

When can unlabeled data improve the learning rate?

Christina Göpfert

Bielefeld University

CGOEPFERT@TECHFAK.DE

Shai Ben-David

University of Waterloo

SHAI@UWATERLOO.CA

Olivier Bousquet

Google Brain

OBOUSQUET@GOOGLE.COM

Sylvain Gelly

Google Brain

SYLVAINGELLY@GOOGLE.COM

Ilya Tolstikhin

Google Brain

TOLSTIKHIN@GOOGLE.COM

Ruth Urner

York University

RUTH@EECS.YORKU.CA

Editors: Alina Beygelzimer and Daniel Hsu

Abstract

In semi-supervised classification, one is given access both to labeled and unlabeled data. As unlabeled data is typically cheaper to acquire than labeled data, this setup becomes advantageous as soon as one can exploit the unlabeled data in order to produce a better classifier than with labeled data alone. However, the conditions under which such an improvement is possible are not fully understood yet. Our analysis focuses on improvements in the *minimax* learning rate in terms of the number of labeled examples (with the number of unlabeled examples being allowed to depend on the number of labeled ones). We argue that for such improvements to be realistic and indisputable, certain specific conditions should be satisfied and previous analyses have failed to meet those conditions. We then demonstrate examples where these conditions can be met, in particular showing rate changes from $1/\sqrt{\ell}$ to $e^{-c\ell}$ and from $1/\sqrt{\ell}$ to $1/\ell$. These results improve our understanding of what is and isn't possible in semi-supervised learning.

Keywords: semi-supervised learning, learning rates, minimax analysis, unlabeled data

1. Introduction

The aim of this paper is to study and clarify different frameworks for analysis of semi-supervised learning (SSL) and formal demonstrations of benefits from unlabeled data. We outline various such frameworks that have been worked with in previous studies and map results from the literature on the theory of SSL to these frameworks. We hereby highlight how rather subtle differences in setup can lead to opposing conclusions about the formal benefits of unlabeled data. As a result, we propose and argue for a clean minimax criterion for identification of learning rate changes between supervised learning (SL) and SSL. That is, classes of data-generating distributions where the minimax expected excess risk of some

SSL algorithm converges at a rate faster than that of *any* SL algorithm as the number of labeled examples grows. Finally, we demonstrate examples of such a strict rate change.

Demonstrating advantages of SSL over supervised learning, requires upper bounding the excess risk of an SSL learner while also providing a lower bound on the risk of any supervised learner under the same distributional assumptions. We find that studies which provide both SSL upper and SL lower bounds with the gap in-between follow one of *three common patterns* in their analysis which we argue are all restrictive in one way or another.

The first common approach to showing such improvements is to grant the SSL algorithm access to the true marginal distribution (Castelli and Cover (1995); Niyogi (2013); Gliberson et al. (2017)). We call this approach *improvements via idealistic SSL*. Several papers demonstrate problems which are not learnable by any SL algorithm, while exact knowledge of the marginal leads to finite rates. Such demonstrations are typically motivated with an argument that access to sufficient amounts of unlabeled data should yield a close enough estimate of the marginal for these phenomena to also occur in a finite data regime. However, we show in Theorem 7 that such a transition from an unlearnable problem to an excess error guarantee is impossible if an SSL algorithm only sees finitely many unlabeled points (rather than the marginal). In this work we thus argue for practical settings where the extra information in SSL is limited to a *finite sample* of unlabeled points and where the amount of unlabeled data is determined by a *fixed function* of the size of the labeled dataset.

The second, and more involved way to demonstrate the desired behavior is to allow the class of distributions to depend on the number of labeled and unlabeled examples (ℓ, u) that the SSL algorithm receives. This will be referred to as *improvements via sample size dependent classes*. Darnstädt (2015, Theorem 20, Lemma 28) provides a binary classification problem over a discrete domain which is not learnable for any SL algorithm but can be successfully learned by SSL if the support of the unknown marginal happens to be bounded by some function of u . Similarly, Singh et al. (2009) consider regression problems in \mathbb{R}^d under a cluster assumption where the target function is smooth over the *decision sets* and show¹ that SSL may achieve rates faster than SL if the minimal *margin* between those sets is larger than roughly $u^{-1/d}$, but smaller than $\ell^{-1/d}$. While these settings are somewhat more practical, they involve learning against a class of distributions that grows with ℓ and u . That is, the set of distributions on which SL and SSL are compared is adapted precisely to the tasks which the SSL algorithm can solve with the data available, while the SL algorithm cannot. We instead argue for comparing SL and SSL on a fixed set of tasks.

Both approaches sketched above used knowledge of certain parameters of the marginal distribution for recovering the target function. These *relevant parameters* (manifold/discrete support, decision sets) are either given to the SSL algorithm or, in the second scenario, are estimated from unlabeled data, but then only classes of distributions for which the amount of unlabeled data was sufficient, are considered in the comparison. A third framework of comparison that has been used in theoretical analysis of SSL fixes the class of distributions by putting strong (but independent of u) assumptions on the class of marginals. For instance, constant upper bounds on the size of support (which allows to quickly find all the domain points where the marginal puts a lot of mass) or constant lower bounds on the margin (which allows to quickly discern the decision sets). We call this approach *improvements*

1. See 3rd and 5th rows in Table on page 6 of Singh et al. (2009)

via *easy marginal estimation*. In such a setup, the overall error naturally decomposes into a part from the marginal estimation and a part from label prediction problem. If the unsupervised (marginal) estimation turns out to be statistically harder than learning the target labeling function with known parameters of the marginal, SSL can use its larger unlabeled budget to eliminate the loss incurred by the former (harder) part of the problem and beat the SL. Singh et al. (2009, last row in Table on page 6) demonstrates such a case.

The three approaches discussed above have one important characteristic in common: SSL is considered on a parameterized set of distributions and it becomes possible for the SSL algorithm to estimate the relevant parameters of the unknown marginal to any given precision. However, parameter estimation is generally not a requirement for strong predictive performance. It is interesting to know whether one can demonstrate improvements of SSL without going through an intermediate step of parameter estimation. In this paper we give a positive answer to that question, demonstrating an example where estimation of the relevant parameters is impossible and yet the SSL algorithm with access to u unlabeled points achieves rates faster than SL.

The paper is structured as follows. In Section 2 we give a formal definition of “unlabeled data helps”, contrast it to previous (implicit and explicit) definitions and take a closer look at what kind of improvements have been shown in the literature. In Section 3 we explore how different types of assumptions affect the ability for SSL to achieve a change of rates. Section 4 uses a simple two-point example to illustrate that improvements via idealistic SSL or sample size dependent classes do not necessarily translate to rate improvements, and how even if SSL increases the convergence speed for *every* fixed distribution, *minimax rates* may remain unchanged, driving home the need for rigorous minimax upper and lower bound analyses. A modification of the example yields a problem where SSL improves minimax rates from $\frac{1}{\sqrt{\ell}}$ to $e^{-\ell}$. The very fast SSL rate indicates that unlabeled data essentially “solves” the problem. We show that it is possible to achieve more interesting rate changes, namely from $\frac{1}{\sqrt{\ell}}$ to $\frac{1}{\ell}$ in Section 5. Section 6 provides a discussion of our results and an outlook on future work.

2. Supervised and semi-supervised rates

2.1. Notation

In the remainder of this paper, let \mathcal{X} be an arbitrary feature space, and $\mathcal{Y} = \{0, 1\}$ the set of labels. Data are distributed according to an unknown distribution P on $\mathcal{X} \times \mathcal{Y}$ with marginal P_X on \mathcal{X} . A *classifier* is a function $f : \mathcal{X} \rightarrow \mathcal{Y}$. The *risk* of f on the distribution P is measured via the misclassification probability:

$$R_P(f) := P(f(X) \neq Y).$$

We call the function $\eta(x) = P(Y = 1|X = x)$ the *labeling function* associated to P and the corresponding classifier $f_P^*(x) = \mathbf{1}_{[\eta(x) \geq \frac{1}{2}]}$, the *Bayes classifier*. This classifier achieves the minimal possible risk among all classifiers. Given a set \mathcal{H} of classifiers of \mathcal{X} , we denote by $R_{P,\mathcal{H}}$ the minimal risk over all $h \in \mathcal{H}$:

$$R_{P,\mathcal{H}} := \inf_{h \in \mathcal{H}} R_P(h).$$

We call such a set \mathcal{H} a *hypothesis class*. Note that we are not assuming that f_P^* is a member of \mathcal{H} . When it is clear from the context, we will drop the dependence on P and write only $R_{\mathcal{H}}$. The learning objective is to identify a classifier h (not necessarily in \mathcal{H}) such that $R_P(h)$ is close to $R_{P,\mathcal{H}}$. In the semi-supervised learning (SSL) setup, we are given a training sample S made up of a set of ℓ pairs $(X_1, Y_1), \dots, (X_\ell, Y_\ell)$ drawn i.i.d. according to P and u elements X'_1, \dots, X'_u drawn i.i.d. according to P_X . The *size* of this sample is (ℓ, u) and its distribution will be denoted by $P^{(\ell, u)} := P^\ell \times P_X^u$. In the supervised learning (SL) setup, S contains only the ℓ pairs $(X_1, Y_1), \dots, (X_\ell, Y_\ell)$ and no unlabeled elements. This corresponds to a semi-supervised sample of size $(\ell, 0)$. A *semi-supervised learning algorithm* is a set of maps

$$A : (\mathcal{X} \times \mathcal{Y})^\ell \times \mathcal{X}^u \rightarrow \mathcal{Y}^{\mathcal{X}}$$

for every $(\ell, u) \in \mathbb{N}$. A supervised algorithm is such a set of maps for $\ell \in \mathbb{N}$ and $u = 0$. We will denote by $A(S)$ the classifier returned by the algorithm A upon receiving S as an input.

