

Efficient Classification for Metric Data

Lee-Ad Gottlieb, Aryeh Kontorovich and Robert Krauthgamer

Abstract—Recent advances in large-margin classification of data residing in general metric spaces (rather than Hilbert spaces) enable classification under various natural metrics, such as string edit and earthmover distance. A general framework developed for this purpose by von Luxburg and Bousquet [JMLR, 2004] left open the questions of computational efficiency and of providing direct bounds on generalization error.

We design a new algorithm for classification in general metric spaces, whose runtime and accuracy depend on the doubling dimension of the data points, and can thus achieve superior classification performance in many common scenarios. The algorithmic core of our approach is an approximate (rather than exact) solution to the classical problems of Lipschitz extension and of Nearest Neighbor Search. The algorithm’s generalization performance is guaranteed via the fat-shattering dimension of Lipschitz classifiers, and we present experimental evidence of its superiority to some common kernel methods. As a by-product, we offer a new perspective on the nearest neighbor classifier, which yields significantly sharper risk asymptotics than the classic analysis of Cover and Hart [IEEE Trans. Info. Theory, 1967].

I. INTRODUCTION

A recent line of work extends the large-margin classification paradigm from Hilbert spaces to less structured ones, such as Banach or even metric spaces, see e.g. [23], [34], [13], [40]. In this metric approach, data is presented as points with distances but lacking the additional structure of inner-products. The potentially significant advantage is that the metric can be precisely suited to the type of data, e.g. earthmover distance for images, or edit distance for sequences.

However, much of the existing machinery of classification algorithms and generalization bounds,

An extended abstract of this work appeared in Proceedings of the 23rd COLT, 2010 [16].

L. Gottlieb is with the Department of Computer Science and Mathematics at Ariel University (email: leead@ariel.ac.il).

A. Kontorovich is with the Department of Computer Science at Ben-Gurion University of the Negev (email: karyeh@cs.bgu.ac.il). His work was supported in part by the Israel Science Foundation (grant No. 1141/12) and a Yahoo Faculty award.

R. Krauthgamer is with the Faculty of Mathematics and Computer Science at the Weizmann Institute of Science (email: robert.krauthgamer@weizmann.ac.il). His work was supported in part by the Israel Science Foundation (grant #452/08), the US-Israel BSF (grant #2010418), and by a Minerva grant.

(e.g. [11], [32]) depends strongly on the data residing in a Hilbert space. This structural requirement severely limits this machinery’s applicability — many natural metric spaces cannot be represented in a Hilbert space faithfully; formally, every embedding into a Hilbert space of metrics such as ℓ_1 , earthmover, and edit distance must distort distances by a large factor [14], [29], [2]. Ad-hoc solutions such as kernelization cannot circumvent this shortcoming, because imposing an inner-product obviously embeds the data in some Hilbert space.

To address this gap, von Luxburg and Bousquet [34] developed a powerful framework of large-margin classification for a general metric space \mathcal{X} . They first show that the natural hypotheses (classifiers) to consider in this context are maximally-smooth Lipschitz functions; indeed, they reduce classification (of points in a metric space \mathcal{X}) with no training error to finding a Lipschitz function $f : \mathcal{X} \rightarrow \mathbb{R}$ consistent with the data, which is a classic problem in Analysis, known as *Lipschitz extension*. Next, they establish error bounds in the form of expected surrogate loss. Finally, the computational problem of evaluating the classification function is reduced, assuming zero training error, to exact nearest neighbor search. This matches a popular classification heuristic, and in retrospect provides a rigorous explanation for this heuristic’s empirical success in general metric spaces, extending the seminal analysis of Cover and Hart [12] for the Euclidean case.

The work of [34] has left open some algorithmic questions. In particular, allowing nonzero training error is apt to significantly reduce the Lipschitz constant, thereby producing classifiers that have lower complexity and are less likely to overfit. This introduces the algorithmic challenge of constructing a Lipschitz classifier that minimizes the 0-1 training error. In addition, exact nearest neighbor search in general metrics has time complexity proportional to the size of the dataset, rendering the technique impractical when the training sample is large. Finally, bounds on the expected surrogate loss may significantly overestimate the generalization error, which is the true quantity of interest.

Our contribution: We solve the problems delineated above by showing that data residing in a metric space of low doubling dimension admits accurate and computationally efficient classification. This is the first result that ties the doubling dimension of the *data* to either classification error or algorithmic runtime.¹ Specifically, we (i) prove generalization bounds for the classification (0-1) error as opposed to surrogate loss, (ii) construct and evaluate the classifier in a computationally-efficient manner, and (iii) perform efficient structural risk minimization by optimizing the tradeoff between the classifier’s smoothness and its training error.

Our generalization bound for Lipschitz classifiers controls the expected classification error directly (rather than expected surrogate loss), and may be significantly sharper than the latter in many common scenarios. We provide this bound in Section III, using an elementary analysis of the fat-shattering dimension. In hindsight, our approach offers a new perspective on the nearest neighbor classifier, with significantly tighter risk asymptotics than the classic analysis of Cover and Hart [12].

We further give efficient algorithms to implement the Lipschitz classifier, both for the training and the evaluation stages. In Section IV we prove that once a Lipschitz classifier has been chosen, the hypothesis can be evaluated quickly on any new point $x \in \mathcal{X}$ using *approximate* nearest neighbor search, which is known to be fast when points have a low doubling dimension. In Section V we further show how to quickly compute a near-optimal classifier (in terms of classification error bound), even when the training error is nonzero. In particular, this necessitates the optimization of the number of incorrectly labeled examples — and moreover, their identity — as part of the structural risk minimization.

Finally, we give in Section VI two exemplary setups. In the first, the data is represented using the earthmover metric over the plane. In the second, the data is a set of time series vectors equipped with a popular distance function. We provide basic theoretical and experimental analysis, which illustrate the potential power of our approach.

II. DEFINITIONS AND NOTATION

Notation: We will use standard $O(\cdot), \Omega(\cdot)$ notation for orders of magnitude. If $f = O(g)$ and $g = O(f)$, we will write $f = \Theta(g)$. Whenever $f =$

¹Previously, the doubling dimension of the *space of classifiers* was used in [8], but this is less relevant to our discussion.

$O(n \text{ polylog } n)$, we will denote this by $f = \tilde{O}(n)$. If $n \in \mathbb{N}$ is a natural number $[n]$ denotes the set $\{1, \dots, n\}$.

Metric spaces: A *metric* ρ on a set \mathcal{X} is a positive symmetric function satisfying the triangle inequality $\rho(x, y) \leq \rho(x, z) + \rho(z, y)$; together the two comprise the metric space (\mathcal{X}, ρ) . The diameter of a set $A \subseteq \mathcal{X}$, is defined by $\text{diam}(A) = \sup_{x, y \in A} \rho(x, y)$ and the distance between two sets $A, B \subset \mathcal{X}$ is defined by $\rho(A, B) = \inf_{x \in A, y \in B} \rho(x, y)$. The Lipschitz constant of a function $f : \mathcal{X} \rightarrow \mathbb{R}$, denoted by $\|f\|_{\text{Lip}}$, is defined to be the smallest $L > 0$ that satisfies $|f(x) - f(y)| \leq L\rho(x, y)$ for all $x, y \in \mathcal{X}$.

Doubling dimension: For a metric space (\mathcal{X}, ρ) , let λ be the smallest value such that every ball in \mathcal{X} can be covered by λ balls of half the radius. λ is the *doubling constant* of \mathcal{X} , the *doubling dimension* of \mathcal{X} is $\text{ddim}(\mathcal{X}) = \log_2 \lambda$. A metric is *doubling* when its doubling dimension is bounded. Note that while a low Euclidean dimension implies a low doubling dimension (Euclidean metrics of dimension d have doubling dimension $\Theta(d)$ [21]), low doubling dimension is strictly more general than low Euclidean dimension.

The following packing property can be demonstrated via repeated applications of the doubling property (see, for example [25]):

Lemma 1. *Let \mathcal{X} be a metric space, and suppose that $S \subset \mathcal{X}$ is finite and has a minimum interpoint distance at least $\alpha > 0$. Then the cardinality of S is*

$$|S| \leq \left(\frac{2\text{diam}(S)}{\alpha} \right)^{\text{ddim}(\mathcal{X})}.$$

Nets: Let (\mathcal{X}, ρ) be a metric space and suppose $S \subset \mathcal{X}$. An ε -*net* of S is a subset $T \subset S$ with the following properties: (i) *Packing:* all distinct $u, v \in T$ satisfy $\rho(u, v) \geq \varepsilon$, which means that T is ε -separated; and (ii) *Covering:* every point $u \in S$ is strictly within distance ε of some point $v \in T$, namely $\rho(y, x) < \varepsilon$.

