

Label Efficient Learning by Exploiting Multi-class Output Codes

Maria Florina Balcan
School of Computer Science
Carnegie Mellon University
ninamf@cs.cmu.edu

Travis Dick
School of Computer Science
Carnegie Mellon University
tdick@cs.cmu.edu

Yishay Mansour
Microsoft Research and
Tel Aviv University
mansour@tau.ac.il

Abstract

We present a new perspective on the popular multi-class algorithmic techniques of one-vs-all and error correcting output codes. Rather than studying the behavior of these techniques for supervised learning, we establish a connection between the success of these methods and the existence of label-efficient learning procedures. We show that in both the realizable and agnostic cases, if output codes are successful at learning from labeled data, they implicitly assume structure on how the classes are related. By making that structure explicit, we design learning algorithms to recover the classes with low label complexity. We provide results for the commonly studied cases of one-vs-all learning and when the codewords of the classes are well separated. We additionally consider the more challenging case where the codewords are not well separated, but satisfy a boundary features condition that captures the natural intuition that every bit of the codewords should be significant.

1 Introduction

Motivation: Large scale multi-class learning problems with an abundance of unlabeled data are ubiquitous in modern machine learning. For example, an in-home assistive robot needs to learn to recognize common household objects, familiar faces, facial expressions, gestures, and so on in order to be useful. Such a robot can acquire large amounts of unlabeled training data simply by observing its surroundings, but it would be prohibitively time consuming (and frustrating) to ask its owner to annotate any significant portion of this raw data. More generally, in many modern learning problems we often have easy and cheap access to large quantities of unlabeled training data (e.g., on the internet) but obtaining high-quality labeled examples is relatively expensive. More examples include text understanding, recommendation systems, or wearable computing [Thrun, 1996, Thrun and Mitchell, 1995b,a, Mitchell et al., 2015]. The scarcity of labeled data is especially pronounced in problems with many classes, since supervised learning algorithms typically require labeled examples from every class. In such settings, algorithms should strive to make the best use of unlabeled data in order to minimize the need for expensive labeled examples.

Overview: We approach label-efficient learning by making the implicit assumptions of popular multi-class learning algorithms explicit and showing that they can also be exploited when learning from limited labeled data. We focus on a family of techniques called *output codes* that work by decomposing a given multi-class problem into a collection of binary classification tasks [Mohri et al., 2012, Dietterich and Bakiri, 1995, Langford and Beygelzimer, 2005, Beygelzimer et al., 2009]. The novelty of our results is to show that the existence of various low-error output codes constrains the distribution of unlabeled data in ways that can be exploited to reduce the label complexity of learning. We consider both the consistent setting, where the output code achieves zero error, and the agnostic setting, where the goal is to compete with the best output code. The most well known output code technique is one-vs-all learning, where we learn one binary classifier for distinguishing each class from the union of the rest. When output codes are successful at learning from labeled data, it often implies geometric structure in the underlying problem. For example, if it is possible to learn an accurate one-vs-all classifier with linear separators, it implies that no three classes can be collinear, since then it would be impossible for a single linear separator to distinguish the middle class from the union of the others. In this work exploit this implicitly assumed structure to design label-efficient algorithms for the commonly assumed cases of one-vs-all and error correcting output codes, as well as a novel boundary features condition that captures the intuition that every bit of the codewords should be significant.

Our results: Before discussing our results, we briefly review the output code methodology. For a problem with L classes, a domain expert designs a code matrix $C \in \{\pm 1\}^{L \times m}$ where each column partitions the classes into two meaningful groups. The number of columns m is chosen by the domain expert. For example, when recognizing household objects we could use the following true/false questions to define the partitions: “is it made of wood?”, “is it sharp?”, “does it have legs?”, “should I sit on it?”, and so on. Each row of the code matrix describes one of the classes in terms of these partitions (or semantic features). For example, the class “table” could be described by the vector $(+1, -1, +1, -1)$, which is called the class’ codeword. Once the code matrix has been designed, we train an output code by learning a binary classifier for each of the binary partitions (e.g., predicting whether an object is made of wood or not). To predict the class of a new example, we predict its codeword in $\{\pm 1\}^m$ and output the class with the nearest codeword under the Hamming distance. Two popular special cases of output codes are one-vs-all learning, where C is the identity matrix (with -1 in the off-diagonal entries), and error correcting output codes, where the Hamming distance between the codewords is large.

In each of our results we assume that there exists a consistent or low-error linear output code classifier and we impose constraints on the code matrix and the distribution that generates the data. We present algorithms and analysis techniques for a wide range of different conditions on the code matrix and data distribution to showcase the variety of implicit structures that can be exploited. For the code matrix, we consider the case when the codewords are well separated (i.e., the output code is error correcting), the case of one-vs-all (where the code matrix is the identity), and a natural boundary features condition. These conditions can loosely be compared in terms of the Hamming distance between codewords. In the case of error correcting output codes, the distance between codewords is large (at least $d + 1$ when the data is d -dimensional), in one-vs-all the distance is always exactly 2, and finally in the boundary features condition the distance can be as small as 1. In the latter cases, the lower Hamming distance requirement is balanced by other structure in the code matrix. For the distribution, we either assume that the data density function satisfies a thick level set condition or that the density is upper and lower bounded on its support. Both regularity conditions are used to ensure that the geometric structure implied by the consistent output code will be recoverable based on a sample of data.

Error correcting output codes: We first showcase how to exploit the implicit structure assumed by the commonly used and natural case of linear output codes where the Hamming distance between codewords is large. In practice, output codes are designed to have this property in order to be robust to prediction errors for the binary classification tasks [Dietterich and Bakiri, 1995]. We suppose that the output code makes at most β errors when predicting codewords and has codewords with Hamming distance at least $2\beta + d + 1$ in a d -dimensional problem. The key insight is that when the code words are well separated, this implies that points belonging to different classes must be geometrically separated as well. This suggests that tight clusters of data will be label-homogeneous, so we should be able to learn an accurate classifier using only a small number of label queries per cluster. The main technical challenge is to show that our clustering algorithm will not produce too many clusters (in order to keep the label complexity controlled), and that with high probability, a new sample from the distribution will have the same label as its nearest cluster. We show that when the data density satisfies a thick-level set condition (requiring that its level sets do not have bridges or cusps that are too thin), then a single-linkage clustering algorithm can be used to recover a small number of label-homogeneous clusters.

One-vs-all: Next, we consider the classic one-vs-all setting for data in the unit ball. This is an interesting setting because of the popularity of one-vs-all classification and because it significantly relaxes the assumption that the codewords are well separated (in a one-vs-all classifier, the Hamming distance between codewords is exactly 2). The main challenge in this setting is that there need not be a margin between classes and a simple single-linkage style clustering might group multiple classes into the same cluster. To overcome this challenge, we show that the classes are probabilistically separated in the following sense: after projecting onto the surface of the unit ball, the level sets of the projected density are label-homogeneous. Equivalently, the high-density regions belonging to different classes must be separated by low-density regions. We exploit this structure by estimating the connected components of the ϵ level set using a robust single-linkage clustering algorithm.

The boundary features condition: Finally, we introduce an interesting and natural condition on the code matrix capturing the intuition that every binary learning task should be significant. This condition has the weakest separation requirement, allowing the codewords to have a Hamming distance of only 1. This setting is our most challenging, since it allows for the classes to be very well connected to one another, which prevents clustering or level set estimation from being used to find a small number of label-homogeneous clusters. Nevertheless, we show that the implicit geometric

structure implied by the output code can be exploited to learn using a small number of label queries. In this case, rather than clustering the unlabeled sample, we apply a novel hyperplane-detection algorithm that uses the *absence* of data to learn local information about the boundaries between classes. We then use the implicit structure of the output code to extend these local boundaries into a globally accurate prediction rule.

Agnostic Setting: Finally, we show that our results for the error correcting, one-vs-all, and boundary features cases can all be extended to an agnostic learning setting, where we do not assume that there exists a consistent output code classifier.

Our results show an interesting trend: when linear output codes are able to learn from labeled data, it is possible to exploit the same underlying structure in the problem to learn using a small number of label requests. Our results hold under several natural assumptions on the output code and general conditions on the data distribution, and employ both clustering and hyperplane detection strategies to reduce the label complexity of learning.

2 Related Work

Reduction to binary classification is one of the most widely used techniques in applied machine learning for attacking multi-class problems. Indeed, the one-vs-all, one-vs-one, and the error correcting output code approaches [Dietterich and Bakiri, 1995] all follow this structure [Mohri et al., 2012, Langford and Beygelzimer, 2005, Beygelzimer et al., 2009, Daniely et al., 2012, Allwein et al., 2000].

There is no prior work providing error bounds for output codes using unlabeled data and interaction. There has been a long line of work for providing provable bounds for semi-supervised learning [Balcan et al., 2004, Balcan and Blum, 2010, Blum and Mitchell, 1998, Chapelle et al., 2010] and active learning [Balcan et al., 2006, Dasgupta, 2011, Balcan and Urner, 2015, Hanneke, 2014]. These works provide bounds on the benefits of unlabeled data and interaction for significantly different semi-supervised and active learning methods that are based different assumptions, often focusing on binary classification, thus the results are largely incomparable. Another line of recent work considers the multi-class setting and uses unlabeled data to consistently estimate the risk of classifiers when the data is generated from a known family of models [Donmez et al., 2010, Balasubramanian et al., 2011a,b]. Their results do not immediately imply learning algorithms and they consider generative assumptions, while in contrast our work explicitly designs learning algorithms under commonly used discriminative assumptions.

Another work related to ours is that of Balcan et al. [2013], where labels are recovered from unlabeled data. The main tool that they use, in order to recover the labels, is the assumption that there are multiple views and an underlying ontology that are known, and restrict the possible labeling. Similarly, Steinhardt and Liang [2016] show how to use the method of moments to estimate the risk of a model from unlabeled data under the assumption that the data has three independent views. Our work is more widely applicable, since it applies when we have only a single view.

The output-code formalism is also used by Palatucci et al. [2009] for the purpose of zero shot learning. They demonstrate that it is possible to exploit the semantic relationships encoded in the code matrix to learn a classifier from labeled data that can predict accurately even classes that *did not appear in the training set*. These techniques make very similar assumptions to our work but require that the code matrix C is known and the problem that they solve is different.

3 Preliminaries

We consider multiclass learning problems over an instance space $\mathcal{X} \subset \mathbb{R}^d$ where each point is labeled by $f^* : \mathcal{X} \rightarrow \{1, \dots, L\}$ to one out of L classes and the probability of observing each outcome $x \in \mathcal{X}$ is determined by a data distribution P on \mathcal{X} . The density function of P is denoted by $p : \mathcal{X} \rightarrow [0, \infty)$. In all of our results we assume that there exists a consistent (but unknown) linear output-code classifier defined by a code matrix $C \in \{\pm 1\}^{L \times m}$ and m linear separators h_1, \dots, h_m . We denote class i 's code word by C_i and define $h(x) = (\text{sign}(h_1(x)), \dots, \text{sign}(h_m(x)))$ to be the predicted code word for point x . We let $d_{\text{Ham}}(c, c')$ denote the Hamming distance between any codewords $c, c' \in \{\pm 1\}^m$. Finally, to simplify notation, we assume that the diameter of \mathcal{X} is at most 1.

Our goal is to learn a hypothesis $\hat{f} : \mathcal{X} \rightarrow \{1, \dots, L\}$ minimizing $\text{err}_P(\hat{f}) = \Pr_{X \sim P}(\hat{f}(X) \neq f(X))$ from an unlabeled sample drawn from the data distribution P together with a small set of actively queried labeled examples.

Finally, we use the following notation throughout the paper: For any set A in a metric space (\mathcal{X}, d) , the σ -interior of A is the set $\text{int}_\sigma(A) = \{x \in A : B(x, \sigma) \subset A\}$. The notation $\tilde{O}(\cdot)$ suppresses logarithmic terms.

