

Fisher Discriminant Triplet and Contrastive Losses for Training Siamese Networks

Benyamin Ghojogh*, Milad Sikaroudi†, Sobhan Shafiei†,

H.R. Tizhoosh†, *Senior Member, IEEE*, Fakhri Karray*, *Fellow, IEEE*, Mark Crowley*

*Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada

†Kimia Lab, University of Waterloo, Waterloo, ON, Canada

Emails: {bghojogh, msikaroudi, s7shafie, tizhoosh, karray, mcrowley}@uwaterloo.ca

Abstract—Siamese neural network is a very powerful architecture for both feature extraction and metric learning. It usually consists of several networks that share weights. The Siamese concept is topology-agnostic and can use any neural network as its backbone. The two most popular loss functions for training these networks are the triplet and contrastive loss functions. In this paper, we propose two novel loss functions, named Fisher Discriminant Triplet (FDT) and Fisher Discriminant Contrastive (FDC). The former uses anchor-neighbor-distant triplets while the latter utilizes pairs of anchor-neighbor and anchor-distant samples. The FDT and FDC loss functions are designed based on the statistical formulation of the Fisher Discriminant Analysis (FDA), which is a linear subspace learning method. Our experiments on the MNIST and two challenging and publicly available histopathology datasets show the effectiveness of the proposed loss functions.

Index Terms—Fisher discriminant analysis, triplet loss, contrastive loss, Siamese neural network, feature extraction.

I. INTRODUCTION

Siamese neural networks have been found very effective for feature extraction [1], metric learning [2], few-shot learning [3], and feature tracking [4]. A Siamese network includes several, typically two or three, backbone neural networks which share weights [5] (see Fig. 1). Different loss functions have been proposed for training a Siamese network. Two commonly used ones are triplet loss [5] and contrastive loss [6] which are displayed in Fig. 1.

We generally start with considering two samples named an anchor and one of its neighbors from the same class, and two more samples named the same anchor and a distant counterpart from a different class. The triplet loss considers the anchor-neighbor-distant triplets while the contrastive loss deals with the anchor-neighbor and anchor-distant pairs of samples. The main idea of these loss functions is to pull the samples of every class toward one another and push the samples of different classes away from each other in order to improve the classification results and hence the generalization capability. We will introduce these losses in Section II-C in more depth.

The Fisher Discriminant Analysis (FDA) [7] was first proposed in [8]. FDA is a linear method based on generalized eigenvalue problem [9] and tries to find an embedding subspace that decreases the variance of each class while increases the variance between the classes. As can be observed, there is a similar intuition behind the concepts of both the Siamese

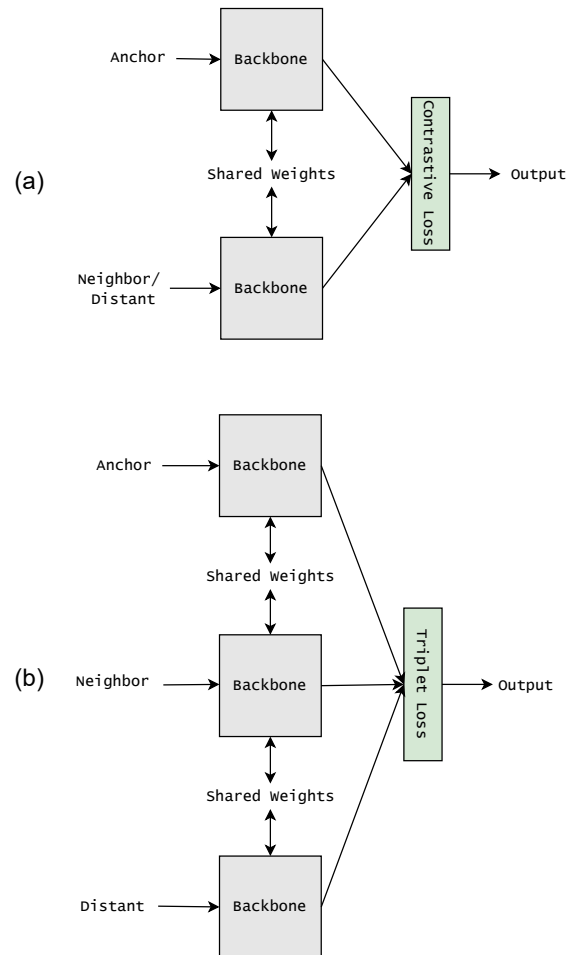


Fig. 1. Siamese network with (a) contrastive and (b) triplet loss functions.

network and the FDA, where they try to embed the data in a way that the samples of each class collapse close together [10] but the classes fall far away. Although FDA is a well-known statistical method, it has been recently noticed in the literature of deep learning [11], [12].

Noticing the similar intuition behind the Siamese network and FDA, we propose two novel loss functions for training Siamese networks, which are inspired by the theory of the FDA. We consider the intra- and inter-class scatters of the

triplets instead of their ℓ_2 norm distances. The two proposed loss functions are Fisher Discriminant Triplet (FDT) and Fisher Discriminant Contrastive (FDC) losses, which correspond to triplet and contrastive losses, respectively. Our experiments show that these loss functions exhibit very promising behavior for training Siamese networks.

