

Generalizing Face Representation with Unlabeled Data

Yichun Shi and Anil K. Jain

Michigan State University, East Lansing, USA
shiyichu@msu.edu, jain@cse.msu.edu

Abstract. In recent years, significant progress has been made in face recognition due to the availability of large-scale labeled face datasets. However, since the faces in these datasets usually contain limited degree and types of variation, the models trained on them generalize poorly to more realistic unconstrained face datasets. While collecting labeled faces with larger variations could be helpful, it is practically infeasible due to privacy and labor cost. In comparison, it is easier to acquire a large number of unlabeled faces from different domains which would better represent the testing scenarios in real-world problems. We present an approach to use such unlabeled faces to learn generalizable face representations, which can be viewed as an unsupervised domain generalization framework. Experimental results on unconstrained datasets show that a small amount of unlabeled data with sufficient diversity can (i) lead to an appreciable gain in recognition performance and (ii) outperform the supervised baseline when combined with less than half of the labeled data. Compared with the state-of-the-art face recognition methods, our method further improves their performance on challenging benchmarks, such as IJB-B, IJB-C and IJB-S.

...

Keywords: Face Recognition, Unlabeled Data, Semi-supervised Learning, Domain Generalization, Image Translation

1 Introduction

Machine learning algorithms typically assumes that training and testing data come from the same underlying distribution. However, in practice, we would often encounter testing domains that are different from the population where the training data is drawn. Since it is non-trivial to collect data for all possible testing domains, learning representations that are generalizable to heterogeneous testing data is desired [28,8,27,22,4]. Particularly for face recognition, this problem is reflected by the domain gap between the semi-constrained training datasets and unconstrained testing datasets. Nearly all of the state-of-the-art deep face networks are trained on large-scale web-crawled face images, most of which are high-quality celebrity photos [50,10]. But in practice, we wish to deploy the trained FR systems for many other scenarios, e.g. unconstrained photos [19,26]

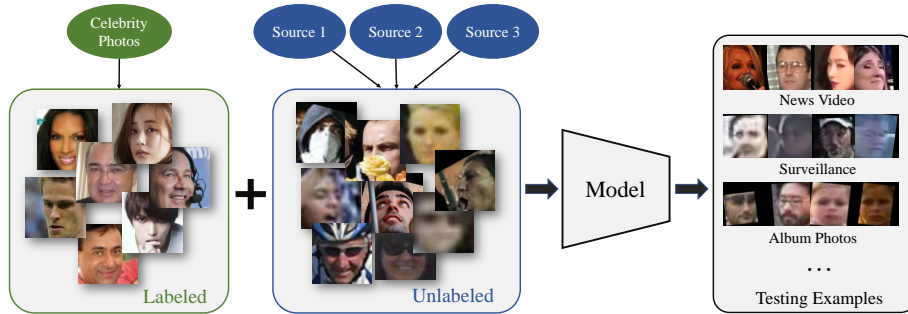


Fig. 1: Illustration of the problem setting in our work.

and surveillance [15]. The large amount of face variation, compared to the training set, in such testing domains results in significant performance drop of the trained face models [26,15].

The simplest solution to the problem is to collect a large number of unconstrained labeled face images from different sources. However, due to privacy issue and human-labeling cost, it is extremely hard to collect such a database. Transfer learning and domain adaptation are alternative methods to the problem. However, these methods require domain-specific data and need to train a model for each of the target domains [29,7,24,38,33,17]. Instead of such domain-specific models, a face representation that is robust to different kinds of variations would be preferred for the purpose of general unconstrained face recognition. We note that compared with building a labeled dataset, it is relatively easier to acquire a large number of unlabeled face images from multiple sources that are representative of testing samples that we would encounter in real-world scenarios. Therefore, it would be helpful if we can devise a method to utilize such unlabeled data to regularize the supervised learning of face representations to make them more generalizable (See Fig. 1).

In this paper, we propose such a semi-supervised framework for learning robust face representations. The unlabeled images are collected from a public face detection, i.e. WiderFace [49], dataset, which involve much larger degree of facial variations compared to typical labeled face datasets used for training. With a constrained labeled face dataset and an unconstrained unlabeled face dataset, we identify three objectives in the learning procedure:

- Maintaining the recognition performance on the labeled data
- Reducing feature domain gap between the labeled data and unlabeled data
- Increasing the discrimination power on the unlabeled data

In order to achieve these goals, the proposed method jointly regularizes the embedding model from feature space and image space. The adversarial regularization in the feature space helps to reduce the holistic domain gap caused by facial variations while increasing the inter-class distance. On other hand, the im-

age augmentation module maps the labeled training samples into the unlabeled domains, thus increasing the discrimination power on difficult face examples. The contributions of this paper is summarized as below:

- A semi-supervised learning framework for generalizing face representations with unlabeled data.
- A representation learning method of joint regularization from both image and feature domains.
- An multi-mode image translation module is proposed to perform data-driven augmentation and increase the diversity of the labeled training samples.
- Empirical results show that the regularization of unlabeled data helps to improve the recognition performance on challenging testing datasets, e.g. IJB-B, IJB-C, and IJB-S.

2 Related Work

2.1 Deep Face Recognition

Deep neural networks are widely adopted in the ongoing research in face recognition [41,39,34,25,23,11,31,44,5]. Taigman et al. [41] were the first to propose using deep convolutional neural network for learning face representations. The subsequent studies have explored different loss functions to improve the discrimination power of the learned feature representation. A number of studies proposed to use metric learning methods for face recognition [34,36]. Recent work has been trying to achieve discriminative embeddings with a single identification loss function where proxy/prototype vectors are used to represent each class in the embedding space [23,44,45,31,5,52,40].

2.2 Semi-supervised Learning

Classic semi-supervised learning involves a small number of labeled images and a large number of unlabeled images [21,32,20,42,47,51,2,37]. The goal is to improve the recognition performance when we don't have sufficient data that are labeled. State-of-the-art semi-supervised learning methods can mainly be classified into four categories. (1) Pseudo-labeling methods generate labels for unlabeled data with the trained model and then use them for training [21]. In spite of its simplicity, it has been shown to be effective primarily for classification tasks where labeled data and unlabeled data share the same label space. (2) Temporal ensemble models maintain different versions of model parameters to serve as teacher models for the current model [20,42]. (3) Consistency-regularization methods apply certain types of augmentation to the unlabeled data while making sure the output prediction remains consistent after augmentation [32,2,37]. (4) Self-supervised learning, originally proposed for unsupervised learning, has recently been shown to be effective for semi-supervised learning as well [51]. Compared with classic semi-supervised learning addressed in the literature, our problem is different in two sense of heterogeneity: different domains and different identities between the labeled and unlabeled data. These differences make many classic semi-supervised learning methods unsuitable for our task.