2.2. Learning Rates

The objective of this paper is to analyze whether SSL algorithms are more powerful than SL algorithms in a statistical learning sense. Performance of an algorithm is measured in terms of the excess risk $R(A(S)) - R_{\mathcal{H}}$ which is a random variable (as it depends on the sample). One can either study the expectation² $\mathbb{E}_{S \sim P^\ell} [(R(A(S)) - R_{\mathcal{H}})_+]$ or the tails of this random variable $P_{S \sim P^\ell}(R(A(S)) - R_{\mathcal{H}} \geq \varepsilon)$.

While it is customary in the SL case to consider the so-called *sample complexity*, that is the function $m(\varepsilon, \delta)$ such that

$$P_S(R(A(S)) - R_{\mathcal{H}} > \varepsilon) \leq \delta \quad \forall \ell \geq m(\varepsilon, \delta),$$

as a way to measure the performance of algorithms, we find it more convenient in the SSL case to work with errors instead. This means we are looking for a bound ε on the excess risk as a function of the confidence δ and the sample size (ℓ, u) . To simplify the analysis, we will also only consider the expected excess risk (this leads to weaker tail bounds in terms of δ but we will leave this as a future research direction) and thus study $\mathbb{E}_{S \sim P^{(\ell, u)}} [(R(A(S)) - R_{\mathcal{H}})_+]$ as a function of (ℓ, u) .

Our aim is to study the so-called minimax behavior of the expected excess risk, that is the performance of the best possible algorithm under the worst possible distribution, over a possibly restricted set \mathcal{P} of *admissible distributions*.

Definition 1 (Minimax expected risk) *The minimax expected excess risk of learning a hypothesis class \mathcal{H} on a sample of size (ℓ, u) over the set of admissible distributions \mathcal{P} is the expected excess risk of the best algorithm A under the worst distribution P :*

$$L(\ell, u, \mathcal{H}, \mathcal{P}) := \inf_A \sup_{P \in \mathcal{P}} \mathbb{E}_{S \sim P^{(\ell, u)}} \left[(R(A(S)) - R_{P,\mathcal{H}})_+ \right]. \quad (1)$$

2. Notice that we only consider the positive part $(x)_+ = \max(x, 0)$ as we are only interested in controlling the excess risk, even if there can be situations where the algorithm performs better than the best member of \mathcal{H} , e.g. when \mathcal{H} is very limited and the algorithm is able to generate more sophisticated classifiers (the so-called non-proper case).

Definition 2 (SL and SSL learnability) *We say that a problem $(\mathcal{H}, \mathcal{P})$ is SL learnable if $L(\ell, 0, \mathcal{H}, \mathcal{P})$ converges to zero as ℓ goes to infinity. We say that a problem $(\mathcal{H}, \mathcal{P})$ is SSL learnable if, for some function $u : \mathbb{N} \rightarrow \mathbb{N}$, $L(\ell, u(\ell), \mathcal{H}, \mathcal{P})$ converges to zero as ℓ goes to infinity.*

As stated, we use the *convergence rate* of the minimax risk as a measure for the hardness of a learning problem. For functions $f, g : \mathbb{N} \rightarrow \mathbb{R}^{\geq 0}$, we say f has *rate* g and write $f(\ell) \asymp g(\ell)$ if there exist $c_1, c_2 \in \mathbb{R}^{>0}$ such that for all ℓ , $c_1 g(\ell) \leq f(\ell) \leq c_2 g(\ell)$. The rate of convergence of $L(\ell, 0, \mathcal{H}, \mathcal{P})$ as a function of ℓ will be called the *supervised rate* (SL rate). In order to compare it to the semi-supervised situation we introduce a function $u : \mathbb{N} \rightarrow \mathbb{N}$ that relates the amount of unlabeled points to the amount of labeled points. The *semi-supervised rate* (SSL rate) for u is the rate of convergence of $L(\ell, u(\ell), \mathcal{H}, \mathcal{P})$ as a function of ℓ .

2.3. Concepts of unlabeled data helping

Since we measure the hardness of a learning problem by the convergence rate of the expected excess risk, we say that unlabeled data helps if the SSL rate is faster than the SL rate.

Definition 3 (Unlabeled data helps) *We say that unlabeled data helps to learn \mathcal{H} over the set of admissible distributions \mathcal{P} if there exists some $u : \mathbb{N} \rightarrow \mathbb{N}$ such that*

$$\lim_{\ell \rightarrow \infty} \frac{\inf_{A_{SSL}} \sup_{P \in \mathcal{P}} \mathbb{E}_{S \sim P^{(\ell, u(\ell))}} [(R(A_{SSL}(S)) - R_{\mathcal{H}})_+]}{\inf_{A_{SL}} \sup_{P \in \mathcal{P}} \mathbb{E}_{S \sim P^\ell} [(R(A_{SL}(S)) - R_{\mathcal{H}})_+]} = 0. \quad (2)$$

Note that this definition is rather restrictive. Indeed, we are looking for cases where the minimax rate (as a function of the number of labeled examples) is strictly improved by making use of a finite number of unlabeled examples. Hence situations where the expected excess error is improved by a constant factor are ruled out. Even if this constant factor is as large as the VC dimension of \mathcal{H} for example (e.g. if the expected excess error goes from d/ℓ to $1/\ell$ where d is the VC dimension), we still wouldn't consider this sufficient with our definition. In the following, we present some other definitions of ‘‘unlabeled data helps’’: the first option is allowing the class of distributions on which SL and SSL compete to depend on the size of the sample.

Definition 4 (Unlabeled data helps non-uniformly) *We say that unlabeled data helps non-uniformly to learn \mathcal{H} over the sequence of distributions $(\mathcal{P}_\ell)_{\ell \in \mathbb{N}}$ if there exists some $u : \mathbb{N} \rightarrow \mathbb{N}$ such that*

$$\lim_{\ell \rightarrow \infty} \frac{\inf_{A_{SSL}} \sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^{(\ell, u(\ell))}} [(R(A_{SSL}(S)) - R_{\mathcal{H}})_+]}{\inf_{A_{SL}} \sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^\ell} [(R(A_{SL}(S)) - R_{\mathcal{H}})_+]} = 0. \quad (3)$$

Alternatively, instead of comparing the performance over each fixed set \mathcal{P}_ℓ , we can compare how the performance of SSL progresses compared to the worst-case performance of the SL, resulting in a weaker definition of non-uniform help;

Definition 5 (Unlabeled data helps weakly non-uniformly) *We say that unlabeled data helps weakly non-uniformly to learn \mathcal{H} over the sequence of distributions $(\mathcal{P}_\ell)_{\ell \in \mathbb{N}}$ with $\mathcal{P}_\ell \subseteq \mathcal{P}_{\ell+1}$ if there exists some $u : \mathbb{N} \rightarrow \mathbb{N}$ such that*

$$\lim_{\ell \rightarrow \infty} \frac{\inf_{A_{SSL}} \sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^{(\ell, u(\ell))}} [(R(A_{SSL}(S)) - R_{\mathcal{H}})_+]}{\inf_{A_{SL}} \sup_{P \in \bigcup_{i \in \mathbb{N}} \mathcal{P}_i} \mathbb{E}_{S \sim P^\ell} [(R(A_{SL}(S)) - R_{\mathcal{H}})_+]} = 0. \quad (4)$$

Finally, if we do not wish to concern ourselves with estimating the marginal from unlabeled data, we can assume that SSL receives the marginal distribution as an additional input.