4 Error Correcting Output Codes

We first consider the implicit structure when there exists a consistent linear *error correcting* output code classifier:

Assumption 1. *There exists a code matrix $C \in \{\pm 1\}^{L \times m}$ and linear functions h_1, \dots, h_m such that: (1) there exists $\beta \geq 0$ such that any point x from class y satisfies $d_{\text{Ham}}(h(x), C_y) \leq \beta$, (2) The Hamming distance between the codewords of C is at least $2\beta + d + 1$; and (3) at most d of the separators h_1, \dots, h_m intersect at any point.*

Part (1) of this condition is a bound on the number of linear separators that can make a mistake when the output code predicts the codeword of a new example, part (2) formalizes the requirement of having well separated codewords, and part (3) requires that the hyperplanes be in general position, which is a very mild condition that can be satisfied by adding an arbitrarily small perturbation to the linear separators.

Despite being very natural, Assumption 1 conveniently implies that there exists a distance $g > 0$ such that any points that f^* assigns to different classes must be at least distance g apart. To see this, fix any pair of points x and x' with $f^*(x) \neq f^*(x')$. By the triangle inequality, we have that $d_{\text{Ham}}(h(x), h(x')) \geq d + 1$, implying that the line segment $[x, x']$ crosses at least $d + 1$ of the linear separators. Since only d linear separators can intersect at a point, the line segment must have non-zero length. Applying this argument to the closest pair of points between all pairs of classes and taking the minimum length gives the result. A formal proof is given Section 9 of the Appendix.

Lemma 1. *Under Assumption 1, there exists $g > 0$ s.t. if points x and x' belong to different classes, then $\|x - x'\| > g$.*

Lemma 1 suggests that we should be able to reduce the label complexity of learning by clustering the data and querying the label of each cluster, since nearby points must belong to the same class. If we use a single-linkage style clustering algorithm that merges clusters whenever their distance is smaller than g , we are guaranteed that the clusters will be label-homogeneous, and therefore we can recover nearly all of the labels by querying one label from the largest clusters. See Algorithm 1 for pseudocode.

Input: Sample $S = \{x_1, \dots, x_n\}$, radius $r_c > 0$, target error $\epsilon > 0$

1. Let $\{\hat{A}_i\}_{i=1}^N$ be the connected components of the graph G with vertex set S and an edge between x_i and x_j if $\|x_i - x_j\| \leq r_c$.
2. In decreasing order of size, query the label of each \hat{A}_i until $\leq \frac{\epsilon}{4}n$ points belong to unlabeled clusters.
3. Output $\hat{f}(x) =$ label of nearest labeled cluster to x .

Algorithm 1: Single-linkage learning.

In order to get a meaningful reduction in label complexity, we need to ensure that when we cluster a sample of data, most of the samples will belong to a small number of clusters. For this purpose, we borrow the following very general and interesting thick level set condition from Steinwart [2015]: a density function p has C -thick level sets if there exists a level $\lambda_0 > 0$ and a radius $\sigma_0 > 0$ such that for every level $\lambda \leq \lambda_0$ and radius $\sigma < \sigma_0$, (1) the σ -interior of $\{p \geq \lambda\}$ is non-empty and (2) every point in $\{p \geq \lambda\}$ is at most distance $C\sigma$ from the σ -interior. This condition elegantly characterizes a large family of distributions for which single-linkage style clustering algorithms succeed at recovering the high-density clusters and only rules out distributions whose level sets have bridges or cusps that are too thin. The thickness parameter C measures how pointed the boundary of the level sets of p can be. For example, in \mathbb{R}^d if the level set of p is a ball then $C = 1$, while if the level set is a cube, then $C = \sqrt{d}$.

Using the thick level set condition to guarantee that our clustering algorithm will not subdivide the high-density clusters of p , we obtain the following result for Algorithm 1

Theorem 1. *Suppose that Assumption 1 holds and that the data distribution has C -thick level sets. For any target error $\epsilon > 0$, let N be the number of connected components of $\{p \geq \epsilon/(2 \text{Vol}(K))\}$. With probability at least $1 - \delta$, running Algorithm 1 with parameter $r_c < g$ on an unlabeled sample of size $n = \tilde{O}(\frac{1}{\epsilon^2}((4C)^{2d}d^{d+1}/r_c^{2d} + N))$ will query at most N labels and output a classifier with error at most ϵ .*

Proof. For convenience, define $\sigma = r_c/(4C)$ and $\lambda = \epsilon/(2 \text{Vol}(K))$. Using a standard VC-bound [Vapnik and Chervonenkis, 1971] together with the fact that balls have VC-dimension $d + 1$, for $n = O((4C)^{2d}d^{d+1}/(\epsilon^2 r_c^{2d}))$ guarantees that with probability at least $1 - \delta/2$ the following holds simultaneously for every center $x \in \mathbb{R}^d$ and radius $r \geq 0$:

$$\left| |B(x, r) \cap S|/n - P(B(x, r)) \right| \leq \frac{1}{2} \lambda \sigma^d v_d, \quad (1)$$

where v_d denotes the volume of the unit ball in \mathbb{R}^d . Assume that this high probability event occurs.

We first show that the sample S forms a $2C\sigma$ -covering of the set $\{p \geq \lambda\}$; that is, for every $x \in \{p \geq \lambda\}$ we have $d(x, S) \leq 2C\sigma$. Let x be any point in $\{p \geq \lambda\}$. Since p has C -thick level sets, we know that there exists a point $y \in \text{int}_\sigma(\{p \geq \lambda\})$ such that $\|x - y\| \leq C\sigma$. Moreover, the ball $B(y, \sigma)$ is contained in $\{p \geq \lambda\}$, which implies that it has probability mass at least $\lambda \sigma^d v_d$ and by (1) we have that $|B(y, \sigma) \cap S|/n \geq \frac{1}{2} \lambda \sigma^d v_d > 0$, so there must exist a point $z \in S \cap B(y, \sigma)$. Now we have that $d(x, S) \leq \|x - z\| \leq \|x - y\| + \|y - z\| \leq C\sigma + \sigma \leq 2C\sigma$, where the final inequality follows from the fact that $C \geq 1$.

Now let A_1, \dots, A_N be the N connected components of $\{p \geq \lambda\}$. We will argue that for each $i \in [N]$, there exists a unique cluster output by step 1 of the algorithm, say \hat{A}_i , such that \hat{A}_i contains $A_i \cap S$ and for any point $x \in A_i$, the closest output cluster is \hat{A}_i .

To see that \hat{A}_i contains $A_i \cap S$, consider any pair of points x and x' in $A_i \cap S$. Since A_i is connected, we know there is a path $\pi : [0, 1] \rightarrow A_i$ such that $\pi(0) = x$ and $\pi(1) = x'$. Since the sample set \mathcal{X} is a $2C\sigma$ covering of $\{p \geq \lambda\}$, it is also a $2C\sigma$ -covering of A_i , which implies that we can find a sequence of points $y_1, \dots, y_M \in \mathcal{X}$ (possibly with repetition) such that the path π passes through the balls $B(y_1, 2C\sigma), \dots, B(y_M, 2C\sigma)$ in order. Since consecutive balls must touch at the point that the path π crosses from one ball to the next, we know that $\|y_i - y_{i+1}\| \leq 4C\sigma = r_c$, and therefore the path $x \rightarrow y_1 \rightarrow \dots \rightarrow y_M \rightarrow x'$ is a path in the graph G connecting x and x' .

Now consider any point $x \in A_i$. We argued above that there exists a sample point $z \in \hat{A}_i$ that was within distance $2C\sigma$ from x . Now let z^* be the closest sample in \mathcal{X} to x . Then we know that $\|x - z^*\| \leq \|x - z\| \leq 2C\sigma$. By the triangle inequality, we have that $d\|z - z^*\| \leq \|z - x\| + \|x - z^*\| \leq 4C\sigma \leq r_c$, and therefore z and z^* are connected in the graph G . Since z belongs to \hat{A}_i , it follows that z^* does too, and therefore the closest cluster to x is \hat{A}_i .

It remains to bound the error of the resulting classification rule. Since there is a margin of width $g > 0$ separating the classes, we know that every connected component of $\{p \geq \lambda\}$ must contain points belonging to exactly one class. Moreover, since we ran the algorithm with connection radius $r_c < g$, we know that the clusters output by step 1 will contain points belonging to exactly one class. It follows that if we query the label of any point in the cluster \hat{A}_i then the algorithm will not error on any test point in A_i . Say that one of the connected components A_i is labeled if we query the label of the corresponding cluster \hat{A}_i .

Applying Hoeffding's inequality and the union bound to all possible 2^N unions of the sets A_1, \dots, A_N , our value of n guarantees that with probability at least $1 - \delta/2$, the following holds simultaneously for all subsets of indices $I \subset [N]$:

$$\left| |S \cap \left(\bigcup_{i \in I} A_i\right)|/n - P\left(\bigcup_{i \in I} A_i\right) \right| \leq \frac{\epsilon}{4}.$$

Since the algorithm queries labels until at most $\frac{\epsilon}{4}n$ points belong to unlabeled clusters, we know that the number of samples belonging to the unlabeled A_i sets is at most $\frac{\epsilon}{4}n$. By the above uniform convergence, it follows that their total probability mass is at most $\epsilon/2$. Finally, since the algorithm only errors on test points in $\{p \leq \lambda\}$, which has probability mass at most $\epsilon/2$ or on unlabeled A_i sets, the error of the resulting classifier is at most ϵ . \square

The exponential dependence on the dimension in Theorem 1 is needed to ensure the sample S will be a fine covering of the level set of p w.h.p, which guarantees that Algorithm 1 will not subdivide its connected components into smaller clusters. When the data has low intrinsic dimensionality, the unlabeled sample complexity is only exponential in the intrinsic dimension. The following result shows that under the common assumption that the distribution is a doubling

measure, then the unlabeled sample complexity is exponential only in the doubling dimension. Recall that a probability measure P is said to have doubling dimension D if for every point x in the support of P and every radius $r > 0$, we have that $P(B(x, 2r)) \leq 2^D P(B(x, r))$ (see, for example, [Dasgupta and Sinha, 2013]).

Theorem 2. *Suppose that Assumption 1 holds the data distribution P has doubling dimension D , and the support of P has N connected components. With probability at least $1 - \delta$, running Algorithm 1 with parameter $r_c < g$ on a sample of size $n = \tilde{O}(d/r_c^{2D} + N/\epsilon^2)$ will query at most N labels and have error at most ϵ .*

Proof. Let x be any point in the support of P . Since we assumed that the diameter of \mathcal{X} is 1, we know that $\mathcal{X} \subset B(x, 1)$ and therefore $P(B(x, 1)) = 1$. Applying the doubling condition $\lg(r)$ times, it follows that for any radius $r > 0$ we have that $P(B(x, r)) \geq r^{-D}$.

As in the proof of Theorem 1, for our choice of n the following holds with probability at least $1 - \delta/2$ uniformly for every center x in \mathcal{X} and radius $r \geq 0$:

$$| |B(x, r) \cap S|/n - P(B(x, r)) | \leq \frac{1}{2} r^{-D}.$$

Assume this high probability event occurs. Since every ball of radius r centered at a point in the support of P has mass at least r^{-D} , each such ball must contain at least one sample point and it follows that the sample S forms an r^{-D} -covering of the support of P .

The rest of the proof now follows identically the proof of Theorem 1 with the A_1, \dots, A_N sets being the connected components of the support of P , since each connected component must be label-homogeneous. \square

The unlabeled sample complexity in Theorem 1 depends on the gap g between classes because we must have $r_c < g$. Such a scale parameter must appear in our results, since Assumption 1 is scale-invariant, yet our algorithm exploits scale-dependent geometric properties of the problem. If we have a conservatively small estimate $\hat{g} \leq g$, then the conclusion of Theorem 1 and Theorem 2 continue to hold if the connection radius and unlabeled sample complexity are set using the estimate \hat{g} . Nevertheless, in some cases we may not have an estimate of g , making it difficult to apply Algorithm 1. The following result shows that if we have an estimate of the number of high-density clusters, and these clusters have roughly balanced probability mass, then we are still able to take advantage of the geometric structure even when the distance g is unknown. The idea is to construct a hierarchical clustering of S using single linkage, and then to use a small number of label queries to find a good pruning.

Input: Sample $S = \{x_1, \dots, x_n\}$, $t \in \mathbb{N}$.

1. Let T be the hierarchical clustering of S obtained by single-linkage.
2. Query the labels of a random subset of S of size t .
3. Let $\{\hat{B}_i\}_{i=1}^M$ be the coarsest pruning of T such that each \hat{B}_i contains labels from one class.
4. Output $\hat{f}(x) = \text{label of nearest } \hat{B}_i \text{ to } x$.

