

Computationally Efficient Wasserstein Loss for Structured Labels

Ayato Toyokuni^{1,3} Sho Yokoi^{2,3} Hisashi Kashima^{1,3} Makoto Yamada^{1,3}

¹ Kyoto University ² Tohoku University ³ RIKEN AIP

{toyokuni.ayato@ml.ist.i, kashima@i, myamada@i}.kyoto-u.ac.jp,

yokoi@ecei.tohoku.ac.jp

Abstract

The problem of estimating the probability distribution of labels has been widely studied as a label distribution learning (LDL) problem, whose applications include age estimation, emotion analysis, and semantic segmentation. We propose a tree-Wasserstein distance regularized LDL algorithm, focusing on hierarchical text classification tasks. We propose predicting the entire label hierarchy using neural networks, where the similarity between predicted and true labels is measured using the tree-Wasserstein distance. Through experiments using synthetic and real-world datasets, we demonstrate that the proposed method successfully considers the structure of labels during training, and it compares favorably with the Sinkhorn algorithm in terms of computation time and memory usage.

1 Introduction

Label distribution learning (LDL), which is a generalized framework for performing single/multi-label classification and estimating the probability distribution over labels, is an important machine-learning problem (Geng, 2016). Its applications include age estimation (Geng et al., 2013), emotion estimation (Zhou et al., 2016), head-pose estimation (Geng and Xia, 2014), and semantic segmentation (Gao et al., 2017). In particular, multi-label classification is an important problem in many NLP areas, and has several applications including multi-label text classification (Banerjee et al., 2019; Chalkidis et al., 2019).

Typically, Kullback-Leibler (KL) divergence is used to measure the similarity between two distributions. However, the KL divergence can tend to infinity if the supports of the two distributions do not overlap, resulting in model failure.

To solve this support problem, Wasserstein distance is used instead of KL divergence (Arjovsky

et al., 2017). Wasserstein distance is defined as the cost of optimally transporting one probability distribution to match another (Villani, 2009; Peyré and Cuturi, 2018). Because it can compare two probability measures while considering the ground metric, it is more powerful than measurements that do not consider geometrical information.

An LDL framework with Wasserstein distance has been recently proposed (Frogner et al., 2015; Zhao and Zhou, 2018). This framework employs the Sinkhorn algorithm (Cuturi, 2013) to calculate the Wasserstein distance, which requires quadratic computational-time. Thus, when we consider extremely large label-sets, for example, 10^5 , the computation cost can be significant. However, the Wasserstein distance on a tree (hereinafter called *tree-Wasserstein distance*) can be written in a closed-form and calculated in linear computation time (Evans and Matsen, 2012; Le et al., 2019).

In this paper, we propose a tree-regularized LDL algorithm with a tree-Wasserstein distance. The key advantage of the tree-Wasserstein distance is that it considers the hierarchical label information explicitly, whereas the Sinkhorn-based algorithm needs a cost matrix using tree-structured data. Moreover, the tree-Wasserstein distance has an analytic form that can be computed in linear time using significantly less memory. We experimentally demonstrate that the proposed algorithm compares favorably with the Sinkhorn-based LDL algorithm (Frogner et al., 2015; Zhao and Zhou, 2018) with considerably lower memory consumption and computational costs. We demonstrate that the calculation is more efficient than that of the existing Wasserstein loss.

Contribution: Our contributions are summarized as follows. (1) We propose training a model by minimizing the tree-Wasserstein distance for hierarchical labels, and (2) we experimentally show

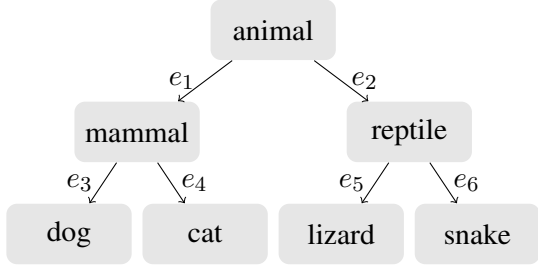


Figure 1: Illustration of a tree-structured label with the root “animal”. $\Gamma(\text{“mammal”}) = \{\text{“mammal”}, \text{“dog”}, \text{“cat”}\}$, $v_{e_2} = \text{“reptile”}$.

that the proposed method is computationally more efficient than the existing methods with Sinkhorn-based methods.

2 Problem Setting

We observe n input and output samples $\{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)\}$ from $(\mathcal{X}, \mathcal{Y})$, where $\mathcal{X} \subset \mathbb{R}^d$. We consider the problem of learning a map from a feature space \mathcal{X} into \mathcal{P} , which is a set of distributions over a finite set \mathcal{Y} .

For example, multi-class classification is included in this problem, \mathbf{y} , which represents the ℓ -th class, and it is expressed as the following one-hot vector:

$$\mathbf{y} = (0, \dots, 0, \underbrace{1}_{\ell}, 0, \dots, 0)^\top \in \mathbb{R}^L,$$

where L denotes the total number of classes, and $\mathbf{y}^\top \mathbf{1}_L = 1$. Additionally, $\mathbf{1}_L \in \mathbb{R}^L$ denotes a vector whose elements are all 1.

When multi-label classification is considered, \mathcal{P} denotes binary vectors that indicate existing labels. For example, if the sample \mathbf{x} belongs to classes ℓ and ℓ' , \mathbf{y} is given as

$$\mathbf{y} = (0, \dots, 0, \underbrace{1}_{\ell}, 0, \dots, 0, \underbrace{1}_{\ell'}, 0, \dots, 0)^\top \in \mathbb{R}^L,$$

where $\mathbf{y}^\top \mathbf{1}_L = 2$. Accordingly, we can transform \mathbf{y} into a probability vector as $\mathbf{p}_y = \mathbf{y}/\mathbf{y}^\top \mathbf{1}_L$. Notably, we assume that \mathcal{Y} is discrete and has a tree structure similar to hierarchical labels.

We aim to estimate the conditional probability vector \mathbf{p}_y for \mathbf{x} by considering the structure information of \mathcal{Y} from $\{(\mathbf{x}_1, \mathbf{p}_{y_1}), \dots, (\mathbf{x}_n, \mathbf{p}_{y_n})\}$.

3 Proposed Method

In this study, we assume \mathcal{Y} has a tree-structure. Accordingly, we propose LDL with tree-Wasserstein distance.

3.1 Wasserstein distance on tree metrics

Let \mathcal{T} be a tree with non-negative weighted edges and $\mathcal{N}_{\mathcal{T}}$ be the set of nodes of \mathcal{T} . A shortest path metric $d_{\mathcal{T}} : \mathcal{N}_{\mathcal{T}} \times \mathcal{N}_{\mathcal{T}} \rightarrow \mathbb{R}$ associated with \mathcal{T} is called the *tree metric*. Let v and v' be the nodes in \mathcal{T} . Accordingly, $d_{\mathcal{T}}(v, v')$ is equal to the sum of the edge weights along the shortest path between v and v' . Next, we know that $\mathcal{M}_{\mathcal{T}} = (\mathcal{N}_{\mathcal{T}}, d_{\mathcal{T}})$ is a metric space and can be naturally derived from \mathcal{T} .

It is assumed that \mathcal{T} is rooted at r . For each node v , the set of nodes in the sub-tree of \mathcal{T} rooted at v is defined as $\Gamma(v) = \{u \in \mathcal{N}_{\mathcal{T}} \mid v \in \mathcal{R}(u)\}$ where $\mathcal{R}(v)$ denotes the set of nodes in a unique path from a node v to the root r in \mathcal{T} . For each edge e , v_e denotes a deeper level node. Figure 1 illustrates a tree-structured label.

Given two probability measures μ, ν supported on $\mathcal{M}_{\mathcal{T}}$, the 1-Wasserstein distance between μ and ν is expressed as follows (Evans and Matsen, 2012; Le et al., 2019):

$$\mathcal{W}_{d_{\mathcal{T}}}^1(\mu, \nu) = \sum_{e \in \mathcal{T}} w_e |\mu(\Gamma(v_e)) - \nu(\Gamma(v_e))|, \quad (1)$$

where w_e denotes the weight of edge e . The key advantage of the tree-Wasserstein distance is that it can be computed with the linear time complexity, whereas the time complexity for the Sinkhorn algorithm is quadratic (Cuturi, 2013).

3.2 LDL with tree-Wasserstein distance

We define the tree-Wasserstein regularizer as follows.

Definition 1 (tree-Wasserstein regularizer). *Let $\mathbf{h}_{\theta} : \mathcal{X} \rightarrow \mathcal{P}$ be a model with learnable parameters θ . Let $T_{\mathcal{Y}} = (V, E, W_E)$ be a tree associated with \mathcal{Y} , where V denotes the set of nodes, E is the set of edges, and $W_E(e)$ is the length of edge $e \in E$. Given input $\mathbf{x} \in \mathcal{X}$ and the ground-truth distribution of $\mathbf{y} \in \mathcal{P}$, then the tree-Wasserstein regularization term $\mathcal{TW}(\mathbf{x}, \mathbf{p}_y)$ is defined as follows:*

$$\begin{aligned} \mathcal{TW}(\mathbf{x}, \mathbf{p}_y) &= \sum_{e \in \mathcal{T}} W_E(e) |(\mathbf{h}_{\theta}(\mathbf{x}))(\Gamma(v_e)) - \mathbf{p}_y(\Gamma(v_e))|, \end{aligned}$$

where \mathbf{h}_{θ} denotes the prediction model.

Using the tree-Wasserstein regularizer, we pro-

Table 1: The results for the Synthetic dataset. The label distributions are given on a random tree with 1000 nodes.

Loss	Wasserstein ↓	KL ↓	Cheby↓	Clark ↓	Canbe ↓	Cos ↑	IntSec ↑
\mathcal{KL}	9.701 ± (.050)	0.431 ± (.001)	0.209 ± (.001)	1.777 ± (.011)	14.512 ± (.060)	0.877 ± (.000)	0.754 ± (.001)
$\mathcal{KL} + \frac{1}{2}\mathcal{W}^1$	10.831 ± (.044)	0.452 ± (.001)	0.230 ± (.001)	1.666 ± (.009)	13.834 ± (.064)	0.868 ± (.000)	0.739 ± (.001)
$\mathcal{KL} + \mathcal{W}^1$	11.631 ± (.048)	0.475 ± (.001)	0.244 ± (.001)	1.618 ± (.008)	13.474 ± (.063)	0.859 ± (.000)	0.727 ± (.001)
$\mathcal{KL} + \frac{1}{2}\mathcal{TW}$	7.257 ± (.110)	0.595 ± (.007)	0.193 ± (.001)	2.098 ± (.040)	19.636 ± (.171)	0.833 ± (.002)	0.729 ± (.003)
$\mathcal{KL} + \mathcal{TW}$	7.158 ± (.117)	0.631 ± (.007)	0.195 ± (.001)	2.143 ± (.030)	19.923 ± (.441)	0.825 ± (.003)	0.721 ± (.004)

Table 2: The results for BlurbGenreCollectionEN.

Loss	Pseudo-Recall	Top5	AUC
\mathcal{KL}	0.679 ± (.008)	1.013 ± (.015)	0.971 ± (.001)
$\mathcal{KL} + \frac{1}{2}\mathcal{W}^1$	0.675 ± (.008)	1.009 ± (.013)	0.970 ± (.002)
$\mathcal{KL} + \mathcal{W}^1$	0.678 ± (.004)	1.008 ± (.018)	0.970 ± (.001)
$\mathcal{KL} + \frac{1}{2}\mathcal{TW}$	0.678 ± (.010)	0.993 ± (.013)	0.971 ± (.002)
$\mathcal{KL} + \mathcal{TW}$	0.678 ± (.009)	0.991 ± (.017)	0.970 ± (.001)