Learning: Our setting in this paper is the *agnostic PAC* learning model [27]. Examples are drawn independently from $\mathcal{X} \times \{-1, 1\}$ according to some unknown probability distribution \mathbb{P} and the learner, having observed n such pairs (x, y) produces a hypothesis $h : \mathcal{X} \rightarrow \{-1, 1\}$. The *generalization error* is the probability of misclassifying a new point drawn from \mathbb{P} :

$$\mathbb{P} \{(x, y) : h(x) \neq y\}.$$

The quantity above is random, since it depends on the n observations, and we wish to upper-bound it in probability. Most bounds of this sort contain a *training error* term, which is the fraction of observed examples misclassified by h and roughly corresponding to bias in Statistics, as well as a *hypothesis complexity* term, which measures the richness of the class of all admissible hypotheses [35], and roughly corresponding to variance in Statistics. Optimizing the tradeoff between these two terms is known as Structural Risk Minimization (SRM).² Keeping in line with the literature, we ignore the measure-theoretic technicalities associated with taking suprema over uncountable function classes.

III. GENERALIZATION BOUNDS

In this section, we derive generalization bounds for Lipschitz classifiers over doubling spaces. As noted by [34] Lipschitz functions are the natural object to consider in an optimization/regularization framework. The basic intuition behind our proofs is that the Lipschitz constant plays the role of the inverse margin in the confidence of the classifier. As in [34], small Lipschitz constant corresponds to large margin, which in turn yields low hypothesis complexity and variance. However, in contrast to [34] (whose generalization bounds rely on Rademacher averages) we use the doubling property of the metric space directly to control the fat-shattering dimension.

We apply tools from generalized Vapnik-Chervonenkis theory to the case of Lipschitz classifiers. Let \mathcal{F} be a collection of functions $f : \mathcal{X} \rightarrow \mathbb{R}$ and recall the definition of the fat-shattering dimension [1], [4]: a set $X \subset \mathcal{X}$ is said to be γ -shattered by \mathcal{F} if there exists some function $r : X \rightarrow \mathbb{R}$ such that for each label assignment $y \in \{-1, 1\}^X$ there is an $f \in \mathcal{F}$ satisfying $y(x)(f(x) - r(x)) \geq \gamma > 0$ for all $x \in X$. The γ -fat-shattering dimension of \mathcal{F} , denoted by $\text{fat}_\gamma(\mathcal{F})$, is the cardinality of the largest set γ -shattered by \mathcal{F} .

For the case of Lipschitz functions, we will show that the notion of fat-shattering dimension may be somewhat simplified. We say that a set $X \subset \mathcal{X}$ is γ -shattered *at zero* by a collection of functions \mathcal{F} if for each $y \in \{-1, 1\}^X$ there is an $f \in \mathcal{F}$ satisfying

$y(x)f(x) \geq \gamma$ for all $x \in X$. (This is the definition above with $r \equiv 0$.) We write $\text{fat}_\gamma^0(\mathcal{F})$ to denote the cardinality of the largest set γ -shattered at zero by \mathcal{F} and show that for Lipschitz function classes the two notions are the same.

Lemma 2. *Let \mathcal{F} be the collection of all $f : \mathcal{X} \rightarrow \mathbb{R}$ with $\|f\|_{\text{Lip}} \leq L$. Then $\text{fat}_\gamma(\mathcal{F}) = \text{fat}_\gamma^0(\mathcal{F})$.*

Proof: We begin by recalling the classic Lipschitz extension result, essentially due to [26] and [36]. Any real-valued function f defined on a subset X of a metric space \mathcal{X} has an extension f^* to all of \mathcal{X} satisfying $\|f^*\|_{\text{Lip}} = \|f\|_{\text{Lip}}$. Thus, in what follows we will assume that any function f defined on $X \subset \mathcal{X}$ is also defined on all of \mathcal{X} via some Lipschitz extension (in particular, to bound $\|f\|_{\text{Lip}}$ it suffices to bound the restricted $\|f|_X\|_{\text{Lip}}$).

Consider some finite $X \subset \mathcal{X}$. If X is γ -shattered at zero by \mathcal{F} then by definition it is also γ -shattered. Now assume that X is γ -shattered by \mathcal{F} . Thus, there is some function $r : X \rightarrow \mathbb{R}$ such that for each $y \in \{-1, 1\}^X$ there is an $f = f_{r,y} \in \mathcal{F}$ such that $f_{r,y}(x) \geq r(x) + \gamma$ if $y(x) = +1$ and $f_{r,y}(x) \leq r(x) - \gamma$ if $y(x) = -1$. Let us define the function \tilde{f}_y on X and as per above, on all of \mathcal{X} , by $\tilde{f}_y(x) = \gamma y(x)$. It is clear that the collection $\{\tilde{f}_y : y \in \{-1, 1\}^X\}$ γ -fat-shatters X at zero; it only remains to verify that $\tilde{f}_y \in \mathcal{F}$, i.e.,

$$\sup_{y \in \{-1, 1\}^X} \|\tilde{f}_y\|_{\text{Lip}} \leq \sup_{y \in \{-1, 1\}^X} \|f_{r,y}\|_{\text{Lip}}.$$

Indeed,

$$\begin{aligned} & \sup_{y \in \{-1, 1\}^X, x, x' \in X} \frac{f_{r,y}(x) - f_{r,y}(x')}{\rho(x, x')} \\ & \geq \sup_{x, x' \in X} \frac{r(x) - r(x') + 2\gamma}{\rho(x, x')} \\ & \geq \sup_{x, x' \in X} \frac{2\gamma}{\rho(x, x')} \\ & = \sup_{y \in \{-1, 1\}^X} \|\tilde{f}_y\|_{\text{Lip}}. \end{aligned}$$

A consequence of Lemma 2 is that in considering the generalization properties of Lipschitz functions we need only bound the γ -fat-shattering dimension at zero. The latter is achieved by observing that the packing number of a metric space controls the fat-shattering dimension of Lipschitz functions defined over the metric space: ■

² Robert Schapire pointed out to us that these terms from Statistics are not entirely accurate in the machine learning setting. In particular, the classifier complexity term does not correspond to the variance of the classifier in any quantitatively precise way. However, the intuition underlying SRM corresponds precisely to the one behind bias-variance tradeoff in Statistics, and so we shall occasionally use the latter term as well.

Theorem 3. Let (\mathcal{X}, ρ) be a metric space. Fix some $L > 0$, and let \mathcal{F} be the collection of all $f : \mathcal{X} \rightarrow \mathbb{R}$ with $\|f\|_{\text{Lip}} \leq L$. Then for all $\gamma > 0$,

$$\text{fat}_\gamma(\mathcal{F}) = \text{fat}_\gamma^0(\mathcal{F}) \leq M(\mathcal{X}, \rho, 2\gamma/L)$$

where $M(\mathcal{X}, \rho, \varepsilon)$ is the ε -packing number of \mathcal{X} , defined as the cardinality of the largest ε -separated subset of \mathcal{X} .

Proof: Suppose that $S \subseteq \mathcal{X}$ is fat γ -shattered at zero. The case $|S| = 1$ is trivial, so we assume the existence of $x \neq x' \in S$ and $f \in \mathcal{F}$ such that $f(x) \geq \gamma > -\gamma \geq f(x')$. The Lipschitz property then implies that $\rho(x, x') \geq 2\gamma/L$, and the claim follows. ■

Corollary 4. Let metric space \mathcal{X} have doubling dimension $\text{ddim}(\mathcal{X})$, and let \mathcal{F} be the collection of real-valued functions over \mathcal{X} with Lipschitz constant at most L . Then for all $\gamma > 0$,

$$\text{fat}_\gamma(\mathcal{F}) \leq \left(\frac{L \text{diam}(\mathcal{X})}{\gamma} \right)^{\text{ddim}(\mathcal{X})}.$$

Proof: The claim follows immediately from Theorem 3 and the packing property of doubling spaces (Lemma 1). ■

Equipped with these estimates for the fat-shattering dimension of Lipschitz classifiers, we can invoke a standard generalization bound stated in terms of this quantity. For the remainder of this section, we take $\gamma = 1$ and say that a function f classifies an example (x_i, y_i) correctly if

$$y_i f(x_i) \geq 1. \quad (1)$$

The following generalization bounds appear in [4].

Theorem 5. Let \mathcal{F} be a collection of real-valued functions over some set \mathcal{X} , define $D = \text{fat}_{1/16}(\mathcal{F})$ and let P be some probability distribution on $\mathcal{X} \times \{-1, 1\}$. Suppose that (x_i, y_i) , $i = 1, \dots, n$ are drawn from $\mathcal{X} \times \{-1, 1\}$ independently according to P and that some $f \in \mathcal{F}$ classifies the n examples correctly, in the sense of (1). Then with probability at least $1 - \delta$,

$$\mathbb{P} \{ (x, y) : \text{sgn}(f(x)) \neq y \} \leq \frac{2}{n} (D \log_2(34en/D) \log_2(578n) + \log_2(4/\delta)).$$

Furthermore, if $f \in \mathcal{F}$ is correct on all but k examples, we have with probability at least $1 - \delta$

$$\mathbb{P} \{ (x, y) : \text{sgn}(f(x)) \neq y \} \leq \frac{k}{n} + \sqrt{\frac{2}{n} (D \ln(34en/D) \log_2(578n) + \ln(4/\delta))}.$$

Applying Corollary 4, we obtain the following consequence of Theorem 5.

