

Sparse Algorithms are not Stable: A No-free-lunch Theorem

Huan Xu, Constantine Caramanis, *Member, IEEE* and Shie Mannor, *Senior Member, IEEE*



Abstract—We consider two desired properties of learning algorithms: *sparsity* and *algorithmic stability*. Both properties are believed to lead to good generalization ability. We show that these two properties are fundamentally at odds with each other: a sparse algorithm cannot be stable and vice versa. Thus, one has to trade off sparsity and stability in designing a learning algorithm. In particular, our general result implies that ℓ_1 -regularized regression (Lasso) cannot be stable, while ℓ_2 -regularized regression is known to have strong stability properties and is therefore not sparse.

Index Terms—Stability, Sparsity, Lasso, Regularization

1 INTRODUCTION

Stability and Sparsity have both emerged as important properties of machine learning algorithms. In a broad sense, stability means that an algorithm is well-posed, so that given two very similar data sets, an algorithm’s output varies little. More specifically, an algorithm is stable if its output is nearly identical on two data sets differing on only one sample (this is known as the leave-one-out error). Stability itself is a desirable property for learning algorithms. For example, in feature-selection, one might seek algorithms that select nearly the same feature set when run on very similar data sets. In addition to inherent application-related reasons that stability is desirable, stability is also known to closely related to statistical property of learning algorithms. Following the landmark work in [1], stability is also an avenue for proving generalization performance of an algorithm. For example, in [2] stability properties of ℓ_2 -regularized Support Vector Machines (SVM) are used to establish consistency. Also see [3], [4], [5] and many others. More recently, in [5], the authors show that a (weak) notion of stability is *necessary and sufficient* for the learnability of an algorithm.

Meanwhile, sparsity has long been important for a variety of reasons: a sparse solution is less complicated and hence generalizes well [6]; sparsity can facilitate interpretability [7], [8], [9], [10]; and sparse algorithms may be computationally much easier to implement, store, compress, etc. Accordingly, numerous sparsity-promoting algorithms have been proposed in signal processing and virtually all fields in machine learning. A partial list includes: Lasso, 1-norm SVM, Deep Belief Network, Sparse PCA [11], [12], [13], [14], [15], [16] and many others.

In this paper, we investigate the mutual relationship of these two concepts. In particular, we show that sparse algorithms are not stable: if an algorithm “encourages sparsity” (in a sense defined precisely below) then its sensitivity to small perturbations of the input data remains bounded away from zero, i.e., it has no uniform stability properties. We define these notions formally in Section 2. We prove this “no-free-lunch” theorem by constructing an instance where the leave-one-out error of the algorithm is bounded away from zero by exploiting the property that a sparse algorithm can have non-unique optimal solutions, and is therefore *ill-posed*. In a broad sense, our result reveals that there is a fundamental tradeoff relationship. Intuitively, sparsity typically requires non-smoothness in the objective; this in turn often results in multiple potentially very different (distant) optimal solutions. Thus, small perturbations can significantly change the output, and hence render the algorithm unstable.

As we detail below, there is a spectrum of notions of stability and of sparsity. The particular definitions we use are arguably the strongest, although they are satisfied by many of the most popular algorithms. While certain relaxations (e.g., to group sparsity) are easy to obtain, our results do not extend to the weakest notion of stability proposed in [5]. Therefore, an interesting but challenging question is to determine the appropriate notion of stability and sparsity that is just strong enough for a similar no-free-lunch theorem to hold. We leave this for future research.

This paper is organized as follows. We start with the necessary definitions in Section 2 and provide the no-free-lunch theorem based on these definitions in Section 3. Sections 2 and 3 are devoted to regression

-
- *H. Xu is with the Department of Mechanical Engineering, The National University of Singapore, SINGAPORE.
E-mail: mpexuh@nus.edu.sg.*
 - *C. Caramanis is with the Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX, USA.
E-mail: caramanis@mail.utexas.edu.*
 - *S. Mannor is with the Department of Electrical Engineering, Technion, Haifa, ISRAEL.
E-mail: shie@ee.technion.ac.il*

algorithms; and in Section 4 we generalize the theorem to arbitrary loss functions. In Section 5 we discuss the justification of the particular notions of stability and sparsity considered in this paper. Brief concluding remarks are given in Section 6.

Notation: Capital letters (e.g., A) and boldface letters (e.g., \mathbf{w}) are used to denote matrices and column vectors, respectively. We use the transpose of a column vector to represent a row vector. Unless otherwise specified, the same letter is used to represent a part of an object. For example, the i^{th} column of a matrix A is denoted by \mathbf{a}_i . Similarly, the i^{th} element of a vector \mathbf{d} is denoted by d_i .

2 SETUP AND ASSUMPTIONS

Consider a training sequence $\{(b_i, \mathbf{a}_i)\}_{i=1}^n$, where \mathbf{a} is the vector of input values of the observation and b is the output — we use (\mathbf{b}, A) to represent the entire sequence. We consider optimization algorithms that seek to minimize the loss given a new observation $(\hat{b}, \hat{\mathbf{a}})$. For a given objective, rather than comparing two solutions $\mathbf{w}^1, \mathbf{w}^2$ by considering their empirical loss, we adopt a somewhat more general framework, considering only the partial ordering induced by any learning algorithm \mathbb{L} and training set (\mathbf{b}, A) . That is, given two candidate solutions, $\mathbf{w}^1, \mathbf{w}^2$, we write

$$\mathbf{w}^1 \preceq_{(\mathbf{b}, A)} \mathbf{w}^2,$$

if on input (\mathbf{b}, A) , the algorithm \mathbb{L} would select \mathbf{w}^2 before \mathbf{w}^1 . In short, given an algorithm \mathbb{L} , each sample set (\mathbf{b}, A) defines an order relationship $\preceq_{(\mathbf{b}, A)}$ among all candidate solutions \mathbf{w} . This order relationship defines a family of “best” solutions, and one of these, \mathbf{w}^* is the output of the algorithm. We denote this by writing $\mathbf{w}^* \in \mathbb{L}_{(\mathbf{b}, A)}$.

Thus, by defining a data-dependent partial order on the space of solutions, we can talk more generally about algorithms, their stability, and their sparsity. As we define below, an algorithm \mathbb{L} is sparse if the set $\mathbb{L}_{(\mathbf{b}, A)}$ of optimal solutions contains a sparse solution, and an algorithm is stable if the sets $\mathbb{L}_{(\mathbf{b}, A)}$ and $\mathbb{L}_{(\hat{\mathbf{b}}, \hat{A})}$ do not contain solutions that are very far apart, when (\mathbf{b}, A) and $(\hat{\mathbf{b}}, \hat{A})$ differ on only one point.

