

AN EMPIRICAL STUDY OF SIMPLICIAL REPRESENTATION LEARNING WITH WASSERSTEIN DISTANCE

Anonymous authors

Paper under double-blind review

ABSTRACT

In this paper, we delve into the problem of simplicial representation learning utilizing the 1-Wasserstein distance on a tree structure (a.k.a., Tree-Wasserstein distance (TWD)), where TWD is defined as the L1 distance between two tree-embedded vectors. Specifically, we consider a framework for simplicial representation estimation employing a self-supervised learning approach based on SimCLR with a negative TWD as a similarity measure. In SimCLR, the cosine similarity with real-vector embeddings is often utilized; however, it has not been well studied utilizing L1-based measures with simplicial embeddings. A key challenge is that training the L1 distance is numerically challenging and often yields unsatisfactory outcomes, and there are numerous choices for probability models. Thus, this study empirically investigates a strategy for optimizing self-supervised learning with TWD and find a stable training procedure. More specifically, we evaluate the combination of two types of TWD (total variation and ClusterTree) and several simplicial models including the softmax function, the ArcFace probability model (Deng et al., 2019), and simplicial embedding (Lavoie et al., 2022). Moreover, we propose a simple yet effective Jeffrey divergence-based regularization method to stabilize the optimization. Through empirical experiments on STL10, CIFAR10, CIFAR100, and SVHN, we first found that the simple combination of softmax function and TWD can obtain significantly lower results than the standard SimCLR (non-simplicial model and cosine similarity). We found that the model performance depends on the combination of TWD and the simplicial model, and the Jeffrey divergence regularization usually helps model training. Finally, we inferred that the appropriate choice of combination of TWD and simplicial models outperformed cosine similarity based representation learning.

1 INTRODUCTION

Unsupervised learning is a widely studied topic, which includes the autoencoder (Kramer, 1991) and variational autoencoders (VAE) (Kingma & Welling, 2013). Self-supervised learning algorithms including SimCLR (Chen et al., 2020b), Bootstrap Your Own Latent (BYOL) (Grill et al., 2020) and MoCo (Chen et al., 2020b; He et al., 2020) can also be regarded as unsupervised learning methods. The key idea of a self-supervised algorithm is to adopt contrastive learning, in which two data points are systematically generated from a common data source, and lower representations are found by maximizing the similarity of these two generated input vectors while minimizing the similarity between a target and negative samples. For example, we generate two images by independently applying image transformation such as rotation and cropping.

In terms of contrastive learning loss functions, the InfoNCE (Oord et al., 2018) is widely utilized. It is a powerful pretraining method because self-supervised learning does not require label information and can be trained with several data points. Most previous studies have focused on obtaining real-valued vector representations. However, a few studies have learned lower representations on a simplex (i.e., simplicial representation), including simplicial embedding (Lavoie et al., 2022).

Wasserstein distance, a widely adopted optimal transport-based distance for measuring distributional difference, is useful in various machine learning tasks including generative adversarial networks (Arjovsky et al., 2017), document classification (Kusner et al., 2015; Sato et al., 2022), and image matching (Liu et al., 2020; Sarlin et al., 2020), to name a few. The Wasserstein distance is also

known as the earth mover’s distance (EMD) in the computer vision community and the word mover’s distance (WMD) in the natural language processing community.

In this paper, we consider the problem of learning a simplicial representation by utilizing a self-supervised framework with 1-Wasserstein distance with tree metric (i.e., Tree-Wasserstein Distance (TWD)). TWD includes the sliced Wasserstein distance and total variation as special cases, and it can be represented by the L1 distance between tree-embedded vectors. Because TWD is given as a non-differentiable function, learning simplicial representation utilizing backpropagation of TWD is challenging. Moreover, there exist several probability representations exist, and it is difficult to determine which one is more suitable for simplicial representation learning. Hence, we investigate a combination of probability models and the structure of TWD. More specifically, we consider the total variation and the ClusterTree for TWD, where we show that the total variation is equivalent to the robust variant of TWD. Then, we propose the utilization of softmax, ArcFace based probability model (Deng et al., 2019), and the simplicial embedding (SEM) (Lavoie et al., 2022). Moreover, to stabilize the training, we propose a Jeffrey divergence based regularization. Through self-supervised learning experiments, we found that the standard softmax formulation with backpropagation yield poor results. For total variation, the combination of the ArcFace-based model performs well. In contrast, SEM is suitable for ClusterTree, whereas the ArcFace-based models achieve modest performance. Moreover, the proposed regularization significantly outperforms when compared to its non-regularized counterparts.

Contribution: The contributions are summarized below:

- We investigated the combination of probability models and TWD (total variation and ClusterTree). Then, we found that the ArcFace model with prior information is suited for total variation, while SEM (Lavoie et al., 2022) is suited for ClusterTree models.
- We proposed a robust variant of TWD (RTWD) and showed that RTWD is equivalent to total variation.
- We proposed the Jeffrey divergence regularization for TWD minimization; the regularization significantly stabilized the training.
- We demonstrated that the combination of TWD and simplicial models can obtain better performance in self-supervised training for CIFAR10, STL10, and SVHN compared to the cosine similarity and real-valued representation combination, while the performance of CIFAR100 can be improved further in future.

2 RELATED WORK

The proposed method involved unsupervised representation learning and optimal transportation.

Unsupervised Representation Learning: Representation learning is an important research topic in machine learning and involves several methods. The autoencoder (Kramer, 1991) and its variational version (Kingma & Welling, 2013) are widely employed in unsupervised representation learning methods. Current mainstream self-supervised learning (SSL) approaches are based on a cross-view prediction framework (Becker & Hinton, 1992), and contrastive learning has emerged as a prominent self-supervised learning paradigm.

In SSL, a model learns by contrasting positive samples (similar instances) with negative samples (dissimilar instances), utilizing methods such as SimCLR (Chen et al., 2020a), BYOL (Grill et al., 2020), and MoCo (Chen et al., 2020b; He et al., 2020). These methods employ data augmentation and similarity metrics to encourage the model to project similar instances close together while pushing dissimilar ones apart. This approach has demonstrated effectiveness across various domains, including computer vision and natural language processing, thereby enabling learning without explicit labels. SimCLR employs InfoNCE loss (Oord et al., 2018). After SimCLR has proposed, several alternative approaches have been proposed, such as Barlow’s twin (Zbontar et al., 2021) and BYOL (Grill et al., 2020). Barlow Twins loss function attempts to maximize the correlation between positive pairs while minimizing the cross-correlation with negative samples. Barlow Twins is closely related to the Hilbert–Schmidt independence criterion, a kernel-based independence measure (Gretton et al., 2005; Tsai et al., 2021).

These approaches mainly focus on learning lower-dimensional representations with real-valued vector embeddings, and cosine similarity is adopted as a similarity measure in contrastive learning. Recently, Lavoie et al. (2022) proposed the simplicial embedding (SEM), which comprises multiple concatenated softmax functions and learns high-dimensional sparse nonnegative representations. This significantly improves the classification accuracy.

Divergence and optimal transport: Measuring the divergence between two probability distributions is a fundamental research problem in machine learning, and it has been utilized for various computer vision and natural language processing applications, including document classification (Kusner et al., 2015; Sato et al., 2022) and image matching (Sarlin et al., 2020). One widely adopted divergence measure is Kullback–Leibler (KL) divergence (Cover & Thomas, 2012). However, KL divergence can diverge to infinity when the supports of the two input probability distributions do not overlap.

The Wasserstein distance, or, as is known in the computer vision community, the EMD, can handle the difference in supports in probability distributions. The Wasserstein distance was computed by linear programming. Moreover, in addition to handling different supports, one of the key advantages of Wasserstein distance over KL is that it can identify matches between data samples. For example, Sarlin et al. (2020) proposed SuperGlue, leveraging optimal transport for correspondence determination in local feature sets. Semantic correspondence investigated optimal transport by Liu et al. (2020). In NLP, Kusner et al. (2015) introduced the WMD, a Wasserstein distance pioneer in textual similarity tasks, widely utilized, including text generation evaluation (Zhao et al., 2019). Sato et al. (2022) further studied the property of WMD thoroughly. Another interesting approach is the word rotator distance (WRD) (Yokoi et al., 2020), which utilizes the normalized norm of word vectors as a simplicial representation and significantly improves WMD’s performance. However, these methods incur cubic-order computational costs, rendering them unsuitable for extensive distribution comparison tasks.

To speed up EMD and Wasserstein distance computation, Cuturi (2013) introduced the Sinkhorn algorithm, solving the entropic regularized optimization problem and achieving quadratic order Wasserstein distance computation ($O(\bar{n}^2)$), where \bar{n} is the number of data points. Moreover, because the optimal solution of the Sinkhorn algorithm can be obtained by an iterative algorithm, it can easily be incorporated into deep-learning applications, which makes optimal transport widely applicable to deep-learning applications. One limitation of the Sinkhorn algorithm is that it still requires quadratic-time computation, and the final solution depends highly on the regularization parameter.

An alternative approach is the sliced Wasserstein distance (SWD) (Rabin et al., 2011; Kolouri et al., 2016), which solves the optimal transport problem within a projected one-dimensional subspace. Given 1D optimal transport’s sortable nature, SWD offers $O(\bar{n} \log \bar{n})$ computations. SWD’s extensions include the generalized sliced Wasserstein distance for multidimensional cases (Kolouri et al., 2019), the max-sliced Wasserstein distance, which determines the optimal transport-enhancing 1D subspace (Mueller & Jaakkola, 2015; Deshpande et al., 2019), and the subspace-robust Wasserstein distance (Paty & Cuturi, 2019).

The 1-Wasserstein distance with a tree metric (also known as the Tree-Wasserstein Distance (TWD)) is a generalization of the sliced Wasserstein distance and total variation (Indyk & Thaper, 2003; Evans & Matsen, 2012; Le et al., 2019). Moreover, TWD is also known as the UniFrac distance (Lozupone & Knight, 2005), which is assumed to have a phylogenetic tree beforehand. An important property of TWD is that TWD has an analytical solution with the L1 distance of tree-embedded vectors. Originally, TWD was studied in theoretical computer science, known as the QuadTree algorithm (Indyk & Thaper, 2003), and was recently extended to unbalanced TWD (Sato et al., 2020; Le & Nguyen, 2021), supervised Wasserstein training (Takezawa et al., 2021), and tree barycenters (Takezawa et al., 2022) in the ML community. These approaches focused on approximating the 1-Wasserstein distance through tree construction, often utilizing constant-edge weights. Backurs et al. (2020) introduced FlowTree, amalgamating QuadTree and cost matrix methods, outperforming QuadTree. They guaranteed that QuadTree and FlowTree approximated the nearest neighbors employing the 1-Wasserstein distance. Dey & Zhang (2022) proposed L1-embedding for approximating 1-Wasserstein distance for persistence diagrams. Yamada et al. (2022) proposed a computationally efficient tree weight estimation technique for TWD and empirically demonstrated that

TWD can have comparable performance to the Wasserstein distance while the computational speed is several orders of magnitude faster than the original Wasserstein distance.

Most existing studies on TWD focus on tree construction (Indyk & Thaper, 2003; Le et al., 2019; Takezawa et al., 2021) and edge weight estimation (Yamada et al., 2022). Takezawa et al. (2022) proposed a Barycenter method based on TWD where the set of simplicial representations are given. Frogner et al. (2015) considered utilizing the Wasserstein distance for multilabel classification. Moreover, Toyokuni et al. (2021) utilized TWD for multilabel classification. These studies focused on supervised learning employing softmax as the probability model. Here, we investigated the Wasserstein distance employing a self-supervised learning framework and evaluated various probability models.

3 BACKGROUND

3.1 SELF-SUPERVISED LEARNING WITH SIMCLR

Let us denote $\mathbf{x} \in \mathbb{R}^d$ as an input vector, and let us have n input vectors $\{\mathbf{x}_i\}_{i=1}^n$. Now we define the data transformation functions $\mathbf{u}^{(1)} = \phi_1(\mathbf{x}) \in \mathbb{R}^d$ and $\mathbf{u}^{(2)} = \phi_2(\mathbf{x}) \in \mathbb{R}^d$, respectively. The neural network model is denoted by $\mathbf{f}_\theta(\mathbf{u}) \in \mathbb{R}^{d_{\text{out}}}$, where θ is a learning parameter. Then, SimCLR attempts to train the model by closing $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ to be close, where $\mathbf{u}^{(1)}$ and/or $\mathbf{u}^{(2)}$ are dissimilar to \mathbf{u}' , which is a negative sample.

To this end, the infoNCE (Oord et al., 2018) is employed in the SimCLR model:

$$\ell_{\text{infoNCE}}(\mathbf{z}_i, \mathbf{z}_j) = -\log \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau)}{\sum_{k=1}^{2R} \delta_{k \neq i} \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_k)/\tau)},$$

where R is the batch size and $\text{sim}(\mathbf{u}, \mathbf{u}')$ is a similarity function that takes a higher positive value if \mathbf{u} and \mathbf{u}' are similar, whereas it takes a smaller positive value or negative value when \mathbf{u} and \mathbf{u}' are dissimilar, τ is the temperature parameter, and $\delta_{k \neq i}$ is a delta function that takes a value of 1 when $k \neq i$ and value 0 otherwise. In SimCLR, the parameter is learned by minimizing the infoNCE loss. Hence, the first term of the infoNCE loss makes \mathbf{z}_i and \mathbf{z}_j similar. The second term is a log sum exp function and is a softmax function. Because we attempt to minimize the maximum similarity between the input \mathbf{z}_i and its negative samples, we can make \mathbf{z}_i and its negative samples dissimilar utilizing the second term.

$$\hat{\theta} := \underset{\theta}{\text{argmin}} \sum_{i=1}^n \ell_{\text{infoNCE}}(\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}).$$

where $\mathbf{z}_i^{(1)} = \mathbf{f}_\theta(\mathbf{u}_i^{(1)})$ and $\mathbf{z}_i^{(2)} = \mathbf{f}_\theta(\mathbf{u}_i^{(2)})$, respectively.

3.2 p -WASSERSTEIN DISTANCE

The p -Wasserstein distance between two discrete measures $\mu = \sum_{i=1}^{\bar{n}} a_i \delta_{\mathbf{x}_i}$ and $\mu' = \sum_{j=1}^{\bar{m}} a'_j \delta_{\mathbf{y}_j}$ is given by

$$\mathcal{W}_p(\mu, \mu') = \left(\min_{\Pi \in U(\mu, \mu')} \sum_{i=1}^{\bar{n}} \sum_{j=1}^{\bar{m}} \pi_{ij} d(\mathbf{x}_i, \mathbf{y}_j)^p \right)^{1/p},$$

where $U(\mu, \mu')$ denotes the set of the transport plans and $U(\mu, \mu') = \{\Pi \in \mathbb{R}_+^{\bar{n} \times \bar{m}} : \Pi \mathbf{1}_{\bar{m}} = \mathbf{a}, \Pi^\top \mathbf{1}_{\bar{n}} = \mathbf{a}'\}$. The Wasserstein distance is computed utilizing a linear program. However, because it includes an optimization problem, the computation of Wasserstein distance for each iteration is computationally expensive.

One of the alternative approaches is to use the entropic regularization (Cuturi, 2013). If we consider the 1-Wasserstein distance, the entropic regularized variant is given as

$$\min_{\Pi \in U(\mathbf{a}, \mathbf{a}')} \sum_{i=1}^{\bar{n}} \sum_{j=1}^{\bar{m}} \pi_{ij} d(\mathbf{x}_i, \mathbf{y}_j) + \lambda \sum_{i=1}^{\bar{n}} \sum_{j=1}^{\bar{m}} \pi_{ij} (\log(\pi_{ij}) - 1).$$




Figure 1: Left tree corresponds to the total variation if we set the weight as $w_i = \frac{1}{2}, \forall i$. Right tree is a ClusterTree (2 class).

This optimization problem can be solved efficiently by using the Sinkhorn algorithm at a computational cost of $O(\bar{n}\bar{m})$. More importantly, the solution of the Sinkhorn algorithm is given as a series of matrix multiplications; the Sinkhorn algorithm is widely employed for deep learning algorithms.

3.3 1-WASSERSTEIN DISTANCE WITH TREE METRIC (TREE-WASSERSTEIN DISTANCE)

Another 1-Wasserstein distance is based on trees (Indyk & Thaper, 2003; Le et al., 2019). The 1-Wasserstein distance with tree metric is defined as the L1 distance between the two probability distributions $\mu = \sum_j a_j \delta_{x_j}$ and $\mu' = \sum_j a'_j \delta_{x_j}$:

$$W_{\mathcal{T}}(\mu, \mu') = \|\text{diag}(\mathbf{w})\mathbf{B}(\mathbf{a} - \mathbf{a}')\|_1 = \|\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a} - \text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}'\|_1,$$

where $\mathbf{B} \in \{0, 1\}^{N_{\text{node}} \times N_{\text{leaf}}}$ and $\mathbf{w} \in \mathbb{R}_+^{N_{\text{node}}}$. Here, we consider a tree with a depth of one and a ClusterTree, as illustrated in Figure 1. If all the edge weights $w_1 = w_2 = \dots = w_N = \frac{1}{2}$ in the left panel of Figure 1, the \mathbf{B} matrix is given as $\mathbf{B} = \mathbf{I}$. Substituting this result into the TWD, we have

$$W_{\mathcal{T}}(\mu, \mu') = \frac{1}{2} \|\mathbf{a} - \mathbf{a}'\|_1 = \|\mathbf{a} - \mathbf{a}'\|_{\text{TV}}.$$

Thus, the total variation is a special case of TWD. In this setting, the shortest path distance in the tree corresponds to the Hamming distance. Note that Raginsky et al. (2013) also assert that the 1-Wasserstein distance with the Hamming metric $d(x, y) = \delta_{x \neq y}$ is equivalent to the total variation (Proposition 3.4.1).

The key advantage of the tree-based approach is that the Wasserstein distance is provided in closed form, which is computationally very efficient. Note that a chain is included as a special case of a tree. Thus, the widely employed sliced Wasserstein distance is included as a special case of TWD. Moreover, it has been empirically reported that TWD and Sinkhorn-based approaches perform similarly in multi-label classification tasks (Toyokuni et al., 2021).

4 SELF-SUPERVISED LEARNING WITH 1-WASSERSTEIN DISTANCE

In this section, we first formulate self-supervised learning utilizing the TWD. We then introduce ArcFace-based probability models and Jeffrey divergence regularization.

4.1 SIMCLR WITH TREE WASSERSTEIN DISTANCE

\mathbf{a} and \mathbf{a}' are the simplicial embedding vectors of x and x' (i.e., $\mathbf{1}^\top \mathbf{a} = 1$ and $\mathbf{1}^\top \mathbf{a}' = 1$) with $\mu = \sum_j a_j \delta_{e_j}$ and $\mu' = \sum_j a'_j \delta_{e_j}$, respectively. Here, e_j is the virtual embedding corresponding to a_j or a'_j , where e is not available in our problem setup in general. Then, we adopt the negative Wasserstein distance between μ and μ' as the similarity score for SimCLR:

$$\text{sim}(\mu, \mu') = -W_{\mathcal{T}}(\mu, \mu') = -\|\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a} - \text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}'\|_1.$$

We cannot determine the tree structure or its weight; we assume that \mathbf{B} and \mathbf{w} are given. We considered the total variation and ClusterTree cases (see Figure 1).

This study attempts to estimate a simplicial representation estimation by adopting self-supervised learning:

$$\hat{\theta} := \operatorname{argmin}_{\theta} \sum_{i=1}^n \left(W_{\mathcal{T}}(\mu_i^{(1)}, \mu_j^{(2)})\tau^{-1} + \log \sum_{k=1}^{2N} \delta_{k \neq i} \exp(-W_{\mathcal{T}}(\mu_i^{(1)}, \mu_j^{(2)})\tau^{-1}) \right),$$

where $\tau > 0$ is the temperature parameter for the InfoNCE loss. We can certainly utilize Barlow Twins instead of infoNCE for simplicial representation estimation.

4.2 ROBUST VARIANT OF TREE WASSERSTEIN DISTANCE

In our setup, it is hard to estimate the tree structure \mathbf{B} and the edge weight \mathbf{w} , because the embedding vectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_{d_{\text{out}}}$ is not available. To deal with the problem, we consider the robust estimation of Wasserstein distance such as the subspace robust Wasserstein distance (SRWD) (Paty & Cuturi, 2019) for TWD. Specifically, with given \mathbf{B} , we propose the robust TWD (RTWD) as follows:

$$\text{RTWD}(\mu, \mu') = \frac{1}{2} \min_{\Pi \in U(\mathbf{a}, \mathbf{a}')} \max_{\mathbf{w} \text{ s.t. } \mathbf{B}^{\top} \mathbf{w} = \mathbf{1} \text{ and } \mathbf{w} \geq 0} \sum_{i=1}^{N_{\text{leafs}}} \sum_{j=1}^{N_{\text{leafs}}} \pi_{ij} d_{\mathcal{T}}(\mathbf{e}_i, \mathbf{e}_j),$$

where $d_{\mathcal{T}}(\mathbf{e}_i, \mathbf{e}_j)$ is the shortest path distance between \mathbf{e}_i and \mathbf{e}_j , where \mathbf{e}_i and \mathbf{e}_j are embedded on a tree \mathcal{T} .

Proposition 1 *The robust variant of TWD (RTWD) is equivalent to total variation:*

$$\text{RTWD}(\mu, \mu') = \|\mathbf{a} - \mathbf{a}'\|_{\text{TV}},$$

where $\|\mathbf{a} - \mathbf{a}'\|_{\text{TV}} = \frac{1}{2} \|\mathbf{a} - \mathbf{a}'\|_1$ is the total variation.

Based on this proposition, the RTWD is equivalent to a total variation and it does not depend on the tree structure \mathbf{B} . That is, if we do not have a prior information about tree structure, using the total variation is a reasonable choice.

4.3 PROBABILITY MODELS

Several probability models were adopted here.

Softmax: The softmax function for simplicial representation is given by

$$\mathbf{a}_{\theta}(\mathbf{x}) = \text{Softmax}(\mathbf{f}_{\theta}^{(\ell)}(\mathbf{x})),$$

where $\mathbf{f}_{\theta}(\mathbf{x}_i)$ is a neural network model.

Simplicial Embedding: Lavoie et al. (2022) proposed a simple yet efficient simplicial embedding method. We assume that the output dimensionality of the output of the neural network model is d_{out} . Then, SEM applies the softmax function to each V -dimensional vector of $\mathbf{f}_{\theta}(\mathbf{x})$, where we have $L = d_{\text{out}}/V$ probability vectors. The ℓ -th softmax function is defined as follows.

$$\mathbf{a}_{\theta}(\mathbf{x}) = [\mathbf{a}_{\theta}^{(1)}(\mathbf{x})^{\top}, \mathbf{a}_{\theta}^{(2)}(\mathbf{x})^{\top}, \dots, \mathbf{a}_{\theta}^{(L)}(\mathbf{x})^{\top}]^{\top} \text{ with } \mathbf{a}_{\theta}^{(\ell)}(\mathbf{x}) = \text{Softmax}(\mathbf{f}_{\theta}^{(\ell)}(\mathbf{x}))/L,$$

where $\mathbf{f}_{\theta}^{(\ell)}(\mathbf{x}) \in \mathbb{R}^V$ is the ℓ -th block of a neural network model. We normalize the softmax function by L because $\mathbf{a}_{\theta}(\mathbf{x})$ must satisfy the sum to one constraint for Wasserstein distance.

ArcFace model (AF): We proposed employing the ArcFace probability model (Deng et al., 2019).

$$\mathbf{a}_{\theta}(\mathbf{x}) = \text{Softmax}(\mathbf{K}^{\top} \mathbf{f}_{\theta}(\mathbf{x})/\eta) = \left(\frac{\exp(\mathbf{k}_1^{\top} \mathbf{f}_{\theta}(\mathbf{x})/\eta)}{\sum_{k=1}^{d_{\text{out}}} \exp(\mathbf{k}_k^{\top} \mathbf{f}_{\theta}(\mathbf{x})/\eta)}, \dots, \frac{\exp(\mathbf{k}_{d_{\text{out}}}^{\top} \mathbf{f}_{\theta}(\mathbf{x})/\eta)}{\sum_{k=1}^{d_{\text{out}}} \exp(\mathbf{k}_k^{\top} \mathbf{f}_{\theta}(\mathbf{x})/\eta)} \right),$$

where $\mathbf{K} = [\mathbf{k}_1, \mathbf{k}_2, \dots, \mathbf{k}_{d_{\text{out}}}] \in \mathbb{R}^{d_{\text{out}} \times d_{\text{prob}}}$ is a learning parameter, $\mathbf{f}_{\theta}(\mathbf{x})$ is a normalized output of a model ($\mathbf{f}_{\theta}(\mathbf{x})^{\top} \mathbf{f}_{\theta}(\mathbf{x}) = 1$), and η is a temperature parameter. Note that AF has a similar structure of attention in transformers (Bahdanau et al., 2014; Vaswani et al., 2017). The key difference from the original attention is normalizing the key matrix \mathbf{K} and the query vector $\mathbf{f}_{\theta}(\mathbf{x})$.

AF with Positional Encoding: The AF can add one more linear layer and then apply the softmax function; the performance can generally be similar to the standard softmax function. Here, we propose replacing the key matrix with a normalized positional encoding matrix $\mathbf{k}_i^\top \mathbf{k}_i = 1, \forall i$:

$$\mathbf{k}_i = \frac{\bar{\mathbf{k}}_i}{\|\bar{\mathbf{k}}_i\|_2}, \quad \bar{\mathbf{k}}_i^{(2j)} = \sin(i/10000^{2j/d_{\text{out}}}), \quad \bar{\mathbf{k}}_i^{(2j+1)} = \cos(i/10000^{2j/d_{\text{out}}}).$$

AF with Discrete Cosine Transform Matrix: Another natural approach would be to utilize an orthogonal matrix as \mathbf{K} . Therefore, we proposed adopting a discrete cosine transform (DCT) matrix as \mathbf{K} . The DCT matrix is expressed as

$$\mathbf{k}_i^{(j)} = \begin{cases} 1/\sqrt{d_{\text{out}}} & (i = 0) \\ \sqrt{\frac{2}{d_{\text{out}}}} \cos \frac{\pi(2j+1)i}{2d_{\text{out}}} & (1 \leq i \leq d_{\text{out}}) \end{cases}.$$

One of the contributions of this study is that combining positional encoding and DCT matrix with the ArcFace model significantly boosts performance, whereas the standard ArcFace model performs similarly to the simple softmax function.

4.4 JEFFREY-DIVERGENCE REGULARIZATION

We empirically observed that optimizing the loss function above is extremely challenging. One problem is that L1 distance is not differentiable at 0. Another issue is that the softmax function has many degrees of freedom and cannot utilize the cost function in our framework. Figure 2 illustrates the learning curve for the standard optimization adopting the softmax function model.

To stabilize the optimization, we proposed the Jeffrey divergence regularization, which is the upper bound of the square of the 1-Wasserstein distance.

Proposition 2 For $\mathbf{B}^\top \mathbf{w} = \mathbf{1}$ and the probability vectors \mathbf{a}_i and \mathbf{a}_j , we have

$$W_{\mathcal{T}}^2(\mu_i, \mu_j) \leq \text{JD}(\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i \parallel \text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j),$$

where $\text{JD}(\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i \parallel \text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j)$ is a Jeffrey divergence.

This result indicates that minimizing the symmetric KL divergence (i.e., Jeffrey divergence) can minimize the tree-Wasserstein distance. Because the Jeffrey divergence is smooth, computation of the gradient of the upper bound is easier.

We propose the following regularized version of InfoNCE.

$$\begin{aligned} \bar{\ell}_{\text{InfoNCE}}(\mathbf{a}_i^{(1)}, \mathbf{a}_j^{(2)}) &= \tau^{-1} W_{\mathcal{T}}(\mathbf{a}_i^{(1)}, \mathbf{a}_j^{(2)}) + \log \sum_{k=1}^{2N} \delta_{k \neq i} \exp(-\tau^{-1} W_{\mathcal{T}}(\mathbf{a}_i^{(1)}, \mathbf{a}_j^{(2)})) \\ &\quad + \lambda \text{JD}(\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i^{(1)} \parallel \text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j^{(2)}). \end{aligned}$$

Note that Frogner et al. (2015) considered a multilabel classification problem utilizing a regularized Wasserstein loss. They proposed utilizing the Kullback–Leibler divergence-based regularization to stabilize the training. We derive the Jeffrey divergence from the TWD, which includes a simple KL divergence-based regularization as a special case. Moreover, we proposed employing JD regularization for self-supervised learning frameworks, which has not been extensively studied.

5 EXPERIMENTS

This section evaluates SimCLR methods with different probability models utilizing STL10, CIFAR10, CIFAR100, and SVHN datasets.

5.1 SETUP

For all experiments, we employed the Resnet18 model with an output dimension of ($d_{\text{out}} = 256$) and coded all methods based on a standard SimCLR implementation¹. We used the Adam optimizer and

¹<https://github.com/sthalles/SimCLR>

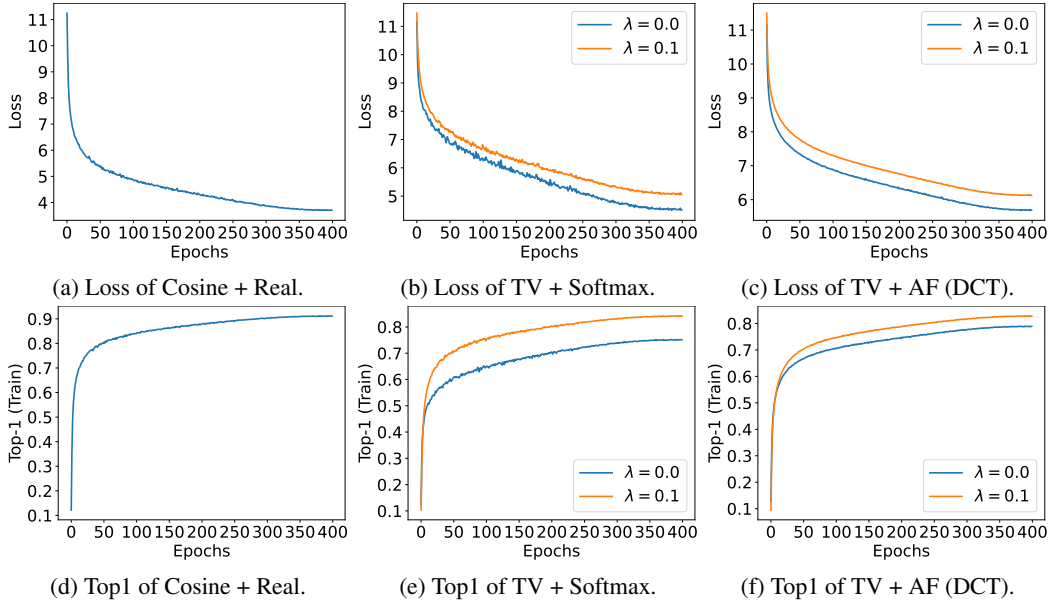


Figure 2: The infoNCE loss and Top-1 (Train) comparisons on STL10 dataset.

set the learning rate as 0.0003, weight decay parameter as $1e-4$, adn temperature τ as 0.07, respectively. For the proposed method, we compared two variants of TWD: total variation and ClusterTree (See Figure 1). As a model, we evaluated the conventional softmax function, attention model (AF), and simplicial embedding (SEM) (Lavoie et al., 2022) and set the temperature parameter $\tau = 0.1$ for all experiments. For SEM, we set $L = 16$ and $V = 16$.

We also evaluated JD regularization, where we set the regularization parameter $\lambda = 0.1$ for all experiments. For reference, we compared the cosine similarity as a similarity function of SimCLR. For all approaches, we utilized the KNN classifier of the scikit-learn package², where the number of nearest neighbor was set to $K = 50$ and utilized the L1 distance for Wasserstein distances and cosine similarity for non-simplicial models. For all experiments, we utilized A6000 GPUs. We ran all experiments three times and reported the average scores.

5.2 PERFORMANCE COMPARISON

Figure 2 illustrates the training loss and top-1 accuracy for three methods: cosine + real-valued embedding, TV + Softmax, and TV + AF (DCT). This experiment revealed that the convergence speed of the loss is nearly identical across all methods. Regarding the training top-1 accuracy, Cosine + Real-valued embedding achieves the highest accuracy, followed by the Softmax function, and AF (DCT) lags. This behavior is expected because real-valued embeddings offer the most flexibility, Softmax comes next, and AF models have the least freedom. For all methods based on Tree Wasserstein Distance (TWD), it is evident that JD regularization significantly aids the training process, especially in the case of the Softmax function. However, for AF (DCT), the improvement was relatively small. This is likely because AF (DCT) can also be considered a form of regularization.

Table 1 presents the experimental results for test classification accuracy using KNN. The first observation is that the simple implementation of the conventional Softmax function performs very poorly (the performance is approximately 10 points lower) compared with the cosine similarity. As expected, AF has just one more layer than the simple Softmax model, performing similarly to Softmax. Compared with Softmax and AF, AF (PE) and AF (DCT) significantly improved the classification accuracy for the total variation and ClusterTree cases. However, for the ClusterTree case, AF (PE) achieved a better classification performance, whereas the AF (DCT) improvement over the softmax model was limited. In the ClusterTree case, SEM significantly improves with a combination of ClusterTree and regularization. Overall, the proposed methods performed better than cosine

²<https://scikit-learn.org/stable/modules/generated/sklearn.neighbors.KNeighborsClassifier.html>

Table 1: KNN classification result with Resnet18 backbone. In this experiment, we set the number of neighbors as $K = 50$ and computed the averaged classification accuracy over three runs. Note that the Wasserstein distance with ($B = I_{d_{out}}$) is equivalent to a total variation.

Similarity Function	simplicial Model	STL10	CIFAR10	CIFAR100	SVHN
Cosine Similarity	N/A	75.77	67.39	32.06	76.35
	Softmax	70.12	63.20	26.88	74.46
	SEM	71.33	61.13	26.08	74.28
	AF (DCT)	72.95	65.92	25.96	76.51
TWD (TV)	Softmax	65.54	59.72	26.07	72.67
	SEM	65.35	56.56	24.31	73.36
	AF	65.61	60.92	26.33	75.01
	AF (PE)	71.71	64.68	26.38	76.44
	AF (DCT)	73.28	67.03	25.85	77.62
	Softmax + JD	72.64	67.08	27.82	77.69
	SEM + JD	71.79	63.60	26.14	75.64
	AF + JD	72.64	67.15	27.45	78.00
	AF (PE) + JD	74.47	67.28	27.01	78.12
	AF (DCT) + JD	76.28	68.60	26.49	79.70
TWD (ClusterTree)	Softmax	69.15	62.33	24.47	74.87
	SEM	72.88	63.82	22.55	77.47
	AF	70.40	63.28	24.28	75.24
	AF (PE)	72.37	65.08	23.33	76.67
	AF (DCT)	71.95	65.89	21.87	77.92
	Softmax + JD	73.52	66.76	24.96	77.65
	SEM + JD	75.93	67.68	22.96	79.19
	AF + JD	73.66	66.61	24.55	77.64
	AF (PE) + JD	73.92	67.00	23.83	77.87
	AF (DCT) + JD	74.29	67.50	22.89	78.31

similarity without real-valued vector embedding when the number of classes was relatively small (i.e., STL10, CIFAR10, and SVHN). In contrast, the performance of the proposed method degraded for CIFAR100, and the results of ClusterTree are particularly poor. Because Wasserstein distance can generally be minimized even if it cannot overfit, it is natural for the Wasserstein distance to make mistakes when the number of classes is large.

Next, we evaluated Jeffrey divergence regularization. Surprisingly, simple regularization dramatically improved the classification performance for all simplicial models. These results also support the idea that the main problem with Wasserstein distance-based representation learning is mainly caused by its numerical problems. Among them, the proposed AF (DCT) + JD (TV) achieves the highest classification accuracy, comparable to the cosine similarity result, and more than ten % in point improvements from the naive implementation with the Softmax function. Moreover, all the simplicial model performances with the cosine similarity combination tended to result in a lower classification error than the combination with TWD and simplicial models. From our empirical study, we propose utilizing TWD (TV) + AF models or TWD (ClusterTree) + SEM for representation learning tasks for simplicial model-based representation learning.

6 CONCLUSION

This study investigated a simplicial representation employing self-supervised learning with TWD. We employed softmax, simplicial embedding (Lavoie et al., 2022), and ArcFace models (Deng et al., 2019). Moreover, to mitigate the intricacies of optimizing the L1 distance, we incorporated an upper bound on the squared 1-Wasserstein distance as a regularization technique. We empirically evaluated benchmark datasets and found that a simple combination of the softmax function and TWD performed poorly. Overall, the combination of ArcFace with DCT outperformed its cosine similarity counterparts. Moreover, we found that the TWD (ClusterTree) and SEM combination yielded favorable performance.

REFERENCES

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.
- Arturs Backurs, Yihe Dong, Piotr Indyk, Ilya Razenshteyn, and Tal Wagner. Scalable nearest neighbor search for optimal transport. In *ICML*, 2020.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Suzanna Becker and Geoffrey E Hinton. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(6356):161–163, 1992.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020a.
- Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NIPS*, 2013.
- Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen Zhao, David Forsyth, and Alexander G Schwing. Max-sliced Wasserstein distance and its use for GANs. In *CVPR*, 2019.
- Tamal K Dey and Simon Zhang. Approximating 1-wasserstein distance between persistence diagrams by graph sparsification. In *ALENEX*, 2022.
- Steven N Evans and Frederick A Matsen. The phylogenetic kantorovich–rubinstein metric for environmental sequence samples. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):569–592, 2012.
- Ky Fan. Minimax theorems. *Proceedings of the National Academy of Sciences*, 39(1):42–47, 1953.
- Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. Learning with a wasserstein loss. *NIPS*, 2015.
- A. Gretton, O. Bousquet, Alex. Smola, and B. Schölkopf. Measuring statistical dependence with Hilbert-Schmidt norms. In *ALT*, 2005.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *NeurIPS*, 2020.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.
- Piotr Indyk and Nitin Thaper. Fast image retrieval via embeddings. In *3rd international workshop on statistical and computational theories of vision*, volume 2, pp. 5. Nice, France, 2003.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Soheil Kolouri, Yang Zou, and Gustavo K Rohde. Sliced wasserstein kernels for probability distributions. In *CVPR*, 2016.
- Soheil Kolouri, Kimia Nadjahi, Umut Simsekli, Roland Badeau, and Gustavo Rohde. Generalized sliced wasserstein distances. In *NeurIPS*, 2019.

- Mark A Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE journal*, 37(2):233–243, 1991.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. From word embeddings to document distances. In *ICML*, 2015.
- Samuel Lavoie, Christos Tsirigotis, Max Schwarzer, Ankit Vani, Michael Noukhovitch, Kenji Kawaguchi, and Aaron Courville. Simplicial embeddings in self-supervised learning and downstream classification. *arXiv preprint arXiv:2204.00616*, 2022.
- Tam Le and Truyen Nguyen. Entropy partial transport with tree metrics: Theory and practice. In *AISTATS*, 2021.
- Tam Le, Makoto Yamada, Kenji Fukumizu, and Marco Cuturi. Tree-sliced approximation of wasserstein distances. *NeurIPS*, 2019.
- Yanbin Liu, Linchao Zhu, Makoto Yamada, and Yi Yang. Semantic correspondence as an optimal transport problem. In *CVPR*, 2020.
- Catherine Lozupone and Rob Knight. Unifrac: a new phylogenetic method for comparing microbial communities. *Applied and environmental microbiology*, 71(12):8228–8235, 2005.
- Jonas W Mueller and Tommi Jaakkola. Principal differences analysis: Interpretable characterization of differences between distributions. *NIPS*, 2015.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- François-Pierre Paty and Marco Cuturi. Subspace robust wasserstein distances. In *ICML*, 2019.
- Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernot. Wasserstein barycenter and its application to texture mixing. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 435–446. Springer, 2011.
- Maxim Raginsky, Igal Sason, et al. Concentration of measure inequalities in information theory, communications, and coding. *Foundations and Trends® in Communications and Information Theory*, 10(1-2):1–246, 2013.
- Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *CVPR*, 2020.
- Ryoma Sato, Makoto Yamada, and Hisashi Kashima. Fast unbalanced optimal transport on tree. In *NeurIPS*, 2020.
- Ryoma Sato, Makoto Yamada, and Hisashi Kashima. Re-evaluating word mover’s distance. *ICML*, 2022.
- Yuki Takezawa, Ryoma Sato, and Makoto Yamada. Supervised tree-wasserstein distance. In *ICML*, 2021.
- Yuki Takezawa, Ryoma Sato, Zornitsa Kozareva, Sujith Ravi, and Makoto Yamada. Fixed support tree-sliced wasserstein barycenter. *AISTATS*, 2022.
- Ayato Toyokuni, Sho Yokoi, Hisashi Kashima, and Makoto Yamada. Computationally efficient Wasserstein loss for structured labels. In *ECAL: Student Research Workshop*, April 2021.
- Yao-Hung Hubert Tsai, Shaojie Bai, Louis-Philippe Morency, and Ruslan Salakhutdinov. A note on connecting barlow twins with negative-sample-free contrastive learning. *arXiv preprint arXiv:2104.13712*, 2021.
- J v. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320, 1928.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NIPS*, 2017.

Makoto Yamada, Yuki Takezawa, Ryoma Sato, Han Bao, Zornitsa Kozareva, and Sujith Ravi. Approximating 1-wasserstein distance with trees. *TMLR*, 2022.