Corollary 6. Let metric space \mathcal{X} have doubling dimension $\text{ddim}(\mathcal{X})$, and let \mathcal{F} be the collection of real-valued functions over \mathcal{X} with Lipschitz constant at most L . Then for any $f \in \mathcal{F}$ that classifies a sample of size n correctly, we have with probability at least $1 - \delta$

$$\mathbb{P} \{ (x, y) : \text{sgn}(f(x)) \neq y \} \leq \frac{2}{n} (D \log_2(34en/D) \log_2(578n) + \log_2(4/\delta)). \quad (2)$$

Likewise, if f is correct on all but k examples, we have with probability at least $1 - \delta$

$$\mathbb{P} \{ (x, y) : \text{sgn}(f(x)) \neq y \} \leq \frac{k}{n} + \sqrt{\frac{2}{n} (D \ln(34en/D) \log_2(578n) + \ln(4/\delta))}. \quad (3)$$

In both cases, $D = \text{fat}_{1/16}(\mathcal{F}) \leq (16L \text{diam}(\mathcal{X}))^{\text{ddim}(\mathcal{X})}$.

A. Comparison with previous generalization bounds

Our generalization bounds are not directly comparable to those of von Luxburg and Bousquet [34]. In general, two approaches exist to analyze binary classification by continuous-valued functions: thresholding by the sign function or bounding some expected surrogate loss function. They opt for the latter approach, defining the surrogate loss function

$$\ell(f(x), y) = \min(\max(0, 1 - yf(x)), 1)$$

and bound the risk $\mathbb{E}[\ell(f(x), y)]$. We take the former approach, bounding the generalization error $\mathbb{P}(\text{sgn}(f(x)) \neq y)$ directly. Although for $\{-1, 1\}$ -valued labels the risk upper-bounds the generalization error, it could potentially be a crude overestimate.

von Luxburg and Bousquet [34] demonstrated that the Rademacher average of Lipschitz functions over the p -dimensional unit cube ($p \geq 3$) is of order $\Theta(n^{-1/p})$, and since the proof uses only covering numbers, a similar bound holds for all metric spaces with bounded diameter and doubling dimension. In conjunction with Theorem 5(b) of [5], this observation yields the following bound.

Lemma 7. Let \mathcal{X} be a metric space with $\text{diam}(\mathcal{X}) \leq 1$, and let \mathcal{F} be the collection of all $f : \mathcal{X} \rightarrow \mathbb{R}$ with $\|f\|_{\text{Lip}} \leq 1$. If $(x_i, y_i) \in \mathcal{X} \times \{-1, 1\}$ are drawn iid with respect to some probability distribution P , then

with probability at least $1 - \delta$ every $f \in \mathcal{F}$ satisfies

$$\mathbb{P}\{(x, y) : f(x) \neq y\} \leq O\left(k_f/n + n^{-1/\text{ddim}(\mathcal{X})} + \sqrt{\ln(1/\delta)/n}\right),$$

where k_f is the number of examples f labels incorrectly.

Our results compare favorably to those of [34] when we assume fixed diameter $\text{diam}(\mathcal{X}) \leq 1$ and Lipschitz constant $L \leq 1$ and the number of observations n goes to infinity. Indeed, Lemma 7 bounds the excess error decay by $O(n^{-1/\text{ddim}(\mathcal{X})})$, whereas Corollary 6 gives a rate of $\tilde{O}(n^{-1/2})$.

B. Comparison with previous nearest-neighbor bounds

Corollary 6 also allows us to significantly sharpen the asymptotic analysis of [12] for the nearest-neighbor classifier. Following the presentation in [33] with an appropriate generalization to general metric spaces, the analysis of [12] implies that the 1-nearest-neighbor classifier h_{NN} achieves

$$\mathbb{E}[\text{err}(h_{\text{NN}})] \leq 2 \text{err}(h^*) + O\left(\frac{\|\eta\|_{\text{Lip}} \text{diam}(\mathcal{X})}{n^{1/(\text{ddim}(\mathcal{X})+1)}}\right), \quad (4)$$

where $\eta(x) = \mathbb{P}(Y = 1 | X = x)$ is the conditional probability of the 1 label, and $h^*(x) = \text{sgn}(\eta(x) - 1/2)$ is the Bayes optimal classifier. The curse of dimensionality exhibited in the term $n^{1/(\text{ddim}(\mathcal{X})+1)}$ is real — for each n , there exists a distribution such that for sample size $n \ll (L + 1)^{\text{ddim}(\mathcal{X})}$, we have $\mathbb{E}[\text{err}(h_{\text{NN}})] \geq \Omega(1)$. However, Corollary 6 shows that this analysis is overly pessimistic. Comparing (4) with (2) in the case where $\text{err}(h^*) = 0$, we see that once the sample size passes a critical number on the order of $(L \text{diam}(\mathcal{X}))^{\text{ddim}(\mathcal{X})}$, the expected generalization error begins to decay as $\tilde{O}(1/n)$, which is much faster than the rate suggested by (4).

IV. LIPSCHITZ EXTENSION CLASSIFIER

Given n labeled points $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \{-1, 1\}$, we construct our classifier in a similar manner to [34], via a Lipschitz extension of the label values y_i to all of \mathcal{X} . Let $S^+, S^- \subset \{x_1, \dots, x_n\}$ be the sets of positive and negative labeled points. Our starting point is the same extension function used in [34], namely, for $\alpha \in [0, 1]$ define $f_\alpha : \mathcal{X} \rightarrow \mathbb{R}$ by

$$f_\alpha(x) = \alpha \min_{i \in [n]} \left(y_i + 2 \frac{\rho(x, x_i)}{\rho(S^+, S^-)} \right) + (1 - \alpha) \max_{j \in [n]} \left(y_j - 2 \frac{\rho(x, x_j)}{\rho(S^+, S^-)} \right). \quad (5)$$

It is easy to verify, see also [34, Lemmas 7 and 12], that $f_\alpha(x_i)$ agrees with the sample label y_i for all $i \in [n]$, and that its Lipschitz constant is identical to the one induced by the labeled points, which in turn is obviously $2/\rho(S^+, S^-)$. However, computing the exact value of $f_\alpha(x)$ for a point $x \in \mathcal{X}$ (or even the sign of $f_\alpha(x)$ at this point) requires an exact nearest neighbor search, and in arbitrary metric space nearest neighbor search requires $\Omega(n)$ time.

In this section, we design a classifier that is evaluated at a point $x \in \mathcal{X}$ using an approximate nearest neighbor search.³ It is known how to build a data structure for a set of n points in time $2^{O(\text{ddim}(\mathcal{X}))} n \log n$, so as to support $(1 + \varepsilon)$ -approximate nearest neighbor searches in time $2^{O(\text{ddim}(\mathcal{X}))} \log n + \varepsilon^{-O(\text{ddim}(\mathcal{X}))}$ [10], [22] (see also [25], [7]). Our classifier below relies only on a given subset of the given n points, which may eventually lead to improved generalization bounds (i.e., it provides a tradeoff between k and L in Theorem 5).

Theorem 8. *Let (\mathcal{X}, ρ) be a metric space, and fix $0 < \varepsilon < \frac{1}{32}$. Let S be a sample consisting of n labeled points $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \{-1, 1\}$. Fix a subset $S_1 \subset S$ of cardinality $n - k$, on which the constructed classifier must agree with the given labels, and partition it into $S_1^+, S_1^- \subset S_1$ according to the labels, letting $L = 2/\rho(S_1^+, S_1^-)$. Then there is a binary classification function $h : \mathcal{X} \rightarrow \{-1, 1\}$ satisfying:*

- (a) $h(x)$ can be evaluated at each $x \in \mathcal{X}$ in time $2^{O(\text{ddim}(\mathcal{X}))} \log n + \varepsilon^{-O(\text{ddim}(\mathcal{X}))}$, after an initial computation of $(2^{O(\text{ddim}(\mathcal{X}))} \log n + \varepsilon^{-O(\text{ddim}(\mathcal{X}))})n$ time.
- (b) With probability at least $1 - \delta$ (over the sampling of S)

$$\mathbb{P}\{(x, y) : h(x) \neq y\} \leq \frac{k}{n} + \sqrt{\frac{2}{n} (D \ln(34en/D) \log_2(578n) + \ln(4/\delta))},$$

where

$$D = \left(\frac{16L \text{diam}(\mathcal{X})}{1 - 32\varepsilon} \right)^{\text{ddim}(\mathcal{X})}.$$

We will use the following simple lemma.