Definition 6 (Knowing the marginal helps) *We say that knowing the marginal helps to learn \mathcal{H} over the set of admissible distributions \mathcal{P} if*

$$\lim_{\ell \rightarrow \infty} \frac{\inf_{A_{SSL}} \sup_{P \in \mathcal{P}} \mathbb{E}_{S \sim P^\ell} [(R(A_{SSL}(S, P_X)) - R_{\mathcal{H}})_+]}{\inf_{A_{SL}} \sup_{P \in \mathcal{P}} \mathbb{E}_{S \sim P^\ell} [(R(A_{SL}(S)) - R_{\mathcal{H}})_+]} = 0. \quad (5)$$

2.4. Related work

Unlabeled data can be useful in many ways. For example, it has been shown to be helpful in active learning, proper learning and co-training, see e.g. [Urner et al. \(2011\)](#); [Urner and Ben-David \(2013\)](#). In this contribution, we are interested in passive semi-supervised learning without properness constraints. A significant body of work investigates how unlabeled data can provide constant-factor improvements, how it can be utilized besides improving minimax rates, and under which models and assumptions it cannot help at all. See [Kääriäinen \(2005\)](#); [Ben-David et al. \(2008\)](#); [Balcan and Blum \(2010\)](#); [Lu \(2009\)](#); [Seeger \(2000\)](#).

There are a number of papers, for example [Ratsaby and Venkatesh \(1995\)](#); [Rigollet \(2007\)](#) that formalize connections between marginal and labeling that could lead to unlabeled data being helpful in the learning process and propose SSL algorithms that exploit this connection. While they typically provide upper bounds on the performance of the proposed algorithms, they do not complement these results with lower bounds for SL algorithms and minimax analyses over worst-case distributions. We will see why this type of analysis is important in [Example 1](#). The papers that do show performance gaps under the proposed assumptions usually follow one of the following three *restrictive patterns* for SSL.

Under the first restrictive pattern, which we call *improvement via idealistic SSL*, improvement is not defined using finite samples. Instead, an SSL algorithm is assumed to have access to the marginal data distribution, and these papers investigate whether knowing the marginal helps, as defined in [Definition 6](#). [Niyogi \(2013\)](#); [Globerson et al. \(2017\)](#); [Castelli and Cover \(1995\)](#) show that under this type of SSL, it is possible to improve from unlearnable to rate $\frac{1}{\ell}$ or $\frac{1}{\sqrt{\ell}}$, depending on whether the labeling is assumed to be realizable. While these are instructive explorations of how unlabeled data *might* help, we prove in [Theorem 7](#) that this kind of result cannot carry over into the realistic setting of limited unlabeled data.

Under the second restrictive pattern, improvements stem from letting the learners fight against different sets of distributions depending on how many labeled and unlabeled samples are available. We call this pattern *improvement via sample size dependent classes*. [Singh et al. \(2009\)](#); [Azizyan et al. \(2013\)](#)³ provide settings where unlabeled data helps non-uniformly ([Definition 4](#)). Here, unlabeled data can be used to partition the feature space into regions on which the regression function is smooth enough to be approximated by a specialized learner with an improved rate. To make this partitioning possible using the unlabeled data available, but not so easy that SL can perform it adequately well using only the labeled samples, the minimal separation between the smooth regions needs to shrink as a function of the labeled and unlabeled sample size. Note that if the unlabeled data provides a constant factor rate improvement for each \mathcal{P}_ℓ , and the constant factor grows

3. 3rd and 5th rows in Table on page 6 of [Singh et al. \(2009\)](#)

without bound, then unlabeled data helps in this definition, even though for each learning problem the improvement is still only a fixed constant factor. This type of example can be constructed by letting \mathcal{H} be of infinite VC-dimension, but setting \mathcal{P}_ℓ such that the covering number of \mathcal{H} for each fixed marginal is small and a ε -cover of \mathcal{H} can be identified for each \mathcal{P}_ℓ using the $u(\ell)$ unlabeled samples, but not the ℓ labeled samples. Then an SSL algorithm can first learn a small ε -cover and perform ERM over that small set, while an SL algorithm lacks this information. This type of construction is used in [Darnstädt et al. \(2013\)](#); [Darnstädt \(2015\)](#), although they only explicitly prove that unlabeled data helps weakly non-uniformly (Definition 5).

A feature of this construction is to restrict the set of admissible distributions \mathcal{P} such that estimating some property of the marginal is beneficial to solving the learning problem, and learning this property is statistically harder than solving the learning problem under knowledge of the property, or solving it without the knowledge in the first place. This construction does not work when the property in question becomes arbitrarily hard to estimate, e.g. because the separation between decision sets goes to zero. Definitions 4 and 5 can be seen as consequences of this problem. It can be countered by bounding the hardness of the marginal estimation problem, in the case of the decision set estimation by bounding the margin width away from zero. This leads to rate improvements in [Singh et al. \(2009\)](#)⁴ in the pattern which we call *improvements via easy marginal estimation*. We give an example free from this particular restriction in Example 3.

In addition, to the best of our knowledge, no previous results show that it is possible to go from rate $\frac{1}{\sqrt{\ell}}$ to $\frac{1}{\ell}$ (the rates typically encountered in VC-class learning) using unlabeled data. We will show one way of constructing such problems in Theorem 10.

3. When Does Unlabeled Data Realistically Help?

3.1. No free learnability

If we do not restrict the set of admissible distributions \mathcal{P} and require SSL to compete with SL on all possible distributions, an improvement of rates is impossible. This is due to the fact that there exist tight lower bounds for learning in the agnostic and realizable case that do not depend on whether or not the learning algorithm has access to the marginal distribution of the data. These bounds demonstrate that for every hypothesis class, there exist marginals for which the best hypothesis is intrinsically hard to learn, so unlabeled data cannot help for unrestricted \mathcal{P} .

If we do restrict \mathcal{P} , what types of improvements are possible? The most interesting type of result would be to demonstrate a “learnability gap”: to show that there exist problems that are not learnable by any supervised learner, but can be learned using a semi-supervised learner. The existence of problem sequences where unlabeled data helps non-uniformly insinuate that a learnability gap is possible. After all, the examples in [Darnstädt et al. \(2013\)](#); [Darnstädt \(2015\)](#) may not be learnable in a supervised fashion, but each \mathcal{P}_ℓ is learnable at the same rate given enough unlabeled data. What these analyses hide is that the *number* of unlabeled samples needed grows without bound as the set of admissible distributions grows. If the unlabeled sample size is a fixed function of the labeled

4. Last row in Table on page 6

sample size (and does not depend on the target distribution), the benefit of unlabeled data vanished. Admitting an unlabeled sample of arbitrary size is clearly not a practical setting. Is it possible to modify these examples to find a learnability gap with the unlabeled sample size limited as a function of ℓ ? As the following Theorem shows, this is not the case.

Theorem 7 *If a problem $(\mathcal{H}, \mathcal{P})$ is unlearnable in the SL setting, i.e. $L(\ell, 0, \mathcal{H}, \mathcal{P})$ does not converge to 0, then it is also unlearnable in the SSL setting, i.e. for any $u : \mathbb{N} \rightarrow \mathbb{N}$, $L(\ell, u(\ell), \mathcal{H}, \mathcal{P})$ does not converge to 0.*

Proof To avoid cluttered notation, we here omit \mathcal{H} and \mathcal{P} from the error rates. We prove that if there is some u such that $\lim L(\ell, u(\ell)) = 0$ then $\lim L(\ell, 0) = 0$. Indeed, $L(\ell, 0) \geq 0$ by definition and for any ℓ and u , $L(\ell + u(\ell), 0) \leq L(\ell, u(\ell))$, since an SL algorithm can simply opt to forget the labels of $u(\ell)$ labeled samples and treat them as unlabeled. Furthermore, $L(\ell, 0)$ is non-increasing since an algorithm receiving an $\ell + 1$ sample can always ignore an example, hence $L(\ell, 0) \leq L(\ell + u(\ell), 0) \leq L(\ell, u(\ell))$. ■

On the one hand, Theorem 7 shows that results obtained under knowledge of the marginal or access to infinite amounts of unlabeled data do not necessarily carry over to realistic settings. On the other hand it demonstrates that unlabeled data can only provide realistic improvements on problems that are already learnable in a supervised setting.

3.2. Why we relate the labeling and the marginal

What *types* of assumptions on \mathcal{P} can lead to rate improvements? The admissible distributions proposed in Darnstädt et al. (2013); Globerson et al. (2017) are “product assumptions” in the sense that they restrict the space of admissible marginals and, in the case of Darnstädt et al. (2013), the admissible labelings to those realizable by the hypothesis class, but they do not restrict how marginals and labelings relate to each other: *any* combination of an admissible marginal and labeling is an admissible distribution. Note that, as evidenced by the constructions in these papers, independence of marginal and labeling *does not* immediately imply that unlabeled data or knowledge of the marginal distribution is useless for learning. While it has been shown e.g. by Seeger (2000) that, if the parameters determining marginal and labeling are independent, unlabeled data is not useful in estimating the parameters that determine the labeling, it may still be helpful in finding a low-risk classifier. For example, in the realizable case, unlabeled data can help determine which of several hypotheses that are compatible with a labeled sample is the most representative, as demonstrated by Kääriäinen (2005). Is it possible to find a set of admissible distributions such that unlabeled data helps in the sense of Definition 3 without relating marginal and labeling?

