A THEORETICALLY GROUNDED CHARACTERIZATION OF FEATURE REPRESENTATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

A large body of work has explored how learned feature representations can be useful for a variety of downstream tasks. This is true even when the downstream tasks differ greatly from the actual objective used to (pre)train the feature representation. This observation underlies the success of, e.g., few-shot learning, transfer learning and self-supervised learning, among others. However, very little is understood about why such transfer is successful, and more importantly, how one should choose the pre-training task. As a first step towards this understanding, we ask: what makes a feature representation good for a target task? We present simple, intuitive measurements of the feature space that are good predictors of downstream task performance. We present theoretical results showing how these measurements can be used to bound the error of the downstream classifiers, and show empirically that these bounds correlate well with actual downstream performance. Finally, we show that our bounds are practically useful for choosing the right pre-trained representation for a target task.

1 INTRODUCTION

Since the (re)-discovery of neural networks for visual recognition, a plethora of work has observed that the feature representations trained on ImageNet generalize to many downstream tasks Donahue et al. (2014); Kornblith et al. (2019); Kolesnikov et al. (2020). Indeed, these representations generalize not only to different tasks in the same domain, but also to completely new domains such as satellite imagery Wallace & Hariharan (2020). This empirical observation is extremely fortuitous, as it allows one to bring the power of neural networks to new tasks and domains where very little training data is available (as in the field of few-shot learning). As such, research has expanded into *other* ways of learning representations, such as *self-supervised learning* Chen et al. (2020) or meta-learning Snell et al. (2017).

This paradigm of *representation learning* is an extremely active area of research. However, very little is understood about why certain learning techniques work and why some others don't. For instance, many meta-learning methods do not produce feature representations that transfer well across domains Guo et al. (2020). Similarly, representations produced by many self-supervised learning techniques struggle when faced with fine-grained recognition problems **?**Wallace & Hariharan (2020). Many papers have sought a deeper understanding into learned representations through empirical measurements Wang & Isola (2020). However, empirical observations are frequently limited by the limited diversity of benchmark datasets that are often heavily curated and closely related to each other (e.g., the CUB, ImageNet and iNaturalist datasets all include bird species). On the other end of the spectrum, theoretical analysis of representation learning often relies on assumptions about the target tasks which may not be true in practice Du et al. (2020); Arora et al. (2019). Because of this limited understanding, when faced with recognition problems in new domains, practitioners have limited visibility into whether any of the many learned representations will work well.

As a first step towards such an understanding, here we ask the question: *what makes a feature representation good for a task?* Unlike previous empirical observations, we propose a theoretically grounded approach that bounds classifier accuracy on a target task. Unlike previous theoretical results, we make minimal assumptions on the target task, and instead propose data-dependent bounds. We look at two kinds of bounds: bounds on *realizability* (i.e., is it even possible to get a good classifier for the task when using this feature representation?) and bounds on *generalizability* (i.e.,

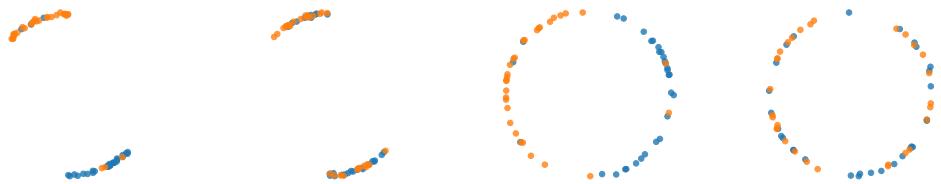

Figure 1: An illustration of the properties of local alignment and congregation. Feature representations on the left are more congregated than those on the right. The first and third feature representations on the top are more aligned locally than the other two. Classical error bounds depend only on the norm of the feature vectors and so cannot distinguish between these. We show theoretically and empirically that the distinctions shown here impact the error of a downstream classifier.

with these features, can classifiers generalize from a few examples?) While superficially these terms appear similar to what is typically studied in machine learning theory, we find that typical generalization bounds ignore the feature space entirely, and as such are poor characterizations of the feature representation. As such, we derive new kinds of bounds that emphasize properties of the feature space, and thus serve as an effective characterization.

The bounds we provide rely on two intuitive properties of the feature representation in relation to the target task (shown in Figure 1). The first property we call *local alignment*: a feature representation is *locally aligned* to the task if data points that are close in feature space have the same labels. The second property we call *congregation*: a feature representation is *congregated* if in general pairs of data points end up being close to each other. If the feature representation is locally aligned to the task, then any smooth classifier will be able to model the task well. If it is additionally congregated, then most test points will have nearby training instances, so the classifier will generalize well from limited data. These intuitions form the crux of our proposed bounds.

Theoretical bounds are useless if they do not impact practical decisions. We also present empirical results on a realistic set of tasks showing that our proposed bounds correlate well with the actual downstream performance of a classifier trained using the feature space in consideration. This correlation is good enough that it can allow us to choose the right pre-trained feature representation for a task *with very little labeled data* and *without training any classifiers for the task*.

In summary, this paper provides a theoretically grounded characterization of feature representations that is based on bounding downstream classifier performance. This characterization is general, intuitive, easy to compute, and has the added benefit of acting as *guarantees* for the downstream classifier. Finally, this characterization reveals novel directions for representation learning research.

2 Related work

Analyzing transferability: The motivations of our work lie in understanding the empirical effectiveness of transferring feature representations from supervised (Donahue et al., 2014; Zhai et al., 2019; Kolesnikov et al., 2020) or self-supervised (Goyal et al.; Wallace & Hariharan, 2020; Chen et al., 2020) tasks. With these successes on transfer, and with the availability of a large number of pretrained features trained on a wide variety of domains, there has been an increasing interest in *predicting transferability*, or in *selecting the right features* for a particular target task. The simplest approach is to train a classifier with every available pre-trained representation and pick the best performer (Zamir et al., 2018; You et al., 2021). This is both computationally expensive and requires lots of labeled data for the target task. If the pre-training task is a classification task with a known label space then the conditional distribution of targets given pre-training task labels is informative of transfer performance (Tran et al., 2019; Nguyen et al., 2020). However, this approach is difficult to apply if the pre-training task is not classification, or is inaccessible (as with models trained on proprietary data). This inaccessibility of pre-training data is also a problem for approaches that match pre-training and target distributions (Gao & Chaudhari, 2021), or explicitly adapt the pre-training images or label space (Cui et al., 2018). In contrast to these approaches, we focus on directly analyzing the pre-trained feature representation, which allows us to be agnostic to the actual task it was trained on. Our work is most similar in setup and motivation to work on measuring the alignment

between feature representations and the target labels (Huang et al., 2021; Bao et al., 2019), but in addition to allowing model selection, yields intuitive and general bounds on accuracy. Our work is orthogonal to work on characterizing tasks and measuring task similarity; this frequently requires pre-trained features in the first place (Achille et al., 2019; Wallace et al., 2021; Song et al., 2020; Dwivedi et al., 2020; Dwivedi & Roig, 2019).