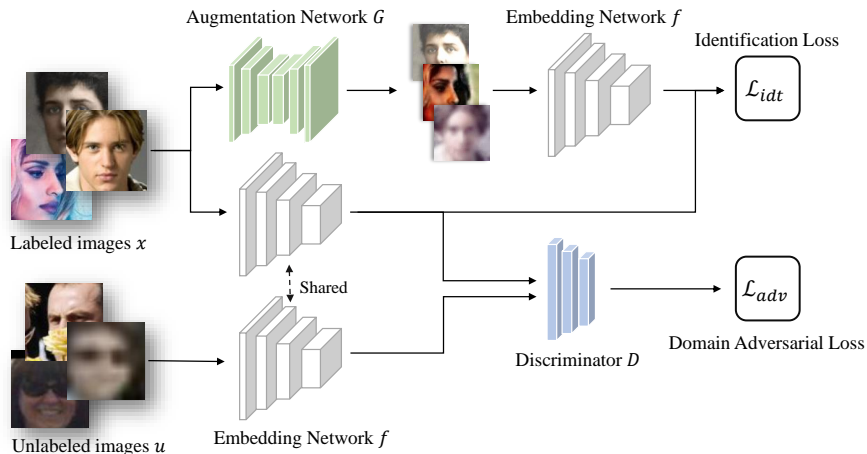


Fig. 2: Overview of the training framework of the embedding network. In each mini-batch, a random subset of labeled data would be augmented by the augmentation network to introduce additional diversity. The non-augmented labeled data are used to train the feature discriminator. The adversarial loss forces the distribution of the unlabeled features to align with the labeled one.

2.3 Domain Adaptation and Generalization

In domain adaptation, the user has a dataset for a source domain and another for a fixed target domain [29,7,24,33,17]. If the target domain is unlabeled, this leads to an *unsupervised domain adaption* setting [7,33,17]. The goal is to improve the performance on the target domain so that it could match the performance on the source domain. This is achieved by reducing the domain gap between the two datasets in feature space. The problem about domain adaption is that one needs to acquire a new dataset and train a new model whenever there is a new target domain. In *domain generalization*, the user is given a set of labeled datasets from different domains. The model is jointly trained on these datasets so that it could better generalize to unseen domains [28,8,27,22,4]. Our problem shares the same goal with domain generalization methods: *we want to increase the model generalizability rather than improving performance on a specific target domain*. However, unlike domain generalization, we do not have identity labels for all the data, which makes our task even more difficult.

3 Methodology

Generally, in face representation learning, we are given a large labeled dataset $\mathcal{X}=\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$, where x_i and y_i are the face images and identity labels, respectively. The goal is to learn an embedding model f such that $f(x)$ would be discriminative enough to distinguish between different identities. However, since f is only trained on the domain defined by \mathcal{X} , which is usually

semi-constrained celebrity photo, it might not generalize to unconstrained settings. In our framework, we assume the availability of another unlabeled dataset $\mathcal{U}=\{u_1, u_2, \dots, u_n\}$, collected from sources different than \mathcal{X} . The face images in \mathcal{U} contain more variations than images in \mathcal{X} and therefore are closer to the images in the unconstrained testing scenarios. We attempt to utilize such an unlabeled dataset \mathcal{U} to regularize the training of f to make it more generalizable. In particular, we wish to simultaneously minimize three types of errors:

- Error due to discrimination power within the labeled domain \mathcal{X} .
- Error due to feature domain gap between \mathcal{X} and \mathcal{U} .
- Error due to discrimination power within the unlabeled domain \mathcal{U} .

The first type of error could be minimized by the identification loss of the labeled images (Sec. 3.1). For the second type of error, we choose to use adversarial learning to minimize the distance between the embedding distributions (Sec. 3.2). For the third type of error, we study two different types of augmentation-based approaches (Sec. 3.3). An overview of the framework is shown in Fig. 2.

3.1 Minimizing Error in the Labeled Domain

The deep representation of a face image is usually a point in a hyper-spherical embedding space, where $\|f(x_i)\|^2 = 1$. State-of-the-art supervised face recognition methods all try to find an objective function to maximize the inter-class margin such that the representation could still be discriminative when tested on unseen identities. In this work, we choose to use CosFace loss function [45][44] for training the labeled images:

$$\mathcal{L}_{idt} = -\mathbb{E}_{x_i, y_i \sim \mathcal{X}} \left[\log \frac{\exp(sW_{y_i}^T f_i - m)}{\exp(sW_{y_i}^T f_i - m) + \sum_{j \neq y_i} \exp(sW_j^T f_i)} \right]. \quad (1)$$

Here s is the hyper-parameter controlling temperature, m is a margin hyper-parameter and W_j is the proxy vector of the j^{th} identity in the embedding space, which is also ℓ_2 normalized. We choose to use CosFace loss function because of its stability and high-performance. It could potentially be replaced by any other supervised identification loss function.

3.2 Minimizing Domain Gap

As shown in Fig. 3 (a), the current face representations usually show a large domain gap between different types of face images, which makes the cross-domain face matching a hard problem. In order to reduce this type of error, we use a feature discriminator network to align the embeddings of the two domains with an adversarial loss, which is given by

$$\mathcal{L}_D = -\mathbb{E}_{x \sim \mathcal{X}} [\log D(y = 0 | f(x))] - \mathbb{E}_{u \sim \mathcal{U}} [\log D(y = 1 | f(u))], \quad (2)$$

$$\mathcal{L}_{adv} = -\mathbb{E}_{x \sim \mathcal{X}} [\log D(y = 1 | f(x))] - \mathbb{E}_{u \sim \mathcal{U}} [\log D(y = 0 | f(u))]. \quad (3)$$

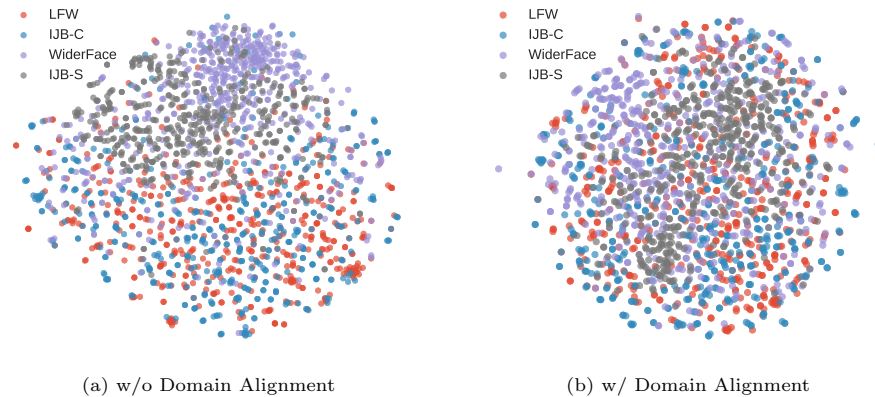


Fig. 3: t-SNE visualization of the face embeddings of different datasets. (a) Clear domain gap can be observed with the baseline model. (b) Domain gap between different datasets are reduced with the domain alignment loss. Here, LFW represents the constrained face images in the labeled domain.

The discriminator D is a multi-layer binary classifier optimized by \mathcal{L}_D . It tries to learn a non-linear classification boundary between the two datasets while the embedding network need to fool the discriminator by reducing the divergence between the distributions of $f(x)$ and $f(u)$. The effect of domain alignment loss is shown in Fig. 3 (b). Note that here Widerface is our unlabeled training set and LFW can represent the domain of the labeled training set. After alignment, the features in the unlabeled domains are forced to be more spread out, leading to a larger inter-class distance (See Fig. 4 (a)). Furthermore, given that WiderFace is a diverse dataset covering different types of faces, the reduction of domain gap also generalizes to other types of datasets, e.g. IJB-S.