Algorithm 2: Hierarchical single-linkage learning.

Theorem 3. *Suppose Assumption 1 holds and the density p has C -thick level sets. For any $0 < \epsilon \leq 1/2$, suppose that $\{A_i\}_{i=1}^N$ are the connected components of $\{p \geq \epsilon/(2 \text{Vol}(K))\}$ and for some $\alpha \geq 1$ we have $P(A_i) \leq \alpha P(A_j)$ for all i, j . With probability $\geq 1 - \delta$, running Algorithm 2 with $t = \tilde{O}(\alpha N)$ on an unlabeled sample of size $n = \tilde{O}(\frac{1}{\epsilon^2}(C^{2d}d^{d+1}/g^{2d} + N))$ will have error $\leq \epsilon$.*

Proof. Define $\lambda = \epsilon/(2 \text{Vol}(K))$ and let A_1, \dots, A_N be the connected components of $\{p \geq \lambda\}$. Suppose that each A_i set has probability mass at least γ . Under the assumption that the probability mass of the largest A_i is at most α times the mass of the smallest, we have that $\gamma \geq (1 - \epsilon)/(\alpha N)$, but the result holds for any arbitrary lower bound γ .

Since we query the labels of points without replacement, the set of labeled examples is an iid sample from the data density p . Whenever $m \geq \frac{2}{\gamma} \ln \frac{2N}{\delta}$, with probability at least $1 - \delta/2$, every set A_i will contain at least one labeled example, since they each have probability mass at least γ . Assume this high probability event holds.




Figure 1: An example problem satisfying Assumption 2 and the projected density q when the density p is uniform on K .

Let g be the margin between classes that is guaranteed by Lemma 1. Whenever samples $x, x' \in S$ have $\|x - x'\| \leq g$, they must belong to the same cluster \hat{B}_i . Applying an identical covering-style argument as in Theorem 1, we have that with probability at least $1 - \delta/2$, for every A_i set there is a cluster, say \hat{B}_i , such that:

1. All samples in $A_i \cap S$ are contained in \hat{B}_i .
2. For every $x \in A_i$, the nearest cluster to x is \hat{B}_i .

Since every A_i set contains at least one labeled example, it follows that whenever two of these high-density clusters belong to different classes, they will contain differently labeled points and therefore will not have been merged by Algorithm 2. It follows that the label of \hat{B}_i must agree with the label of A_i . At this point, the error analysis follows identically as in Theorem 1. \square

In Section 7 we describe a meta-argument that can be used to extend our results into the agnostic setting, where we no longer require that the output code is consistent. Details for the error correcting case are given in Section 9.

In this section we showed that when there exists linear error correcting correcting output code with low error, then it is possible to reduce the label complexity of learning to the number of high-density clusters, which are the connected components of $\{p \geq \epsilon\}$. The label complexity of our algorithms is always linear in the number of high density clusters, while the worst-case unlabeled complexity of our algorithms is exponential in the dimension (or intrinsic dimension).

5 One-Versus-All on the Unit Ball

In this section we show that even when the codewords are not well separated, we can still exploit the implicit structure of output codes to reduce the label complexity of learning by clustering the data. Specifically, we consider the implicit structure of a linear one-vs-all classifier over the unit ball:

Assumption 2. *The instance space \mathcal{X} there exist L linear separators h_1, \dots, h_L such that: (1) point x belongs to class i iff $h_i(x) > 0$, and (2) for all i , $h_i(x) = w_i^\top x - b_i$ with $\|w_i\| = 1$ and $b_i \geq b_{\min} > 0$.*

See Figure 1 for an example problem satisfying this condition. Since a one-vs-all classifier is an output code where the code matrix is the identity, the Hamming distance between any pair of codewords is exactly 2. Therefore, in this setting we do not have a result similar to Lemma 1 to ensure that the classes are geometrically separated. Instead, we exploit the one-vs-all structure to show the classes are probabilistically separated and employ a robust clustering algorithm.

As before, we study this problem under a mild constraint on the data distribution. For each class i denote the set of points in class i by $K_i = \{x : \|x\| \leq 1, h_i(x) > 0\}$ and $K = \bigcup_{i=1}^L K_i$. In this section, we assume that the density p is supported on K with upper and lower bounds:

Assumption 3. *There exist constants $0 < c_{\text{lb}} \leq c_{\text{ub}}$ s.t. for $x \in K$ we have $c_{\text{lb}} \leq p(x) \leq c_{\text{ub}}$ and otherwise $p(x) = 0$.*

This distributional constraint is quite general: it only requires that we will not observe examples for which the one-vs-all classifier would be confused (i.e., where none of its linear separators claim the point) and that the density does not take extreme values. When K is compact, every continuous density supported on K satisfies Assumption 3.

Our algorithm for this setting first projects the data onto the unit sphere $\mathcal{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\| = 1\}$ and then applies a robust clustering algorithm to the projected data. The projection does not introduce any errors, since the label

of an example is independent of its distance to the origin. This is because each linear separator carves out a spherical cap for its class, and no two class caps overlap. Since we assume that no class contains the origin, it follows that an examples label depends only on its projection to the sphere. We show that projecting to the sphere has the useful property that the projected density goes to zero at the boundary of the classes, which suggests that we can use a robust single-linkage style clustering algorithm to find label-homogeneous clusters. Algorithm 3 gives pseudocode, using the notation $\theta(u, v) = \arccos(u^\top v)$ for the angle between u and v and $V^d(r)$ is the probability that a uniformly random sample from S^{d-1} lands in a given spherical cap of angular radius r .

Input: Sample $S = \{x_1, \dots, x_n\}$, radius $r_c > 0$.

1. Define $r_a = r_c/2$ and $\tau = \frac{c_{\text{lb}}}{2c_{\text{ub}}} V^d(r_a) \epsilon$.
2. Let $v_i = \frac{x_i}{\|x_i\|}$ be the projection of x_i to the sphere.
3. Mark v_i active if $|\{v_j : \theta(v_i, v_j) \leq r_a\}| \geq \tau n$ and inactive otherwise for $i \in [n]$.
4. Let $\hat{A}_1, \dots, \hat{A}_N$ be the connected components of the graph G whose vertices are the active v_i with an edge between v_i and v_j if $\theta(v_i, v_j) < r_c$.
5. In decreasing order of size, query the label of each \hat{A}_i until $\leq \frac{\epsilon}{4} n$ points belong to unlabeled clusters.
6. Output $\hat{f}(x) = \text{label of nearest cluster to } x/\|x\|$.

Algorithm 3: Robust single-linkage learning.

Our first result characterizes the density of the projected data (defined relative to the uniform distribution on S^{d-1}).

Lemma 2. *Suppose Assumptions 2 and 3 hold and let $q : S^{d-1} \rightarrow [0, \infty)$ be the density function of the data projected onto the unit sphere. Then $q_{\text{lb}}(v) \leq q(v) \leq q_{\text{ub}}(v)$, where*

$$q_{\text{lb}}(v) = \begin{cases} c_{\text{lb}} dv_d (1 - (b_i/w_i^\top v)^d) & \text{if } v \in K_i \\ 0 & \text{otherwise,} \end{cases}$$

and $q_{\text{ub}}(v) = c_{\text{ub}}/c_{\text{lb}} \cdot q_{\text{lb}}(v)$, where v_d is the volume of the unit ball in d dimensions.

Proof. Let $X \sim p$ be and set $V = X/\|X\|_2$ so that V is a sample from q . For any set $A \subset S^{d-1}$, we know that $\Pr(V \in A) = \Pr(X \in \text{cone}(A))$, where $\text{cone}(A) = \{rv : r > 0, v \in A\}$, which gives

$$\Pr(V \in A) = \Pr(X \in \text{cone}(A)) = \int_{x \in \text{cone}(A)} p(x) dx = \int_{v \in A} dv_d \int_{r=0}^{\infty} p(rv) r^{d-1} dr d\mu_{\circ}(v),$$

where the last inequality follows by a change of variables x to (r, v) where $r = \|x\|_2$ and $v = x/\|x\|_2$. The term r^{d-1} is the determinant of the Jacobian of the change of variables, and the term dv_d , which is the surface area of S^{d-1} , appears since μ_{\circ} is normalized so that $\mu_{\circ}(S^{d-1}) = 1$. From this, it follows that the density function q can be written as

$$q(v) = dv_d \int_{r=0}^{\infty} p(rv) r^{d-1} dr, \quad (2)$$

since integrating this function over any set A gives the probability that V will land in A . From our assumptions on p , we know that

$$p(rv) \geq \sum_{i=1}^L \mathbb{I}\{rv \in K_i\} c_{\text{lb}}.$$

Moreover, we can rewrite the indicator as $\mathbb{I}\{rv \in K_i\} = \mathbb{I}\{\frac{b_i}{w_i^\top v} < r \leq 1\}$. Substituting this into (2) gives

$$\begin{aligned}
q(v) &\geq \sum_{i=1}^L c_{\text{lb}} dv_d \int_{r=0}^{\infty} \mathbb{I}\left\{\left[\frac{b_i}{w_i^\top v} < r \leq 1\right] r^{d-1} dr\right. \\
&= \sum_{i=1}^L \mathbb{I}\{v \in K_i\} c_{\text{lb}} dv_d \int_{r=b_i/(w_i^\top v)}^1 r^{d-1} dr \\
&= \sum_{i=1}^L \mathbb{I}\{v \in K_i\} c_{\text{lb}} dv_d (1 - b_i/(w_i^\top v)^d) \\
&= q_{\text{lb}}(v)
\end{aligned}$$

Note that the indicator $\mathbb{I}\{v \in K_i\}$ appears in line 2 because the integral is only non-zero when $b_i/(w_i^\top v) < 1$, which is exactly the condition that $v \in K_i$. The upper bound on q follows by an identical argument using the upper bound on $p(rv)$. \square

Both bounds are defined piece-wise with one piece for each class. Restricted to class i , both the $q_{\text{lb}}(v)$ and $q_{\text{ub}}(v)$ are decreasing functions of $\theta(w_i, v)$, which implies that their λ -level sets are spherical caps. Therefore, each class contributes one large connected component to the level set of q that is roughly a spherical cap centered at the point w_i and the density of q goes to zero at the boundary of each class. Our main result is as follows:

Theorem 4. *Suppose Assumptions 2 and 3 hold and that f^* is consistent. There exists an r_c satisfying $r_c = \Omega(\epsilon c_{\text{lb}} / (c_{\text{ub}}^2 b_{\text{min}}))$ such that with probability at least $1 - \delta$, running Algorithm 3 with parameter r_c on an unlabeled sample of size $n = \tilde{O}((c_{\text{ub}}^4 d / (\epsilon^2 c_{\text{lb}}^2 b_{\text{min}}^2))^d)$ will query at most L labels and output a classifier with error at most ϵ .*

Note that if the scale parameter b_{min} is unknown, the conclusion of Theorem 4 continues to hold if the connection radius r_c and unlabeled sample complexity n are set using a conservatively small estimate $\widehat{b_{\text{min}}}$ satisfying $\widehat{b_{\text{min}}} \leq b_{\text{min}}$. This comes at the cost of an increased unlabeled sample complexity.

Before proving Theorem 4, we develop some general results for the robust linkage clustering algorithm. More generally, Algorithm 3 can be applied in any metric space (\mathcal{X}, d) by replacing θ with the distance metric d and suitable settings for the internal parameters r_a and τ . For the robust linkage approach to have low error, each class should have one large connected component in the graph G constructed by the algorithm so that: (1) with high probability a new point in class i will be nearest to that largest component, and (2) the large components of different classes are separated. Intuitively, G will have these properties if each positive region K_i has a connected high-density inner region A_i covering most of its probability mass and when it is rare to observe a point that is close to two or more classes. This notion is formalized below.

Let S be any set in \mathcal{X} . We say that a path $\pi : [0, 1] \rightarrow \mathcal{X}$ *crosses* S if the path starts and ends in different connected components of the complement of S in \mathcal{X} and we say that the *width* of S is the length of the shortest path that crosses S .

