

Unimodal-uniform Constrained Wasserstein Training for Medical Diagnosis

Xiaofeng Liu^{1,2}, Xu Han^{3†}, Yukai Qiao^{4†}, Yi Ge^{2†}, Site Li², Jun Lu^{1*}

¹BIDMC, Harvard Medical School; ²Carnegie Mellon University;

³Johns Hopkins University; ⁴Beijing University of Posts and Telecommunications

†Contribute equally *Corresponding Author: jlu@bidmc.harvard.edu

Abstract

The labels in medical diagnosis task are usually discrete and successively distributed. For example, the Diabetic Retinopathy Diagnosis (DR) involves five health risk levels: no DR (0), mild DR (1), moderate DR (2), severe DR (3) and proliferative DR (4). This labeling system is common for medical disease. Previous methods usually construct a multi-binary-classification task or propose some re-parameter schemes in the output unit. In this paper, we target on this task from the perspective of loss function. More specifically, the Wasserstein distance is utilized as an alternative, explicitly incorporating the inter-class correlations by pre-defining its ground metric. Then, the ground metric which serves as a linear, convex or concave increasing function w.r.t. the Euclidean distance in a line is explored from an optimization perspective.

Meanwhile, this paper also proposes of constructing the smoothed target labels that model the inlier and outlier noises by using a unimodal-uniform mixture distribution. Different from the one-hot setting, the smoothed label endues the computation of Wasserstein distance with more challenging features. With either one-hot or smoothed target label, this paper systematically concludes the practical closed-form solution. We evaluate our method on several medical diagnosis tasks (e.g., Diabetic Retinopathy and Ultrasound Breast dataset) and achieve state-of-the-art performance.

1. Introduction

In the realm of medical diagnosis, there are numerous prediction tasks in which the output labels demonstrate high discrete and successive features. The problem of the health risk level could be a very example. Although it can be a continuous variable, it is often discretized *i.e.*, at several intervals in practices. In parallel with previous statements, the Diabetic Retinopathy Diagnosis (DR) labels the disease level into five rankings: no DR (0), mild DR (1),

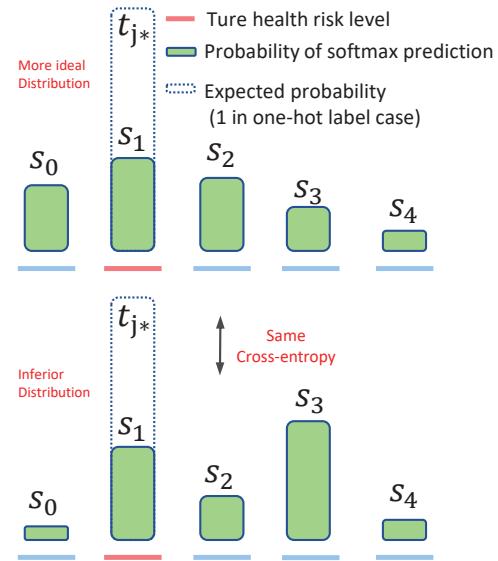


Figure 1. The limitation of CE loss for health risk level estimation. The ground truth direction of the car is t_{j*} . Two possible softmax predictions (green bar) of the health risk level estimator have the same probability at t_{j*} position. Therefore, both predicted distributions have the same CE loss. However, the top prediction is preferable to the bottom, since we desire the predicted probability distribution to be larger and closer to the ground truth class.

moderate DR (2), severe DR (3) and proliferative DR (4) [35, 22]. This kind of labeling system has been widely accepted by The Breast Imaging-Reporting and Data System (BIRADS), liver (LIRADS), gynecology (GIRADS), colonography (CRADS), etc.

Due to the successively nature of risk level, the error that misclassifying a proliferative DR (4) image to mild DR (1) is considerably severe than the counterpart that misclassifying to severe DR (3). According to earlier literature, this kind of medical diagnosis task often casts as a multi-class classification problem or a metric regression problem.

The multi-class classification formulation using the cross-entropy (CE) loss, the class labels are assumed to be independent of each other [33]. Therefore, the inter-class

similarity is not properly exploited. For instance, in Fig. 1, the histogram prediction is most well-liked to be focused close to the truth category, whereas the Cross entropy loss doesn't encourage that.

On the opposite hand, regression treats the discrete risk as continuous value. As discussed in [33, 40, 38, 15], training the regression model with discrete labels can result in over-fitting and has similar or lower accuracy than classification. Therefore, it's necessary to contemplate the sequential and discrete nature of the ordinal medical diagnosis.

Recent works generally use the $N - 1$ binary classification sub-tasks using sigmoid output with MSE loss or softmax output with cross entropy loss, where N is the number of levels. Unfortunately, the cumulative probabilities $p(y > 1|\mathbf{x}), \dots, p(y > N - 1|\mathbf{x})$ are calculated by several independent branches, failing to guarantee the that they are monotonically decreasing. This result in $p(y = i|\mathbf{x})$ don't seem to be ensured to be strictly positive, additionally the poor learning potency within the early stage of training. Moreover, $N - 1$ weights need to be manually fine-tuned to balance the CE loss of each branch. [35, 22] proposes to use a stick-breaking process to re-parametrize the outputs of $N - 1$ units that is associated with the Bayesian non-parametric Dirichlet process. Embarking on this, the cumulative probabilities can achieve the expected monotonical decrease, but it couldn't be unheeded that it's considerably more sophisticate than standard CE-loss.

Furthermore, [13] propose to use one output unit to calculate the parameter of a unimodal distribution, and define the $p(y = i|\mathbf{x})$ follows a Poisson or Binomial distribution, which suffers from lacking the ability to control the variance. Since the peak (also the mean and variance) of a Poisson distribution is equal to a designated λ , the peak cannot be assigned to the first or last class, and its variance is intended to be rather high when the peak is needed in the very later classes.

In this paper, we employ the Wasserstein loss as an alternative for empirical risk minimization. The 1^{st} Wasserstein distance is defined as the cost of optimal transport for moving the mass in one distribution to match the target distribution [5, 46, 47]. Specifically, the Wasserstein distance, between softmax prediction and its target label that are separately normalized as histograms, is measured. By defining the ground metric as class similarity, prediction performance earns the measure room in a sensitive way to correlations between the classes.