Analyzing feature representations: Our work is also related to research on understanding feature representations in general. Some of this work has focused on understanding invariance properties of convnet features (Aubry & Russell, 2015). Others have looked at what individual feature dimensions mean (Agrawal et al., 2014; Zeiler & Fergus, 2014; Szegedy et al., 2013). Still others have explored if and when features from pre-trained networks transfer well between tasks as a function of the layer chosen (Yosinski et al., 2014) or the architecture (Kornblith et al., 2019). Recently these explorations have been extended to other representation learning techniques, notably self supervised techniques (Wang & Isola, 2020; Wallace & Hariharan, 2020). The insights from these explorations inspire our results. However, these explorations have been primarily empirical, and are therefore limited by the diversity of real world benchmarks that they experiment with. On the other end of the spectrum, there has been a substantial volume of work on theoretically characterizing representation learning (Frosst et al., 2019; Arora et al., 2019) and transfer learning based approaches. Theoretical analysis of transfer learning requires making assumptions about the relationship between tasks (Baxter, 2000; Maurer, 2009; Du et al., 2020). In comparison to these theoretical approaches, we assume that the target task is given; our bounds are thus *data-dependent* bounds.

3 PROBLEM SETUP

Suppose we are interested in mapping a space of inputs $\mathcal{X}$ to a space of targets $\mathcal{Y}$. There is an underlying distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$. We have a feature representation $\phi : \mathcal{X} \to \mathbb{R}^f$ which might be pre-trained on another dataset or task, which is inaccessible to us. We will assume that $\|\phi(x)\| \leq B$.

For ease of exposition in the paper, we focus on the task of binary classification; the case of multi-class classification is similar. For binary classification, we write the label space as $\mathcal{Y} = \{-1, 1\}$. Our classifier will use a scoring function that operates on feature space, $h : \mathbb{R}^f \to \mathbb{R}$. The predicted label will then be $\text{sign}(h(\phi(x)))$, where if $h(\phi(x))$ is 0, we will arbitrarily assign a label of -1. The set of possible functions h defines the hypothesis class for the classifier; denote this by $\mathcal{H}$. For most of our analysis, we primarily care about the *smoothness* of the functions in $\mathcal{H}$. We will assume that all functions in $\mathcal{H}$ are Lipschitz continuous with Lipschitz constant less than W. Thus:

$$h(\phi(x)) - h(\phi(x')) \leq W\|\phi(x) - \phi(x')\| \quad \forall h \in \mathcal{H} \tag{1}$$

We note that one such hypothesis class which is commonly used in practice is the class of *linear classifiers of bounded norm*: $\mathcal{W} = \left\{\mathbf{v} \mapsto \mathbf{w}^T\mathbf{v}; \|\mathbf{w}\| \leq W\right\}$.

Because the zero-one loss, $l^*(h(\phi(x)), y) = \mathbb{I}[\text{sign}(h(\phi(x))) \neq y]$ is difficult to analyze, we will use the following continuous upper bound which is standard in theoretical treatments (Mohri et al., 2018)

$$l(h(\phi(x)), y) = \min(1, \max(0, 1 - yh(\phi(x)))) \geq l^*(h(\phi(x)), y) \tag{2}$$

Our focus in this paper is to understand how the properties of the feature extractor ϕ affect the loss of the classifiers $l(h(\phi(x)))$. In particular, we wish to understand how ϕ impacts (a) the *lowest* average error that one can achieve, and (b) the *true error* when one generalizes from a small training set.

We begin by first identifying the key properties of feature representations.

4 TWO PROPERTIES OF FEATURE REPRESENTATIONS

What properties should we use to characterize feature representations? First, we should use properties that are easy to measure, potentially with limited labeled data. Second, these properties should be easy for human developers and practitioners to reason about. Third, they should correlate well

with the final accuracy of downstream classifiers. In sum, we want simple, intuitive measurements of the feature space that are predictive of downstream accuracy.

One kind of intuitive measurement is to look at what the feature representation considers as similar. In particular, we could look at pairs of examples that embed close to each other in feature space and ask if they are indeed similar in terms of their ground truth labels. We call this property **local alignment**. In particular, we make the following definition

Definition 1. *Suppose $\alpha > 0$. The **local alignment** of the feature space ϕ, denoted by $p_a^\phi(\alpha)$ is the probability that two data points (x, y) and (x', y') share a label given that they embed within a distance of α:*

$$p_a^\phi(\alpha) \triangleq P(y = y' \mid \|\phi(x) - \phi(x')\| \leq \alpha) \tag{3}$$

α here is a hyperparameter which governs the resolution at which we do the analysis. Note that this notion of local alignment is different from the alignment proposed by Wang & Isola (2020): the latter looks at how often two data points that are semantically similar embed close to each other, while the former looks at how often two data points that embed close to each other are semantically similar. A feature space that is locally aligned per our definition may not actually be aligned per the definition of Wang & Isola (2020) because two very similar images may be embedded far away from each other.

Local alignment alone may be meaningless if data points do not generally embed close to each other. We also need the feature space to be such that data points generally **congregate**:

Definition 2. *Suppose $\alpha > 0$. The **degree of congregation** of the feature space ϕ, denoted by p_c^ϕ, is the probability that two points sampled from $\mathcal{D}$ embed within a distance of α:*

$$p_c^\phi(\alpha) \triangleq P(\|\phi(x) - \phi(x')\| \leq \alpha) \tag{4}$$

We have α as a hyperparameter here too.

We now use these notions of *local alignment* and *congregation* to analyze the downstream error of classifiers.