The remainder of the paper is organized as follows: Section II reviews the foundation of the Fisher criterion, FDA, Siamese network, triplet loss, and contrastive loss. In Sections III-B and III-C, we propose the FDT and FDC loss functions for training Siamese networks, respectively. In Section IV, we report multiple experiments on different benchmark datasets to demonstrate the effectiveness of the proposed losses. Finally, Section V concludes the paper.

II. BACKGROUND

A. Fisher Criterion

Assume the data include c classes where the k -th class, with the sample size n_k , is denoted by $\{\mathbf{x}_i^{(k)}\}_{i=1}^{n_k}$. Let the dimensionality of data be d . Consider a p -dimensional subspace (with $p \leq d$) onto which the data are projected. We can define intra- (within) and inter-class (between) scatters as the scatter of projected data in and between the classes. The Fisher criterion is increasing and decreasing with the intra- and inter-class scatters, respectively; hence, by maximizing it, one can aim to maximize the inter-class scatter of projected data while minimizing the intra-class scatter. There exist different versions of the Fisher criterion [13]. Suppose $\mathbf{U} \in \mathbb{R}^{d \times p}$ is the projection matrix onto the subspace, then the trace of the matrix $\mathbf{U}^\top \mathbf{S} \mathbf{U}$ can be interpreted as the variance of the projected data [14]. Based on this interpretation, the most popular Fisher criterion is defined as follows [8], [15]

$$J := \frac{\mathbf{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U})}{\mathbf{tr}(\mathbf{U}^\top \mathbf{S}_W \mathbf{U})}, \quad (1)$$

where $\mathbf{tr}(\cdot)$ denotes the trace of matrix and $\mathbf{S}_B \in \mathbb{R}^{d \times d}$ and $\mathbf{S}_W \in \mathbb{R}^{d \times d}$ are the inter- and intra-class scatter matrices, respectively, defined as

$$\mathbf{S}_W := \sum_{k=1}^c \sum_{i=1}^{n_k} \sum_{j=1}^{n_k} (\mathbf{x}_i^{(k)} - \mathbf{x}_j^{(k)})(\mathbf{x}_i^{(k)} - \mathbf{x}_j^{(k)})^\top, \text{ and} \quad (2)$$

$$\mathbf{S}_B := \sum_{k=1}^c \sum_{\ell=1, \ell \neq k}^c \sum_{i=1}^{n_k} \sum_{j=1}^{n_\ell} (\mathbf{x}_i^{(k)} - \mathbf{x}_j^{(\ell)})(\mathbf{x}_i^{(k)} - \mathbf{x}_j^{(\ell)})^\top. \quad (3)$$

Some other versions of Fisher criterion are [13]

$$J := \mathbf{S}_W^{-1} \mathbf{S}_B, \text{ and} \quad (4)$$

$$J := \mathbf{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U}) - \mathbf{tr}(\mathbf{U}^\top \mathbf{S}_W \mathbf{U}), \quad (5)$$

where the former is because the solution to maximizing (1) is the generalized eigenvalue problem $(\mathbf{S}_B, \mathbf{S}_W)$ (see Section II-B) whose solution can be the eigenvectors of $\mathbf{S}_W^{-1} \mathbf{S}_B$ [16]. The reason for latter is because (1) is a Rayleigh-Ritz quotient [17] and its denominator can be set to a constant [14]. The Lagrange relaxation of the optimization would be similar to (5).

B. Fisher Discriminant Analysis

FDA [7], [8] is defined as a linear transformation which maximizes the criterion function (1). This criterion is a generalized Rayleigh-Ritz quotient [17] and we may recast the problem to [14]

$$\begin{aligned} & \underset{\mathbf{U}}{\text{maximize}} && \mathbf{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U}), \\ & \text{subject to} && \mathbf{U}^\top \mathbf{S}_W \mathbf{U} = \mathbf{I}, \end{aligned} \quad (6)$$

where \mathbf{I} is the identity matrix. The Lagrange relaxation of the problem can be written as follows

$$\mathcal{L} = \mathbf{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U}) - \mathbf{tr}(\mathbf{\Lambda}^\top (\mathbf{U}^\top \mathbf{S}_W \mathbf{U} - \mathbf{I})), \quad (7)$$

where $\mathbf{\Lambda}$ is a diagonal matrix which includes the Lagrange multipliers [18]. Setting the derivative of Lagrangian to zero gives

$$\frac{\partial \mathcal{L}}{\partial \mathbf{U}} = 2\mathbf{S}_B \mathbf{U} - 2\mathbf{S}_W \mathbf{U} \mathbf{\Lambda} \stackrel{\text{set}}{=} \mathbf{0} \implies \mathbf{S}_B \mathbf{U} = \mathbf{S}_W \mathbf{U} \mathbf{\Lambda}, \quad (8)$$

which is the generalized eigenvalue problem $(\mathbf{S}_B, \mathbf{S}_W)$ where the columns of \mathbf{U} and the diagonal of $\mathbf{\Lambda}$ are the eigenvectors and eigenvalues, respectively [16]. The column space of \mathbf{U} is the FDA subspace.