pose the following LDL:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \sum_{i=1}^n \lambda \mathcal{TW}(\mathbf{h}_{\theta}(\mathbf{x}_i), \mathbf{p}_{y_i}) + \mathcal{KL}(\mathbf{h}_{\theta}(\mathbf{x}_i), \mathbf{p}_{y_i}), \quad (2)$$

where

$$\mathcal{KL}(\mathbf{h}_{\theta}(\mathbf{x}_i), \mathbf{p}_{y_i}) = \sum_{\ell=1}^L \mathbf{p}_{y_i}^{(\ell)} \log \frac{\mathbf{p}_{y_i}^{(\ell)}}{\mathbf{h}_{\theta}(\mathbf{x}_i)^{(\ell)}}, \quad (3)$$

is the multi-class Kullback-Leibler loss function, and $\lambda \geq 0$ is its regularization parameter.

Notably, $\mathcal{TW}(\mathbf{h}_{\theta}(\mathbf{x}_i), \mathbf{p}_{y_i})$ is calculated in $O(L)$ time, where L denotes the number of labels. Unlike the Sinkhorn-Knopp algorithm, we need not compute and hold a distance matrix. For tree-structured labels, including hierarchical labels, the tree structure can be used directly as a tree metric. If we have prior knowledge about labels (e.g., similarity), we can set edge-weights using the prior knowledge.

4 Related Work

4.1 Label distribution learning

LDL (Geng, 2016) is the task of estimating the distribution of labels from each input. While age estimation (Geng et al., 2013), head-pose estimation (Geng and Xia, 2014), and semantic segmentation (Gao et al., 2017) are well known LDL tasks, in this study, we consider the task of estimating a distribution on a hierarchical structure. The key difference between LDL and a generative model is that the “true” distribution on labels is given in LDL.

4.2 Wasserstein distance

Given two probability vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}_{\geq 0}^n$ and a distance matrix $\mathbf{D} \in \mathbb{R}_{\geq 0}^{n \times n}$, the 1-Wasserstein distance $\mathcal{W}^1(\mathbf{a}, \mathbf{b})$ between \mathbf{a} and \mathbf{b} is defined as:

$$\mathcal{W}^1(\mathbf{a}, \mathbf{b}) = \min_{\mathbf{P} \in \Pi} \langle \mathbf{D}, \mathbf{P} \rangle, \quad (4)$$

where Π denotes the set of transport plans such that $\Pi = \{\mathbf{P} \in \mathbb{R}_{\geq 0}^{n \times n} \mid \mathbf{P}\mathbf{1}_n = \mathbf{a}, \mathbf{P}^{\top}\mathbf{1}_n = \mathbf{b}\}$.

Because Wasserstein distance can incorporate the ground metric in the comparison of the probability distributions, it has been widely used in applications, including domain adaptation (Courty et al., 2017), generative models (Arjovsky et al., 2017), and natural language processing (Kusner et al., 2015). A loss function that uses the Wasserstein distance can improve predictions based on a structure of labels (Frogner et al., 2015; Zhao and Zhou, 2018). Additionally, an entropic optimal transport loss can provide a robustness against noise labels by finding the coupling of the data samples and propagating their labels according to the coupling weight (Damodaran et al., 2020).

Frogner et al. (2015) proposed learning using a Wasserstein loss to consider the geometric information in predicting a probability distribution. Because computing a sub-gradient of the exact Wasserstein loss is expensive, they estimated the sub-gradient by introducing an entropic-regularization term and using the Sinkhorn-Knopp algorithm. Although they also suggested extending the Wasserstein loss to unnormalized measures, we do not consider this case. Zhao and Zhou (2018) showed that Wasserstein loss influenced LDL in terms of simultaneously learning label correlations and distribution. We proposed learning using an exact Wasserstein distance with efficient computations when the ground metric is represented by a tree.

