

Cross-Domain Similarity Learning for Face Recognition in Unseen Domains

Masoud Faraki¹ Xiang Yu¹ Yi-Hsuan Tsai¹ Yumin Suh¹ Manmohan Chandraker^{1,2}
¹NEC Labs America
²University of California, San Diego
 {mfaraki, xiangyu, ytsai, yumin, manu}@nec-labs.com

Abstract

Face recognition models trained under the assumption of identical training and test distributions often suffer from poor generalization when faced with unknown variations, such as a novel ethnicity or unpredictable individual make-ups during test time. In this paper, we introduce a novel cross-domain metric learning loss, which we dub Cross-Domain Triplet (CDT) loss, to improve face recognition in unseen domains. The CDT loss encourages learning semantically meaningful features by enforcing compact feature clusters of identities from one domain, where the compactness is measured by underlying similarity metrics that belong to another training domain with different statistics. Intuitively, it discriminatively correlates explicit metrics derived from one domain, with triplet samples from another domain in a unified loss function to be minimized within a network, which leads to better alignment of the training domains. The network parameters are further enforced to learn generalized features under domain shift, in a model-agnostic learning pipeline. Unlike the recent work of Meta Face Recognition [18], our method does not require careful hard-pair sample mining and filtering strategy during training. Extensive experiments on various face recognition benchmarks show the superiority of our method in handling variations, compared to baseline and the state-of-the-art methods.

1. Introduction

Face recognition using deep neural networks has shown promising outcomes on popular evaluation benchmarks [21, 25, 26, 23]. Many current methods base their approaches on the assumption that the training data – CASIA-Webface [46] or MS-Celeb-1M [19] being the widely used ones – and the testing data have similar distributions. However, when deployed to real-world scenarios, those models often do not generalize well to test data with unknown statistics. In face recognition applications, this may mean a shift in attributes such as ethnicity, gender or age between the training and evaluation data. On the other hand, collecting and labelling more data along the underrepresented attributes

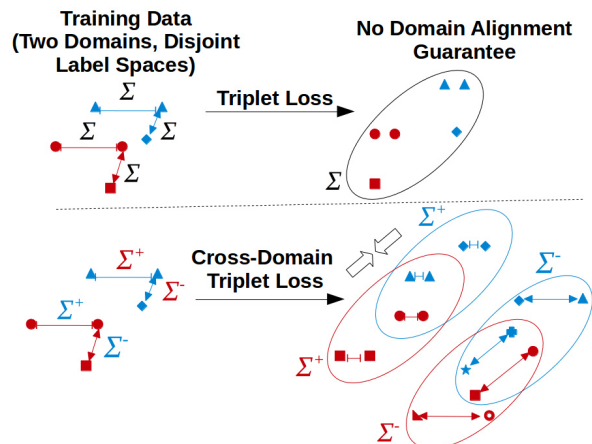


Figure 1. Comparison between the conventional triplet and our Cross-Domain Triplet losses. **Top:** The standard triplet loss is domain agnostic and utilizes a shared metric matrix, Σ , to measure distances of all (anchor, positive) and (anchor, negative) pairs. **Bottom:** Our proposed Cross-Domain Triplet loss, takes into account Σ^+ and Σ^- , i.e., the similarity metrics obtained from positive and negative pairs in one domain, to make compact clusters of triplets that belong to another domain. This, results to better alignment of the two domains. Here, colors indicate domains.

is costly. Therefore, given existing data, learning algorithms are needed to yield universal face representations and in turn, be applicable across such diverse scenarios.

Domain generalization has recently emerged to address the same challenge, but mainly for object classification with limited number of classes [3, 9, 32]. It aims to employ multiple labeled source domains with different distributions to learn a model that generalizes well to unseen target data at test time. However, many domain generalization methods are tailored to closed-set scenarios and hence, not directly applicable if the label spaces of the domains are disjoint. Generalized face recognition is indeed a prominent example of open-set applications with very large number of categories, encouraging the need for further research in this area.

In this paper, we introduce an approach to improve the problem of face recognition from unseen domains by learning semantically meaningful representations. To this end,

we are motivated by recent works in few-shot learning [31], domain generalization [9] and face recognition [22], revealing a general fact that, in training a model, it is beneficial to exploit notions of semantic consistency between training data coming from various sources. Therefore, we introduce Cross-Domain Triplet (CDT) loss based on the triplet objective [36], that learns useful features by considering two domains, where the similarity metrics provided by one domain are utilized in another domain to learn compact feature clusters of identities (Fig 1).

Such similarity metrics are encoded by means of covariance matrices, borrowing the idea from [31]. Different from [31], however, instead of using class-specific covariance matrices, we cast the problem in domain alignment regime, where our model first estimates feature distributions between the anchor and positive/negative samples, namely the similarity metrics derived from positive/negative pairs in one domain (*i.e.*, Σ^+ and Σ^- in Fig 1). Then, we utilize these similarity metrics and apply them to triplets of another domain to learn compact clusters. As supported by theoretical insights and experimental evaluations, our CDT loss aligns distributions of two domains in a discriminative manner. Furthermore, by leveraging a meta-learning framework, our network parameters are further enforced to learn generalized features under domain shift, following recent studies [9, 18].

Our experiments demonstrate the effectiveness of our approach equipped with the Cross-Domain Triplet loss, consistently outperforming the state-of-the-art on practical scenarios of face recognition for unknown ethnicity using the Cross-Ethnicity Faces (CEF) [39] and Racial Faces in-the-Wild (RFW) [44] benchmark datasets. Furthermore, it can satisfactorily handle face recognition across other variations as shown by empirical evaluations.

To summarize, we introduce an effective Cross-Domain Triplet loss function which utilizes explicit similarity metrics existing in one domain, to learn compact clusters of identities from another domain. This, results to learning semantically meaningful representations for face recognition from unseen domains. To further expose the network parameters to domain shift, under which more generalized features are obtained, we also incorporate the new loss in a model-agnostic learning pipeline. Our experiments show that our proposed method achieves state-of-the-art results on the standard face recognition from unseen domain datasets.

2. Related Work

Face Recognition. Following the success of deep neural networks, recent works on face recognition have tremendously improved the performances [46, 36, 35, 43, 7, 37, 6], thanks to large amount of labeled data being at the disposal. Many loss designs are shown to be effective in large scale network training, *i.e.*, CosFace [43] proposes to squeeze the classification boundary with a loss margin in cosine man-

ifold, ArcFace [7] combines the boundary margin with an angular margin to have better classification effect. Very recently, URFace [37] proposes sample-level confidence weighted cosine loss with an adversarial de-correlation loss to achieve better feature representations.

Generally, face recognition algorithms conjecture that the training data (e.g. CASIA-WebFace [46] or MS-Celeb-1M [19]) and testing data follow similar distributions. Recent works [18, 39] however, demonstrate unsatisfactory performance of such systems due to their poor generalization ability to handle unseen data in practice. This, makes face recognition models that are aware of test-time distribution changes, be more favorable. Furthermore, in minimizing distribution discrepancy, it is crucial to consider class information to avoid mis-alignment of samples from different categories in different domains [24].

Meta-learning and Domain Generalization. It is now an accepted fact that meta-learning (a.k.a. learning to learn [41]) can boost generalization ability of a model. The episodic training scheme originated from Model-Agnostic Meta-Learning (MAML) [16] has been widely used to address few shot learning [40, 38], domain generalization for object recognition from unseen distributions [29, 3, 32] and very recently face recognition from unseen domains [18]. The underlying idea is to simulate the train/test distribution gap in each round of training through creating episodic train/test splits out of the available training data. Some other examples include Model-Agnostic learning of Semantic Features (MASF) [9] which aligns a soft confusion matrix to retain knowledge about inter-class relationships, Feature Critic Networks [32] which proposes learning an auxiliary loss to help generalization and Meta-Learning Domain Generalization (MLDG) [28] which generates domain shift during training by synthesizing virtual domains within each batch.

Efforts in addressing unseen scenarios are either on transferring existing class variances to under-represented classes [47], or learning universal features through various augmentations [37]. A recent effort is Meta Face Recognition (MFR) [18], in which a loss is composed of distances of hard samples, identity classification and the distances between domain centers. However, simply enforcing alignment of the centers of training domains does not necessarily align their distributions and may lead to undesirable effects, *e.g.*, aligning different class samples from different domains [24]. As a result, this loss component does not always improve recognition (see w/o da. rows in Asian and Caucasian sections of Table 10 in [18]).

3. Proposed Method

In this section, we present our approach to improve the problem of face recognition from unseen domains by learning semantically meaningful representations. To do so, we

are inspired by recent works [31, 9, 22], showing that, in training a model, it is beneficial to exploit notions of semantic consistency between data coming from different distributions. We learn semantically meaningful features by enforcing compact clusters of identities from one domain, where the compactness is measured by underlying similarity metrics that belong to another domain with different statistics. In fact, we distill the knowledge encoded as similarity metrics across the domains with different label spaces.

We start with introducing the overall network architecture. Our architecture closely follows a typical image/face recognition design. It consists of a representation-learning network $f_r(\cdot; \theta_r)$, parametrized by θ_r , an embedding network $f_e(\cdot; \theta_e)$, parametrized by θ_e and a classifier network $f_c(\cdot; \theta_c)$, parametrized by θ_c . Following the standard setting [3, 9, 18], both $f_c(\cdot; \theta_c)$ and $f_e(\cdot; \theta_e)$ are light networks, a couple of Fully Connected (FC) layers, which take inputs from $f_r(\cdot; \theta_r)$. More specifically, forwarding an image I through $f_r(\cdot)$ outputs a tensor $f_r(I) \in \mathbb{R}^{H \times W \times D}$ that, after being flattened, acts as input to both the classifier $f_c(\cdot)$ and the embedding network $f_e(\cdot)$.

Before delving into more details, we first review some basic concepts used in our formulation. Then, we provide our main contribution to learn generalized features from multiple source domains with disjoint label spaces. Finally, we incorporate the solution into a model-agnostic algorithm, originally based on Model-Agnostic Meta-Learning (MAML) [16].

3.1. Notation and Preliminaries

Throughout the paper, we use bold lower-case letters (*e.g.*, \mathbf{x}) to show column vectors and bold upper-case letters (*e.g.*, \mathbf{X}) for matrices. The $d \times d$ identity matrix is denoted by \mathbf{I}_d . By a *tensor* \mathcal{X} , we mean a multi-dimensional array of order k , *i.e.*, $\mathcal{X} \in \mathbb{R}^{d_1 \times \dots \times d_k}$. $[\mathcal{X}]_{i,j,\dots,k}$ denotes the element at position $\{i, j, \dots, k\}$ in \mathcal{X} .

In Riemannian geometry, the Euclidean space \mathbb{R}^d is a Riemannian manifold equipped with the inner product defined as $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T \boldsymbol{\Sigma} \mathbf{y}$, $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ [11, 14]. The class of Mahalanobis distances in \mathbb{R}^d , $d : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^+$, is denoted by

$$d_{\boldsymbol{\Sigma}}(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x} - \mathbf{y})^T \boldsymbol{\Sigma} (\mathbf{x} - \mathbf{y})}, \quad (1)$$

where $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$ is a Positive Semi-Definite (PSD) matrix [27], region covariance matrix of features being one instance [10, 15]. This, boils down to the Euclidean (l_2) distance when the metric matrix, $\boldsymbol{\Sigma}$, is chosen to be \mathbf{I}_d . The motivation behind Mahalanobis metric learning is to determine $\boldsymbol{\Sigma}$ such that $d_{\boldsymbol{\Sigma}}(\cdot, \cdot)$ endows certain useful properties by expanding or shrinking axes in \mathbb{R}^d .

In a general deep neural network for metric learning, one relies on a FC layer with weight matrix $\mathbf{W} \in \mathbb{R}^{D \times d}$ immediately before a loss layer (*e.g.* contrastive [20] or triplet [36])

to provide the embeddings of the data to a reduced dimension space [12, 13, 36]. Then, given the fact that $\boldsymbol{\Sigma}$ is a PSD matrix and can be decomposed as $\boldsymbol{\Sigma} = \mathbf{W}^T \mathbf{W}$, the squared l_2 distance between two samples x and y (of a batch) passing through a network is computed as

$$\begin{aligned} d_{\boldsymbol{\Sigma}}^2(x, y) &= \|\mathbf{W}(f(x) - f(y))\|_2^2 \\ &= (f(x) - f(y))^T \boldsymbol{\Sigma} (f(x) - f(y)), \quad (2) \end{aligned}$$

where $f(x) \in \mathbb{R}^d$ denotes functionality of the network on an image x .

In this work, we use positive (negative) image pairs to denote face images with equal (different) identities. Moreover, a triplet, (anchor, positive, negative), consists of one anchor face image, another sample from the same identity and one image from a different identity.

3.2. Cross-Domain Similarity Learning

Here, we tackle the face recognition scenario where during training we observe k source domains, each with different attributes like ethnicity. At test time, a new target domain is presented to the network which has samples of individuals with different identities and attributes. We formulate this problem as optimizing a network using a novel loss based on the triplet loss [36] objective function, which we dub Cross-Domain Triplet loss. Cross-Domain Triplet loss, accepts inputs from two domains ${}^i\mathcal{D}$ and ${}^j\mathcal{D}$, estimates underlying distributions of positive and negative pairs from one domain (*e.g.*, ${}^j\mathcal{D}$), to measures the distances between (anchor, positive) and (anchor, negative) samples of the other domain (*e.g.*, ${}^i\mathcal{D}$), respectively. Then, using the computed distances and a pre-defined margin, the standard triplet loss function is applied.

Let ${}^j\mathbb{T} = \{{}^j(a, p, n)_b\}_{b=1}^{B_j}$ represents a batch of B_j triplets from the j -th domain, $j \in 1 \dots k$, from which we can consider positive samples ${}^j\mathbb{I}^+ = \{{}^j(a, p)_b\}_{b=1}^{B_j}$. For simplicity we drop the superscript j here. We combine all local descriptors of each image to estimate the underlying distribution by a covariance matrix. Specifically, we forward each positive image pair (a, p) , through $f_r(\cdot)$ to obtain the feature tensor representation $f_r(a), f_r(p) \in \mathbb{R}^{H \times W \times D}$. We cast the problem in the space of pairwise differences. Therefore, we define the tensor $\mathcal{R}^+ = f_r(a) - f_r(p)$. Next, we flatten the resulting tensor \mathcal{R}^+ into vectors $\{\mathbf{r}_i^+\}_{i=1}^{HW}$, $\mathbf{r}_i^+ \in \mathbb{R}^D$. This allows us to calculate a covariance matrix of positive pairs in pairwise difference space as

$$\boldsymbol{\Sigma}^+ = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{r}_i^+ - \boldsymbol{\mu}^+) (\mathbf{r}_i^+ - \boldsymbol{\mu}^+)^T, \quad (3)$$

where $N = BHW$ and $\boldsymbol{\mu}^+ = \frac{1}{N} \sum_{i=1}^N \mathbf{r}_i^+$.

Similarly, using B negative pairs $\mathbb{I}^- = \{(a, n)_b\}_{b=1}^B$, we find $\mathcal{R}^- = f_r(a) - f_r(n)$ for each (a, n) and flatten \mathcal{R}^- into vectors $\{\mathbf{r}_i^-\}_{i=1}^{HW}$, $\mathbf{r}_i^- \in \mathbb{R}^D$. This enables us to define a covariance matrix of negative pairs as

$$\Sigma^- = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{r}_i^- - \mu^-) (\mathbf{r}_i^- - \mu^-)^T, \quad (4)$$

where $N = BHW$ and $\mu^- = \frac{1}{N} \sum_{i=1}^N \mathbf{r}_i^-$.

Considering that a batch of images has adequate samples, this will make sure a valid PSD covariance matrix is obtained, since each face image has HW samples in covariance computations. Furthermore, samples from a large batch-size can satisfactorily approximate domain distributions according to [1, 8].

Our Cross Domain Triplet loss, l_{cdt} , works in a similar fashion by utilizing the distance function $d_{\Sigma}^2(\cdot, \cdot)$ defined in (2), to compute distance of samples using the similarity metrics from another domain. Given triplet images ${}^i\mathbb{T}$ from domain ${}^i\mathcal{D}$ and ${}^j\Sigma^+$, ${}^j\Sigma^-$ from domain ${}^j\mathcal{D}$ computed via (3) and (4), respectively, it is defined as

$$l_{cdt}({}^i\mathbb{T}, {}^j\mathbb{T}; \theta_r) = \frac{1}{B} \sum_{b=1}^B \left[\frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W d_{\Sigma^+}^2([f_r({}^i a_b)]_{h,w}, [f_r({}^j p_b)]_{h,w}) - \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W d_{\Sigma^-}^2([f_r({}^i a_b)]_{h,w}, [f_r({}^j n_b)]_{h,w}) + \tau \right]_+ \quad (5)$$

where τ is a pre-defined margin and $[\cdot]_+$ is the hinge function. We utilize class balanced sampling to provide inputs to both covariance and Cross-Domain Triplet loss calculations as this has been shown to be more effective in long-tailed recognition problems [18, 33].

Insights Behind our Method. Central to our proposal, is distance of the form $\mathbf{r}^T \Sigma \mathbf{r}$, defined on samples of two domains with different distributions. If \mathbf{r} is drawn from a distribution then the multiplications with Σ results in a distance according to the empirical covariance matrix, which optimizing over the entire points translates to alignment of the domains. More specifically, assuming that Σ is PSD, then eigendecomposition exists, i.e., $\Sigma = \mathbf{V} \Lambda \mathbf{V}^T$. Expanding the term leads to

$$\mathbf{r}^T \Sigma \mathbf{r} = (\Lambda^{\frac{1}{2}} \mathbf{V}^T \mathbf{r})^T (\Lambda^{\frac{1}{2}} \mathbf{V}^T \mathbf{r}) = \|\Lambda^{\frac{1}{2}} \mathbf{V}^T \mathbf{r}\|_2^2$$

which correlates \mathbf{r} with the eigenvectors of Σ weighted by the corresponding eigenvalues. This, attains its maximum when \mathbf{r} is in the direction of leading eigenvectors of the empirical covariance matrix Σ . In other words, as the eigenvectors of Σ are directions where its input data has

maximal variance, minimizing this term over \mathbf{r} vectors results to alignment of the two data sources. Fig 2 depicts the underlying process in our loss.