Definition 1. *The sets A_1, \dots, A_L are (r_c, r_a, τ, γ) -clusterable under probability P if there exists a separating set S of width at least r_c such that: (1) Each A_i is connected; (2) If $x \in \mathcal{X}$ satisfies $d(x, A_i) \leq r_c/3$ then $\Pr_{X \sim P}(X \in B(x, r_a)) > \tau + \gamma$; (3) If $x \in A_i$ then $\Pr_{X \sim P}(X \in B(x, r_c/3)) > \gamma$; (4) Every path from A_i to A_j crosses S ; and (5) If $x \in S$ then $\Pr_{X \sim P}(X \in B(x, r_a)) < \tau - \gamma$.*

Note that typically there must be a gap between the set A_i and the set S in order to satisfy the probability requirements (i.e., the set S will be smaller than $\mathcal{X} - \bigcup_{i=1}^L A_i$). The first three properties ensure that each set A_i will have one large connected component and the remaining two properties ensure that these connected components will be disconnected. Following an analysis similar to that of the cluster tree algorithm of Chaudhuri and Dasgupta [2010] gives the following result.

Lemma 3. *Suppose that the sets A_1, \dots, A_N are (r_c, r_a, τ, γ) -clusterable with respect to distribution P . For any failure probability $\delta > 0$, let G be the graph constructed by Algorithm 3 run on a sample S of size $O(\frac{1}{\gamma^2}(D + \ln \frac{1}{\delta}))$, where D is the VC-dimension of balls in (\mathcal{X}, d) , with parameters r_c, r_a , and τ . Define $K_i = \{x \in S : d(x, A_i) \leq r_c/3\}$ for each $i \in [N]$. With probability at least $1 - \delta$, the graph G has the following properties:*

1. **Complete:** For each i , all samples in K_i are active and included in the graph G .
2. **Separated:** For any $i \neq j$, there is no path in G from K_i to K_j .
3. **Connected:** For every i , the set K_i is connected in G .
4. **Extendible:** For any point $x \in A_i$, the nearest connected component of G to x contains K_i .

Proof. The proof technique used here follows a similar argument as Chaudhuri and Dasgupta [2010].

We use a standard VC bound [Vapnik and Chervonenkis, 1971] to relate the probability constraints in the clusterability definition to the empirical measure \hat{P} . For our value of n we have

$$\Pr(\sup_{x,r} |\hat{P}(B(x,r)) - P(B(x,r))| > \gamma) < \delta.$$

This implies that with probability at least $1 - \delta$ for all points x we have: (1) if $d(x, A_i) \leq \frac{r_c}{3}$ for any i then $\hat{P}(B(x, r_a)) > \tau$; (2) if $x \in S$ then $\hat{P}(B(x, r_a)) < \tau$; and (3) if $x \in A_i$ for any i then $\hat{P}(B(x, \frac{r_c}{3})) > 0$. We now use these facts to prove that the graph G has the completeness, separation, and connectedness properties.

Completeness follows from the fact that every sample $x \in \hat{K}_i$ is within distance $r_c/3$ of A_i and therefore $\hat{P}(B(x, r_a)) > \tau$.

To show separation, first observe that every sample $z \in S$ that belongs to S will be marked as inactive, since $\hat{P}(B(z, r_a)) < \tau$. Now let $x \in \hat{K}_i$ and $x' \in \hat{K}_j$ for $i \neq j$. Since the graph G does not contain any samples in the set S , any path in G from x to x' must have one edge that crosses S . Since the width of S is at least r_c , this edge would not be included in the graph G , and therefore G does not include a path from x to x' .

To show connectedness, let x and x' be any pair of samples in \hat{K}_i and let v and v' be their nearest points in A_i , respectively. By definition of \hat{K}_i , we know that $d(x, v) < r_c/3$ and $d(x', v') < r_c/3$. Since A_i is a connected set, there is a path $\pi : [0, 1] \rightarrow A_i$ in A_i starting at v and ending at v' . Cover the path π with a sequence of points z_1, \dots, z_k such that $d(z_j, z_{j+1}) < r_c/3$ for all j and the path π is covered by the balls $B(z_j, r_c/3)$. Further, choose $z_1 = v$ and $z_k = v'$. Since each point z_j belongs to A_i , the empirical probability mass of the ball $B(z_j, r_c/3)$ is non-zero, which implies that it must contain at least one sample point, say $y_j \in S$. We may take $y_1 = x$ and $y_k = x'$. Since every sample y_1, \dots, y_k is within distance $r_c/3$ of A_i , they are all active and included in the graph G . Moreover, since $d(y_j, y_{j+1}) < r_c$, we have that the path $x = y_1 \rightarrow \dots \rightarrow y_k = x'$ is a path connecting x and x' in G , as required.

Finally to show extensibility, let $x \in A_i$ be any point. By the uniform convergence for balls, $P(x, r_c/3)$ has non-zero empirical probability mass and therefore contains at least one active sample, say z . Since z is within distance $r_c/3$ of A_i , it belongs to the set K_i . Now let z^* be the closest active sample to x . We must have $d(x, z^*) \leq d(x, z) \leq r_c/3$ and it follows that $d(z, z^*) \leq d(z, x) + d(x, z^*) \leq 2r_c/3 < r_c$. Therefore, z^* also belongs to K_i , as required. \square

We now prove Theorem 4 by combining Lemmas 2 and 3:

Proof of Theorem 4. For each class $i \in [L]$, define $A_i = \{q_{\text{ub}}^{(i)} \geq \epsilon\}$. We will show that the sets A_1, \dots, A_L are (r_c, r_a, γ, τ) -clusterable for appropriate choices of the parameters. Then Lemma 3 will guarantee that with high probability, the clustering produced by Algorithm 3 will approximate the connected components of the ϵ -level of $\{q_{\text{ub}} \geq \epsilon\}$.

Recall that for each class $i \in [L]$, the sets $\{q_{\text{ub}}^{(i)} \geq \epsilon\}$ and $\{q_{\text{lb}}^{(i)} \geq \epsilon\}$ are spherical caps. To simplify notation, let $C(u, r) = \{v \in \mathcal{S}^{d-1} : \theta(v, u) \leq r\}$ denote the spherical cap of angular radius r centered at u . Let $\rho_{\text{ub}}^{(i)}(\lambda) = \arccos(b_i(1 - \lambda/(c_{\text{ub}}dv_d)))^{-1/d}$ denote the angular radius of $\{q_{\text{ub}}^{(i)} \geq \epsilon\}$, so that $\{q_{\text{ub}}^{(i)} \geq \epsilon\} = C(w_i, \rho_{\text{ub}}^{(i)}(\epsilon))$, and $\rho_{\text{lb}}^{(i)}(\lambda)$, defined similarly, be the angular radius of $\{q_{\text{lb}}^{(i)} \geq \epsilon\}$. Define $\tilde{\epsilon} = \frac{c_{\text{lb}}}{c_{\text{ub}}}\epsilon$ and suppose for the moment that we can find an activation radius $r_a > 0$ small enough so that the following inequalities hold for all classes $i = 1, \dots, L$:

$$\frac{5}{3}r_a \leq \rho_{\text{lb}}^{(i)}\left(\frac{3\tilde{\epsilon}}{4}\right) - \rho_{\text{ub}}^{(i)}(\epsilon) \quad \text{and} \quad 2r_a \leq \rho_{\text{ub}}^{(i)}(0) - \rho_{\text{ub}}^{(i)}\left(\frac{\tilde{\epsilon}}{4}\right).$$

Given such an activation radius, we will show that the sets A_1, \dots, A_L are (r_c, r_a, τ, γ) -clusterable with $r_c = 2r_a$, $\tau = \frac{\tilde{\epsilon}V^d(r_a)}{2}$, and $\gamma = \frac{\tilde{\epsilon}V^d(r_c/3)}{4}$ and the separating set is $S = \{v \in \mathcal{S}^{d-1} : \theta(v, w_i) \geq \rho_{\text{ub}}^{(i)}(0) - r_a \text{ for all } i\}$:

1. *Connection:* Each A_i set is a spherical cap and therefore connected.
2. *High-density near A_i :* Let $v \in \mathcal{S}^{d-1}$ be such that $\theta(v, A_i) < r_c/3$ and let $u \in C(v, r_a)$ be any point in the spherical cap of angular radius r_a about v . By the triangle inequality, we know that $\theta(w_i, u) \leq \theta(w_i, v) + \theta(v, u) \leq \rho_{\text{ub}}^{(i)}(\epsilon) + \frac{5}{3}r_a \leq \rho_{\text{lb}}^{(i)}(\frac{3\tilde{\epsilon}}{4})$. This implies that $q(u) \geq \frac{3\tilde{\epsilon}}{4}$ for all points in $C(v, r_a)$ and therefore $\Pr_{V \sim q}(V \in C(v, r_a)) \geq \frac{4\tilde{\epsilon}}{3}V^d(r_a) \geq \tau + \gamma$.
3. *High-density inside A_i :* Now let $v \in A_i$. Since $r_c/3 < r_a$, the above arguments show that $q(u) \geq \frac{4\tilde{\epsilon}}{3}$ for all points $u \in C(v, r_c/3)$ and therefore $\Pr_{V \sim q}(V \in C(v, r_c/3)) \geq \frac{4\tilde{\epsilon}}{3}V^d(r_c/3) \geq \gamma$.
4. *Separation by the set S :* For each class i , the set S contains the annulus $\{v \in \mathcal{S}^{d-1} : \rho_{\text{ub}}^{(i)}(0) - r_a \leq \theta(w_i, v) \leq \rho_{\text{ub}}^{(i)}(0)\}$ which has width r_a . Any path from one A_i to another A_j must cross two such annuli, each of width r_a , so the length of the path crossing S is at least $2r_a = r_c$.
5. *Low density inside S :* Finally, let v be any point in the set S and let $u \in C(v, r_a)$. For any class i , the reverse triangle inequality gives that $\theta(w_i, v) \geq \theta(v, w_i) - \theta(u, w_i) \geq \rho_{\text{ub}}^{(i)}(0) - 2r_a \geq \rho_{\text{ub}}^{(i)}(\frac{\tilde{\epsilon}}{4})$. Since this is true for all classes i , we have $q(v) \leq \frac{\tilde{\epsilon}}{4}$ and therefore $\Pr_{V \sim q}(V \in C(v, r_a)) \leq \frac{\tilde{\epsilon}}{4}V^d(r_a) \leq \tau - \gamma$.

It follows that the sets A_1, \dots, A_L are (r_c, r_a, τ, γ) -clusterable and it only remains to find an activation radius r_a that satisfies the above inequalities. Since the robust linkage algorithm needs to estimate the probability mass of balls to within error $\gamma = \frac{\tilde{\epsilon}V^d(2r_a/3)}{4}$, we want this activation radius to be not too small.

Taking the first order Taylor expansion of the $\rho_{\text{lb}}^{(i)}$ and $\rho_{\text{ub}}^{(i)}$ functions, we have:

$$\begin{aligned}\rho_{\text{lb}}^{(i)}(\lambda) &= \arccos(b_i) - \frac{b_i}{\sqrt{1-b_i^2}} \frac{1}{c_{\text{lb}} dv_d} \lambda + O(\lambda^2) \\ \rho_{\text{ub}}^{(i)}(\lambda) &= \arccos(b_i) - \frac{b_i}{\sqrt{1-b_i^2}} \frac{1}{c_{\text{ub}} dv_d} \lambda + O(\lambda^2),\end{aligned}$$

as $\lambda \rightarrow 0$. Therefore, we have that

$$\rho_{\text{lb}}^{(i)}(3\tilde{\epsilon}/4) - \rho_{\text{ub}}^{(i)}(\epsilon) = \frac{1}{4dv_d c_{\text{ub}}} \cdot \frac{b_i}{\sqrt{1-b_i^2}} \epsilon + O(\epsilon^2)$$

and

$$\rho_{\text{ub}}^{(i)}(0) - \rho_{\text{ub}}^{(i)}(\tilde{\epsilon}/4) = \frac{c_{\text{lb}}}{4dv_d c_{\text{ub}}^2} \cdot \frac{b_i}{\sqrt{1-b_i^2}} \epsilon + O(\epsilon^2),$$

which shows that it is sufficient to set $r_a = \frac{3c_{\text{lb}}}{20dv_d c_{\text{ub}}^2} \cdot \frac{b_i}{\sqrt{1-b_i^2}} \epsilon + O(\epsilon^2) = \Omega(\frac{c_{\text{lb}}}{c_{\text{ub}}^2} b_{\min} \epsilon)$ as $\epsilon \rightarrow 0$ and it follows that $n = O(\frac{1}{\gamma^2}(d + \ln \frac{1}{\delta})) = \tilde{O}((c_{\text{ub}}^4 d / (\epsilon^2 c_{\text{lb}}^2 b_{\min}^2))^d)$.