Le et al. (2019) suggested the tree-sliced Wasserstein distance, where the Wasserstein distance is approximated on a continuous space by averaging the Wasserstein distances on tree metrics constructed by dividing that space. An unbalanced variant of

the tree-Wasserstein distance has been recently proposed (Sato et al., 2020).

5 Experiments

We applied our proposed method to LDL on trees based on a synthetic dataset and to multi-label text classification of a hierarchical structure based on a real dataset. We implemented all the methods using Pytorch (Paszke et al., 2019). Our models were optimized using a gradient method with the *Adam* (Kingma and Ba, 2015) optimizer.

Baselines: We compared our proposed method to the Wasserstein-loss-based LDL framework (Frogner et al., 2015; Zhao and Zhou, 2018) and a multi-class KL loss mentioned in (3). Notably, in the original paper (Zhao and Zhou, 2018), they did not include KL loss and used only Wasserstein loss, but (Frogner et al., 2015) used a linear combination of KL divergence and Wasserstein distance as the loss. To ensure fair comparison, we also report the combination of Wasserstein loss and multi-class KL loss as a strong baseline. Therefore, we set the combination parameter $\lambda = \{0, \frac{1}{2}, 1\}$ defined in Eq 2 and the weight of all edges to 1. The Wasserstein loss was computed using the Sinkhorn-Knopp algorithm in the log domain (Schmitzer, 2019; Peyré and Cuturi, 2018) on GPUs. For the proposed method, we computed the tree-Wasserstein loss on the CPU and then passed it to the GPU to compute the gradient. Then, we set the number of iterations of the Sinkhorn-Knopp algorithm to 10 and the regularization parameter to 50, respectively.

5.1 Synthetic data

We generated a synthetic dataset that comprises pairs of a real vector and a target probability distribution on the nodes of a randomly generated tree. This dataset was created as follows: First, we defined a parametric distribution on a graph. Given a graph, $G = (V, E)$, the shortest path metric, d_G , and the probability distribution, $F_{vu\sigma}$, over V parameterized by $v, u \in V, \sigma > 0$ is defined as:

$$F_{vu\sigma}(s) = \frac{1}{C} \left(\exp \frac{d_G(v, s)}{\sigma^2} + \exp \frac{d_G(u, s)}{\sigma^2} \right)$$

$$C = \sum_{s \in V} \left(\exp \frac{d_G(v, s)}{\sigma^2} + \exp \frac{d_G(u, s)}{\sigma^2} \right).$$

Algorithm 1 shows the algorithm used to generate the dataset used in the experiments. In this experiment, we prepared datasets with the distribution on a random tree with 1000 nodes using NetworkX

(Hagberg et al., 2008). The size of each of the training and testing datasets is 1000. We set the number of epochs to 500 and the batch size to 10, and we fixed the learning rate at .001. We reported the average scores of the experiments using 10 different random seeds.

Predictive model: We adopted the following model for class ℓ :

$$h_{\theta}(\mathbf{x})^{(\ell)} = \frac{\exp(\mathbf{w}_{\ell}^{\top} \mathbf{x} + b_{\ell})}{\sum_j \exp(\mathbf{w}_j^{\top} \mathbf{x} + b_j)},$$

where \mathbf{w}_i, b_i are learnable parameters.

Evaluation Metric: To evaluate predictions from various perspectives, we used the metric listed in Table 3. Notably we adopted the **exact** Wasserstein distance, called *Wasserstein*, between the prediction and ground-truth label distributions to assess the extent to which the ground metric was considered in the prediction. In these experiments, we used the Python Optimal Transport (POT) library (Flamary and Courty, 2017) to calculate the exact Wasserstein distance, and the weights of all the edges were set to 1. The other evaluation metrics are the same as those used in (Geng, 2016).

The scores of the experiment with synthetic data are presented in Table 1. The proposed linear combinations of \mathcal{KL} and \mathcal{TW} outperformed the others in terms of *Wasserstein* and *Chebyshev* metric, but they performed poorly in terms of the other metrics.