3.3. A Solution in a Model Agnostic Framework

Following recent trends in domain generalization tasks, we employ gradient-based meta-train/meta-test episodes under a model-agnostic learning framework to further expose the optimization process to distribution shift [9, 32, 18]. Algorithm (1) summarizes our overall training process. More specifically, in each round of training, we split input source domains into one meta-test and the remaining meta-train domains. We randomly sample B triplets from each domain to calculate our losses. First, we calculate two covariance matrices, Σ^+ and Σ^- , as well as temporary set of parameters, Θ' , based on the summation of a classification and triplet losses, L_s . The network is trained to semantically perform well on the held-out meta-test domain, hence Σ^+ , Σ^- and Θ' are used to compute the loss on the meta-test domain, L_t . This loss has additional CDT loss, l_{cdt} , to also involve cross-domain similarity for domain alignment. In the end, the model parameters are updated by accumulated gradients of both L_s and L_t , as this has been shown to be more effective than the original MAML [2, 18]. Here, the accumulated L_t loss provides extra regularization to update the model with higher-order gradients. In the following, we provide details of the identity classification and the triplet losses.

Having a classification training signal is crucial to face recognition applications. Hence, we use the standard Large Margin Cosine Loss (LMCL) [43] as our identity classification loss which is as follows

$$l_{cls}(I_i; \theta_r, \theta_c) = -\log \frac{\exp(s \mathbf{w}_{y_i}^T f_c(I_i) - m)}{\exp(s \mathbf{w}_{y_i}^T f_c(I_i) - m) + \sum_{y_j \neq y_i} \exp(s \mathbf{w}_{y_j}^T f_c(I_i))}, \quad (6)$$

where y_i is the ground-truth identity of the image I_i , $f_c(\cdot)$ is the classifier network, \mathbf{w}_{y_i} is the weight vector of the identity y_i in θ_c , s is a scaling multiplier and m is the margin.

We further encourage f_r network to learn locally compact semantic features according to identities from one domain. To this end, we use the triplet loss. Using the standard l_2 distance function $\|\cdot\|_2$, the triplet loss function provides training signal such that for each triplet, the distance between a and n becomes greater than the distance between a and p plus a predefined margin ρ . More formally,

$$l_{trp}(\mathbb{T}; \theta_r, \theta_e) = \frac{1}{B} \sum_{b=1}^B \left[\|f_e(a_b) - f_e(p_b)\|_2^2 - \|f_e(a_b) - f_e(n_b)\|_2^2 + \rho \right]_+. \quad (7)$$

Note that \mathbf{w}_y , $f_c(I)$ and $f_e(I)$ are l_2 normalized prior to computing the loss. Furthermore, $f_c(\cdot)$ and $f_e(\cdot)$ operate on the extracted representation by $f_r(\cdot)$.

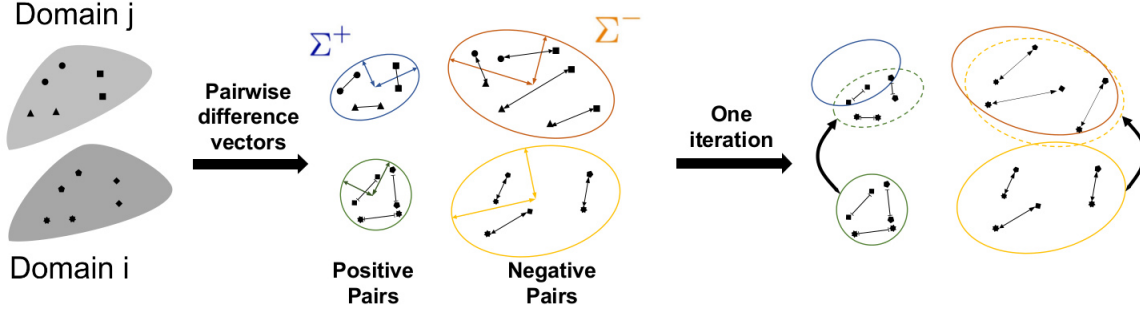


Figure 2. A schematic view of our Cross-Domain Triplet loss. In one iteration, given samples of two domains i and j with their associated (disjoint) labels are available, covariance matrices in difference spaces of positive and negative pairs of domain j are calculated and used to make positive and negative pairs of domain i , close and far away, respectively. This, makes compact clusters of identities while aligning the distributions. Note that alignments of positive and negative pairs are done simultaneously in a unified manner based on the triplet loss.

Algorithm 1 : Learning Generalized Features for Face Recognition.

Input:
Source domains $\mathcal{D} = [{}^1\mathcal{D}, {}^2\mathcal{D}_2, \dots, {}^k\mathcal{D}]$;
Batch-size B ;
Hyper-parameters α, β, λ ;

Output:
Learned parameters: $\hat{\Theta} = \{\hat{\theta}_r, \hat{\theta}_c, \hat{\theta}_e\}$