Sho Yokoi, Ryo Takahashi, Reina Akama, Jun Suzuki, and Kentaro Inui. Word rotator’s distance. *EMNLP*, 2020.

Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021.

Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M Meyer, and Steffen Eger. MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance. *EMNLP-IJCNLP*, 2019.

A APPENDIX

A.1 PROOF OF PROPOSITION 1

(Proof) Let $\mathbf{B} \in \{0, 1\}^{N \times N_{\text{leaf}}} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_{N_{\text{leaf}}}]$ and $\mathbf{b}_i \in \{0, 1\}^N$. The shortest path distance between leaves i and j can be represented as (Yamada et al., 2022)

$$d_{\mathcal{T}}(\mathbf{e}_i, \mathbf{e}_j) = \mathbf{w}^\top (\mathbf{b}_i + \mathbf{b}_j - 2\mathbf{b}_i \circ \mathbf{b}_j).$$

$d_{\mathcal{T}}(\mathbf{e}_i, \mathbf{e}_j)$ is represented by a linear function with respect to \mathbf{w} with given \mathbf{B} and the constraint of \mathbf{w} and $\mathbf{\Pi}$ are convex. Thus, we have the following representation by the Minimax theorem (v. Neumann, 1928; Fan, 1953):

$$\begin{aligned} \text{RTWD}(\mu, \mu') &= \frac{1}{2} \max_{\mathbf{w} \text{ s.t. } \mathbf{B}^\top \mathbf{w} = \mathbf{1} \text{ and } \mathbf{w} \geq \mathbf{0}} \min_{\mathbf{\Pi} \in U(\mathbf{a}, \mathbf{a}')} \sum_{i=1}^{N_{\text{leaf}}} \sum_{j=1}^{N_{\text{leaf}}} \pi_{ij} \mathbf{w}^\top (\mathbf{b}_i + \mathbf{b}_j - 2\mathbf{b}_i \circ \mathbf{b}_j) \\ &= \frac{1}{2} \max_{\mathbf{w} \text{ s.t. } \mathbf{B}^\top \mathbf{w} = \mathbf{1} \text{ and } \mathbf{w} \geq \mathbf{0}} \|\text{diag}(\mathbf{w})\mathbf{B}(\mathbf{a} - \mathbf{a}')\|_1, \end{aligned}$$

where we used $\text{TWD}(\mu, \mu') = \min_{\mathbf{\Pi} \in U(\mathbf{a}, \mathbf{a}')} \sum_{i=1}^{N_{\text{leaf}}} \sum_{j=1}^{N_{\text{leaf}}} \pi_{ij} d_{\mathcal{T}}(\mathbf{x}_i, \mathbf{x}_j) = \|\text{diag}(\mathbf{w})\mathbf{B}(\mathbf{a} - \mathbf{a}')\|_1$.

Without loss of generality, we consider $w_0 = 0$. We first re-write the norm $\|\text{diag}(\mathbf{w})\mathbf{B}(\mathbf{a} - \mathbf{a}')\|_1$ as

$$\|\text{diag}(\mathbf{w})\mathbf{B}(\mathbf{a} - \mathbf{a}')\|_1 = \sum_{j=1}^N w_j \left| \sum_{k \in [N_{\text{leaf}}], k \in de(j)} (a_k - a'_k) \right|,$$

where $de(j)$ is the set of descendants of node $j \in [N]$ (including itself). Using the triangle inequality, we have

$$\begin{aligned} \|\text{diag}(\mathbf{w})\mathbf{B}(\mathbf{a} - \mathbf{a}')\|_1 &\leq \sum_{j=1}^N w_j \sum_{k \in [N_{\text{leaf}}], k \in de(j)} |a_k - a'_k| \\ &= \sum_{k \in [N_{\text{leaf}}]} |a_k - a'_k| \sum_{j \in [N], j \in pa(k)} w_j, \end{aligned}$$

where $pa(k)$ is the set of ancestors of leaf k . Re-writing the constraint $\mathbf{B}^\top \mathbf{w} = \mathbf{1}$ as $\sum_{j \in [N], j \in pa(k)} w_j = 1$ for any $k \in [N_{\text{leaf}}]$, we obtain that

$$\|\text{diag}(\mathbf{w})\mathbf{B}(\mathbf{a} - \mathbf{a}')\|_1 \leq \sum_{k \in [N_{\text{leaf}}]} |a_k - a'_k| = \|\mathbf{a} - \mathbf{a}'\|_1.$$

The latter inequality holds for any weight vector \mathbf{w} , therefore, considering the vector such that $w_j = 1$ if $j \in [N_{\text{leaf}}]$ and 0 otherwise, which satisfies the constraint $\mathbf{B}^\top \mathbf{w} = \mathbf{1}$, we obtain

$$\|\text{diag}(\mathbf{w})\mathbf{B}(\mathbf{a} - \mathbf{a}')\|_1 = \sum_{k=1}^{N_{\text{leaf}}} |a_k - a'_k| = \|\mathbf{a} - \mathbf{a}'\|_1.$$

This terminates the proof of this proposition.

A.2 PROOF OF PROPOSITION 2

(Proof) The following holds if $\mathbf{B}^\top \mathbf{w} = \mathbf{1}$ with a probability vector \mathbf{a} (such that $\mathbf{a}^\top \mathbf{1} = 1$).

$$\mathbf{1}^\top \text{diag}(\mathbf{w})\mathbf{B}\mathbf{a} = 1.$$

Then, by using the Pinsker's inequality we derive the following inequalities:

$$W_{\mathcal{T}}(\mu_i, \mu_j) = \|\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i - \text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j\|_1 \leq \sqrt{2\text{KL}(\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i \|\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j)},$$

and

$$W_{\mathcal{T}}(\mu_i, \mu_j) = \|\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j - \text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i\|_1 \leq \sqrt{2\text{KL}(\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j \|\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i)},$$

Thus,

$$W_{\mathcal{T}}^2(\mu_i, \mu_j) \leq \text{KL}(\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i \|\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j) + \text{KL}(\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_j \|\text{diag}(\mathbf{w})\mathbf{B}\mathbf{a}_i)$$