Lemma 9. *For any function class \mathcal{F} mapping \mathcal{X} to \mathbb{R} , define its ε -perturbation \mathcal{F}_ε to be*

$$\mathcal{F}_\varepsilon = \left\{ \tilde{f} \in \mathbb{R}^{\mathcal{X}} : \|f - \tilde{f}\|_\infty \leq \varepsilon, f \in \mathcal{F} \right\},$$

³If x^* is the nearest neighbor for a test point x , then any point \tilde{x} satisfying $\rho(x, \tilde{x}) \leq (1 + \varepsilon)\rho(x, x^*)$ is called a $(1 + \varepsilon)$ -approximate nearest neighbor of x .

where $\|f - \tilde{f}\|_\infty = \sup_{x \in \mathcal{X}} |f(x) - \tilde{f}(x)|$. Then for $0 < \varepsilon < \gamma$,

$$\text{fat}_\gamma(\mathcal{F}_\varepsilon) \leq \text{fat}_{\gamma-\varepsilon}(\mathcal{F}).$$

Proof: Suppose that \mathcal{F}_ε is able to γ -shatter the finite subset $X \subset \mathcal{X}$. Then there is an $r \in \mathbb{R}^X$ so that for all $y \in \{-1, 1\}^X$, there is an $\tilde{f}_y \in \mathcal{F}_\varepsilon$ such that

$$y(x)(\tilde{f}_y(x) - r(x)) \geq \gamma, \quad \forall x \in X. \quad (6)$$

Now by definition, for each $\tilde{f}_y \in \mathcal{F}_\varepsilon$ there is some $f_y \in \mathcal{F}$ such that $\sup_{x \in \mathcal{X}} |f_y(x) - \tilde{f}_y(x)| \leq \varepsilon$. We claim that the collection $\{f_y : y \in \{-1, 1\}^X\}$ is able to $(\gamma - \varepsilon)$ -shatter X . Indeed, replacing $\tilde{f}_y(x)$ with $f_y(x)$ in (6) perturbs the left-hand side by an additive term of at most ε . ■

Proof of Theorem 8: Without loss of generality, assume $S_1 \subset S$ corresponds to points indexed by $i = 1, \dots, n - k$. We begin by observing that since all of the sample labels have values in $\{\pm 1\}$, any Lipschitz extension may be truncated to the range $[-1, 1]$. Formally, if g is a Lipschitz extension of the labels y_i from the sample S to all of \mathcal{X} , then so is $\text{T}_{[-1,1]} \circ g$, where

$$\text{T}_{[a,b]}(z) = \max\{a, \min\{b, z\}\}$$

is the truncation operator. In particular, take g to be as in (5) with $\alpha = 1$ and write

$$r_i(x) = 2 \frac{\rho(x, x_i)}{\rho(S_1^+, S_1^-)}.$$

Now defining

$$\begin{aligned} f(x) &= \text{T}_{[-1,1]} \left(\min_{i \in [n-k]} \{y_i + r_i(x)\} \right) \\ &= \min_{i \in [n-k]} \{ \text{T}_{[-1,1]}(y_i + r_i(x)) \}, \end{aligned} \quad (7)$$

where the second equality is by monotonicity of the truncation operator, we conclude that f is a Lipschitz extension of the data, with the same Lipschitz constant $L = 2/\rho(S_1^+, S_1^-)$.

Now precompute⁴ in time $2^{O(\text{ddim}(\mathcal{X}))} n \log n$ a data structure that supports $(1 + \varepsilon)$ -approximate nearest neighbor searches on the point set S_1^+ , and a similar one for the point set S_1^- . Now compute (still during the learning phase) an estimate $\tilde{\rho}(S_1^+, S_1^-)$ for $\rho(S_1^+, S_1^-)$, by searching the second data structure for

⁴The word *precompute* underscores the fact that this computation is done during the “offline” learning phase. Its result is then used to achieve fast “online” evaluation of the classifier on any point $x \in \mathcal{X}$ during the testing phase.

each of the points in S_1^+ , and taking the minimum of all the resulting distances. This estimate satisfies

$$1 \leq \frac{\tilde{\rho}(S_1^+, S_1^-)}{\rho(S_1^+, S_1^-)} \leq 1 + \varepsilon, \quad (8)$$

and this entire precomputation process takes $(2^{O(\text{ddim}(\mathcal{X}))} \log n + \varepsilon^{-O(\text{ddim}(\mathcal{X}))})n$ time.

Given a test point $x \in \mathcal{X}$ to be classified, search for x in the two data structures (for S_1^+ and for S_1^-), and denote the indices of the points answered by them by $a^+, a^- \in [n - k]$, respectively. The $(1 + \varepsilon)$ -approximation guarantee means that

$$1 \leq \frac{\rho(x, x_{a^+})}{\rho(x, S_1^+)} \leq 1 + \varepsilon, \quad \text{and} \quad 1 \leq \frac{\rho(x, x_{a^-})}{\rho(x, S_1^-)} \leq 1 + \varepsilon.$$

Define, as a computationally-efficient estimate of f , the function

$$\tilde{f}(x) = \min_{a \in \{a^+, a^-\}} \left\{ \text{T}_{[-1,1]} \left(y_a + 2 \frac{\rho(x, x_a)}{\tilde{\rho}(S_1^+, S_1^-)} \right) \right\},$$

and let our classifier be $h(x) = \text{sgn}(\tilde{f}(x))$. We remark that the case $a = a^-$ always attains the minimum in the definition of \tilde{f} (because $a = a^+$ only produces values greater or equal than $y_{a^+} = 1$), and therefore one can avoid the computation of a^+ , and even the construction of a data structure for S_1^+ . In fact, the same argument shows that also in the definition of f in (7) we can omit from the minimization points with label $y_i = +1$.

This classifier $h = \text{sgn}(\tilde{f})$ can be evaluated on a new point $x \in \mathcal{X}$ in time $2^{O(\text{ddim}(\mathcal{X}))} \log n + \varepsilon^{-O(\text{ddim}(\mathcal{X}))}$, and it thus remains to bound the generalization error of h . To this end, we will show that

$$\sup_{x \in \mathcal{X}} |f(x) - \tilde{f}(x)| \leq 2\varepsilon. \quad (9)$$

This means that \tilde{f} is a 2ε -perturbation of f , as stipulated by Lemma 9, and the generalization error of h will follow from Corollary 6 using the fact that f has Lipschitz constant L .

To prove (9), fix an $x \in \mathcal{X}$. Now let $i^* \in [n - k]$ be an index attaining the minimum in the definition of $f(x)$ in (7), and similarly $a^* \in [n - k]$ for $\tilde{f}(x)$. Using the remark above, we may assume that their labels are $y_{i^*} = y_{a^*} = -1$. Moreover, by inspecting the definition of f we may further assume that i^* attains the minimum of $r_i(x)$ (over all points labeled -1) and thus also of its numerator $\rho(x, x_i)$. And since index a^* was chosen as an approximate nearest neighbor (among all points labeled -1), we get $1 \leq \frac{\rho(x, x_{a^*})}{\rho(x, x_{i^*})} \leq 1 + \varepsilon$. Together with (8), we have

$$\frac{1}{1 + \varepsilon} r_{i^*}(x) \leq 2 \frac{\rho(x, x_{a^*})}{\tilde{\rho}(S_1^+, S_1^-)} \leq (1 + \varepsilon) r_{i^*}(x). \quad (10)$$

We now need the following simple claim:

$$0 \leq B \leq (1 + \varepsilon)C \implies \\ \mathbb{T}_{[-1,1]}(-1 + B) \leq 2\varepsilon + \mathbb{T}_{[-1,1]}(-1 + C).$$

To verify the claim, assume first that $C \leq 2$; then $B \leq C + \varepsilon C \leq C + 2\varepsilon$, and now use the fact that adding -1 and truncating are both monotone operations, to get $\mathbb{T}_{[-1,1]}(-1 + B) \leq \mathbb{T}_{[-1,1]}(-1 + C + 2\varepsilon)$, and the right-hand side is clearly at most $\mathbb{T}_{[-1,1]}(-1 + C) + 2\varepsilon$. Assume next that $C \geq 2$; then obviously $\mathbb{T}_{[-1,1]}(-1 + B) \leq 1 = \mathbb{T}_{[-1,1]}(-1 + C)$. The claim follows.