C. Siamese Network and Loss Functions

1) *Siamese Network*: Siamese network is a set of several (typically two or three) networks which share weights with each other [5] (see Fig. 1). The weights are trained using a loss based on anchor, neighbor (positive), and distant (negative) samples, where anchor and neighbor belong to the same class, but the anchor and distant tiles are in different classes. We denote the anchor, neighbor, and distant samples by \mathbf{x}_a , \mathbf{x}_n , and \mathbf{x}_d , respectively. The loss functions used to train a Siamese network usually make use of the anchor, neighbor, and distant samples, trying to pull the anchor and neighbor towards one another and simultaneously push the anchor and distant tiles away from each other. In the following, two different loss functions are introduced for training Siamese networks.

2) *Triplet Loss*: The triplet loss uses anchor, neighbor, and distant. Let $\mathbf{f}(\mathbf{x})$ be the output (i.e., embedding) of the network for the input \mathbf{x} . The triplet loss tries to reduce the distance of anchor and neighbor embeddings and desires to increase the distance of anchor and distant embeddings. As long as the distances of anchor-distant pairs get larger than the distances of anchor-neighbor pairs by a margin $\alpha \geq 0$, the desired embedding is obtained. The triplet loss, to be minimized, is defined as [5]

$$\ell_t = \sum_{i=1}^b \left[\|\mathbf{f}(\mathbf{x}_a^i) - \mathbf{f}(\mathbf{x}_n^i)\|_2^2 - \|\mathbf{f}(\mathbf{x}_a^i) - \mathbf{f}(\mathbf{x}_d^i)\|_2^2 + \alpha \right]_+ \quad (9)$$

where \mathbf{x}^i is the i -th triplet sample in the mini-batch, b is the mini-batch size, $[z]_+ := \max(z, 0)$, and $\|\cdot\|_2$ denotes the ℓ_2 norm.




Fig. 2. The network structure for the proposed loss functions.

3) *Contrastive Loss*: The contrastive loss uses pairs of samples which can be anchor and neighbor or anchor and distant. If the samples are anchor and neighbor, they are pulled towards each other; otherwise, their distance is increased. In other words, the contrastive loss performs like the triplet loss but one by one rather than simultaneously. The desired embedding is obtained when the anchor-distant distances get larger than the anchor-neighbor distances by a margin of α . This loss, to be minimized, is defined as [6]

$$\ell_c = \sum_{i=1}^b \left[(1-y) \|\mathbf{f}(\mathbf{x}_1^i) - \mathbf{f}(\mathbf{x}_2^i)\|_2^2 + y \left[-\|\mathbf{f}(\mathbf{x}_1^i) - \mathbf{f}(\mathbf{x}_2^i)\|_2^2 + \alpha \right]_+ \right] \quad (10)$$

where y is zero and one when the pair $\{\mathbf{x}_1^i, \mathbf{x}_2^i\}$ is anchor-neighbor and anchor-distant, respectively.

III. THE PROPOSED LOSS FUNCTIONS

A. Network Structure

Consider any arbitrary neural network as the backbone. This network can be either a multi-layer perception or a convolutional network. Let q be the number of its output neurons, i.e., the dimensionality of its embedding space. We add a fully connected layer after the q -neurons layer to a new embedding space (output layer) with $p \leq q$ neurons. Denote the weights of this layer by $\mathbf{U} \in \mathbb{R}^{q \times p}$. We name the first q -dimensional embedding space as the *latent space* and the second p -dimensional embedding space as the *feature space*. Our proposed loss functions are network-agnostic as they can be used for any network structure and topology of the backbone. The overall network structure for the usage of the proposed loss functions is depicted in Fig. 2.

Consider a triplet $\{\mathbf{x}_a, \mathbf{x}_n, \mathbf{x}_d \in \mathbb{R}^d\}$ or a pair $\{\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^d\}$. We feed the triplet or pair to the network. We denote the latent embedding of data by $\{\mathbf{o}_a, \mathbf{o}_n, \mathbf{o}_d \in \mathbb{R}^q\}$ and $\{\mathbf{o}_1, \mathbf{o}_2 \in \mathbb{R}^q\}$ while the feature embedding of data is denoted by $\{\mathbf{f}(\mathbf{x}_a), \mathbf{f}(\mathbf{x}_n), \mathbf{f}(\mathbf{x}_d) \in \mathbb{R}^p\}$ and $\{\mathbf{f}(\mathbf{x}_1), \mathbf{f}(\mathbf{x}_2) \in \mathbb{R}^p\}$. The last layer of network is projecting the latent embedding to the feature space where the activation function of the last layer is linear because of unsupervised feature extraction. Hence, the last layer acts as a linear projection $\mathbf{f}(\mathbf{x}) = \mathbf{U}^\top \mathbf{o}$.

During the training, the latent space is adjusted to extract some features; however, the last-layer projection fine-tunes the

latent features in order to have better discriminative features. In Section IV-D, we show the results of experiments to demonstrate this.