3.3 Minimizing Error in the Unlabeled Domains

As shown in Figure 4, the domain alignment loss helps to increase the inter-class variance on the unconstrained domains. Thus, the remaining task to improve the discrimination power is to reduce the intra-class variance, for which multiple images of each class is necessary. Many semi-supervised classification methods address this problem by using pseudo-labeling of unlabeled data [21,2,37], but this is not applicable to our problem since our unlabeled dataset do not share the same label space with the labeled one. Furthermore, because of data collection protocols, it is expected that there is very little chance that one identity would have multiple unlabeled images. Thus, clustering-based methods also becomes infeasible for our task. Here, we consider two methods to address this issue:

Consistency Regularization The first solution we consider is a self-supervised method that is recently in semi-supervised learning [32,37]. The idea is to gen-

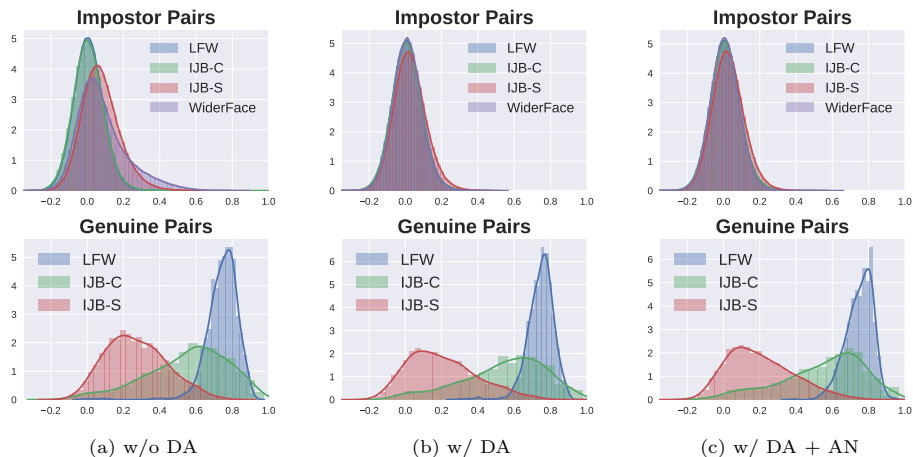


Fig. 4: The histogram of pairwise scores within different datasets. Fig (a) shows the impostor (up) and genuine (down) scores of the baseline method. Fig (b) shows the scores of the model trained with domain alignment (DA). Domain alignment helps to shift the score distribution of impostor pairs to be lower, i.e. it increases the inter-class variance. Fig (c) shows the score distribution of further incorporating the augmentation network (AN).

erate ground-truth positive pairs from unlabeled data using data perturbation. Given the same input sample, the model should give consistent outputs regardless of different perturbation. Since our focus is on learning the embedding instead of classification, we would like to keep the consistency in the feature space. In particular, let us define a stochastic perturbation function \mathcal{A} , which is manually designed. The consistency regularization loss would be:

$$\mathcal{L}_{CR} = \mathbb{E}_{u \sim \mathcal{U}}[\|f(u) - f(\mathcal{A}(u))\|_2^2], \quad (4)$$

Data-driven Augmentation Opposite to the consistency regularization, another approach we consider is to leverage the labeled images to create positive pairs in the unlabeled domains. In particular, we need a function G which maps labeled samples x into the image space defined by the unlabeled faces, i.e. $p(G(x)) \rightarrow p(u)$. Then, training the embedding f on $G(x)$ could make it more discriminative in the image space defined by U . There are two requirements of the function G : (1) it should not change the identity of the input image and (2) it should be able to capture different styles that are present in the unlabeled images. Inspired by recent progress in image translation frameworks [53,14], we propose to train G as a style-transfer network that learns the texture styles in the unlabeled domain and transfer it to the labeled images. The network G can then be used as a data-driven augmentation module that generates diverse samples given an input from the labeled dataset. During the training, we randomly replace a subset

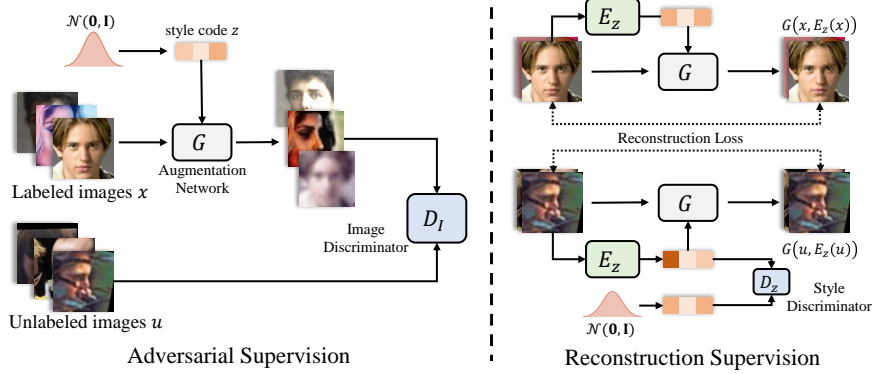


Fig. 5: Training framework of the augmentation network G . The two pipelines are optimized jointly during training.

of the labeled images to be augmented and put them into our identification learning framework. The details of training the augmentation network G is given in Section 3.4.

The overall loss function for the embedding network is given by:

$$\mathcal{L} = \lambda_{idt}\mathcal{L}_{idt} + \lambda_{adv}\mathcal{L}_{adv} + \lambda_{CR}\mathcal{L}_{CR} \quad (5)$$

where \mathcal{L}_{idt} also includes the augmented labeled samples.

3.4 Multi-mode Augmentation Network

The augmentation network G is a fully convolutional network that maps an image to another. To preserve the geometric structure, our architecture does not involve any downsampling or upsampling. In order to generate styles similar to the unlabeled images, an image discriminator D_I is trained to distinguish between the texture styles of unlabeled images and generated images:

$$\mathcal{L}_{D_I} = -\mathbb{E}_{x \sim \mathcal{X}}[\log D_I(y = 0 | G(x, z))] - \mathbb{E}_{u \sim \mathcal{U}}[\log D_I(y = 1 | u)], \quad (6)$$

$$\mathcal{L}_{adv}^G = -\mathbb{E}_{x \sim \mathcal{X}}[\log D_I(y = 1 | G(x, z))]. \quad (7)$$

Here $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is a random style vector to control the styles of the output image, which is injected into the generation process via Adaptive Instance Normalization (AdaIN) [13]. Although the adversarial learning could make sure the output are in the unlabeled space, but it cannot ensure that (1) the content of the input is maintained in the output image and (2) the random style z is being used to generate diverse visual styles. Inspired by recent work on unsupervised image-to-image translation [53,14], we propose to use use an additional reconstruction pipeline to simultaneously satisfy these two requirements. First, we introduce an additional style encoder E_z to capture the corresponding style in



Fig. 6: Example generated images of the augmentation network.

the input image. A reconstruction loss is then enforced to keep the consistency of the image content:

$$\mathcal{L}_{rec}^G = \mathbb{E}_{x \sim \mathcal{X}}[\|x - G(x, E_z(x))\|^2] + \mathbb{E}_{u \sim \mathcal{U}}[\|u - G(u, E_z(u))\|^2], \quad (8)$$

Second, during the reconstruction, we add another latent style discriminator D_z to guarantee the distribution of $E_z(u)$ align with prior distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$:

$$\mathcal{L}_{D_z} = -\mathbb{E}_{u \sim \mathcal{U}}[\log D_z(y = 0 | E_z(u))] - \mathbb{E}_{z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})}[\log D_z(y = 1 | z)], \quad (9)$$

$$\mathcal{L}_{adv}^z = -\mathbb{E}_{u \sim \mathcal{U}}[\log D_z(y = 1 | E_z(u))], \quad (10)$$

The overall loss function of the generator is given by:

$$\mathcal{L}^G = \lambda_{adv}^G \mathcal{L}_{adv}^G + \lambda_{rec}^G \mathcal{L}_{rec}^G + \lambda_{adv}^z \mathcal{L}_{adv}^z \quad (11)$$

An overview of the training framework of G is given in Fig. 5 and example generated images are shown in Fig. 6. The architecture details of different modules are given in the supplementary file.