Finally, we show that the algorithm correctly recovers the labels of the large clusters. For $n = \tilde{O}(L/\epsilon^2)$, we have that with probability at least $1 - \delta$ the following holds simultaneously for all 2^L subsets $I \subset [L]$: $|\hat{P}(A_I) - P(A_I)| \leq \epsilon/4$, where $A_I = \bigcup_{i \in I} A_i$. Since all samples in A_I are marked as active (by Lemma 3), this implies that all but at most $\frac{\epsilon}{4}n$ of the active points will belong to the A_i sets. It follows that if the algorithm queries the labels of the L largest clusters, they will also contain all but $\frac{\epsilon}{4}n$ active samples.

On the other hand, whenever we query the label of one of the A_i sets, we know that we will correctly classify every test point belonging to A_i , so the error of the resulting classifier is at most the probability mass of $\{q_{\text{ub}} \leq \epsilon\}$ together with the probability mass of the A_i sets for which we did not query the label. Since the unqueried A_i sets have empirical probability mass at most $\epsilon/4$ and we have uniform convergence for all unions of A_i sets to within error $\epsilon/4$, it follows that the total probability mass of the unlabeled A_i sets is at most $\epsilon/2$ and it follows that the error of the resulting classifier is at most ϵ . \square

There are two main differences between the sample complexity of Theorem 4 and the results from Section 4. First, the unlabeled sample complexity now has an ϵ^{-2d} dependence, rather than only ϵ^{-2} . This is because the distance between the connected components of $\{p \geq \epsilon\}$ goes to zero (in the worst case) as $\epsilon \rightarrow 0$, so our algorithm must be able to detect low-density regions of small width. In contrast, Lemma 1 allowed us to establish a non-diminishing gap $g > 0$ between the classes when the codewords were well separated. On the other hand, the label complexity in this setting is better, scaling with L instead of N , since we are able to establish that each class will have one very large cluster containing nearly all of its data.

Theorem 7 in the appendix gives an analysis of Algorithm 3 in the agnostic setting of Section 7.

6 The Boundary Features Condition

Finally, in this section we introduce a novel condition on the code matrix called the boundary features condition that captures the intuition that every binary classification task should be significant. Assumption 4 formalizes this intuition.

Assumption 4. *There exists a code matrix $C \in \{\pm 1\}^{L \times m}$, linear functions h_1, \dots, h_m , and a scale parameter $R > 0$ so that: (1) for any point x in class y , we have $h(x) = C_y$; (2) for each h_j , there exists a class i such that negating the j^{th} entry of C_i produces a codeword C'_i not in C and there exists a point x on the hyperplane $h_j = 0$ such that every point in $B(x, R)$ has either code word C_i or C'_i ; and (3) any pair of points $x, x' \in \mathcal{X}$ such that $h(x)$ and $h(x')$ are not codewords in C and $h(x) \neq h(x')$ must have $\|x - x'\| \geq R$.*

Part (1) of this assumption requires that the output code classifier is consistent, part (2) is a condition that guarantees every linear separator h_j separates at least one class i from a region of space that does not belong to any class, and part (3) requires that points with codewords not in the code matrix must either have the same codeword or be separated by distance R . Part (3) allows us to simplify our algorithm and analysis and is trivially satisfied in cases where all points in \mathcal{X} that do not belong to any class have the same codeword, as is the case for one-vs-all classification and the problem in Figure 2.

Problems in this setting are more challenging than those of the previous sections because they may not be amenable to clustering-based learning strategies. Whenever the Hamming distance between a pair of codewords is only 1, this implies that one of the linear separators h_j forms a shared boundary between the classes, and therefore these classes may be connected by a large and high-density region. Instead, Assumption 4 guarantees that for every linear separator h_j , there is some ball $B(x, R)$ centered on h_j that is half-contained in the set of points belonging to some class i and the other half belongs to the set of points that do not belong to any class. Therefore, by looking for hyperplanes that locally separate sample data from empty regions of space, we can recover the linear separator h_j from the local *absence* of data. Define $K_i = \{x \in \mathcal{X} : h(x) = C_i\}$ to be the set of points that belong to class i and $K = \bigcup_{i=1}^L K_i$. Under the condition that the density p is supported on K and is upper and lower bounded, we exploit this structure in an algorithm that directly learns the linear separators h_1, \dots, h_m .

Our hyperplane detection algorithm works by searching for balls of radius r whose centers are sample points such that one half of the ball contains very few samples. If a half-ball contains very few sample points then it must be mostly disjoint from the set K . But since its center point belongs to the set K , this means that the hyperplane defining the half-ball is a good approximation to at least one of the true hyperplanes. See Figure 3 for examples of half-balls that would pass and fail this test. The collection H of hyperplanes produced in this way partition the space into cells. Our algorithm queries the labels of the cells containing the most sample points and classifies test points





Figure 2: An example of the boundary features problem. The arrows indicate the positive side of the linear functions.

Input: Sample $S = \{x_1, \dots, x_n\}$, $r > 0$, $\tau > 0$.

1. Initialize set of candidate hyperplanes $H = \emptyset$.
2. For all samples $\hat{x} \in S$ with $B(\hat{x}, r) \subset \mathcal{X}$:
 - (a) Let $\hat{w} = \operatorname{argmin}_{w \in \mathcal{S}^{d-1}} |B^{1/2}(\hat{x}, r, w) \cap S|$.
 - (b) If $|B^{1/2}(\hat{x}, r, \hat{w}) \cap S|/n < \tau$, add (\hat{x}, \hat{w}) to H .
3. Let $\{\hat{C}_i\}_{i=1}^N$ be the partitioning of \mathcal{X} induced by H .
4. Query the label of the L cells with the most samples.
5. Output $\hat{f}(x) = \text{label of } C_i \text{ containing } x$.



Algorithm 4: Plane-detection algorithm.

Figure 3: Examples of half-balls that would be included (green) or excluded (red) by the plane detection algorithm.

based on the label of their cell in the partition (and if the label is unknown, we output a random label). Pseudocode is given in Algorithm 4 using the following notation: for any center $x \in \mathcal{X}$, radius $r \geq 0$, and direction $w \in \mathcal{S}^{d-1}$, let $B^{1/2}(x, r, w) = \{y \in B(x, r) : w^\top(y - x) > 0\}$ and define $p^{1/2}(r) = \frac{1}{2}c_{1b}r^d v_d$.

Each candidate hyperplane produced by Algorithm 4 is associated with a half-ball that caused it to be included in H . In fact, we can think of the pairs (\hat{x}, \hat{w}) in H as either encoding the linear function $\hat{h}(x) = w^\top(x - \hat{x})$ or the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$, where r is the scale parameter of the algorithm. Most of our arguments will deal with the half-balls directly, so we adopt the second interpretation. The analysis of Algorithm 4 has two main steps. First, we show that the face of every half-ball in the set H is a good approximation to at least one of the true hyperplanes, and that every true hyperplane is well approximated by the face of at least one half-ball in H . Second, using the fact that the half-balls in H are good approximations to the true hyperplanes, we argue that the output classifier will only be inconsistent with the true classification rule in a small margin around each of the true linear separators. Then the error of the classification rule is easily bounded by bounding the probability mass of these margins.

To measure the approximation quality, we say that the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ is an α -approximation to the linear function h if $\Pr_{X \sim B^{1/2}}(\operatorname{sign}(h(X)) \neq \operatorname{sign}(h(\hat{x}))) \leq \alpha$, where $\Pr_{X \sim B^{1/2}}$ denotes the probability when X is sampled uniformly from the half-ball $B^{1/2}$. The motivation for this definition is as follows: given any point $\hat{x} \in \mathcal{X}$, the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ will be an α -approximation to h_i only if \hat{x} is on one side of the decision surface of h_i and all but an α -fraction of the half-ball's volume is on the other side. Intuitively, this means that the face of the half-ball must approximate the decision surface of the function h_i .

The following Lemma shows that when Algorithm 4 is run with appropriate parameters and on a large enough sample drawn from the data distribution, then with high probability the algorithm will include at least one half-ball in H α -approximating each true hyperplane h_i and every half-ball in H will be an α -approximation to at least one true hyperplane. Recall that $p^{1/2}(r) = \frac{1}{2}c_{1b}r^d v_d$ is a lower bound on the probability mass of a half-ball of radius r contained in the set K .

Lemma 4. *Fix any $\alpha > 0$ and confidence parameter $\delta > 0$. Let H be the set of half-balls produced by Algorithm 4 when run with parameters $r = R/2$ and $\tau = \frac{1}{2}\alpha p^{1/2}(r)$ on a sample of size $n = O(\frac{1}{\gamma^2}(\ln^2 \frac{d}{\gamma} + \ln \frac{1}{\delta}))$ where $\gamma = \frac{2}{5}\tau = \frac{1}{5}\alpha p^{1/2}(r)$. Then with probability at least $1 - \delta$, every half-ball in H will be an α -approximation to at least one true hyperplane h_i , and every true hyperplane h_i will be α -approximated by at least one half-ball in H .*

Proof. Since the VC-dimension of both balls and half-spaces in \mathbb{R}^d is $d+1$, the VC-dimension of the set of intersections of balls and up to two half-spaces is $O(d \ln d)$. Therefore, by a standard VC-bound [Vapnik and Chervonenkis, 1971], if we see an iid sample S of size $n = O(\frac{1}{\gamma^2}(\ln^2 \frac{d}{\gamma} + \ln \frac{1}{\delta}))$, then with probability at least $1 - \delta$ the empirical measure of any ball intersected with up to two half-spaces will be within γ of its true probability mass. In other words, the fraction of the sample set S that lands in any ball intersected with up to two half-spaces will be within γ of the probability that a sample X drawn from P will land in the same set. For the remainder of the proof, assume that this high-probability event holds.

First, we show that every half-ball in the set H is an α -approximation to at least one true hyperplane. Suppose

otherwise, then there is a half-ball $B^{1/2} = B^{1/2}(\hat{x}, r, \hat{w})$ with $(\hat{x}, \hat{w}) \in H$ that is not an α approximation to any true hyperplane h_i . The center \hat{x} of the half-ball must belong to the positive region K , since it is one of the sample points. If the half-ball $B^{1/2}$ is contained entirely in the set K , then the probability that a new sample X drawn from P will land in the half-ball $B^{1/2}$ is $p^{1/2}(r)$ and therefore the fraction of samples that landed in the half-ball is at least $p^{1/2}(R/2) - \gamma$. But since $p^{1/2}(r) - \gamma \geq \frac{4}{5}\alpha p^{1/2}(r) > \tau$, this contradicts the half-ball being included in the set H . Otherwise, the half-ball contains at least one point y that does not belong to the set K (i.e., it does not belong to any class). Since \hat{x} is in the set K , there is at least one true hyperplane h_i that separates \hat{x} from y . Since $r = R/2 < R$, every other point y' in the half-ball that does not belong to any class must have the same code word as y (since, by assumption, points outside of K that do not belong to any class must have the same code word when they are closer than R), and therefore must be on the same side of h_i as y . It follows that all points in the half-ball on the same side of h_i as \hat{x} (i.e., those points for which the sign of h_i matches the sign of $h_i(\hat{x})$) belong to the set K . But, since the half-ball is not an α -approximation to h_i , this implies that at least an α fraction of the half-ball's volume must belong to the set K . Therefore, the probability that a new sample x drawn from the data distribution p belongs to the half-ball can be lower bounded as follows:

$$\Pr_{x \sim p}(x \in B^{1/2}) \geq c_{\text{lb}} \text{Vol}(B^{1/2} \cap K) = c_{\text{lb}} \text{Vol}(B^{1/2}) \frac{\text{Vol}(B^{1/2} \cap K)}{\text{Vol}(B^{1/2})} \geq \alpha p^{1/2}(r).$$

By the uniform convergence argument, the fraction of the samples in S contained in the half-ball $B^{1/2}$ is at least $\alpha p^{1/2}(r) - \gamma > \tau$, which contradicts the half-ball being in H . In either case we arrived at a contradiction and it follows that every half-ball in H is an α -approximation to at least one true hyperplane h_i .