B. Fisher Discriminant Triplet Loss

As in neural networks the loss function is usually minimized, we minimize the negative of Fisher criterion where we use (5) as the criterion:

$$\underset{\mathbf{U}}{\text{minimize}} \quad -J = \text{tr}(\mathbf{U}^\top \mathbf{S}_W \mathbf{U}) - \text{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U}). \quad (11)$$

This problem is ill-defined because by increasing the total scatter of embedded data, the inter-class scatter also gets larger and this objective function gets decreased. Therefore, the embedding space scales up and explodes gradually to increase the term $\text{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U})$. In order to control this issue, we penalize the total scatter of the embedded data, denoted by $\mathbf{S}_T \in \mathbb{R}^{d \times d}$:

$$\underset{\mathbf{U}}{\text{min}} \quad \text{tr}(\mathbf{U}^\top \mathbf{S}_W \mathbf{U}) - \text{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U}) + \epsilon \text{tr}(\mathbf{U}^\top \mathbf{S}_T \mathbf{U}), \quad (12)$$

where $\epsilon \in (0, 1)$ is the regularization parameter. The total scatter can be considered as the summation of the inter- and intra-class scatters [19]:

$$\mathbf{S}_T := \mathbf{S}_B + \mathbf{S}_W. \quad (13)$$

Hence, we have:

$$\begin{aligned} & \text{tr}(\mathbf{U}^\top \mathbf{S}_W \mathbf{U}) - \text{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U}) + \epsilon \text{tr}(\mathbf{U}^\top \mathbf{S}_T \mathbf{U}) \\ &= \text{tr}(\mathbf{U}^\top (\mathbf{S}_W - \mathbf{S}_B + \epsilon \mathbf{S}_T) \mathbf{U}) \\ &\stackrel{(13)}{=} \text{tr}(\mathbf{U}^\top ((\epsilon + 1) \mathbf{S}_W + (\epsilon - 1) \mathbf{S}_B) \mathbf{U}) \\ &\stackrel{(a)}{=} (2 - \lambda) \text{tr}(\mathbf{U}^\top \mathbf{S}_W \mathbf{U}) - \lambda \text{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U}), \end{aligned} \quad (14)$$

where (a) is because $(0, 1) \ni \lambda := 1 - \epsilon$. It is recommended for ϵ and λ to be close to one and zero, respectively because the total scatter should be controlled not to explode. For example, a good value can be $\lambda = 0.1$.

We want the inter-class scatter term to get larger than the intra-class scatter term by a margin $\alpha > 0$. Hence, the FDT loss, to be minimized, is defined as:

$$\ell_{\text{fdt}} = \left[(2 - \lambda) \text{tr}(\mathbf{U}^\top \mathbf{S}_W \mathbf{U}) - \lambda \text{tr}(\mathbf{U}^\top \mathbf{S}_B \mathbf{U}) + \alpha \right]_+ \quad (15)$$

where we defer the mathematical definition of intra- and inter-class scatter matrices in our loss functions to Section III-D.

C. Fisher Discriminant Contrastive Loss

Rather than the triplets of data, we can consider the pairs of samples. For this goal, we propose the FDC loss function defined as

$$\ell_{\text{fdc}} = (2 - \lambda) \text{tr}(\mathbf{U}^\top \tilde{\mathbf{S}}_W \mathbf{U}) + \left[-\lambda \text{tr}(\mathbf{U}^\top \tilde{\mathbf{S}}_B \mathbf{U}) + \alpha \right]_+ \quad (16)$$

where the intra- and inter-class scatter matrices, which will be defined in Section III-D, consider the anchor-neighbor and anchor-distant pairs.

D. Intra- and Inter-Class Scatters

1) *Scatter Matrices in FDT*: Let the output embedding of the backbone, i.e. the second-to-last layer of total structure, be denoted by $\mathbf{o} \in \mathbb{R}^q$. We call this embedding the *latent embedding*. Consider the latent embeddings of anchor, neighbor, and distant, denoted by \mathbf{o}_a , \mathbf{o}_n , and \mathbf{o}_d , respectively. If we have a mini-batch of b triplets, we can define $\mathbb{R}^{q \times b} \ni \mathbf{O}_W := [\mathbf{o}_a^1 - \mathbf{o}_n^1, \dots, \mathbf{o}_a^b - \mathbf{o}_n^b]$ and $\mathbb{R}^{q \times b} \ni \mathbf{O}_B := [\mathbf{o}_a^1 - \mathbf{o}_d^1, \dots, \mathbf{o}_a^b - \mathbf{o}_d^b]$ where \mathbf{o}^i is the i -th sample in the mini-batch. The intra- and inter-class scatter matrices are, respectively, defined as

$$\mathbb{R}^{q \times q} \ni \mathbf{S}_W := \sum_{i=1}^b (\mathbf{o}_a^i - \mathbf{o}_n^i)(\mathbf{o}_a^i - \mathbf{o}_n^i)^\top = \mathbf{O}_W \mathbf{O}_W^\top, \quad (17)$$

$$\mathbb{R}^{q \times q} \ni \mathbf{S}_B := \sum_{i=1}^b (\mathbf{o}_a^i - \mathbf{o}_d^i)(\mathbf{o}_a^i - \mathbf{o}_d^i)^\top = \mathbf{O}_B \mathbf{O}_B^\top. \quad (18)$$

The ranks of the intra- and inter-class scatters are $\min(q, b-1)$. As the subspace of FDA can be interpreted as the eigenspace of $\mathbf{S}_W^{-1} \mathbf{S}_B$, the rank of the subspace would be $\min(q, b-1) = b-1$ because we usually have $b < q$. In order to improve the rank of the embedding subspace, we slightly strengthen the main diagonal of the scatter matrices [20]

$$\mathbf{S}_W := \mathbf{O}_W \mathbf{O}_W^\top + \mu_W \mathbf{I}, \quad (19)$$

$$\mathbf{S}_B := \mathbf{O}_B \mathbf{O}_B^\top + \mu_B \mathbf{I}, \quad (20)$$

where $\mu_W, \mu_B > 0$ are small positive numbers, e.g., 10^{-4} . Hence, the embedding subspace becomes full rank with $q \geq p$.