The ground metric can be predefined when the similarity structure is known as a priori to incorporate the inter-class correlation, *e.g.*, the Euclidean distance in a line. We further extend the Euclidean distance to its increasing function from an optimization perspective [31, 26]. The exact Wasserstein distance in a one-hot target label setting can be formulated as a soft-attention scheme of all prediction prob-

abilities and be rapidly computed.

Another challenge of health risk level estimation stems from the label quality. For instance, the agreement rate of the radiologists for malignancy is usually less than 80%, resulting in a noisy labeled dataset [39, 48]. Despite the often-uncleared distinction between adjacent labels, it is more possible that a well-trained annotator will mislabel a Severe DR (3) sample to Moderate DR (2) rather than No DR (0). This requires modeling the noise for robust training [19, 10].

The mis-labeled data can misleading the training [49, 3, 4, 42]. In here, two kinds of the label noise are investigated. The outlier noise refer to a sample is different from all of the others which is caused by the random error. It can be sculpturesque by a uniform distribution [49]. However, it can be more common that the ordinal case is a lot of probably to possess the inlier noise, in which the labels are wrong annotated as the near one. We propose to model it with a unimodal distribution. We give a solution that use a smoothed target histogram by smoothing the one-hot label with an uniform-unimodal mixture model.

In contrast to the one-hot setting, the smoothed label makes the calculation of Wasserstein loss more sophisticated. This is due to the various attainable transportation plans. The $\mathcal{O}(N^3)$ computational complexity for N classes has long been difficulty in using Wasserstein distance for large-scale applications. Rather than approximate its Wasserstein distance with a $\mathcal{O}(N^2)$ complexity algorithm [12, 16], we propose to systematically conclude the fast closed-form computation of Wasserstein distance in the setting of smoothed label, when the ground metric is a linear, convex, or concave increasing function *w.r.t.* the Euclidean distance. We show that the linear and convex cases can be solved with a linear complexity of $\mathcal{O}(N)$. In comparison to its approximate counterpart, the *exact* solutions that this paper proposes are more effective.

The main contributions of this paper are summarized as follows:

- The health risk level estimation casts as a Wasserstein training problem. The inter-class relationship of health risk level data is explicitly incorporated as prior information in the gained ground metric which can be pre-defined (*e.g.*, a function *w.r.t.* Euclidean distance in a line).
- The inlier and outlier error of health risk level is modeled with a discrete unimodal-uniform mixture distribution, and regularizes the target confidence by transforming one-hot label to the smoothed target label.
- For either one-hot or smoothed target label, this paper systematically concludes the possible fast closed-form solution when a non-negative linear, convex or concave increasing mapping function is applied in ground metric.

We validate the the proposed method on several medical diagnosis datasets. Our method achieves the state-

of-the-art performance over the current methods, especially choosing the convex mapping function for ground metric, smoothed target, and closed-form solution.

2. Related Works

2.1. Health risk level estimation

The traditional ordinal level estimation methods can be classified to 3 categories, *i.e.*, naive, binary decomposition and threshold methods [17, 37, 52]. Moreover, the health risk isn't the sole area that has the successive and discrete label. The age prediction and atheistic rating also tightly related to this task. Following the development of deep learning, several works have been proposed to target the successive data. [35, 44] put forward the multi-task learning framework. However, the percentages of each class are not guaranteed to be positive, which may hurt the training, especially that in the early stage. Besides, there are $N - 1$ weights to balance the branches, which is a hard task for manually tuning. [35, 22] proposes a sophisticated stick-breaking process to reparameterise the $N - 1$ outputs to alleviate this issue. [36] incorporate the metric learning for data relationship analysis. Different from these methods, we propose to use the Wasserstein distance as the optimization objective to inherit the label similarity.

2.2. Wasserstein distance

Wasserstein distance is a distance function defined between probability distributions on a given metric space [33]. In these few years, it attracts a lot attention in adversarial generative models *etc* [1, 30, 34, 27, 28]. It also been used for multi-class multi-label task with a linear model [16]. Since the significant amount of computation cost required to solve its exact solution, these methods usually choose the approximate solution, of which the complexity is still in $\mathcal{O}(N^2)$ [12]. The fast computing of discrete Wasserstein distance is also closely related to SIFT [41, 9] descriptor, hue in HSV or LCH space [8, 14]. Based on these works, we further adapt this idea to the health risk level estimation, and encode the correlation of label classes using the ground matrix [33]. We show that the fast algorithms exist in our health risk label structure with the one-hot or smoothed target label. Moreover, the ground metric used in here does not limited to the Euclidean distance.

2.3. Robust training with noise data

Robust training with noise data has been studied for a long time in general classification area [19, 3]. The possible solution can be smoothing the one-hot label [49] with a uniform distribution or regularizing the entropy of softmax output [42]. However, for the discrete successive label, the studies speak little voice.

	0	1	2	3	4
0	0	1	2	3	4
1	1	0	1	2	3
2	2	1	0	1	2
3	3	2	1	0	1
4	4	3	2	1	0

Figure 2. The ground matrix using Euclidean distance as ground metric.

2.4. Unimodality of Discrete and Successive Data

[13] propose to predict a Poisson distribution. In their parametric version, the output of the neural network is a single sigmoid unit, which predict the parameter λ in a Poisson distribution. However, requiring the output strictly follows a specific distribution could be a very strong assumption. Besides, it is difficult to control the variance of the resulting Poisson distribution. [2] introduces an additional temperature parameter to control the variance, but results in more complicate hyper-parameter tuning. Here, we propose to use an exponential function following the softmax to flexibly adjust the shape of target label distribution. Noticing this modification works on the target label instead of the output distribution as [2].

3. Methodology

In this paper, we consider the task of learning a health risk level estimator h_θ , parameterized by θ , with N -dimensional softmax output unit. It maps a medical image \mathbf{x} to a vector $\mathbf{s} \in \mathbb{R}^N$. We perform learning over a hypothesis space \mathcal{H} of h_θ . With the input \mathbf{x} and its target ground truth one-hot label \mathbf{t} , typically, learning is performed via empirical risk minimization to solve $\min_{h_\theta \in \mathcal{H}} \mathcal{L}(h_\theta(\mathbf{x}), \mathbf{t})$, with a loss $\mathcal{L}(\cdot, \cdot)$ acting as a surrogate of performance measure.

