

Agnostic Sample Compression for Linear Regression

Steve Hanneke

Toyota Technological Institute at Chicago

STEVE.HANNEKE@GMAIL.COM

Aryeh Kontorovich

Ben-Gurion University

KARYEH@CS.BGU.AC.IL

Menachem Sadigurschi

Ben-Gurion University

SADIGURS@POST.BGU.AC.IL

Editor: Satyen Kale and Aurélien Garivier

Abstract

We obtain the first positive results for bounded sample compression in the agnostic regression setting. We show that for $p \in \{1, \infty\}$, agnostic linear regression with ℓ_p loss admits a bounded sample compression scheme. Specifically, we exhibit efficient sample compression schemes for agnostic linear regression in \mathbb{R}^d of size $d + 1$ under the ℓ_1 loss and size $d + 2$ under the ℓ_∞ loss. We further show that for every other ℓ_p loss ($1 < p < \infty$), there does not exist an agnostic compression scheme of bounded size. This refines and generalizes a negative result of David, Moran, and Yehudayoff (2016a) for the ℓ_2 loss. We close by posing a general open question: for agnostic regression with ℓ_1 loss, does every function class admit a compression scheme of size equal to its pseudo-dimension? This question generalizes Warmuth’s classic sample compression conjecture for realizable-case classification (Warmuth, 2003).

Keywords: Sample Compression, Linear Regression, Agnostic Learning

1. Introduction

Sample compression is a central problem in learning theory, whereby one seeks to retain a “small” subset of the labeled sample that uniquely defines a “good” hypothesis. Quantifying *small* and *good* specifies the different variants of the problem. For instance, in the classification setting, taking *small* to mean “constant size” (i.e., depending only on the VC-dimension d of the concept class but not on the sample size m) and *good* to mean “consistent with the sample” specifies the classic realizable sample compression problem for VC classes. The feasibility of the latter was an open problem between its being posed by Littlestone and Warmuth (1986) and its positive resolution by Moran and Yehudayoff (2016), with various intermediate steps inbetween (Floyd, 1989; Helmbold, Sloan, and Warmuth, 1992; Floyd and Warmuth, 1995; Ben-David and Litman, 1998; Kuzmin and Warmuth, 2007; Rubinstein, Bartlett, and Rubinstein, 2009; Rubinstein and Rubinstein, 2012; Chernikov and Simon, 2013; Livni and Simon, 2013; Moran, Shpilka, Wigderson, and Yehudayoff, 2017).

A stronger form of this problem, where *small* means $O(d)$ (or even exactly d), remains open (Warmuth, 2003).

David, Moran, and Yehudayoff (2016a) recently generalized the definition of *compression scheme* to the agnostic case, where it is required that the function reconstructed from the compression set obtains an average loss on the full data set nearly as small as the function in the class that minimizes this quantity. Below, we give a strong motivation for this criterion by arguing an equivalence to the generalization ability of the compression-based learning algorithm. Under this definition, David et al. (2016a) extended the realizable-case result for VC classes to cover the agnostic case as well: a bounded-size compression scheme for the former implies such a scheme (in fact, of the same size) for the latter. They also generalized from binary to multiclass concept families, with the graph dimension in place of VC-dim. Proceeding to real-valued function classes, David et al. came to a starkly negative conclusion: they established that there is *no* constant-size agnostic sample compression scheme for linear functions under the ℓ_2 loss. (*Realizable* linear regression in \mathbb{R}^d trivially admits sample compression of size $d + 1$, under any loss, by selecting a minimal subset that spans the data.)

Main results. The negative result of David et al. (2016a) raises a general doubt over whether sample compression is ever a viable approach to agnostic learning of real-valued functions. In this work, we address this concern by proving that, if we replace the ℓ_2 loss with the ℓ_1 loss, then there *is* a simple agnostic compression scheme of size $d + 1$ for linear regression in \mathbb{R}^d . This is somewhat surprising, given the above negative result for the ℓ_2 loss. We also construct an agnostic compression scheme of size $d + 2$ under the ℓ_∞ loss. However, interestingly, we also generalize the argument of David et al. (2016a) to show that these are the *only two* ℓ_p losses ($1 \leq p \leq \infty$) for which there exists a constant-size compression scheme. Computationally, our compression schemes for ℓ_1 and ℓ_∞ amount to solving a polynomial (in fact, linear) size linear program. These appear to be the first positive results for bounded agnostic sample compression for real-valued function classes. We close by posing an intriguing open question generalizing our result to arbitrary function classes: under the ℓ_1 loss, does *every* function class admit an agnostic compression scheme of size equal to its pseudo-dimension? We argue that this represents a generalization of Warmuth’s classic sample compression problem, which asks whether every space of classifiers admits a compression scheme of size VC-dimension in the realizable case.