We make a few assumptions on the preference order:

Assumption 1: (i) The value of a column corresponding to a non-selected feature has no effect on the ordering: given j , \mathbf{b} , A , \mathbf{w}^1 and \mathbf{w}^2 , suppose that $\mathbf{w}^1 \preceq_{(\mathbf{b}, A)} \mathbf{w}^2$, and $w_j^1 = w_j^2 = 0$. Then for any $\hat{\mathbf{a}}$,

$$\mathbf{w}^1 \preceq_{(\mathbf{b}, \hat{A})} \mathbf{w}^2,$$

where

$$\hat{A} = (\mathbf{a}_1, \dots, \mathbf{a}_{j-1}, \hat{\mathbf{a}}, \mathbf{a}_{j+1}, \dots, \mathbf{a}_m).$$

(ii) Adding a sample that is perfectly predicted by a particular solution, cannot decrease its place in

the partial ordering: given \mathbf{b} , A , \mathbf{w}^1 , \mathbf{w}^2 , b' and \mathbf{z} , suppose that $\mathbf{w}^1 \preceq_{(\mathbf{b}, A)} \mathbf{w}^2$, and $b' = \mathbf{z}^\top \mathbf{w}^2$. Then

$$\mathbf{w}^1 \preceq_{(\bar{\mathbf{b}}, \bar{A})} \mathbf{w}^2,$$

where

$$\bar{\mathbf{b}} = \begin{pmatrix} \mathbf{b} \\ b' \end{pmatrix}; \quad \bar{A} = \begin{pmatrix} A \\ \mathbf{z}^\top \end{pmatrix}.$$

(iii) The order relationship is preserved when a trivial (all zeros) feature is added: given j , \mathbf{b} , A , \mathbf{w}^1 and \mathbf{w}^2 , suppose that $\mathbf{w}^1 \preceq_{(\mathbf{b}, A)} \mathbf{w}^2$. Then

$$\hat{\mathbf{w}}^1 \preceq_{(\mathbf{b}, \hat{A})} \hat{\mathbf{w}}^2,$$

where

$$\hat{\mathbf{w}}^i = \begin{pmatrix} \mathbf{w}^i \\ 0 \end{pmatrix}, \quad i = 1, 2; \quad \hat{A} = (A, \mathbf{0}).$$

(iv) The partial ordering and hence the algorithm, is feature-wise symmetric: given \mathbf{b} , A , \mathbf{w}^1 , \mathbf{w}^2 and $P \in \mathbb{R}^{m \times m}$ a permutation matrix, if $\mathbf{w}^1 \preceq_{(\mathbf{b}, A)} \mathbf{w}^2$, then

$$P^\top \mathbf{w}^1 \preceq_{(\mathbf{b}, AP)} P^\top \mathbf{w}^2.$$

Many popular algorithms, e.g., standard and regularized regression, satisfy these assumptions. See also Section 5.

Next, we define precisely what we mean by stability and sparsity. We recall the definition of uniform (algorithmic) stability first, as given in [1]. We let \mathcal{Z} denote the space of points and labels (typically this will either be \mathbb{R}^{m+1} or a closed subset of it) so that $S \in \mathcal{Z}^n$ denotes a collection of n labelled training points. For regression problems, therefore, we have $S = (\mathbf{b}, A) \in \mathcal{Z}^n$. We let \mathbb{L} denote a learning algorithm, and for $(\mathbf{b}, A) \in \mathcal{Z}^n$, we let $\mathbb{L}_{(\mathbf{b}, A)}$ denote the output of the learning algorithm (i.e., the regression function it has learned from the training data). Then given a loss function l , and a labelled point $s = (b, \mathbf{z}) \in \mathcal{Z}$, $l(\mathbb{L}_{(\mathbf{b}, A)}, s)$ denotes the loss of the algorithm that has been trained on the set (\mathbf{b}, A) , on the data point s . Thus in the regression setup, we would have $l(\mathbb{L}_{(\mathbf{b}, A)}, s) = |\mathbb{L}_{(\mathbf{b}, A)}(\mathbf{z}) - b|$.

Definition 1: [1] An algorithm \mathbb{L} has uniform stability β_n with respect to the loss function l if the following holds:

$$\forall (\mathbf{b}, A) \in \mathcal{Z}^n, \forall i \in \{1, \dots, n\} : \\ \max_{\mathbf{z}' \in \mathcal{Z}} |l(\mathbb{L}_{(\mathbf{b}, A)}, \mathbf{z}') - l(\mathbb{L}_{(\mathbf{b}, A) \setminus i}, \mathbf{z}')| \leq \beta_n.$$

Here $\mathbb{L}_{(\mathbf{b}, A) \setminus i}$ stands for the learned solution with the i^{th} sample removed from (\mathbf{b}, A) , i.e., with the i^{th} row of A and the i^{th} element of \mathbf{b} removed.

At first glance, this definition may seem too stringent for any reasonable algorithm to exhibit good stability properties. However, as shown in [1], many algorithms have uniform stability with β_n going to zero. In particular, Tikhonov regularized regression (i.e., ℓ_2 -regularized regression) has stability that goes to zero as $1/n$. Indeed, recent work [17] shows that for $p > 1$,

ℓ_p regularization has uniform stability with β_n going to zero as $1/n$. Stability can be used to establish strong PAC bounds. For example, [1] shows that if the loss is bounded by M , then with n samples, the following holds with probability at least $(1 - \delta)$,

$$R \leq R_{\text{emp}} + 2\beta_n + (4n\beta_n + M)\sqrt{\frac{\ln 1/\delta}{2n}},$$

where β_n denotes the uniform stability, R the expected loss, and R_{emp} the empirical (i.e., training) loss.

Since Lasso is an example of an algorithm that yields sparse solutions, one implication of the results of this paper is that while ℓ_p -regularized ($p > 1$) regression yields stable solutions, ℓ_1 -regularized regression does not. We show that the stability parameter of Lasso does not decrease in the number of samples (compared to the $O(1/n)$ decay for ℓ_p -regularized regression). In fact, we show that Lasso's stability is, in the following sense, the worst possible stability. To this end, we define the notion of the Pseudo Maximal Error (PME), which is the worst possible error a training algorithm can have for arbitrary training set and testing sample labelled by zero.

Definition 2: Given the sample space $\mathcal{Z} = \mathcal{Y} \times \mathcal{X}$ where $\mathcal{Y} \subseteq \mathbb{R}$, $\mathcal{X} \subseteq \mathbb{R}^m$, and $0 \in \mathcal{Y}$, the pseudo maximal error for a learning algorithm \mathbb{L} w.r.t. \mathcal{Z} is

$$\mathfrak{b}_n(\mathbb{L}, \mathcal{Z}) \triangleq \max_{(\mathbf{b}, A) \in \mathcal{Z}^n, \mathbf{z} \in \mathcal{X}} l(\mathbb{L}_{(\mathbf{b}, A)}, (0, \mathbf{z})).$$

As above, $l(\cdot, \cdot)$ is a given loss function.

As an example, if \mathcal{X} is the unit ball, and $W(\mathbb{L})$ is the set of vectors \mathbf{w} that are optimal with respect to at least one training set, then $\mathfrak{b}_n(\mathbb{L}, \mathcal{Z}) = \max_{\mathbf{w} \in W(\mathbb{L})} \|\mathbf{w}\|$. Thus, unless \mathbb{L} is a trivial algorithm which always outputs $\mathbf{0}$, the PME is bounded away from zero.

Observe that $\mathfrak{b}_n(\mathbb{L}, \mathcal{Z}) \geq \mathfrak{b}_1(\mathbb{L}, \mathcal{Z})$, since by repeatedly choosing the worst sample (for \mathfrak{b}_1), the algorithm will yield the same solution. Hence the PME does not diminish as the number of samples, n , increases.

We next define the notion of sparsity of an algorithm which we use.

Definition 3: A weight vector \mathbf{w}^* Identifies Redundant Features of A if

$$\forall i \neq j, \quad \mathbf{a}_i = \mathbf{a}_j \Rightarrow w_i^* w_j^* = 0.$$

An algorithm \mathbb{L} is said to be *able to Identify Redundant Features* (IRF for short) if $\forall (\mathbf{b}, A)$ there exists $\mathbf{w}^* \in \mathbb{L}_{(\mathbf{b}, A)}$ that identifies redundant features of A .

Being IRF means that at least one solution of the algorithm does not select both features if they are identical. We note that this is a rather weak notion of sparsity. An algorithm that achieves reasonable sparsity (such as Lasso) should be IRF. Notice that IRF is a property that is typically easy to check.

Before concluding this section, we comment on the two definitions that we considered, namely, the uniform stability and IRF.

Stability is generally a desirable characteristic of learning algorithms. Recently, a notion of stability termed **all-i-LOO** stable, has been shown to be a necessary and sufficient condition for learnability [5]. An algorithm is **all-i-LOO** stable, if

$$\forall i \in \{1, \dots, n\} : \quad \mathbb{E}_{S \sim \mu^n} |l(\mathbb{L}_S, s_i) - l(\mathbb{L}_{S \setminus i}, s_i)| \leq \beta_n,$$

where μ is the generating distribution. This definition (and also others defined in the literature, e.g., [3]) requires knowledge of the distribution that generates samples and thus may be hard to verify, or simply inappropriate when this distribution is unknown. In contrast, the notion of *uniform stability* does not involve the *unknown* generating distribution and thus can be evaluated. This is a principal reason that uniform stability, while more restrictive, has seen wide application, particularly for deriving generalization bounds of learning algorithms.

The notion of IRF is proposed as an easily verifiable property that most sparse algorithms satisfy. While there are different notions of sparsity proposed in the literature, the most widely applied, recently popularized in the compressed sensing literature (and around in myriad other places) calls sparsity the number of non-zero elements of a vector. Accordingly, an algorithm is called sparse if it finds the sparsest or nearly-sparsest solution subject to performance constraints (e.g., small regression error). Under this definition, it is clear that IRF is a necessary property for an algorithm to be sparse.

3 THE MAIN THEOREM

The main contribution of this paper establishes an incompatibility between stability and sparsity: if an algorithm is sparse, in the sense that it identifies redundant features, then that algorithm *is not stable*. Notably, this theorem applies to Lasso.

Theorem 1: Let $\mathcal{Z} = \mathcal{Y} \times \mathcal{X}$ be the sample space with m features, where $\mathcal{Y} \subseteq \mathbb{R}$, $\mathcal{X} \subseteq \mathbb{R}^m$, $0 \in \mathcal{Y}$ and $\mathbf{0} \in \mathcal{X}$. Let $\hat{\mathcal{Z}} = \mathcal{Y} \times \mathcal{X} \times \mathcal{X}$ be the sample space with $2m$ features. If a learning algorithm \mathbb{L} (trained on points in $\hat{\mathcal{Z}}$) satisfies Assumption 1 and identifies redundant features, its uniform stability bound β is lower bounded by $\mathfrak{b}_n(\mathbb{L}, \mathcal{Z})$, and in particular does not go to zero with n .

Proof: Note that in light of the definition of uniform stability, it suffices to provide one example that algorithm \mathbb{L} fails to achieve a small stability bound. We construct such an (somewhat extreme) example as follows.

Let (\mathbf{b}, A) and $(0, \mathbf{z}^\top)$ be the sample set and the new observation such that they jointly achieve $\mathfrak{b}_n(\mathbb{L}, \mathcal{Z})$, i.e., for some $\mathbf{w}^* \in \mathbb{L}_{(\mathbf{b}, A)}$, we have

$$\mathfrak{b}_n(\mathbb{L}, \mathcal{Z}) = l(\mathbf{w}^*, (0, \mathbf{z})). \quad (1)$$

Let $0^{n \times m}$ be the $n \times m$ 0-matrix, and $\mathbf{0}$ stand for the zero vector of length m . Further, let

$$\begin{aligned} \hat{\mathbf{z}} &\triangleq (\mathbf{0}^\top, \mathbf{z}^\top); & \hat{A} &\triangleq (A, A); \\ \tilde{\mathbf{b}} &\triangleq \begin{pmatrix} \mathbf{b} \\ 0 \end{pmatrix}; & \tilde{A} &\triangleq \begin{pmatrix} A & A \\ \mathbf{0}^\top & \mathbf{z}^\top \end{pmatrix}. \end{aligned}$$

Observe that $(\mathbf{b}, \hat{A}) \in \hat{\mathcal{Z}}^n$ and $(\tilde{\mathbf{b}}, \tilde{A}) \in \hat{\mathcal{Z}}^{n+1}$. We first show that

$$\begin{pmatrix} \mathbf{0} \\ \mathbf{w}^* \end{pmatrix} \in \mathbb{L}_{(\mathbf{b}, \hat{A})}; \quad \begin{pmatrix} \mathbf{w}^* \\ \mathbf{0} \end{pmatrix} \in \mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})}. \quad (2)$$

Notice that \mathbb{L} is feature-wise symmetric (Assumption 1 (iv)) and I.R.F, hence there exists a \mathbf{w}' such that

$$\begin{pmatrix} \mathbf{0} \\ \mathbf{w}' \end{pmatrix} \in \mathbb{L}_{(\mathbf{b}, \hat{A})}.$$

Since $\mathbf{w}^* \in \mathbb{L}_{(\mathbf{b}, A)}$, we have

$$\begin{aligned} &\mathbf{w}' \preceq_{(\mathbf{b}, A)} \mathbf{w}^* \\ \Rightarrow &\begin{pmatrix} \mathbf{0} \\ \mathbf{w}' \end{pmatrix} \preceq_{(\mathbf{b}, (0^{n \times m}, A))} \begin{pmatrix} \mathbf{0} \\ \mathbf{w}^* \end{pmatrix} \\ \Rightarrow &\begin{pmatrix} \mathbf{0} \\ \mathbf{w}' \end{pmatrix} \preceq_{(\mathbf{b}, \hat{A})} \begin{pmatrix} \mathbf{0} \\ \mathbf{w}^* \end{pmatrix} \\ \Rightarrow &\begin{pmatrix} \mathbf{0} \\ \mathbf{w}^* \end{pmatrix} \in \mathbb{L}_{(\mathbf{b}, \hat{A})}. \end{aligned}$$

The first implication follows from Assumption 1 part (iii), and the second from part (i).

By Assumption 1 (iv) (feature-wise symmetry), we have

$$\begin{pmatrix} \mathbf{w}^* \\ \mathbf{0} \end{pmatrix} \in \mathbb{L}_{(\mathbf{b}, \hat{A})}.$$

Furthermore,

$$0 = (\mathbf{0}^\top, \mathbf{z}^\top) \begin{pmatrix} \mathbf{w}^* \\ \mathbf{0} \end{pmatrix},$$

and thus by Assumption 1(ii) we have

$$\begin{pmatrix} \mathbf{w}^* \\ \mathbf{0} \end{pmatrix} \in \mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})}.$$

Hence (2) holds. This leads to (recall that $l(\mathbf{w}^*, (\hat{b}, \hat{\mathbf{a}})) = |\hat{b} - \hat{\mathbf{a}}^\top \mathbf{w}^*|$)

$$l(\mathbb{L}_{(\mathbf{b}, \hat{A})}, (0, \hat{\mathbf{z}})) = l(\mathbf{w}^*, (0, \mathbf{z})); \quad l(\mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})}, (0, \hat{\mathbf{z}})) = 0.$$

By definition of the uniform bound, we have

$$\beta \geq l(\mathbb{L}_{(\mathbf{b}, \hat{A})}, (0, \hat{\mathbf{z}})) - l(\mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})}, (0, \hat{\mathbf{z}})).$$

Hence by (1) we have $\beta \geq \mathfrak{b}_n(\mathbb{L}, \mathcal{Z})$, which establishes the theorem. \square

Theorem 1 not only means that a sparse algorithm is not stable, it also states that if an algorithm is stable, there is no hope that it will be sparse, since it cannot even identify redundant features. For instance, ℓ_2 regularized regression is stable (see Example 3 with a linear kernel), and does not identify redundant features.

4 GENERALIZATION TO ARBITRARY LOSS

Thus far our focus has been on the regression problem, i.e., with loss function $l(\mathbf{w}^*, (\hat{b}, \hat{\mathbf{a}})) = |\hat{b} - \hat{\mathbf{a}}^\top \mathbf{w}^*|$. Our results are more general, and apply more broadly, e.g., to the ϵ -insensitive loss function $l(\mathbf{w}^*, (\hat{b}, \hat{\mathbf{a}})) = \max(|\hat{b} - \hat{\mathbf{a}}^\top \mathbf{w}^*| - \epsilon, 0)$ or the classification error $l(\mathbf{w}^*, (\hat{b}, \hat{\mathbf{a}})) = \mathbf{1}_{\hat{b} \neq \text{sign}(\hat{\mathbf{a}}^\top \mathbf{w}^*)}$. We can generalize the results derived to algorithms with loss function having the form $l(\mathbf{w}^*, (\hat{b}, \hat{\mathbf{a}})) = f_m(\hat{b}, \hat{a}_1 w_1^*, \dots, \hat{a}_m w_m^*)$ for any f_m (here, \hat{a}_i and w_i^* denote the i^{th} component of $\hat{\mathbf{a}} \in \mathbb{R}^m$ and $\mathbf{w}^* \in \mathbb{R}^m$, respectively) that satisfies the following conditions:

- (a) $f_m(b, v_1, \dots, v_i, \dots, v_j, \dots, v_m)$
 $= f_m(b, v_1, \dots, v_j, \dots, v_i, \dots, v_m); \quad \forall b, \mathbf{v}, i, j.$
- (b) $f_m(b, v_1, \dots, v_m) = f_{m+1}(b, v_1, \dots, v_m, 0); \quad \forall b, \mathbf{v}.$

In words, (a) means that the loss function is feature-wise symmetric, and (b) means that a dummy feature does not change the loss. Observe that both the ϵ -insensitive loss and the classification error satisfy these conditions.

In contrast to the regression setup, under an arbitrary loss function, there may not exist a sample that can be perfectly predicted by the zero vector, rendering Definition 2 unnecessarily restrictive. We thus modify Definition 2 to accommodate this setting.

Definition 4: Let $\mathcal{Z} = \mathcal{Y} \times \mathcal{X}$ where $\mathcal{Y} \subseteq \mathbb{R}$ and $\mathcal{X} \subseteq \mathbb{R}^m$. Then the pseudo maximal error for a learning algorithm \mathbb{L} w.r.t. \mathcal{Z} is given by

$$\hat{\mathfrak{b}}_n(\mathbb{L}, \mathcal{Z}) \triangleq \max_{(\mathbf{b}, A) \in \hat{\mathcal{Z}}^n, (b, \mathbf{z}) \in \mathcal{Z}} \left\{ l(\mathbb{L}_{(\mathbf{b}, A)}, (b, \mathbf{z})) - l(\mathbf{0}, (b, \mathbf{z})) \right\}.$$

The PME in the arbitrary loss case is thus defined as the largest (w.r.t. all possible testing samples) performance gap of outputs of a learning algorithm and the zero vector. Observe that Definition 4 is a relaxation of Definition 2 in the sense that if the loss function is indeed the regression error, then the PME defined by Definition 4 is no smaller than that of Definition 2.

To account for the modification of Definition 2, we need to make Assumption 1 slightly stronger: we replace Assumption 1(ii) with the following one.

Assumption 2: (ii) Given $\mathbf{b}, A, \mathbf{w}^1, \mathbf{w}^2, b'$ and \mathbf{z} if

$$\mathbf{w}^1 \preceq_{(\mathbf{b}, A)} \mathbf{w}^2, \quad l(\mathbf{w}^2, (b', \mathbf{z})) \leq l(\mathbf{w}^1, (b', \mathbf{z}))$$

then

$$\mathbf{w}^1 \preceq_{(\bar{\mathbf{b}}, \bar{A})} \mathbf{w}^2, \quad \text{where } \bar{\mathbf{b}} = \begin{pmatrix} \mathbf{b} \\ b' \end{pmatrix}; \quad \bar{A} = \begin{pmatrix} A \\ \mathbf{z}^\top \end{pmatrix}.$$

Assumption 2(ii) means that adding a sample that is better predicted (i.e., smaller loss) cannot make a candidate solution less preferred.

With these modifications, we have a generalization of Theorem 1.

Theorem 2: Let $\mathcal{Z} = \mathcal{Y} \times \mathcal{X}$ be the sample space with m features, where $\mathcal{Y} \subseteq \mathbb{R}$, $\mathcal{X} \subseteq \mathbb{R}^m$, and $\mathbf{0} \in \mathcal{X}$.

Let $\hat{\mathcal{Z}} = \mathcal{Y} \times \mathcal{X} \times \mathcal{X}$ be the sample space with $2m$ features. If a learning algorithm \mathbb{L} (trained on points in $\hat{\mathcal{Z}}$) satisfies Assumption 2 and identifies redundant features, its uniform stability bound β is lower bounded by $\hat{\mathbf{b}}_n(\mathbb{L}, \mathcal{Z})$, and in particular does not go to zero with n .

Proof: This proof follows a similar line of reasoning as the proof of Theorem 1. Let (\mathbf{b}, A) and (b', \mathbf{z}^\top) be the sample set and the new observation such that they jointly achieve $\hat{\mathbf{b}}_n(\mathbb{L}, \mathcal{Z})$, i.e., there exists $\mathbf{w}^* \in \mathbb{L}(\mathbf{b}, A)$ such that:

$$\begin{aligned} \hat{\mathbf{b}}_n(\mathbb{L}, \mathcal{Z}) &= l(\mathbf{w}^*, (b', \mathbf{z})) - l(\mathbf{0}, (b', \mathbf{z})) \\ &= f_m(b', w_1^* z_1, \dots, w_m^* z_m) - f(b', 0, \dots, 0). \end{aligned}$$

Let $0^{n \times m}$ be the $n \times m$ 0-matrix, and $\mathbf{0}$ stand for the zero vector of length m . We denote

$$\begin{aligned} \hat{\mathbf{z}} &\triangleq (\mathbf{0}^\top, \mathbf{z}^\top); & \hat{A} &\triangleq (A, A); \\ \tilde{\mathbf{b}} &\triangleq \begin{pmatrix} \mathbf{b} \\ b' \end{pmatrix}; & \tilde{A} &\triangleq \begin{pmatrix} A & A \\ \mathbf{0}^\top & \mathbf{z}^\top \end{pmatrix}. \end{aligned}$$

Observe that $(\mathbf{b}, \hat{A}) \in \hat{\mathcal{Z}}^n$ and $(\tilde{\mathbf{b}}, \tilde{A}) \in \hat{\mathcal{Z}}^{n+1}$. To prove the theorem, it suffices to show that there exist $\mathbf{w}^1, \mathbf{w}^2$ such that

$$\mathbf{w}^1 \in \mathbb{L}_{(\mathbf{b}, \hat{A})}, \quad \mathbf{w}^2 \in \mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})},$$

and

$$l(\mathbf{w}^1, (b', \hat{\mathbf{z}})) - l(\mathbf{w}^2, (b', \hat{\mathbf{z}})) \geq \hat{\mathbf{b}}_n(\mathbb{L}, \mathcal{Z})$$

where again,

$$\hat{\mathbf{b}}_n(\mathbb{L}, \mathcal{Z}) = f_m(b', w_1^* z_1, \dots, w_m^* z_m) - f_m(b', 0, \dots, 0).$$

By an identical argument to the proof of Theorem 1, Assumption 1(i), (iii) and (iv) imply that:

$$\begin{pmatrix} \mathbf{0} \\ \mathbf{w}^* \end{pmatrix} \in \mathbb{L}_{(\mathbf{b}, \hat{A})}.$$

Hence there exists $\mathbf{w}^1 \in \mathbb{L}_{(\mathbf{b}, \hat{A})}$ such that

$$\begin{aligned} l(\mathbf{w}^1, (b', \hat{\mathbf{z}})) &= l\left(\begin{pmatrix} \mathbf{0} \\ \mathbf{w}^* \end{pmatrix}, (b', \hat{\mathbf{z}})\right) \\ &= f_m(b', w_1^* z_1, \dots, w_m^* z_m). \end{aligned} \quad (4)$$

The last equality follows from Equation (3) easily. By feature-wise symmetry (Assumption 1(iv)), we have

$$\begin{pmatrix} \mathbf{w}^* \\ \mathbf{0} \end{pmatrix} \in \mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})}. \quad (5)$$

Hence there exists $\mathbf{w}^2 \in \mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})}$ such that

$$\begin{aligned} l(\mathbf{w}^2, (b', \hat{\mathbf{z}})) &\leq l\left(\begin{pmatrix} \mathbf{w}^* \\ \mathbf{0} \end{pmatrix}, (b', \hat{\mathbf{z}})\right) \\ &= f_m(b', 0, \dots, 0). \end{aligned} \quad (6)$$

The last equality follows from Equation (3). The inequality here holds because by Assumption 2(ii), if

there is no $\mathbf{w}^2 \in \mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})}$ that satisfies the inequality, then by (5) and definition of $\tilde{\mathbf{b}}$ and \tilde{A} we have

$$\mathbf{w}^2 \preceq_{(\tilde{\mathbf{b}}, \tilde{A})} \begin{pmatrix} \mathbf{w}^* \\ \mathbf{0} \end{pmatrix}$$

which implies that

$$\begin{pmatrix} \mathbf{w}^* \\ \mathbf{0} \end{pmatrix} \in \mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})},$$

from the optimality of \mathbf{w}^2 . However, this is a contradiction of the assumption that there is no $\mathbf{w}^2 \in \mathbb{L}_{(\tilde{\mathbf{b}}, \tilde{A})}$ that satisfies the inequality of (6).

Combining (4) and (6) proves the theorem. \square

Other Generalizations

While this paper focuses on the case where a learned solution takes a vector form, it is straightforward to generalize the setup to the matrix case and show that a similar no-free-lunch theorem between stability and group sparsity holds. As an example, consider the following group-sparse algorithm: Minimize: $\|B - AW\|_F + \|W\|_{1,2}$, where $\|W\|_{1,2}$ is the summation of the ℓ_2 norm of each row of W . Then, treating *each row* of W as the value of a feature of the solution and following a similar argument as the proof of Theorem 1, one can show that such a group sparse algorithm cannot be stable. Due to space constraints, we do not elaborate.

5 DISCUSSION

To see that the two notions of stability and sparsity that we consider are not too restrictive, we list in this section some algorithms that either admit a diminishing uniform stability bound or identify redundant features. Thus, by applying Theorem 2 we conclude that they are either non-sparse or non-stable.

5.1 Stable algorithms

All algorithms listed in this section have a uniform stability bound that decreases as $O(\frac{1}{n})$, and are hence stable. Examples 1 to 5 and adapted from [1].

Example 1 (Bounded SVM regression): Assume k is a bounded kernel, that is $k(\mathbf{x}, \mathbf{x}) \leq \kappa^2$. Let \mathcal{F} denote the RKHS space of k . Consider $\mathcal{Y} = [0, B]$ and the loss function

$$\begin{aligned} l(f, (y, \mathbf{x})) &= |f(\mathbf{x}) - y|_\epsilon \\ &= \begin{cases} 0 & \text{if } |f(\mathbf{x}) - y| \leq \epsilon; \\ |f(\mathbf{x}) - y| - \epsilon & \text{otherwise.} \end{cases} \end{aligned}$$

The SVM regression algorithm with kernel k is defined as

$$\mathbb{L}_S = \arg \min_{g \in \mathcal{F}} \left\{ \sum_{i=1}^n l(g, (y_i, \mathbf{x}_i)) + \lambda n \|g\|_\kappa^2 \right\}$$

where $S = ((y_1, \mathbf{x}_1), \dots, (y_n, \mathbf{x}_n))$. Then, its uniform stability satisfies

$$\beta_n \leq \frac{\kappa^2}{2\lambda n}.$$

Example 2 (Soft-margin SVM classification): Assume k is a bounded kernel, that is $k(\mathbf{x}, \mathbf{x}) \leq \kappa^2$. Let \mathcal{F} denote the RKHS space of k . Consider $\mathcal{Y} = \{0, 1\}^1$ and the loss function

$$l(f, (y, \mathbf{x})) = (1 - (2y - 1)f(\mathbf{x}))^+ \\ = \begin{cases} 1 - (2y - 1)f(\mathbf{x}) & \text{if } 1 - (2y - 1)f(\mathbf{x}) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

The soft-margin SVM (without bias) algorithm with kernel k is defined as

$$\mathbb{L}_S = \arg \min_{g \in \mathcal{F}} \left\{ \sum_{i=1}^n l(g, (y_i, \mathbf{x}_i)) + \lambda n \|g\|_\kappa^2 \right\},$$

where $S = ((y_1, \mathbf{x}_1), \dots, (y_n, \mathbf{x}_n))$. Its uniform stability satisfies

$$\beta_n \leq \frac{\kappa^2}{2\lambda n}.$$

Example 3 (RKHS regularized least square regression): Assume k is a bounded kernel, that is $k(\mathbf{x}, \mathbf{x}) \leq \kappa^2$. Let \mathcal{F} denote the RKHS space of k . Consider $\mathcal{Y} = [0, B]$ and the loss function

$$l(f, (y, \mathbf{x})) = (f(\mathbf{x}) - y)^2.$$

The regularized least square regression algorithm with kernel k is defined as

$$\mathbb{L}_S = \arg \min_{g \in \mathcal{F}} \left\{ \sum_{i=1}^n l(g, (y_i, \mathbf{x}_i)) + \lambda n \|g\|_\kappa^2 \right\},$$

where $S = ((y_1, \mathbf{x}_1), \dots, (y_n, \mathbf{x}_n))$. Its uniform stability satisfies

$$\beta_n \leq \frac{2\kappa^2 B^2}{\lambda n}.$$

The next example is relative entropy regularization. In this case, we are given a class of base hypotheses, and the output of the algorithm is a mixture of them, or more precisely a probability distribution over the class of base hypotheses.

Example 4 (Relative Entropy Regularization): Let $\mathcal{H} = \{h_\theta : \theta \in \Theta\}$ be the class of base hypotheses, where Θ is a measurable space with a reference measure. Let \mathcal{F} denote the set of probability distributions over Θ dominated by the reference measure. Consider the loss function for $f \in \mathcal{F}$

$$l(f, \mathbf{z}) = \int_{\Theta} r(h_\theta, \mathbf{z}) f(\theta) d\theta;$$

where $r(\cdot, \cdot)$ is a loss function bounded by M . Further, let f_0 be a fixed element of \mathcal{F} and $K(\cdot, \cdot)$ denote

1. This is slightly different from but equivalent to the standard setup where $\mathcal{Y} = \{-1, 1\}$.

the Kullback-Leibler divergence. The relative entropy regularized algorithm is defined as

$$\mathbb{L}_S = \arg \min_{g \in \mathcal{F}} \left\{ \sum_{i=1}^n l(g, \mathbf{z}_i) + \lambda n K(g, f_0) \right\},$$

where $S = (\mathbf{z}_1, \dots, \mathbf{z}_n)$. Then, its uniform stability satisfies

$$\beta_n \leq \frac{M^2}{\lambda n}.$$

A special case of relative entropy regularization is the following *maximum entropy discrimination* proposed in [18].

Example 5 (Maximum entropy discrimination): Let $\mathcal{H} = \{h_{\theta, \gamma} : \theta \in \Theta, \gamma \in \mathbb{R}\}$. Let \mathcal{F} denote the set of probability distributions over $\Theta \times \mathbb{R}$ dominated by the reference measure. Consider $\mathcal{Y} = \{0, 1\}$ and the loss function

$$l(f, \mathbf{z}) = \left(\int_{\Theta, \mathbb{R}} [\gamma - (2y - 1)h_{\theta, \gamma}(\mathbf{x})] f(\theta, \gamma) d\theta d\gamma \right)_+,$$

where $[\gamma - (2y - 1)h_{\theta, \gamma}(\mathbf{x})]$ is bounded by M . The maximum entropy discrimination is a real-valued classifier defined as

$$\mathbb{L}_S = \arg \min_{g \in \mathcal{F}} \left\{ \sum_{i=1}^n l(g, \mathbf{z}_i) + \lambda n K(g, f_0) \right\};$$

where $S = (\mathbf{z}_1, \dots, \mathbf{z}_n)$. Its uniform stability satisfies

$$\beta_n \leq \frac{M}{\lambda n}.$$

If an algorithm is not stable, one way to stabilize it is to average its solutions trained on small bootstrap subsets of the training set, a process called subbagging [19], which we recall in the following example.

Example 6 (Subbagging, see Theorem 5.2 of [19]): Let \mathbb{L} be a learning algorithm with a stability β_n , and consider the following algorithm

$$\hat{\mathbb{L}}_D^k(\mathbf{x}) \triangleq \mathbb{E}_S(\mathbb{L}_S(\mathbf{x}))$$

where \mathbb{E}_S is the expectation with respect to k points sampled in \mathcal{D} uniformly *without replacement*. Then $\hat{\mathbb{L}}^k$ has a stability $\hat{\beta}_n$ satisfying

$$\hat{\beta}_n \leq \frac{k}{n} \beta_k.$$

In a recent work, [17] consider the uniform stability of ℓ_p regularization for $1 < p \leq 2$ and elastic net proposed in [20]. Their results imply the following examples.

