-521-88427-3. URL <http://www.cambridge.org/gb/knowledge/isbn/item2327542/>.
- Richard M. Dudley. *Uniform Central Limit Theorems*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1999. ISBN 9780521461023.
- Prerona Dutta and Khai T. Nguyen. Covering numbers for bounded variation functions. *J. Math. Anal. Appl.*, 468(2):1131–1143, 2018. ISSN 0022-247X. doi: 10.1016/j.jmaa.2018.08.062. URL <https://doi.org/10.1016/j.jmaa.2018.08.062>.
- Klim Efremenko, Aryeh Kontorovich, and Moshe Noivirt. Fast and bayes-consistent nearest neighbors. In *International Conference on Artificial Intelligence and Statistics, AISTATS*, 2020.
- Ariel Elperin and Aryeh Kontorovich. Bounded variation separates weak and strong average Lipschitz, in progress. 2021+.

- Evarist Giné and Joel Zinn. Some limit theorems for empirical processes. *Ann. Probab.*, 12(4):929–998, 1984. ISSN 0091-1798. With discussion.
- Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Efficient classification for metric data (extended abstract: COLT 2010). *IEEE Transactions on Information Theory*, 60(9):5750–5759, 2014. doi: 10.1109/TIT.2014.2339840. URL <http://dx.doi.org/10.1109/TIT.2014.2339840>.
- Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Adaptive metric dimensionality reduction (extended abstract: ALT 2013). *Theoretical Computer Science*, pages 105–118, 2016.
- Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Efficient regression in metric spaces via approximate Lipschitz extension (extended abstract: SIMBAD 2013). *IEEE Transactions on Information Theory*, 63(8):4838–4849, 2017.
- Lee-Ad Gottlieb, Aryeh Kontorovich, and Pinhas Nisnevitch. Near-optimal sample compression for nearest neighbors (extended abstract: NIPS 2014). *IEEE Trans. Information Theory*, 64(6):4120–4128, 2018. doi: 10.1109/TIT.2018.2822267. URL <https://doi.org/10.1109/TIT.2018.2822267>.
- Steve Hanneke, Aryeh Kontorovich, Sivan Sabato, and Roi Weiss. Universal Bayes consistency in metric spaces. *The Annals of Statistics*, 49(4):2129 – 2150, 2021. doi: 10.1214/20-AOS2029. URL <https://doi.org/10.1214/20-AOS2029>.
- Steve Hanneke, Aryeh Kontorovich, and Guy Kornowski. Efficient agnostic learning with average smoothness. In *Neural Information Processing Systems (NIPS)*, 2023.
- Steve Hanneke, Aryeh Kontorovich, and Guy Kornowski. Near-optimal learning with average hölder smoothness. In *Algorithmic Learning Theory, ALT 2024*, Proceedings of Machine Learning Research. PMLR, 2024.
- Sariel 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. doi: 10.1137/S0097539704446281. URL <http://link.aip.org/link/?SMJ/35/1148/1>.
- Juha Heinonen. *Lectures on analysis on metric spaces*. Universitext. Springer-Verlag, New York, 2001. ISBN 0-387-95104-0. doi: 10.1007/978-1-4613-0131-8. URL <http://dx.doi.org/10.1007/978-1-4613-0131-8>.
- Petri Juutinen. Absolutely minimizing Lipschitz extensions on a metric space. *Ann. Acad. Sci. Fenn. Math.*, 27(1):57–67, 2002. ISSN 1239-629X.
- Andreï N. Kolmogorov and Vladimir M. Tihomirov. ε -entropy and ε -capacity of sets in functional space. *Amer. Math. Soc. Transl. (2)*, 17:277–364, 1961.
- Aryeh Kontorovich and Iosif Pinelis. Exact lower bounds for the agnostic probably-approximately-correct (pac) machine learning model. *Ann. Statist.*, 47(5):2822–2854, 2019. ISSN 0090-5364. doi: 10.1214/18-AOS1766.

