Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/122840/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Jiang, Luo, Zhang, Juyong and Deng, Bailin ORCID: <https://orcid.org/0000-0002-0158-7670> 2020. Robust RGB-D face recognition using attribute-aware loss. IEEE Transactions on Pattern Analysis and Machine Intelligence 42 (10) , pp. 2552-2566. 10.1109/TPAMI.2019.2919284 file

Publishers page: <https://doi.org/10.1109/TPAMI.2019.2919284>
<<https://doi.org/10.1109/TPAMI.2019.2919284>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.

See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.

Robust RGB-D Face Recognition Using Attribute-Aware Loss

Luo Jiang, Juyong Zhang[†], *Member, IEEE*, and Bailin Deng, *Member, IEEE*

Abstract—Existing convolutional neural network (CNN) based face recognition algorithms typically learn a discriminative feature mapping, using a loss function that enforces separation of features from different classes and/or aggregation of features within the same class. However, they may suffer from bias in the training data such as uneven sampling density, because they optimize the adjacency relationship of the learned features without considering the proximity of the underlying faces. Moreover, since they only use facial images for training, the learned feature mapping may not correctly indicate the relationship of other attributes such as gender and ethnicity, which can be important for some face recognition applications. In this paper, we propose a new CNN-based face recognition approach that incorporates such attributes into the training process. Using an attribute-aware loss function that regularizes the feature mapping using attribute proximity, our approach learns more discriminative features that are correlated with the attributes. We train our face recognition model on a large-scale RGB-D data set with over 100K identities captured under real application conditions. By comparing our approach with other methods on a variety of experiments, we demonstrate that depth channel and attribute-aware loss greatly improve the accuracy and robustness of face recognition.

Index Terms—Face Recognition, RGB-D images, uneven sampling density, attribute-aware loss.

1 INTRODUCTION

CONVOLUTIONAL neural networks (CNNs) play a significant role in face analysis tasks, such as landmarks detection [38], [52], face recognition [12], [15], [47] and 3D face reconstruction [6], [31]. With the emergence of large public face data sets [50] and sophisticated networks [8], [39], the problem of face recognition has gained lots of attention and developed rapidly. At present, some mainstream methods already outperform humans on certain benchmark datasets such as [12]. These methods usually map faces to discriminative feature vectors in a high-dimensional Euclidean space, to determine whether a pair of faces belong to the same category. For example, deep metric learning methods (such as contrastive loss [7] or triplet loss [35]) usually train a CNN by comparing pairs or triplets of facial images to learn discriminative features. Later, different variants of the softmax loss [23], [29], [40], [43], [46] are used as supervision signals in CNNs to extract discriminative features, which achieve excellent performance under the protocol of small training set. These methods [23], [35], [46] utilize CNNs to learn strong discriminative deep features, using loss functions that enforce either intra-class compactness or inter-class dispersion.

Although the above two categories of methods have achieved remarkable performance, they still have their own limitations. First, the contrastive loss and triplet loss suffer from slow convergence due to the construction of a large number of pairs or triplets. To accelerate convergence, [37] proposed a $(N + 1)$ -tuple loss that increases the number

of negative examples. However, this loss still requires complex recombination of training samples. In comparison, the softmax loss and its variants have no such requirement on the training data and converge more quickly. The center loss [46] is the first to add soft constraints on deep features in the softmax loss to minimize the intra-class variations, significantly improving the performance of softmax loss. Afterward, the angular softmax loss [23] imposed discriminative constraints on a hypersphere manifold, which further improved the performance of softmax loss. However, by enforcing intra-class aggregation and inter-class separation among the training data, existing variants of softmax loss encourage the uniform distribution of feature vectors for the training data, even though the training data may not be sampled uniformly. As a result, the proximity between the learned feature vectors for two test data may not correctly indicate the proximity between their underlying faces, which can affect the accuracy of face recognition algorithms based on feature proximity. To address this issue, we propose an attribute-aware loss function that regularizes the learned feature mapping using other attributes such as gender, ethnicity, and age. The proposed loss function imposes a global linear relation between the feature difference and the attribute difference between nearby training data, such that feature vectors for facial data with similar attributes are driven towards each other. In addition, as these attributes are correlated with facial geometry and appearance, the attribute-aware loss also implicitly regularizes the feature proximity with respect to the facial proximity, which helps to account for potential sampling bias in the training set.

In addition, although existing RGB image-based face recognition methods have achieved great success, they rely solely on the appearance information and may suffer from poor lighting conditions such as dark environments. On the other hand, the depth image captured by RGB-D sensors

- L. Jiang and J. Zhang are with School of Mathematical Sciences, University of Science and Technology of China. The research work was done when Luo Jiang did his internship at Beijing Dilusense Technology Corporation.
- B. Deng is with School of Computer Science and Informatics, Cardiff University.

[†]Corresponding author. Email: juyong@ustc.edu.cn.

such as PrimeSense sensors provides additional geometric information that is independent of illumination, which can help to improve the robustness of recognition. To this end, we develop a CNN-based RGB-D face recognition approach, by first aligning the depth map with the RGB image grid and normalizing the depth values to the same range as the RGB values, and then feeding the resulting RGB-D values into CNNs for training and testing. