
Learning Rules-First Classifiers

Deborah Cohen
Google Research

Amit Daniely
Google Research
Hebrew University

Amir Globerson
Google Research
Tel-Aviv University

Gal Elidan
Google Research
Hebrew University

Abstract

Complex classifiers may exhibit “embarrassing” failures in cases where humans can easily provide a justified classification. Avoiding such failures is obviously of key importance. In this work, we focus on one such setting, where a label is perfectly predictable if the input contains certain features, or rules, and otherwise it is predictable by a linear classifier. We define a hypothesis class that captures this notion and determine its sample complexity. We also give evidence that efficient algorithms cannot achieve this sample complexity. We then derive a simple and efficient algorithm and show that its sample complexity is close to optimal, among efficient algorithms. Experiments on synthetic and sentiment analysis data demonstrate the efficacy of the method, both in terms of accuracy and interpretability.

1 Introduction

The accuracy of machine learning algorithms has dramatically improved since the re-emergence of deep learning models. However, in many machine learning applications, the model will make “embarrassing” mistakes. Namely, mistakes on examples that a human would classify easily, and have a clear explanation for her decision. As a motivating example, consider a medical diagnosis system that, on average, performs better than the family doctor. However, every now and then, the system makes an embarrassing mistake and fails in a scenario where a simple mechanism can provide the correct diagnosis. As another example, consider an online streaming platform where it would be “embarrassing” not to recommend episode $i + 1$ to someone who is watching episode i in a series.

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

Clearly, we would like to avoid such mistakes. This is important for improving usability of learned models, and for making them more interpretable. A key challenge in addressing the above problem is defining the notion of an embarrassing mistake. From the viewpoint of standard statistical learning theory, all mistakes are identical, and one is not more embarrassing than the other. But, we can structure our hypothesis class such that “easy” cases are processed in an explainable way.

We take the first step toward an explicit formalization of this goal by considering easy examples to be those whose label is deterministic given certain values of a single feature (e.g., in the streaming example above, if we observed episode 3 of a series, we will want to watch episode 4). However, we clearly do not expect all samples to be classified using rules, and therefore allow the label to also result from a different classifier over the other features, when no rule applies. We call such hybrid models *rules-first classifiers*. Specifically, we consider the case where for a set \mathcal{K} of k “rule” features, the label is 1 if any feature in \mathcal{K} is non-negative. Otherwise, the label is determined by a linear classifier whose norm is bounded by B . We call such distributions (k, B) -realizable.

We investigate the computational and sample complexity of learning (k, B) -realizable distributions, and contrast these with related hypothesis classes defined by a bound on ℓ_1 or ℓ_2 norms. Specifically, we prove that the sample complexity of the problem is $\tilde{\Theta}\left(\frac{k+B^2}{\epsilon}\right)$. Interestingly, we show that this sample complexity is substantially better compared to that of the natural convex relaxation, which is $\tilde{\Omega}\left(\frac{kB^2}{\epsilon}\right)$.

After settling the statistical complexity for the problem, we investigate its computational complexity. We derive an efficient greedy algorithm for the classification task, and show that it enjoys a sample complexity of $\tilde{\Theta}\left(\frac{kB+B^2}{\epsilon}\right)$. While this sample complexity is much better compared to the natural convex relaxation, it is still inferior to the information theoretic limit of $\tilde{\Theta}\left(\frac{k+B^2}{\epsilon}\right)$.

Can better sample complexity be achieved by efficient algorithms? We give evidence that the answer is negative. Indeed, we show that an efficient algorithm whose sample complexity is better than $\tilde{\Theta}\left(\frac{kB+B^2}{\epsilon}\right)$ would lead to efficient algorithms for problems that are hypothesized to be hard.

The topic of rule learning has been studied in the past of course (e.g., see Rivest, 1987; Zhang and Zhang, 2002, and Section 6 for more details). Most of these approaches consider the case where *every* classification decision corresponds to activating a rule (e.g., for decision trees). Here we focus on the arguably more realistic setting whereby rules only apply to a subset of the cases, i.e. the easily explained examples, and other cases are covered by a function of all the features. To the best of our knowledge, we provide the first theoretical characterization of this rules-first setting.

2 Preliminaries

We begin with notation and relevant background. Throughout the paper, the following notations are used. The set of integers in $[1, d]$ is denoted by $[d]$ and the complement of a set \mathcal{K} is \mathcal{K}^c . We denote column vectors by boldface letters. The j th feature of \mathbf{x} is denoted $\mathbf{x}(j)$. The vector \mathbf{z} restricted to the set \mathcal{K} is $\mathbf{z}_{|\mathcal{K}}$ and $\langle \cdot, \cdot \rangle$ stands for the inner product. For $k \leq d$, we denote $\binom{[d]}{k} = \{A \subseteq [d] \mid |A| = k\}$ and $\binom{[d]}{\leq k} = \{A \subseteq [d] \mid |A| \leq k\}$. The ℓ_1 , ℓ_2 and ℓ_∞ norms are $\|\cdot\|_1$, $\|\cdot\|_2$ and $\|\cdot\|_\infty$, respectively. We also use the ℓ_0 -pseudonorm $\|\mathbf{x}\|_0 = |\{j \mid \mathbf{x}(j) \neq 0\}|$. We denote $\mathbb{B}_B^{d,p} = \{\mathbf{x} \in \mathbb{R}^d \mid \|\mathbf{x}\|_p \leq B\}$. For $p = 2$, let $\mathbb{B}_B^d = \mathbb{B}_B^{d,2}$.

Regularized Linear Classification Models: Consider the standard supervised classification problem. Let $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^m$ be a set of m training samples, drawn i.i.d. from some distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$, with $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{Y} = \{\pm 1\}$. To avoid measure theoretic subtleties, we assume that the support of \mathcal{D} is finite (none of the results will depend on its cardinality and all will hold for \mathcal{D} with infinite support). The goal is to find a classifier $h : \mathcal{X} \rightarrow \mathcal{Y}$ whose error $\text{Err}_{\mathcal{D}}(h) = \Pr_{(\mathbf{x}, y) \sim \mathcal{D}}(h(\mathbf{x}) \neq y)$ is as small as possible.

We consider classes of linear classifiers. Namely, classes of the form $\mathcal{H} = \{\mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle \mid \mathbf{w} \in \mathcal{W}\}$, for some $\mathcal{W} \subset \mathbb{R}^d$. Two typical choices of \mathcal{W} are the ℓ_1 and ℓ_2 balls, as well as combinations thereof such as the elastic-net ball (Zou and Hastie, 2005) where $\{\mathbf{w} : \|\mathbf{w}\|_1 \leq B_1 \text{ and } \|\mathbf{w}\|_2^2 \leq B_2\}$. Recently, Zadorozhnyi et al. (2016) also proposed the Huber-norm ball $\{\mathbf{w}_a + \mathbf{w}_b : \|\mathbf{w}_a\|_1 \leq B_1 \text{ and } \|\mathbf{w}_b\|_2^2 \leq B_2\}$ for $B_1, B_2 > 0$.

A popular approach for the classification problem is to

minimize a surrogate loss function. Namely, given a class \mathcal{H} of functions from \mathbb{R}^d to \mathbb{R} and a loss function $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow [0, \infty)$, solve $\min_{h \in \mathcal{H}} \ell(h, S)$, where $\ell(h, S) = \frac{1}{m} \sum_{i=1}^m \ell(h(\mathbf{x}_i), y_i)$. The expected loss ℓ of $h \in \mathcal{H}$ with respect to \mathcal{D} is $\ell(h, \mathcal{D}) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \ell(h(\mathbf{x}), y)$. The optimal true loss is $\ell(\mathcal{H}, \mathcal{D}) = \inf_{h \in \mathcal{H}} \ell(h, \mathcal{D})$, and the optimal empirical loss is $\ell(\mathcal{H}, S) = \inf_{h \in \mathcal{H}} \ell(h, S)$.

Popular loss functions are the mis-classification (zero-one) loss:

$$\ell_{\text{mis}}(\hat{y}, y) \triangleq \begin{cases} 0 & \hat{y} \cdot y > 0 \\ 1 & \hat{y} \cdot y \leq 0, \end{cases}$$

the margin loss where the above 0 threshold is relaxed to 1, the hinge loss $\ell_{\text{hinge}}(\hat{y}, y) \triangleq \max\{0, 1 - \hat{y} \cdot y\}$, and the ramp loss $\ell_{\text{ramp}}(\hat{y}, y) \triangleq \llbracket 1 - \hat{y} \cdot y \rrbracket$, where $\llbracket r \rrbracket \triangleq \max(0, \min(r, 1))$.

Note that the ramp loss is upper-bounded by the margin loss and lower-bounded by the mis-classification error. Therefore, whenever \mathcal{D} has a low large-margin loss, it also has low ramp loss. Likewise, once we find a hypothesis with small ramp loss, we also find a hypothesis with small mis-classification loss.

Sample Complexity Definitions: We now define the sample complexity of an algorithm and a hypothesis class with respect to a loss function, which we use later on to evaluate and compare different algorithms.

Definition 1 (Sample Complexity of Algorithm). Fix a hypothesis class \mathcal{H} . The sample complexity of an algorithm \mathcal{A} is the function $m_{\mathcal{A}} : (0, 1) \times (0, 1) \rightarrow \mathbb{N}$ so that $m_{\mathcal{A}}(\epsilon, \delta)$ is the minimal number for which the following holds: If $m \geq m_{\mathcal{A}}(\epsilon, \delta)$, then w.p. $\geq 1 - \delta$ over the choice of S and the internal randomness of \mathcal{A} , we have that $\ell_{\text{ramp}}(\mathcal{A}(S), \mathcal{D}) \leq \ell_{\text{ramp}}(\mathcal{H}, \mathcal{D}) + \epsilon$.

Definition 2 (Sample Complexity of Hypothesis Class). Fix a hypothesis class \mathcal{H} and a loss ℓ . The sample complexity of \mathcal{H} is $m_{\mathcal{H}}(\epsilon, \delta) = \min_{\mathcal{A}} m_{\mathcal{A}}(\epsilon, \delta)$.

We say that \mathcal{D} is realizable if $\ell_{\text{ramp}}(\mathcal{H}, \mathcal{D}) = 0$. Likewise, \mathcal{D} is η -realizable if $\ell_{\text{ramp}}(\mathcal{H}, \mathcal{D}) \leq \eta$. The realizable sample complexity of an algorithm and a class is defined similarly to the standard sample complexity, but restricted to realizable \mathcal{D} . We note that our definitions of sample complexity consider the ramp loss. This is motivated by the properties of the ramp loss noted above. From now on, we fix δ to be a small constant, and omit it from the complexity measures.

3 The Rules-First Learning Problem

We are now ready to formalize our learning problem. Recall that we would like to learn rules-first classifiers. Namely, classifiers whose outcome is either determined

via a small set of features, which are referred to as rules, or a bounded norm linear classifier on the remaining features. A simple such rule based case is when we have a set \mathcal{K} of k features such that the label is 1 if one of these features is positive, i.e.,

$$\Pr_{(\mathbf{x}, y) \sim \mathcal{D}} (y = 1 \mid \mathbf{x}(j) > 0 \text{ for some } j \in \mathcal{K}) = 1, \quad (1)$$

and otherwise the label is determined by a bounded norm linear classifier, i.e.,

$$\Pr_{(\mathbf{x}, y) \sim \mathcal{D}} (y \langle \mathbf{w}, \mathbf{x} \rangle \geq 1 \mid \mathbf{x}(j) \leq 0 \text{ for all } j \in \mathcal{K}) = 1, \quad (2)$$

where $\mathbf{w} \in \mathbb{R}^d$ with $\|\mathbf{w}\|_2^2 \leq B^2$.

Definition 3. A distribution \mathcal{D} is (k, B) -realizable if there is a set $\mathcal{K} \in \binom{[d]}{\leq k}$ and a weight vector \mathbf{w} for which (1) and (2) hold.

An equivalent notion with ℓ_1 regularization over \mathbf{w} may be defined, such that the results presented in the following sections transfer in the expected way.

In the above definition, a single rule can determine the label. We next consider a broader set of distributions, which we will use for deriving the sample complexity of (k, B) -realizable \mathcal{D} . Begin by noting that if \mathcal{D} is (k, B) -realizable then there are vectors \mathbf{w}_a with $\|\mathbf{w}_a\|_2^2 \leq B^2$ and \mathbf{w}_b with $\|\mathbf{w}_b\|_0 \leq k$ such that:¹

$$\Pr_{(\mathbf{x}, y) \sim \mathcal{D}} (y \langle \mathbf{w}_a + \mathbf{w}_b, \mathbf{x} \rangle \geq 1) = 1. \quad (3)$$

Motivated by this observation, we say that \mathcal{D} is (k, B) -weakly realizable if there exist norm bounded $\mathbf{w}_a, \mathbf{w}_b$ as above, such that (3) holds. A (k, B) -weakly realizable distribution can be realized by the following hypothesis class (we omit the dependence on k, B):

$$\mathcal{H}_{2,0} = \{\mathbf{x} \in \mathbb{B}_1^{d,2} \mapsto \langle \mathbf{w}_a + \mathbf{w}_b, \mathbf{x} \rangle \mid \|\mathbf{w}_a\|_2^2 \leq B^2, \|\mathbf{w}_b\|_0 \leq k\}. \quad (4)$$

Namely, \mathcal{D} is realizable by $\mathcal{H}_{2,0}$ if and only if it is (k, B) -weakly realizable. The hypothesis class $\mathcal{H}_{2,0}$ induces weight vectors composed of k unbounded entries (rules) and a remaining $d - k$ entries with bounded ℓ_2 norm. This drives the prediction to be dictated by the k features with highest weights, or rules, and in their absence, to be determined by a bounded linear classifier on the remaining features. Similarly to $\mathcal{H}_{2,0}$, we define:

$$\mathcal{H}_{1,0} = \{\mathbf{x} \in \mathbb{B}_1^{d,\infty} \mapsto \langle \mathbf{w}_a + \mathbf{w}_b, \mathbf{x} \rangle \mid \|\mathbf{w}_a\|_1 \leq B, \|\mathbf{w}_b\|_0 \leq k\}. \quad (5)$$

As we shall see, these *rules-first learning* formulations lead to sample complexity reduction as well as practical advantages. Specifically, the contributions of this work are as follows (ignoring logarithmic factors):

- We show that the sample complexity of (k, B) -realizable distributions is $\frac{k+B^2}{\epsilon}$.
- We derive an efficient and simple greedy algorithm for learning (k, B) -realizable distributions, with somewhat inferior sample complexity of $\frac{Bk+B^2}{\epsilon}$.
- We give evidence that the sample complexity of our greedy algorithm is close to optimal among efficient algorithms and show that it is better than that of the natural convex relaxation of the problem.
- We experiment with algorithms for the aforementioned scenario, comparing the greedy approach to the traditional ℓ_1 and ℓ_2 regularization approaches.

Taken together, our results indicate that the problem of learning rules-first classifiers exhibits an interesting statistical computational trade-off, and that efficient algorithms work well in practice.

4 Sample Complexity

In this section, we derive the sample complexity of the rule-based hypothesis classes $\mathcal{H}_{2,0}$ and $\mathcal{H}_{1,0}$ and use the former to obtain the sample complexity of (k, B) -realizable distributions.

Theorem 1. The sample complexity of $\mathcal{H}_{2,0}$ is $\tilde{O}\left(\frac{k \log d + B^2}{\epsilon^2}\right)$.

Theorem 2. The sample complexity of $\mathcal{H}_{1,0}$ is $\tilde{O}\left(\frac{k \log d + B^2 \log d}{\epsilon^2}\right)$.

To prove Theorem 1, we rely on the following result from Sabato et al. (2013), which considers the problem of distribution-dependent sample complexity. In their setting, the distribution of the input features has few directions in which the variance is high, but the combined variance in all other directions is small. With this assumption, they show that the sample complexity is characterized by the sum of the number of high-variance dimensions k and the squared norm in the other directions B^2 .

Formally, for $\mathcal{K} \in \binom{[d]}{k}$, let

$$\mathcal{Z}_{\mathcal{K}, B} = \{\mathbf{z} \in \mathbb{R}^d \mid \|\mathbf{z}_{|\mathcal{K}^c}\|_2^2 \leq B^2\}.$$

Consider the class

$$\mathcal{H}_{\mathcal{K}, B} = \{\mathbf{x} \in \mathcal{Z}_{\mathcal{K}, B} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle \mid \mathbf{w} \in \mathcal{Z}_{\mathcal{K}, 1}\}. \quad (6)$$

Then, Sabato et al. (2013) show the following result.

Proposition 1. For any $\delta \in (0, 1)$, with probability $1 - \delta$, every $h \in \mathcal{H}_{\mathcal{K}, B}$ satisfies

$$\ell_{\text{ramp}}(h, \mathcal{D}) \leq \ell_{\text{ramp}}(h, S) + \sqrt{\frac{O(k + B^2) \ln(m)}{m}} + \sqrt{\frac{8 \ln(2/\delta)}{m}}. \quad (7)$$

¹To see this, note that one can take $\mathbf{w}_a = \mathbf{w}$ and \mathbf{w}_b to be the indicator vector of \mathcal{K} , multiplied by a large enough scalar since we assume that \mathcal{D} has finite support.

The above result focuses on the class $\mathcal{H}_{\mathcal{K},B}$ which makes an assumption on the input features \mathbf{x} . In our setting, we make a similar distributional assumption but on the conditional distribution of the *target variable* given the special set of features, i.e. the rules. Specifically, we wish to derive sample complexity bounds for the rule-based hypothesis classes (4) as well as (5).

To do so, we associate each example $\mathbf{x} \in \mathcal{Z}_{B,\mathcal{K}}$ with the example $\mathbf{x}' \in \mathcal{Z}_{1,\mathcal{K}}$ obtained by dividing each coordinate $i \in \mathcal{K}$ by B . We then have that the sample complexity of $\mathcal{H}_{\mathcal{K},B}$ is the same as that of the class

$$\mathcal{G}_{\mathcal{K},B} = \{\mathbf{x} \in \mathcal{Z}_{\mathcal{K},1} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle \mid \mathbf{w} \in \mathcal{Z}_{\mathcal{K},B}\}. \quad (8)$$

Now, since $\mathcal{H}_{2,0} \subset \cup_{\mathcal{K} \in \binom{[d]}{k}} (\mathcal{G}_{\mathcal{K},B})|_{\mathbb{B}^d}$, Proposition 1 and a union bound imply Theorem 1. A detailed proof as well as an adaptation for Theorem 2 are provided in the supplementary materials.

We note that both theorems are tight, up to logarithmic factors. Indeed, both $\mathcal{H}_{2,0}$ and $\mathcal{H}_{1,0}$ realize the class of k -disjunctions, which has sample complexity $\Omega\left(\frac{k}{\epsilon^2}\right)$. Likewise, $\mathcal{H}_{2,0}$ (respectively $\mathcal{H}_{1,0}$) contains the class of linear classifiers with ℓ_2 norm (respectively ℓ_1 norm) smaller than B , which has sample complexity $\Omega\left(\frac{B^2}{\epsilon^2}\right)$ (Anthony and Bartlett, 2009). Hence, both rule-based classes have sample complexity of $\Omega\left(\frac{B^2+k}{\epsilon^2}\right)$.

We also note that boosting (Freund and Schapire, 1997) implies that the realizable sample complexities of $\mathcal{H}_{2,0}$ and $\mathcal{H}_{1,0}$ are $\tilde{O}\left(\frac{k \log d + B^2}{\epsilon}\right)$ and $\tilde{O}\left(\frac{k \log d + B^2 \log d}{\epsilon}\right)$, respectively. Indeed, once we fix ϵ , the general sample complexity result yields a weak learner with sample complexity of $\tilde{O}(k + B^2)$. Applying boosting on top of it yields a strong learner with the above mentioned sample complexity guarantees in the realizable case.

As a corollary to Theorem 1 we obtain the sample complexity of learning (k, B) -realizable distributions. This follows from the equivalence of weak (k, B) -realizability to learning in $\mathcal{H}_{2,0}$, and the fact that (k, B) -realizability implies weak (k, B) -realizability.

Corollary 1. *The sample complexity of (k, B) -realizable distributions is $\tilde{O}\left(\frac{k \log d + B^2}{\epsilon}\right)$.*

5 Efficient Algorithms

The sample complexity obtained in the previous section may be achieved by using an ERM algorithm. Unfortunately, in the sequel we argue that it is unlikely that there is an efficient implementation of such an algorithm. Thus, we begin by proposing an efficient learning procedure, and provide corresponding sample complexity results. Further, we show our proposed algorithm dominates the natural regularization based approach to the problem.

5.1 An Efficient Greedy Algorithm

We start with the description of a greedy based algorithm and analysis of its sample complexity. Let S be a training sample. A *rule* is a coordinate $j \in [d]$ such that $y_i = 1$ whenever $\mathbf{x}_i(j) > 0$. We say that a rule j *covers* an example (\mathbf{x}, y) if $\mathbf{x}(j) > 0$. Consider the *GreedyRule* algorithm in Figure 1. Defining a similar algorithm with ℓ_1 regularization over \mathbf{w} is straightforward.

```

Initialize  $\mathcal{K} = \emptyset$  and let  $S_{\text{non-covered}} = S$ .
while there is a rule  $j \in [d]$  that covers more than
 $\frac{m}{100k(B+1)}$  examples from  $S_{\text{non-covered}}$  do
    - Add  $j$  to  $\mathcal{K}$ 
    - Discard samples covered by  $j$  from  $S_{\text{non-covered}}$ 
Find  $\mathbf{w}$  that minimizes the hinge loss on  $S_{\text{non-covered}}$ 
such that  $\|\mathbf{w}\|_2^2 \leq B^2$ 
    
```

Figure 1: The *GreedyRule* Algorithm

Now define *BoostRule* to be a boosting algorithm that uses *GreedyRule* as a weak learner.

Theorem 3. *BoostRule can learn (k, B) -realizable distributions with a sample complexity of $\tilde{O}\left(\frac{kB+B^2}{\epsilon}\right)$.*

We will prove this by showing, in the following lemma, that *GreedyRule* (Figure 1) is a weak learner. Namely, it is guaranteed to return a hypothesis with error $\frac{1}{2} - \Omega(1)$ whenever it runs on (k, B) -realizable distributions. The theorem will then be implied by boosting (Freund and Schapire, 1997). Indeed, applying boosting on top of a weak learner with sample complexity of M results in a strong learner with sample complexity of $\tilde{O}\left(\frac{M}{\epsilon}\right)$.

Lemma 1. *If \mathcal{D} is (k, B) -realizable and $m = \tilde{\Omega}(kB + B^2)$, then w.h.p. the greedy algorithm will return a hypothesis with error $\leq 1/4$.*

Proof. (sketch) We first note that upon termination of the algorithm, \mathcal{K} contains at most $100k(B+1)$ rules. Hence, the hypothesis returned by the algorithm belongs to $\mathcal{H}_{2,0}$ with $100k(B+1)$ instead of k . By Theorem 1 and the assumption that $m = \tilde{\Omega}(kB + B^2)$, it is enough to show that the empirical error is $\leq 1/5$. Indeed, for this amount of examples, Theorem 1 guarantees a generalization error smaller than $\frac{1}{4} - \frac{1}{5}$.

Since there are no mistakes on the covered examples, it is enough to show that at most $0.2m$ of the non-covered examples are mis-classified by the vector \mathbf{w} that was found in step 3. We will show an even stronger property. Namely, that

$$\sum_{(\mathbf{x}_i, y_i) \in S_{\text{non-covered}}} l_{\text{hinge}}(\langle \mathbf{w}, \mathbf{x}_i \rangle, y_i) \leq 0.2m.$$

Let $\mathcal{K}^* \in \binom{[n]}{\leq k}$ and \mathbf{w}^* be respectively a set and a vector given which S is (k, B) realizable. It is enough to show that

$$\sum_{(\mathbf{x}_i, y_i) \in S_{\text{non-covered}}} l_{\text{hinge}}(\langle \mathbf{w}^*, \mathbf{x}_i \rangle, y_i) \leq 0.2m.$$

To see that the last equation holds, let S_{covered}^* be the examples in S that are covered by the rules in \mathcal{K}^* . Denoting $U = S_{\text{non-covered}} \cap S_{\text{covered}}^*$, we have that

$$\begin{aligned} \sum_{\substack{(\mathbf{x}_i, y_i) \in \\ S_{\text{non-covered}}}} l_{\text{hinge}}(\langle \mathbf{w}^*, \mathbf{x}_i \rangle, y_i) &= \sum_{(\mathbf{x}_i, y_i) \in U} l_{\text{hinge}}(\langle \mathbf{w}^*, \mathbf{x}_i \rangle, y_i) \\ &\leq |U|(\|\mathbf{w}^*\| + 1) \leq |U|(B + 1), \end{aligned}$$

The first equality follows from the fact that, since S is (k, B) realizable for \mathcal{K}^* and \mathbf{w}^* , then there are no mistakes in $S_{\text{non-covered}}^*$. In other words, the only mistakes in $S_{\text{non-covered}}^*$ are in U . The result follows by noting that $|U| \leq \frac{m}{100(B+1)}$, since each rule in \mathcal{K}^* covers at most $\frac{m}{100(B+1)k}$ examples from $S_{\text{non-covered}}^*$, or step 2 would not terminate. \square

5.2 Theoretical Limitation of Regularization-based Approaches

An alternative approach to efficiently learning sparse classifiers is to replace the sparsity (i.e., ℓ_0) constraint with an ℓ_1 constraint, and show that the distribution at hand can be realized by low-norm linear classifiers (Ng, 2004). This suggests that we can try and learn (k, B) distributions by optimizing over $\mathcal{H}_{2,0}$ with the ℓ_0 norm replaced by ℓ_1 . Refer to this class as $\mathcal{H}_{2,1}$. The following lemma proves that this strategy is inferior to the greedy algorithm. Specifically, it results in lower bounded sample complexity $\Omega(\frac{kB^2}{\epsilon})$, which is larger than the upper bound on the greedy sample complexity.

To show that an algorithm has sample complexity of at most $\frac{C}{\epsilon}$, it suffices to show that there exists a distribution which can be realized by a linear classifier of squared norm at most C . Namely, there is a C -bounded norm linear function that is greater than 1 on positive points and smaller than -1 on negative points.

Lemma 2. *Let $B \geq 1$. There exists a (k, B) -realizable distribution \mathcal{D} such that*

1. *The marginal distribution of \mathcal{D} on \mathbb{R}^d is supported in $\mathbb{B}_1^{d,2}$ (and hence also in $\mathbb{B}_1^{d,\infty}$).*
2. *Any linear classifier that realizes \mathcal{D} with margin has squared ℓ_1 norm $\Omega(k^2B^2)$ and squared ℓ_2 norm $\Omega(kB^2)$.*

Proof. Let $d = k + B^2$ and let $\mathbf{a} = \frac{1}{B} \sum_{i=1}^{B^2} \mathbf{e}_{k+i}$. Consider the uniform distribution on

$$\left(\frac{\mathbf{e}_1 + \mathbf{a}}{\sqrt{2}}, 1 \right), \dots, \left(\frac{\mathbf{e}_k + \mathbf{a}}{\sqrt{2}}, 1 \right), (\mathbf{e}_{k+1}, -1), \dots, (\mathbf{e}_{k+B^2}, -1),$$

Clearly, the distribution is (k, B) -realizable. Likewise, if \mathbf{w} realizes \mathcal{D} , we must have $\mathbf{w}_i \leq -1$ for any $i =$

Table 1: Comparison of sample complexities. In the table, Comp. stands for complexity.

METHOD	EFFICIENT	SAMPLE COMP.
ERM	NO	$\tilde{O}\left(\frac{k+B^2}{\epsilon}\right)$
CONVEX RELAXATION	YES	$\tilde{\Omega}\left(\frac{kB^2}{\epsilon}\right)$
GREEDY	YES	$\tilde{O}\left(\frac{kB+B^2}{\epsilon}\right)$

$k+1, \dots, k+B^2$. It follows that $\langle \mathbf{w}, \mathbf{a} \rangle \leq -B$. Hence, we must have $\mathbf{w}_i \geq B$ for any $i = 1, \dots, k$. \square

To conclude the lower bound argument, we note that for learning in $\mathcal{H}_{2,1}$ we need to restrict the ℓ_2 norm to at least kB^2 to achieve the minimal sample complexity. The latter is thus lower bounded by $\Omega(\frac{kB^2}{\epsilon})$ as the upper bound on the sample complexity with respect to the ℓ_2 norm is tight (Anthony and Bartlett, 2009). The sample complexity results of the above algorithmic variants are summarized in Table 1. Note that these still hold for the non-boosting version of the listed algorithms when replacing ϵ with ϵ^2 in the sample complexity expressions (Freund and Schapire, 1997). This holds for the hardness result formulated in the next section as well.

5.3 Hardness

Having shown that our greedy approach is better in terms of sample complexity than a natural regularization based approach, we now show that in some sense we cannot do better than this greedy approach. In particular, we provide evidence that its sample complexity, namely $O\left(\frac{B^2+Bk \log(d)}{\epsilon}\right)$, is close to optimal among all efficient (poly(B, d) runtime) algorithms. Concretely, we will show that an efficient algorithm with sample complexity of $O\left(\frac{B^2+B^{1-\alpha}k \log(d)}{\epsilon}\right)$ for any $\alpha > 0$ would lead to a breakthrough in the extensively studied problem (e.g., see Shalev-Shwartz et al., 2010; Birnbaum and Shwartz, 2012; Daniely et al., 2014) of learning large margin classifiers with noise. To do so, we require a few additional definitions. We say that a distribution \mathcal{D} on $\mathbb{B}_1^d \times \{\pm 1\}$ is (η, B) -realizable if there exists $\mathbf{w} \in \mathbb{B}_B^d$ such that

$$\Pr_{(\mathbf{x}, y) \sim \mathcal{D}} (y \langle \mathbf{w}, \mathbf{x} \rangle \leq 1) \leq \eta(B).$$

The notion of (η, B) -realizable sample is defined similarly. We next describe the problem of learning large-margin classifiers with noise rate of $\eta: \mathbb{N} \rightarrow [0, \frac{1}{4}]$. We are given a norm bound $B \in \{1, 2, \dots\}$ and access to an $(\eta(B), B)$ -realizable distribution \mathcal{D} on $\mathbb{B}_1^d \times \{\pm 1\}$. The goal is to find a classifier with 0-1 error $\leq \frac{1}{4}$ in time poly(B).

This problem and variants have been studied extensively. Clearly, the problem becomes easier as η gets smaller. The best known algorithms (Birbaum and Schwartz, 2012) can tolerate noise of rate $\frac{\text{poly log}(B)}{B}$. Furthermore, there are lower bounds (Daniely et al., 2014) that show that for a large family of algorithms (specifically, generalized linear methods), better bounds cannot be achieved. Likewise, there are hardness results (Daniely, 2016) that show that, under certain complexity assumptions, no algorithm can tolerate a noise rate of $2^{-\log^{1-\alpha}(B)}$.

We will next show that algorithms for learning (k, B) -realizable distributions with sample complexity of $O(B^2 + B^{1-2\alpha}k \log(d))$ would lead to an algorithm for learning large margin classifiers with noise rate $\leq \frac{1}{B^{1-\alpha}}$, improving on the current state of the art. By boosting, this is true even if the algorithm is only required to return a hypothesis with non trivial performance (say, error at most 0.499) for (k, B) -realizable distributions. This serves as an indication that the sample complexity of $O(B^2 + Bk \log(d))$, achieved by our greedy algorithm, is close to optimal among efficient algorithms. A similar argument would rule out, under the complexity assumption from Daniely (2016), efficient algorithms that enjoy a sample complexity of $O(B^2 + 2^{\log^{1-\alpha}(B)}k \log(d))$.

We next sketch the argument. Suppose that \mathcal{A} is a learner for the problem of learning (k, B) -realizable distributions, with sample complexity of $O(B^2 + B^{1-2\alpha}k \log(d))$. Suppose now that \mathcal{D} is (η, B) -realizable with $\eta \leq \frac{1}{B^{1-\alpha}}$ and S is a sample consisting of $m = B^{10}$ points. We will generate a new sample $\Psi(S)$ by replacing (\mathbf{x}_i, y_i) with $((\mathbf{x}_i, \mathbf{e}_i), y_i)$, where $\mathbf{e}_i \in \mathbb{R}^m$ is the i th vector in the standard basis. It is not hard to verify that, with constant probability, $\Psi(S)$ is $(2\eta m, B)$ -realizable. Indeed, the original vector that testifies that \mathcal{D} is (η, B) -realizable will correctly classify about $m - \eta m$ examples with margin 1. The remaining examples can be handled using $k \approx \eta m$ rules. Now, since $m = \tilde{\Omega}(B^2 + B^{1-2\alpha}\eta m \log(m))$, \mathcal{A} will have non-trivial performance. This translates into a non-trivial performance on the original distribution for the large margin with noise problem.

We have thus shown in the last few sections that learning with rules, while inherently hard, does lead to sample complexity improvements, and can be learned in practice using a greedy algorithm that trades-off computational and statistical efficiency. As we shall see below, learning with rules is also beneficial in practice.

6 Related Work

A long history of works in machine learning is devoted to learning rules. Association rule learning (Zhang

and Zhang, 2002; Agrawal et al., 1993) is a rule-based method for discovering relations between variables, or rules, in large databases. Rules lists (Rivest, 1987; Clearwater and Provost, 1990; Letham et al., 2015) which consist of a series of *if... then...* statements, are a type of associative classifier, as the lists are formed from association rules. The *if* statements define a partition of a set of features, or rules, and the *then* statements correspond to the predicted outcome. Rules lists, or decision lists, generalize decision trees (Quinlan, 1993), in the sense that any decision tree can be expressed as a decision list, and any decision list is a one-sided decision tree (Letham et al., 2015).

All the above works assume that the data may be explained and *perfectly* classified via a set of relatively simple rules. In contrast, we propose a hybrid and more realistic framework, where labels are determined either by a set of simple rules or by a bounded-norm classifier in examples where the rules are not applicable. To the best of our knowledge, our work is the first to investigate the computational and sample complexity of this natural setting. In principle, one can augment a decision tree with linear classifier nodes (i.e., oblique decision trees) to handle such cases (Murthy et al., 1994). However, this would result in a different linear classifier for each rule. Furthermore, learning such trees cannot be done optimally, and does not result in performance guarantees like we have here.

The works of (Calderon et al., 2018; Juba, 2017) are concerned with a somewhat different goal as they seek to find subsets in the data, determined by k-Disjunctive Normal Form “rules” over some features, for which a good linear predictor can be found. In contrast, we aim at learning a linear classifier for all of the data unexplained by the rules. Our approach is also different from (Viola and Jones, 2001) that combines increasingly more complex classifiers in a cascade in contrast to our joint learning approach. The latter work does not provide theoretical results on the proposed method. Another relevant body of work considers learning with constraints. Abu-Mostafa (1993) deals with incorporating hints, or prior knowledge, such as invariance or oddness, in the learning process under the form of artificially generated examples.

An alternative approach to rule learning is to consider a sparse linear classifier. Since sparsity constraints are hard to enforce, a typical approach is to use ℓ_1 regularization as a surrogate for the ℓ_0 sparsity constraint. Under some conditions it can be shown (Ng, 2004) that this may result in tractable learning of rule based classifier. Similar results are available for online learning with the Winnow algorithm and its variants (Littlestone, 1988). However, these guarantees will no longer hold for the case of mixed rules and bounded

norm classifier as we consider here. An additional related line of work is on mixed norm regularization (e.g., see (e.g., see Zadorozhnyi et al., 2016; Zou and Hastie, 2005)), which uses both norms ℓ_1 and ℓ_2 . However, as we saw, such mixed regularization results in sample complexity bounds that are inferior to those obtained by our greedy algorithm.

7 Experimental Evaluation

We now empirically demonstrate the merit of our approach. We compare the performance of our ℓ_1 and ℓ_2 *GreedyRule* to traditional ℓ_1 and ℓ_2 penalties. We first consider binary classification on a synthetic dataset, generated with perfect rule features, and then turn to a real-life Twitter sentiment analysis based on the SemEval '17 task (Rosenthal et al., 2017).

Our greedy rule-based approach, described in Section 5, iteratively selects the feature that minimizes the current evaluation loss when added to the rules set. At each step, a regularized linear classifier is trained after removing the rule features. Prediction is then carried out first using these rules, and then by the learned classifier for examples where none of the rules apply.

For the non-rule part of our classifier, as well as the baseline classifiers, we consider a standard constrained logistic regression objective:

$$\min_{\mathbf{w}, c} \frac{1}{m} \sum_{i=1}^m \log(\exp(-y_i(\mathbf{w}^T \mathbf{x}_i + c)) + 1) + \frac{1}{C} R(\mathbf{w}), \quad (9)$$

where $1/C$ is the regularization strength parameter that trades-off training accuracy and regularization and the penalty $R(\mathbf{w})$ can be either $R(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|_2^2$ or $R(\mathbf{w}) = \|\mathbf{w}\|_1$. We use the logistic regression implementation of the scikit-learn library (Pedregosa et al., 2011) for both our greedy approach and the baseline linear classifier.

7.1 Synthetic Dataset

We generate m training samples with 400 standard features and $k = 20$ rule features. The rule features are i.i.d. Bernoulli random variables with parameter $p = 1/60$. The remaining features are i.i.d. random variables generated from a Gaussian distribution with $\mu = -0.02$ and $\sigma^2 = 1$. For each sample, $y = 1$ if one of the rule features is non zero and $y = \text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle)$ with $\mathbf{w}_j = 1$, for $1 \leq j \leq d$, otherwise. For the greedy algorithms, we use $2m/3$ samples for training and $m/3$ samples for evaluation to select the k rule features. We then retrain the chosen classifier on the m training samples. The test set is composed of 2000 samples, generated similarly to the training samples. The results are averaged over 20 realizations.




Figure 2: Accuracy comparison of the *GreedyRule* based approach and the corresponding baseline on the synthetic dataset as a function of the number of training sample m (left) and the number of rules $|\mathcal{K}|$ in the generating distribution (right).




Figure 3: Comparison of our method and the baseline classifier for two settings: using a base linear classifier or a neural network (NN). Shown is accuracy vs. the threshold used to pre-select the rule candidates.

Figure 2(left) shows the test accuracy of the different algorithms as a function of the number of training samples m . It can be seen that our greedy ℓ_2 and ℓ_1 algorithms outperform the traditional ℓ_2 and ℓ_1 regularized classifiers, as they succeed in finding the rule features. Appealingly, the gap between our approach and classic regularization is greater when m is smaller. In Figure 2(right), we show the accuracy of our greedy approaches as a function of $|\mathcal{K}|$, the number of rules allowed in *GreedyRule* (Figure 1). It can be clearly seen that increasing $|\mathcal{K}|$ towards the true k improves the performance while values beyond k decrease accuracy.

7.2 Sentiment Analysis - Twitter

We now turn to the SemEval-2017 Task 4 sub-task A (Rosenthal et al., 2017) of message polarity prediction. That is, given a Twitter message, the goal is to classify whether it has a positive, negative or neutral sentiment. We note that the task cannot be reconstructed precisely since some tweets become unavailable with time. We report results on binary polarity prediction (positive vs. negative), as they better demonstrate the effectiveness of rules. Results for three way classification (not shown) were similar in trend, and resemble state-of-the-art results on this problem (Rosenthal et al., 2017).

Figure 4: Examples of tweet prediction. The last column provides the stem which was considered as a rule by the greedy algorithm. 'Pos.' and 'Neg.' stand for positive and negative, respectively.

TWEET	GREEDY ℓ_2	ℓ_2	STEM
BOKO HARAM ON SATURDAY CLAIMED RESPONSIBILITY FOR ATTACKS IN CHAD'S CAPITAL	NEG.	POS.	ATTACK
DUSTIN JOHNSON DEALT WITH ANOTHER MAJOR DISAPPOINTMENT	NEG.	POS.	DISAPPOINT
RANDY ORTON IS ONE OF MY FAVORITES, DESPITE EVERYTHING. WHY? BECAUSE HE'S SO DAMN GOOD IN THE RING.	POS.	NEG.	GOOD
WE MAY BELIEVE WHATEVER WE WANT ABOUT GAY MARRIAGE. BUT GOD CALLS US TO LOVE, NOT TO CONDEMN.	POS.	NEG.	LOVE

The reduced dataset is composed of 3K training tweets, 1K evaluation tweets, and 8K are held out as test tweets. As a pre-processing step, we clean the text by removing links and special characters. We then use the SnowBall stemmer to transform each token into a stem. We adopt a bag-of-words representation, where the features of each example are a binary vector of appearance of tokens, or stems, in the tweet. The resulting stem dictionary is constructed with respect to the training and evaluation examples and contains about $d = 6.5K$ tokens.