For classes of finite VC dimension, under realizability assumptions on the labeling, Darnstädt et al. (2013) answer this question in the negative up to log factors. They show that if \mathcal{H} is fixed and finite, then there is a supervised learner whose sample complexity for *every* marginal is at most $\mathcal{O}(\ln |\mathcal{H}|)$ times worse than that of a semi-supervised learner with access to the marginal. If \mathcal{H} is not finite, but of finite VC-dimension, the sample complexity benefit of using a learner specialized to the marginal can be at most $\mathcal{O}(\text{VCdim}(\mathcal{H}) \cdot \log(1/\varepsilon))$, though they do not provide examples that yield such a gap. Golovnev et al. (2019) provide lower bounds that show that when \mathcal{H} is the class of projections over $\{0, 1\}^d$, for

any supervised algorithm there exists a subset of marginals where the algorithm requires $\Omega(\text{VCdim}(\mathcal{H}))$ as many samples as an algorithm with knowledge of the marginal.

Without any restriction on the possible labeling functions, an improvement of SSL rates is equally out of the question: even a small portion of the input space that can be labeled arbitrarily leads to slow SSL rates, as shown by Theorem 9.

Definition 8 *We say that a family of probability distributions \mathcal{P} is rich for a class \mathcal{H} if there exist hypotheses $h, h' \in \mathcal{H}$ and marginal P_X such that $P_X(\{x : h(x) \neq h'(x), h(x) = 0\}) \neq P_X(\{x : h(x) \neq h'(x), h(x) = 1\})$, and for every $\alpha \in (0, \frac{1}{2})$, \mathcal{P} contains P_α and $P_{-\alpha}$ which consist of P_X paired with labeling functions η_α and $\eta_{-\alpha}$, that agree with h where $h = h'$ and take values $\frac{1}{2} + \alpha$ and $\frac{1}{2} - \alpha$ respectively where $h \neq h'$.*

Theorem 9 *Let \mathcal{H} be a class of finite VC dimension. Then for every set of probability distributions \mathcal{P} that is rich for \mathcal{H} , knowing the marginal does not help to learn \mathcal{H} over the set of admissible distributions \mathcal{P} .*

Proof We give a detailed proof of Theorem 9 in Appendix A. The main idea is that for a suitably chosen constant a , the distributions $P_{a/\sqrt{\ell}}$ and $P_{-a/\sqrt{\ell}}$ have the same marginal, and deciding between them using ℓ labeled samples has expected excess risk of order $\frac{1}{\sqrt{\ell}}$. ■

4. Relating marginal and labeling – a helpful and a non-helpful example

In the following we provide an example where knowing the marginal helps (Definition 6) and it is possible to construct a sequence of distributions such that unlabeled data helps non-uniformly and weakly non-uniformly (Definitions 4 and 5), but unlabeled data does not improve minimax rates (Definition 3).

Example 1 *Let $\mathcal{X} = \{x_1, x_2\}$ and $\mathcal{H} = \{0, 1\}^{\mathcal{X}}$. Then every marginal distribution P_X on \mathcal{X} can be parameterized by $\beta \in (-\frac{1}{2}, \frac{1}{2})$ with $P_X^\beta(x_1) = \frac{1}{2} + \beta$, $P_X^\beta(x_2) = \frac{1}{2} - \beta$. Now, for each P_X^β , restrict \mathcal{P} to contain only those P , denoted by $P^{\alpha\beta}$, such that $P(Y = 1|x_1) = \frac{1}{2} + \alpha = P(Y = 0|x_2)$ with $\alpha\beta > 0$, i.e. restrict the possible labelings such that the Bayes classifier assigns opposite labels to the two points, and labels with 1 the point that is seen more often.*

In Example 1, the Bayes classifier is completely determined by the marginal distribution. As such, this is an example where we can observe improvements via idealistic SSL, i.e. knowing the marginal helps (Definition 6): an SSL algorithm that knows the marginal has expected excess risk zero, while the SL rate of learning $(\mathcal{H}, \mathcal{P})$ is $\frac{1}{\sqrt{\ell}}$.

It is also the case that unlabeled data helps non-uniformly and weakly non-uniformly (Definitions 4 and 5). Proofs for the following lower and upper bounds can be found in Appendix B. Let $u(\ell) = \ell^2$ and set

$$\mathcal{P}_\ell := \{P^{\alpha\beta} \mid |\beta| \geq \frac{1}{\sqrt{\ell}}\}. \tag{6}$$

Then the SSL algorithm A that discards the labeled sample and uses the unlabeled sample to determine which point is more likely to appear is essentially predicting the flip of a coin with bias at least $\frac{1}{\sqrt{\ell}}$, and its expected excess risk can be bounded by

$$\sup_{P=P^{\alpha\beta} \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^{(\ell, \ell^2)}} \left[(R(A(S)) - R_{P, \mathcal{H}})_+ \right] \leq \sup_{\alpha, \beta} 2\alpha e^{-2\ell^2\beta^2} \leq e^{-2\ell} \quad (7)$$

using Hölder's inequality, while

$$\sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^\ell} \left[(R(A(S)) - R_{P, \mathcal{H}})_+ \right] \geq \frac{1}{\sqrt{\ell}}, \quad (8)$$

so

$$\lim_{\ell \rightarrow \infty} \frac{\inf_A \sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^{(\ell, u(\ell))}} \left[(R(A(S)) - R_{P, \mathcal{H}})_+ \right]}{\inf_A \sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^\ell} \left[(R(A(S)) - R_{P, \mathcal{H}})_+ \right]} = 0. \quad (9)$$

We may be tempted to think that unlabeled data helps learn \mathcal{H} for $\mathcal{P} = \bigcup_\ell \mathcal{P}_\ell$ in accordance with Definition 3. After all, the correct labeling can be estimated using only unlabeled samples. However, for *no* function u will access to $u(\ell)$ unlabeled samples change *minimax rates*. The *relevant parameter* of the marginal that determines the labeling is $\text{sgn}(\beta)$. While β itself can be estimated with high accuracy independently of its value, $\text{sgn}(\beta)$ is difficult to estimate when β is close to zero. No matter how many unlabeled samples are available, there exists some value of β such that the probability of incorrectly estimating $\text{sgn}(\beta)$ is high. Indeed, a minimax analysis shows that the SL and SSL rate are both lower bounded by $\frac{1}{\sqrt{\ell}}$ for any u . This difficulty can be overcome as is suggested by the definition of \mathcal{P}_ℓ : by bounding $|\beta|$ away from zero.

Example 2 Let $(\mathcal{H}, \mathcal{P})$ be the problem defined in Example 1. For $c \in (0, \frac{1}{4})$, let

$$\mathcal{P}' := \{P^{\alpha\beta} \mid |\beta| \geq c\}. \quad (10)$$

The SSL rate on $(\mathcal{H}, \mathcal{P}')$ is arbitrarily fast the faster u grows, while the SL rate is lower-bounded by $\frac{c}{2}e^{-32\ell c^2}$.

This is an example of *improvements via easy marginal estimation*. It yields a problem where unlabeled data helps improve from one exponential rate to any faster rate. Both SL and SSL can take advantage of the information provided by marginal estimation, but SSL overpowers SL due to the vast amounts of data it can leverage.

Can we give an example where unlabeled data helps that is free from the common patterns favoring SSL? The difficulty with using SSL to estimate parameters of the marginal is that, in contrast to what occurs e.g. when estimating the bias of a coin in supervised learning, if the properties are hard to estimate this does not imply that the excess risk of making the wrong choice is small. Indeed, as β approaches 0, the probability of incorrectly guessing $\text{sgn}(\beta)$ approaches $\frac{1}{2}$ for each fixed number of unlabeled samples, but the cost of guessing incorrectly is fixed at α . This problem can be circumvented by linking the hardness of the learning problem for the supervised learner to the excess risk of the wrong choice: by setting $\alpha = \beta$. By limiting \mathcal{P} in this way, we can achieve an improvement of minimax learning rate, without bounding the hardness of the marginal estimation.

Example 3 Let $\mathcal{X} = \{x_1, x_2\}$ and $\mathcal{H} = \{0, 1\}^{\mathcal{X}}$. Now, restrict \mathcal{P} to those distributions such that $P(Y = 1|x_i) = P_X(x_i)$. That is, the Bayes classifier labels x_1 and x_2 with opposite labels, the noise of each labeling is equal to the noise in choosing x_i , and the point that is less likely to be seen is the one which the Bayes classifier labels with 0.