Example 7 (ℓ_p regularization): Consider a collection of feature functions $(\varphi_\gamma(\cdot))_{\gamma \in \Gamma}$, where Γ is a countable set, such that for every $\mathbf{x} \in \mathcal{X}$,

$$\sum_{\gamma \in \Gamma} |\varphi_\gamma(\mathbf{x})|^2 \leq \kappa.$$

Let \mathcal{F} denote the linear span of the feature functions, i.e.,

$$\mathcal{F} = \left\{ \sum_{\gamma \in \Gamma} \alpha_\gamma \varphi_\gamma(\cdot) : \alpha \in \ell_2(\Gamma) \right\}.$$

Further assume that the loss function is such that $l(f, (y, \mathbf{x})) = V(f(\mathbf{x}), y)$, for some $V(\cdot, \cdot)$ that is convex, bounded, and Lipschitz continuous. That is,

- 1) V is convex.
- 2) For all y, y' we have $0 \leq V(y', y) \leq B$.
- 3) For all y_1, y_2, y , we have $|V(y_1, y) - V(y_2, y)| \leq L|y_1 - y_2|$.

Then, the ℓ_p regularization algorithm, defined as

$$\mathbb{L}_S = \arg \min_{\alpha \in \ell_2(\Gamma)} \left\{ \sum_{i=1}^n l\left(\sum_{\gamma \in \Gamma} \alpha_\gamma \varphi_\gamma(\cdot), (y_i, \mathbf{x}_i)\right) + \lambda n \sum_{\gamma \in \Gamma} |\alpha_\gamma|^p \right\};$$

where $S = ((y_1, \mathbf{x}_1), \dots, (y_n, \mathbf{x}_n))$, is uniformly stable with

$$\beta_n = \frac{1}{p(p-1)} \left(\frac{B}{\lambda}\right)^{(2-p)/p} \frac{4L^2\kappa}{n\lambda}.$$

Up to a constant, Examples 1 to 3 are special cases of Example 7 with $p = 2$. One interesting observation is that when $p = 1$ the stability bound breaks. As we know from previous sections, this is due to the sparsity of ℓ_1 regularization.

Example 8 (Elastic Net): Under the same assumptions as Example 7, the elastic-net regularization algorithm, defined as

$$\mathbb{L}_S = \arg \min_{\alpha \in \ell_2(\Gamma)} \left\{ \sum_{i=1}^n l\left(\sum_{\gamma \in \Gamma} \alpha_\gamma \varphi_\gamma(\cdot), (y_i, \mathbf{x}_i)\right) + \lambda n \sum_{\gamma \in \Gamma} (w_\gamma |\alpha_\gamma| + \epsilon \alpha_\gamma^2) \right\};$$

where $S = ((y_1, \mathbf{x}_1), \dots, (y_n, \mathbf{x}_n))$, for some $w_\gamma \geq 0$, is uniformly stable with

$$\beta_n = \frac{2L^2\kappa}{\epsilon n \lambda}.$$

Note that the weights $(w_\gamma)_{\gamma \in \Gamma}$ have no effect in the stability bound. This is easily expected as ℓ_1 regularization itself is not stable. Indeed, the stability bound of the elastic net coincides with that of a ℓ_2 regularization algorithm. One may easily check that because of the extra ℓ_2 norm, elastic nets do not enjoy the property of IRF.

We briefly comment on the last example. In [20] the authors proposed the elastic net and used the terminology ‘‘sparsity,’’ but with somewhat different meaning than the one we adopt here. Motivated by biomedical applications, the intention in [20] is precisely opposite from the goal of sparsity algorithms like compressed sensing (and the present paper). In contrast to compressed-sensing style sparsity that seeks a solution with the least nonzero elements, work in [20] in fact seeks algorithms that spread out weight among all similar features.² Therefore, this example

2. Indeed, because of the extra ℓ_2 term, in almost all instances, the elastic net would output a solution with at least the same number of non-zero coefficients as the ℓ_1 regularization, and sometimes output a much denser solution.

does not contradict the results of our main result, that sparse algorithms are not stable.

5.2 Sparse Algorithms

Next we list some algorithms that identify redundant features.

Example 9 (ℓ_0 Minimization): Subset selection algorithms based on minimizing ℓ_0 norm identify redundant features. One example of such an algorithm is the *canonical selection procedure* [21], which is defined as

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathbb{R}^m} \{ \|\mathbf{A}\mathbf{w} - \mathbf{b}\|_2 + \lambda \|\mathbf{w}\|_0 \}. \quad (7)$$

Proof: Note that if a solution \mathbf{w}^* achieves the minimum of (7) and has non-zero weights on two redundant features i and i' , then by constructing a $\hat{\mathbf{w}}$ such that $\hat{w}_i = w_i^* + w_{i'}^*$ and $\hat{w}_{i'} = 0$ we get a strictly better solution, which is a contradiction. Hence ℓ_0 minimizing algorithms are IRF. \square

Since in general finding the minimum of (7) is NP-hard [22], many algorithms rely on the ℓ_1 norm as a convex relaxation. Most such algorithms either minimize the ℓ_1 norm of the solution under the constraint of a regression error, or minimize the convex combination of some regression error and the ℓ_1 norm of the solution.

Example 10 (ℓ_1 Minimization): The following subset selection algorithms are based on minimizing the ℓ_1 norm to identify redundant features.

- 1) *Lasso* [11] defined as

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathbb{R}^m} \{ \|\mathbf{A}\mathbf{w} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{w}\|_1 \}.$$

And equivalently, the LARS algorithm [23] that solves Lasso.

- 2) *Basis Pursuit* [24] defined as the solution of the following optimization problem on $\mathbf{w} \in \mathbb{R}^m$:

$$\begin{aligned} \min : & \|\mathbf{w}\|_1 \\ \text{s.t.} : & \mathbf{A}\mathbf{w} = \mathbf{b}. \end{aligned}$$

- 3) *Dantzig Selector* [25] defined as

$$\begin{aligned} \text{Minimize:} & \quad \|\mathbf{w}\|_1 \\ \text{Subject to:} & \quad \|A^*(\mathbf{A}\mathbf{w} - \mathbf{b})\|_\infty \leq c. \end{aligned}$$

Here, A^* is the complex conjugate of A , and c is some positive constant.

- 4) 1-norm SVM [12], [13] defined as the solution of the following optimization problem on α, ξ, γ .

$$\begin{aligned} \min : & \|\alpha\|_1 + C \sum_{i=1}^n \xi_i \\ \text{s.t.} : & y_i \left\{ \sum_{j=1}^n \alpha_j k(\mathbf{x}_i, \mathbf{x}_j) + \gamma \right\} \geq 1 - \xi_i; \quad \forall i; \\ & \xi_i \geq 0; \quad \forall i. \end{aligned}$$

5) ℓ_1 norm SVM regression [26] defined as the solution of the following optimization problem on α , ξ and γ :

$$\begin{aligned} \min : & \|\alpha\|_1 + C \sum_{i=1}^n \xi_i \\ \text{s.t.} : & \left\{ \sum_{j=1}^n \alpha_j k(\mathbf{x}_i, \mathbf{x}_j) + \gamma \right\} - y_i \leq \varepsilon + \xi_i; \quad \forall i; \\ & y_i - \left\{ \sum_{j=1}^n \alpha_j k(\mathbf{x}_i, \mathbf{x}_j) + \gamma \right\} \leq \varepsilon + \xi_i; \quad \forall i; \\ & \xi_i \geq 0; \quad \forall i, \end{aligned}$$

where $\varepsilon > 0$ is a fixed constant.