5 What is the best accuracy we can achieve?

Given a feature space ϕ, what is the best accuracy we can hope to get using scoring functions from $\mathcal{H}$? Recall that these scoring functions are all W-Lipschitz. Our key intuition is that this Lipschitz continuity will force any scoring function in $\mathcal{H}$ to produce very similar scores for nearby data points. If these have different labels, then the classifier will be forced to err on at least one of them. To formalize this intuition, we begin with the following claim

Claim 1. *Consider two data points (x, y) and (x', y'). If $\|\phi(x) - \phi(x')\| < \frac{1}{W}$ and $y \neq y'$, then:*

$$l(h(\phi(x)), y) + l(h(\phi(x')), y') \geq 1 \tag{5}$$

Proof. Observe first that due to Lipschitz continuity of h:

$$|h(\phi(x)) - h(\phi(x'))| \leq W\|\phi(x) - \phi(x')\| \leq 1 \tag{6}$$

$$\Rightarrow h(\phi(x')) > h(\phi(x)) - 1 \tag{7}$$

Now, since $y \neq y'$, it follows that $y' = -y$. Without loss of generality, let us assume that $y = 1$. Denote $l(h(\phi(x)), y)$ by a and $l(h(\phi(x)), y)$ by b. We will prove this claim by contradiction. Suppose, if possible, that the clam is not true, and $a + b < 1$. Since the loss is always non-negative, it follows that $a < 1$ and $b < 1$. Thus:

$$a = \max(0, 1 - yh(\phi(x))) \tag{8}$$

$$\geq 1 - yh(\phi(x)) = 1 - h(\phi(x) \tag{9}$$

$$\Rightarrow h(\phi(x)) \geq 1 - a \tag{10}$$

Taking the other data point:

$$
\begin{aligned}
b &= \max(0, 1 - y' h(\phi(x'))) && (11)\\
&\geq 1 - y' h(\phi(x')) = 1 + h(\phi(x')) && (12)\\
&\geq 1 + h(\phi(x)) - 1 = h(\phi(x)) \qquad (\because \textit{equation } 7) && (13)\\
&\geq 1 - a && (14)\\
\Rightarrow a + b &\geq 1 && (15)
\end{aligned}
$$

thus yielding a contradiction. □

Thus, if two data points are close, at least one of them must be incorrectly classified. We can thus bound the error of the best possible classifier from below by estimating how often this happens:

Theorem 1. *Let l be the loss function defined above, and $\mathcal{H}$ be a hypothesis class of Lipschitz functions with Lipschitz constant at most W. Let $\mathcal{D}$ be a distribution over $\mathcal{X} \times \mathcal{Y}$, and p_a^ϕ and p_c^ϕ defined as above. Then*

$$\inf_{h \in \mathcal{H}} \mathbb{E}_{x,y \sim \mathcal{D}}[l(\mathbf{h}(\phi(x)), y)] > \frac{1}{2}\left(1 - p_a^\phi\left(\frac{1}{W}\right)\right) p_c^\phi\left(\frac{1}{W}\right) \tag{16}$$

Proof. We observe that if $(x, y), (x', y') \sim \mathcal{D}$, then the probability that $\|\phi(x) - \phi(x')\| < \frac{1}{W}$ and $y \neq y'$ is given by $\tilde{p} = (1 - p_a^\phi(1/W))p_c^\phi(1/W)$. Then, for all $h \in \mathcal{H}$.

$$
\begin{aligned}
\mathbb{E}_{x,y \sim \mathcal{D}}[l(h(\phi(x)), y)] &= \frac{1}{2}\mathbb{E}_{(x,y),(x',y') \sim \mathcal{D}}[l(h(\phi(x)), y) + l(h(\phi(x')), y')] && (17)\\
&\geq \frac{\tilde{p}}{2}\mathbb{E}[l(h(\phi(x)), y) + l(h(\phi(x')), y') \;\Big|\; \|\phi(x) - \phi(x')\| < \frac{1}{2W} \text{ and } y \neq y'] && (18)\\
&\geq \frac{\tilde{p}}{2} && (19)\\
& && (20)
\end{aligned}
$$

where the second step follows because the loss is non-negative, so only focusing on the case when x and x' are close but have different labels yields a lower bound. This is true for all $h \in \mathcal{H}$, so it is true for the infimum as well. □

5.1 Implications

The lower bound suggests that if we want a feature space where a good classifier can separate the data, we want to *maximize p_a^ϕ or minimize p_c^ϕ*. Importantly, we can zero out this lower bound simply by pushing p_c^ϕ to 0, which corresponds to pushing data points far away from each other. This matches what is done by contrastive learning approaches which try to spread out a dataset of unlabeled images in feature space (Chen et al., 2020) and might potentially explain the success of these approaches in producing a useful feature representation. However, as we will see below, simply pushing all examples far apart might hurt generalization.

6 How well will classifiers generalize?

There may be classifiers that yield low error, but it can be challenging to find these from small training sets. Therefore we now ask if we can use the feature space to analyze how well classifiers trained on small training sets might *generalize*.

While generalization bounds abound, they typically do not include any reasoning about the underlying feature space beyond the maximum norm of feature vectors. As such, these bounds are insufficient for understanding the impact of the feature space on generalization.

Here, we present two kinds of bounds. The first bound uses a similar proof strategy of relying on the Lipschitz continuity of the classifier. The second bound leverages a more traditional Rademacher complexity-based analysis to bound the excess risk (i.e., the difference between the test loss and training loss).

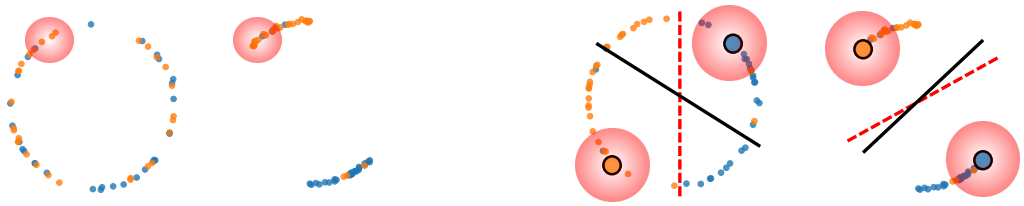

Figure 2: Intuition behind our analysis for the lower bound on error (Theorem 1, Left) and upper bound on generalization (Theorem 2, Right). In each case we look at pairs of data points that fall within a distance α of each other (shown as red balls). The lower bound notes that examples with different labels in these balls yield high error. The upper bound notes that outside the red balls centered on training points (large dots) generalization is not guaranteed. Also shown are the optimal (red) and obtained (black) classifiers.