5.2 BlurbGenreCollectionEN

In this study, we used the BlurbGenreCollectionEN¹ (Cortes and Vapnik, 1995; Lewis et al., 2004) dataset for performing experiments with real data. It comprises advertising descriptions of books from the Penguin Random House webpage. Each instance has one or multiple labels that are hierarchically structured. Because the hierarchical structure of these data is a *forest* and not a *tree*, we added a root node to the hierarchical tree. Of the total 91,892 data samples 64%, 16% and 20% were used in the train, validation, and test sets, respectively. We set the number of epochs to 100 and the batch size to 100, and we fixed the learning rate to .001. We reported the average scores and standard deviations of the experiments using 10 different random seeds.

¹<https://www.inf.uni-hamburg.de/en/inst/ab/lt/resources/data/blurb-genre-collection.html>

Canberra	$\sum_{\ell=1}^L \frac{ \mathbf{h}_{\theta}(\mathbf{x})^{(\ell)} - \mathbf{p}_{\mathbf{y}}^{(\ell)} }{\mathbf{h}_{\theta}(\mathbf{x})^{(\ell)} + \mathbf{p}_{\mathbf{y}}^{(\ell)}}$
Chebyshev	$\max_i \mathbf{h}_{\theta}(\mathbf{x})^{(\ell)} - \mathbf{p}_{\mathbf{y}}^{(\ell)} $
Clark	$\sqrt{\sum_{\ell=1}^L \frac{(\mathbf{h}_{\theta}(\mathbf{x})^{(\ell)} - \mathbf{p}_{\mathbf{y}}^{(\ell)})^2}{(\mathbf{h}_{\theta}(\mathbf{x})^{(\ell)} + \mathbf{p}_{\mathbf{y}}^{(\ell)})^2}}$
Cosine	$\frac{\sum_{\ell=1}^L \mathbf{h}_{\theta}(\mathbf{x})^{(\ell)} \mathbf{p}_{\mathbf{y}}^{(\ell)}}{\sqrt{\sum_{\ell=1}^L (\mathbf{h}_{\theta}(\mathbf{x})^{(\ell)})^2} \sqrt{\sum_{\ell=1}^L (\mathbf{p}_{\mathbf{y}}^{(\ell)})^2}}$
Intersection	$\sum_{\ell=1}^L \min(\mathbf{h}_{\theta}(\mathbf{x})^{(\ell)}, \mathbf{p}_{\mathbf{y}}^{(\ell)})$
Kullback-Leibler	$\sum_{\ell=1}^L \mathbf{p}_{\mathbf{y}}^{(\ell)} \ln \frac{\mathbf{p}_{\mathbf{y}}^{(\ell)}}{\mathbf{h}_{\theta}(\mathbf{x})^{(\ell)}}$

Table 3: Evaluation metrics for LDL. $\mathbf{h}_{\theta}(\mathbf{x})$ is the predicted distribution of \mathbf{x} , and $\mathbf{p}_{\mathbf{y}}$ is the ground truth distribution of a label \mathbf{y} .

Predictive model: We adopted a *long-short-term-memory* (LSTM) (Hochreiter and Schmidhuber, 1997) model with a hidden state size of 200. Because LSTM can efficiently learn long-term dependencies of time-series data, it has often been used in the natural-language processing domain (Yin et al., 2017; Kuncoro et al., 2018). Additionally, we used *fastText* (Bojanowski et al., 2017; Joulin et al., 2017) for word embeddings. A fully connected layer exists before the output layer, and the output function is a softmax function.

Evaluation metric: We evaluated prediction accuracy using three metrics, namely pseudo-recall, top- k cost, and receiver operating characteristic area under the curve (ROC-AUC). Pseudo-recall is defined as $\frac{|\mathcal{P} \cup \mathcal{L}|}{|\mathcal{L}|}$, where \mathcal{L} denotes the set of ground-truth labels, and \mathcal{P} is a set that comprises $L = |\mathcal{L}|$ labels in descending order of the probability score.