Related work. David et al. (2016a, Theorem 4.1) obtained the aforementioned negative result for ℓ_2 agnostic linear regression, as well as an $\tilde{O}(\log(d/\varepsilon))$ -size compression scheme for *approximate* ℓ_2 agnostic linear regression (the latter model is not considered here).

Hanneke et al. (2018) showed how to convert *consistent* real-valued learners into constant-size (i.e., independent of sample size) efficiently computable compression schemes for the realizable (or nearly realizable) case, for a notion of compression scheme that allows an ϵ slack in the empirical ℓ_∞ loss of the reconstructed function. This result was obtained via a weak-to-strong boosting procedure, coupled with a generic construction of weak learners out of abstract regressors. The *agnostic* variant of this problem remains open in its full generality.

Ashtiani et al. (2018) adapted the notion of a compression scheme to the distribution learning problem. They showed that if a class of distributions admits robust compressibility then it is agnostically learnable.

2. Problem setting, definitions and notation

Our instance space is $\mathcal{X} = \mathbb{R}^d$, label space is $\mathcal{Y} = \mathbb{R}$, and hypothesis class is $\mathcal{F} \subseteq \mathcal{Y}^{\mathcal{X}}$, consisting of all $h_{\mathbf{a},b} : \mathcal{X} \rightarrow \mathcal{Y}$ given by $h_{\mathbf{a},b}(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle + b$, indexed by $\mathbf{a} \in \mathbb{R}^d, b \in \mathbb{R}$. For $1 \leq p < \infty$, the loss incurred by a hypothesis $h \in \mathcal{F}$ on a labeled sample $S = ((\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m))$ is given by

$$L_p(h, S) := \frac{1}{m} \sum_{i=1}^m |h(\mathbf{x}_i) - y_i|^p,$$

while for $p = \infty$,

$$L_\infty(h, S) := \max_{1 \leq i \leq m} |h(\mathbf{x}_i) - y_i|.$$

Now let us introduce a formal definition of sample compression, and a criterion we require of any valid *agnostic compression scheme*. Following the definition, we provide a strong motivation for this criterion in terms of an equivalence to the generalization ability of the learning algorithm under general conditions. Following David et al. (2016b), we define a *selection scheme* (κ, ρ) for a hypothesis class $\mathcal{F} \subset \mathcal{Y}^{\mathcal{X}}$ is defined as follows. A *k-selection* function κ maps sequences $((\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)) \in \bigcup_{\ell \geq 1} (\mathcal{X} \times \mathcal{Y})^\ell$ to elements in $\mathcal{K} = \bigcup_{\ell \leq k'} (\mathcal{X} \times \mathcal{Y})^\ell \times \bigcup_{\ell \leq k''} \{0, 1\}^\ell$, where $k' + k'' \leq k$. A *reconstruction* is a function $\rho : \mathcal{K} \rightarrow \mathcal{Y}^{\mathcal{X}}$. We say that (κ, ρ) is a *k-size agnostic sample compression scheme* for \mathcal{F} if κ is a *k-selection* and for all $S = ((\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m))$, $f_S := \rho(\kappa(S))$ achieves \mathcal{F} -competitive empirical loss:

$$L_p(f_S, S) \leq \inf_{f \in \mathcal{F}} L_p(f, S).$$

In principle, the *size k* of an agnostic compression scheme may depend on the data set size m , in which case we may denote this dependence by $k(m)$. However, in this work we are primarily interested in the case when $k(m)$ is *bounded*: that is, $k(m) \leq k$ for some m -independent value k . Note that the above definition is fully general, in that it defines a notion of agnostic compression scheme for *any* function class \mathcal{F} and loss function L , though in the present work we focus on \mathcal{F} as linear functions in \mathbb{R}^d and the loss as L_p for $1 \leq p \leq \infty$.