4 Experiments

4.1 Implementation Details

Training Details of the Recognition Models All the models are implemented with Pytorch v1.1. We use the RetinaFace [6] for face detection and alignment. All images are transformed into a size of 112×112 . A modified 50-layer ResNet in [5] is used as our architecture. The embedding size is 512 for all models. By default, all the models are trained with 150,000 steps with a batch size of 256. For semi-supervised models, we use 64 unlabeled images and 192 labeled images in each mini-batch. For models which uses the augmentation module, 20% of the labeled images are augmented by the generator network. The scale parameter s and margin parameter m are set to 30 and 0.5, respectively. We empirically set λ_{idt} , λ_{adv} as 1.0 and 0.01. For models that utilizes the consistency regularization, λ_{CR} is set to 0.2. Random image translation, flipping, occlusion and downsampling are used as data perturbation for those models.



Fig. 7: Examples face images of different datasets used in this work.

Training Details of the Generator Models The generator is trained for 160,000 steps with a batch size of 8 images (4 from each dataset). Adam optimizer is used with $\beta_1 = 0.5$ and $\beta_2 = 0.99$. The learning rate starts with $1e - 4$ and drops to $1e - 5$ after 80,000 steps. The detailed architectures are provided in the supplementary material. λ_{adv}^G , λ_{rec}^G and λ_{adv}^z are set to as 1.0, 10.0 and 1.0, respectively.

4.2 Datasets

We use **MS-Celeb-1M** [10] as our labeled training dataset. MS-Celeb-1M is a large-scale public face datasets of celebrity photos. The original dataset is known to contain a large number of noisy labels [3], so we use a cleaned version from ArcFace [5] as training data. After removing the overlapped subjects with the testing sets and duplicate images, we are left 3.9M images of 85.7K classes. As for unlabeled images, we choose **WiderFace** [49] as our training data. WiderFace is dataset collected by retrieving images from search engines with different event keywords. As a face detection dataset, it includes a much wider domain of photos and the faces (See Fig. 7). Many faces in this dataset still cannot be detected by state-of-the-art detection methods [6]. We only keep the detectable faces in the WiderFace training set as our training data. Our goal is to close the gap between face detection and recognition engine and improve the recognition performance on a general settings with any detectable faces. At the end, we were able to detect about 70K faces from WiderFace, less than 2% of our labeled training data.

To evaluate the face representation models, we test on three benchmarks, namely IJB-B, IJB-C and IJB-S. These datasets represent real-world testing scenarios where face images show different characteristics from celebrity photos in the training set. The details of these datasets are as follows:

- **IJB-B** [46] includes both high quality celebrity photos taken from the wild and low quality photos or video frames with large variations of illumination, occlusion, head pose, etc. There are 68,195 images of 1,845 identities in all. We test on both verification and identification protocols of the IJB-B benchmark.
- **IJB-C** [26] is a newer version of IJB-B dataset. It has a similar protocol but with 140,732 images of 3,531 identities.

- **IJB-S** [15] is an extremely challenging benchmarks where the images were collected from surveillance cameras. There are in all 202 identities with an average of 12 videos per person. Each person also has 7 high-quality enrollment photos (with different poses) which constitute the gallery. We test on two protocols of the IJB-S dataset, Surveillance-to-Still (**V2S**) and Surveillance-to-Booking (**V2B**), both of which are identification protocols. The difference between them is that in Surveillance-to-Still (V2S) the gallery of each person is a single frontal photo while Surveillance-to-Booking (V2B) uses all 7 registration photos as gallery. To Reduce the evaluation time, the experiments in Sec. 4.3 and Sec. 4.4 are conducted with subsampled frames from each video, whose performance is close to using the whole video (Sec. 4.5).

Although our goal is to improve recognition performance on domains that are different from the training set, we would not like to lose the discrimination power in the original domain (high-quality photos) either. Therefore, during ablation we also evaluate our models on the standard **LFW** [12] protocol, which is a celebrity photo dataset, similar to the labeled training data (MS-Celeb-1M). Note that the accuracy on the LFW protocol is highly saturated, so the main goal is just to check whether there is a significant performance drop on the constrained faces while increasing the generalizability to unconstrained ones. Example images of different datasets are shown in Figure 7.

4.3 Ablation Study

In this section, we conduct an ablation study to quantitatively evaluate the effect of different modules proposed in this paper. In particular, we have three modules to study: Domain Alignment (DA), Consistency Regularization (CR) and Augmentation Network (AN). The performance is shown in Table 1. As we already showed in Fig. 3 and Fig. 4, domain adversarial loss is able to force the embeddings of face images to increase the intra-class distance. Consequently, we observe the performance improvement on the most of the protocols on IJB-C and IJB-S. Being Combined with the domain alignment module, both the consistency regularization strategy and augmentation network method could lead to further improvement. In particular, augmentation network leads to a more stable improvement across the protocols. However, when we attempt to combine these two modules, we observe that the performance actually becomes lower, which means the constraints enforced by the two modules might not be independent from each other. Therefore, for the following experiment, we use the “DA+AN” as our default method.

4.4 Quantity vs. Diversity

Although we have shown in Sec. 4.3 that utilizing unlabeled data leads to better performance on challenging testing benchmarks, generally it shall be expected that simply increasing the number of labeled training data can also have a similar effect. Therefore, in this section, we conduct a more detailed study to answer

Table 1: Ablation study over different training methods of the embedding network. All models has identification loss by default. “DA”, “CR”, “AN” refer to “Domain Alignment”, “Consistency Regularization” and “Augmentation Network”, respectively.

Method	IJB-C (Vrf)			IJB-C (Idt)		IJB-S (V2S)		LFW
	1e-7	1e-6	1e-5	Rank1	Rank5	Rank1	Rank5	Accuracy
Baseline	62.90	82.94	90.73	94.90	96.77	53.23	62.91	99.80
DA	72.74	85.33	90.52	94.99	96.75	56.35	66.77	99.82
DA + CR	73.36	87.53	91.65	95.34	96.97	55.53	65.45	99.73
DA + AN (proposed)	77.39	87.92	91.86	95.61	97.13	57.33	65.37	99.75
DA + CR + AN	73.39	86.60	91.41	95.48	97.03	57.14	65.69	99.72




Fig. 8: Evaluation Results on IJB-C and IJB-S with different protocols and different number of labeled training data.

such a question: *which is more important for feature generalizability: quantity or diversity of the training data?* In particular, we train several supervised models by adjusting the number of labeled training data. For each such model, we also train a corresponding model with additional unlabeled data. The evaluation results are shown in Figure 8.