Finally, we show that the set H will contain at least one half-ball that is an α -approximation to each true hyperplane h_i . Fix any true hyperplane h_i . By assumption, there is a class ℓ and a point x_0 on the decision surface of h_i so that one half-ball of $B(x_0, R)$ with face h_i is contained in $K_\ell \subset K$ and the other half-ball is disjoint from K . Suppose WLOG that the half-ball on the negative side of h_i is contained in K (the case when the half-ball on the positive side is contained in K is identical). Define $\rho > 0$ to be the width such that the probability that a new sample X from P lands in the slice of the ball $S = \{x \in B(\hat{x}_0, r) : h_i(x) \in [-\rho, 0]\}$ is equal to $\tau - \gamma$. Note that, since the half-ball on the negative side of h_i is a subset of K and $\tau - \gamma = \frac{3}{10}\alpha p^{1/2}(r) < p^{1/2}(r)$, such a value of ρ always exists. Since $\tau - \gamma > \gamma$, the uniform convergence argument guarantees that there will be at least one sample point in the slice, say $\hat{x} \in S$. Since \hat{x} is within distance $r = R/2$ of the point x_0 , the ball $B(\hat{x}, r)$ is contained in $B(x_0, R)$. Therefore, the ball of radius r centered at \hat{x} only contains points that either belong to class ℓ or no class, since only the linear separator h_i passes through this ball. By construction, the half-ball $B^{1/2}(\hat{x}, r, w_i)$ (where w_i is the coefficient vector defining $h_i(x) = w_i^\top x - b_i$) with face parallel to h_i intersects the set K in a slice of width at most ρ and therefore has probability mass at at most $\tau - \gamma$. It follows that the direction \hat{w} that minimizes the number of samples in the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ will result in the half-ball containing at most a τ fraction of the sample set, and therefore the pair (\hat{x}, \hat{w}) will be included in H , and this will be an α -approximation to h_i . \square

Naturally, if a half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ is an α -approximation to the linear function h , we expect that the decision surface of $\hat{h}(x) = \hat{w}^\top(x - \hat{x})$ is similar to the decision surface of h . In turn, this suggests that either $\hat{h}(x)$ or $-\hat{h}(x)$ should take similar function values to $h(x)$ (since the coefficient vectors are normalized). We first give a simple probability lemma that bounds the fraction of a ball contained between two parallel hyperplanes, one passing through the ball's center. The proof of Lemma 5 is in Section 11 of the appendix.

Lemma 5. *Let $r > 0$ be any radius and X be a random sample drawn uniformly from the ball of radius r centered at the origin. For any width $0 \leq \rho \leq r/\sqrt{2}$, the probability that the first coordinate of X lands in $[0, \rho]$ can be bounded as follows:*

$$\sqrt{\frac{d}{2^d \pi}} \frac{\rho}{r} \leq \Pr_{X \sim B(r, 0)}(X_1 \in [0, \rho]) \leq \sqrt{\frac{d+1}{2\pi}} \frac{\rho}{r}.$$

Using Lemma 5, we show the following:

Lemma 6. Let the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ be an α -approximation to the linear function $h(x) = w^\top x - b$ with $\|w\| = 1$, $\hat{x} \in \mathcal{X}$, and $\alpha < \frac{1}{2}$. Let D be the diameter of \mathcal{X} . If $h(\hat{x}) < 0$ then for all $x \in \mathcal{X}$ we have

$$|h(x) - \hat{h}(x)| \leq \left(2D + \sqrt{\frac{2^d \pi}{d} \frac{r}{2}}\right) \sqrt{\alpha},$$

where $\hat{h}(x) = \hat{w}^\top (x - \hat{x})$. Otherwise, if $h(\hat{x}) > 0$ then the same upper bound holds for $|h(x) + \hat{h}(x)|$.

Proof. Suppose that $h(\hat{x}) < 0$ and let X be a uniformly random sample from the half-ball $B^{1/2} = B^{1/2}(\hat{x}, r, \hat{w})$. By assumption, we know that $\Pr(h(X) < 0) \leq \alpha$.

First we show that $\|w - \hat{w}\|$ is small. Since $\alpha < 1/2$ we have that $w^\top \hat{w} > 0$. To see this, notice that we must have $h(\hat{x} + r\hat{w}) \geq 0$, since otherwise at least half of the half-ball would be on the negative side of h . Define $g(x) = w^\top (x - \hat{x})$ to be the linear function whose decision surface runs parallel to that of h but passes through the point \hat{x} . Since $h(x) = g(x) + h(\hat{x}) \leq g(x)$, we have that $\alpha > \Pr(h(X) < 0) \geq \Pr(g(X) < 0)$. Moreover, since the decision surface of g passes through the center of the half-ball $B^{1/2}$ and the uniform distribution on the half-ball is radially symmetric about the point \hat{x} , we have that $\Pr(g(X) < 0) = \frac{\theta(w, \hat{w})}{\pi}$. It follows that $\theta(w, \hat{w}) \leq \pi\alpha$. Using this fact, we can bound $\|w - \hat{w}\|$ as follows:

$$\|w - \hat{w}\|^2 = \|w\|^2 + \|\hat{w}\|^2 - 2w^\top \hat{w} = 2(1 - w^\top \hat{w}).$$

Since $w^\top \hat{w} = \cos(\theta(w, \hat{w}))$ and on the interval $[0, \pi/2]$, the $\cos(\theta)$ function is decreasing and lower bounded by $1 - \frac{2}{\pi}\theta$, we have that $2(1 - w^\top \hat{w}) \leq 4\alpha$. Taking the square root gives that $\|w - \hat{w}\| \leq 2\sqrt{\alpha}$.

Next we show that $|h(\hat{x})|$ (the distance from \hat{x} to the decision surface of h) is not too large. The half-ball $B^{1/2}(\hat{x}, r, w)$, whose direction w matches the coefficient vector of h is one half-ball centered at \hat{x} of radius r minimizing the fraction of its volume contained on the same side of h as \hat{x} . This is because every point in the ball $B(\hat{x}, r)$ not on the same side as \hat{x} is contained in $B^{1/2}(\hat{x}, r, w)$. Let Y be uniformly sampled from $B^{1/2}(\hat{x}, r, w)$. By construction of the half-ball Y is sampled from, we have that $\Pr(h(X) < 0) \geq \Pr(h(Y) < 0)$, which gives

$$\alpha \geq \Pr_{X \sim B^{1/2}(\hat{x}, r, \hat{w})}(h(X) < 0) \geq \Pr_{Y \sim B^{1/2}(\hat{x}, r, w)}(h(Y) < 0) \geq \sqrt{\frac{d}{2^d \pi}} \frac{2|h(\hat{x})|}{r},$$

which implies that

$$|h(\hat{x})| \leq \sqrt{\frac{2^d \pi}{d}} \frac{r\alpha}{2}.$$

Finally, let x' be any point on the decision surface of h , so that $h(x) = w^\top (x - x')$. Combining the above calculations we have

$$\begin{aligned} |h(x) - \hat{h}(x)| &= |w^\top (x - x') - \hat{w}^\top (x - \hat{x})| \\ &= |w^\top (x - \hat{x}) + w^\top (\hat{x} - x') - \hat{w}^\top (x - \hat{x})| \\ &= |(w - \hat{w})(x - \hat{x}) + w^\top (\hat{x} - x')| \\ &\leq \|w - \hat{w}\| \|x - \hat{x}\| + |h(\hat{x})| \\ &\leq 2\sqrt{\alpha}D + \sqrt{\frac{2^d \pi}{d}} \frac{r\alpha}{2} \\ &\leq (2D + \sqrt{\frac{2^d \pi}{d}} \frac{r}{2}) \sqrt{\alpha}, \end{aligned}$$

as required. The proof of the case when $h(x) > 0$ follows by applying the above arguments to the function $-h$. \square

Recall that for any hyperplane $h(x) = w^\top x - b$ with $\|w\|_2 = 1$, the distance from point x to the decision surface of h is $|h(x)|$. The above lemma implies that if $B^{1/2}(\hat{x}, r, \hat{w})$ is an α -approximation to h , then either \hat{h} or $-\hat{h}$ will have the same sign as h for all points in \mathcal{X} except those in a margin of width $O(\sqrt{\alpha})$ around h . Under the uniform distribution on K , the probability mass of the margins surrounding the true hyperplanes isn't large, which results in low error for the classification rule.

Theorem 5. *Suppose Assumptions 3 and 4 hold. For any desired error $\epsilon > 0$, with probability at least $1 - \delta$, running Algorithm 4 with parameters $r \leq R/2$ and $\tau = \alpha p^{1/2}(r)/2$ for a known constant α on a sample of size $n = \tilde{O}(dm^2 c_{\text{ub}}^2 R^d / (c_{\text{lb}}^2 \epsilon^4))$ will have error at most ϵ .*

Proof. By Lemma 4, for the parameter settings τ and r and the given sample size, with probability at least $1 - \delta$ every half-ball included in the set H will be an α -approximation to some true hyperplane h_i , and every true hyperplane h_i is α -approximated by at least one half-ball in H . Assume that this high probability event occurs.

Let $H = \{(\hat{x}_1, \hat{w}_1), \dots, (\hat{x}_M, \hat{w}_M)\}$ be the set of half-balls produced by the algorithm and define $\hat{h}_i(x) = \hat{w}_i^\top (x - \hat{x}_i)$ for $i = 1, \dots, M$ to be the corresponding linear functions. Algorithm 4 uses these hyperplanes to partition the space \mathcal{X} into a collection of polygonal regions and assigns a unique class label to each cell in the partition. Notice that negating any of the \hat{h}_i functions does not change the partitioning of the space. Therefore, negating any subset of the \hat{h}_i will not change the permutation-invariant error of the resulting classifier.

Let I_1, \dots, I_m be a partition of the set of indices $\{1, \dots, M\}$ such that for all $j \in I_i$, we have that $B^{1/2}(\hat{x}_j, \hat{w}_j, r)$ is an α -approximation to h_i . By Lemma 6, we know that for at least one $g \in \{\hat{h}_j, -\hat{h}_j\}$, we have that

$$|h_i(x) - g(x)| \leq \left(2D + \sqrt{\frac{2^d \pi}{d} \frac{r}{2}}\right) \sqrt{\alpha}$$

Since negating any of the functions \hat{h}_j does not change the error of the resulting classifier, assume WLOG that the above holds for $g = \hat{h}_j$.

This implies that whenever $|h_i(x)| > c\sqrt{\alpha}$, where $c = 2D + \sqrt{\frac{2^d \pi}{d} \frac{R}{4}}$, then for every $j \in I_i$, the sign of $\hat{h}_j(x)$ is the same as the sign of $h_i(x)$. It follows that for points that are not within a margin of $c\sqrt{\alpha}$ of any of the true hyperplanes, every \hat{h}_j function with $j \in I_i$ will have the same sign as h_i for all $i = 1, \dots, m$. It follows that the classifier can only error on points that are within a $c\sqrt{\alpha}$ margin of one of the true hyperplanes.

Using Lemma 5 we can bound the probability that a sample X drawn uniformly from K lands in the $c\sqrt{\alpha}$ -margin of hyperplane h_i as follows:

$$\Pr(X \text{ in } c\sqrt{\alpha}\text{-margin of } h_i) \leq 2\sqrt{\frac{d+1}{2\pi}} \frac{c\sqrt{\alpha}}{D} D^d v_d c_{\text{ub}},$$

where D is the diameter of X . We can make this upper bound equal to ϵ/m by setting

$$\alpha = \frac{\pi}{2(d+1)} \left(\frac{\epsilon D}{m c D^d v_d c_{\text{lb}}}\right)^2 = \Omega\left(\frac{\epsilon^2}{m^2 2^d R^2 D^{2d} v_d^2 c_{\text{lb}}^2}\right)$$

Applying the union bound to the m hyperplanes h_1, \dots, h_m shows that the error of \hat{f} is at most ϵ . \square

Note that if the scale parameter R is unknown, the conclusions of Theorem 5 continue to hold when the parameter r and the unlabeled sample complexity n are set using a conservatively small estimate \hat{R} satisfying $\hat{R} \leq R$.