Let A be the algorithm that disregards all labels, assigns 1 to the x_i that appears more often in the sample and 0 to the x_i that appears less often. The expected excess loss of this algorithm is bounded by $\frac{1}{\sqrt{\ell+u(\ell)}}$, so the SSL minimax rate for $u(\ell) = \ell^2$ is at least $f(\ell) = \frac{1}{\ell}$, the SSL minimax rate for $u(\ell) = \ell^4$ is at least $f(\ell) = \frac{1}{\ell^2}$ and the SSL minimax rate for $u(\ell) = \exp(\ell)$ is exponential. The SL rate of $(\mathcal{H}, \mathcal{P})$ is $\frac{1}{\sqrt{\ell}}$.⁵ Proofs of lower and upper bounds can be found in Appendix B.

Example 3 is a simple demonstration of a problem where SL and SSL have different *minimax* rates even though the amount of unlabeled data available to the SSL learner is limited by a fixed function. In order to achieve this improvement, we made the marginal fully determine the labeling, and connected the difficulty of learning the marginal to the excess risk of choosing the wrong hypothesis. This allows us to achieve arbitrarily fast learning rates if the number of unlabeled samples available is very large. This artificial example gives some insight into how unlabeled data can help, or fail to help. However, the restrictions on \mathcal{P} are so strong that observing labels becomes completely unnecessary. We are interested in problems where unlabeled data provides a significant, but not arbitrary advantage, and collecting at least some labels is unavoidable. The types of rates most common to supervised PAC learning are $\frac{1}{\sqrt{\ell}}$ and $\frac{1}{\ell}$. Is it possible to construct a problem that has the slower SL rate and the faster SSL rate, but does not become arbitrarily fast? We will show that this is indeed possible. Our proof uses a general-purpose construction that can be used to create further problems with intermediate rate pairs.

5. A non-trivial rate change

In the following, we show that there exist problems of learning over classes of finite VC dimension where SSL leads to “intermediate rates”. That is, we provide an example of a problem where unlabeled data helps improve rates from $\frac{1}{\sqrt{\ell}}$ to $\frac{1}{\ell}$, but no $u : \mathbb{N} \rightarrow \mathbb{N}$ can make the rate improve beyond $\frac{1}{\ell}$.

Theorem 10 *There exists a problem with SL rate $\frac{1}{\sqrt{\ell}}$ and SSL rate $\frac{1}{\ell}$ that cannot be improved by increasing the amount of unlabeled data.*

Proof The proof of the Theorem relies on mixing two types of problems: one with SL and SSL rate $\frac{1}{\ell}$ (i.e. unlabeled data does not help), and another with SL rate $\frac{1}{\sqrt{\ell}}$, but very fast SSL rate (i.e. unlabeled data helps a lot). This mixture is defined as follows: Let $(\mathcal{H}_A, \mathcal{P}_A)$ be a problem with domain \mathcal{X}_A and $(\mathcal{H}_B, \mathcal{P}_B)$ be a problem with domain \mathcal{X}_B . Without loss of generality, let $\mathcal{X}_A \cap \mathcal{X}_B = \emptyset$. Then the *mixture problem* $(\mathcal{H}, \mathcal{P})$ is given by

$$\begin{aligned} \mathcal{H} &:= \{h : \mathcal{X}_A \cup \mathcal{X}_B \rightarrow \{0, 1\} : h|_{\mathcal{X}_A} \in \mathcal{H}_A, h|_{\mathcal{X}_B} \in \mathcal{H}_B\}, \\ \mathcal{P} &:= \left\{ \frac{1}{2} \cdot P_A + \frac{1}{2} \cdot P_B : P_A \in \mathcal{P}_A, P_B \in \mathcal{P}_B \right\}. \end{aligned}$$

5. A similar result could be achieved by requiring that $|\frac{1}{2} - P(Y = 1|X = x_i)| \leq |\frac{1}{2} - P(X = x_i)|$.

In Lemmas 12 and 13 in Appendix C, we prove that if we assume that for every $P \in \mathcal{P}$ the Bayes classifier is in \mathcal{H} , then the SL and SSL rates of the mixture problem are dictated by the respective slower rates of the component problems. Now let problem A be Example 3 and let problem B be the problem of learning a hypothesis class with finite VC-dimension under the realizability assumption. According to (Darnstädt et al., 2013, Theorem 1), B has SL and SSL rate $\frac{1}{\ell}$ for any u . The two-point example from Example 3 has SL rate $\frac{1}{\sqrt{\ell}}$ and SSL rate faster than $\frac{1}{\ell}$ for $u(\ell) \geq \ell^2$. Then the mixture problem has SL rate $\frac{1}{\sqrt{\ell}}$ and SSL rate $\frac{1}{\ell}$ for any u with $u(\ell) \geq \ell^2$. ■

5.1. How much unlabeled data is enough?

Our definition of “unlabeled data helps” allows the algorithm to have access to unlabeled samples of arbitrary size, as long as it is controlled by some function of the labeled sample size. Is it possible to estimate *how much* unlabeled data is needed for a rate improvement? It is possible to give a lower bound on the required growth of u , depending on the SL rate.

Proposition 11 *If the SL rate of $(\mathcal{H}, \mathcal{P})$ is $\ell^{-\alpha}$ with $\alpha > 0$, and unlabeled data helps for $u : \mathbb{N} \rightarrow \mathbb{N}$, then $\frac{u(\ell)}{\ell} \rightarrow \infty$, i.e. u must grow superlinearly.*

Proof Since $L(\ell + u(\ell), 0) < L(\ell, u(\ell))$, $\frac{L(\ell, u(\ell))}{L(\ell, 0)} \rightarrow 0$ implies that $\frac{L(\ell + u(\ell), 0)}{L(\ell, 0)} \rightarrow 0$. Applying these conditions to $L(\ell, 0) = \ell^{-\alpha}$ yields the result. ■

6. Discussion

We have highlighted how subtly different concepts of SSL improvements can lead to opposing results on the value of unlabeled data for classification problems in a VC framework. In particular, we have highlighted three common patterns of SSL analyses and demonstrated their restrictions on simple examples. We have also demonstrated that different types of rate improvements through SSL in a finite VC setting are possible in examples free of those restrictive patterns. Our examples are artificial and meant to clarify conditions under which unlabeled data helps. Whether such conditions can arise naturally, or indeed what natural conditions mean in this context, remains an interesting open question. Moreover, it would be interesting to further investigate the difference between settings where SSL helps because it is assumed that certain margin parameters are easy to estimate, versus settings where this is not assumed. Finally, it would be interesting to further examine under what types of assumptions finite amounts of unlabeled data can affect rates for hypothesis classes of infinite VC dimension, and whether this case differs markedly from non-parametric learning.

Acknowledgments

Part of this work was completed while the first author was at Google. In addition, funding by the CITEC cluster of excellence and by the BMBF under grant number 01IS18041A is gratefully acknowledged.

References

- Martin Azizyan, Aarti Singh, and Larry Wasserman. Density-sensitive semisupervised inference. *The Annals of Statistics*, 41(2):751–771, 2013.
- Maria-Florina Balcan and Avrim Blum. A discriminative model for semi-supervised learning. *Journal of the ACM (JACM)*, 57(3):19, 2010.
- Shai Ben-David, Tyler Lu, and Dávid Pál. Does unlabeled data provably help? worst-case analysis of the sample complexity of semi-supervised learning. In *COLT*, pages 33–44, 2008.
- Vittorio Castelli and Thomas M Cover. On the exponential value of labeled samples. *Pattern Recognition Letters*, 16(1):105–111, 1995.
- Malte Darnstädt. *An investigation on the power of unlabeled data*. Doctoral thesis, Ruhr-Universität Bochum, 2015.
- Malte Darnstädt, Hans Ulrich Simon, and Balázs Szörényi. Unlabeled data does provably help. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 20. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.
- Amir Globerson, Roi Livni, and Shai Shalev-Shwartz. Effective semisupervised learning on manifolds. In *COLT*, pages 978–1003, 2017.
- Alexander Golovnev, Dávid Pál, and Balázs Szörényi. The information-theoretic value of unlabeled data in semi-supervised learning. *arXiv preprint arXiv:1901.05515*, 2019.
- Matti Kääriäinen. Generalization error bounds using unlabeled data. In *COLT*, pages 127–142, 2005.
- Tyler Tian Lu. *Fundamental limitations of semi-supervised learning*. Doctoral thesis, University of Waterloo, 2009.
- Partha Niyogi. Manifold regularization and semi-supervised learning: some theoretical analyses. *Journal of Machine Learning Research*, 14(1):1229–1250, 2013.
- Joel Ratsaby and Santosh S Venkatesh. Learning from a mixture of labeled and unlabeled examples with parametric side information. In *COLT*, pages 412–417, 1995.
- Philippe Rigollet. Generalization error bounds in semi-supervised classification under the cluster assumption. *Journal of Machine Learning Research*, 8(Jul):1369–1392, 2007.
- Matthias W. Seeger. Input-dependent regularization of conditional density models. 2000.
- Aarti Singh, Robert Nowak, and Xiaojin Zhu. Unlabeled data: Now it helps, now it doesn't. In *NIPS*, pages 1513–1520, 2009.
- Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer series in statistics. Springer, 2009.