On the IJB-S dataset, which is significantly different from the labeled training data, we see that the models trained with unlabeled data consistently outperforms the supervised baselines with a large margin. In particular, the proposed method achieves better performance than the supervised baseline even when there is only one-fourth of the overall labeled training data (1M vs 4M), indicating the value of data diversity during training. Note that there is a significant performance boost when increasing the number of labeled samples from 0.5M to 1M. However, after that, the benefit of acquiring more labeled data plateaus and in fact it is more helpful to introduce 70K unlabeled data than 3M additional labeled data.

Table 2: Performance comparison with state-of-the-art methods on the IJB-C dataset.

Method	Data	Model	Verification				Identification	
			1e-7	1e-6	1e-5	1e-4	Rank1	Rank5
Cao et al. [3]	13.3M	SE-ResNet-50	-	-	76.8	86.2	91.4	95.1
PFE [35]	4.4M	ResNet-64	-	-	89.64	93.25	95.49	97.17
ArcFace [5]	5.8M	ResNet-50	67.40	80.52	88.36	92.52	93.26	95.33
Ranjan et al. [30]	5.6M	ResNet-101	67.4	76.4	86.2	91.9	94.6	97.5
AFRN [16]	3.1M	ResNet-101	-	-	88.3	93.0	95.7	97.6
Baseline	3.9M	ResNet-50	62.90	82.94	90.73	94.57	94.90	96.77
Proposed	4.0M	ResNet-50	77.39	87.92	91.86	94.66	95.61	97.13

Table 3: Performance comparison with state-of-the-art methods on the IJB-B dataset.

Method	Data	Model	Verification				Identification	
			1e-6	1e-5	1e-4	1e-3	Rank1	Rank5
Cao et al. [3]	13.3M	SE-ResNet-50	-	70.5	83.1	90.8	90.2	94.6
Comparator [48]	3.3M	ResNet-50	-	-	84.9	93.7	-	-
ArcFace [5]	5.8M	ResNet-50	40.77	84.28	91.66	94.81	92.95	95.60
Ranjan et al. [30]	5.6M	ResNet-101	48.4	80.4	89.8	94.4	93.3	96.6
AFRN [16]	3.1M	ResNet-101	-	77.1	88.5	94.9	97.3	97.6
Baseline	3.9M	ResNet-50	40.12	84.38	92.79	95.90	93.85	96.55
Proposed	4.0M	ResNet-50	43.38	88.19	92.78	95.86	94.62	96.72

On the IJB-C dataset, for both verification and identification protocols, we observe a similar trend as the IJB-S dataset. In particular, larger improvement is achieved at lower FARs. This is because the verification threshold at lower FARs is affected by the low quality test data (difficult impostor pairs), which is more similar to our unlabeled data. Another interesting observation is that the improvement margin increases when there is more labeled data. Note that in general semi-supervised learning, we would expect less improvement by using unlabeled data when there is more labeled data. But it is the opposite in our case because the unlabeled data has different characteristics than the labeled data. So when the performance of supervised model saturates with sufficient labeled data, transferring the knowledge from diverse unlabeled data becomes more helpful.

For both IJB-S and IJB-C (TAR@FAR=1e-7), we observe that after a certain point, adding more labeled data does not boost performance any more and the performance starts to fluctuate. This happens because the new labeled data does not necessarily help with those hard cases. Based on these results, we conclude that *when the number of labeled training data is small, it is more important to increase the quantity of the labeled dataset. Once there is sufficient labeled training data, the generalizability of the representation tends to saturate while the diversity of the training data becomes more important.* Additional experimental results on the choice of the unlabeled dataset can be found in Section A.

4.5 Comparison with State-of-the-Art FR Methods

In Table 2 we show more complete results on IJB-C dataset and compare our method with other state-of-the-art methods. In generally, we observe that with fewer labeled training samples and number of parameters, we are able to achieve

Table 4: Performance on the IJB-S benchmark.

Method	Surveillance-to-Still					Surveillance-to-Booking				
	Rank1	Rank5	Rank10	1%	10%	Rank1	Rank5	Rank10	1%	10%
MARN [9]	58.14	64.11	-	21.47	-	59.26	65.93	-	32.07	-
PFE [35]	50.16	58.33	62.28	31.88	35.33	53.60	61.75	62.97	35.99	39.82
ArcFace[5]	50.39	60.42	64.74	32.39	42.99	52.25	61.19	65.63	34.87	43.50
Baseline	53.23	62.91	67.83	31.88	43.32	54.26	64.18	69.26	32.39	44.32
Proposed	59.29	66.91	69.63	39.92	50.49	60.58	67.70	70.63	40.80	50.31

state-of-the-art performance on most of the protocols. Particularly at low FARs, the proposed method outperforms the baseline methods with a good margin. This is because at a low FAR, the verification threshold is mainly determined by low quality impostor pairs, which are instances of the difficult face samples that we are targeting with additional unlabeled data. Similar trend is observed for IJB-B dataset (Table 3). Note that because of fewer number of face pairs, we are only able to test at higher FARs for IJB-B dataset.

In Table 4 we show the results on two different protocols of IJB-S. Both the Surveillance-to-Still (V2S) and Surveillance-to-Booking (V2B) protocols use surveillance videos as probes and mugshots as gallery. Therefore, IJB-S results represent a cross domain comparison problem. Overall, the proposed system achieve new state-of-the-art performance on both protocols.

5 Conclusions

We have proposed a semi-supervised framework of learning robust face representation that could generalize to domains beyond the labeled training dataset. Without collecting domain specific data, we utilized a relatively small unlabeled dataset containing diverse styles of face images. The images in the unlabeled dataset are assumed to come from different sources and hence span more testing domains. In order to fully utilize the unlabeled dataset, two methods are proposed. First, we showed that the domain adversarial learning, which is common in adaptation methods, can also be applied in our setting to reduce the holistic feature gap and to increase the inter-class distance between the unconstrained faces. Second, we propose to use an image translation network as a data-driven augmentation module. The augmentation network can capture different visual styles in the unlabeled dataset and apply them to the labeled images during training, making the face representation more discriminative in the unlabeled domains. Our experimental results show that as the number of labeled images increases, the performance of the supervised baseline tends to saturate on the challenging testing scenarios. Instead, introducing more diverse training data becomes more important and helpful. In a few challenging protocols, we showed that the proposed method can outperform the supervised baseline with less than half of the labeled data. By training on the labeled MS-Celeb-1M dataset and unlabeled WiderFace dataset, our final model achieves state-of-the-art performance on challenging benchmarks such as IJB-B, IJB-C and IJB-S.