- 1: Initialize parameters $\Theta = \{\theta_r, \theta_c, \theta_e\}$
- 2: **repeat**
- 3: Initialize the gradient accumulator: $G_\Theta \leftarrow 0$
- 4: **for** each ${}^i\mathcal{D}$ (meta-test domain) in \mathcal{D} **do**
- 5: **for** each ${}^j\mathcal{D}$, $i \neq j$ (meta-train domain) in \mathcal{D} **do**
- 6: Sample B triplets ${}^i\mathbb{T}$, from B identities of ${}^i\mathcal{D}$
- 7: Sample B triplets ${}^j\mathbb{T}$, from B identities of ${}^j\mathcal{D}$
- 8: Compute $L_s \leftarrow \mathbb{E}[l_{cls}({}^j\mathbb{T}; \theta_r, \theta_c)] + l_{trp}({}^j\mathbb{T}; \theta_r, \theta_e)$
- 9: Compute $\Theta' \leftarrow \Theta - \alpha \nabla_{(\Theta)} L_s$
- 10: Compute ${}^j\Sigma^+$ and ${}^j\Sigma^-$ using positive and negative pairs of ${}^j\mathbb{T}$
- 11: Compute $L_t \leftarrow \mathbb{E}[l_{cls}({}^i\mathbb{T}; \theta'_r, \theta'_c)] + l_{trp}({}^i\mathbb{T}; \theta'_r, \theta'_e) + l_{cdt}({}^i\mathbb{T}, {}^j\mathbb{T}; \theta'_r)$
- 12: **end for**
- 13: $G_\Theta \leftarrow G_\Theta + \lambda \nabla_{\Theta} L_s + (1 - \lambda) \nabla_{\Theta} L_t$
- 14: **end for**
- 15: Update model parameters: $\Theta \leftarrow \Theta - \frac{\beta}{k} G_\Theta$
- 16: **until** convergence

4. Experiments

In this part, we first provide our implementation details for the sake of reproducibility. Then, we provide experimental results for face recognition from unseen domains. We conclude this section by ablation analysis of important components of our algorithm. To the best of our knowledge, Meta Face Recognition (MFR) [18] is the very recent work that addresses the same problem. Therefore, we compare our method against MFR in all evaluations. We also consider the performances of CosFace [43], Arcface [7] and URFace [37]

as our baseline.

Implementation Details. Our implementation was done in PyTorch [34]. For our baseline model, we utilized a 28-layer ResNet with a final FC layer that generates an embedding space of \mathbb{R}^{256} , *i.e.*, the final generalized feature space at test time. In our design, $f_r(\cdot; \theta_r)$ is the backbone immediately before this FC layer. $f_c(\cdot)$ generates logits immediately after it, while $f_e(\cdot)$ stacks an additional FC layer, to map inputs to a lower dimensional space of \mathbb{R}^{128} . As for the optimizer, we adopted stochastic gradient descent with the weight decay of 0.0005 and momentum of 0.9. In Algorithm (1), the batch-size B is set to 128, α and β to $1e^{-4}$ decaying to half every 1K steps and λ to 0.7. We set the classifier margin m to 0.5 and both margins τ and ρ to 1.

4.1. Cross Ethnicity Face Verification/Identification

As our first set of experiments, we tackled the problem of face recognition from unseen ethnicity. Here, the evaluation protocol is leave-one-ethnicity-out, *i.e.*, excluding samples of one ethnicity (domain) from training and then evaluating the performances on the samples of the held-out domain.

Datasets. For evaluation, there exist two recent datasets recommended for studying facial cross ethnicity recognition performances, namely the Cross-Ethnicity Faces (CEF) [39] and Racial Faces in-the-Wild (RFW) [44] datasets. CEF dataset is selected from MS-Celeb-1M [19], consisting of four types of ethnicity images, Caucasian, African-American, East-Asian and South-Asian. We combined the last two sets into one ethnicity domain, Asian. There are 200 identities with 10 different images from each ethnicity. Note that, the identities in CEF are disjoint from the MS-Celeb-1M dataset, which we use to train our model. Similarly, RFW is another cross-ethnicity evaluation benchmark made from the MS-Celeb-1M dataset with four ethnicity subsets: Caucasian, African, Asian and Indian. CEF has only test samples, while

Training Domain(s)	Verification			Identification		
	CA	AA	AS	CA	AA	AS
CA	97.8	92.0	93.2	90.0	69.2	75.7
AA	91.5	96.4	92.0	60.4	83.8	62.5
AS	93.3	91.8	95.6	63.9	64.5	84.1
All	98.4	97.1	97.0	90.2	84.0	84.4

Table 1. Verification and Identification accuracy numbers in % on the CEF [39] testing set using our annotated MS-Celeb-1M dataset with ethnicity labels. We report the evaluation metrics when either a single domain or all domains (*i.e.*, the last row) are used to train a CosFace model [43]. Testing data is either from CA: Caucasian, AA: African-American or AS: Asian domain, from disjoint identities to those of the training data. As shown in the results, regardless of the amount of data, a model trained on samples from a specific domain, performs much better on the testing data from the same domain. This, suggests the importance of having bias free training data when evaluating generalized face recognition algorithms. The top two results are highlighted in bold.

RFW provides both training and testing data across the domains.

Effect of Ethnicity Bias in Training Data. Note that, most public face datasets are collected from the web by querying celebrities. This, leads to significant performance bias towards the Caucasian ethnicity. For example, 82% of the data are Caucasian images in the MS-Celeb-1M dataset, while there are only 9.7% African-American, 6.4% East-Asian and less than 2% Latino and South-Asian combined altogether [39].

To further highlight the influence of having ethnicity bias in training data of face recognition models, we performed an experiment using our annotated MS-Celeb-1M dataset with ethnicity labels and the CEF test set. To train a model, we considered two cases either 1) only with training samples of a single test ethnicity or 2) all training samples from all of the ethnicities. We report two standard evaluation metrics, verification accuracy and identification accuracy. To calculate the verification accuracy, we follow the standard protocol suggested by [21], from which 10 splits are constructed and the overall average number is reported. Each split contains 900 positive and 900 negative pairs and the accuracy on each split is computed using the threshold found from the remaining 9 splits. As suggested by [39], in facial recognition systems, the drop in the identification accuracy is larger when dealing with face images from a novel ethnicity during test time. Therefore, we also report identification accuracy numbers suggested by [39]. More specifically, we consider a positive pair of reference and query to be correct, if there is no other image from a different identity that is closer in distance to the reference, than the query.