6.1 Bounding the probability of high-loss data points

Classical generalization bounds are based primarily on concentration inequalities and the law of large numbers. This produces substantially coarse estimates for applications like *few-shot learning*, where the training dataset has very few labels.

We propose a simpler alternative bound here. We observe that owing to the Lipschitz continuity of the classifier, test data points that are close to training data points will have very similar loss as the corresponding training points, which can be no worse than the loss on the worst performing training example.

Theorem 2. *Let $\mathcal{H}_\phi = \{z = x \mapsto h(\phi(x)); h \in \mathcal{H}\}$. Suppose S is a sampled training set of m points. For any $z \in \mathcal{H}_\phi$, let $l_{max}(z, S) = \max_{(x,y)\in S} l(z(x), y)$ be the maximum loss z incurs on S. Then, for all $\epsilon > 0$ and $(x, y) \sim \mathcal{D}$:*

$$P\left(l(z(x), y) > l_{max}(z, S) + \epsilon\right) \leq \left(1 - p_a^\phi\left(\frac{\epsilon}{W}\right) p_c^\phi\left(\frac{\epsilon}{W}\right)\right)^m. \tag{21}$$

Proof. It can be shown that for any y, $l(z(x), y) - l(z(x'), y) \leq |z(x) - z(x')|$ (see Appendix). Further, $\forall h \in \mathcal{H}$, h is W-Lipschitz. So we have, for all x, x', y:

$$l(z(x), y) - l(z(x'), y) \leq |z(x) - z(x')| \leq W\|\phi(x) - \phi(x')\| \tag{22}$$

It follows that for any $(x, y), (x', y')$:

$$\left.\begin{array}{c}\|\phi(x) - \phi(x')\| < \frac{\Delta}{W} \\ y = y'\end{array}\right\} \Rightarrow l(z(x), y) \leq l(z(x'), y') + \Delta \qquad \forall z \in \mathcal{H}_\phi \tag{23}$$

Note that the probability of sampling such an (x', y') is $p_a^\phi(\Delta/W)p_c^\phi(\Delta/W)$. If (x', y') is in S, then $l(z(x'), y') \leq l_{max}(z, S)$. Thus, for any $(x, y) \sim \mathcal{D}$

$$P\left(l(z(x), y) \leq l_{max}(z, S) + \epsilon\right) \geq P\left(\exists (x', y') \in S \text{ s.t } \|\phi(x) - \phi(x')\| < \frac{\Delta}{W} \text{ and } y = y'\right) \tag{24}$$

$$= 1 - \left(1 - p_c^\phi\left(\frac{\epsilon}{W}\right) p_a^\phi\left(\frac{\epsilon}{W}\right)\right)^m \tag{25}$$

□

6.2 Bounding the excess risk

We can also perform a more traditional analysis of the excess risk, which uses the law of large numbers and various concentration inequalities. It is well known that the difference between train loss and the test loss is upper bounded by the Rademacher complexity of the hypothesis class plus

a small constant (Mohri et al., 2018). Concretely, with probability greater than $1-\delta$, the following holds for all $h \in \mathcal{H}$:

$$R(h) - \hat{R}(h, S) \leq \mathcal{R}_m(\mathcal{H}) + \sqrt{\frac{\log \frac{1}{\delta}}{2m}} \tag{26}$$

where $R(h)$ is the test loss (or expected risk), $\hat{R}(h, S)$ is the training loss (or empirical risk), and $\mathcal{R}_m(\mathcal{H})$ is the *Rademacher complexity* of $\mathcal{H}$.

Thus, to bound the generalization error, it suffices to bound the Rademacher complexity of $\mathcal{H}$, which is defined as $\mathcal{R}_m(\mathcal{H}) = \mathbb{E}_{S\sim\mathcal{D}^m,\sigma} \sup_{h\in\mathcal{H}} \frac{1}{m}\sum_{i=1}^{m} h(x_i)$. Here σ_i are Rademacher random variables (i.e., $\sigma_i \in \{-1, 1\}, \mathbb{E}\sigma_i = 0$). The Rademacher complexity is a measure of the size of the hypothesis class, and can be used to bound the difference between a hypothesis' performance on two sampled datasets. To allow us to define a concrete bound, we will analyze the class of *linear classifiers with bounded norm*, $\mathcal{W}$.

The key insight here is that if, for a particular α, $p_c^\phi(\alpha)$ is large, then any two datasets (e.g., a train and a test dataset) that we sample from the underlying distribution will be "similar", in terms of having nearby data points. As such, linear classifiers will be forced to produce similar scores for both datasets, yielding similar loss. This insight is expressed as the following bound on the Rademacher complexity of the set of linear classifiers composed with ϕ:

Theorem 3. *Let $\mathcal{W}_\phi = \{z = x \mapsto h(\phi(x)); h \in \mathcal{W}\}$. Let p_c^ϕ be defined as above. Then the Rademacher complexity of $\mathcal{W}_\phi$ is bounded above by:*

$$\mathcal{R}_m(\mathcal{W}_\phi) \leq \frac{W\left(\alpha\sqrt{p_c^\phi(\alpha)/2} + 2B\sqrt{(1-p_c^\phi(\alpha)) + 2(1-p_c^\phi(\alpha))^2}\right)}{2\sqrt{m}} \tag{27}$$

Proof sketch: (Full proof presented in Appendix). For any $x_i \in S$, there is a probability $p_c^\phi(\alpha)/2$ that the next point x_{i+1} is (a) within a distance α in feature space, and (b) is multiplied with a Rademacher variable of the opposite sign. Such pairs will effectively cancel out each other except for a small quantity of α, resulting in the first term. The second term is obtained from all the rest of the data points in S, which have to be analyzed in the usual way.

6.3 Implications

Importance of congregation: Both bounds involve the degree of congregation, p_c^ϕ, and lead to lower values if p_c^ϕ is high for low values of α. Thus, for generalization from a few examples, we wish the feature representation to at least push some examples closer together.

Importance of local alignment: Here the two bounds diverge. The second Rademacher complexity bound will remain low even if data points with different labels are pushed closer together (so p_a^ϕ is low). This however will generally make even the training error go up (although the generalization error may go down). In contrast the first bound actually requires that p_a^ϕ should also be high. The first bound thus suggests that for generalization, we need to solve the hard problem of pushing together examples with the same labels and pushing apart those with different labels.