References

1. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv:1607.06450 (2016)
2. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: NeurIPS (2019)
3. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: A dataset for recognising faces across pose and age. In: IEEE FG (2018)
4. Carlucci, F.M., D’Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalization by solving jigsaw puzzles. In: CVPR (2019)
5. Deng, J., Guo, J., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. CVPR (2019)
6. Deng, J., Guo, J., Zhou, Y., Yu, J., Kotsia, I., Zafeiriou, S.: Retinaface: Single-stage dense face localisation in the wild. arXiv preprint arXiv:1905.00641 (2019)
7. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation (2015)
8. Ghifary, M., Bastiaan Kleijn, W., Zhang, M., Balduzzi, D.: Domain generalization for object recognition with multi-task autoencoders. In: ICCV (2015)
9. Gong, S., Shi, Y., Jain, A.: Low quality video face recognition: Multi-mode aggregation recurrent network (marn). In: ICCV Workshops (2019)
10. Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J.: Ms-celeb-1m: A dataset and benchmark for large scale face recognition. In: ECCV (2016)
11. Hasnat, A., Bohné, J., Milgram, J., Gentric, S., Chen, L.: Deepvisage: Making face recognition simple yet with powerful generalization skills. ICCV (2017)
12. Huang, G.B., Ramesh, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Tech. Rep. 07-49, University of Massachusetts, Amherst (October 2007)
13. Huang, X., Belongie, S.J.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV (2017)
14. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: ECCV (2018)
15. Kalka, N.D., Maze, B., Duncan, J.A., O’Connor, K.J., Elliott, S., Hebert, K., Bryan, J., Jain, A.K.: IJB-S : IARPA Janus Surveillance Video Benchmark . In: BTAS (2018)
16. Kang, B.N., Kim, Y., Jun, B., Kim, D.: Attentional feature-pair relation networks for accurate face recognition. In: ICCV (2019)
17. Kang, G., Jiang, L., Yang, Y., Hauptmann, A.G.: Contrastive adaptation network for unsupervised domain adaptation. In: CVPR (2019)
18. Kemelmacher-Shlizerman, I., Seitz, S.M., Miller, D., Brossard, E.: The megaface benchmark: 1 million faces for recognition at scale. In: CVPR (2016)
19. Klare, B.F., Klein, B., Taborsky, E., Blanton, A., Cheney, J., Allen, K., Grother, P., Mah, A., Jain, A.K.: Pushing the frontiers of unconstrained face detection and recognition: IARPA Janus Benchmark A. In: CVPR (2015)
20. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning (2017)
21. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: ICML Workshop (2013)
22. Li, H., Jialin Pan, S., Wang, S., Kot, A.C.: Domain generalization with adversarial feature learning. In: CVPR (2018)
23. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Sphereface: Deep hypersphere embedding for face recognition. In: CVPR (2017)

24. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Deep transfer learning with joint adaptation networks. In: ICML (2017)
25. Masi, I., Tran, A.T., Hassner, T., Leksut, J.T., Medioni, G.: Do we really need to collect millions of faces for effective face recognition? In: ECCV (2016)
26. Maze, B., Adams, J., Duncan, J.A., Kalka, N., Miller, T., Otto, C., Jain, A.K., Niggel, W.T., Anderson, J., Cheney, J., et al.: Iarpa janus benchmark-c: Face dataset and protocol. In: ICB (2018)
27. Motiian, S., Piccirilli, M., Adjeroh, D.A., Doretto, G.: Unified deep supervised domain adaptation and generalization. In: ICCV (2017)
28. Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalization via invariant feature representation. In: ICML (2013)
29. Pan, S.J., Tsang, I.W., Kwok, J.T., Yang, Q.: Domain adaptation via transfer component analysis. *IEEE Trans. on Neural Networks* (2010)
30. Ranjan, R., Bansal, A., Zheng, J., Xu, H., Gleason, J., Lu, B., Nanduri, A., Chen, J.C., Castillo, C.D., Chellappa, R.: A fast and accurate system for face detection, identification, and verification. *IEEE Trans. on Biometrics, Behavior, and Identity Science* (2019)
31. Ranjan, R., Castillo, C.D., Chellappa, R.: L2-constrained softmax loss for discriminative face verification. arXiv:1703.09507 (2017)
32. Rasmus, A., Berglund, M., Honkala, M., Valpola, H., Raiko, T.: Semi-supervised learning with ladder networks. In: NeurIPS (2015)
33. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: CVPR (2018)
34. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: CVPR (2015)
35. Shi, Y., Jain, A.K.: Probabilistic face embeddings. In: ICCV (2019)
36. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: NIPS (2016)
37. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. arXiv:2001.07685 (2020)
38. Sohn, K., Shang, W., Yu, X., Chandraker, M.: Unsupervised domain adaptation for face recognition in unlabeled videos. In: ICLR (2019)
39. Sun, Y., Chen, Y., Wang, X., Tang, X.: Deep learning face representation by joint identification-verification. In: NIPS (2014)
40. Sun, Y., Cheng, C., Zhang, Y., Zhang, C., Zheng, L., Wang, Z., Wei, Y.: Circle loss: A unified perspective of pair similarity optimization. arXiv:2002.10857 (2020)
41. Taigman, Y., Yang, M., Ranzato, M., Wolf, L.: Deepface: Closing the gap to human-level performance in face verification. In: CVPR (2014)
42. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: NeurIPS (2017)
43. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. arXiv:1607.08022 (2016)
44. Wang, F., Liu, W., Liu, H., Cheng, J.: Additive margin softmax for face verification. arXiv:1801.05599 (2018)
45. Wang, H., Wang, Y., Zhou, Z., Ji, X., Li, Z., Gong, D., Zhou, J., Liu, W.: Cosface: Large margin cosine loss for deep face recognition. CVPR (2018)
46. Whitelam, C., Taborsky, E., Blanton, A., Maze, B., Adams, J., Miller, T., Kalka, N., Jain, A.K., Duncan, J.A., Allen, K., et al.: Iarpa janus benchmark-b face

- dataset. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 90–98 (2017)
47. Xie, Q., Dai, Z., Hovy, E., Luong, M.T., Le, Q.V.: Unsupervised data augmentation for consistency training. arXiv:1904.12848 (2019)
 48. Xie, W., Shen, L., Zisserman, A.: Comparator networks. In: ECCV (2018)
 49. Yang, S., Luo, P., Loy, C.C., Tang, X.: Wider face: A face detection benchmark. In: CVPR (2016)
 50. Yi, D., Lei, Z., Liao, S., Li, S.Z.: Learning face representation from scratch. arXiv:1411.7923 (2014)
 51. Zhai, X., Oliver, A., Kolesnikov, A., Beyer, L.: S4l: Self-supervised semi-supervised learning. In: ICCV (2019)
 52. Zhang, X., Zhao, R., Qiao, Y., Wang, X., Li, H.: Adacos: Adaptively scaling cosine logits for effectively learning deep face representations. In: CVPR (2019)
 53. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)

A Choice of the Unlabeled Dataset

In Section 4.4 of the main paper, we discussed on the impact of the quantity/diversity of training data on feature generalizability, where we conducted the experiments by adjusting the number of labeled faces. Here, we extend the discussion by showing more experiments on the choice of unlabeled dataset. In addition to the WiderFace dataset that we used in the main paper, we consider to utilize two other datasets: MegaFace [18] and CASIA-WebFace [50]. For MegaFace, we only use the distractor images in their identification protocol, which are crawled from album photos on Flickr and present a larger degree of variation compared with the faces in MS-Celeb-1M. CASIA-WebFace, similar to MS-Celeb-1M, is mainly composed of celebrity photos, and therefore it should not introduce much additional diversity. Note that CASIA-WebFace is a labeled dataset but we ignore its labels for this experiment. The diversity (facial variation) of the three datasets can be ranked as: *WiderFace* > *MegaFace* > *CASIA-WebFace*. Example images of the three datasets are shown in Fig. 9. For both MegaFace and CASIA-Webface, we choose a random subset to match the number of the WiderFace. Furthermore, to see the impact of the quantity of unlabeled dataset, we also train the models with different numbers of unlabeled data. Then, we evaluate all the models on IJB-S, IJB-C and LFW. The reason to evaluate on LFW here is to see the impact of different unlabeled datasets on the performance in the original domain. The results are shown in Fig. 10. Note that due to the large number of experiments, we do not use augmentation network here. But empirically we found the trends would be similar.