Remark 1 *At first, it might seem unclear why this is an appropriate generalization of sample compression to the agnostic setting. To see that it is so, we note that one of the main interests in sample compression schemes is their ability to generalize: that is, to achieve low excess risk under a distribution P on $\mathcal{X} \times \mathcal{Y}$ when the data S are sampled iid according to P (Littlestone and Warmuth, 1986; Floyd and Warmuth, 1995; Graepel, Herbrich, and Shawe-Taylor, 2005). Also, as mentioned, in this work we are primarily interested in sample compression schemes that have bounded size: $k(m) \leq k$ for an m -independent value k . Furthermore, we are also focusing on the most-general case, where*

this size bound should be independent of everything else in the scenario, such as the data S or the underlying distribution P . Given these interests, we claim that the above definition is essentially the only reasonable choice. More specifically, for L_p loss with $1 \leq p < \infty$, any compression scheme with $k(m)$ bounded such that its expected excess risk under any P converges to 0 as $m \rightarrow \infty$ necessarily satisfies the above condition (or is easily converted into one that does). To see this, note that for any data set S for which such a compression scheme fails to satisfy the above \mathcal{F} -competitive empirical loss criterion, we can define a distribution P that is simply uniform on S , and then the compression scheme’s selection function would be choosing a bounded number of points from S and a bounded number of bits, while guaranteeing that excess risk under P approaches 0, or equivalently, excess empirical loss approaches 0. To make this argument fully formal, only a slight modification is needed, to handle having multiple copies of points from S in the compression set; given that the size is bounded, these repetitions can be encoded in a bounded number of extra bits, so that we can stick to strictly distinct points in the compression set.

In the converse direction, we also note that any bounded-size agnostic compression scheme (in the sense of the above definition) will be guaranteed to have excess risk under P converging to 0 as $m \rightarrow \infty$, in the case that S is sampled iid according to P , for losses L_p with $1 \leq p < \infty$, as long as P guarantees that $(X, Y) \sim P$ has Y bounded (almost surely). This follows from classic arguments about the generalization ability of compression schemes, which includes results for the agnostic case (Graepel, Herbrich, and Shawe-Taylor, 2005). For unbounded Y one cannot, in general, obtain distribution-free generalization bounds. However, one can still obtain generalization under certain broader restrictions (see, e.g., Mendelson, 2015 and references therein). The generalization problem becomes more subtle for the L_∞ loss: this cannot be expressed as a sum of pointwise losses and there are no standard techniques for bounding the deviation of the sample risk from the true risk. One recently-studied guarantee achieved by minimizing empirical L_∞ loss is a kind of “hybrid error” generalization, developed in Hanneke et al. (2018, Theorem 9). We refer the interested reader to that work for the details of those results, which can easily be extended to apply to our notion of agnostic compression scheme.

We denote set cardinality by $|\cdot|$ and $[m] := \{1, \dots, m\}$. Vectors $\mathbf{v} \in \mathbb{R}^d$ are denoted by boldface, and their j th coordinate is indicated by $\mathbf{v}(j)$. (Thus, $\mathbf{v}_i(j)$ indicates the j th coordinate of the i th vector in a sequence.)

3. Impossibility results for ℓ_p , $1 < p < \infty$

David et al. (2016b, Theorem 4.1) proved an impossibility result for the ℓ_2 loss:

Theorem 2 (David et al. (2016b)) *There is no agnostic sample compression scheme for zero-dimensional linear regression with size $k(m) \leq m/2$.*

We show that constant-size compression is impossible for all ℓ_p losses with $1 < p < \infty$:

Theorem 3 *There is no agnostic sample compression scheme for zero-dimensional linear regression under ℓ_p loss, $1 < p < \infty$, with size $k(m) < \log(m)$.*

Proof Consider a sample $(y_1, \dots, y_m) \in \{0, 1\}^m$. Partition the indices $i \in [m]$ into $S_0 := \{i \in [m] : y_i = 0\}$ and $S_1 := \{i \in [m] : y_i = 1\}$. The empirical risk minimizer is given by

$$\hat{r} := \operatorname{argmin}_{s \in \mathbb{R}} \sum_{i=1}^m |y_i - s|^p.$$