However, cross entropy-based loss treat the output dimensions independently [16]. Therefore, it ignores the similarity structure on label space.

$\mathbf{s} = \{s_i\}_{i=0}^{N-1}$ is the output of $h_\theta(\mathbf{x})$, *i.e.*, softmax prediction with N classes (health risk levels). $\mathbf{t} = \{t_j\}_{j=0}^{N-1}$ is the target label distribution, where $i, j \in \{0, \dots, N-1\}$ be the index of dimension (class). Assume that the class label possesses a ground metric $\mathbf{D}_{i,j}$, which measures the similarity of i -th and j -th dimensions of the output. There are N^2 possible $\mathbf{D}_{i,j}$ in a N class dataset and form a ground distance matrix $\mathbf{D} \in \mathbb{R}^{N \times N}$. When \mathbf{s} and \mathbf{t} are both his-

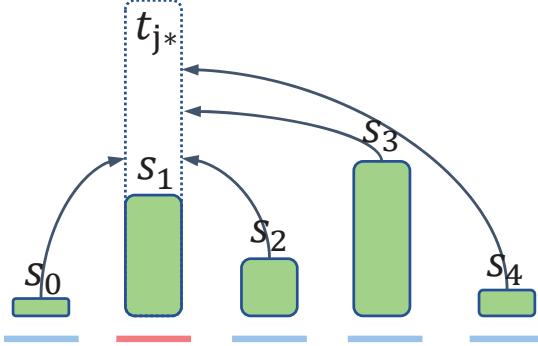


Figure 3. The only possible transport plan in one-hot target case.

tograms, the discrete measure of exact Wasserstein loss is defined as

$$\mathcal{L}_{\mathbf{D}_{i,j}}(\mathbf{s}, \mathbf{t}) = \inf_{\mathbf{W}} \sum_{j=0}^{N-1} \sum_{i=0}^{N-1} \mathbf{D}_{i,j} \mathbf{W}_{i,j} \quad (1)$$

where \mathbf{W} is the transportation matrix with $\mathbf{W}_{i,j}$ indicating the mass moved from the i^{th} point in source distribution to the j^{th} target position. A valid transportation matrix \mathbf{W} satisfies: $\mathbf{W}_{i,j} \geq 0$; $\sum_{j=0}^{N-1} \mathbf{W}_{i,j} \leq s_i$; $\sum_{i=0}^{N-1} \mathbf{W}_{i,j} \leq t_j$; $\sum_{j=0}^{N-1} \sum_{i=0}^{N-1} \mathbf{W}_{i,j} = \min(\sum_{i=0}^{N-1} s_i, \sum_{j=0}^{N-1} t_j)$.

The entries of \mathbf{D} in Wasserstein distance are usually unknown, but they have clear meanings in our task. The i, j -th entry $\mathbf{D}_{i,j}$ indicates the geometrical distance between the i -th and j -th points in a line. A possible choice is using the Euclidean distance $d_{i,j}$ of a line (*i.e.*, ℓ_1 distance between the i -th and j -th points in a line) as the ground metric $\mathbf{D}_{i,j} = d_{i,j}$.

$$d_{i,j} = |i - j| \quad (2)$$

The Wasserstein distance can be the Earth mover's distance when the two distributions have the same total masses (*i.e.*, $\sum_{i=0}^{N-1} s_i = \sum_{j=0}^{N-1} t_j$) and using the symmetric distance $d_{i,j}$ as $\mathbf{D}_{i,j}$. The ground matrix using Euclidean distance is shown in Fig. 2.

The previous efficient algorithms to solve Wasserstein distance usually holds only for $\mathbf{D}_{i,j} = d_{i,j}$ [9, 41, 46, 6, 51] and do not consider the neural network optimization. Regarding this, this paper proposes extending the ground metric in $\mathbf{D}_{i,j}$ as $f(d_{i,j})$, where f is a positive increasing function *w.r.t.* $d_{i,j}$.

3.1. Wasserstein training with one-hot target

In the multi-class single-label dataset, the one-hot encoding is typically used. The distribution of a target label probability is $\mathbf{t} = \delta_{j,j^*}$, where j^* is the ground truth class, δ_{j,j^*} is a Dirac delta, which equals to 1 for $j = j^*$ ¹, and 0

¹We use i, j interlaced for \mathbf{s} and \mathbf{t} , since they index the same group of positions in a line.

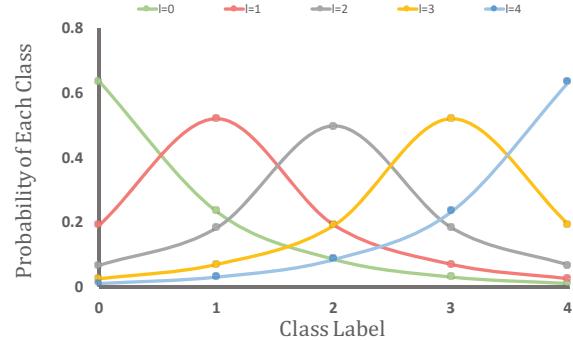


Figure 4. The distribution of normalized exponential function $e^{-|i-l|}$ for a dataset with 5 classes.

otherwise.

Theorem 1. Assume that $\sum_{j=0}^{N-1} t_j = \sum_{i=0}^{N-1} s_i$, and \mathbf{t} is a one-hot distribution with $t_{j^*} = 1$ (*or* $\sum_{i=0}^{N-1} s_i$)², there is only one feasible optimal transport plan.