Ruth Urner and Shai Ben-David. Probabilistic lipschitzness a niceness assumption for deterministic labels. In *Learning Faster from Easy Data-Workshop@ NIPS*, 2013.

Ruth Urner, Shai Shalev-Shwartz, and Shai Ben-David. Access to unlabeled data can speed up prediction time. In *ICML*, pages 641–648, 2011.

Appendix A. Proof of Theorem 9

Since \mathcal{H} has finite VC dimension, the SL rate of $(\mathcal{H}, \mathcal{P})$ is upper-bounded by $\frac{1}{\sqrt{\ell}}$. Showing that the rate of any algorithm with access to the marginal distribution is lower-bounded by $\frac{1}{\sqrt{\ell}}$ proves that both the SL and SSL rates are of order $\frac{1}{\sqrt{\ell}}$, since the SSL rate cannot be slower than the SL rate. Let $C := \{x : h(x) \neq h'(x)\}$, let $c := P_X(C)$ and $c' := P_X(\{x : x \in C \wedge h(x) = 1\})$. By requirement, $c' \neq c/2$. Without loss of generality, assume $c' > c/2$. A simple calculation shows that

$$\text{KL}(P_\alpha^\ell, P_{-\alpha}^\ell) = 2cl\alpha \log\left(\frac{1+2\alpha}{1-2\alpha}\right). \quad (11)$$

Using

$$\frac{1+x}{1-x} = 1 + \frac{2x}{1-x} \quad \text{and} \quad \log(1+x) < x \quad \forall x > 0, \quad (12)$$

we find that

$$\text{KL}(P_\alpha^\ell, P_{-\alpha}^\ell) \leq 2cl\alpha \frac{4\alpha}{1-2\alpha}. \quad (13)$$

For $\alpha < \frac{1}{4}$, this can be bounded by

$$\text{KL}(P_\alpha^\ell, P_{-\alpha}^\ell) \leq 16cl\alpha^2. \quad (14)$$

(Tsybakov, 2009, Theorem 2.2, (i)) shows that for any hypothesis test that decides between P_α and $P_{-\alpha}$, the probability of choosing incorrectly is lower-bounded by $\frac{1-a}{2}$, where $a \geq \text{TV}(P_\alpha, P_{-\alpha})$ and TV denotes the total variation distance. Since

$$\text{TV}(P_\alpha^\ell, P_{-\alpha}^\ell) \leq \sqrt{\frac{1}{2} \text{KL}(P_\alpha^\ell, P_{-\alpha}^\ell)}, \quad (15)$$

the probability of a test choosing incorrectly between P_α and $P_{-\alpha}$ can be lower-bounded by

$$\frac{1 - \sqrt{8cl\alpha^2}}{2}. \quad (16)$$

Since P_α and $P_{-\alpha}$ have the same marginal distribution, this probability is independent of whether or not the marginal is known to the learner. If P_α is the true underlying distribution, any hypothesis that does not majorly label C with 1 incurs an excess risk of at least $2\alpha(c' - c/2)$. Likewise, if $P_{-\alpha}$ is the true underlying distribution, any hypothesis that does not majorly label C with 0 incurs an excess risk of at least $2\alpha(c' - c/2)$. Let $\alpha = \frac{1}{\sqrt{32lc}}$. Then Equation (16) shows that the probability of choosing incorrectly between P_α and $P_{-\alpha}$ can be lower-bounded by $\frac{1}{4}$, and the expected excess risk of any algorithm can be lower-bounded by

$$\frac{2c' - c}{16\sqrt{2c}} \frac{1}{\sqrt{\ell}} \succsim \frac{1}{\sqrt{\ell}}. \quad (17)$$

Appendix B. Two-point example bounds

Let $\mathcal{X} = \{x_1, x_2\}$ and $\mathcal{H} = \{h_{01}, h_{10}\}$ with $h_{01}(x_1) = h_{10}(x_2) = 0$ and $h_{01}(x_2) = h_{10}(x_1) = 1$. Let $\alpha, \beta \in (0, \frac{1}{2})$. Let $P_{\alpha\beta+}$ denote the distribution on $\mathcal{X} \times \{0, 1\}$ with

$$\begin{aligned} P_{\alpha\beta+}(x_1, 0) &= \left(\frac{1}{2} + \beta\right) \left(\frac{1}{2} - \alpha\right) & P_{\alpha\beta+}(x_2, 0) &= \left(\frac{1}{2} - \beta\right) \left(\frac{1}{2} + \alpha\right) \\ P_{\alpha\beta+}(x_1, 1) &= \left(\frac{1}{2} + \beta\right) \left(\frac{1}{2} + \alpha\right) & P_{\alpha\beta+}(x_2, 1) &= \left(\frac{1}{2} - \beta\right) \left(\frac{1}{2} - \alpha\right) \end{aligned}$$

and let $P_{\alpha\beta-}$ denote the distribution on $\mathcal{X} \times \{0, 1\}$ with

$$\begin{aligned} P_{\alpha\beta-}(x_1, 0) &= \left(\frac{1}{2} - \beta\right) \left(\frac{1}{2} + \alpha\right) & P_{\alpha\beta-}(x_2, 0) &= \left(\frac{1}{2} + \beta\right) \left(\frac{1}{2} - \alpha\right) \\ P_{\alpha\beta-}(x_1, 1) &= \left(\frac{1}{2} - \beta\right) \left(\frac{1}{2} - \alpha\right) & P_{\alpha\beta-}(x_2, 1) &= \left(\frac{1}{2} + \beta\right) \left(\frac{1}{2} + \alpha\right). \end{aligned}$$

We can think of x_1 and x_2 as two coins. Then $P_{\alpha\beta+}$ and $P_{\alpha\beta-}$ are both distributions where the coins are biased in opposite directions with bias size α . For $P_{\alpha\beta+}$, the first coin is the one we flip more often, and also the one more likely to show heads, while for $P_{\alpha\beta-}$ this is true for the second coin. In the following, we will show that

- For the set of admissible distributions

$$\mathcal{P}_0 = \{P_{\alpha\beta+}, P_{\alpha\beta-} \mid \alpha, \beta \in \left(0, \frac{1}{2}\right)\}, \quad (18)$$

the SL and SSL rate are both $\asymp \frac{1}{\sqrt{\ell}}$ (for any function u relating labeled and unlabeled sample size)

- For the set of admissible distributions

$$\mathcal{P}_1 = \{P_{\alpha\beta+}, P_{\alpha\beta-} \mid \alpha, \beta \in \left(0, \frac{1}{2}\right), \alpha = \beta\}, \quad (19)$$

the SL rate is $\asymp \frac{1}{\sqrt{\ell}}$ while the SSL rate can become arbitrarily fast.

- For the set of admissible distributions

$$\mathcal{P}_\ell = \{P_{\alpha\beta+}, P_{\alpha\beta-} \mid \alpha \in \left(0, \frac{1}{2}\right), \beta \in \left(\frac{1}{\sqrt{\ell}}, \frac{1}{2}\right)\}, \quad (20)$$

$$\inf_{A_{SSL}} \sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^{(\ell, \ell^2)}} [(R(A_{SSL}(S)) - R_{\mathcal{H}})_+] \leq e^{-\ell}$$

$$\inf_{A_{SL}} \sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^\ell} [(R(A_{SL}(S)) - R_{\mathcal{H}})_+] \geq \frac{1}{\sqrt{\ell}}$$

- For the set of admissible distributions

$$\mathcal{P}_c = \{P_{\alpha\beta+}, P_{\alpha\beta-} \mid \alpha \in \left(0, \frac{1}{2}\right), \beta \in \left(c, \frac{1}{2}\right)\}, \quad (21)$$

with $c \in (0, \frac{1}{4})$, the supervised rate is lower-bounded by $\frac{c}{2}e^{-32\ell c^2}$.