To obtain an explicit expression for \hat{r} , define

$$F(s) = \sum_{i=1}^m |y_i - s|^p = |S_1|(1-s)^p + |S_0|s^p =: N_1(1-s)^p + N_0s^p.$$

We then compute

$$F'(s) = pN_0s^{p-1} - pN_1(1-s)^{p-1}$$

and find that $F'(s) = 0$ occurs at

$$\hat{s} = \frac{\mu^{1/(p-1)}}{1 + \mu^{1/(p-1)}},$$

where $\mu = N_1/N_0$. A straightforward analysis of the second derivative shows that $\hat{s} = \hat{r}$ is indeed the unique minimizer of F .

Thus, given a sample of size m , the unique minimizer \hat{r} is uniquely determined by N_0 — which can take on any of integer $m + 1$ values between 0 and m . On the other hand, every output of a k -selection function κ outputs a multiset $\hat{S} \subseteq S$ of size k' and a binary string of length $k'' = k - k'$. Thus, the total number of values representable by a k -selection scheme is at most

$$\sum_{k'=0}^k k' 2^{k-k'} < 2^{k+1} - k,$$

which, for $k < \log m$, is less than m . ■

Remark 4 *A more refined analysis, along the lines of David et al. (2016b, Theorem 4.1), should yield a lower bound of $k = \Omega(m)$. A technical complication is that unlike the $p = 2$ case, whose empirical risk minimizer has a simple explicit form, the general ℓ_p loss does not admit a closed-form solution and uniqueness must be argued from general convexity principles. We leave this for the extended version.*

4. Compressibility results for ℓ_1 and ℓ_∞

In sharp contrast with the $1 < p < \infty$ case, we show that in \mathbb{R}^d , agnostic linear regression admits a compression scheme of size $d + 1$ under ℓ_1 and $d + 2$ under ℓ_2 .

Theorem 5 *There exists an efficiently computable compression scheme for agnostic linear regression in \mathbb{R}^d under the ℓ_1 loss of size $d + 1$.*

Proof We start with $d = 0$. The sample then consists of (y_1, \dots, y_m) [formally: pairs (x_i, y_i) , where $x_i \equiv 0$], and $\mathcal{F} = \mathbb{R}$ [formally, all functions $h : 0 \mapsto \mathbb{R}$]. We define f_S to be the median of (y_1, \dots, y_m) , which for odd m is defined uniquely and for even m can be taken arbitrarily as the smaller of the two midpoints. It is well-known that such a choice minimizes the empirical ℓ_1 risk, and it clearly constitutes a compression scheme of size 1.

The case $d = 1$ will require more work. The sample consists of $(x_i, y_i)_{i \in [m]}$, where $x_i, y_i \in \mathbb{R}$, and $\mathcal{F} = \{\mathbb{R} \ni x \mapsto ax + b : a, b \in \mathbb{R}\}$. Let (a^*, b^*) be a (possibly non-unique) minimizer of

$$L(a, b) := \sum_{i \in [m]} |(ax_i + b) - y_i|, \tag{1}$$

achieving the value L^* . We claim that we can always find two indices $\hat{i}, \hat{j} \in [m]$ such that the line determined by $(x_{\hat{i}}, y_{\hat{i}})$ and $(x_{\hat{j}}, y_{\hat{j}})$ also achieves the optimal empirical risk L^* . More precisely, the line (\hat{a}, \hat{b}) induced by $((x_{\hat{i}}, y_{\hat{i}}), (x_{\hat{j}}, y_{\hat{j}}))$ via¹ $\hat{a} = (y_{\hat{j}} - y_{\hat{i}})/(x_{\hat{j}} - x_{\hat{i}})$ and $\hat{b} = y_{\hat{i}} - \hat{a}x_{\hat{i}}$, verifies $L(\hat{a}, \hat{b}) = L^*$.