Applying this simple claim twice, once for each inequality in (10), we obtain that

$$-2\varepsilon \leq \mathbb{T}_{[-1,1]}(y_{i^*} + r_{i^*}(x)) \\ - \mathbb{T}_{[-1,1]}(y_{a^*} + 2 \frac{\rho(x, x_{a^*})}{\rho(S_1^+, S_1^-)}) \leq 2\varepsilon,$$

which proves (9), and completes the proof of the theorem. \blacksquare

V. STRUCTURAL RISK MINIMIZATION

In this section, we show how to efficiently construct a classifier that optimizes the ‘‘bias-variance tradeoff’’ implicit in Corollary 6, equation (3). Let \mathcal{X} be a metric space, and assume we are given a labeled sample $S = (x_i, y_i) \in \mathcal{X} \times \{-1, 1\}$. For any Lipschitz constant L , let $k(L)$ be the minimal training error of S over all classifiers with Lipschitz constant L . We rewrite the generalization bound as follows:

$$\mathbb{P}\{(x, y) : \text{sgn}(f(x)) \neq y\} \leq \frac{k(L)}{n} + \sqrt{\frac{2}{n}} (D \ln(34en/D) \log_2(578n) + \ln(4/\delta)) \\ =: G(L) \quad (11)$$

where $D = (16L \text{diam}(\mathcal{X}))^{\text{ddim}(\mathcal{X})}$. This bound contains a free parameter, L , which may be tuned in the course of structural risk minimization. More precisely, decreasing L drives the ‘‘bias’’ term (number of mistakes) up and the ‘‘variance’’ term (fat-shattering dimension) down. We thus seek an (optimal) value of L where $G(L)$ achieves a minimum value, as described in the following theorem, which is our SRM result.

Theorem 10. *Let \mathcal{X} be a metric space and $0 < \varepsilon < \frac{1}{32}$. Given a labeled sample $S = (x_i, y_i) \in \mathcal{X} \times \{-1, 1\}$, $i = 1, \dots, n$, there exists a binary classification function $h : \mathcal{X} \rightarrow \{-1, 1\}$ satisfying the following properties:*

- (a) $h(x)$ can be evaluated at each $x \in \mathcal{X}$ in time $2^{O(\text{ddim}(\mathcal{X}))} \log n + \varepsilon^{-O(\text{ddim}(\mathcal{X}))}$, after an initial computation of $2^{O(\text{ddim}(\mathcal{X}))} n \log n + \left(\frac{\text{ddim}(\mathcal{X})}{\varepsilon}\right)^{O(\text{ddim}(\mathcal{X}))} n$ time.
- (b) The generalization error of h is bounded by

$$\mathbb{P}\{(x, y) : \text{sgn}(f(x)) \neq y\} \leq c \cdot \inf_{L>0} \left[\frac{k(L)}{n} + \sqrt{\frac{2}{n}} \left(D \ln(34en/D) \log_2(578n) + \ln \frac{4}{\delta} \right) \right].$$

for some constant $c \leq 2(1 + \varepsilon)$, and where

$$D = D(L) = \left(\frac{16L \text{diam}(\mathcal{X})}{1 - 32\varepsilon} \right)^{\text{ddim}(\mathcal{X})}.$$

We proceed with a description of our algorithm. We first give an algorithm with runtime $O(n^{4.373})$, and then improve the runtime, first to $O\left(\frac{\text{ddim}(\mathcal{X})}{\varepsilon} n^{2.373} \log n\right)$, then to $O(\text{ddim}(\mathcal{X}) n^2 \log n)$, and finally to $2^{O(\text{ddim}(\mathcal{X}))} n \log n + \left(\frac{\text{ddim}(\mathcal{X})}{\varepsilon}\right)^{O(\text{ddim}(\mathcal{X}))} n$.

Algorithm description: We start by giving a randomized algorithm that finds a value L^* that is optimal, namely, $G(L^*) = \inf_{L>0} G(L)$ for $G(L)$ that was defined in (11). The runtime of this algorithm is $O(n^{4.373})$ with high probability. First note the behavior of $k(L)$ as L increases. $k(L)$ decreases only when the value of L crosses certain critical values, each determined by a pair $x_i \in S^+$, $x_j \in S^-$ (that is, $L = \frac{2}{\rho(x_i, x_j)}$); for such L , the classification function h can correctly classify both these points. There are $O(n^2)$ critical values of L , and these can be determined by enumerating all interpoint distances between subsets $S^+, S^- \subset S$.

Below, we will show that for any given L , the value $k(L)$ can be computed in randomized time $O(n^{2.373})$. More precisely, we will show how to compute a partition of S into sets S_1 (with Lipschitz constant L) and S_0 (of size $k(L)$) in this time. Given sets $S_0, S_1 \subset S$, we can construct the classifier of Corollary 6. Since there are $O(n^2)$ critical values of L , we can calculate $k(L)$ for all critical values in $O(n^{4.373})$ total time, and thereby determine L^* . Then by Corollary 6, we may compute a classifier with a bias-variance tradeoff arbitrarily close to optimal.

To compute $k(L)$ for a given L in randomized time $O(n^{2.373})$, consider the following algorithm: Construct a bipartite graph $G = (V^+, V^-, E)$. The vertex sets V^+, V^- correspond to the labeled sets $S^+, S^- \in S$, respectively. The length of edge $e =$

(u, v) connecting vertices $u \in V^+$ and $v \in V^-$ is equal to the distance between the points, and E includes all edges of length less than $\frac{2}{L}$. (This E can be computed in $O(n^2)$ time.) Now, for all edges $e \in E$, a classifier with Lipschitz constant L necessarily misclassifies at least one endpoint of e . Hence, finding a classifier with Lipschitz constant L that misclassifies a minimum number of points in S is exactly the problem of finding a minimum vertex cover for the bipartite graph G . (This is an unweighted graph – the lengths are used only to determine E .) By König’s theorem, the minimum vertex cover problem in bipartite graphs is equivalent to the maximum matching problem, and a maximum matching in bipartite graphs can be computed in randomized time $O(n^{2.373})$ [28], [37]. This maximum matching immediately identifies a minimum vertex cover, which in turn gives the subsets S_0, S_1 , allowing us to compute a classifier achieving nearly optimal SRM.

First improvement: The runtime given above can be reduced from randomized $O(n^{4.373})$ to randomized $O(\text{ddim}(\mathcal{X})/\varepsilon \cdot n^{2.373} \log n)$, if we are willing to settle for a generalization bound $G(L)$ within a $(1 + \varepsilon)$ factor of the optimal $G(L^*)$, for any $\varepsilon \in (0, 1)$. To achieve this improvement, we discretize the candidate values of L , and evaluate $k(L)$ only for $O(\text{ddim}(\mathcal{X})/\varepsilon \cdot \log n)$ values of L , rather than all $\Theta(n^2)$ values as above. In the extreme case where the optimal hypothesis fails on all points of a single label, the classifier h is a constant function and $L^* = 0$. In all other cases, L^* must take values in the range $\left[\frac{2}{\text{diam}(\mathcal{X})}, \frac{n}{\text{diam}(\mathcal{X})}\right]$; indeed, every hypothesis correctly classifying a pair of opposite labelled points has Lipschitz constant at least $\frac{2}{\text{diam}(\mathcal{X})}$, and if $L^* > \frac{n}{\text{diam}(\mathcal{X})}$ then the complexity term (and $G(L^*)$) is greater than 1.

Our algorithm evaluates $k(L)$ for values of $L = \frac{2}{\text{diam}(\mathcal{X})} \left(1 + \frac{\varepsilon}{\text{ddim}(\mathcal{X})}\right)^i$ for $i = 0, 1, \dots, \left\lceil \log_{1+\varepsilon/\text{ddim}(\mathcal{X})} \frac{n}{2} \right\rceil$, and uses the candidate that minimizes $G(L)$. The number of candidate values for L is $O(\text{ddim}(\mathcal{X})/\varepsilon \cdot \log n)$, and one of these values — call it L' — satisfies $L^* \leq L' < \left(1 + \frac{\varepsilon}{\text{ddim}(\mathcal{X})}\right)L^*$. Observe that $k(L') \leq k(L^*)$ and that the complexity term for L' is greater than that for L^* by at most a factor $\sqrt{\left(1 + \frac{\varepsilon}{\text{ddim}(\mathcal{X})}\right)^{\text{ddim}(\mathcal{X})}} \leq e^{\varepsilon/2} \leq 1 + \varepsilon$ (where the final inequality holds since $\varepsilon < 1$). It follows that $G(L') < (1 + \varepsilon)G(L^*)$, implying that this algorithm

achieves a $(1 + \varepsilon)$ -approximation to $G(L^*)$.

Second improvement: The runtime can be further reduced from randomized $O(\text{ddim}(\mathcal{X})/\varepsilon \cdot n^{2.373} \log n)$ to deterministic $O(\text{ddim}(\mathcal{X})n^2 \log n)$, if we are willing to settle for a generalization bound $G(L)$ within a constant factor 2 of the optimal $G(L^*)$. The improvement comes from a faster vertex-cover computation. It is well known that a 2-approximation to vertex cover can be computed (in arbitrary graphs) by a greedy algorithm in time linear in the graph size $O(|V^+ \cup V^-| + |E|) = O(n^2)$, see e.g. [3]. Hence, we can compute in $O(n^2)$ time a function $k'(L)$ that satisfies $k(L) \leq k'(L) \leq 2k(L)$. We replace the randomized $O(n^{2.373})$ algorithm with this $O(n^2)$ time greedy algorithm. Then $k'(L) \leq 2k(L)$, and because we can approximate the complexity term to a factor smaller than 2 (as above, by choosing a constant $\varepsilon < 1$), our resulting algorithm finds a Lipschitz constant L' for which $G(L') \leq 2 \cdot G(L^*)$.

