
The Tree Loss: Improving Generalization with Many Classes

Yujie Wang

Claremont Graduate University

Mike Izbicki

Claremont McKenna College

Abstract

Multi-class classification problems often have many semantically similar classes. For example, 90 of ImageNet’s 1000 classes are for different breeds of dog. We should expect that these semantically similar classes will have similar parameter vectors, but the standard cross entropy loss does not enforce this constraint.

We introduce the tree loss as a drop-in replacement for the cross entropy loss. The tree loss re-parameterizes the parameter matrix in order to guarantee that semantically similar classes will have similar parameter vectors. Using simple properties of stochastic gradient descent, we show that the tree loss’s generalization error is asymptotically better than the cross entropy loss’s. We then validate these theoretical results on synthetic data, image data (CIFAR100, ImageNet), and text data (Twitter).

1 Introduction

The cross entropy loss is the most widely used loss function for multi-class classification problems, and stochastic gradient descent (SGD) is the most common algorithm for optimizing this loss. Standard results show that the generalization error decays at a rate of $O(\sqrt{k/n})$, where k is the number of class labels, and n is the number of data points. These results (which we review later) make no assumptions about the underlying distribution of classes.

In this paper, we assume that the class labels have an underlying metric structure, and we introduce the tree loss for exploiting this structure. The doubling dimension c of a metric space is a common measure

of the complexity of the metric, and we show that SGD applied to the tree loss will converge at a rate of $O(\sqrt{\log k/n})$ when $c \leq 1$ or $O(\sqrt{k^{1-1/c}/n})$ when $c \geq 1$. These improvements are obviously most dramatic for small c , and the tree loss outperforms the cross entropy loss best in this regime. The tree loss is the first multi-class loss function with provably better generalization error than the cross entropy loss in any regime.

Our paper is organized as follows. We begin in Section 2 by discussing limitations of related loss functions and how the tree loss addresses those limitations. Section 3 then formally defines the problem setting, and Section 4 formally defines the tree loss. We emphasize that the tree loss is simply a reparameterization of the standard cross entropy loss, and so it is easy to implement in common machine learning libraries. Section 5 reviews standard results on the convergence of stochastic gradient descent, and uses those results to prove the convergence bounds for the tree loss. Finally, Section 6 conducts experiments on synthetic, real world image (CIFAR100, ImageNet) and text (Twitter) data. We show that in practice, the tree loss essentially always outperforms other multi-class loss functions.

2 Related Work

Previous work on multi-class loss functions has focused on improving either the statistical or computational performance. Statistical work includes loss functions designed for hierarchichally structured labels (Cesa-Bianchi et al., 2006; Wu et al., 2019; Bertinetto et al., 2020), loss functions for improving top- k accuracy (Lapin et al., 2016), or focusing on noisy labels (Sukhbaatar et al., 2015; Zhang and Sabuncu, 2018). The work most similar to ours is the SimLoss (Kobs et al., 2020), which also makes a metric assumption about the class label structure. Our work improves on all of this statistical work by being the first to provide convergence bounds that improve on the bounds of the cross entropy loss.

Other work focuses on improving the speed of multi-class classification. The hierarchical softmax (Morin

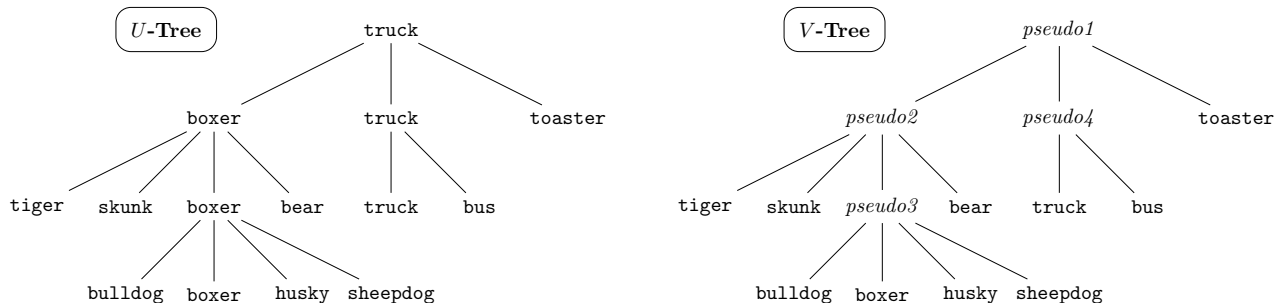


Figure 1: Example label tree structures for a subset of 10 ImageNet classes. The U -tree has class labels that repeat at multiple levels, and the V -tree introduces “pseudoclasses”. The pseudoclass $pseudo3$ represents the class of “dogs”, and the pseudoclass $pseudo2$ represents the class of “animals”.

and Bengio, 2005) is a particularly well studied modification of the cross entropy loss with many variations (e.g. Peng et al., 2017; Jiang et al., 2017; Yang et al., 2017; Mohammed and Umaashankar, 2018). It is easily confused with our tree loss because both loss functions involve a tree structure. The difference, however, is that the hierarchical softmax focuses on improving runtime performance, and most variants actually sacrifice statistical performance to do so. The tree loss, in contrast, maintains the runtime performance of the standard cross entropy loss but improves the statistical performance.

3 Problem Setting

We consider the multi-class classification setting with k classes and d input feature dimensions. The cross entropy loss is the standard loss function for this setting. It is defined to be

$$\ell(W; (\mathbf{x}, y)) = -\log \frac{\exp(-\mathbf{w}_y^T \mathbf{x})}{\sum_{j=1}^k \exp(-\mathbf{w}_j^T \mathbf{x})} \quad (1)$$

where for each class $i \in [k]$, $\mathbf{w}_i \in \mathbb{R}^d$ is the parameter vector associated with class i ; the variable $W \in \mathbb{R}^{k \times d} = (\mathbf{w}_1; \mathbf{w}_2; \dots; \mathbf{w}_k)$ is the full parameter matrix; $\mathbf{x} \in \mathbb{R}^d$ is the input feature vector; and $y \in [k]$ is the input class label.

The cross entropy loss has no constraints on the weight matrix that cause similar classes to have similar parameter vectors. The tree loss adds these constraints, resulting in faster convergence and better generalization.

4 The Tree Loss

The main idea of the tree loss is that similar classes should be forced to “share” some of their parameters. The tree loss refactors the cross entropy loss’s weight matrix in order to enforce this sharing.

We propose two variants of the tree loss which enforce this sharing in different ways. We begin by introducing the U -tree loss, which is simpler to explain and analyze. Then, we introduce the V -tree loss, which improves on the U -tree loss and is the loss we suggest using in practice. Whenever we refer to the tree loss without a qualifier, we mean the V -tree loss.

4.1 The U -Tree Loss

Explaining our new U -tree parameterization requires introducing some notation.

We define a U -tree over the class labels to be a tree where each leaf node is represented by a label, and each internal node shares a label with a child node. Figure 1 (left) shows an example U -tree structure. The definition of a U -tree is motivated by the cover tree data structure (Beygelzimer et al., 2006; Izbicki and Shelton, 2015), which generates U -tree structures given a metric over the class labels.

For each class i , we let the sequence P_i denote the path from the leaf node to the root with duplicates removed. For example, using the U -tree in Figure 1, we have

$$\begin{aligned} P_{\text{bear}} &= (\text{bear}, \text{boxer}, \text{truck}) \\ P_{\text{sheepdog}} &= (\text{sheepdog}, \text{boxer}, \text{truck}) \\ P_{\text{truck}} &= (\text{truck}) \end{aligned}$$

For each class i , we define $\text{parent}(i)$ to be the first node in P_i not equal to i . In other words, $\text{parent}(i)$ is the parent of the upper-most node for class i .

We are now ready to introduce the U -tree parameterization. We associate for each class i a new parameter vector \mathbf{u}_i defined to be

$$\mathbf{u}_i = \begin{cases} \mathbf{w}_i & \text{if } i \text{ is the root} \\ \mathbf{w}_{\text{parent}(i)} - \mathbf{w}_i & \text{otherwise} \end{cases}, \quad (2)$$

and we define the parameter matrix $U \in \mathbb{R}^{k \times d}$ to be $(\mathbf{u}_1; \mathbf{u}_2; \dots; \mathbf{u}_k)$. We can recursively rewrite the original

parameter vectors in terms of this new parameterization as

$$\mathbf{w}_i = \sum_{j \in P_i} \mathbf{u}_j. \quad (3)$$

The U -tree loss is then defined by substituting (3) into (1) to get

$$\ell(U; (\mathbf{x}, y)) = -\log \frac{\exp(-\sum_{j \in P_y} \mathbf{u}_j^T \mathbf{x})}{\sum_{i=1}^k \exp(-\sum_{j \in P_i} \mathbf{u}_j^T \mathbf{x})}$$

We use the same ℓ notation for the standard cross entropy loss function and the U -Tree loss because the function is exactly the same; the only difference is how we represent the parameters during the training procedure.

4.2 The V -Tree Loss

The V -tree is constructed from the U -tree by replacing all non-leaf nodes with “pseudoclasses”. Let k' denote the total number of classes plus the total number of pseudoclasses.

We let P_i be defined as the path from leaf node i to the root as before, but note that there are now no duplicates to remove and so the paths are longer. For example,

$$\begin{aligned} P_{\text{bear}} &= (\text{bear}, \text{pseudo2}, \text{pseudo1}) \\ P_{\text{sheepdog}} &= (\text{sheepdog}, \text{pseudo3}, \text{pseudo2}, \text{pseudo1}) \\ P_{\text{truck}} &= (\text{truck}, \text{pseudo4}, \text{pseudo1}) \end{aligned}$$

The \mathbf{v}_i and V parameters are now defined analogously to the \mathbf{u}_i and U parameters, but over the V -tree structure instead of the U -tree structure. An important distinction between the V and U parameter matrices is that they will have different shapes. The U matrix has shape $k \times d$, which is the same as the W matrix, but the V matrix has shape $k' \times d$, which is potentially a factor of up to 2 times larger.

The V -tree loss is defined by using the V matrix to parameterize the cross entropy loss, giving

$$\ell(V; (\mathbf{x}, y)) = -\log \frac{\exp(-\sum_{k \in P_y} \mathbf{v}_k^T \mathbf{x})}{\sum_{j=1}^{k'} \exp(-\sum_{k \in P_j} \mathbf{v}_k^T \mathbf{x})}. \quad (4)$$

We continue to use the ℓ function to represent both the standard cross entropy loss and the V -tree loss function to emphasize that these are the same loss functions, just with different parameterizations of the weight matrix.

4.3 Intuition

The intuition behind the V -tree reparameterization is that whenever we perform a gradient update on a data

point with class i , we will also be “automatically” updating the parameter vectors of similar classes. To see this, note that when we have a `sheepdog` data point, we will perform gradient updates on all \mathbf{v}_i in P_{sheepdog} ; i.e. $\mathbf{v}_{\text{sheepdog}}$, $\mathbf{v}_{\text{pseudo3}}$, $\mathbf{v}_{\text{pseudo2}}$, and $\mathbf{v}_{\text{pseudo1}}$. This will cause the $\mathbf{w}_{\text{husky}}$ and \mathbf{w}_{bear} parameters (among many others) to update because they also depend on the pseudoclass vectors $\mathbf{v}_{\text{pseudo3}}$ and $\mathbf{v}_{\text{pseudo2}}$. Because P_{husky} has a larger overlap with P_{sheepdog} than P_{bear} , the parameters of this more similar class will be updated more heavily.

This reparameterization is reminiscent of the way that `fastText` (Bojanowski et al., 2017) reparameterizes the `word2vec` (Mikolov et al., 2013) model to improve statistical efficiency. Two notable differences, however, are that `fastText` is a domain-specific adaptation and provides no theoretical guarantees; our tree loss works on any domain and provides theoretical guarantees.

4.4 Implementation Notes

Both the U -tree and V -tree losses are easy to implement in practice using the built-in cross entropy loss function in libraries like PyTorch (Paszke et al., 2019) or Tensorflow (Abadi et al., 2015). The only change needed is to represent the W parameter matrix in your code as an appropriate sum over the U or V parameter matrices. The magic of automatic differentiation will then take care of the rest and ensure that the underlying U or V matrix is optimized. Our `TreeLoss` library¹ provides easy-to-use functions for generating these matrices, and so modifying code to work with the tree loss is a 1 line code change.

In practice, the tree loss is slightly slower than the standard cross entropy loss due to the additional summation over the paths. In our experiments, we observed an approximately 2x slowdown. Modifications to the cross entropy loss that improve runtime (such as the hierarchical softmax or negative sampling) could also be used to improve the runtime of the tree loss, but we do not explore this possibility in depth in this paper.

5 Theoretical Results

We use standard properties of stochastic gradient descent to bound the generalization error of the tree loss. In this section we formally state our main results, but we do not review the details about stochastic gradient descent or prove the results. The Appendix reviews stochastic gradient descent in detail and uses the results to prove the theorems below.

¹<https://github.com/cora1021/TreeLoss>

To state our results, we first must introduce some learning theory notation. Define the true loss of our learning problem as

$$L_D(A) = \mathbb{E}_{(\mathbf{x}, y) \sim D} \ell(A; (\mathbf{x}, y)) \quad (5)$$

where $A \in \{W, U, V\}$ is a parameterization of the cross entropy loss. We define \bar{A} to be the result of running SGD on n data points to optimize parameterization A , and

$$A^* = \arg \min_A L_D(A) \quad (6)$$

to be the optimal parameterization. Our goal is to bound the generalization error

$$\mathbb{E} L_D(\bar{A}) - L_D(A^*). \quad (7)$$

In order to bound this error, we make the following standard assumption about our data.

Assumption 1. For each feature vector \mathbf{x} in the data set, $\|\mathbf{x}\|_2 \leq \rho$.

This assumption is equivalent to stating that L_D (or ℓ) is ρ -Lipschitz.

Now our key observation is that $\|A^*\|_F$ bounds the generalization error, as formalized in the following lemma.

Lemma 1. Under Assumption 1, we have that for any parameterization A of the cross entropy loss,

$$\mathbb{E} L_D(\bar{A}) - L_D(A^*) \leq \frac{\|A^*\|_F \rho}{\sqrt{n}}. \quad (8)$$

We will next show how to use Lemma 1 to recover the standard convergence rate of multi-class classification by bounding $\|W^*\|_F$. We make the following assumption.

Assumption 2. For each class i , the optimal parameter vector \mathbf{w}_i^* satisfies $\|\mathbf{w}_i^*\|_2 \leq B$.

It follows that

$$\|W^*\|_F^2 = \sum_{i=1}^k \|\mathbf{w}_i^*\|_2^2 \leq kB^2. \quad (9)$$

Substituting Eq (30) into (29) gives the following bound.

Corollary 1. Under assumptions 1 and 2, then the generalization error of the standard cross entropy parameterization when trained with SGD is

$$\mathbb{E} L_D(\bar{W}) - L_D(W^*) \leq \frac{\sqrt{k}B\rho}{\sqrt{n}}. \quad (10)$$

Next, we bound $\|U^*\|_F$ and $\|V^*\|_F$ in order to bound the generalization error of the U/V parameterizations. The analysis is divided into two parts. First,

we make no assumptions that the U -tree structure is good, and we show that even in the worst case $\|U^*\|_F = O(\|W^*\|_F)$. This implies that using the tree loss cannot significantly hurt our performance.

Lemma 2. Under assumption 2, the following bound holds for all U/V -tree structures:

$$\|V^*\|_F \leq \|U^*\|_F \leq 2\sqrt{k}B. \quad (11)$$

Now we consider the more important case of when we have a tree structure that meaningfully captures the similarity between classes. This idea is captured in our final assumption.

Assumption 3. Let $\lambda \geq 1$, and d be a distance metric over the labels such that for all labels i and j ,

$$\frac{1}{\lambda}d(i, j) \leq \|\mathbf{w}_i^* - \mathbf{w}_j^*\|_2 \leq \lambda d(i, j). \quad (12)$$

We denote by c the doubling dimension of the metric d , and we assume that the U -tree structure is built using a cover tree (Beygelzimer et al., 2006).

The λ parameter above measures the quality of our metric. When $\lambda = 1$, the metric d is good and perfectly predicts the distance between parameter vectors; when λ is large, the metric d is bad.

The Cover Tree was originally designed for speeding up nearest neighbor queries in arbitrary metric spaces. The definition is rather technical, so we do not restate it here. Instead, we mention only two important properties of the cover tree. First, it can be constructed in time $O(k)$. This is independent of the number of training examples n , so building the U -tree structure is a cheap operation that has no meaningful impact on training performance. Second, the cover tree has a hyperparameter which we call **base**. This hyperparameter controls the fanout and depth of the tree because at any node i at depth $\text{depth}(i)$ in the tree, the cover tree maintains the invariant that $d(i, \text{parent}(i)) \leq \text{base}^{-\text{depth}(i)}$. Increasing the **base** hyperparameter results in shorter, fatter trees, and decreasing it results in taller, narrower trees. Our analysis follows the convention of setting **base** = 2, but we show in the experiments below that good performance is achieved for a wide range of base values.

The following Lemma uses Assumption 3 and properties of the cover tree to bound the norm of $\|U^*\|_F$. It is our most difficult technical result.

Lemma 3. Under Assumptions 2 and 3, when $c \leq 1$, we have that

$$\|U^*\|_F \leq \frac{1}{\sqrt{2}}\lambda B \sqrt{\log_2 k}, \quad (13)$$

and when $c > 1$, we have that

$$\|U^*\|_F \leq \sqrt{5}\lambda B \sqrt{k^{(1-1/c)}}. \quad (14)$$

We note that embedding techniques can be used to reduce the intrinsic dimension of the metric (c) at the expense of increasing the metric’s distortion (λ), but we make no attempt to optimize this tradeoff in this paper.

We now state our main result. It is an immediate consequence Lemmas 1, 2 and 3.

Corollary 2. *Under assumptions 1, 2, and 3, when $c \leq 1$, the generalization error of the tree loss is bounded by*

$$\mathbb{E} L_D(\bar{V}) - L_D(V^*) \leq \frac{\lambda B \rho \sqrt{\log_2 k}}{\sqrt{2n}}. \quad (15)$$

When $c > 1$, the generalization error of the tree loss is bounded by

$$\mathbb{E} L_D(\bar{V}) - L_D(V^*) \leq \frac{\lambda B \rho \sqrt{5k^{(1-1/c)}}}{\sqrt{n}}. \quad (16)$$

These convergence rates are asymptotically better than the convergence rates for the standard parameterization of the cross entropy loss.

6 Experiments

We evaluate the tree loss on 4 synthetic and 3 real world experiments. We compare the tree loss to the standard cross entropy loss, the recently proposed Sim-Loss (Kobs et al., 2020), and the hierarchical softmax (Morin and Bengio, 2005). The results confirm our theoretical findings from Section 5 and demonstrate that the tree loss significantly improves the performance of these baseline losses in a wide range of scenarios.

6.1 Synthetic data

Experiments on synthetic data let us control various hyperparameters of the dataset in order to see how the tree loss behaves in a wide range of scenarios. The next subsection introduces our data generation procedure, and then subsequent subsections describe 4 different experiments based on this procedure.

6.1.1 Data Generation Procedure

Our data generation procedure has 4 key hyperparameters: the number of data points n , the number of feature dimensions d , the number of classes k , and a randomness parameter σ .

Let \mathcal{N} be the standard normal distribution. Then sample the true parameter matrix as

$$W^* \sim \mathcal{N}^{k \times d}. \quad (17)$$

Standard results on random matrices show that $\|W^*\|_F = O(\sqrt{kd})$ with high probability, as assumed by our theory.

For each data point $i \in [n]$, we sample the data point according to the following rules:

$$y_i \sim \text{Uniform}([k]), \text{ and} \quad (18)$$

$$\mathbf{x}_i \sim \mathcal{N}(\mathbf{w}_{y_i}^*; \sigma). \quad (19)$$

Observe that larger σ values result in more noise in the data points, making the classes harder to distinguish, and increasing the bayes error of the problem. Also observe that for any two classes i and j , the distance between the two classes $\|\mathbf{w}_i^* - \mathbf{w}_j^*\|_2 = O(\sqrt{d})$. This implies that as d increases, the separation between data points from different classes will also increase, reducing the bayes error of the problem.

6.1.2 Experiment I: Data Regimes

Our first experiment investigates the tree loss’s performance in a wide range of statistical settings controlled by the dataset hyperparameters (n , d , k , and σ). We use a baseline experimental setting with $n = 100$, $d = 64$, $k = 10$, and $\sigma = 1$. For each of the hyperparameters, we investigate the effect of that hyperparameter’s performance on the problem by varying it over a large range. For each value in the range, we: (1) randomly sample 50 different W^* using the procedure described in Section 6.1.1; (2) for each W^* , we: (2a) sample a training set with n data points and a test set with 10000 data points²; (2b) train a separate model for each of our loss functions; (2c) report the average test accuracy across 50 samples from step (1). Figure 2 shows the results of these experiments. As our theory suggests, the tree loss outperforms all other losses in all regimes.

Consider the top-left plot. This plot has a low bayes error, and so all the variation in performance we see is due to the generalization error of the models. As the number of data points n grows large, its influence on the standard cross entropy’s convergence rate of $O(\sqrt{k/n})$ dominates, and the model generalizes well. When the number of data points is small, then the dependence on k becomes more important. The tree loss’s dependence of $O(\sqrt{k^{1-1/c}/n})$ is strictly better than the cross entropy loss’s, explaining the observed improved performance. A similar effect explains our model’s observed improved performance in the bottom-left plot as k varies.

Now consider the top-right plot. The bayes error of our problem setup is inversely related to the problem

²The size of the test set is fixed and large to ensure that our results have high statistical significance.

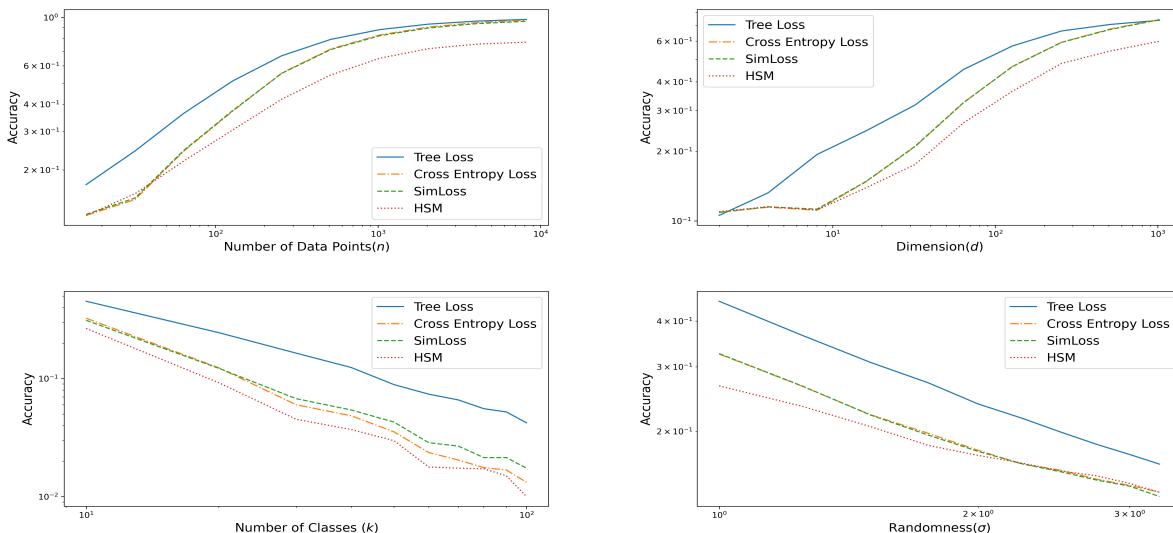


Figure 2: Results of Synthetic Experiment I. The x-axis of each plot shows which problem parameter is being varied. As our theory predicts, the Tree Loss outperforms the baseline loss functions in all data regimes.

dimension d (see observations in Section 6.1.1), so this plot compares performance on different problem difficulties. On maximally easy (large d) and maximally hard (small d) problems, the tree loss performs similarly to the other loss functions because the performance is dominated by the bayes error and no improvement is possible. The main advantage of the tree loss is in the mid-difficulty problems. In this regime, performance is dominated by the generalization ability of the different losses, and the tree loss’s improved generalization results in noticeable improvements. The bottom-right plot tells a similar story but controls the difficulty of the problem directly by tuning the randomness σ .

6.1.3 Experiment II: Parameter Norms

Lemma 1 suggests that the norm of the parameter matrix is what controls convergence rate of SGD, and the proof of our main result in Corollary 2 relies on bounding this norm. In this experiment, we directly measure the norm of the parameter matrices and show that $\|V\|_F$ is significantly less than $\|W\|_F$, justifying the good convergence rates we observed in Synthetic Experiment I above. We set $n = 1000$, $d = 10$, and $\sigma = 1$, and vary the number of classes k . Figure 3 shows that $\|V^*\|_F$ grows much slower than $\|W^*\|_F$. Notice in particular that $\|W^*\|_F$ grows at the rate $O(\sqrt{k})$ and $\|V^*\|_F$ grows at a rate strictly less than $o(\sqrt{k})$ as our theory predicts.

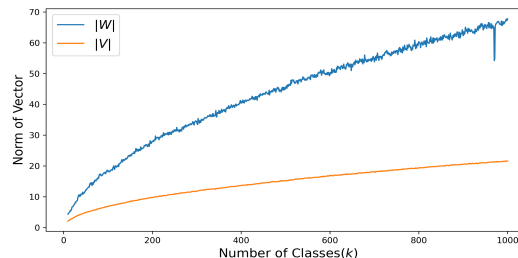


Figure 3: Results of Synthetic Experiment II. As we increase k , $\|V^*\|_F$ grows significantly slower than $\|W^*\|_F$. Since the convergence rate of SGD is proportional to $\|\cdot\|_F$ (by Lemma 1), the tree loss converges faster than the standard cross entropy loss.

6.1.4 Experiment III: Tree Shapes

In this experiment, we study how the shape of our tree structure impacts performance. We set $n = 1000$, $d = 10$, $k = 100$, and $\sigma = 1$. We use a relatively large number of classes k compared to the standard problem in Section 6.1.2 above in order to ensure that we have enough classes to generate meaningfully complex tree structures. The basic trends we observe are consistent for all other values of n , d , and σ .

Recall that the cover tree has a parameter `base` which controls the rate of expansion between layers of the tree. Reducing `base` increases the height of the tree and increasing `base` reduces the height of the tree. This will affect the performance of our tree loss because taller trees result in more parameter sharing. Figure 4 plots the accuracy of the tree loss as a func-

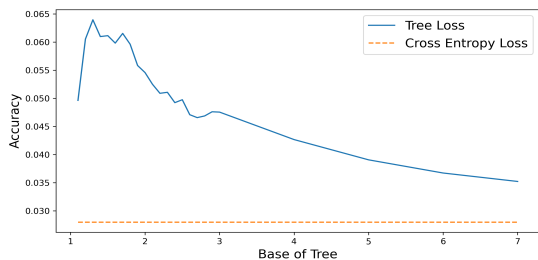


Figure 4: Results of Synthetic Experiment III. The tree loss outperforms the standard cross entropy loss for all values of `base`.

tion of the `base` hyperparameter. Across all ranges, the tree loss outperforms the cross entropy loss. Interestingly, the tree loss’s accuracy is maximized when `base` ≈ 1.3 . The original cover tree paper (Beygelzimer et al., 2006) also found that a `base` of 1.3 resulted in the fastest nearest neighbor queries. We do not know of a good theoretical explanation for this phenomenon.

6.1.5 Experiment IV: Metric Quality

Recall that our theoretical results state that the convergence rate of SGD depends on a factor λ that measures how well our metric over the class labels is able to predict the distances between the true parameter vectors. (See Assumption 3.) In this experiment, we study the effect of this λ parameter on prediction accuracy.

We fix the following problem hyperparameters: $n = 100$, $d = 64$, $k = 10$, and $\sigma = 1$. Then we construct a bad parameter matrix W^{bad} using the same procedure we used to construct the optimal parameter matrix; that is,

$$W^{bad} \sim \mathcal{N}^{k \times d}. \quad (20)$$

Each row \mathbf{w}_i^{bad} of W^{bad} is now a d dimensional random vector that has absolutely no relation to the true parameter vector. Next, we define a family of “ ϵ -bad” parameter vectors that mix between the bad and optimal parameter vectors:

$$\mathbf{w}_i^\epsilon = (1 - \epsilon)\mathbf{w}_i^* + \epsilon\mathbf{w}_i^{bad}. \quad (21)$$

Finally, we define our ϵ -bad distance metric as

$$d_\epsilon(i, j) = \|\mathbf{w}_i^\epsilon - \mathbf{w}_j^\epsilon\|_2 \quad (22)$$

and build our cover tree structure using d_ϵ . When $\epsilon = 0$, this cover tree structure will be the ideal structure and λ will be equal to 1; when $\epsilon = 1$, this cover tree structure will be maximally bad, and λ will be large.

Figure 5 shows the performance of the tree loss as a function of ϵ . Remarkably, the tree loss outperforms




Figure 5: Results for Synthetic Experiment IV. As $\epsilon \rightarrow 1$, the tree structure becomes perfectly random. The tree loss still outperforms the cross entropy loss in this worst-case scenario.

the standard cross entropy loss even when using a perfectly bad tree structure (i.e. when $\epsilon = 1$). This surprising empirical finding actually agrees with our theoretical results in two ways. First, Lemma 2 implies that $\|V^*\|_F = O(\|W^*\|_F)$, and so no tree structure can perform asymptotically worse than the standard cross entropy loss. Second, when $\epsilon = 1$ and we have a perfectly random tree structure, the distance d_ϵ can still be embedded into the ideal metric d with some large (but finite!) λ . Asymptotically, Corollary 2 then implies that SGD converges at the rate $O(\sqrt{k^{1-1/\epsilon}})$, which is less than the convergence rate of the standard cross entropy loss.

6.2 Real World Data

We now validate the tree loss on real world image (CIFAR100, ImageNet) and text (Twitter) data. For each dataset, we report standard top-1 accuracy scores, and the less well known *similarity accuracy* (SA) score. SA is a top-1 accuracy considering similarity. The idea is that misclassifying a **sheepdog** as a **husky** should be penalized less than misclassifying a **sheepdog** as a **bus** because the **sheepdog** and **husky** classes are more similar to each other.

Table 1 shows that on each of these datasets and for each metric, the tree loss outperforms the baseline cross entropy loss and SimLoss.³ The remainder of this section describes the experimental procedures for each dataset in detail.

6.2.1 CIFAR100

CIFAR100 is a standard image dataset with 100 classes (Krizhevsky et al., 2009). It is more difficult than MNIST and CIFAR10, but small and so less com-

³We do not evaluate against the hierarchical softmax due to limited computational resources. The results on the synthetic data, however, suggest that the hierarchical softmax would have performed poorly.

	CIFAR100		ImageNet		Twitter	
	Top1	Acc	SA	Top1	Acc	SA
tree loss	52.78	65.90	76.52	88.22	20.96	50.54
cross entropy loss	51.87	65.58	75.30	87.51	19.95	48.33
simloss	52.43	65.82	76.24	87.98	19.96	48.32

Table 1: Experimental results on real world datasets. For all performance measures, larger numbers are better. The tree loss achieves the best results in all cases.

putationally demanding than ImageNet. In this experiment, we exactly follow the procedure used in the SimLoss paper (Kobs et al., 2020), and find that under their experimental conditions, the tree loss has the best performance.

First, we generate our distance metric over the class labels using a word2vec model (Mikolov et al., 2013) pre-trained on GoogleNews. The distance between class labels is defined to be the distance between their corresponding word2vec embeddings. There are 4 class labels with no corresponding word2vec embedding, and so following the SimLoss paper we discard these classes.

We train 3 ResNet20 models on CIFAR100 (He et al., 2016) using hyperparameters specified in the original paper. The only difference between the three models is the final layer: one model uses the standard cross entropy loss, one uses the tree loss, and one uses the SimLoss. The results shown in Table 1 shows that the tree loss significantly outperforms the other losses.

6.2.2 ImageNet

ImageNet is the gold-standard dataset for image classification tasks (Russakovsky et al., 2015). It has 1000 class labels and 1.2 million images.

We generate our distance metric over the image labels using a pretrained fastText model (Bojanowski et al., 2017). We use fastText rather than word2vec because many of the class labels contain words not in the word2vec model’s vocabulary; since fastText is naturally able to handle out-of-vocabulary words, there is no need to discard class labels like we did for CIFAR100. Many ImageNet class label names contain multiple words, and for these classes we generate the embedding with a simple average of the fastText embedding over all words.

We again train a ResNet50 model (He et al., 2016) using standard hyperparameters, replacing only the last layer of the model. As with CIFAR100, Table 1 shows that the tree loss significantly outperforms the other losses. Since the publication of ResNet50 in 2016, there have been many newer network archi-

tectures with better performance on ImageNet (e.g. Howard et al., 2017; Huang et al., 2017; Tan and Le, 2019). We unfortunately did not have sufficient computational resources to train these more modern models with the tree loss, and we emphasize that we are not trying to compete directly with these network architectures to achieve state-of-the-art performance. Instead, our goal is to show that for a given network architecture, replacing the cross entropy loss with the tree loss results in improved performance.

6.2.3 Twitter Data

We now consider the problem of predicting the emotions of Twitter text. This is a very different domain than the image problems considered above, and demonstrates that the tree loss works across many domains. We use the dataset collected by Izbicki et al. (2019) and Stoikos and Izbicki (2020). This dataset contains all geolocated tweets written in more than 100 languages sent over the 4 year period between October 2017 and October 2021—approximately 5.5×10^9 tweets in total. Tweets in the dataset are preprocessed to have all emojis, usernames, and urls removed from the text, and then the goal is to predict the emoji given only the remaining text. Since most emojis represent emotions, the task of emoji prediction serves as a proxy for emotion prediction.

We generate our distance metric over the class labels using the pretrained emoji2vec model (Eisner et al., 2016), which associates each emoji with a vector embedding. We then follow the procedure in Stoikos and Izbicki (2020) to train multilingual BERT models (Feng et al., 2020) with the last layers replaced by the three different loss functions. Table 1 shows the tree loss significantly outperforming the baseline models.

7 Conclusion

The tree loss is a drop-in replacement for the cross entropy loss for multiclass classification problems. It can take advantage of background knowledge about the underlying class structure in order to improve the convergence rate of SGD. Both theoretical and empir-

ical results suggest there is no disadvantage for using the tree loss, but potentially large advantages.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Luca Bertinetto, Romain Mueller, Konstantinos Tzitikas, Sina Samangooei, and Nicholas A Lord. Making better mistakes: Leveraging class hierarchies with deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12506–12515, 2020.
- Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd international conference on Machine learning*, pages 97–104, 2006.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017.
- Nicolo Cesa-Bianchi, Claudio Gentile, and Luca Zani-boni. Incremental algorithms for hierarchical classification. *The Journal of Machine Learning Research*, 7:31–54, 2006.
- Ben Eisner, Tim Rocktäschel, Isabelle Augenstein, Matko Bosnjak, and S. Riedel. emoji2vec: Learning emoji representations from their description. *ArXiv*, abs/1609.08359, 2016.
- Fangxiaoyu Feng, Yin-Fei Yang, Daniel Matthew Cer, N. Arivazhagan, and Wei Wang. Language-agnostic bert sentence embedding. *ArXiv*, abs/2007.01852, 2020.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- Mike Izbicki. *Divide and Conquer Algorithms for Machine Learning*. PhD thesis, University of California, Riverside, 2017.
- Mike Izbicki and Christian Shelton. Faster cover trees. In *International Conference on Machine Learning*, pages 1162–1170. PMLR, 2015.
- Mike Izbicki, Vagelis Papalexakis, and Vassilis Tso-tras. Geolocating tweets in any language at any location. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 89–98, 2019.
- Nan Jiang, Wenge Rong, M. Gao, Yikang Shen, and Z. Xiong. Exploration of tree-based hierarchical softmax for recurrent language models. In *IJCAI*, 2017.
- Konstantin Kobs, M. Steininger, Albin Zehe, Florian Lautenschlager, and A. Hotho. Simloss: Class similarities in cross entropy. *ArXiv*, abs/2003.03182, 2020.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Maksim Lapin, Matthias Hein, and Bernt Schiele. Loss functions for top-k error: Analysis and insights. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1468–1477, 2016.
- Tomas Mikolov, Kai Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. In *ICLR*, 2013.
- Abdul Arfat Mohammed and Venkatesh Umaashankar. Effectiveness of hierarchical softmax in large scale classification tasks. *2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, pages 1090–1094, 2018.
- Frederic Morin and Yoshua Bengio. Hierarchical probabilistic neural network language model. In *Aistats*, volume 5, pages 246–252. Citeseer, 2005.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward

- Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. 2019.
- Hao Peng, Jianxin Li, Y. Song, and Yaopeng Liu. Incrementally learning the hierarchical softmax function for neural language models. In *AAAI*, 2017.
- Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. *ICML*, 2012.
- Olga Russakovsky, J. Deng, Hao Su, J. Krause, S. Satheesh, S. Ma, Zhiheng Huang, A. Karpthy, A. Khosla, Michael S. Bernstein, A. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115:211–252, 2015.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Stefanos Stoikos and Mike Izbicki. Multilingual emotion prediction of tweets about covid-19. In *Proceedings of the Third Workshop on Computational Modeling of People’s Opinions, Personality, and Emotion’s in Social Media*, pages 109–118, 2020.
- Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. Training convolutional networks with noisy labels. *ICLR*, 2015.
- Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/tan19a.html>.
- Cinna Wu, Mark Tygert, and Yann LeCun. A hierarchical loss and its problems when classifying non-hierarchically. *PLoS ONE*, 2019.
- Zhixuan Yang, Chong Ruan, Caihua Li, and J. Hu. Optimize hierarchical softmax with word similarity knowledge. *Polytech. Open Libr. Int. Bull. Inf. Technol. Sci.*, 55:11–16, 2017.
- Zhilu Zhang and Mert R Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *32nd Conference on Neural Information Processing Systems (NeurIPS)*, 2018.

Supplementary Material: The Tree Loss: Improving Generalization with Many Classes

A Theoretical Results (Supplement)

This appendix is an expanded explanation of our theoretical results from Section 5 in the main paper. We begin with a short background explaining the standard properties of stochastic gradient descent that our proofs rely on, then we restate our theoretical results from before along with detailed proofs.