In order to satisfy the criteria of \mathbf{W} , all masses should be transferred to the cluster of the ground truth label j^* , as shown in Fig. 3. Then, the Wasserstein distance between softmax prediction \mathbf{s} and one-hot target \mathbf{t} can be simplified to

$$\mathcal{L}_{\mathbf{D}_{i,j}^f}(\mathbf{s}, \mathbf{t}) = \sum_{i=0}^{N-1} s_i f(d_{i,j^*}) \quad (3)$$

where $\mathbf{D}_{i,j}^f = f(d_{i,j})$. f is an increasing function proper, *e.g.*, p^{th} power of $d_{i,j}$ and Huber function. The exact solution of Eq. (3) can be computed with a complexity of $\mathcal{O}(N)$. The ground metric term $f(d_{i,j^*})$ works as the weights *w.r.t.* s_i , which takes all classes into account following a soft attention scheme [29, 24, 25]. It explicitly encourages the probabilities distributing on the neighboring classes of j^* . Since each s_i is a function of the network parameters, differentiating $\mathcal{L}_{\mathbf{D}_{i,j}^f}$ *w.r.t.* network parameters yields $\sum_{i=0}^{N-1} s'_i f(d_{i,j^*})$.

In contrast, the cross-entropy loss in one-hot setting can be formulated as $-\log s_{j^*}$, which only considers a single class prediction as hard attention [29, 24]. Similarly, the regression loss using softmax prediction could be $f(d_{i^*,j^*})$, where i^* is the class with maximum prediction probability.

3.2. Unimodal-uniform label smoothing

The outlier noise exists commonly in data-driven tasks, and can be modeled by a uniform distribution [49, 32, 23]. However, ordinal labels are more possible to be mislabeled as a close class of the true class. It is more reasonable to form a unimodal distribution to depict the inlier noise in

²We note that softmax cannot strictly guarantee the sum of its outputs to be 1 considering the rounding operation. However, the difference of setting t_{j^*} to 1 or $\sum_{i=0}^{N-1} s_i$ is not significant in our experiments using the typical format of softmax output which is accurate to 8 decimal places.

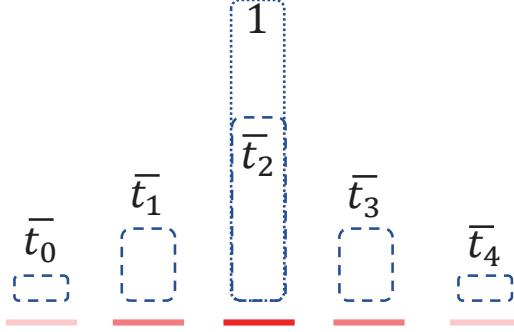


Figure 5. The unimodal-uniform smoothed target label distribution when the ground-of-truth class is 2.

health risk level estimation, which has a peak at class j^* while decreasing its value for farther classes. This paper has sampled a continuous unimodal distribution and follow by normalization.

In here, we propose to sample on an exponential function $e^{\frac{-|i-t|}{\tau}}$ and processed by a softmax normalization. Discrete distributions with five classes are illustrated in Fig. 4.

The normalized unimodal value is denoted as p_j . The unimodal-uniform smoothed target distribution $\bar{\mathbf{t}}$ is constructed by replacing t_j in \mathbf{t} with $(1 - \xi - \eta)t_j + \xi p_j + \eta \frac{1}{N}$, which can be regarded as the weighted sum of the original label distribution \mathbf{t} and a unimodal-uniform mixture distribution. In the context that the uniform distribution target is utilized for the CE loss, it is equivalent to label smoothing [49], a typical mechanism for outlier noisy label training, which encourages the model to accommodate less-confident labels. The smoothed distribution is shown in Fig. 5

Noticing that the smoothed target label can also be adapted into CE loss which is formulated as

$$\mathcal{L} = \sum_i^{N-1} \bar{t}_i [-\log(p(y = i | \mathbf{x}))] \quad (4)$$

By enforcing \mathbf{s} to form a unimodal-uniform mixture distribution, we also implicitly encourage the probabilities to distribute on the neighbor classes of j^* .

Since \bar{t}_i is monotonically decreasing w.r.t. the farther distance from the true class l , we can regard it as a weight of $-\log(p(y = i | \mathbf{x}))$. Since the target label regularization can be processed in advance, the training time does not increase by adding the unimodal-uniform mixture distribution regularization.

3.3. Wasserstein training with smoothed target

When using the unimodal-uniform smoothed label, the fast solution in Eq. (3) does not apply. The possible solution is to regard it as a general case and solve its exact solution in the complexity higher than $\mathcal{O}(N^3)$ or approximate its

exact solution with a complexity in $\mathcal{O}(N^2)$. In this section, a series of analytic formulation when the ground metric is a nonnegative increasing linear/convex/concave function w.r.t. Euclidean distance with reasonable complexity.

3.3.1 Linear and Convex function w.r.t. $d_{i,j}$ as the ground metric.

Choosing the ground metric as $d_{i,j}$ or extend it as an nonnegative increasing and convex function of $d_{i,j}$, the analytic formulation of our loss function $\mathcal{L}_{d_{i,j}}(\mathbf{s}, \bar{\mathbf{t}})$ can be formulated as

$$\mathcal{L}_{d_{i,j}}(\mathbf{s}, \bar{\mathbf{t}}) = \sum_{j=0}^{N-1} \left| \sum_{i=0}^j (s_i - \bar{t}_i) \right| \quad (5)$$

Eq. (5) was developed in [51], where it is proved for sets of points with unitary masses on a line. A similar conclusion for the Kantorovich-Rubinstein problem was derived in [6, 7], which is known to be identical to the Wasserstein distance problem when $\mathbf{D}_{i,j}$ is a distance. We note that this is true for $\mathcal{L}_{d_{i,j}}$ (but not hold for $\mathcal{L}_{\mathbf{D}^\rho}(\mathbf{s}, \bar{\mathbf{t}})$ with $\rho > 1$). An equivalent calculation is proposed from the cumulative distribution perspective [43]. All of these works notice that computing Eq. (5) can be solved in linear time ($\mathcal{O}(N)$). See [50] for a comprehensive review).

Noticing that the partial derivative of Eq. (5) w.r.t. s_n is $\sum_{j=0}^{N-1} \text{sgn}(\varphi_j) \sum_{i=0}^j (\delta_{i,n} - s_i)$, where $\varphi_j = \sum_{i=0}^j (s_i - \bar{t}_i)$, and $\delta_{i,n} = 1$ when $i = n$.

Here, we give some measures ³ using the typical convex ground metric function.

Using d^ρ as ground metric, i.e., $\mathbf{D}_{i,j}^\rho = d_{i,j}^\rho$. The loss function is written as $\mathcal{L}_{\mathbf{D}_{i,j}^\rho}(\mathbf{s}, \bar{\mathbf{t}})$, with $\rho = 2, 3, \dots$. When set $\rho = 2$, the Wasserstein distance is equivalent to the Cramér distance [45]. Note that the Cramér distance is not a distance metric proper. However, its square root is.

We can also use Huber cost function with a parameter τ and denote it as $\mathcal{L}_{\mathbf{D}_{i,j}^{H\tau}}(\mathbf{s}, \bar{\mathbf{t}})$.

$$\mathbf{D}_{i,j}^{H\tau} = \begin{cases} d_{i,j}^2 & \text{if } d_{i,j} \leq \tau \\ \tau(2d_{i,j} - \tau) & \text{otherwise.} \end{cases} \quad (6)$$

3.3.3 Concave function w.r.t. $d_{i,j}$ as the ground metric

In practice, it may be useful to define the ground metric as a nonnegative, concave and increasing function w.r.t. $d_{i,j}$. Although the general computation speed of the concave function is not satisfactory, the step function $f(t) = \mathbb{1}_{t \neq 0}$ (one every where except at 0) can be a special case, which has significantly less complexity [50]. Assuming that the $f(t) = \mathbb{1}_{t \neq 0}$, the Wasserstein metric between two normalized discrete histograms on N bins can be simplified to

³We refer to “measure”, since a ρ^{th} -root normalization is required to get a distance [50], which satisfies positive definiteness, symmetry and triangle inequality.

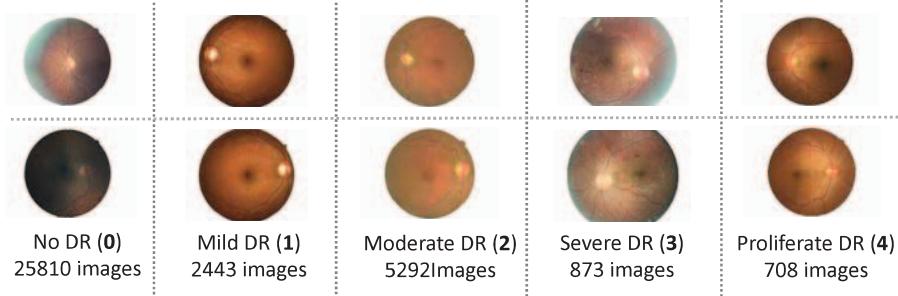


Figure 6. Some samples with different retinopathy level in the DR dataset. The top row is the left retinopathy image while the bottom row is the right retinopathy image. The samples show a large inner-class variation and small inter-class variation.

the ℓ_1 distance.

$$\mathcal{L}_{1d_{i,j} \neq 0}(\mathbf{s}, \bar{\mathbf{t}}) = \frac{1}{2} \sum_{i=0}^{N-1} |\mathbf{s}_i - \bar{\mathbf{t}}_i| = \frac{1}{2} \|\mathbf{s} - \bar{\mathbf{t}}\|_1 \quad (7)$$

where $\|\cdot\|_1$ is the discrete ℓ_1 norm.

However, its fast computation is at the cost of losing the ability to discriminate the difference of probability in a different position of bins.

4. Experiments

To evaluate the effectiveness of our Wasserstein loss, we show implementation details and experimental results on the two widely used health risk level diagnosis datasets, *i.e.*, Diabetic Retinopathy and Ultrasound BIRADS datasets. To manifest the effectiveness of each setting choice and their combinations, we give a serial of elaborate ablation studies along with the standard measures. For the fair comparison, we choose the same neural network backbones as in previous works. All of networks in our training use the \mathcal{L}_2 norm of 10^{-4} , ADAM optimizer [20] with 128 training batch-size and initial learning rate of 10^{-3} . The learning rate will be divided by ten when either the validation loss or the valid set QWK plateaus. There is no significant difference in the training time of Wasserstein loss and CE-loss based multi-class classification, and the smoothed unimodal target label is constructed before the training stage. All the experiments are implemented in deep learning platform Pytorch⁴.

We use the prefix \approx denote the approximate computation of Wasserstein distance [12, 16]. (\mathbf{s}, \mathbf{t}) and $(\mathbf{s}, \bar{\mathbf{t}})$ refer to using one-hot or smoothed target label. For instance, $\mathcal{L}_{d_{i,j}}(\mathbf{s}, \mathbf{t})$ means choosing Wasserstein loss with Euclidean distance in a line as ground metric and using one-hot target label.

4.1. Evaluations

Since the health risk level has a discrete label, the performance of a system can be simply measured by the aver-

⁴<https://pytorch.org/>

age classification accuracy as the conventional classification problem. [44] further utilized the Mean True Negative Rate (TNR) at True Positive Rate (TPR) of 0.95. The relatively high TPR used here is fitted for strict TPR requirements of medical applications to avoid misdiagnosing diseased cases as healthy. However, they do not consider the severity of different misclassification.

Considering the inherent ordered label relationship, the Mean Absolute Error (MAE) metric, *i.e.*, L_1 loss, can also be used as an evaluation metric in related risk evaluation datasets [40], which is computed using the average of the absolute errors between the ground truth and the estimated result. Here, we also propose its use in evaluating the proposed method on two medical health risk evaluation benchmarks.

Moreover, as defined in previous Kaggle DR competition, we also evaluate the quadratic weighted kappa (QWK)⁵. It can punish the misclassification proportional to the distance between the predicted label of the network and the ground-of-truth label [11]. The QWK is formulated as:

$$k = 1 - \frac{\sum_{i,j} \mathbf{W}_{i,j} \mathbf{O}_{i,j}}{\sum_{i,j} \mathbf{W}_{i,j} \mathbf{E}_{i,j}} \quad (8)$$

We evaluate on Diabetic Retinopathy (DR) and Ultrasound BIRADS datasets which are suitable for deep learning implementations in the medical area.

4.2. Diabetic Retinopathy (DR)

The Diabetic Retinopathy (DR) dataset⁶ contains a large amount of high-resolution fundus (*i.e.*, interior surface at the back of the eye) images which have been labeled as five levels of DR. The level 0 to 4 representing the No DR, Mild DR, Moderate DR, Severe DR, and Proliferative DR, respectively. The left and right fundus images from 17563 pa-

⁵<https://www.kaggle.com/c/diabetic-retinopathy-detection/overview/evaluation>

⁶<https://www.kaggle.com/c/diabetic-retinopathy-detection>

Table 1. Performance on the DR dataset.

Evaluations	Mean TNR@TPR=0.95			Valid Acc	Valid QWK	MAE
	0 vs 1-4	0-1 vs 2-4	0-2 vs 3-4			
MC	41.5%	30.9%	31.1%	82.4%	0.724	0.37
RG	40.3%	30.6%	30.8 %	76.2%	0.705	0.38
Poisson [2]	38.8%	30.0%	29.6 %	77.1%	0.713	0.38
MT [44]	42.7%	31.7%	31.3%	82.8%	0.726	0.36
SB[35]	44.0%	33.1%	32.6%	84.2%	0.743	0.32
$\mathcal{L}_{d_{i,j}}(\mathbf{s}, \mathbf{t})$	46.9%	37.1%	34.4%	87.3%	0.768	0.29
$\mathcal{L}_{\mathbf{D}_{i,j}^2}(\mathbf{s}, \mathbf{t})$	47.2%	37.3%	34.6%	87.4%	0.769	0.28
$\mathcal{L}_{\mathbf{D}_{i,j}^{H,\tau}}(\mathbf{s}, \mathbf{t})$	47.2%	37.4%	34.5%	87.6%	0.769	0.28
MC(\mathbf{s}, \mathbf{t})	42.4%	31.2%	31.8%	82.7%	0.728	0.35
$\approx \mathcal{L}_{d_{i,j}}(\mathbf{s}, \bar{\mathbf{t}})$	45.8%	36.4%	33.8%	86.6%	0.759	0.29
$\approx \mathcal{L}_{\mathbf{D}_{i,j}^2}(\mathbf{s}, \bar{\mathbf{t}})$	45.8%	36.5%	33.9%	86.5%	0.760	0.29
$\approx \mathcal{L}_{\mathbf{D}_{i,j}^{H,\tau}}(\mathbf{s}, \bar{\mathbf{t}})$	45.9%	36.6%	34.0%	86.6%	0.760	0.28
$\mathcal{L}_{d_{i,j}}(\mathbf{s}, \bar{\mathbf{t}})$	47.3%	37.5%	34.8%	87.8%	0.771	0.27
$\mathcal{L}_{\mathbf{D}_{i,j}^2}(\mathbf{s}, \bar{\mathbf{t}})$	47.6%	37.7%	34.9%	87.8%	0.772	0.26
$\mathcal{L}_{\mathbf{D}_{i,j}^{H,\tau}}(\mathbf{s}, \bar{\mathbf{t}})$	47.5%	37.7%	34.8%	88.0%	0.773	0.26

tients are publicly available. The ResNet [18] style model with 11 ResBlocks as in [2, 35] has been adopted for DR dataset. We use five neurons with softmax normalization as our output to represent the probability of each level.

In our experiments, we follow the setting of [2, 35]. The subject-independent 10-fold cross-validation is adopted, *i.e.*, the validation set consisting of 10% of the patients is set aside. The images belonging to a subject will only appear in a single fold. By doing this, we can avoid contamination. The images are also preprocessed as in [2, 35, 22] and subsequently resized as 256×256 size images. Some examples can be found in Fig. 6.

We show the results in the DR dataset in Table 1. The evaluation metrics discussed earlier is utilized. Several baseline methods are chosen for comparisons. For example, the CE-loss based multi-class classification (MC), MSE-loss based metric regression (MSE), Poisson distribution output with CE-loss (Poisson), multi-task network using a series of binary CE loss (MT), and the stick-breaking with CE-loss (SB).

The MC usually outperforms MSE in most of the metric. However, MSE usually appears to be competitive w.r.t. MAE, since MSE optimizes a similar metric as MAE in its training phase. The Poisson does not manage to achieve performance improvements in most of the evaluations due to its uncontrollable variance. The MT is more promising than MC as it considers the successive relationship, despite it has a lot to be tuned hyper-parameters. By addressing some limitation in MT, the SB has a better performance than MT.

Our Wasserstein training outperforms all of the previous methods, especially the $\mathcal{L}_{\mathbf{D}_{i,j}^2}$ and $\mathcal{L}_{\mathbf{D}_{i,j}^{H,\tau}}$ which use the convex function of the Euclidean distance in a line as the ground metric.

Moreover, the unimodal-uniform smoothed target label can efficiently improve the performance without additional training costs. The smoothing process is benefiting to both conventional CE loss and Wasserstein loss. Besides, the exact solution can outperform its approximate counterpart consistently.

We set our hyper-parameters $\xi = 0.15$, $\eta = 0.05$ and $\tau = 1$. QWK is not sensitive to the $\tau \in \{0.8, 0.9, 1, 1.1\}$ when we fix the $\xi = 0.15$. Similarly, the QWK keep at the same level when we adjust ξ from 0.12 to 0.18.

4.3. Ultrasound BIRADS

The second medical dataset is the Ultrasound BIRADS (US-BIRADS) [44]. It consists of 4904 breast images with the BIRADS system label. Considering the relatively limited number of samples in level 4, we usually regard the 3-4 as a single level [44]. That results in 2700 healthy (0) images, 1113 benign (1) images, 359 probably benign (2) images, and 732 may contain/contain malignant (3-4) images. We divide this dataset into 5 subsets for subject-independent five-fold cross-validation. We show some samples at different levels in Fig. 7.

AlexNet style architecture [21] with six convolution layers and following two dense layers is used for US-BIRADS image dataset as in [44]. We set $\xi = 0.15$, $\eta = 0.05$, and



Figure 7. Some samples with different malignant risk in the US-BIRADS.