Final improvement: We can further improve the runtime from $O(\text{ddim}(\mathcal{X})n^2 \log n)$ to $2^{O(\text{ddim}(\mathcal{X}))}n \log n + (\text{ddim}(\mathcal{X})/\varepsilon)^{O(\text{ddim}(\mathcal{X}))}n$, at the cost of increasing the approximation factor to $2(1 + \varepsilon)$. The idea is to work with a sparser representation of the vertex cover problem. Recall that we discretized the values of L to powers of $\left(1 + \frac{\varepsilon}{\text{ddim}(\mathcal{X})}\right)$. As was already observed by [25], [10] in the context of hierarchies for doubling metric, \mathcal{X} contains at most $\varepsilon^{-1}2^{O(\text{ddim}(\mathcal{X}))}n$ of these distinct rounded critical values. After constructing a standard hierarchy (in time $2^{O(\text{ddim}(\mathcal{X}))}n \log n$), these ordered values may be extracted with $\left(\frac{\text{ddim}(\mathcal{X})}{\varepsilon}\right)^{O(\text{ddim}(\mathcal{X}))}n$ more work.

Let L be a discretized value considered above. We extract from S a subset $S' \subset S$ that is a $\left(\frac{\varepsilon}{\text{ddim}(\mathcal{X})} \cdot \frac{2}{L}\right)$ -net for S . Map each point $p \in S$ to its closest net point $p' \in S'$, and maintain for each net point two lists of points of S that are mapped to it — one list for positively labeled points and one for negatively labeled points. We now create an instance of vertex cover for the points of S : An edge $e = (u, v)$ for $u \in V^+$ and $v \in V^-$ is added to the edge set E' if the distance between the respective net points u' and v' is at most $\left(1 - \frac{2\varepsilon}{\text{ddim}(\mathcal{X})}\right) \frac{2}{L}$. Notice that $E' \subset E$, because the distance between such u, v is at most $\left(1 - \frac{2\varepsilon}{\text{ddim}(\mathcal{X})}\right) \frac{2}{L} + \frac{2\varepsilon}{\text{ddim}(\mathcal{X})} \cdot \frac{2}{L} = \frac{2}{L}$. Moreover, the edge set E' can be stored *implicitly* by recording every pair of net points that are within distance $\left(1 - \frac{2\varepsilon}{\text{ddim}(\mathcal{X})}\right) \frac{2}{L}$ — oppositely labeled point pairs that map (respectively) to this net-point pair is

considered (implicitly) to have an edge in E' . By the packing property, the number of net-point pairs to be recorded is at most $(\text{ddim}(\mathcal{X})/\varepsilon)^{O(\text{ddim}(\mathcal{X}))} n$, and by employing a hierarchy, the entire (implicit) construction may be done in time $2^{O(\text{ddim}(\mathcal{X}))} n \log n + (\text{ddim}(\mathcal{X})/\varepsilon)^{O(\text{ddim}(\mathcal{X}))} n$.

Now, for a given L , the run of the greedy algorithm for vertex cover can be implemented on this graph in time $(\text{ddim}(\mathcal{X})/\varepsilon)^{O(\text{ddim}(\mathcal{X}))} n$, as follows. The greedy algorithm considers a pair of net points within distance $\left(1 - \frac{2\varepsilon}{\text{ddim}(\mathcal{X})}\right) \frac{2}{L}$. If there exist $u \in V^+$ and $v \in V^-$ that map to these net points, then u, v are deleted from S and from the respective lists of the net points. (And similarly if u, v map to the same net point.) The algorithm terminates when there are no more points to remove, and correctness follows.

We now turn to the analysis. Since $E' \subset E$, the guarantees of the earlier greedy algorithm still hold. The resulting point set may contain opposite labeled points within distance $\left(1 - \frac{2\varepsilon}{\text{ddim}(\mathcal{X})}\right) \frac{2}{L} - \frac{2\varepsilon}{\text{ddim}(\mathcal{X})} \cdot \frac{2}{L} = \left(1 - \frac{4\varepsilon}{\text{ddim}(\mathcal{X})}\right) \frac{2}{L}$, resulting in a Lipschitz constant $L/(1 - \frac{4\varepsilon}{\text{ddim}(\mathcal{X})})$. This Lipschitz constant is slightly larger than the given L , which has the effect of increasing the complexity term in $G(L)$ by factor $\left(1 - \frac{4\varepsilon}{\text{ddim}(\mathcal{X})}\right)^{-\text{ddim}(\mathcal{X})/2} = 1 + \Theta(\varepsilon)$. The final result is achieved by scaling down ε to remove the leading constant.

VI. EXAMPLE: EARTHMOVER AND TIME-SERIES METRICS

To illustrate the potential power of our approach, we analyze its potential for two well-known metrics, the earthmover distance which operates on geometric sets, and Edit Distance with Real Penalty, which operates on time-series. We use the earthmover distance again in Section VII for our experiments.

Earthmover distance: We will analyze the doubling dimension of an earthmover metric, which is a natural metric for comparing two *sets* of k geometric features. It is often used in computer vision; for instance, an image can be represented as a set of pixels in a color space, yielding an accurate measure of dissimilarity between color characteristics of the images [31]. In an analogous manner, an image can be represented as a set of representative geometric features, such as object contours [19], other features [20], and SIFT descriptions [30]. In these contexts, $k \geq 2$ is usually a parameter which models the number of geometric features identified inside each image.

We use a simple yet common version, denoted $(\mathcal{X}_k, \text{EMD})$, where each point in \mathcal{X}_k is a multiset of size k in the unit square in the Euclidean plane, formally $S \subset [0, 1]^2$ and $|S| = k$ (allowing and counting multiplicities). The distance between such sets $S, T \in \mathcal{X}_k$ is given by

$$\text{EMD}(S, T) = \min_{\pi: S \rightarrow T} \left\{ \frac{1}{k} \sum_{s \in S} \|s - \pi(s)\|_2 \right\}, \quad (12)$$

where the minimum is over all one-to-one mappings $\pi: S \rightarrow T$. In other words, $\text{EMD}(S, T)$ is the minimum-cost bipartite matching between the two sets S, T , where costs correspond to Euclidean distance.

Lemma 11. *The earthmover metric over \mathcal{X}_k satisfies $\text{diam}(\mathcal{X}_k) \leq \sqrt{2}$ and $\text{ddim}(\mathcal{X}_k) \leq O(k \log k)$.*

Proof: For the rest of this proof, a point refers to one in the unit square, not \mathcal{X}_k . Consider a ball in \mathcal{X}_k of radius $r > 0$ around some S . Let N be an $r/2$ -net of the unit square $[0, 1]^2$, according to the definition in Section II. Now consider all multisets $T \subset [0, 1]^2$ of size k that satisfy the following condition: every point in T belongs to the net N and is within (Euclidean) distance $(k + 1/2)r$ from at least one point of S . Points in such a multiset T are chosen from a collection of size at most $k \cdot \left\lceil \frac{(k+1/2)r}{r/2} \right\rceil^{O(1)} \leq k^{O(1)}$ (by the packing property of the net points in the Euclidean plane). Thus, the number of such multisets T is $\lambda \leq (k^{O(1)})^k = k^{O(k)}$.

To complete the proof of the lemma, it suffices to show that the radius r ball (in \mathcal{X}_k) around S is covered by the λ balls of radius $r/2$ whose centers are given by the above multisets T . To see this, consider a multiset S' such that $\text{EMD}(S, S') \leq r$, and let us show that S' is contained in an $r/2$ -ball around one of the above multisets T . Observe that every point in S' is within distance at most kr from at least one point of S . Now “map” each point in S' to its nearest point in the net N , which must be less than $r/2$ away, by the covering property of the net. The result is a multiset T as above with $\text{EMD}(S', T) \leq r/2$. ■

Time-series distance metric: To present another example of the utility of our classification algorithms, we show that a commonly used metric model of sparse time-series vectors (with unbounded real coordinates) actually has a low doubling dimension. A widely used similarity function for time series is the Dynamic Time Warp (DTW) [39], a non-metric distance function between two time series which is similar to the ℓ_1 norm, except that it also allows

coordinate deletions or insertions in order to align the two series. The latter operations are used to ensure that the resulting series are of equal length, and these operations can also correlate the respective peaks and troughs of the series. We will consider a simple and popular metric version of DTW known as Edit Distance with Real Penalty (ERP) [9], which allows for insertions of zero-valued elements only.