Theorem 8 in the appendix extends the above result to the agnostic setting considered in Section 7.

7 Extensions to the Agnostic Setting

The majority of our algorithms have two phases: first, we extract a partitioning of the unlabeled data into groups that are likely label-homogeneous, and second, we query the label of the largest groups. We can extend our results for these algorithms to the agnostic setting by querying multiple labels from each group and using the majority label.

Specifically, suppose that the data is generated according to a distribution P over $\mathcal{X} \times [L]$ and there exists a labeling function f^* such that $\Pr_{(x,y) \sim P}(f^*(x) \neq y) \leq \eta$ and our assumptions hold when the unlabeled data is drawn from the marginal $P_{\mathcal{X}}$ but the labels are assigned by f^* . That is, the true distribution over class labels disagrees with a function f^* satisfying our assumptions with probability at most η . In this setting, the first phase of our algorithms, which deals with only unlabeled data, behaves exactly as in the realizable setting. The only difference is that we will need to query

multiple labels from each group of data to ensure that the majority label is the label predicted by f^* . Suppose that the training data is $(x_1, y_1), \dots, (x_n, y_n)$ drawn from P (where the labels y_i are initially unobserved). For $n = \tilde{O}(1/\eta^2)$, we are guaranteed that on at most $2\eta n$ of the training points we have that $y_i \neq f^*(x_i)$. Moreover, if we only need to guess the label of large groups of samples, say those containing at least $8\eta n$ points, then we are guaranteed that within each group at least $1/4$ of the sample points will have labels that agree with f^* . Therefore, after querying $O(\log(1/\delta))$ labeled examples from each group, the majority label will agree with f^* . If we use these labels in the second phase of the algorithm, we would be guaranteed that the error of our algorithm would be at most ϵ had the labels been produced by f^* , and therefore the error under the distribution P is at most $\eta + \epsilon$. The appendix contains agnostic versions of Theorems 1, 4, and 5.

Similarly, modifying Algorithm 2 to require that the each cluster in the pruning have a majority label that accounts for at least $3/4$ of the cluster’s data can be used to extend the corresponding results to the agnostic setting.

8 Conclusion and Discussion

In this work we showed how to exploit the implicit geometric assumptions made by output code techniques under the well studied cases of one-vs-all and well separated codewords, and for a novel boundary features condition that captures the intuition that every binary learning task should be significant. We provide label-efficient learning algorithms for both the consistent and agnostic learning settings with guarantees when the data density has thick level sets or upper and lower bounds. In all cases, our algorithms show that the implicit assumptions of output code learning can be used to learn from very limited labeled data.

In this work we focused on linear output codes, which have been in several practical works. For example Palatucci et al. [2009] use linear output codes for neural decoding of thoughts from fMRI data, Berger [1999] used them successfully for text classification, and Crammer and Singer [2000] show that they perform well on MNIST and several UCI datasets. Many other works use non-linear output codes, and it is a very interesting research direction to extend our work to such cases.

The unlabeled sample complexity of our algorithms is exponential in the dimension because our algorithms require the samples to cover high-density regions. It is common for semi-supervised algorithms to require exponentially more unlabeled data than labeled, e.g. [Singh et al., 2008, Castelli and Cover, 1995]. Our results also show that the unlabeled sample complexity only scales exponentially with the intrinsic dimension, which may be significantly lower than the ambient dimension for real-world problems. An interesting direction for future work is to determine further conditions under which the unlabeled sample complexity can be drastically reduced.

Acknowledgments

This work was supported in part by NSF grants CCF-1422910, CCF-1535967, IIS-1618714, a Sloan Research Fellowship, a Microsoft Research Faculty Fellowship, and a Google Research Award.

References

- E. Allwein, R. Schapire, and Y. Singer. Reducing multiclass to binary: A unifying approach for margin classifiers. In *Journal of Machine Learning Research*, 2000.
- K. Balasubramanian, P. Donmez, and G. Lebanon. Unsupervised supervised learning ii: Margin-based classification without labels. In *AISTATS*, pages 137–145, 2011a.
- K. Balasubramanian, P. Donmez, and G. Lebanon. Unsupervised supervised learning ii: Margin-based classification without labels. In *Journal of Machine Learning Research*, volume 12, pages 3119–3145, 2011b.
- M-F. Balcan and A. Blum. A discriminative model for semi-supervised learning. In *Journal of the ACM*, 2010.
- M-F. Balcan and R. Urner. Active learning. In *Survey in the Encyclopedia of Algorithms*, 2015.
- M-F. Balcan, A. Blum, and K. Yang. Co-training and expansion: Towards bridging theory and practice. In *NIPS*, 2004.
- M-F. Balcan, A. Beygelzimer, and J. Lanford. Agnostic active learning. In *ICML*, 2006.
- M.-F. Balcan, A. Blum, and Y. Mansour. Exploiting ontology structures and unlabeled data for learning. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 1112–1120, 2013.
- A. Berger. Error-correcting output coding for text classification. In *IJCAI Workshop on machine learning for information filtering*, 1999.
- A. Beygelzimer, J. Langford, and P. Ravikumar. Solving multiclass learning problems via error-correcting output codes. *ALT*, 2009.
- A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *COLT*, 1998.
- V. Castelli and T. Cover. On the exponential value of labeled samples. In *Pattern Recognition Letters*, 1995.
- O. Chapelle, B. Schölkopf, and A. Zien. *Semi-Supervised Learning*. The MIT Press, 1st edition, 2010. ISBN 0262514125, 9780262514125.
- K. Chaudhuri and S. Dasgupta. Rates of convergence for the cluster tree. In *Advances in Neural Information Processing 23 (NIPS)*, pages 343–351, 2010.
- K. Crammer and Y. Singer. Improved output coding for classification using continuous relaxation. In *NIPS*, 2000.
- A. Daniely, M. Schapira, and G. Shahaf. Multiclass learning approaches: A theoretical comparison with implications. In *NIPS*, 2012.
- S. Dasgupta. Two faces of active learning. In *Theoretical Computer Science*, 2011.
- S. Dasgupta and K. Sinha. Randomized partition trees for exact nearest neighbor search. In *COLT*, 2013.
- T. G. Dietterich and G. Bakiri. Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research*, pages 263–286, 1995.
- P. Donmez, G. Lebanon, and K. Balasubramanian. Unsupervised supervised learning i: Estimating classification and regression errors without labels. In *Journal of Machine Learning Research*, volume 11, pages 1323–1351, 2010.
- S. Hanneke. Theory of active learning. *Foundations and Trends in Machine Learning*, 7(2–3), 2014.
- J. Langford and A. Beygelzimer. Sensitive error correcting output codes. *COLT*, 2005.

- T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. Never-ending learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, 2015.
- M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. MIT press, 2012.
- M. Palatucci, D. Pomerleau, G. Hinton, and T. Mitchell. Zero-shot learning with semantic output codes. In *NIPS*, 2009.
- A. Singh, X. Zhu, and R. Nowak. Unlabeled data: Now it helps, now it doesn't. In *NIPS*, 2008.
- J. Steinhardt and P. Liang. Unsupervised risk estimation with only structural assumptions. 2016. (Preprint from <http://cs.stanford.edu/~jsteinhardt/publications/risk-estimation/preprint.pdf>).
- I. Steinwart. Fully adaptive density based clustering. In *Annals of Statistics*, volume 43, pages 2132–2167, 2015.
- S. Thrun. *Explanation-Based Neural Network Learning: A Lifelong Learning Approach*. Kluwer Academic Publishers, Boston, MA, 1996.
- S. Thrun and T. Mitchell. Learning one more thing. In *Proc. 14th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1217–1225, 1995a.
- Sebastian Thrun and Tom M. Mitchell. Lifelong robot learning. *Robotics and Autonomous Systems*, 15(1-2):25–46, 1995b.
- V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.

9 Appendix for Error Correcting Output Codes

First, we show that the line segment $[x, y]$ crosses the decision surface of the linear separator h_k if and only if $h(x)$ and $h(y)$ differ on the k^{th} entry.

Lemma 7. *Let $i \neq j$ be any pair of classes whose codewords disagree on the k^{th} bit. Then for any points $x \in K_i$ and $y \in K_j$, the line segment $[x, y]$ intersects with the line $h_k = 0$.*

Proof. Without loss of generality, suppose that $C_{ik} = 1$ and $C_{jk} = -1$. Then, from the definition of K_i and K_j , we have that $h_k(x) > 0$ and $h_k(y) < 0$. The function $f(t) = h_k((1-t)x + ty)$ is continuous and satisfies $f(0) = h_k(x) > 0$ and $f(1) = h_k(y) < 0$. It follows that there must be some $t_0 \in (0, 1)$ such that $f(t_0) = 0$. But this implies that the point $z = (1-t_0)x + t_0y \in [x, y]$ satisfies $h_k(z) = 0$ and it follows that $h_k = 0$ intersects with $[x, y]$ at the point z . \square

Next, we show that when the consistent linear output code makes at most β errors when predicting the code word of a new example and the Hamming distance of the code words is at least $2\beta + d + 1$, then there must be a minimum gap $g > 0$ between any pair of points belonging to different classes.

Lemma 1. *Under Assumption 1, there exists $g > 0$ s.t. if points x and x' belong to different classes, then $\|x - x'\| > g$.*

Proof. For sets A and B , let $d(A, B) = \min_{a \in A, b \in B} \|a - b\|$ denote the distance between them and recall that for each $i = 1, \dots, L$, we defined $K_i = \{x \in \mathcal{X} : d_{\text{Ham}}(h(x), C_i) \leq \beta\}$ to be the set of points that belong to class i .

Fix any pair of classes i and j and suppose for contradiction that $d(K_i, K_j) = 0$. This implies that there are two code words $c, c' \in \{\pm 1\}^m$ such that $d_{\text{Ham}}(c, C_i) \leq \beta$, $d_{\text{Ham}}(c', C_j) \leq \beta$, and the distance between $A = \{x \in \mathcal{X} : h(x) = c\}$ and $B = \{x \in \mathcal{X} : h(x) = c'\}$ is 0. First, we construct a point x that belongs to $\overline{A} \cap \overline{B}$, where \overline{A} and \overline{B} denote the closure of A and B , respectively. Since $d(A, B) = 0$, there exists a sequence of points $x_1, x_2, \dots \in A$ such that $d(x_n, B) \rightarrow 0$ as $n \rightarrow \infty$. But, since A is bounded, so is the sequence (x_n) , and therefore by the Bolzano-Weierstrass theorem, (x_n) has a convergent subsequence. Without loss of generality, suppose that (x_n) itself converges to the point x . Then x is a limit point of A and therefore belongs to the closure of A . On the other hand, since the function $z \mapsto d(z, B)$ is continuous, it follows that $d(x, B) = \lim_{n \rightarrow \infty} d(x_n, B) = 0$ and therefore x is also in the closure of B .

Now let k be any index such that the code words c and c' differ on the k^{th} entry. Next, we show that $h_k(x) = 0$. For each integer $n > 0$, let $C_n = B(x, 1/n)$ be the ball of radius $1/n$ centered at x . Since x belongs to the closure of A and C_n is a neighborhood of x , we can find some point, say x_n that belongs to the intersection $A \cap C_n$. Similarly, we can find a point y_n belonging to $B \cap C_n$. Since the line segment $[x_n, y_n]$ passes from A to B , Lemma 7 guarantees that there is a point $z_n \in [x_n, y_n] \subset C_n$ such that $h_k(z_n) = 0$. But, by construction, the sequence z_n is converging to x and, since linear functions are continuous, it follows that $h_k(x) = \lim_{n \rightarrow \infty} h_k(z_n) = 0$.

But this leads to a contradiction: since the codewords c and c' must disagree on at least $d + 1$ entries, at least $d + 1$ of the linear separators h_1, \dots, h_m intersect at the point x , which contradicts our assumption that at most d lines intersect at any point $x \in \mathcal{X}$. Therefore, we must have $d(K_i, K_j) > 0$. Since there are finitely many classes, taking $g = \min_{i,j} d(K_i, K_j)$ completes the proof. \square

Next, we prove a similar result to Theorem 1 that holds in the agnostic setting of Section 7.