From Fig. 10, it can be seen that in general, the more diverse the unlabeled dataset is, the more performance boost it leads to. In particular, using CASIA-WebFace as the unlabeled dataset hardly improves performance on any protocol. This is expected because CASIA-WebFace is very similar to MS-Celeb-1M and hence it cannot introduce additional diversity to regularize the training of face representations. Using MegaFace distractors as the unlabeled dataset improves the performance on both IJB-C and IJB-S, both of which have more variations than the MS-Celeb-1M. Using WiderFace as the unlabeled dataset further improves the performance on the IJB-S dataset. Note that all the models in this experiment maintain the high performance on the LFW dataset. In other words, *using a more diverse unlabeled dataset would not deteriorate the performance*




Fig. 9: Examples face images of different unlabeled datasets used in this experiment.

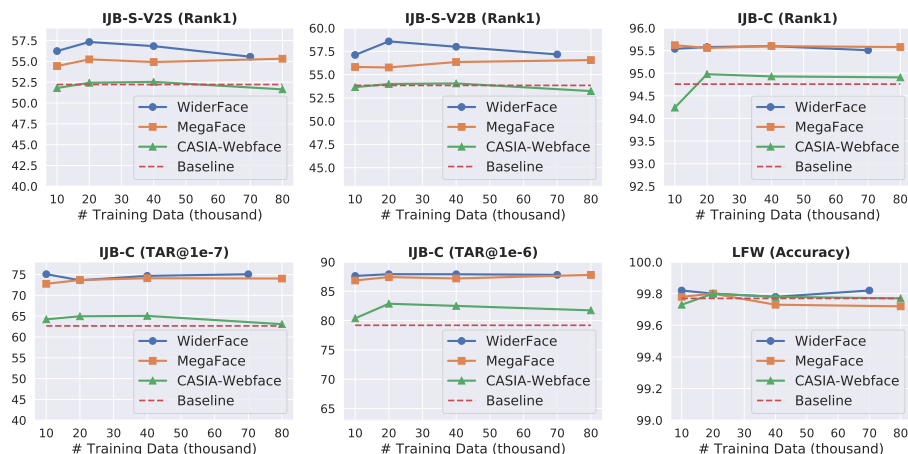


Fig. 10: Evaluation Results on IJB-S, IJB-C and LFW with different protocols and different number and choice of unlabeled training data. The red line here refers the performance of the supervised baseline which does not use any unlabeled data.

on the original domain and safely improves the performance on the challenging new domains. An additional result that we can observe is that the size of the unlabeled dataset does not have a clear effect compared to its diversity.

B Discussion on Binary Domain Adversarial Loss

In our work, the unlabeled dataset should be a diverse dataset collected from different sources, i.e. covering different sub-domains (types) of face images. Thus, if we have the access to such sub-domain labels, a natural choice of the domain adversarial loss would be aligning each of the sub-domain with feature distribution of the labeled images. However, in practice, we might not have access to such domain labels and therefore we proposed a more generic solution using a binary domain alignment loss. In this loss function, we consider all the unlabeled images as one distribution and try to minimize the holistic domain gap between the unlabeled images and labeled ones. Such a choice is made with two reasons: (1) the binary domain alignment is a prerequisite for the alignment of sub-domains and therefore, as an upper bound, this loss function should also reduce the domain gap for sub-domains; (2) Empirically we found the binary domain alignment loss is sufficient to reduce the sub-domain gap. To show the second point, we conduct a controlled experiments with a toy dataset. We split the MS-Celeb-1M into labeled images and unlabeled images (no identity overlap). The unlabeled images are then processed with one of the three degradation: random Gaussian noise, random occlusion and downsampling. Thus, we create three sub-domains in the unlabeled dataset. The corresponding domain shift can be observed in the t-SNE plot in Fig. 11 (a), where the model is trained only on the labeled

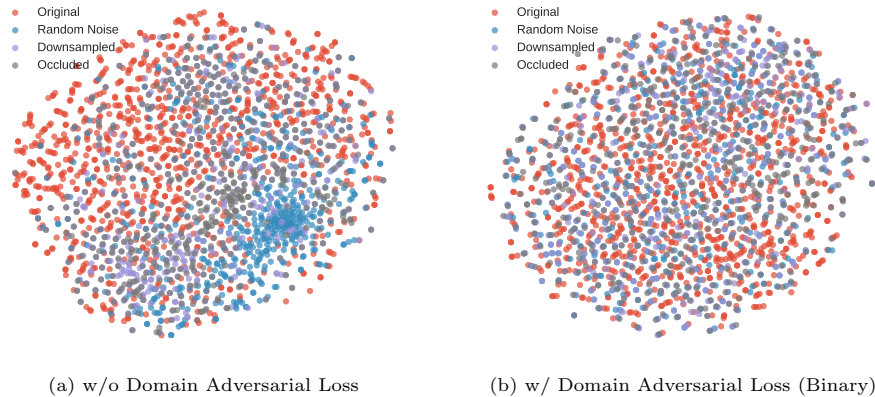


Fig. 11: t-SNE visualization of the face embeddings using synthesized unlabeled images. Using part of the MS-Celeb-1M as unlabeled dataset, we create three sub-domains by processing the images with either random Gaussian noise, random occlusion or downsampling. (a) different sub-domains show different domain shift in the embedding space of the supervised baseline. (b) with the holistic binary domain adversarial loss, each of the sub-domain is aligned with the distribution of the labeled data.

split. Then, we incorporate the augmented unlabeled images into training with the binary domain adversarial loss. In Fig. 11 (b), we observe that with such a holistic regularization, the distribution of each of sub-domain is aligned with the original domain.

C Numerical Results on IJB-B, IJB-C

In Table 5 and Table 6, we show more numerical results on the IJB-C and IJB-B dataset, respectively. Since all the baseline methods (from other papers) are trained on different number of labeled images, we report the performance of our models trained on different labeled subsets for a more fair comparison. From the tables, we could observe that our models outperform most of the baselines with equal or less than 2M labeled data.

D Architecture of Augmentation Network

The architecture of our augmentation network is based on MUNIT [14]. Let c5s1-k be a 5×5 convolutional layer with k filters and stride 1. dk-IN denotes a 3×3 convolutional layer with k filters and dilation 2, where IN means Instance Normalization [43]. Similarly, AdaIN means Adaptive Instance Normalization [13] and LN denotes Layer Normalization [1]. fc8 denotes a fully

Table 5: Performance comparison with state-of-the-art methods on the IJB-C dataset.