The ERP distance is formally defined as follows. Given time-series vectors r and s with unbounded real coordinates and where the length of the longer series is exactly m , we may insert into r and s any number of zero-valued coordinates (called *gaps*) to produce new series \tilde{r} and \tilde{s} of equal length. Let R_p be the set of all time series of length $p \geq m$ which may be derived from r via gap insertions, and similarly S_p for s . Then $d_{\text{ERP}}(r, s) = \min_{p \geq m, \tilde{r} \in R_p, \tilde{s} \in S_p} \|\tilde{r} - \tilde{s}\|_1$. The ERP distance can be computed in quadratic time [9].

Our contribution is twofold: We show in Lemma 13 that a set of time series of length at most m may have doubling dimension $\Omega(m)$ under ERP, even when the coordinate range is limited to $\{1, 2\}$. This dimension is quite high, and motivates us to consider *sparse* time series, which form an active field of study [15], [41], [18], [24], [38]. We show in Lemma 14 that the set of time series vectors with only k non-zero elements — that is, k -sparse vectors — has doubling dimension $O(k \log k)$ under ERP, irrespective of the vector length m and even when the coordinates are real and unbounded. We first prove the claim below, and then proceed to the lower-bound on the dimension of length m vectors under ERP.

Claim 12. *Consider the set $T = \{1, 2\}^m$, and an integer $d \in [4, m/2]$. Then every vector $r \in T$ is within ERP distance d of fewer than $(\frac{3em}{d})^{2d}$ other vectors of T .*

Proof: We may view ERP on the vectors of T as a procedure transforming a vector $r \in T$ into some vector $s \in T$ as follows: The procedure inserts gaps in r to produce vector \tilde{r} , uses substitutions to transform \tilde{r} to \tilde{s} , and then deletes all gaps (i.e., zero-valued elements) from \tilde{s} to produce vector s . Here, the cost of a substitution from \tilde{r} to \tilde{s} is considered to be $|\tilde{r}_i - \tilde{s}_i| \in \{1, 2\}$, while the insertions and deletions entail no cost. But without loss of generality, we may assume that \tilde{r}_i and \tilde{s}_i are not both gaps. It follows that whenever \tilde{r}_i is produced by a gap insertion or \tilde{s}_i is a gap coordinate to be deleted, there must be a substitution from \tilde{r}_i to \tilde{s}_i . Thus, if $d_{\text{ERP}}(r, s) \leq j$ for $r, s \in T$, then the ERP procedure includes at most j

substitutions, and consequently at most j insertions and at most j deletions.

For a fixed r , the vector \tilde{r} can be produced in one of at most

$$\sum_{j=0}^d \binom{m+j}{j} < d \binom{3m/2}{d} \leq d \left(\frac{3em}{2d}\right)^d$$

possible ways of inserting $j \leq d$ gaps elements among the m coordinates of r . (Here we used the standard formula for combination with replacement.) Having produced \tilde{r} , the vector \tilde{s} can be produced in at most

$$\sum_{j=0}^d \binom{m+d}{j} 2^d < d \binom{3m/2}{d} 2^d \leq d \left(\frac{3em}{d}\right)^d$$

possible ways via substitutions in $j \leq d$ elements, where a single substitution sets some \tilde{s}_i to one of two possible values different from \tilde{r}_i . Having produced \tilde{s} , the vector s is produced by simply removing all gaps in \tilde{s} . It follows that there are fewer than $d^2 \left(\frac{3em}{2d}\right)^d \left(\frac{3em}{d}\right)^d \leq \left(\frac{3em}{d}\right)^{2d}$ vectors of T within distance d of r , as claimed. ■

Lemma 13. *There exists a set $S \subset \{1, 2\}^m$ whose doubling dimension under the ERP metric is $\Omega(m)$.*

Proof: We will demonstrate that for all $m \geq 4 \cdot 35 = 140$, there exists a set $S \subset \{1, 2\}^m$ of cardinality $2^{m/2}$ with diameter at most m and minimum interpoint distance at least $d = \lfloor \frac{m}{35} \rfloor$. As a consequence of Lemma 1, this S has doubling dimension $\Omega(m)$. Our proof uses a neighborhood counting argument similar to the one presented in [6, Lemma 8].

We begin with the set $T = \{1, 2\}^m$ of cardinality 2^m . The maximum interpoint distance under ERP in T is at most m , as this is the maximum distance under ℓ_1 in T . By Claim 12 with $d = \lfloor \frac{m}{35} \rfloor$, each vector $r \in T$ is within distance d of fewer than $(\frac{3em}{d})^{2d} \leq \left(3e \frac{m}{m/35-1}\right)^{2m/35} = \left(3 \cdot 35e \frac{1}{1-35/m}\right)^{2m/35} \leq (4 \cdot 35e)^{2m/35} < 2^{m/2}$ other points of T . We now use T to construct S greedily, starting with the empty set and repeatedly placing in S a point of T at distance at least $d = \lfloor \frac{m}{35} \rfloor$ from all points currently in S . Each point added to S invalidates fewer than $2^{m/2}$ other points in T from appearing in S in the future — these are the points within distance d of the new point. It follows that S contains at least $\frac{2^m}{2^{m/2}} = 2^{m/2}$ points, as claimed. ■

Although the doubling dimension of time-series vectors under ERP is large, the situation for sparse

vectors is much better, irrespective of the vector length and even when the coordinates are unbounded reals.

Lemma 14. *Every set S of k -sparse time-series vectors of real unbounded coordinates and arbitrary length has doubling dimension $O(k \log k)$ under the ERP metric.*

Proof: Similar to [17], let the *density constant* $\mu(S)$ of S be the smallest number such that for every $p > 0$, every ball of radius p contains at most $\mu(S)$ points of S at mutual distance strictly greater than $\frac{p}{2}$. It is known that the doubling constant of S is at most the density constant, i.e., $\lambda(S) \leq \mu(S)$. Indeed, for each radius p ball in S , take a maximal set (with respect to containment) of points at mutual distance strictly greater than $\frac{p}{2}$, and let each point be the center of a ball of radius $\frac{p}{2}$; the maximality implies that the small balls cover all points of the larger one. We get that every ball in S can be covered by at most $\mu(S)$ balls of half the radius.

It thus suffices to prove an upper bound on the density constant $\mu(S)$. To this end, consider a subset $T \subset S$ such that for some $p > 0$, all interpoint distances in T are in the range $(4p, 8p]$. In what follows, we prove an upper bound on $|T|$. First, for each vector $r \in T$, discretize the vector by rounding down each coordinate r_i to the nearest multiple of $\frac{p}{k}$, producing new set T' . Done over all coordinates, the rounding alters interpoint ERP distances by less than $2k \cdot \frac{p}{k} = 2p$ in total, and therefore all interpoint ERP distances in T' are in the range $(2p, 10p]$. We further remove from each $r \in T'$ all zero-valued coordinates, and this has no effect on interpoint ERP distances. The resulting set is T' of the same size as T , i.e., $|T'| = |T|$.

To bound $|T'|$, fix an arbitrary vector $r \in T'$, and consider the number of distinct discretized k -sparse vectors at ERP distance at most $10p$ from r . The ERP procedure may add up to k gaps to r to produce \tilde{r} , and there are $\sum_{j=0}^k \binom{k+j}{j} \leq \sum_{j=0}^k \binom{2k}{j} < 2^{2k}$ possible gap configurations. Note that the length of \tilde{r} is at most $2k$. Moving to the substitutions, since all vectors of T' are discretized into multiples of $\frac{p}{k}$ and are at distance at most $10p$, we can view the substitutions as adding or removing from these coordinates weight in units (multiples) of $\frac{p}{k}$, and there are in total $\frac{10p}{p/k} = 10k$ such units. If we view each substitution as accounting for a unit weight, and associate each coordinate with a sign that encodes whether the weight will be added or subtracted from that coordinate, then there are at most $(2k)^{10k} \cdot 2^{2k}$

possible substitution configurations to produce \tilde{s} . Having produced \tilde{s} , the vector s is produced from it by removing all gap elements. Altogether, $2^{\text{ddim}(S)} \leq \mu(S) = |T| = |T'| \leq 2^{2k} (2k)^{10k} 2^{2k} = k^{O(k)}$, from which the lemma follows. ■

VII. EXPERIMENTS

We considered the task of distinguishing five-petaled flowers from six-petaled ones. The images were taken from a shape matching/retrieval database called MPEG-7 Test Set⁵, and are displayed in Fig. 1. The original images were represented as 512×512 black and white matrices; we sampled these down to 128×128 . To render the task nontrivial, we retained only the image contour, as otherwise, it would suffice to consider the ratio of black/white pixels to achieve 100% accuracy. To illustrate the relative advantage of earthmover distance over the Euclidean one, we translated each image in the plane by various random shifts. We ran four classification algorithms on this data:

- **Euclidean Nearest Neighbor.** The images were treated as vectors in $\mathbb{R}^{128 \times 128}$, endowed with the Euclidean metric ℓ_2 .
- **EMD Nearest Neighbor.** The images were cut up into $16 \times 16 = 256$ square blocks, where each block b is viewed as a vector $b \in \mathbb{R}^{8 \times 8}$. Each image is thus represented as a sequence of 256 blocks, and over these sequences, EMD is defined as in (12) — except we used ℓ_1 instead of ℓ_2 as the base distance.
- **Euclidean SVM.** The Support Vector Machine (SVM) algorithm [11] was used, operating on vectors in $\mathbb{R}^{128 \times 128}$ with the Euclidean kernel $\langle x, y \rangle = x^T y$ and the regularization constant tuned by cross-validation.
- **SVM with RBF kernel.** The SVM algorithm was used, operating on vectors in $\mathbb{R}^{128 \times 128}$ with the Radial Basis Function (RBF) kernel $\langle x, y \rangle = \exp(-\|x - y\|_2^2 / \sigma^2)$, where the regularization constant and σ were tuned by cross-validation.