Theorem 6. *Assume Assumption 1, $\text{err}_P(f^*) \leq \eta$, and p has C -thick level sets. For $0 < \epsilon \leq \eta$, suppose $\{p \geq \epsilon/(2 \text{Vol}(K))\}$ has N connected components, each with probability at least 7η . With probability at least $1 - \delta$, running Algorithm 1 with parameter $r_c < g$ on an unlabeled sample of size $n = \tilde{O}(\frac{1}{\epsilon^2}((4C)^{2d}d^{d+1}/r_c^2d + N))$ and querying $t = O(\ln N/\delta)$ labels per cluster will have error at most $\eta + \epsilon$ after querying at most Nt labels.*

Proof. Define $\lambda = \epsilon/(2 \text{Vol}(K))$ and let A_1, \dots, A_N be the connected components of $\{p \geq \lambda\}$. Since Assumption 1 holds, Lemma 1 guarantees that there is a distance $g > 0$ such that whenever $f^*(x) \neq f^*(x')$, we must have $\|x - x'\| \geq g$. This implies that for any $\lambda' > 0$, f^* must be constant on the connected components of $\{p \geq \lambda'\}$, since otherwise we could construct a pair of points closer than g with $f^*(x) \neq f^*(x')$. In particular, we know that f^* is constant on each of the A_i sets.

Since the clustering produced by Algorithm 1 does not see the labeled examples, an identical covering argument to the one in the proof of Theorem 1 shows that for $n = O((4C)^{2d} d^{d+1} / (\epsilon^2 r_c^{2d}))$ with probability at least $1 - \delta$, for each set A_i there is a unique cluster, say \hat{A}_i , such that \hat{A}_i contains $S \cap A_i$, the closest cluster to every point in A_i is \hat{A}_i . Assume this high probability event occurs.

Similarly to the proof of Theorem 1, for $n = O(\frac{N}{\epsilon^2} \ln \frac{1}{\delta})$, we have that with probability at least $1 - \delta$, for any subset of indices $I \subset [N]$, we have that

$$||S \cap A_I|/n - P_{\mathcal{X}}(A_I)| \leq \epsilon,$$

where $A_I = \bigcup_{i \in I} A_i$. Assume this high probability event occurs.

Now let y_1, \dots, y_n be the (unobserved) labels corresponding to the unlabeled sample x_1, \dots, x_n . Since $\Pr_{(x,y) \sim P}(f^*(x) \neq y) \leq \eta$, if $n = O(\frac{1}{\eta^2} \ln \frac{1}{\delta})$, then with probability at least $1 - \delta$, we have that $f^*(x_i) \neq y_i$ for at most $2\eta n$ of the sample points.

Now, for any connected component A_i , let \hat{A}_i be the cluster containing $A_i \cap S$. Since we have uniform convergence for all unions of the A_i sets, and $P_{\mathcal{X}}(A_i) \geq 7\eta$, we know that the set A_i contains at least $6\eta n$ sample points. Therefore, even if every point whose label y_i disagrees with f^* belongs to \hat{A}_i , we know that at most a $2\eta n / (6\eta n) = 1/3$ fraction of the points belonging to the cluster \hat{A}_i will have labels other than $f^*(A_i)$. If we query the label of $t = 32 \ln \frac{2N}{\delta} = O(\ln \frac{N}{\delta})$ points belonging to cluster \hat{A}_i , then with probability at least $1 - \delta/N$ the majority label will agree with f^* on A_i . Applying the union bound over the connected components A_1, \dots, A_N gives the same guarantee for all connected components with probability at least $1 - \delta$.

Let \hat{f} be the classifier output by Algorithm 1 and $Q \subset [N]$ be the indices of the A_i sets for which the algorithm queried the label of the corresponding cluster \hat{A}_i . The above arguments show that with probability at least $1 - 4\delta$, we have that $\hat{f}(x) = f^*(x)$ for any $x \in \bigcup_{i \in Q} A_i$ and, as in Theorem 1, we know that $P_{\mathcal{X}}(\bigcup_{i \notin Q} A_i) \leq \epsilon/2$. This gives the following bound on the error of \hat{f} : Let $(x, y) \sim P$, then

$$\begin{aligned} \Pr(\hat{f}(x) \neq y) &= \Pr(\hat{f}(x) \neq y, x \in \{p < \lambda\}) \\ &\quad + \Pr(\hat{f}(x) \neq y, x \in \bigcup_{i \in Q} A_i) \\ &\quad + \Pr(\hat{f}(x) \neq y, x \in \bigcup_{i \notin Q} A_i) \\ &\leq \Pr(x \in \{p < \lambda\}) \\ &\quad + \Pr(f^*(x) \neq y) \\ &\quad + \Pr(x \in \bigcup_{i \notin Q} A_i). \end{aligned}$$

By our choice of λ , the first term is at most $\epsilon/2$, by assumption the second term is at most η , and the last term is at most $\epsilon/2$, giving the final error bound of $\eta + \epsilon$. \square

10 Appendix For One-vs-all on the Unit Ball

The following result is similar to Theorem 4 and shows that Algorithm 3 continues to work in the agnostic setting of Section 7.

Theorem 7. *Suppose the data is drawn from distribution P over $\mathcal{X} \times [L]$ and that there exists a labeling function f^* such that $\Pr_{(x,y) \sim P}(f^*(x) \neq y) \leq \eta$ and Assumptions 2 and 3 hold when labels are assigned by f^* . Assume that $\Pr_{x \sim P_{\mathcal{X}}}(f^*(x) = i) \geq 19\eta$ for all classes i . For any excess error ϵ , There exists an r_c satisfying $r_c = \Omega(\epsilon c_{\text{lb}} / (c_{\text{ub}}^2 b_{\text{min}}))$ such that with probability at least $1 - \delta$, running Algorithm 3 with parameter r_c on an unlabeled sample of size $n = \tilde{O}((c_{\text{ub}}^4 d / (\epsilon^2 c_{\text{lb}}^2 b_{\text{min}}^2))^d)$ and querying $t = O(\ln \frac{N}{\delta})$ labels from each cluster will output a classifier with error at most $\eta + \epsilon$ and query at most tL labels.*

Proof. For small enough ϵ , we know that at least half of the probability mass of the points assigned to class i will belong to the ϵ -level set of $\{q_{\text{ub}}^{(i)} \geq \epsilon\}$ (in the notation of Theorem 4). Therefore, the probability mass of each of the sets

A_1, \dots, A_L in the proof of Theorem 4 is at least 9η . It follows that if we see an unlabeled set of size $n = \tilde{O}(\frac{1}{\eta^2})$, then with probability at least $1 - \delta$ every A_i set will contain at least $8\eta n$ points. Since these points belong to A_i , we know that they will be active, included in the graph G , and connected to the cluster that contains samples belonging to A_i . Moreover, under the same high probability event, we know that there are at most $2\eta n$ points whose labels disagree with f^* . Therefore, the cluster that contains samples from A_i must have at least $8\eta n$ points, at most $2\eta n$ of which can have labels that disagree with f^* , so the label assigned by f^* will account for at least a $3/4$ fraction of the points belonging to the cluster containing A_i . It follows that if we query $O(\log(L/\delta))$ labels from each A_i set then with probability at least $1 - \delta$, we will output a classification rule that agrees with f^* except with probability ϵ . It follows that the error with respect to P at most $\eta + \epsilon$. \square

11 Appendix for Boundary Features Condition

We begin by proving the probability bounds for slices of a d -dimensional ball under the uniform distribution.

Lemma 5. *Let $r > 0$ be any radius and X be a random sample drawn uniformly from the ball of radius r centered at the origin. For any width $0 \leq \rho \leq r/\sqrt{2}$, the probability that the first coordinate of X lands in $[0, \rho]$ can be bounded as follows:*

$$\sqrt{\frac{d}{2d\pi}} \frac{\rho}{r} \leq \Pr_{X \sim B(r,0)}(X_1 \in [0, \rho]) \leq \sqrt{\frac{d+1}{2\pi}} \frac{\rho}{r}.$$

Proof. Let B be the ball of radius r centered at the origin and $S = \{x \in B : x_1 \in [0, \rho]\}$ be the slice of B for which the first coordinate is in the interval $[0, \rho]$. The probability that a uniformly random sample from B lands in the subset S is given by $\text{Vol}(S)/\text{Vol}(B)$, where Vol denotes the (Lebesgue) volume of a set.

We bound the volume of the set S by writing the volume as a double integral over the first coordinate x_1 and the remaining $d - 1$ coordinates x_R .

$$\text{Vol}(S) = \int_0^\rho \int_{\mathbb{R}^{d-1}} \mathbb{I}\{\|x_R\|_2^2 \leq r^2 - x_1^2\} dx_R dx_1$$

Noticing that the inner integral is actually the volume of a $d - 1$ dimensional ball of radius $\sqrt{r^2 - x_1^2}$, and using the fact that for any d , the volume of a d -dimensional ball of radius r is $r^d v_d$, where v_d is the volume of the d -dimensional unit ball, we have

$$\text{Vol}(S) = v_{d-1} \int_0^\rho (r^2 - x_1^2)^{(d-1)/2} dx_1.$$

Upper bounding the integrand by r^{d-1} gives that $\text{Vol}(S) \leq v_{d-1} \rho r^{d-1}$. Lower bounding the integrand by $(r^2 - \rho^2)^{(d-1)/2}$ and using the fact that $\rho \leq \frac{r}{\sqrt{2}}$ we have that $\text{Vol}(S) \geq v_{d-1} \frac{1}{\sqrt{2^{d-1}}} \rho r^{d-1}$. Dividing both inequalities by the volume of B , which is $r^d v_d$, and using the fact that for all d we have $\frac{v_{d-1}}{v_d} \in [\sqrt{\frac{d}{2\pi}}, \sqrt{\frac{d+1}{2\pi}}]$ gives

$$\sqrt{\frac{d}{2d\pi}} \frac{\rho}{r} \leq \Pr_{X \sim B}(X \in S) \leq \sqrt{\frac{d+1}{2\pi}} \frac{\rho}{r},$$

as required. \square

The following is an extension of Theorem 5 to the agnostic setting described in Section 7.

Theorem 8. *Suppose the data is drawn from distribution P over $\mathcal{X} \times [L]$ and that there exists a labeling function f^* such that $\Pr_{(x,y) \sim P}(f^*(x) \neq y) \leq \eta$ and Assumptions 3 and 4 hold when labels are assigned by f^* . Moreover, assume that $\Pr_{x \sim P_{\mathcal{X}}}(f^*(x) = i) \geq 10\eta$ for all classes i . For any excess error $0 < \epsilon \leq \eta$, with probability at least $1 - \delta$, running Algorithm 4 with parameters $r \leq R/2$ and $\tau = \alpha \rho^{1/2}(r)/2$ for a known constant α on a sample of size $n = \tilde{O}(dm^2 c_{\text{ub}}^2 R^d / (c_{\text{lb}}^2 \epsilon^4))$ and querying $t = O(\ln(N/\delta))$ labels from the L largest clusters will have error at most $\eta + \epsilon$.*

Proof. In the proof of Theorem 5 we argued that with the set of hyperplanes produced by Algorithm 4 will be good approximations to the true hyperplanes. We additionally showed that the set of hyperplanes approximating one of the linear separators h_i defining the output code will agree with high probability with h_i except in a small margin and we bounded the probability mass of these margins around each h_i by ϵ . It follows that for each class i , the probability mass of the set of points in that class not contained in these margins is at least $10\eta - \epsilon \geq 9\eta$, and it follows that if our unlabeled sample is of size at least $\tilde{O}(\frac{1}{\eta^2})$ that with probability at least $1 - \delta$, we will see at least $8\eta n$ points from each class which are not contained in the small margins. Under the same high probability event, we know that at most $2\eta n$ of the labels we query can disagree with f^* , which implies that the majority label within the L largest cells will be the label predicted by f^* on these cells. It follows that if we query the labels of $O(\ln N/\delta)$ labels from each class, then with probability at least $1 - \delta$ the resulting classifier will predict labels that disagree with f^* with probability at most ϵ . It follows that the error of the classifier with respect to the distribution P is at most $\eta + \epsilon$. \square