2) *Scatter Matrices in FDC*: As in the regular contrastive loss, we consider the pairs of anchor-neighbor and anchor-distant for the FDC loss. Let y be zero and one when the pair $\{\mathbf{x}_1^i, \mathbf{x}_2^i\}$ is an anchor-neighbor or anchor-distant pair, respectively. The latent embedding of this pair is denoted by $\{\mathbf{o}_1^i, \mathbf{o}_2^i\}$. The intra- and inter-class scatter matrices in the FDC loss are, respectively, defined as

$$\begin{aligned} \tilde{\mathbf{S}}_W &:= \sum_{i=1}^b (1-y)(\mathbf{o}_1^i - \mathbf{o}_2^i)(\mathbf{o}_1^i - \mathbf{o}_2^i)^\top + \mu_W \mathbf{I} \\ &= \tilde{\mathbf{O}}_W \tilde{\mathbf{O}}_W^\top + \mu_W \mathbf{I}, \end{aligned} \quad (21)$$

$$\begin{aligned} \tilde{\mathbf{S}}_B &:= \sum_{i=1}^b y(\mathbf{o}_1^i - \mathbf{o}_2^i)(\mathbf{o}_1^i - \mathbf{o}_2^i)^\top + \mu_B \mathbf{I} \\ &= \tilde{\mathbf{O}}_B \tilde{\mathbf{O}}_B^\top + \mu_B \mathbf{I}, \end{aligned} \quad (22)$$

where $\tilde{\mathbf{O}}_W := [\{\mathbf{o}_1^i - \mathbf{o}_2^i \mid y = 0\}]$ and $\tilde{\mathbf{O}}_B := [\{\mathbf{o}_1^i - \mathbf{o}_2^i \mid y = 1\}]$.

Note that in both FDT and FDC loss functions, there exist the weight matrix \mathbf{U} and the intra- and inter-class scatter matrices. By back-propagation, both the last layer and the previous layers are trained because \mathbf{U} in loss affects the last layer, and the scatter matrices in loss impact all the layers.

IV. EXPERIMENTS

A. Datasets

For the experiments, we used three public datasets, i.e., MNIST and two challenging histopathology datasets. In the

following, we introduce these datasets. The MNIST images have one channel, but the histopathology images exhibit color in three channels.

MNIST dataset – The MNIST dataset [21] includes 60,000 training images and 10,000 test images of size 28×28 pixels. We created a dataset of 500 triplets from the training data to test the proposed loss functions for a small training sample size.

CRC dataset – The first histopathology dataset is the Colorectal Cancer (CRC) dataset [22]. It contains tissue patches from different tissue types of colorectal cancer tissue slides. The tissue types are background (empty), adipose tissue, mucosal glands, debris, immune cells (lymphoma), complex stroma, simple stroma, and tumor epithelium. Some sample patches of CRC tissue types can be seen in Fig. 4. We split data into train/test sets with 60%–40% portions. Using the training set, we extracted 22,528 triplets by considering the tissue types as the classes.

TCGA dataset – The second histopathology dataset is The Cancer Genome Atlas (TCGA) dataset [23]. TCGA Whole Slide Images (WSIs) come from 25 different organs for 32 different cancer subtypes. We use the three most common sites, which are prostate, gastrointestinal, and lung [23], [24]. These organs have a total of 9 cancer subtypes, i.e., Prostate adenocarcinoma (PRAD), Testicular germ cell tumors (TGCT), Oesophageal carcinoma (ESCA), Stomach adenocarcinoma (STAD), Colonic adenocarcinoma (COAD), Rectal adenocarcinoma (READ), Lung adenocarcinoma (LUAD), Lung squamous cell carcinoma (LUSC), and Mesothelioma (MESO). By sampling patches from slides, we extracted 22,528 triplets to test the proposed losses with a large triplet sample size. The anchor and neighbor patches were selected from one WSI, but we used four ways of extraction of the distant patch, i.e., from the same WSI but far from the anchor, from another WSI of the same cancer subtype as an anchor, from another cancer subtype but the same anatomic site as anchor, and from another anatomic site.

B. Visual Comparison of Embeddings

In our experiments, we used ResNet-18 [25] as the backbone in our Siamese network structure (see Fig. 2). In our experiments, we set $q = 300$, $p = 128$, $b = 32$, and $\alpha = 0.25$. The learning rate was set to 10^{-5} in all experiments.

Embedding of MNIST data – The embeddings of the train/test sets of the MNIST dataset in the feature spaces of different loss functions are illustrated in Fig. 3 where $\lambda = 0.1$ was used for FDT and FDC. We used the Uniform Manifold Approximation and Projection (UMAP) [26] for visualizing the 128-dimensional embedded data. As can be seen, both embeddings of train and test data by the FDT loss are much more discriminating than the embedding of triplet loss. On the other hand, comparing the embeddings of contrastive and FDC losses shows that their performances are both good enough as the classes are well separated. Interestingly, the similar digits usually are embedded as close classes in the feature space, and this shows the meaningfulness of the trained subspace.

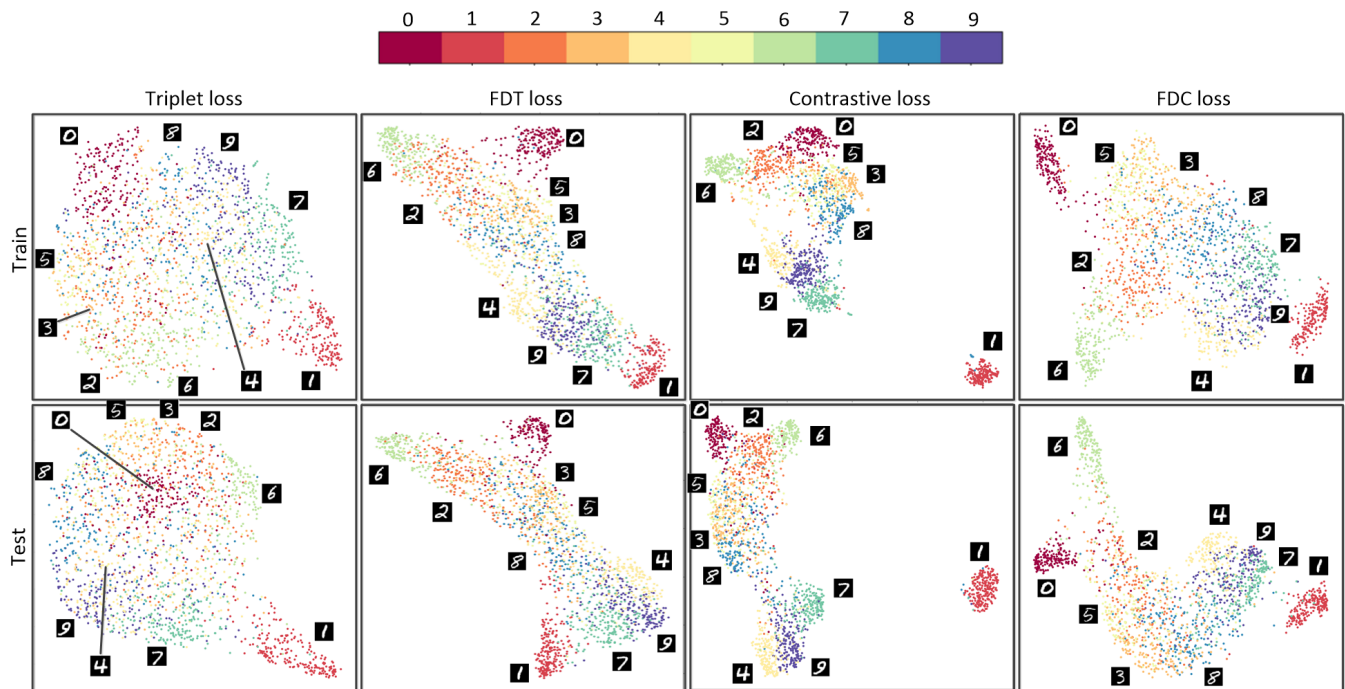


Fig. 3. Embedding of the training and test sets of MNIST dataset in the feature spaces of different loss functions.