Our experimental results are listed in Table I. The relative magnitudes carry more significance than the absolute values, as the latter fluctuate with experiment design choices, such as the magnitude of the image translations, the thickness of the contour retained, and so forth. These results exhibit a natural setting in which classification algorithms for a non-Hilbertian

⁵<http://www.dabi.temple.edu/~shape/MPEG7/dataset.html>; the 5-petaled flowers are under device0-1 and 6-petaled are under device1-1.

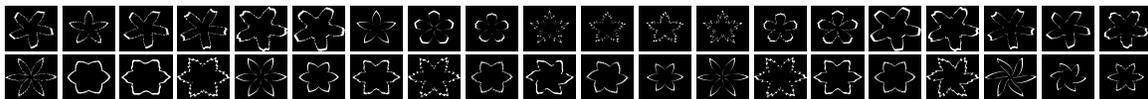


Fig. 1. The raw flower-contour data, 512×512 pixel black-and-white images. The two kinds of flowers (five- and six-petaled) are displayed in separate rows.

Method	Error
EMD nearest-neighbor	0.13
Euclidean nearest-neighbor	0.39
Euclidean SVM	0.43
SVM with RBF kernel	0.39

TABLE I
EXPERIMENTS CLASSIFYING THE FLOWER IMAGES. EACH METHOD IS AVERAGED OVER HUNDREDS OF EXPERIMENTS.

metric significantly outperforms the Hilbert-space algorithms, which is the main point we wished to illustrate here.

ACKNOWLEDGEMENTS

We thank the editor and anonymous referees for helpful suggestions for the manuscript. Thanks also to Shai Ben-David and Shai Shalev-Shwartz for sharing their book draft and to Ulrike von Luxburg for enlightening discussions.

REFERENCES

- [1] Noga Alon, Shai Ben-David, Nicolò Cesa-Bianchi, and David Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *Journal of the ACM*, 44(4):615–631, 1997.
- [2] Alexandr Andoni and Robert Krauthgamer. The computational hardness of estimating edit distance. *SIAM J. Comput.*, 39(6):2398–2429, April 2010.
- [3] R. Bar-Yehuda and S. Even. A linear-time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2(2):198 – 203, 1981.
- [4] Peter Bartlett and John Shawe-Taylor. Generalization performance of support vector machines and other pattern classifiers. In *Advances in kernel methods: support vector learning*, pages 43–54. MIT Press, 1999.
- [5] Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- [6] T. Batu, F. Ergun, J. Kilian, A. Magen, S. Raskhodnikova, R. Rubinfeld, and R. Sami. A sublinear algorithm for weakly approximating edit distance. In *Proceedings of the 35th ACM Symposium on Theory of Computing*, pages 316–324, 2003.
- [7] Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *ICML '06: Proceedings of the 23rd international conference on Machine learning*, pages 97–104. ACM, 2006.
- [8] Nader H. Bshouty, Yi Li, and Philip M. Long. Using the doubling dimension to analyze the generalization of learning algorithms. *Journal of Computer and System Sciences*, 75(6):323 – 335, 2009.
- [9] Lei Chen and Raymond Ng. On the marriage of L_p -norms and edit distance. In *Proceedings of the Thirtieth International Conference on Very Large Data Bases - Volume 30, VLDB '04*, pages 792–803, 2004.
- [10] Richard Cole and Lee-Ad Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *STOC*, pages 574–583, 2006.
- [11] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.
- [12] T. M. Cover and P. E. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13:21–27, 1967.
- [13] Ricky Der and Daniel Lee. Large-Margin Classification in Banach Spaces. In *JMLR Workshop and Conference Proceedings Volume 2: AISTATS 2007*, pages 91–98, 2007.
- [14] P. Enflo. On the nonexistence of uniform homeomorphisms between L_p -spaces. *Ark. Mat.*, 8:103–105, 1969.
- [15] Marco Franciosi and Giulia Menconi. Multi-dimensional sparse time series: feature extraction. *CoRR*, 2008.
- [16] Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Efficient classification for metric data. In *COLT*, pages 433–440, 2010.
- [17] Lee-Ad Gottlieb and Robert Krauthgamer. Proximity algorithms for nearly doubling spaces. *SIAM J. Discrete Math.*, 27(4):1759–1769, 2013.
- [18] Josif Grabocka, Alexandros Nanopoulos, and Lars Schmidt-Thieme. Classification of sparse time series via supervised matrix factorization. In *AAAI'12*, 2012.
- [19] K. Grauman and T. Darrell. Fast contour matching using approximate Earth Mover’s Distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Washington DC, June 2004.
- [20] K. Grauman and T. Darrell. The pyramid match kernel: Discriminative classification with sets of image features. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Beijing, China, October 2005.
- [21] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS*, pages 534–543, 2003.
- [22] Sarel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006.
- [23] Matthias Hein, Olivier Bousquet, and Bernhard Schölkopf. Maximal margin classification for metric spaces. *J. Comput. Syst. Sci.*, 71(3):333–359, 2005.
- [24] Vahid Khanagha and Khalid Daoudi. An efficient solution to sparse linear prediction analysis of speech. *EURASIP Journal on Audio, Speech, and Music Processing*, 2013(1):1–9, 2013.
- [25] R. Krauthgamer and J. R. Lee. Navigating nets: Simple algorithms for proximity search. In *15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 798–807, January 2004.
- [26] E. J. McShane. Extension of range of functions. *Bull. Amer. Math. Soc.*, 40(12):837–842, 1934.
- [27] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations Of Machine Learning*. The MIT Press, 2012.
- [28] Marcin Mucha and Piotr Sankowski. Maximum matchings via gaussian elimination. In *FOCS '04: Proceedings of the*

- 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 248–255, 2004.
- [29] Assaf Naor and Gideon Schechtman. Planar earthmover is not in l_1 . *SIAM J. Comput.*, 37:804–826, June 2007.
 - [30] Ofir Pele and Michael Werman. A linear time histogram metric for improved sift matching. In *10th European Conference on Computer Vision, ECCV '08*, pages 495–508. Springer-Verlag, 2008.
 - [31] Y. Rubner, C. Tomassi, and L. J. Guibas. The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, 2000.
 - [32] Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, 2002.
 - [33] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
 - [34] Ulrike von Luxburg and Olivier Bousquet. Distance-based classification with lipschitz functions. *Journal of Machine Learning Research*, 5:669–695, 2004.
 - [35] Larry Wasserman. *All of nonparametric statistics*. Springer Texts in Statistics. Springer, New York, 2006.
 - [36] Hassler Whitney. Analytic extensions of differentiable functions defined in closed sets. *Transactions of the American Mathematical Society*, 36(1):63–89, 1934.
 - [37] Virginia Vassilevska Williams. Breaking the Coppersmith-Winograd barrier. In *STOC '12: Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, New York, NY, USA, 2012. ACM Press.
 - [38] Bin Yang, Chenjuan Guo, and Christian S. Jensen. Travel cost inference from sparse, spatio temporally correlated time series using markov models. *Proc. VLDB Endow.*, 6(9):769–780, July 2013.
 - [39] Byoung-Kee Yi, H. V. Jagadish, and Christos Faloutsos. Efficient retrieval of similar time sequences under time warping. In *Proceedings of the Fourteenth International Conference on Data Engineering, ICDE '98*, pages 201–208. IEEE Computer Society, 1998.
 - [40] Haizhang Zhang, Yuesheng Xu, and Jun Zhang. Reproducing kernel banach spaces for machine learning. *Journal of Machine Learning Research*, 10:2741–2775, 2009.
 - [41] J. Ziniel, L.C. Potter, and P. Schniter. Tracking and smoothing of time-varying sparse signals via approximate belief propagation. In *Signals, Systems and Computers (ASILOMAR), 2010 Conference Record of the Forty Fourth Asilomar Conference on*, pages 808–812, Nov 2010.