To prove this claim, we begin by recasting (1) as a linear program:

$$\begin{aligned} \min_{(\varepsilon_1, \dots, \varepsilon_m, a, b) \in \mathbb{R}^{m+2}} \quad & \sum_{i=1}^m \varepsilon_i \quad \text{s.t.} \\ \forall i \in [m] \quad & \varepsilon_i \geq 0 \\ \forall i \in [m] \quad & ax_i + b - y_i \leq \varepsilon_i \\ \forall i \in [m] \quad & -ax_i - b + y_i \leq \varepsilon_i. \end{aligned} \tag{2}$$

We observe that the linear program in (2) is feasible with a finite solution (and actually, the constraints $\varepsilon_i \geq 0$ are redundant). Furthermore, any optimal value is achievable at one of the extreme points of the constraint-set polytope $\mathcal{P} \subset \mathbb{R}^{m+2}$. Next, we claim that the extreme points of the polytope \mathcal{P} are all of the form $v \in \mathcal{P}$ with two (or more) of the ε_i s equal to 0. This suffices to prove our main claim, since $\varepsilon_i = 0$ in $v \in \mathcal{P}$ iff the (a, b) induced by v verifies $ax_i + b = y_i$; in other words, the line induced by (a, b) contains the point (x_i, y_i) . If a line contains two data points, it is uniquely determined by them: these constitute a compression set of size 2. (See illustration in Figure 1.)

Now we prove our claimed property of the extreme points. First, we claim that any extreme point of \mathcal{P} must have least one ε_i equal to 0. Indeed, let (a, b) define a line. Define

$$b^+ := \min \left\{ \tilde{b} \in [b, \infty) : \exists i \in [m], ax_i + \tilde{b} = y_i \right\}$$

and analogously,

$$b^- := \max \left\{ \tilde{b} \in (-\infty, b] : \exists i \in [m], ax_i + \tilde{b} = y_i \right\}.$$

In words, (a, b^+) is the line obtained by increasing b to a maximum value of b^+ , where the line (a, b^+) touches a datapoint, and likewise, (a, b^-) is the line obtained by decreasing b to a minimum value of b^- , where the line (a, b^-) touches a datapoint.

1. We ignore the degenerate possibility of vertical lines, which reduces to the 0-dimensional case.




Figure 1: A sample S of $m = 20$ points (x_i, y_i) was drawn iid uniformly from $[0, 1]^2$. On this sample, ℓ_1 regression was performed by solving the LP in (2), shown on the left, and ℓ_∞ regression was performed by solving the LP in (3), on the right. In each case, the regressor provided by the LP solver is indicated by the thick (red) line. Notice that for ℓ_1 , the line contains exactly 2 datapoints. For ℓ_∞ , the regressor contains no datapoints; rather, the $d + 2 = 3$ “support vectors” are indicated by ●.

Define by $S_{a,b}^+ := \{i : |ax_i + b < y_i|\}$ the points above the line defined by (a, b) and $S_{a,b}^- := \{i : |ax_i + b > y_i|\}$ the points below the line defined by (a, b) . For a line (a, b) which does not contain a data point we can rewrite the sample loss as

$$\begin{aligned} L(a, b) &= \sum_{i \in S_{a,b}^+} (y_i - (ax_i + b)) + \sum_{i \in S_{a,b}^-} ((ax_i + b) - y_i) \\ &= \left(\sum_{i \in S_{a,b}^-} x_i - \sum_{i \in S_{a,b}^+} x_i \right) a + (|S_{a,b}^-| - |S_{a,b}^+|) b + \left(\sum_{i \in S_{a,b}^+} y_i - \sum_{i \in S_{a,b}^-} y_i \right) \\ &=: \lambda a + \mu b + \nu. \end{aligned}$$

Since for fixed a and $b \in [b^-, b^+]$, the quantities $S_{a,b}^-, S_{a,b}^+$ are constant, it follows that the function $L(a, \cdot)$ is affine in b , and hence minimized at $b^\pm \in \{b^-, b^+\}$. Thus, there is no loss of generality in taking $b^* = b^\pm$, which implies that the optimal solution’s line (a^*, b^*) contains a data point (x_i, y_i) . If the line (a^*, b^\pm) contains other data points then we are done, so assume to the contrary that ε_i is the only ε_i that vanishes in the corresponding solution $v^* \in \mathcal{P}$.

Let $\mathcal{P}_i \subset \mathcal{P}$ consist of all v for which $\varepsilon_i = 0$, corresponding to all feasible solutions whose line contains the data point (x_i, y_i) . Let us say that two lines $(a_1, b_1), (a_2, b_2)$ are *equivalent* if they induce the same partition on the data points, in the sense of linear separation in the plane. The formal condition is $S_{a_1, b_1}^- = S_{a_2, b_2}^-$, which is equivalent to $S_{a_1, b_1}^+ = S_{a_2, b_2}^+$.