Top- k cost is defined as:

$$\frac{1}{K} \sum_{k=1}^K \min_{\ell \in \mathcal{L}} d(\ell_{p_k}, \ell),$$

where ℓ_{p_k} denotes the label with the k -th highest probability score. This metric measures how close the predicted top- k labels are to the ground-truth labels. We calculate ROC-AUC using the output distribution of each model as a score vector, which is assigned 1 on the ground truth labels or 0 on the other labels. Table 2 presents the comparison results. Both regularization terms (\mathcal{W}^1 and \mathcal{TW}) did not have a significant impact on the results.

5.3 Computational-efficiency comparison

In the computational efficiency experiment, distributions with 10^2 , 10^3 , 10^4 , and 10^5 supports were prepared. Subsequently, the computation time and memory required to calculate the loss of pairs of

Algorithm 1: Generating a synthetic dataset

```

1 Generate a random tree :  $G = (V, E)$ ,
   where  $V = \{s_1, \dots, s_l\}$ 
2  $W_1 \leftarrow (n \times m)$ -dim random matrix
3  $W_2 \leftarrow (m \times (l + 1))$ -dim random matrix
4 for  $i = 1$  to  $N$  do
5    $x_i \leftarrow n$ -dimensional random vector
6    $x_i \leftarrow \frac{1}{1 + \exp(-W_1 x_i)}$ 
7    $x_i \leftarrow \frac{1}{1 + \exp(-W_2 x_i)}$ 
8    $\sigma \leftarrow 10x_i^{(l+1)}$ 
9    $j \leftarrow \operatorname{argmax}_{1 \leq j \leq l} x_i^{(j)}$ ;  $v \leftarrow s_j$ 
10   $k \leftarrow \operatorname{argmin}_{1 \leq k \leq l} x_i^{(k)}$ ;  $u \leftarrow s_k$ 
11   $p_G(s) \leftarrow F_{vu\sigma}(s), \forall s \in V$ 
12 return  $\{(x_i, p_G(V))\}_{i=1}^N$ 

```

Table 4: Comparison of computational efficiency.

L	Loss	Time(s)	Memory
10^2	\mathcal{TW}	0.0024	1.58 MB
	\mathcal{W}^1 with GPU	0.0062	3.32 MB
	\mathcal{W}^1 with CPU	0.0528	2.98 MB
10^3	\mathcal{TW}	0.0126	2.44 MB
	\mathcal{W}^1 with GPU	0.0071	16.94 MB
	\mathcal{W}^1 with CPU	0.1279	7.08 MB
10^4	\mathcal{TW}	0.1204	9.82 MB
	\mathcal{W}^1 with GPU	0.5277	766.88 MB
	\mathcal{W}^1 with CPU	25.7985	1148.22 MB
10^5	\mathcal{TW}	1.6454	66.00 MB
	\mathcal{W}^1 with GPU	-	(37.25 GB)
	\mathcal{W}^1 with CPU	-	(40.00 GB)

random probability distributions on the supports were measured. To avoid calculating a shortest-path distance matrix, we used the matrix $(\mathbf{1}\mathbf{1}^\top - \mathbb{I})$, where \mathbb{I} denotes an identity matrix, as the distance matrix while computing the Wasserstein loss. Additionally, we used a random tree, with edge weights of 1, as a tree metric while computing the tree-Wasserstein loss. We report the average scores of three measurements.

Table 4 presents the time and memory required to calculate the losses for various numbers of nodes. \mathcal{TW} outperforms the other Wasserstein losses in terms of computation time and is significantly superior in terms of memory consumption. Although \mathcal{W}^1 that uses a GPU is faster than the others with 10^3 supports, it cannot calculate the loss with 10^5 supports because the required memory cannot be allocated.