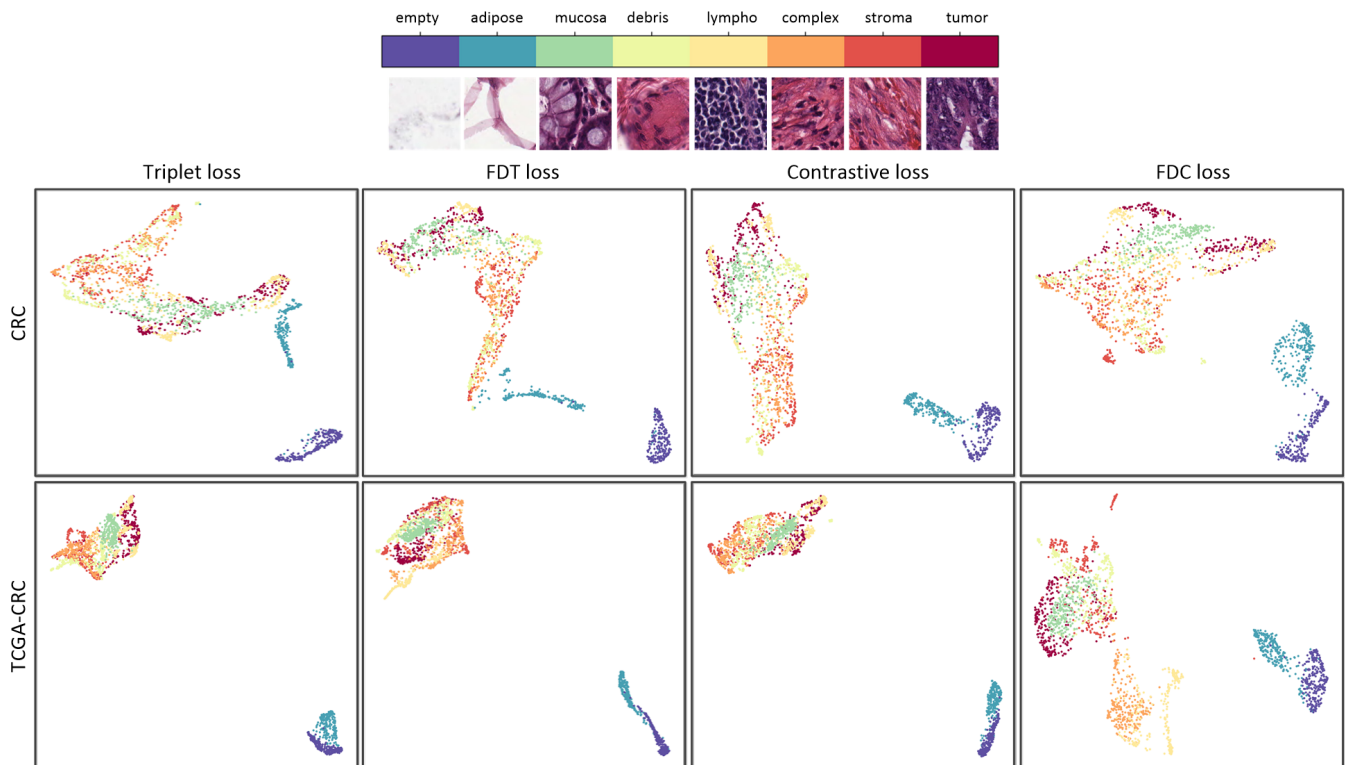


Fig. 4. Embedding of the CRC test data for different loss functions (top row: CRC, bottom row: TCGA).

For example, the digit pairs (3, 8), (1, 7), and (4, 9) with the second writing format of digit four can transition into each other by slight changes, and that is why they are embedded close together.

Embedding of histopathology data – For embedding of the histopathology data, we performed two different experiments. In the first experiment, we trained and tested the Siamese network using the CRC data. The second experiment was to

TABLE I
ACCURACY OF 1-NN SEARCH FOR DIFFERENT LOSS FUNCTIONS.

	MNIST	CRC	TCGA-CRC
triplet	82.21%	95.75%	95.50%
FDT ($\lambda = 0.01$)	82.76%	96.45%	97.60%
FDT ($\lambda = 0.1$)	85.74%	96.05%	96.40%
FDT ($\lambda = 0.8$)	79.59%	95.35%	95.95%
contrastive	89.99%	95.55%	96.55%
FDC ($\lambda = 0.01$)	78.47%	94.25%	96.55%
FDC ($\lambda = 0.1$)	89.00%	96.40%	98.10%
FDC ($\lambda = 0.8$)	87.71%	97.00%	97.05%

train the Siamese network using TCGA data and test it using the CRC test set. The latter, which we denote by TCGA-CRC, is more difficult because it tests generalization of the feature space, which is trained by different data from the test data, although with a similar texture. Figure 4 shows the embeddings of the CRC test sets in the feature spaces trained by CRC and TCGA data. The embeddings by all losses, including FDT and FDC, are acceptable, noticing that the histopathology data are hard to discriminate even by a human (see the sample patches in Fig. 4). As expected, the empty and adipose data, which are similar, are embedded closely. Comparing the TCGA-CRC embeddings of contrastive and FDC losses shows FDC has discriminated classes slightly better. Overall, the good embedding in TCGA-CRC shows that the proposed losses can train a generalizing feature space, which is very important in histopathology analysis because of the lack of labeled data [27].

C. Numerical Comparison of Embeddings

In addition to visualization, we can assess the embeddings numerically. For the evaluation of the embedded subspaces, we used the 1-Nearest Neighbor (1-NN) search because it is useful to evaluate the subspace by the closeness of the projected data samples. The accuracy rates of the 1-NN search for the embedding test data by different loss functions are reported in Table I. We report the results for different values of $\lambda \in \{0.01, 0.1, 0.8\}$ in order to analyze the effect of this hyper-parameter. As the results show, in most cases, the FDT and FDC losses have outperformed the triplet and contrastive losses, respectively. Moreover, we see that $\lambda = 0.1$ is often better performing. This can be because the large value of λ (e.g., 0.8) imposes less penalty on the total scatter, which may cause the embedding space to expand gradually. The very small value of λ (e.g., 0.01), on the other hand, puts too much emphasis on the total scatter where the classes do not tend to separate well enough, so they do not increase the total scatter.

D. Comparison of the Latent and Feature Spaces

As explained in Section III-A, the last layer behaves as a linear projection of the latent space onto the feature space. This projection fine-tunes the embeddings for better discrimination of classes. Figure 5 shows the latent embedding of the MNIST train set for both FDT and FDC loss functions. Comparing them to the feature embeddings of the MNIST train set in Fig.

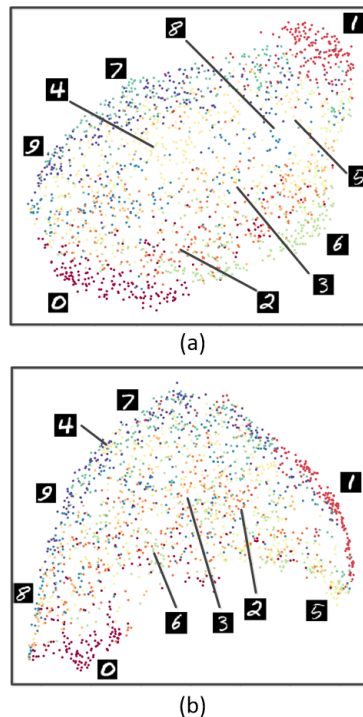


Fig. 5. The latent embedding of MNIST training for (a) FDT and (b) FDC.