- Christos Koufogiannakis and Neal E. Young. A nearly linear-time ptas for explicit fractional packing and covering linear programs. *Algorithmica*, 70(4):648–674, 2014.
- Samory Kpotufe. k-NN regression adapts to local intrinsic dimension. In *Advances in Neural Information Processing Systems 24*, pages 729–737, 2011. URL http://books.nips.cc/papers/files/nips24/NIPS2011_0498.pdf.
- Samory Kpotufe and Sanjoy Dasgupta. A tree-based regressor that adapts to intrinsic dimension. *J. Comput. Syst. Sci.*, 78(5):1496–1515, September 2012. ISSN 0022-0000. doi: 10.1016/j.jcss.2012.01.002. URL <http://dx.doi.org/10.1016/j.jcss.2012.01.002>.
- Samory Kpotufe and Vikas K. Garg. Adaptivity to local smoothness and dimension in kernel regression. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 3075–3083, 2013. URL <http://papers.nips.cc/paper/5103-adaptivity-to-local-smoothness-and-dimension-in-kernel-regression>.
- Samory Kpotufe, Ruth Urner, and Shai Ben-David. Hierarchical label queries with data-dependent partitions. In Peter Grünwald, Elad Hazan, and Satyen Kale, editors, *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, volume 40 of *JMLR Workshop and Conference Proceedings*, pages 1176–1189. JMLR.org, 2015. URL <http://proceedings.mlr.press/v40/Kpotufe15.html>.
- 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.
- L. Kuipers and H. Niederreiter. *Uniform distribution of sequences*. Wiley-Interscience [John Wiley & Sons], New York-London-Sydney, 1974. Pure and Applied Mathematics.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces*. Springer-Verlag, 1991.
- Philip M. Long. Efficient algorithms for learning functions with bounded variation. *Inf. Comput.*, 188(1):99–115, 2004. doi: 10.1016/S0890-5401(03)00164-0. URL [https://doi.org/10.1016/S0890-5401\(03\)00164-0](https://doi.org/10.1016/S0890-5401(03)00164-0).
- Yu. V. Malykhin. Averaged modulus of continuity and bracket compactness. *Mat. Zametki*, 87(3):468–471, 2010. ISSN 0025-567X. doi: 10.1134/S0001434610030181. URL <https://doi.org/10.1134/S0001434610030181>.
- Enno Mammen and Alexandre B. Tsybakov. Smooth discrimination analysis. *Ann. Statist.*, 27(6):1808–1829, 12 1999. doi: 10.1214/aos/1017939240. URL <https://doi.org/10.1214/aos/1017939240>.
- Pasin Manurangsi. Almost-polynomial ratio eth-hardness of approximating densest k-subgraph. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 954–961, 2017.

- Edward James McShane. Extension of range of functions. *Bull. Amer. Math. Soc.*, 40(12): 837–842, 1934. ISSN 0002-9904. doi: 10.1090/S0002-9904-1934-05978-0. URL <http://dx.doi.org/10.1090/S0002-9904-1934-05978-0>.
- S. Mendelson and R. Vershynin. Entropy and the combinatorial dimension. *Invent. Math.*, 152(1):37–55, 2003. ISSN 0020-9910. doi: 10.1007/s00222-002-0266-3. URL <http://dx.doi.org/10.1007/s00222-002-0266-3>.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations Of Machine Learning*. The MIT Press, 2012.
- Richard Nickl and Benedikt M. Pötscher. Bracketing metric entropy rates and empirical central limit theorems for function classes of Besov- and Sobolev-type. *J. Theoret. Probab.*, 20(2):177–199, 2007. ISSN 0894-9840. doi: 10.1007/s10959-007-0058-1. URL <https://doi.org/10.1007/s10959-007-0058-1>.
- Harald Niederreiter and Denis Talay, editors. *Monte Carlo and quasi-Monte Carlo methods 2004*, 2006. Springer-Verlag, Berlin. ISBN 978-3-540-25541-3; 3-540-25541-9. doi: 10.1007/3-540-31186-6. URL <https://doi.org/10.1007/3-540-31186-6>.
- Adam M. Oberman. An explicit solution of the Lipschitz extension problem. *Proc. Amer. Math. Soc.*, 136(12):4329–4338, 2008. ISSN 0002-9939. doi: 10.1090/S0002-9939-08-09457-4. URL <https://doi.org/10.1090/S0002-9939-08-09457-4>.
- Yuval Peres, Oded Schramm, Scott Sheffield, and David B. Wilson. Tug-of-war and the infinity Laplacian. *J. Amer. Math. Soc.*, 22(1):167–210, 2009. ISSN 0894-0347. doi: 10.1090/S0894-0347-08-00606-1. URL <https://doi.org/10.1090/S0894-0347-08-00606-1>.
- Blagovest Sendov and Vasil A. Popov. *The averaged moduli of smoothness*. Pure and Applied Mathematics (New York). John Wiley & Sons, Ltd., Chichester, 1988. ISBN 0-471-91952-7. Applications in numerical methods and approximation, A Wiley-Interscience Publication.
- 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.
- Pavel Shvartsman. Whitney-type extension theorems for jets generated by Sobolev functions. *Adv. Math.*, 313:379–469, 2017. ISSN 0001-8708. doi: 10.1016/j.aim.2017.04.009. URL <https://doi.org/10.1016/j.aim.2017.04.009>.
- Alexandre B. Tsybakov. *Introduction à l'estimation non-paramétrique*, volume 41 of *Mathématiques & Applications (Berlin) [Mathematics & Applications]*. Springer-Verlag, Berlin, 2004. ISBN 3-540-40592-5.
- Roman Vershynin. *High-dimensional probability*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2018. ISBN 978-1-108-41519-4. doi: 10.1017/9781108231596. URL <https://doi.org/10.1017/9781108231596>. An introduction with applications in data science, With a foreword by Sara van de Geer.

Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint (Cambridge Series in Statistical and Probabilistic Mathematics)*. Cambridge University Press, 2019. ISBN 1108498027. URL <https://www.amazon.com/High-Dimensional-Statistics-Non-Asymptotic-Statistical-Probabilistic/dp/1108498027?SubscriptionId=AKIAIOBINVZYXZQZ2U3A&tag=chimbori05-20&linkCode=xm2&camp=2025&creative=165953&creativeASIN=1108498027>.

Hassler Whitney. Analytic extensions of differentiable functions defined in closed sets. *Transactions of the American Mathematical Society*, 36(1):63–89, 1934. ISSN 00029947. URL <http://www.jstor.org/stable/1989708>.