A.1 SGD Background

SGD is a standard algorithm for optimizing convex losses, and it has many variations. Our analysis uses the results from Chapter 14 of [Shalev-Shwartz and Ben-David \(2014\)](#) which are based off of a version of SGD that uses parameter averaging.⁴ In this section we review these results. We follow the notation of [Shalev-Shwartz and Ben-David \(2014\)](#) as closely as possible, but we occasionally change the names of variables to avoid conflicts with variable names used elsewhere in this paper.⁵

SGD is an iterative procedure for optimizing a function $f : \Theta \rightarrow \mathbb{R}$ given only noisy samples of the function’s gradient. We let $\theta^{(t)}$ denote the parameter value at iteration t , and $\eta : \mathbb{R}$ be the step size. Let ν_t be a noisy gradient of f satisfying

$$\mathbb{E}[\nu_t | \theta^{(t)}] = \nabla f(\theta^{(t)}). \quad (23)$$

Then the parameter update for step t is defined to be

$$\theta^{(t+1)} = \theta^{(t)} - \eta \nabla \nu_t. \quad (24)$$

When using parameter averaging, the final model parameters are defined to be

$$\bar{\theta} = \frac{1}{T} \sum_{i=1}^T \theta^{(i)}. \quad (25)$$

Parameter averaging is not often used in practice because it has a large memory overhead, but it is widely used in theoretical analysis of SGD. Theorems using

⁴ Other variants of SGD have similar results, and we refer the reader to [Shalev-Shwartz and Ben-David \(2014\)](#) for a discussion of these variants.

⁵ Specifically, we replace all Latin variables with Greek alternatives. We use θ instead of \mathbf{w} , ν instead of \mathbf{v} , and β instead of B . The Greek letters ρ and η have the same meaning in both documents.

parameter averaging are typically much easier to state and prove, and equivalent results without parameter averaging typically require an additional $\log T$ factor in their convergence rates ([Rakhlin et al., 2012](#)).

The main result we rely on is a bound on the convergence rate of SGD for convex Lipschitz functions. We reproduce Theorem 14.8 of [Shalev-Shwartz and Ben-David \(2014\)](#) in the theorem below.

Theorem 1. *Let $\beta, \rho > 0$. Let f be a convex, ρ -Lipschitz function and let $\theta^* = \arg \min_{\theta: \|\theta\|_2 \leq \beta} f(\theta)$. Assume that SGD is run for T iterations with $\eta = \sqrt{\frac{\beta^2}{\rho^2 T}}$. Then,*

$$\mathbb{E} f(\bar{\theta}) - f(\theta^*) \leq \frac{\beta \rho}{\sqrt{T}}. \quad (26)$$

A.2 SGD and Statistical Learning

We now introduce more statistical learning theory notation, and apply Theorem 1 above to bound the rate of learning.

Define the true loss of our learning problem as

$$L_D(A) = \mathbb{E}_{(\mathbf{x}, y) \sim D} \ell(A; (\mathbf{x}, y)) \quad (27)$$

where $A \in \{W, U, V\}$ is a parameterization of the cross entropy loss. The optimal parameter vector is defined to be the minimum of the true loss:

$$A^* = \arg \min_A L_D(A). \quad (28)$$

In order to apply Theorem 1 to a learning problem, we make the following assumption about the size of our data points.

Assumption 1. *For each feature vector \mathbf{x} in the data set, $\|\mathbf{x}\|_2 \leq \rho$.*

Under this assumption, the cross entropy loss ℓ is ρ -Lipschitz for all parameterizations, and the true loss L_D is therefore also ρ -Lipschitz.

We use SGD to optimize L_D . We let $f = L_D$, and $\theta = A$, so ν_t is the gradient of ℓ on the t th data point. Applying Theorem 1 gives the following convergence bound.

Lemma 1. *Under Assumption 1, we have that for any parameterization A of the cross entropy loss, SGD*

converges at the rate

$$\mathbb{E} L_D(\bar{A}) - L_D(A^*) \leq \frac{\|A^*\|_F \rho}{\sqrt{n}}. \quad (29)$$

We emphasize that this convergence rate above is both computational and statistical. The expression on the left hand side of (29) is equivalent to the expected generalization error of the parameter \bar{A} , and so we use the terms ‘‘convergence rate’’ and ‘‘generalization error’’ essentially interchangeably.

The important intuition of Lemma 1 is that the generalization error of the cross entropy loss depends on the Frobenius norm of the optimal parameter matrix. We can therefore compare the convergence of the standard parameterization with the U/V -tree parameterizations by comparing these norms.

Our next task is to recover the standard convergence rate of multi-class classification, and so we will now bound $\|W^*\|_F$. We make the following assumption.

Assumption 2. For each class i , the optimal parameter vector \mathbf{w}_i^* satisfies $\|\mathbf{w}_i^*\|_2 \leq B$.

It follows that

$$\|W^*\|_F^2 = \sum_{i=1}^k \|\mathbf{w}_i^*\|_2^2 \leq kB^2. \quad (30)$$

Substituting Eq (30) into (29) gives the following bound.

Corollary 1. Under assumptions 1 and 2, then the generalization error of the standard cross entropy parameterization when trained with SGD is

$$\mathbb{E} L_D(\bar{W}) - L_D(W^*) \leq \frac{\sqrt{k}B\rho}{\sqrt{n}}. \quad (31)$$

A.3 Bounding the Tree Loss

We now bound the Frobenius norm of U^* in order to bound the generalization of the U -tree loss. By construction, $\|V^*\|_F \leq \|U^*\|_F$, so this will also bound the convergence rates of the V -tree loss.

The analysis is divided into two parts. First, we make no assumptions that the U -tree structure is good, and we show that even in the worst case $\|U^*\|_F = O(\|W^*\|_F)$. This implies that using the tree loss cannot significantly hurt our performance.

Lemma 2. Under assumption 2, the following bound holds for all tree structures:

$$\|V^*\|_F \leq \|U^*\|_F \leq 2\sqrt{k}B. \quad (32)$$

Proof. We have

$$\|U^*\|_F^2 = \sum_{i=1}^k \|\mathbf{u}_i^*\|_2^2 \quad (33)$$

$$= \sum_{i=1}^k \|\mathbf{w}_i^* - \mathbf{w}_{\text{parent}(i)}^*\|_2^2 \quad (34)$$

$$\leq \sum_{i=1}^k \left(\|\mathbf{w}_i^*\|_2 + \|\mathbf{w}_{\text{parent}(i)}^*\|_2 \right)^2 \quad (35)$$

$$\leq \sum_{i=1}^k (2B)^2 \quad (36)$$

$$\leq k(2B)^2. \quad (37)$$

□

Now we consider the more important case of when we have a useful U -tree structure that meaningfully captures the similarity between classes. This idea is captured in the following assumption.