Table 2. Performance on the US-BIRADS dataset.*Our implementations have slightly higher TNR using MC baseline than the results reported in [44]

Evaluations	Mean TNR@TPR=0.95			Valid Acc	Valid QWK	MAE
	0 vs 1-4	0-1 vs 2-4	0-2 vs 3-4			
MC	33.2%*	28.7%*	29.8%*	73.3%	0.678	0.42
RG	31.6%	28.5%	29.5%	73.0%	0.677	0.44
Poisson [2]	29.6%	27.2%	29.5%	72.2%	0.665	0.45
MT [44]	38.5%	29.2%	31.3%	76.5%	0.685	0.41
SB[35]	39.1%	30.2%	32.0%	78.3%	0.694	0.39
$\mathcal{L}_{d_{i,j}}(\mathbf{s}, \mathbf{t})$	42.5%	33.6%	35.7%	80.1%	0.712	0.36
$\mathcal{L}_{\mathbf{D}_{i,j}^2}(\mathbf{s}, \mathbf{t})$	42.6%	33.8%	35.9%	80.2%	0.714	0.35
$\mathcal{L}_{\mathbf{D}_{i,j}^{H,\tau}}(\mathbf{s}, \mathbf{t})$	42.6%	33.7%	35.9%	80.3%	0.715	0.35
$\text{MC}(\mathbf{s}, \mathbf{t})$	33.4%	29.0%	30.4%	73.6%	0.682	0.40
$\mathcal{L}_{d_{i,j}}(\mathbf{s}, \bar{\mathbf{t}})$	42.9%	34.0%	36.2%	80.5%	0.715	0.34
$\mathcal{L}_{\mathbf{D}_{i,j}^2}(\mathbf{s}, \bar{\mathbf{t}})$	43.0%	34.2%	36.3%	80.5%	0.716	0.34
$\mathcal{L}_{\mathbf{D}_{i,j}^{H,\tau}}(\mathbf{s}, \bar{\mathbf{t}})$	43.0%	34.1%	36.3%	80.6%	0.716	0.33

$\tau = 1$ for the unimodal distribution.

The leading performance of our method is also observed in the US-BIRADS dataset (Table 2). Since its labels are noisier (more severe annotator-dependent problem), the unimodal-uniform smoothing usually offers a more appealing contribution to the results. The Wasserstein training with convex ground metric function and smoothed target label achieves state-of-the-art performance consistently.

5. Conclusions

Based on the Wasserstein distances, we proposed an efficient loss function as an alternative to the CE loss for health risk level estimation. The ground metric inherits the inter-class correlation and can be predefined by with an increasing function w.r.t. the Euclidean distance of a line. The intlier and outlier noise in health risk data can be incorporated in a unimodal-uniform mixture distribution to form the smoothed target. We systematically discusses the fast closed-form solutions in one-hot and conservative label cases. The results show that the best performance can be

achieved by choosing convex function, unimodal-uniform distribution for smoothing and solving its exact solution. Although it was originally developed for health risk level estimation, it is essentially applicable to other problems with discrete and ordinal labels. In the future, we intend to develop an adaptive ground metric learning scheme, and adjust the shape of conservative target distribution automatically.

6. Acknowledgement

The funding support from National Institute of Health (NIH), National Institute of Neurological Disorders and Stroke (NINDS) (NS061841, NS095986), Youth Innovation Promotion Association, CAS (2017264), Innovative Foundation of CIOMP, CAS (Y586320150) and Hong Kong Government General Research Fund GRF (Ref. No.152202/14E) are greatly appreciated.