Naturally, the dataset does not contain perfect rules, which requires us to pre-select candidates that are near rules. We first discard features that appear in less than 4 negative (16 positive) training tweets or for which $\hat{p}_{1,j} < 0.75$ ($\hat{p}_{0,j} < 0.9$), where $\hat{p}_{1|0,j}$ is the empirical probability that the label has value 1 or 0 given feature $x(j)$. Note that in this case, we consider rules for both labels $y = 1$ and $y = 0$. The discrepancy between the chosen thresholds reflect the data bias, which contains four times more positive than negative training examples. We then choose the top k rules, ordered by $\sqrt{M}p_{1|0,j}$, where M is the number of samples containing $x(j)$. In words, this balances between the nearness to rules and the coverage of the feature.

Figure 7.1 shows the overall accuracy of the standard and greedy ℓ_2 methods as a function of that threshold. Again, to take into account the data bias, we use the threshold as presented in the figure for rules inducing negative tweets, and four times that threshold for rules inducing positive tweets. The figure also presents the accuracy of a neural network classifier as well as a *GreedyRule* variation using the same non-linear classifier instead of a regularized logistic regression. For both our greedy approach and the baseline non-linear classifier, we use the neural network implementation of the scikit-learn library (Pedregosa et al., 2011) with 2 hidden layers with 5 and 2 neurons, respectively. We considered adding token pairs as features, e.g., to cope with the issue of negation. However, this did not improve the *GreedyRule* classifier's performance.

It may be observed that, as we allow more candidate rules by lowering the threshold, accuracy improves for both linear and non-linear *GreedyRule* classifiers. Below a certain threshold (approximately 1.5) accuracy begins

to decrease since non-rule features begin to be mistaken for rules, due to the data sparsity. We note that cross-validation on the evaluation data yields a threshold of 1.7, which corresponds to the highest accuracy in the test set as well. Although no theoretical results have been provided for the case of non-linear classifiers, the *GreedyRule* non-linear variation behaves similarly to its linear counterpart, as might have been expected.

The results are also quite appealing qualitatively. Stem rules chosen by the greedy algorithm have a clear sentiment semantics and include stems such as: happi, danger, evil, fail, excit, annoy, blame, loser, thank, magic, disappoint, failure, ruin, shame, stupid, love, terribl, worst, great, ridicul, disagree. Figure 7.2 shows a few test tweets for which our greedy ℓ_2 linear model does well but that are misclassified by traditional ℓ_2 .

8 Summary

In this work, we tackled the problem of learning rules-first classifiers. These, in addition to achieving high accuracy, do not make "embarrassing mistakes" where a simple explanation to the true label is possible, i.e., when the label can be accurately predicted from a single feature or rule. We formalized the notion of rules-based hypothesis classes, characterized the sample and computational complexity of learning with such classes and proposed an efficient greedy algorithm that trades-off computational and statistical complexity. Appealingly, its sample complexity is better than that of standard convex relaxation, and is likely optimal among all efficient algorithms. Finally, we demonstrated the benefit of our approach on simulated data as well as on a real-life tweets sentiment analysis task.

Our work is a first step toward an explicit formalization of the desideratum that the learning model does not make mistakes where good predictions can easily be achieved as well as explained. There are many intriguing directions for future developments, such as the obviously needed but non-trivial extension to soft rules and the adoption of a loss function that weights training points to account for embarrassing misclassification. More generally, we would like to learn under more flexible "embarrassment" requirements, such as ensuring the learn model does not make mistakes where simpler models do well.

References

- Yaser S Abu-Mostafa. A method for learning from hints. In *Advances in Neural Information Processing Systems*, pages 73–80, 1993.
- Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. Mining association rules between sets of items in large databases. In *ACM SIGMOD record*, volume 22, pages 207–216. ACM, 1993.
- Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. Cambridge University Press, 2009.
- Aharon Birnbaum and Shai S Shwartz. Learning halfspaces with the zero-one loss: time-accuracy tradeoffs. In *Advances in Neural Information Processing Systems*, pages 926–934, 2012.
- Diego Calderon, Brendan Juba, Zongyi Li, and Lisa Ruan. Conditional linear regression. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Scott H Clearwater and Foster J Provost. R14: A tool for knowledge-based induction. In *Tools for Artificial Intelligence, 1990., Proceedings of the 2nd International IEEE Conference on*, pages 24–30. IEEE, 1990.
- Amit Daniely. Complexity theoretic limitations on learning halfspaces. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 105–117. ACM, 2016.
- Amit Daniely, Nati Linial, and Shai Shalev-Shwartz. The complexity of learning halfspaces using generalized linear methods. In *Conference on Learning Theory*, pages 244–286, 2014.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Brendan Juba. Conditional sparse linear regression. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- Benjamin Letham, Cynthia Rudin, Tyler H McCormick, David Madigan, et al. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *The Annals of Applied Statistics*, 9(3):1350–1371, 2015.
- Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning*, 2(4):285–318, 1988.
- Sreerama K. Murthy, Simon Kasif, and Steven Salzberg. A system for induction of oblique decision trees. *Journal of Artificial Intelligence Research*, 2:1–32, 1994.
- Andrew Y Ng. Feature selection, ℓ_1 vs. ℓ_2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, page 78. ACM, 2004.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- J. R. Quinlan. *C4.5: Programs for Machine Learning*. Elsevier, 1993.
- Ronald L Rivest. Learning decision lists. *Machine learning*, 2(3):229–246, 1987.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. SemEval-2017 task 4: Sentiment analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation, SemEval '17*, Vancouver, Canada, August 2017. Association for Computational Linguistics.
- S. Sabato, N. Srebro, and Naftali Tishby. Distribution-dependent sample complexity of large margin learning. *The Journal of Machine Learning Research*, 14(1):2119–2149, 2013.
- Shai Shalev-Shwartz, Ohad Shamir, and Karthik Sridharan. Learning kernel-based halfspaces with the zero-one loss. *arXiv preprint arXiv:1005.3681*, 2010.
- Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. In *Conference on Computer Vision and Pattern Recognition*, page 511. IEEE, 2001.
- Oleksandr Zadorozhnyi, Gunthard Benecke, Stephan Mandt, Tobias Scheffer, and Marius Kloft. Huber-norm regularization for linear prediction models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 714–730. Springer, 2016.
- Chengqi Zhang and Shichao Zhang. *Association rule mining: models and algorithms*. Springer-Verlag, 2002.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.