We show the results of this experiment in Table 1. Several conclusions can be drawn here. First, the results indicate that a network trained only on samples of a specific eth-

Unseen Ethnicity	Method	TAR@FAR's of			Rank@
		0.001	0.01	0.1	1
Caucasian	CosFace [43]	63.94	77.26	91.44	91.05
	Arcface [7]	63.98	77.43	91.92	91.15
	URFace [37]	64.10	77.81	92.27	91.76
	MFR [18]	64.31	79.89	93.44	92.90
	Ours	65.74	82.48	95.18	94.25
African-American	CosFace [43]	89.48	92.75	95.12	95.02
	Arcface [7]	89.69	92.31	95.55	95.66
	URFace [37]	89.75	92.76	96.23	95.70
	MFR [18]	90.66	95.01	97.69	97.07
	Ours	94.89	96.98	98.17	97.23
Asian	CosFace [43]	81.98	91.63	95.41	93.94
	Arcface [7]	82.96	91.81	96.09	93.92
	URFace [37]	84.33	92.44	96.10	94.37
	MFR [18]	85.54	94.73	96.86	95.41
	Ours	88.74	96.41	98.56	97.27
Indian	CosFace [43]	80.54	87.39	94.55	93.24
	Arcface [7]	82.17	88.58	95.24	93.41
	URFace [37]	83.32	89.11	95.30	93.53
	MFR [18]	84.17	89.42	96.08	93.94
	Ours	87.84	92.08	97.12	95.18

Table 2. Comparative results of our method against the state of the art on the four leave-one-ethnicity-out scenarios on the CEF test set [39]. Note that our method consistently outperforms the competitors.

nicity, tends to perform highly better on the same ethnicity, regardless of the amount of training samples. Second, the second best numbers in each column are very close to their corresponding upper bound shown in the last row ("All" here means when all training samples from all of the ethnicities are considered during training). Comparatively, the gaps between the seen domain and unseen test domains under the identification score are very larger, exceeding 20% between CA and AA when only Caucasian samples are used during training. This, clearly shows that such a bias may invalidate conclusions drawn from performance of facial recognition systems in unseen domains. A related study [42] shows the impact of noise in training data in face recognition.

We note that, for pre-training, the method of MFR utilizes the MS-Celeb-1M dataset, but without removing target ethnicity samples from training data. As we have experimentally observed and discussed above, this, causes difficulties when addressing generalized face recognition problem. Therefore, in this section, we adapt MFR, when considering face recognition from unseen domains. Thus, we first aim to remove such a bias from our training data. To this end, we train an ethnicity classifier network on manually annotated face images with ground-truth ethnicity label. The number of images per ethnicity varies between 5K to 8K. We then, make use of the network to find correct ethnicity label of each individual in the dataset. This allow us to remove all known samples of an specific domain when addressing cross-ethnicity face recognition.

Unseen Ethnicity	Method	TAR@FAR's of		
		0.001	0.01	0.1
Caucasian	CosFace [43]	61.15	70.74	85.50
	Arcface [7]	61.18	70.85	85.93
	URFace [37]	62.54	72.94	88.24
	MFR [18]	63.81	76.06	90.43
	Ours	65.20	78.80	91.87
African	CosFace [43]	71.55	83.40	90.11
	Arcface [7]	71.82	84.07	92.11
	URFace [37]	73.93	86.34	93.26
	MFR [18]	75.31	88.94	93.67
	Ours	77.90	91.17	96.87
Asian	CosFace [43]	66.33	80.04	90.26
	Arcface [7]	66.49	80.97	91.73
	URFace [37]	67.55	81.04	90.89
	MFR [18]	69.02	82.54	91.94
	Ours	69.17	82.60	94.10
Indian	CosFace [43]	65.55	80.17	91.03
	Arcface [7]	67.39	81.15	91.44
	URFace [37]	69.26	83.03	91.97
	MFR [18]	71.44	83.11	92.22
	Ours	76.63	86.70	95.47

Table 3. Comparative results of our method against the state of the art on the four leave-one-ethnicity-out scenarios of the RFW dataset [44]. Note that our method consistently outperforms the competitors.

Training Data. To train our model, we used our MS-Celeb-1M dataset annotated with ethnicity labels. In the case of RFW experiments, we further trained the model using RFW train set, while following the leave-one-ethnicity-out testing protocol. Where there is only one training sample per identity, we made use of random augmentation of the image, to generate a positive pair. The augmentation is a random combination of Gaussian blur and occlusion suggested by [37]. Since RFW has overlapping identities with MS-Celeb-1M, following MFR, we first removed samples of such identities and made MS-C-w/o-RFW, *i.e.*, MS-Celeb-without-RFW.

Evaluation Metrics. At test time, the extracted feature vectors of each image and its flipped version are concatenated and considered as the final representation of the image. We used cosine distance to compute the distances. As for evaluating performances, we report the True Acceptance Rate (TAR) at different levels of False Acceptance Rate (FAR) such as 0.001, 0.01 and 0.1 using the Receiver Operating Characteristic (ROC) curve. Moreover, we report Rank-1 accuracy where a pair of probe and gallery image is considered correct if after matching the probe image to all gallery images, the top-1 result is within the same identity.

Results on the CEF Testing Set. In Table 2, we compare our results with those of the state-of-the-art techniques across all four leave-one-domain-out scenarios. Note that there is

Method	TAR@FAR's of		Rank@1	AUC	Acc
	0.001	0.01			
CosFace [43]	92.55	96.55	96.85	99.53	99.35
LF-CNNs [45]	-	-	-	99.3	98.5
MFR [18]	94.05	97.25	97.8	99.8	99.78
Ours	94.20	97.22	97.75	99.8	99.76

Table 4. Comparative results of our method against the state of the art on cross-age face recognition using the CACD-VS dataset [4]. Our method works on par with MFR.

no annotated face image as Indian in the source domains here. Our method, equipped with the Cross-Domain Triplet loss, comfortably outperforms the recent method of MFR. This, we believe, clearly shows the benefits of our approach, which allows us to learn cross-domain similarities from labeled source face images, thus yielding a better representation for unseen domains.