6 Conclusions

This study proposed the use of a tree-Wasserstein regularizer for learning. The experimental results indicate that our proposed method can successfully predict the distributions of structured labels and that it outperforms existing Wasserstein loss calculation methods in terms of both computational speed and memory consumption.

Acknowledgments

This work was supported by the JSPS KAKENHI Grant Number 20H04243 and 20H04244. This work was also supported by JST, ACT-X Grant Number JPMJAX200S, Japan.

References

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 214–223.
- Siddhartha Banerjee, Cem Akkaya, Francisco Perez-Sorrosal, and Kostas Tsioutsoulouklis. 2019. Hierarchical transfer learning for multi-label text classification. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, volume 1, pages 6295–6300.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2019. Large-scale multi-label text classification on EU legislation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, volume 1, pages 6314–6322. Association for Computational Linguistics.
- Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning*, 20(3):273–297.
- Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. 2017. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1853–1865.
- Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300.
- Bharath Bhushan Damodaran, Rémi Flamary, Vivien Seguy, and Nicolas Courty. 2020. An entropic optimal transport loss for learning deep neural networks under label noise in remote sensing images. *Computer Vision and Image Understanding*, 191:102863.
- Steven N Evans and Frederick A Matsen. 2012. The phylogenetic kantorovich–rubinstein metric for environmental sequence samples. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):569–592.
- R’emi Flamary and Nicolas Courty. 2017. POT python optimal transport library. Web: <https://pythonot.github.io/>.
- Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. 2015. Learning with a wasserstein loss. In *Advances in Neural Information Processing Systems*, pages 2053–2061.
- Bin-Bin Gao, Chao Xing, Chen-Wei Xie, Jianxin Wu, and Xin Geng. 2017. Deep label distribution learning with label ambiguity. *IEEE Transactions on Image Processing*, 26(6):2825–2838.
- Xin Geng. 2016. Label distribution learning. *IEEE Transactions on Knowledge and Data Engineering*, 28(7):1734–1748.
- Xin Geng and Yu Xia. 2014. Head pose estimation based on multivariate label distribution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1837–1842.
- Xin Geng, Chao Yin, and Zhi-Hua Zhou. 2013. Facial age estimation by learning from label distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(10):2401–2412.
- Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. 2008. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference*, pages 11–15.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, volume 2, pages 427–431.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*.
- Adhiguna Kuncoro, Chris Dyer, John Hale, Dani Yogatama, Stephen Clark, and Phil Blunsom. 2018. LSTMs can learn syntax-sensitive dependencies well, but modeling structure makes them better. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1426–1436.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 957–966.

- Tam Le, Makoto Yamada, Kenji Fukumizu, and Marco Cuturi. 2019. Tree-sliced variants of wasserstein distances. In *Advances in Neural Information Processing Systems*, pages 12283–12294.
- David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. 2004. Rcv1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research*, 5(Apr):361–397.
- 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. 2019. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, pages 8024–8035. Curran Associates, Inc.
- Gabriel Peyré and Marco Cuturi. 2018. Computational optimal transport. *arXiv preprint arXiv:1803.00567*.
- Ryoma Sato, Makoto Yamada, and Hisashi Kashima. 2020. Fast unbalanced optimal transport on tree. In *Advances in Neural Information Processing Systems*.
- Bernhard Schmitzer. 2019. Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM Journal on Scientific Computing*, 41(3):A1443–A1481.
- Cédric Villani. 2009. *Optimal transport: old and new*. Grundlehren der mathematischen Wissenschaften. Springer.
- Wenpeng Yin, Katharina Kann, Mo Yu, and Hinrich Schütze. 2017. Comparative study of cnn and rnn for natural language processing. *arXiv preprint arXiv:1702.01923*.
- Peng Zhao and Zhi-Hua Zhou. 2018. Label distribution learning by optimal transport. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 4506–4513.
- Deyu Zhou, Xuan Zhang, Yin Zhou, Quan Zhao, and Xin Geng. 2016. Emotion distribution learning from texts. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 638–647.