3 shows that the feature embedding discriminates the classes much better than the latent embedding.

V. CONCLUSIONS

In this paper, we proposed two novel loss functions for training Siamese networks. These losses were based on the theory of the FDA, which attempts to decrease the intra-class scatter but increase the inter-class scatter of the projected data. The FDT and FDC losses make use of triplets and pairs of samples, respectively. By experimenting on MNIST and two histopathology datasets, we showed that the proposed losses mostly perform better than the well-known triplet and contrastive loss functions for Siamese networks.

REFERENCES

- [1] M. Ye, X. Zhang, P. C. Yuen, and S.-F. Chang, "Unsupervised embedding learning via invariant and spreading instance feature," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6210–6219.
- [2] B. Kumar, G. Carneiro, and I. Reid, "Learning local image descriptors with deep siamese and triplet convolutional networks by minimising global loss functions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5385–5394.
- [3] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition," in *ICML deep learning workshop*, vol. 2, 2015.
- [4] A. He, C. Luo, X. Tian, and W. Zeng, "A twofold siamese network for real-time object tracking," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4834–4843.
- [5] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, 2015, pp. 815–823.

- [6] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, vol. 2. IEEE, 2006, pp. 1735–1742.
- [7] J. Friedman, T. Hastie, and R. Tibshirani, *The elements of statistical learning*. Springer series in statistics New York, NY, USA, 2009, vol. 2.
- [8] R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Annals of eugenics*, vol. 7, no. 2, pp. 179–188, 1936.
- [9] B. Ghojogh, F. Karray, and M. Crowley, "Roweis discriminant analysis: A generalized subspace learning method," *arXiv preprint arXiv:1910.05437*, 2019.
- [10] A. Globerson and S. T. Roweis, "Metric learning by collapsing classes," in *Advances in neural information processing systems*, 2006, pp. 451–458.
- [11] D. Díaz-Vico, A. Omari, A. Torres-Barrán, and J. R. Dorronsoro, "Deep Fisher discriminant analysis," in *International Work-Conference on Artificial Neural Networks*. Springer, 2017, pp. 501–512.
- [12] D. Díaz-Vico and J. R. Dorronsoro, "Deep least squares Fisher discriminant analysis," *IEEE transactions on neural networks and learning systems*, 2019.
- [13] K. Fukunaga, *Introduction to statistical pattern recognition*. Elsevier, 2013.
- [14] B. Ghojogh, F. Karray, and M. Crowley, "Fisher and kernel Fisher discriminant analysis: Tutorial," *arXiv preprint arXiv:1906.09436*, 2019.
- [15] Y. Xu and G. Lu, "Analysis on Fisher discriminant criterion and linear separability of feature space," in *2006 International Conference on Computational Intelligence and Security*, vol. 2. IEEE, 2006, pp. 1671–1676.
- [16] B. Ghojogh, F. Karray, and M. Crowley, "Eigenvalue and generalized eigenvalue problems: Tutorial," *arXiv preprint arXiv:1903.11240*, 2019.
- [17] B. N. Parlett, "The symmetric eigenvalue problem," *Classics in Applied Mathematics*, vol. 20, 1998.
- [18] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [19] J. Ye, "Least squares linear discriminant analysis," in *Proceedings of the 24th international conference on machine learning*. ACM, 2007, pp. 1087–1093.
- [20] S. Mika, G. Ratsch, J. Weston, B. Scholkopf, and K.-R. Mullers, "Fisher discriminant analysis with kernels," in *Neural networks for signal processing IX: Proceedings of the 1999 IEEE signal processing society workshop (cat. no. 98th8468)*. IEEE, 1999, pp. 41–48.
- [21] Y. LeCun, C. Cortes, and C. J. Burges, "MNIST handwritten digits dataset," <http://yann.lecun.com/exdb/mnist/>, accessed: 2019.
- [22] J. N. Kather *et al.*, "Multi-class texture analysis in colorectal cancer histology," *Scientific reports*, vol. 6, p. 27988, 2016.
- [23] L. A. Cooper, E. G. Demicco, J. H. Saltz, R. T. Powell, A. Rao, and A. J. Lazar, "Pancancer insights from the cancer genome atlas: the pathologist's perspective," *The Journal of pathology*, vol. 244, no. 5, pp. 512–524, 2018.
- [24] S. Kalra *et al.*, "Pan-cancer diagnostic consensus through searching archival histopathology images using artificial intelligence," *npj Digital Medicine*, vol. 3, no. 1, pp. 1–15, 2020.
- [25] 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*, 2016, pp. 770–778.
- [26] L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform manifold approximation and projection for dimension reduction," *arXiv preprint arXiv:1802.03426*, 2018.
- [27] O. Jimenez-del Toro *et al.*, "Analysis of histopathology images: From traditional machine learning to deep learning," in *Biomedical Texture Analysis*. Elsevier, 2017, pp. 281–314.