B.1. Lower bounds

Let $P_0 = P_{\alpha\beta+}^\ell \times P_{\alpha\beta+|\mathcal{X}}^u$ and $P_1 = P_{\alpha\beta-}^\ell \times P_{\alpha\beta-|\mathcal{X}}^u$. Then a simple calculation shows that

$$\text{KL}(P_0, P_1) = 2\ell\alpha \log\left(\frac{1+2\alpha}{1-2\alpha}\right) + 2(\ell+u)\beta \log\left(\frac{1+2\beta}{1-2\beta}\right). \quad (22)$$

Using

$$\frac{1+x}{1-x} = 1 + \frac{2x}{1-x} \quad \text{and} \quad \log(1+x) < x \quad \forall x > 0, \quad (23)$$

we find that

$$\text{KL}(P_0, P_1) \leq 2\ell\alpha \frac{4\alpha}{1-2\alpha} + 2(\ell+u)\beta \frac{4\beta}{1-2\beta} \quad (24)$$

Now assume that $\alpha, \beta < \frac{1}{4}$, then

$$\text{KL}(P_0, P_1) \leq 16\ell\alpha^2 + 16(\ell+u)\beta^2. \quad (25)$$

We will now lower-bound

$$\inf_A \sup_{P \in \{P_0, P_1\}} \mathbb{E}_{S \sim P^\ell \times P_X^u} \left[(R(A(S)) - R_{P, \mathcal{H}})_+ \right]. \quad (26)$$

Since the hypothesis class contains only two hypotheses, one of them optimal, the expected excess loss for each distribution is equal to the excess loss of the wrong hypothesis times the probability that the algorithm chooses the wrong hypothesis. The excess loss of the non-optimal hypothesis is 2α . The minimal probability of choosing the wrong hypothesis is equal to the minimal probability of error in hypothesis testing between P_0 and P_1 . (Tsybakov, 2009, Theorem 2.2, (i)) shows that for any hypothesis test between P_0 and P_1 , the probability of guessing incorrectly is lower-bounded by $\frac{1-a}{2}$, where $a \geq \text{TV}(P_0, P_1)$ and TV denotes the total variation distance. Since

$$\text{TV}(P_0, P_1) \leq \sqrt{\frac{1}{2} \text{KL}(P_0, P_1)}, \quad (27)$$

the probability of a test choosing incorrectly between P_0 and P_1 can be lower-bounded by

$$\frac{1 - \sqrt{8\ell\alpha^2 + 8(\ell+u)\beta^2}}{2}. \quad (28)$$

Let $\alpha = \frac{1}{8\sqrt{\ell}}$ and $\beta = \frac{1}{8\sqrt{\ell+u}}$. Then Equation (28) shows that the probability of choosing incorrectly between P_0 and P_1 can be lower-bounded by $\frac{1}{4}$, so the minimax expected excess risk over the set of admissible distributions $\{P_0, P_1\}$ can be bounded by

$$\inf_A \sup_{P \in \{P_0, P_1\}} \mathbb{E}_{S \sim P^\ell \times P_X^u} \left[(R(A(S)) - R_{P, \mathcal{H}})_+ \right] \geq 2 \frac{1}{8\sqrt{\ell}} \frac{1}{4} = \frac{1}{16\sqrt{\ell}}. \quad (29)$$

If $P_0, P_1 \in \mathcal{P}$,

$$\inf_A \sup_{P \in \mathcal{P}} \mathbb{E}_{S \sim P^\ell \times P_X^u} \left[(R(A(S)) - R_{P, \mathcal{H}})_+ \right] \geq \inf_A \sup_{P \in \{P_0, P_1\}} \mathbb{E}_{S \sim P^\ell \times P_X^u} \left[(R(A(S)) - R_{P, \mathcal{H}})_+ \right], \quad (30)$$

so Equation (29) also provides lower bounds for any set of admissible distributions that contains P_0 and P_1 . In particular,

- $P_0, P_1 \in \mathcal{P}_0$, so Equation (29) shows that the semi-supervised rate (and thus the supervised rate) on \mathcal{P}_0 are lower-bounded by $\frac{1}{16\sqrt{\ell}}$,
- if $u = 0$, $\frac{1}{8\sqrt{\ell}} = \frac{1}{8\sqrt{\ell+u}}$, so that $P_0, P_1 \in \mathcal{P}_1$ and Equation (29) shows that the supervised rate on \mathcal{P}_1 is lower-bounded by $\frac{1}{16\sqrt{\ell}}$.
- $P_0, P_1 \in \mathcal{P}_\ell$, so

$$\inf_{ASL} \sup_{P \in \mathcal{P}_\ell} \mathbb{E}_{S \sim P^\ell} [(R(ASL(S)) - R_{\mathcal{H}})_+] \geq \frac{1}{16\sqrt{\ell}} \quad (31)$$

Applying [Tsybakov \(2009\)](#)[Theorem 2.2 (iii)] with the bound on the KL divergence in Equation (25) with $\alpha = \beta = c$ yields the lower bound on the supervised rate for \mathcal{P}_c .

B.2. Upper bounds

Since \mathcal{H} is finite, ERM learns \mathcal{H} in a supervised fashion over the set of all possible distributions at rate $\frac{1}{\sqrt{\ell}}$. The supervised and semi-supervised rates for \mathcal{P}_0 and \mathcal{P}_1 cannot be slower than $\frac{1}{\sqrt{\ell}}$, yielding upper bounds for \mathcal{P}_0 and \mathcal{P}_1 .

On the other hand, for each fixed distribution $P_{\alpha\beta}$, an algorithm may try to discard all labels and attempt to solve the learning problem by guessing $\text{sgn}(\beta)$ based on which point appears more often in the sample. The probability that this algorithm chooses the incorrect hypothesis is equal to the probability of a binomial variable with parameters $n = \ell + u$ and $p = \frac{1}{2} + \beta$ taking a value of more than $\frac{\ell+u}{2}$. This probability can be bounded by $e^{-2\beta^2(\ell+u)}$ using Hoeffding's inequality. The excess loss of choosing the wrong hypothesis is 2α . If $\alpha = \beta$, as is the case for all distributions in \mathcal{P}_1 , the expected excess risk is bounded by $2\alpha e^{-2\alpha^2(\ell+u)}$ and reaches its maximum of $\frac{1}{\sqrt{\ell+u}} e^{-\frac{1}{2}}$ for $\alpha = \frac{1}{2\sqrt{\ell+u}}$. This can be made smaller than any desired $r(\ell)$ by choosing $u(\ell) > e^{-\frac{1}{2}} \frac{1}{r(\ell)}$. Applying this algorithm to \mathcal{P}_ℓ with $u = \ell^2$ yields the upper bound $2 \cdot 1 \cdot e^{-2\ell}$.

Appendix C. Bounds for mixture problems

Let $(\mathcal{H}_A, \mathcal{P}_A)$ be a problem with domain \mathcal{X}_A and $(\mathcal{H}_B, \mathcal{P}_B)$ be a problem with domain \mathcal{X}_B . Assume that \mathcal{H}_A contains the Bayes classifier for all distributions in \mathcal{P}_A and \mathcal{H}_B contains the Bayes classifier for all distributions in \mathcal{P}_B . Without loss of generality, let $\mathcal{X}_A \cap \mathcal{X}_B = \emptyset$. Let the mixture problem $(\mathcal{H}, \mathcal{P})$ be given by

$$\begin{aligned} \mathcal{H} &:= \{h : \mathcal{X}_A \cup \mathcal{X}_B \rightarrow \{0, 1\} : h|_{\mathcal{X}_A} \in \mathcal{H}_A, h|_{\mathcal{X}_B} \in \mathcal{H}_B\}, \\ \mathcal{P} &:= \left\{ \frac{1}{2} P_A + \frac{1}{2} \cdot P_B : P_A \in \mathcal{P}_A, P_B \in \mathcal{P}_B \right\}. \end{aligned}$$

In the following, we prove upper and lower bounds on the SL and SSL rates of the mixture problem. Since a supervised learning rate $r(\ell)$ is equivalent to an SSL rate $r(\ell)$ with $u(\ell) = 0$, we only need to consider SSL rates in our proofs. All results transfer directly to SL rates.

Lemma 12 *If $L(\ell, u(\ell), \mathcal{H}_A, \mathcal{P}_A) \geq r_A(\ell)$ for u and $L(\ell, u(\ell), \mathcal{H}_B, \mathcal{P}_B) \geq r_B(\ell)$, then*

$$L(\ell, u(\ell), \mathcal{H}, \mathcal{P}) \geq \max\left(\frac{1}{2}r_A(\ell), \frac{1}{2}r_B(\ell)\right). \quad (32)$$

In particular, the SSL rate of $(\mathcal{H}, \mathcal{P})$ for u is at least as slow as the slower rate of the mixture components.