Method	Data	Model	Verification				Identification	
			1e-7	1e-6	1e-5	1e-4	Rank1	Rank5
Cao et al. [3]	13.3M	SE-ResNet-50	-	-	76.8	86.2	91.4	95.1
PFE [35]	4.4M	ResNet-64	-	-	89.64	93.25	95.49	97.17
ArcFace [5]	5.8M	ResNet-50	67.40	80.52	88.36	92.52	93.26	95.33
Ranjan et al. [30]	5.6M	ResNet-101	67.4	76.4	86.2	91.9	94.6	97.5
AFRN [16]	3.1M	ResNet-101	-	-	88.3	93.0	95.7	97.6
Baseline	500K	ResNet-50	51.13	66.44	77.58	87.73	90.90	94.50
Proposed	500K+70K	ResNet-50	60.33	71.24	80.31	88.18	91.81	94.96
Baseline	1.0M	ResNet-50	59.53	77.70	86.16	92.13	93.62	95.93
Proposed	1.0M+70K	ResNet-50	61.87	79.76	87.16	92.39	94.19	96.30
Baseline	2.0M	ResNet-50	67.64	78.66	88.16	93.48	94.34	96.34
Proposed	2.0M+70K	ResNet-50	78.62	84.91	90.61	93.77	95.04	96.80
Baseline	3.0M	ResNet-50	62.65	79.20	89.20	94.20	94.76	96.49
Proposed	3.0M+70K	ResNet-50	78.38	85.91	91.56	94.48	95.51	97.04
Baseline	3.9M	ResNet-50	62.90	82.94	90.73	94.57	94.90	96.77
Proposed	3.9M+70K	ResNet-50	77.39	87.92	91.86	94.66	95.61	97.13

Table 6: Performance comparison with state-of-the-art methods on the IJB-B dataset.

Method	Data	Model	Verification				Identification	
			1e-6	1e-5	1e-4	1e-3	Rank1	Rank5
Cao et al. [3]	13.3M	SE-ResNet-50	-	70.5	83.1	90.8	90.2	94.6
Comparator [48]	3.3M	ResNet-50	-	-	84.9	93.7	-	-
ArcFace [5]	5.8M	ResNet-50	40.77	84.28	91.66	94.81	92.95	95.60
Ranjan et al. [30]	5.6M	ResNet-101	48.4	80.4	89.8	94.4	93.3	96.6
AFRN [16]	3.1M	ResNet-101	-	77.1	88.5	94.9	97.3	97.6
Baseline	500K	ResNet-50	39.35	71.14	84.37	92.12	89.74	94.16
Proposed	500K+70K	ResNet-50	45.39	72.35	84.75	92.00	90.46	94.42
Baseline	1.0M	ResNet-50	45.75	80.11	90.19	94.48	92.37	95.78
Proposed	1.0M+70K	ResNet-50	41.59	82.10	90.09	94.64	92.88	95.91
Baseline	2.0M	ResNet-50	47.62	82.30	91.82	95.46	93.25	96.05
Proposed	2.0M+70K	ResNet-50	44.76	86.26	91.92	95.27	94.01	96.23
Baseline	3.0M	ResNet-50	42.77	82.86	92.48	95.78	93.80	96.23
Proposed	3.0M+70K	ResNet-50	43.09	87.31	92.80	95.70	94.35	96.53
Baseline	3.9M	ResNet-50	40.12	84.38	92.79	95.90	93.85	96.55
Proposed	3.9M+70K	ResNet-50	43.38	88.19	92.78	95.86	94.62	96.72

connected layer with 8 filters. `avgpool` denotes a global average pooling layer. No normalization is used in the style encoder. We use Leaky ReLU with slope 0.2 in the discriminator and ReLU activation everywhere else. The architectures of different modules are as follows:

- Style Encoder:
c5s1-32, c3s2-64, c3s2-128, avgpool, fc8
- Generator:
c5s1-32-IN, d32-IN, d32-AdaIN, d32-LN, d32-LN, c5s1-3
- Discriminator:
c5s1-32, c3s2-64, c3s2-128

The length of the latent style code is set to 8. A style decoder (multi-layer perceptron) has two hidden fully connected layers of 128 filters without normalization, which transforms the latent style code to the parameters of the AdaIN layer.

E Ablation over the Settings of Augmentation Network

In this section, we ablate over the training modules of the augmentation network. In particular, we consider to remove the following modules for different variants: Latent-style code for multi-mode generation (MM), Image Discriminator (D_I), Reconstruction Loss (Rec), Style Discriminator (D_z) and the architecture without downsampling (ND). The qualitative results of different models are shown in Fig. 12. Without the latent style code (Model a), the augmentation network can only output one deterministic image for each input, which mainly applies blurring to the input image. Without the image adversarial loss (Model b), the model cannot capture the realistic variations in the unlabeled dataset and the style code can only change the color channel in this case. Without the Reconstruction Loss (Model c), the model is trained only with adversarial loss but without the regularization of content preservation. And therefore, we see clear artifacts on the output images. However, adding reconstruction loss alone hardly helps, since the latent code used in the reconstruction of the unlabeled images could be very different from the prior distribution $p(z)$ that we use for generation. Therefore, similar artifacts can be observed if we do not add latent code adversarial loss (Model d). As for the architecture, if we choose to use an encoder-decoder style network as in the original MUNIT [14], with downsampling and upsampling (Model e), we observe that the output images are always blurred due to the loss of spatial information. In contrast, with our architecture (Model f), the network is capable of augmenting images with diverse color, blurring and illumination styles but without clear artifacts.

Furthermore, we incorporate these different variants of augmentation networks into training and show the results in Table 7. The baseline model here is a model that only uses domain alignment loss without augmentation network. In fact, compared with this baseline, using all different variants of the augmentation network achieves performance improvement in spite of the artifacts in the generated images. But a more stable improvement is observed for the proposed augmentation network across different evaluation protocols. We also show more examples of augmented images in Figure 13.




Fig. 12: Ablation study of the augmentation network. Input images are shown in the first column. The subsequent columns show the results of different models trained without a certain module or loss. The texture style codes are randomly sampled from the normal distribution.

Table 7: Ablation study over different training methods of the augmentation network. “MM”, “ D_I ”, “ D_Z ”, “rec”, “ND” refer to “Multi-mode”, “Image Discriminator”, “Reconstruction Loss”, “Latent Style Discriminator” and “No Downsampling”, respectively. The first row is a baseline that uses only the domain adversarial loss but no augmentation network. “Model (a)” is a single-mode translation network that does not use latent style code.

Model	Modules					IJB-C (Vrf)			IJB-C (Idt)		IJB-S (V2S)		LFW
	MM	D_I	Rec	D_Z	ND	1e-7	1e-6	1e-5	Rank1	Rank5	Rank1	Rank5	Accuracy
						72.74	85.33	90.52	94.99	96.75	56.35	66.77	99.82
(a)		✓			✓	74.80	87.58	91.94	95.51	97.09	56.98	65.66	99.80
(b)	✓		✓	✓	✓	75.32	88.00	91.71	95.42	97.04	57.54	66.72	99.75
(c)	✓	✓			✓	74.51	87.49	91.97	95.61	97.18	57.17	66.24	99.78
(d)	✓	✓	✓		✓	75.07	88.11	92.19	95.66	97.12	56.85	64.87	99.78
(e)	✓	✓	✓	✓		73.99	86.52	91.33	95.33	97.04	58.47	66.00	99.73
(f)	✓	✓	✓	✓	✓	77.39	87.92	91.86	95.61	97.13	57.33	65.37	99.75



Fig. 13: More examples of augmented images. The photos in the first column are the input images. The remaining images in each row are generated by the augmentation network with different style code.