References

- [1] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017. 3
- [2] C. Beckham and C. Pal. Unimodal probability distributions for deep ordinal classification. *arXiv preprint arXiv:1705.05278*, 2017. 3, 7, 8
- [3] A. J. Bekker and J. Goldberger. Training deep neural networks based on unreliable labels. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 2682–2686. IEEE, 2016. 2, 3
- [4] V. Belagiannis, C. Rupprecht, G. Carneiro, and N. Navab. Robust optimization for deep regression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2830–2838, 2015. 2
- [5] V. I. Bogachev and A. V. Kolesnikov. The monge-kantorovich problem: achievements, connections, and perspectives. *Russian Mathematical Surveys*, 67(5):785–890, 2012. 2
- [6] C. A. Cabrelli and U. M. Molter. The kantorovich metric for probability measures on the circle. *Journal of Computational and Applied Mathematics*, 57(3):345–361, 1995. 4, 5
- [7] C. A. Cabrelli and U. M. Molter. A linear time algorithm for a matching problem on the circle. *Information processing letters*, 66(3):161–164, 1998. 5
- [8] S.-H. Cha. A fast hue-based colour image indexing algorithm. *Machine Graphics & Vision International Journal*, 11(2/3):285–295, 2002. 3
- [9] S.-H. Cha and S. N. Srihari. On measuring the distance between histograms. *Pattern Recognition*, 35(6):1355–1370, 2002. 3, 4
- [10] T. Che, X. Liu, S. Li, Y. Ge, R. Zhang, C. Xiong, and Y. Bengio. Deep verifier networks: Verification of deep discriminative models with deep generative models. In *ArXiv*, 2019. 2
- [11] J. Cohen. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin*, 70(4):213, 1968. 6
- [12] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013. 2, 3, 6
- [13] J. F. P. da Costa, H. Alonso, and J. S. Cardoso. The unimodal model for the classification of ordinal data. *Neural Networks*, 21(1):78–91, 2008. 2, 3
- [14] J. Delon, J. Salomon, and A. Sobolevski. Fast transport optimization for monge costs on the circle. *SIAM Journal on Applied Mathematics*, 70(7):2239–2258, 2010. 3
- [15] G. Divon and A. Tal. Viewpoint estimation insights & model. *arXiv preprint arXiv:1807.01312*, 2018. 2
- [16] C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. A. Poggio. Learning with a wasserstein loss. In *Advances in Neural Information Processing Systems*, pages 2053–2061, 2015. 2, 3, 6
- [17] P. A. Gutierrez, M. Perez-Ortiz, J. Sanchez-Monedero, F. Fernandez-Navarro, and C. Hervas-Martinez. Ordinal regression methods: survey and experimental study. *IEEE Transactions on Knowledge and Data Engineering*, 28(1):127–146, 2016. 3
- [18] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [19] P. J. Huber. Robust statistics. In *International Encyclopedia of Statistical Science*, pages 1248–1251. Springer, 2011. 2, 3
- [20] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [21] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 7
- [22] X. Liu, F. Fan, Y. Ge, S. Li, J. You, and J. Lu. Unimodal regularized neuron stick-breaking for ordinal regression. In *Neurocomputing*, 2019. 1, 2, 3, 7
- [23] X. Liu, Y. Ge, C. Yang, and P. Jia. Adaptive metric learning with deep neural networks for video-based facial expression recognition. *Journal of Electronic Imaging*, 27(1):013022, 2018. 4
- [24] X. Liu, Z. Guo, J. Jia, and B. Kumar. Dependency-aware attention control for imageset-based face recognition. In *IEEE TIFS*, 2019. 4
- [25] X. Liu, Z. Guo, S. Li, J. You, and K. B.V.K. Permutation-invariant feature restructuring for correlation-aware image set-based recognition. In *ICCV*, 2019. 4
- [26] X. Liu, L. Kong, Z. Diao, and P. Jia. Line-scan system for continuous hand authentication. *Optical Engineering*, 56(3):033106, 2017. 2
- [27] X. Liu, B. V. Kumar, Y. Ge, C. Yang, J. You, and P. Jia. Normalized face image generation with perceptron generative adversarial networks. In *2018 IEEE 4th International Conference on Identity, Security, and Behavior Analysis (ISBA)*, pages 1–8, 2018. 3
- [28] X. Liu, B. V. Kumar, P. Jia, and J. You. Hard negative generation for identity-disentangled facial expression recognition. *Pattern Recognition*, 88:1–12, 2019. 3
- [29] X. Liu, B. V. K. V. Kumar, C. Yang, Q. Tang, and J. You. Dependency-aware attention control for unconstrained face recognition with image sets. 2018. 4
- [30] X. Liu, S. Li, L. Kong, W. Xie, P. Jia, J. You, and B. Kumar. Feature-level frankenstein: Eliminating variations for discriminative recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 637–646, 2019. 3
- [31] X. Liu, Z. Li, L. Kong, Z. Diao, J. Yan, Y. Zou, C. Yang, P. Jia, and J. You. A joint optimization framework of low-dimensional projection and collaborative representation for discriminative classification. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 1493–1498. 2
- [32] X. Liu, B. Vijaya Kumar, J. You, and P. Jia. Adaptive deep metric learning for identity-aware facial expression recognition. In *CVPR Workshops*, pages 20–29, 2017. 4
- [33] X. Liu, Y. Zou, T. Che, J. You, and K. B.V.K. Conservative wasserstein training for pose estimation. In *ICCV*, 2019. 1, 2, 3

[34] X. Liu, Y. Zou, L. Kong, Z. Diao, J. Yan, J. Wang, S. Li, P. Jia, and J. You. Data augmentation via latent space interpolation for image classification. In *24th International Conference on Pattern Recognition (ICPR)*, pages 728–733, 2018. 3

[35] X. Liu, Y. Zou, Y. Song, C. Yang, J. You, and B. V. Kumar. Ordinal regression with neuron stick-breaking for medical diagnosis. In *European Conference on Computer Vision*, pages 335–344. Springer, 2018. 1, 2, 3, 7, 8

[36] Y. Liu, A. W.-K. Kong, and C. K. Goh. Deep ordinal regression based on data relationship for small datasets. In *IJCAI*, pages 2372–2378, 2017. 3

[37] Z. Ma and S. Chen. A convex formulation for multiple ordinal output classification. *Pattern Recognition*, 86:73–84, 2019. 3

[38] S. Mahendran, H. Ali, and R. Vidal. A mixed classification-regression framework for 3d pose estimation from 2d images. *BMVC*, 2018. 2

[39] R. M. Nishikawa, C. E. Comstock, M. N. Linver, G. M. Newstead, V. Sandhir, and R. A. Schmidt. Agreement between radiologists’ interpretations of screening mammograms. In *International Workshop on Digital Mammography*, pages 3–10. Springer, 2016. 2

[40] Z. Niu, M. Zhou, L. Wang, X. Gao, and G. Hua. Ordinal regression with multiple output cnn for age estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4920–4928, 2016. 2, 6

[41] O. Pele and M. Werman. A linear time histogram metric for improved sift matching. In *European conference on computer vision*, pages 495–508. Springer, 2008. 3, 4

[42] G. Pereyra, G. Tucker, J. Chorowski, L. Kaiser, and G. Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*, 2017. 2, 3

[43] J. Rabin, J. Delon, and Y. Gousseau. A statistical approach to the matching of local features. *SIAM Journal on Imaging Sciences*, 2(3):931–958, 2009. 5

[44] V. Ratner, Y. Shoshan, and T. Kachman. Learning multiple non-mutually-exclusive tasks for improved classification of inherently ordered labels. *arXiv preprint arXiv:1805.11837*, 2018. 3, 6, 7, 8

[45] M. L. Rizzo and G. J. Székely. Energy distance. *Wiley Interdisciplinary Reviews: Computational Statistics*, 8(1):27–38, 2016. 5

[46] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000. 2, 4

[47] L. Rüschendorf. The wasserstein distance and approximation theorems. *Probability Theory and Related Fields*, 70(1):117–129, 1985. 2

[48] A. J. Salazar, J. A. Romero, O. A. Bernal, A. P. Moreno, and S. C. Velasco. Reliability of the bi-rads final assessment categories and management recommendations in a telemammography context. *Journal of the American College of Radiology*, 14(5):686–692, 2017. 2

[49] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 2, 3, 4, 5

[50] C. Villani. *Topics in optimal transportation*. Number 58. American Mathematical Soc., 2003. 5

[51] M. Werman, S. Peleg, R. Melter, and T. Y. Kong. Bipartite graph matching for points on a line or a circle. *Journal of Algorithms*, 7(2):277–284, 1986. 4, 5

[52] H. Zhao, Z. Wang, and P. Liu. The ordinal relation preserving binary codes. *Pattern Recognition*, 48(10):3169–3179, 2015. 3