Results on the RFW Dataset. In Table 3, using the RFW dataset, we compare our results with those of the state-of-the-art methods described above in the leave-one-ethnicity-out protocol. As in the previous experiment, our method is the top performing one in all cases. Under FAR@0.1 measure, the gap between our algorithm and the closest competitor, MFR, exceeds 3% when African-American and Indian images are considered as the test domain. CosFace method is considered as the baseline and is the least performing method, indicating the importance of learning generalized features in face recognition.

4.2. Handling Other Variations

To demonstrate that our approach is generic and can handle other common variations during test time, we considered additional experiments. Following the standard protocol suggested by MFR [18], the full train data annotated with ethnicity is used as our source domains in all experiments.

Cross age face verification and identification. First, we considered an experiment using the Cross-Age Celebrity Dataset (CACD-VS) [4]. The dataset provides 4k cross-age image pairs with equal number of positive and negative pairs to form a testing subset for face verification task. The results of this experiment are shown in Table 4, where we have also reported the Area Under the ROC Curve (AUC) as well as the verification accuracy numbers. The table indicates, while our method outperforms the recent method of MFR under the FAR@0.001 measure, works competitively very close to it under other evaluation metrics. Under the AUC measure, we obtain the exact same performance as the MFR. In all comparisons, our method outperforms the previous work of LF-CNNs [45].

Method	TAR@FAR's of			Rank@
	0.001	0.01	0.1	1
CosFace [43]	69.03	71.35	90.47	94.29
Arcface [7]	69.15	71.62	91.22	94.36
URFace [37]	70.79	74.67	93.40	94.72
MFR [18]	72.33	81.92	95.97	96.92
Ours	76.41	84.53	97.68	97.24

Table 5. Comparative results of our method against the state of the art on CASIA NIR-VIS 2.0 dataset [30].

Method	TAR@FAR's of			Rank@
	0.001	0.01	0.1	1
CosFace [43]	68.49	98.96	99.92	99.82
ArcFace [7]	69.53	99.0	99.96	99.86
URFace [37]	69.77	99.54	99.96	99.90
MFR [18]	74.54	99.96	100	99.92
Ours	76.62	99.96	100	99.92

Table 6. Comparative results of our method against the state of the art on Multi-PIE dataset [17].

Method	TAR@FAR's of		
	0.001	0.01	0.1
w/o l_{cls}	64.24	76.58	89.96
w/o l_{trp}	64.65	77.95	91.02
w/o l_{cdt}	63.75	75.90	89.84
Ours (full)	65.20	78.80	91.87

Table 7. Ablation study of the effects of different loss components of our proposal.

Near infrared vs. visible light face recognition. Next, we considered an experiment using the CASIA NIR-VIS 2.0 face database [30]. Here, the gallery face images are captured under visible light while the probe ones are under near infrared light. Table 5 shows that our method consistently outperforms the competitors.

Cross pose face recognition. Finally, we considered an experiment using the Multi-PIE cross pose dataset [17, 35]. Similar to the previous experiment, our method outperforms the competitors as shown in Table 6.

4.3. Ablation Analysis

In this section, we conduct further experiments to show the impact of important components of Algorithm 1 on verification accuracy. For this analysis, we made use of the Caucasian subset of the RFW dataset as the unseen domain and the rest of the data from all other domains as the source domains.

We first investigate how final recognition accuracy numbers vary, when different loss components are excluded from our algorithm. Different components here are identity classification, triplet and cross-domain triplet losses, denoted by l_{cls} , l_{trp} and l_{cdt} , respectively. The results of this experiment are shown in Table 7. As the table indicates, every component of our overall loss, contributes importantly to

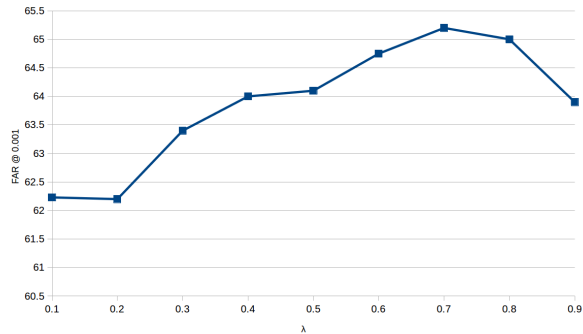


Figure 3. Ablation study over the hyper-parameter λ . This indicates the effect of meta-test loss which has our Cross-Domain Triplet loss. We empirically observed that a value close to 0.7 gives the best results.

the final performance, as excluding either of them leads to a consistent performance drop. Further comparing across the different loss terms, we observe that our proposed CDT loss plays a more important role here, since the performances significantly drop without this loss component. For instance, FAR@0.1 drops by more than 2% without l_{cdt} .

Furthermore, one hyper-parameter in our approach is the contribution ratio of meta-train and meta-test losses, *i.e.*, L_s and L_t . This is determined by hyper-parameter λ in Algorithm 1. In Fig 3, we show the effect of varying λ on the FAR@0.001 measure, when the Caucasian subset of the RFW dataset is considered the target domain. We observe that a value close to 0.7 gives the best results.

5. Conclusions

We have introduced a cross-domain metric learning loss, which is dubbed Cross-Domain Triplet (CDT) loss, that leverages the information jointly contained in two observed domains to provide better alignment of the domains. In essence, it first, takes into account similarity metrics of one data distribution, and then in a similar fashion to the triplet loss, uses the metrics to enforce compact feature clusters of identities that belong to another domain. Intuitively, CDT loss discriminatively correlates explicit metrics obtained from one domain with triplet samples from another domain in a unified loss function to be minimized within a network, which leads to better alignment of the training domains. We have also incorporated the loss in a meta-learning pipeline, to further enforce the network parameters to learn generalized features under domain shift. Extensive experiments on various face recognition benchmarks have shown the superiority of our method in handling variations, ethnicity being the most important one. In the future, we will investigate how a universal form of covariance matrix (such as the one used in [5]) could be utilized in our framework.