Define $\mathcal{P}_i^* \subset \mathcal{P}_i$ to consist of those feasible solutions whose line is equivalent to (a^*, b^\pm) . Denote by $a^+ := \max \{a : (\varepsilon_1, \dots, \varepsilon_m, a, b) \in \mathcal{P}_i^*\}$ and define v^+ to be a feasible solution in \mathcal{P}_i^* with slope a^+ , and analogously, $a^- := \min \{a : (\varepsilon_1, \dots, \varepsilon_m, a, b) \in \mathcal{P}_i^*\}$ and $v^- \in \mathcal{P}_i^*$ with

slope a^- . Geometrically this corresponds to rotating the line (a^*, b^*) about the point (x_i, y_i) until it encounters a data point above and below.

Writing, as above, the sample loss in the form $L(a, b)$, we see that $L(\cdot, b^\pm)$ is affine in a over the range $a \in [a^-, a^+]$ and hence is minimized at one of the endpoints. This furnishes another datapoint (x_j, y_j) verifying $\hat{a}x_j + \hat{b} = y_j$ for $L(\hat{a}, \hat{b}) = L^*$, and hence proves compressibility into two points for $d = 1$.

Generalizing to $d > 1$ is quite straightforward. We define

$$L(\mathbf{a}, b) = \sum_{i \in [m]} |(\langle \mathbf{a}, \mathbf{x}_i \rangle + b) - y_i|$$

and express it as a linear program analogous to (2), where the minimization is over $(\varepsilon_1, \dots, \varepsilon_m, \mathbf{a}, b) \in \mathbb{R}^{m+d+1}$ and the expression ax_i in the constraints is replaced by $\langle \mathbf{a}, \mathbf{x}_i \rangle$. Given an optimal solution (\mathbf{a}^*, b^*) , we argue exactly as above that b^* may be chosen so that the optimal regressor contains some datapoint — say, (\mathbf{x}_1, y_1) . Holding b^* and $\mathbf{a}(j)$, $j \neq 1$ fixed, we argue, as above, that $\mathbf{a}(1)$ may be chosen so that the optimal regressor contains another datapoint (say, (\mathbf{x}_2, y_2)). Proceeding in this fashion, we inductively argue that the optimal regressor may be chosen to contain some $d + 1$ datapoints, which provides the requisite compression scheme. ■

Theorem 6 *There exists an efficiently computable compression scheme for agnostic linear regression in \mathbb{R}^d under the ℓ_∞ loss of size $d + 2$.*

Proof

Given m labeled points in $\mathbb{R}^d \times \mathbb{R}$, $S = (\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)$ and any $\mathbf{a} \in \mathbb{R}^d$, $b \in \mathbb{R}$ define the empirical risk

$$L(\mathbf{a}, b) := \max \{ |\langle \mathbf{a}, \mathbf{x}_i \rangle + b - y_i| : i \in [m] \}.$$

We cast the risk minimization problem as a linear program:

$$\begin{aligned} \min_{(\varepsilon, \mathbf{a}, b) \in \mathbb{R}^{d+2}} &: \quad \varepsilon & (3) \\ \text{s.t. } \forall i &: \quad \varepsilon - \langle \mathbf{a}, \mathbf{x}_i \rangle - b + y_i \geq 0 \\ & \quad \varepsilon + \langle \mathbf{a}, \mathbf{x}_i \rangle + b - y_i \geq 0. \end{aligned}$$

(As before, the constraint $\varepsilon \geq 0$ is implicit in the other constraints.) Introducing the Lagrange multipliers $\lambda_i, \mu_i \geq 0$, $i \in [m]$, we cast the optimization problem in the form of a Lagrangian:

$$\mathcal{L}(\varepsilon, \mathbf{a}, b, \mu_1, \dots, \mu_m, \lambda_1, \dots, \lambda_m) = \varepsilon - \sum_{i=1}^m \lambda_i (\varepsilon - \langle \mathbf{a}, \mathbf{x}_i \rangle - b + y_i) - \sum_{i=1}^m \mu_i (\varepsilon + \langle \mathbf{a}, \mathbf{x}_i \rangle + b - y_i).$$