Dependence on the Lipschitz constant: Another difference between the two bounds is their dependence on the Lipschitz constant W. In general, the classical Rademacher complexity-based bound, and the second bound above, will decrease if W is decreased. However, this is not true for the first bound proposed above, since as W is decreased, Δ/W will increase and $p_a^\phi(\Delta/W)$ might in fact start to fall. This in turn suggests the prospect of there being a sweet spot; a particular level of regularization of the classifier that is optimal, and that depends on where both p_a^ϕ and p_c^ϕ are high.

7 Empirical analysis

Above we have shown that p_a^ϕ and p_c^ϕ yield bounds on classifier performance. But are they useful enough in practice to make decisions? We first look to see whether they are correlated with actual loss values, and then present potential applications.

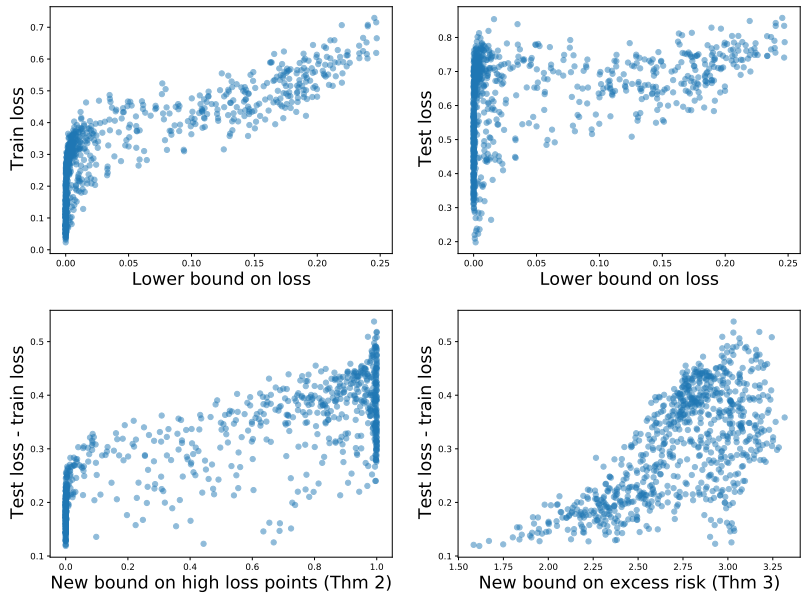

Figure 3: Correlations of our proposed bounds with actual performance. Top: Lower bound vs train loss (left) and test loss (right). Bottom: Upper bounds from Theorem 2(left) and Theorem 3 (right) vs actual excess risk.

7.1 Correlations with classifier accuracy

Experimental setup: We used CIFAR-10(Krizhevsky, 2009) as our test bed. We sampled pairs of classes from CIFAR-10 as "recognition tasks", resulting in 45 different tasks. For each task we ran 10 different trials. In each trial, we sampled a small training set to train a linear classifier, and then evaluated the classifier on a large validation set. The training and validation accuracies were averaged over the 10 trials. We used 18 different representations, all with the architecture of a ResNet-18 that were each trained on ImageNet(Russakovsky et al., 2015), iNaturalist(Van Horn et al., 2018), or CIFAR-100 using various loss functions (see Appendix). The 45 tasks and 18 representations together yielded 810 sets of measurements.

For each case, we estimated p_a^ϕ and p_c^ϕ using all pairs of data points from the validation set with corresponding ground truth labels, and used these to estimate our lower and upper bounds. We also estimated the classical textbook upper bound on excess risk.

Results: We find that the lower bound is extremely well correlated with the training loss (Spearman's ρ 0.90, Figure 3, top), suggesting that it can be used to estimate the training loss without any training. Its correlation with the test loss is much weaker: there are many models with high test loss but a zero lower bound. This is of course expected because the lower bound does not take into account the perils of classifiers trained on small training sets. The difference between the training and validation losses are correlated very well with our two upper bounds on generalization (Spearman's ρ for Thm 2 0.80, Spearman's ρ for Thm 3 0.61, Figure 3 bottom), much more so than with the classical Rademacher complexity-based generalization bound (Spearman's ρ 0.56, shown in appendix). Since Theorem 2 yields a stronger correlation, we use this upper bound in conjunction with the lower bound for a more practical experiment on model selection below.

7.2 Applications to few-shot transfer

We now see if our bounds can help choose the right feature representation for problems with very limited labeled data. In particular, we assume that we have a very small labeled set S (e.g., 10 labeled examples per class) and a larger unlabeled set U . Because labeled data is very limited, one cannot afford to use labeled data for held-out validation sets. Therefore, a way of characterizing feature representations without training and evaluating classifiers is valuable.

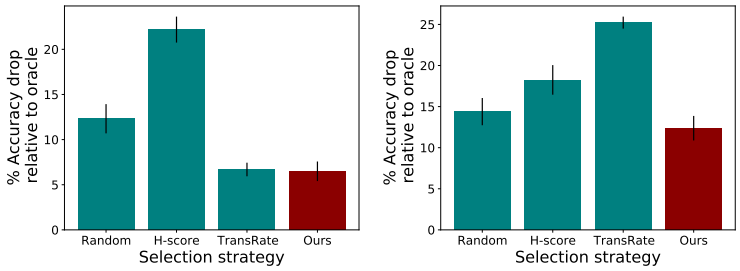

Figure 4: Our characterization of the feature space allows us to choose the best pre-trained representation for a given few-shot task. Lower numbers are better. Left: results on CUB-200. Right: results on CIFAR-10.

Baselines: Much of prior work on estimating transferability assumes that the representations come from classification tasks Nguyen et al. (2020) on accessible datasets Cui et al. (2018), or that an initial appropriate feature representation is available Achille et al. (2019); Dwivedi et al. (2020). With the advent of self-supervised learning Chen et al. (2020) and large proprietary datasets Kolesnikov et al. (2020), and the wide variety of visual domains Wallace & Hariharan (2020), these assumptions are no longer appropriate. We therefore consider only baselines that make minimal assumptions about the pre-trained feature representation and the target task: H-score Bao et al. (2019) and TransRate Huang et al. (2021). Both approaches are based on information-theoretic arguments justifying the use of intra- and inter-class covariance matrices, which are global statistics of the feature representation, in contrast to our more local measurements. In addition, we also include a baseline that chooses a representation at random.

Experimental setup: We run few-shot experiments on two datasets: CIFAR-10 and CUB-200(Welinder et al., 2010). The pre-trained representations were obtained by training on iNaturalist, ImageNet and CIFAR-10 (only for CUB-200) with various loss functions. In each case, we sample 50 binary classification problems with 10 labeled examples each; all other examples from the classes involved are unlabeled. To estimate $p_a^{\phi}(\alpha)$ using just 10 labeled examples, we choose α to be large enough that for at least one class, every pair of labeled examples is at most α away in feature space. Our estimate of transferability involves linearly combining our lower bound on error (with a weight of 100 because of its smaller scale) and the upper bound on generalization from Theorem 2; lower values indicate better representations. For the baselines, the unconditional covariance matrix was estimated using the unlabeled data and other label-dependent measurements used the labeled data. We evaluate all approaches in terms of the average accuracy drop relative to the oracle representation (the one which post-hoc yields the highest test accuracy). Thus lower drops are better.

Results: We show results in figure 4. Interestingly, TransRate and H-score seem to have opposite trends on these two datasets, with H-score performing worse than random selection for CUB and TransRate performing worse than random on CIFAR-10 Our approach consistently yields better choices than both baselines on both benchmark datasets thus suggesting the usefulness of our characterization for practical problems.

8 Conclusion

We have presented a new characterization of feature representations that yields bounds on downstream accuracy, and in practice can be used to choose better representations for a given target task. We hope that these results will also illuminate new directions for directly learning representations that minimize these bounds.

Reproducibility: All theorems and proofs are included in the main paper, with the exception of one claim and one theorem proof (proof of Theorem 3) which are presented in the Appendix. Experimental details regarding the representations used and the data pre-processing are also presented in the appendix. The code, pre-trained representations, and the raw data from the experiments will be released upon acceptance.

Ethics statement: This work aims to improve our understanding of feature representations, thus paving the way for more widespread use of pre-trained feature representations. We hope that this makes the power of sophisticated visual recognition techniques available to a much broader community who may not have the data or resources to train big networks from scratch. However, we note that our work focuses only on the *average accuracy* one can obtain from these representations in a downstream task; we do not delve into the biases that these representations might perpetuate Steed & Caliskan (2021), or privacy or consent issues in the original datasets used to train these representations Birhane & Prabhu (2021). We suggest that when choosing a pre-trained representation, practitioners should consider not just our characterization of downstream accuracy, but also a broader characterization in terms of the many ethical implications of choosing a representation pre-trained on questionably collected or opaque datasets.

REFERENCES

Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charless C Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6430–6439, 2019.

Pulkit Agrawal, Ross Girshick, and Jitendra Malik. Analyzing the performance of multilayer neural networks for object recognition. In *European conference on computer vision*, pp. 329–344. Springer, 2014.

Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. In *36th International Conference on Machine Learning, ICML 2019*, pp. 9904–9923. International Machine Learning Society (IMLS), 2019.

Mathieu Aubry and Bryan C Russell. Understanding deep features with computer-generated imagery. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2875–2883, 2015.

Yajie Bao, Yang Li, Shao-Lun Huang, Lin Zhang, Lizhong Zheng, Amir Zamir, and Leonidas Guibas. An information-theoretic approach to transferability in task transfer learning. In *2019 IEEE International Conference on Image Processing (ICIP)*, pp. 2309–2313. IEEE, 2019.

Jonathan Baxter. A model of inductive bias learning. *Journal of artificial intelligence research*, 12: 149–198, 2000.

Abeba Birhane and Vinay Uday Prabhu. Large image datasets: A pyrrhic win for computer vision? In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1536–1546. IEEE, 2021.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.

Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4109–4118, 2018.

Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pp. 647–655. PMLR, 2014.

Simon Shaolei Du, Wei Hu, Sham M Kakade, Jason D Lee, and Qi Lei. Few-shot learning via learning the representation, provably. In *International Conference on Learning Representations*, 2020.

Kshitij Dwivedi and Gemma Roig. Representation similarity analysis for efficient task taxonomy & transfer learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12387–12396, 2019.

Kshitij Dwivedi, Jiahui Huang, Radoslaw Martin Cichy, and Gemma Roig. Duality diagram similarity: a generic framework for initialization selection in task transfer learning. In *European Conference on Computer Vision*, pp. 497–513. Springer, 2020.

Nicholas Frosst, Nicolas Papernot, and Geoffrey Hinton. Analyzing and improving representations with the soft nearest neighbor loss. In *International Conference on Machine Learning*, pp. 2012–2020. PMLR, 2019.

Yansong Gao and Pratik Chaudhari. An information-geometric distance on the space of tasks. In *International Conference on Machine Learning*, pp. 3553–3563. PMLR, 2021.

Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6390–6399. IEEE.

Yunhui Guo, Noel C Codella, Leonid Karlinsky, James V Codella, John R Smith, Kate Saenko, Tajana Rosing, and Rogerio Feris. A broader study of cross-domain few-shot learning. In *European Conference on Computer Vision*, pp. 124–141. Springer, 2020.

Long-Kai Huang, Ying Wei, Yu Rong, Qiang Yang, and Junzhou Huang. Frustratingly easy transferability estimation. *arXiv preprint arXiv:2106.09362*, 2021.

Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pp. 491–507. Springer, 2020.

Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2661–2671, 2019.

Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.

Andreas Maurer. Transfer bounds for linear feature learning. *Machine learning*, 75(3):327–350, 2009.

Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.

Cuong Nguyen, Tal Hassner, Matthias Seeger, and Cedric Archambeau. Leep: A new measure to evaluate transferability of learned representations. In *International Conference on Machine Learning*, pp. 7294–7305. PMLR, 2020.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.

Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 4080–4090, 2017.

Jie Song, Yixin Chen, Jingwen Ye, Xinchao Wang, Chengchao Shen, Feng Mao, and Mingli Song. Depara: Deep attribution graph for deep knowledge transferability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3922–3930, 2020.

Ryan Steed and Aylin Caliskan. Image representations learned with unsupervised pre-training contain human-like biases. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 701–713, 2021.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

Anh T Tran, Cuong V Nguyen, and Tal Hassner. Transferability and hardness of supervised classification tasks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1395–1405, 2019.

Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8769–8778, 2018.

Bram Wallace and Bharath Hariharan. Extending and analyzing self-supervised learning across domains. In *ECCV*, 2020.

Bram Wallace, Ziyang Wu, and Bharath Hariharan. Can we characterize tasks without labels or features? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1245–1254, 2021.

Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pp. 9929–9939. PMLR, 2020.

P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010.

Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *arXiv preprint arXiv:1411.1792*, 2014.

Kaichao You, Yong Liu, Jianmin Wang, and Mingsheng Long. Logme: Practical assessment of pre-trained models for transfer learning. In *International Conference on Machine Learning*, pp. 12133–12143. PMLR, 2021.

Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3712–3722, 2018.

Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pp. 818–833. Springer, 2014.

Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruyssen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, et al. The visual task adaptation benchmark. 2019.

A Appendix

A.1 Proofs

Claim 2. *(Used in proof of Theorem 2) Let* $l(z, y) = \min(1, \max(0, 1 - zy))$. *Then* $l(z, y) - l(z', y) \leq |z - z'|$.

Proof. Write $f(x) = \min(1, \max(0, x))$. Then $l(z, y) = f(1 - yz)$. To prove the claim, it suffices to show that $f(x) - f(x') \leq |x - x'|$. We do a case-by-case analysis.

1. $x \geq 0, x' \leq 1$. In this case, $f(x') \geq x'$ and $f(x) \leq x$. So $f(x) - f(x') \leq x - x' \leq |x - x'|$.
2. $x < 0$. In this case $f(x) = 0$. Since $f(x') \geq 0$, it follows that $f(x) - f(x') < 0 < |x - x'|$.
3. $x' > 1$. In this case $f(x') = 1$. Since $f(x) \leq 1$ it follows that $f(x) - f(x') < 0 < |x - x'|$.

□

Theorem 4. *Let $\mathcal{W}_\phi = \{z = x \mapsto h(\phi(x)); h \in \mathcal{W}\}$. Let p_c^ϕ be defined as above. Then the Rademacher complexity of $\mathcal{W}_\phi$ is bounded above by:*

$$\mathcal{R}_m(\mathcal{W}_\phi) \leq \frac{W\left(\alpha\sqrt{p_c^\phi(\alpha)/2} + 2B\sqrt{(1-p_c^\phi(\alpha)) + 2(1-p_c^\phi(\alpha))^2}\right)}{2\sqrt{m}} \tag{28}$$

Proof. First note that:

$$\begin{aligned}
\mathcal{R}_m(\mathcal{W}_\phi) &= \mathbb{E}_{S,\sigma} \sup_{z\in\mathcal{W}_\phi} \frac{1}{m}\sum_{i=1}^m \sigma_i z(x_i) && (29)\\
&= \frac{1}{m}\mathbb{E}_{S,\sigma} \sup_{h\in\mathcal{W}} \sum_{i=1}^m \sigma_i \mathbf{w}_h^T \phi(x_i) && (30)\\
&= \frac{1}{m}\mathbb{E}_{S,\sigma} \sup_{h\in\mathcal{W}} \mathbf{w}_h^T (\sum_{i=1}^m \sigma_i phi(x_i)) && (31)\\
&\leq \frac{W}{m}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^m \sigma_i\phi(x_i)\| \qquad (\because \|\mathbf{w}_h\| \leq W) && (32)\\
& && (33)
\end{aligned}$$

Thus we need to bound $\mathbb{E}_{S,\sigma}\|\sum_{i=1}^m \sigma_i\phi(x_i)\|$

Consider any set S of m points sampled from $\mathcal{D}$. Let π be the mapping $i \mapsto (i+1) \mod m$, and define $\sigma_i' = \sigma_{\pi(i)}$ and $x_i' = x_\pi(i)$. Then, for any $z \in \mathcal{H}_\phi$ we have that $\sum_{i=1}^m \sigma_i\phi(x_i) = \sum_{i=1}^m \sigma_i'\phi(x_i') = \frac{1}{2}\sum_{i=1}^m(\sigma_i\phi(x_i) + \sigma_i'\phi(x_i'))$.

Note that $x_i \perp\!\!\!\perp x_i'$ and $\sigma_i \perp\!\!\!\perp \sigma_i'$. Define $a_i = \mathbb{I}[(\sigma_i = -\sigma_i')$ and $(\|\phi(x_i) - \phi(x_i')\| < \alpha)]$. Then $\mathbb{E}_{S,\sigma}a_i = \mathbb{E}_{S,\sigma}a_i^2 = \frac{p_c^\phi(\alpha)}{2}$, which we will denote as p_α.

Therefore:

$$\begin{aligned}
&\mathbb{E}_{S,\sigma}\|\sum_{i=1}^m \sigma_i\phi(x_i)\| && (34)\\
&= \frac{1}{2}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^m(\sigma_i\phi(x_i) + \sigma_i'\phi(x_i'))\| && (35)\\
&= \frac{1}{2}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^m(a_i\sigma_i(\phi(x_i) - \phi(x_i')) + \sum_{i=1}^m \bar{a}_i(\sigma_i\phi(x_i) + \sigma_i'\phi(x_i')))\| && (36)\\
&\leq \frac{1}{2}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^m a_i\sigma_i(\phi(x_i) - \phi(x_i'))\|\\
&\quad + \frac{1}{2}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^m \bar{a}_i\sigma_i\phi(x_i)\|\\
&\quad + \frac{1}{2}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^m \bar{a}_i\sigma_i'\phi(x_i')\| && (37)
\end{aligned}$$

Observe that $a_i \perp\!\!\!\perp \sigma_i$, and $a_i \perp\!\!\!\perp \sigma_i'$. Using this, the second and third terms are similar. The second term yields:

$$\frac{1}{2}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^{m}\bar{a}_i\sigma_i\phi(x_i)\| \tag{38}$$

$$=\frac{1}{2}\mathbb{E}_{S,\sigma}\sqrt{\sum_{i,j=1}^{m}\bar{a}_i\bar{a}_j\sigma_i\sigma_j\phi(x_i)^T\phi(x_j)} \tag{39}$$

$$\leq\frac{1}{2}\sqrt{\sum_{i,j=1}^{m}\mathbb{E}_{S,\sigma}\bar{a}_i\bar{a}_j\sigma_i\sigma_j\phi(x_i)^T\phi(x_j)} \tag{40}$$

We note that if $i \neq j$ and $i \neq \pi(j)$ and $j \neq \pi(i)$ then $\bar{a}_i\sigma_i \perp\!\!\!\perp \bar{a}_j\sigma_j$. Thus, looking at the expression under the square root:

$$\sum_{i,j=1}^{m}\mathbb{E}_{S,\sigma}\bar{a}_i\bar{a}_j\sigma_i\sigma_j\phi(x_i)^T\phi(x_j) \tag{41}$$

$$\leq\sum_{i}\mathbb{E}_{S,\sigma}\bar{a}_i^2\sigma_i^2\|\phi(x_i)\|^2 \tag{42}$$

$$+\sum_{i,j:i=\pi(j)\text{or}j=\pi(i)}\mathbb{E}_{S,\sigma}\bar{a}_i\bar{a}_j\sigma_i\sigma_j\phi(x_i)^T\phi(x_j) \tag{43}$$

$$\leq B^2\sum_{i}\mathbb{E}_{S,\sigma}\bar{a}_i^2\sigma_i^2 \qquad (\because \mathbb{E}\bar{a}_i^2=\mathbb{E}\bar{a}_i=1-p_\alpha) \tag{44}$$

$$+B^2\sum_{i,j:i=\pi(j)\text{or}j=\pi(i)}\mathbb{E}_{S,\sigma}\bar{a}_i\bar{a}_j\sigma_i\sigma_j \tag{45}$$

$$\leq B^2m(1-p_\alpha)+B^2\sum_{i,j:i=\pi(j)\text{or}j=\pi(i)}\mathbb{E}_{S,\sigma}\bar{a}_i\bar{a}_j \qquad (\because \mathbb{E}\sigma_i\leq 1) \tag{46}$$

$$\leq B^2m(1-p_\alpha)+B^2 2m(1-p_\alpha)^2 \tag{47}$$

. Thus:

$$\frac{1}{2}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^{m}\bar{a}_i\sigma_i\phi(x_i)\| \tag{48}$$

$$\leq\frac{B}{2}\sqrt{m(1-p_\alpha)+2m(1-p_\alpha)^2} \tag{49}$$

Through similar reasoning, the first term yields:

$$\frac{1}{2}\mathbb{E}_{S,\sigma}\|\sum_{i=1}^{m}a_i\sigma_i(\phi(x_i)-\phi(x_i')\| \tag{50}$$

$$\leq\frac{1}{2}\sqrt{\sum_{i=1}^{m}\mathbb{E}_{S,\sigma}a_i^2\sigma^2\|\phi(x_i)-\phi(x_i')\|^2} \tag{51}$$

$$\leq\frac{1}{2}\sqrt{\sum_{i=1}^{m}\mathbb{E}_{S,\sigma}a_i^2\alpha^2} \tag{52}$$

$$=\frac{1}{2}\sqrt{mp_\alpha\alpha^2} \qquad (\because \mathbb{E}a_i^2=\mathbb{E}a_i=p_\alpha) \tag{53}$$

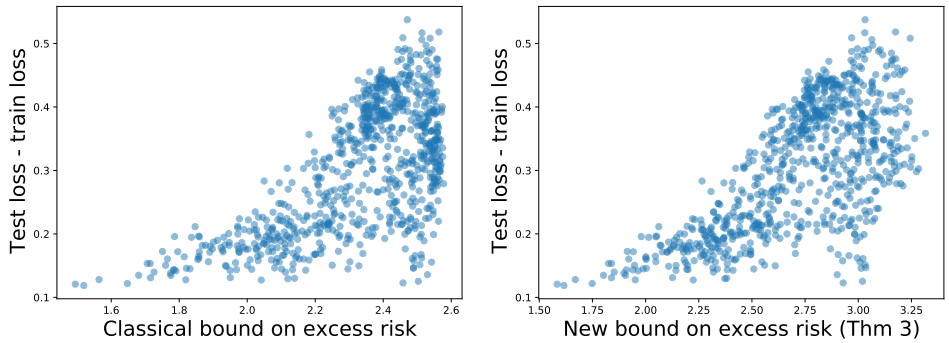

Figure 5: Comparison between our proposed Rademacher complexity-based bound (Theorem 3, right) and the classical textbook bound (left). Our proposed bound is better correlated with the generalization error, especially for high values of the generalization error.

Putting these expressions back in equation 37, we get:

$$\mathcal{R}_m(\mathcal{H}_\phi) \leq \frac{W}{2m}(\sqrt{mp_\alpha\alpha^2} + B\sqrt{m(1-p_\alpha)+2m(1-p_\alpha)^2} + B\sqrt{m(1-p_\alpha)+2m(1-p_\alpha)^2}) \tag{54}$$

$$= \frac{W\left(\alpha\sqrt{p_\alpha} + 2B\sqrt{m(1-p_\alpha)+2m(1-p_\alpha)^2}\right)}{2\sqrt{m}} \tag{55}$$

□

A.2 Our proposed generalization bound vs the classical bound

See Figure 5.

A.3 Representations used in experiments

In general it is difficult to do statistically rigorous experiments for choosing representations since publicly available pre-trained representations vary wildly in architecture and dimensionality. However, we want to ensure that the characterization we produce is nuanced enough to distinguish between fairly similar representations that are based on the same backbone architecture and the same feature dimensionality. Therefore we used multiple feature representations that are all trained on ResNet 18 on the same 3 datasets (ImageNet, CIFAR-100 and iNaturalist) but with very different loss functions. These representations were trained in the context of an unrelated exploration on the use of coarsely labeled data (see supplementary). For completeness, we present a brief overview of these representations here.

All representations were trained on a combination of a coarsely labeled dataset C and a finely labeled dataset F. C may have some classes that are not present in F. The representations we used are:

1. A representation trained with simple supervised learning on F alone.
2. A representation trained with simple supervised learning on C alone.
3. A representation trained on $C \cup F$ with two classifier heads, one for coarse categorization and another for fine-grained categorization
4. A representation trained on $C \cup F$ with *fine-grained soft pseudo-labels* for C obtained using a classifier trained on F
5. A representation trained on $C \cup F$ with *fine-grained soft pseudo-labels* for C obtained using a classifier trained on F, and with the pseudo-labels filtered to be consistent with the coarse label,
6. A representation trained on $C \cup F$ with ground-truth fine-grained. labels for images in C.

6 representations trained on 3 datasets gives us 18 representations. Details on the training are in the paper linked in the supplementary. **These pre-trained representations will be released upon acceptance**.