References

- [1] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *International conference on machine learning*, pages 1247–1255. PMLR, 2013. 4
- [2] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your maml. *arXiv preprint arXiv:1810.09502*, 2018. 4
- [3] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. In *Advances in Neural Information Processing Systems*, pages 998–1008, 2018. 1, 2, 3
- [4] Bor-Chun Chen, Chu-Song Chen, and Winston H Hsu. Cross-age reference coding for age-invariant face recognition and retrieval. In *European conference on computer vision*, pages 768–783. Springer, 2014. 7
- [5] Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu, and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 627–636, 2019. 8
- [6] Aruni Roy Chowdhury, Xiang Yu, Kihyuk Sohn, Erik Learned-Miller, and Manmohan Chandraker. Improving deep face recognition by clustering unlabeled faces in the wild. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. 2, 5, 6, 7, 8
- [8] Matthias Dorfer, Rainer Kelz, and Gerhard Widmer. Deep linear discriminant analysis. 2015. 4
- [9] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In *NIPS*, pages 6450–6461, 2019. 1, 2, 3, 4
- [10] Masoud Faraki, Mehrtash T Harandi, and Fatih Porikli. Material classification on symmetric positive definite manifolds. In *2015 IEEE Winter Conference on Applications of Computer Vision*, pages 749–756. IEEE, 2015. 3
- [11] Masoud Faraki, Mehrtash T Harandi, and Fatih Porikli. More about vlad: A leap from euclidean to riemannian manifolds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4951–4960, 2015. 3
- [12] Masoud Faraki, Mehrtash T Harandi, and Fatih Porikli. Large-scale metric learning: A voyage from shallow to deep. *IEEE transactions on neural networks and learning systems*, 29(9):4339–4346, 2017. 3
- [13] Masoud Faraki, Mehrtash T Harandi, and Fatih Porikli. No fuss metric learning, a hilbert space scenario. *Pattern Recognition Letters*, 98:83–89, 2017. 3
- [14] Masoud Faraki, Mehrtash T Harandi, and Fatih Porikli. A comprehensive look at coding techniques on riemannian manifolds. *IEEE transactions on neural networks and learning systems*, 29(11):5701–5712, 2018. 3
- [15] Masoud Faraki, Mehrtash T Harandi, Arnold Wiliem, and Brian C Lovell. Fisher tensors for classifying human epithelial cells. *Pattern Recognition*, 47(7):2348–2359, 2014. 3
- [16] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017. 2, 3
- [17] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-pie. *Image and vision computing*, 28(5):807–813, 2010. 8
- [18] Jianzhu Guo, Xiangyu Zhu, Chenxu Zhao, Dong Cao, Zhen Lei, and Stan Z Li. Learning meta face recognition in unseen domains. In *CVPR*, pages 6163–6172, 2020. 1, 2, 3, 4, 5, 6, 7, 8
- [19] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *European conference on computer vision*, pages 87–102. Springer, 2016. 1, 2, 5
- [20] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006. 3
- [21] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. 2008. 1, 6
- [22] Yuge Huang, Pengcheng Shen, Ying Tai, Shaoxin Li, Xiaoming Liu, Jilin Li, Feiyue Huang, and Rongrong Ji. Improving face recognition from hard samples via distribution distillation loss. In *ECCV*, pages 138–154. Springer, 2020. 2, 3
- [23] Nathan D Kalka, Brianna Maze, James A Duncan, Kevin O'Connor, Stephen Elliott, Kaleb Hebert, Julia Bryan, and Anil K Jain. Ijb-s: Iarpa janus surveillance video benchmark. In *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–9. IEEE, 2018. 1
- [24] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4893–4902, 2019. 2
- [25] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4873–4882, 2016. 1
- [26] Brendan F Klare, Ben Klein, Emma Taborsky, Austin Blanton, Jordan Cheney, Kristen Allen, Patrick Grother, Alan Mah, and Anil K Jain. Pushing the frontiers of unconstrained face detection and recognition: Iarpa janus benchmark a. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1931–1939, 2015. 1
- [27] Brian Kulis et al. Metric learning: A survey. *Foundations and trends in machine learning*, 5(4):287–364, 2012. 3
- [28] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. 2017. 2

- [29] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5400–5409, 2018. 2
- [30] Stan Li, Dong Yi, Zhen Lei, and Shengcai Liao. The casia nirvis 2.0 face database. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 348–353, 2013. 8
- [31] Wenbin Li, Jinglin Xu, Jing Huo, Lei Wang, Yang Gao, and Jiebo Luo. Distribution consistency based covariance metric networks for few-shot learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8642–8649, 2019. 2, 3
- [32] Yiyang Li, Yongxin Yang, Wei Zhou, and Timothy M Hospedales. Feature-critic networks for heterogeneous domain generalization. *arXiv preprint arXiv:1901.11448*, 2019. 1, 2, 4
- [33] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019. 4
- [34] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 5
- [35] Xi Peng, Xiang Yu, Kihyuk Sohn, Dimitris N. Metaxas, and Manmohan Chandraker. Reconstruction-based disentanglement for pose-invariant face recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 2, 8
- [36] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 2, 3
- [37] Yichun Shi, Xiang Yu, Kihyuk Sohn, Manmohan Chandraker, and Anil K Jain. Towards universal representation learning for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6817–6826, 2020. 2, 5, 6, 7, 8
- [38] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017. 2
- [39] Kihyuk Sohn, Wendy Shang, Xiang Yu, and Manmohan Chandraker. Unsupervised domain adaptation for distance metric learning. In *International conference on Learning Representation*, 2019. 2, 5, 6
- [40] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018. 2
- [41] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012. 2
- [42] Fei Wang, Liren Chen, Cheng Li, Shiyao Huang, Yanjie Chen, Chen Qian, and Chen Change Loy. The devil of face recognition is in the noise. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 765–780, 2018. 6
- [43] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018. 2, 4, 5, 6, 7, 8
- [44] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yao-hai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 692–702, 2019. 2, 5, 7
- [45] Yandong Wen, Zhifeng Li, and Yu Qiao. Latent factor guided convolutional neural networks for age-invariant face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4893–4901, 2016. 7
- [46] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014. 1, 2
- [47] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. Feature transfer learning for face recognition with under-represented data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2