The KKT conditions imply, in particular, that

$$\begin{aligned} \forall i &: \quad \lambda_i (\varepsilon - \langle \mathbf{a}, \mathbf{x}_i \rangle - b + y_i) = 0 \\ & \quad \mu_i (\varepsilon + \langle \mathbf{a}, \mathbf{x}_i \rangle + b - y_i) = 0. \end{aligned}$$

Geometrically, this means that either the constraints corresponding to the i th datapoint are inactive — in which case, omitting the datapoint does not affect the solution — or otherwise, the i th datapoint induces the active constraint

$$\langle \mathbf{a}, \mathbf{x}_i \rangle + b - y_i = \varepsilon. \tag{4}$$

On analogy with SVM, let us refer to the datapoints satisfying (4) as the *support vectors*; clearly, the remaining sample points may be discarded without affecting the solution. Solutions to (3) lie in \mathbb{R}^{d+2} and hence $d + 2$ linearly independent datapoints suffice to uniquely pin down an optimal $(\varepsilon, \mathbf{a}, b)$ via the equations (4). ■

5. Open Problem: Compressing to Pseudo-dimension Number of Points

The above positive results for ℓ_1 loss may also lead us to wonder how general of a result might be possible. In particular, noting that the pseudo-dimension (Pollard, 1984, 1990; Anthony and Bartlett, 1999) of linear functions in \mathbb{R}^d is precisely $d + 1$ (Anthony and Bartlett, 1999), there is an intriguing possibility for the following generalization. For any class \mathcal{F} of real-valued functions, denote by $d(\mathcal{F})$ the pseudo-dimension of \mathcal{F} .

Open Problem: Under the ℓ_1 loss, does every class \mathcal{F} of real-valued functions admit an agnostic compression scheme of size $d(\mathcal{F})$?

It is also interesting, and perhaps more approachable as an initial aim, to ask whether there is an agnostic compression scheme of size at most *proportional to* $d(\mathcal{F})$. Even falling short of this, one can ask the more-basic question of whether classes with $d(\mathcal{F}) < \infty$ always have *bounded* agnostic compression schemes (i.e., independent of sample size m), and more specifically whether the bound is expressible purely as a function of $d(\mathcal{F})$ (Moran and Yehudayoff, 2016 have shown this is always possible in the realizable classification setting).

These questions are directly related to (and inspired by) the well-known long-standing conjecture of Floyd and Warmuth (1995); Warmuth (2003), which asks whether, for realizable-case binary classification, there is always a compression scheme of size at most linear in the VC dimension of the concept class. Indeed, it is clear that a positive solution of our open problem above would imply a positive solution to the original sample compression conjecture, since in the realizable case with a function class \mathcal{F} of $\{0, 1\}$ -valued functions, the minimal empirical ℓ_1 loss on the data is zero, and any function obtaining zero empirical ℓ_1 loss on a data set labeled with $\{0, 1\}$ values must be $\{0, 1\}$ -valued on that data set, and thus can be thought of as a sample-consistent classifier.² Noting that, for \mathcal{F} containing $\{0, 1\}$ -valued functions, $d(\mathcal{F})$ is equal the VC dimension, the implication is clear.

The converse of this direct relation is not necessarily true. Specifically, for a set \mathcal{F} of real-valued functions, consider the set \mathcal{H} of subgraph sets: $h_f(x, y) = \mathbb{I}[y \leq f(x)]$, $f \in \mathcal{F}$. In particular, note that the VC dimension of \mathcal{H} is precisely $d(\mathcal{F})$. It is *not* true that any realizable classification compression scheme for \mathcal{H} is also an agnostic compression scheme for \mathcal{F} under ℓ_1 loss. Nevertheless, this reduction-to-classification approach seems intuitively

2. To make such a function actually binary-valued everywhere, it suffices to threshold at $1/2$.

appealing, and it might possibly be the case that there is some way to *modify* certain types of compression schemes for \mathcal{H} to convert them into agnostic compression schemes for \mathcal{F} . Following up on this line of investigation seems the natural next step toward resolving the above general open question.