Proof Since the Bayes classifier is contained in the hypothesis class, $R_P(\mathcal{A}(S)) - R_{P, \mathcal{H}} > 0$ for all P and we can omit $(\cdot)_+$ in the minimax expected excess risk. Assume that $L(\ell, u(\ell), \mathcal{H}, \mathcal{P})$ is not lower bounded by $\frac{1}{2}r_A(\ell)$. Then there exists an algorithm \mathcal{A} such that for all $P_A \in \mathcal{P}_A, P_B \in \mathcal{P}_B$, if $P \sim CP_A + (1 - C)P_B$, then

$$\mathbb{E}_{S \sim P(\ell, u(\ell))} [R_P(\mathcal{A}(S)) - R_{P, \mathcal{H}}] < \frac{1}{2}r_A(\ell). \quad (33)$$

Recall that

$$R_P(\mathcal{A}(S)) - R_{P, \mathcal{H}} = \frac{1}{2}(R_{P_A}(\mathcal{A}(S)|_{\mathcal{X}_A}) - R_{P_A, \mathcal{H}_A}) + \frac{1}{2}(R_{P_B}(\mathcal{A}(S)|_{\mathcal{X}_B}) - R_{P_B, \mathcal{H}_B}), \quad (34)$$

so Equation (33) implies that

$$\frac{1}{2}\mathbb{E}_{S \sim P^\ell \times P_X^{u(\ell)}} [R_{P_A}(\mathcal{A}(S)|_{\mathcal{X}_A}) - R_{P_A, \mathcal{H}_A}] < \frac{1}{2}r_A(\ell). \quad (35)$$

Let \mathcal{A}_A be the algorithm that learns $(\mathcal{H}_A, \mathcal{P}_A)$ from a sample S' by

1. discarding part of the sample S' , where the number of discarded points follows a binomial distribution,
2. sampling from P_B as many labeled and unlabeled points as were previously discarded, resulting in a sample S'' from P_A and P_B ,
3. applying the algorithm \mathcal{A} to the new sample S'' , and
4. restricting the resulting mapping $\mathcal{A}(S'')$ to \mathcal{X}_A .

Then Equation (35) implies that $L(\ell, u(\ell), \mathcal{H}_A, \mathcal{A}_A, \mathcal{P}_A) < r_A(\ell)$, which contradicts the assumption that $r_A(\ell)$ is a lower bound on the minimax expected excess loss of learning $(\mathcal{H}_A, \mathcal{P}_A)$. The analogous argument for $r_B(\ell)$ proves the Lemma. \blacksquare

Lemma 13 *Let $u(\ell) \geq \ell$. If $L(\ell, u(\ell), \mathcal{H}_A, \mathcal{P}_A) \leq r_A(\ell)$ and $L(\ell, u(\ell), \mathcal{H}_B, \mathcal{P}_B) \leq r_B(\ell)$ for u , then $L(\ell, u(\ell), \mathcal{H}, \mathcal{P}) \leq \max(r_A(\frac{\ell}{4}), r_B(\frac{\ell}{4})) + \exp(-C\ell) + \exp(-Cu(\ell))$. In particular, the SSL rate of $(\mathcal{H}, \mathcal{P})$ for u is at most as slow as the slower rate of the mixture components, if the slower rate is not faster than exponential.*

Proof Since the Bayes classifier is contained in the hypothesis class, $R_P(\mathcal{A}(S)) - R_{P, \mathcal{H}} > 0$ for all P and we can omit $(\cdot)_+$ in the minimax expected excess risk. Let \mathcal{A}_A and \mathcal{A}_B be algorithms corresponding to problems A and B whose rates are bounded by r_A and r_B , respectively. Given a sample S , let $S_{\bullet}^l := S \cap \mathcal{X}_{\bullet} \times \mathcal{Y}$ and $S_{\bullet}^u := S \cap \mathcal{X}_{\bullet}$ be the labeled and

unlabeled examples for problems $\bullet \in \{A, B\}$, respectively. Further, let $S|_A = S|_A^l \cup S|_A^u$ and $S|_B = S|_B^l \cup S|_B^u$. Then $(\mathcal{H}, \mathcal{P})$ can be solved by the algorithm

$$\mathcal{A}(S)(x) := \begin{cases} \mathcal{A}_A(S|_A)(x) & x \in \mathcal{X}_A \\ \mathcal{A}_B(S|_B)(x) & x \in \mathcal{X}_B. \end{cases} \quad (36)$$

The excess risk of this algorithm trained on the sample S is

$$\frac{1}{2} (R(\mathcal{A}_A(S|_A)) - R_{\mathcal{H}_A}) + \frac{1}{2} (R(\mathcal{A}_B(S|_B)) - R_{\mathcal{H}_B}). \quad (37)$$

The expected excess risk can be bounded by

$$2e^{-\frac{\ell}{8}} + 2e^{-\frac{u(\ell)}{8}} + \frac{1}{2}r_A \binom{\ell}{\frac{\ell}{4}} + \frac{1}{2}r_B \binom{\ell}{\frac{\ell}{4}}. \quad (38)$$

The exponential terms arise because there is always some (albeit small) possibility that the sample S contains only few examples for one of the problems, in which case the excess loss may be large. The number of labeled and unlabeled samples for problem A in S can be thought of as realizations of binomial random variables $C_\ell \sim B(\ell, 0.5)$ and $C_u \sim B(u(\ell), 0.5)$. The probability that we see less than $\frac{\ell}{4}$ labeled examples for either problem, or less than $\frac{u(\ell)}{4}$ unlabeled samples for either problem, can be bounded by $2e^{-\frac{\ell}{8}}$ and $2e^{-\frac{u(\ell)}{8}}$, respectively, using Hoeffding's inequality. The expected excess risk over samples with at least $\frac{\ell}{4}$ labeled and $\frac{u(\ell)}{4}$ unlabeled examples for each problem can be bounded using risk bounds for the lowest possible number of samples, and observing that the probabilities of all sample number combinations sum to (less than) one:

$$\begin{aligned} & \sum_{i=\frac{\ell}{4}}^{\frac{3\ell}{4}} \sum_{j=\frac{u(\ell)}{4}}^{\frac{3u(\ell)}{4}} P(C_\ell = i, C_u = j) \frac{1}{2} \left(\mathbb{E}_{S_A \sim P_A^{(i,j)}} [R(\mathcal{A}_A(S_A))] + \mathbb{E}_{S_B \sim P_B^{(\ell-i, u(\ell)-j)}} [R(\mathcal{A}_B(S_B))] \right) \\ & \leq \sum_{i=\frac{\ell}{4}}^{\frac{3\ell}{4}} \sum_{j=\frac{u(\ell)}{4}}^{\frac{3u(\ell)}{4}} P(C_\ell = i, C_u = j) \frac{1}{2} \left(\mathbb{E}_{S_A \sim P_A^{\left(\frac{\ell}{4}, \frac{u(\ell)}{4}\right)}} [R(\mathcal{A}_A(S_A))] + \mathbb{E}_{S_B \sim P_B^{\left(\frac{\ell}{4}, \frac{u(\ell)}{4}\right)}} [R(\mathcal{A}_B(S_B))] \right) \\ & \leq 1 \cdot \frac{1}{2} \left(\mathbb{E}_{S_A \sim P_A^{\left(\frac{\ell}{4}, \frac{u(\ell)}{4}\right)}} [R(\mathcal{A}_A(S_A))] + \mathbb{E}_{S_B \sim P_B^{\left(\frac{\ell}{4}, \frac{u(\ell)}{4}\right)}} [R(\mathcal{A}_B(S_B))] \right). \end{aligned}$$

If $u(\ell) \geq \ell$, $\frac{u(\ell)}{4} \geq u\left(\frac{\ell}{4}\right)$, so the previous quantity can be bounded by

$$\begin{aligned} & \frac{1}{2} \left(\mathbb{E}_{S_A \sim P_A^{\frac{\ell}{4}} \times P_{X_A}^{u\left(\frac{\ell}{4}\right)}} [R(\mathcal{A}_A(S_A))] + \mathbb{E}_{S_B \sim P_B^{\frac{\ell}{4}} \times P_{X_B}^{u\left(\frac{\ell}{4}\right)}} [R(\mathcal{A}_B(S_B))] \right) \\ & \leq \frac{1}{2} L \left(\frac{\ell}{4}, u \left(\frac{\ell}{4} \right), \mathcal{H}_A, \mathcal{A}_A, \mathcal{P}_A \right) + \frac{1}{2} L \left(\frac{\ell}{4}, u \left(\frac{\ell}{4} \right), \mathcal{H}_B, \mathcal{A}_B, \mathcal{P}_B \right) \\ & \leq \frac{1}{2} r_A \binom{\ell}{\frac{\ell}{4}} + \frac{1}{2} r_B \binom{\ell}{\frac{\ell}{4}}. \end{aligned}$$

■