Assumption 3. Let $\lambda \geq 1$, and d be a distance metric over the labels such that for all labels i and j ,

$$\frac{1}{\lambda}d(i, j) \leq \|\mathbf{w}_i^* - \mathbf{w}_j^*\|_2 \leq \lambda d(i, j). \quad (38)$$

We denote by c the doubling dimension of the metric d , and we assume that the U -tree structure is built using a cover tree (Beygelzimer et al., 2006).

The Cover Tree was originally designed for speeding up nearest neighbor queries in arbitrary metric spaces. The definition is rather technical, so we do not restate it here. Instead, we mention only two important properties of the cover tree. First, it can be constructed in time $O(k)$, so building the cover tree is a cheap operation that has no meaningful impact on training performance. It also only has to be done in the initialization step of the learning process, and not on every iteration. Second, the cover tree construction has a hyperparameter which we call **base**. This hyperparameter controls the fanout and depth of the tree because at any node i at depth $\text{depth}(i)$ in the tree, the cover tree maintains the invariant that $d(i, \text{parent}(i)) \leq \text{base}^{-\text{depth}(i)}$. Increasing the **base** hyperparameter results in shorter, fatter trees, and decreasing it results in taller, narrower trees. Our analysis follows the convention of setting **base** = 2, but we show in the experiments below that good performance is achieved for a wide range of base values.

The following Lemma uses Assumption 3 and properties of the cover tree to bound the norm of $\|U^*\|_F$. It is our main technical result and the proof is deferred to the appendix.

Lemma 3. *Under Assumptions 2 and 3, when $c \leq 1$, we have that*

$$\|U^*\|_F \leq \frac{1}{\sqrt{2}} \lambda B \sqrt{\log_2 k}, \quad (39)$$

and when $c > 1$, we have that

$$\|U^*\|_F \leq \sqrt{5} \lambda B \sqrt{k^{(1-1/c)}}. \quad (40)$$

Proof. First, we review some simple properties of the cover tree structure that will be needed in our proof. Let q denote the maximum number of children of any node in the cover tree. A standard property of cover trees is that $q \leq (4^c)$ (Izbicki, 2017, Lemma 25). Let Q_i denote the number of nodes at level i in the cover tree. Then we have that $|Q_i| \leq q^i$. Finally, let r be the height of the cover tree. Without loss of generality, we analyze the bound when the cover tree is perfectly balanced and so $r = \log_q k = \log_2 k / (2c)$.

We have that

$$\|U^*\|_F^2 = \sum_{i=1}^k \|\mathbf{u}_i^*\|_2^2 \quad (41)$$

$$= \sum_{i=0}^r \sum_{j \in Q_i} \|\mathbf{u}_j^*\|_2^2 \quad (42)$$

$$= \sum_{i=0}^r \sum_{j \in Q_i} \|\mathbf{w}_j^* - \mathbf{w}_{\text{parent}(j)}^*\|_2^2 \quad (43)$$

$$\leq \sum_{i=0}^r \sum_{j \in Q_i} \lambda^2 d(j, \text{parent}(j))^2 \quad (44)$$

$$\leq \sum_{i=0}^r \sum_{j \in Q_i} \lambda^2 (B/2^i)^2 \quad (45)$$

$$\leq \sum_{i=0}^r q^i \lambda^2 (B/2^i)^2 \quad (46)$$

$$\leq \sum_{i=0}^r \lambda^2 B^2 (q/4)^i \quad (47)$$

We proceed by considering the case of $c \leq 1$ and $c > 1$ separately.

When $c \leq 1$, $q \leq 4$. Substituting into Eq (47), we have

$$\|U^*\|_F^2 \leq \sum_{i=0}^r \lambda^2 B^2 \quad (48)$$

$$= (r+1) \lambda^2 B^2 \quad (49)$$

$$= \frac{\log_2 k}{2c} \lambda^2 B^2 \quad (50)$$

$$\leq \frac{\log_2 k}{2} \lambda^2 B^2. \quad (51)$$

This matches Eq (39) in the lemma's statement.

When $c > 1$, then $q > 4$. We first show the following equality.

$$\left(\frac{q}{4}\right)^r = \left(\frac{q}{4}\right)^{\log_q k} \quad (52)$$

$$= \frac{k}{4^{\log_q k}} \quad (53)$$

$$= \frac{k}{4^{\log_4 k / \log_4 q}} \quad (54)$$

$$= \frac{k}{k^{1/\log_4 q}} \quad (55)$$

$$= k^{(1-1/\log_4 q)} \quad (56)$$

$$= k^{(1-1/c)}. \quad (57)$$

Substituting into Eq (47), we have

$$\|U^*\|_F^2 = \lambda^2 B^2 \left(\frac{(q/4)^{r+1} - 1}{q/4 - 1} \right) \quad (58)$$

$$= \lambda^2 B^2 \left(\frac{(q/4)k^{(1-1/c)} - 1}{q/4 - 1} \right) \quad (59)$$

$$\leq \lambda^2 B^2 \left(\frac{(q/4)k^{(1-1/c)}}{q/4 - 1} \right) \quad (60)$$

$$\leq 5 \lambda^2 B^2 k^{(1-1/c)} \quad (61)$$

Inequality (61) follows because $\frac{q/4}{q/4-1} \leq 5$ for integer values of $q > 4$. \square

We note that embedding techniques can be used to reduce the intrinsic dimension of the metric (c) at the expense of increasing the metric's distortion, but we make no attempt to optimize this tradeoff in this paper.

We now state our main result. It is an immediate consequence of the fact that $\|V^*\|_F \leq \|U^*\|_F$ and Lemmas 1 and 3.

Corollary 2. *Under assumptions 1, 2, and 3, when $c \leq 1$, the generalization error of the tree loss is bounded by*

$$\mathbb{E} L_D(\bar{V}) - L_D(V^*) \leq \frac{\lambda B \rho \sqrt{\log_2 k}}{\sqrt{2n}}. \quad (62)$$

When $c > 1$, the generalization error of the tree loss is bounded by

$$\mathbb{E} L_D(\bar{V}) - L_D(V^*) \leq \frac{\lambda B \rho \sqrt{5k^{(1-1/c)}}}{\sqrt{n}}. \quad (63)$$

These convergence rates are asymptotically better than the convergence rates for the standard parameterization of the cross entropy loss.