References

- 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>.
- Hassan Ashtiani, Shai Ben-David, Nick Harvey, Christopher Liaw, Abbas Mehrabian, and Yaniv Plan. Settling the sample complexity for learning mixtures of gaussians. In *NIPS*, 2018.
- Shai Ben-David and Ami Litman. Combinatorial variability of vapnik-chervonenkis classes with applications to sample compression schemes. *Discrete Applied Mathematics*, 86(1): 3–25, 1998.
- Artem Chernikov and Pierre Simon. Externally definable sets and dependent pairs. *Israel J. Math.*, 194(1):409–425, 2013. ISSN 0021-2172. URL <https://doi.org/10.1007/s11856-012-0061-9>.
- Ofir David, Shay Moran, and Amir Yehudayoff. Supervised learning through the lens of compression. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2784–2792, 2016a.
- Ofir David, Shay Moran, and Amir Yehudayoff. Supervised learning through the lens of compression. In *Advances in Neural Information Processing Systems*, pages 2784–2792, 2016b.
- Sally Floyd. Space-bounded learning and the vapnik-chervonenkis dimension. In *Proceedings of the second annual workshop on Computational learning theory*, pages 349–364. Morgan Kaufmann Publishers Inc., 1989.
- Sally Floyd and Manfred K. Warmuth. Sample compression, learnability, and the Vapnik-Chervonenkis dimension. *Machine Learning*, 21(3):269–304, 1995.
- Thore Graepel, Ralf Herbrich, and John Shawe-Taylor. PAC-bayesian compression bounds on the prediction error of learning algorithms for classification. *Machine Learning*, 59(1-2):55–76, 2005.
- Steve Hanneke, Aryeh Kontorovich, and Menachem Sadigurschi. Sample compression for real-valued learners. *CoRR*, abs/1805.08254, 2018. URL <http://arxiv.org/abs/1805.08254>.
- David Helmbold, Robert Sloan, and Manfred K Warmuth. Learning integer lattices. *SIAM Journal on Computing*, 21(2):240–266, 1992.

- Dima Kuzmin and Manfred K. Warmuth. Unlabeled compression schemes for maximum classes. *Journal of Machine Learning Research*, 8:2047–2081, 2007. URL <http://dl.acm.org/citation.cfm?id=1314566>.
- Nick Littlestone and Manfred K. Warmuth. Relating data compression and learnability. Technical report, Department of Computer and Information Sciences, Santa Cruz, CA, Ju, 1986.
- Roi Livni and Pierre Simon. Honest compressions and their application to compression schemes. In *Conference on Learning Theory*, pages 77–92, 2013.
- Shahar Mendelson. Learning without concentration. *J. ACM*, 62(3):21:1–21:25, 2015. doi: 10.1145/2699439. URL <http://doi.acm.org/10.1145/2699439>.
- Shay Moran and Amir Yehudayoff. Sample compression schemes for VC classes. *J. ACM*, 63(3):21:1–21:10, 2016. doi: 10.1145/2890490. URL <http://doi.acm.org/10.1145/2890490>.
- Shay Moran, Amir Shpilka, Avi Wigderson, and Amir Yehudayoff. Teaching and compressing for low vc-dimension. In *A Journey Through Discrete Mathematics*, pages 633–656. Springer, 2017.
- David Pollard. *Convergence of Stochastic Processes*. Springer-Verlag, 1984.
- David Pollard. *Empirical processes: theory and applications*. NSF-CBMS Regional Conference Series in Probability and Statistics, 2. Institute of Mathematical Statistics, Hayward, CA, 1990. ISBN 0-940600-16-1.
- Benjamin I. P. Rubinstein and J. Hyam Rubinstein. A geometric approach to sample compression. *Journal of Machine Learning Research*, 13:1221–1261, 2012. URL <http://dl.acm.org/citation.cfm?id=2343686>.
- Benjamin I. P. Rubinstein, Peter L. Bartlett, and J. Hyam Rubinstein. Shifting: One-inclusion mistake bounds and sample compression. *J. Comput. Syst. Sci.*, 75(1):37–59, 2009. doi: 10.1016/j.jcss.2008.07.005. URL <https://doi.org/10.1016/j.jcss.2008.07.005>.
- Manfred K. Warmuth. Compressing to VC dimension many points. In *Proceedings of the 16th Conference on Learning Theory*, 2003.