Proof: Given an optimal \mathbf{w}^* we construct a new solution $\hat{\mathbf{w}}$ such that for any subset of redundant features I , $\sum_{i \in I} \mathbf{1}(\hat{w}_i \neq 0) \leq 1$ and $\sum_{i \in I} \hat{w}_i = \sum_{i \in I} w_i^*$. Thus, $\hat{\mathbf{w}}$ and \mathbf{w}^* are equally good, which implies that any ℓ_1 minimizing algorithm has at least one optimal solution that is IRF. Hence such algorithm is IRF by definition. \square

6 CONCLUSION

In this paper, we prove that sparsity and stability are at odds with each other. We show that if an algorithm is sparse, then its uniform stability is lower bounded by a nonzero constant. This also shows that any algorithmically stable algorithm cannot be sparse. Thus, we show that these two widely used concepts, namely *sparsity* and *algorithmic stability* contradict each other. At a high level, this theorem provides us with additional insight into these concepts and their interrelation, and it furthermore implies that a tradeoff between these two concepts is unavoidable in designing learning algorithms. Given that both sparsity and stability are desirable properties, one interesting direction is to understand the full implications of having one of them. That is, what other properties must a sparse solution have? Given that sparse algorithms often perform well, one may further ask for meaningful and computable notions of stability that are not in conflict with sparsity.

ACKNOWLEDGMENTS

This work was supported by NUS startup grant R-265-000-384-133, NSF (grants EFRI-0735905, CNS-0721532, CNS-0831580), DTRA (grant HDTRA1-08-0029), and Israel Science Foundation (contract 890015). A preliminary version of this work was presented in the Forty-Sixth Allerton Conference on Communication, Control, and Computing.

REFERENCES

- [1] O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of Machine Learning Research*, 2:499–526, 2002.
- [2] I. Steinwart. Consistency of support vector machines and other regularized kernel classifiers. *IEEE Transactions on Information Theory*, 51(1):128–142, 2005.
- [3] T. Poggio, R. Rifkin, S. Mukherjee, and P. Niyogi. General conditions for predictivity in learning theory. *Nature*, 428(6981):419–422, 2004.
- [4] S. Mukherjee, P. Niyogi, T. Poggio, and R. Rifkin. Learning theory: Stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization. *Advances in Computational Mathematics*, 25(1-3):161–193, 2006.
- [5] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability and stability in the general learning setting. In *Proceedings of 22nd Annual Conference of Learning Theory*, 2009.
- [6] F. Girosi. An equivalence between sparse approximation and support vector machines. *Neural Computation*, 10(6):1445–1480, 1998.
- [7] R. R. Coifman and M. V. Wickerhauser. Entropy-based algorithms for best basis selection. *IEEE Transactions on Information Theory*, 38(2):713–718, 1992.
- [8] S. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, 1993.
- [9] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.
- [10] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [11] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- [12] O. L. Mangasarian. Generalized support vector machines. In A. J. Smola, P. L. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 135–146. MIT Press, 2000.
- [13] J. Zhu, S. Rosset, T. Hastie, and R. Tibshirani. 1-norm support vector machines. In *Advances in Neural Information Processing Systems 16*, 2003.
- [14] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313:504–507, 2006.
- [15] A. d’Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *SIAM Review*, 49(3):434–448, 2007.
- [16] A. d’Aspremont, F. Bach, and L. El Ghaoui. Full regularization path for sparse principal component analysis. In *Proceedings of the Twenty-fourth International Conference on Machine Learning*, 2007.
- [17] A. Wibisono, L. Rosasco, and T. Poggio. Sufficient conditions for uniform stability of regularization algorithms. Technical Report MIT-CSAIL-TR-2009-060, Massachusetts Institute of Technology, 2009.
- [18] T. Jaakkola, M. Meila, and T. Jebara. Maximum entropy discrimination. In *Advances in Neural Information Processing Systems 12*, pages 470–476. MIT Press, 1999.
- [19] T. Evgeniou, M. Pontil, and A. Elisseeff. Leave one out error, stability, and generalization of voting combinations of classifiers. *Machine Learning*, 55(1):71–97, 2004.
- [20] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B*, 67(2):301–320, 2005.
- [21] D. P. Foster and E. I. George. The risk inflation criterion for multiple regression. *The Annals of Statistics*, 22:1947–1975, 1994.
- [22] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal of Computation*, 24:227–234, 1995.
- [23] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of Statistics*, 32(2):407–499, 2004.
- [24] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1999.
- [25] E. J. Candès and T. Tao. The Dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics*, 35(6):2313–2351, 2007.
- [26] B. Schölkopf and A. J. Smola. *Learning with Kernels*. MIT Press, 2002.



Huan Xu received the B.Eng. degree in automation from Shanghai Jiaotong University, Shanghai, China in 1997, the M.Eng. degree in electrical engineering from the National University of Singapore in 2003, and the Ph.D. degree in electrical engineering from McGill University, Canada in 2009. From 2009 to 2010, he was a postdoctoral associate at The University of Texas at Austin. Since 2011, he has been an assistant professor at the Department of Mechanical Engineering at the National University of Singapore. His research interests include statistics, machine learning, robust optimization, and planning and control.



Constantine Caramanis (M'06) received his Ph.D. in EECS from the Massachusetts Institute of Technology in 2006. Since then, he has been on the faculty in Electrical and Computer Engineering at The University of Texas at Austin. He received the NSF CAREER award in 2011. His current research interests include robust and adaptable optimization, machine learning and high-dimensional statistics, with applications to large scale networks.



Shie Mannor (S'00-M'03-SM'09) received the B.Sc. degree in electrical engineering, the B.A. degree in mathematics, and the Ph.D. degree in electrical engineering from the Technion-Israel Institute of Technology, Haifa, Israel, in 1996, 1996, and 2002, respectively. From 2002 to 2004, he was a Fulbright scholar and a postdoctoral associate at M.I.T. He was with the Department of Electrical and Computer Engineering at McGill University from 2004 to 2010 where he was a Canada Research chair in Machine Learning. He has been an associate professor at the Faculty of Electrical Engineering at the Technion since 2008. His research interests include machine learning and pattern recognition, planning and control, multi-agent systems, and communications.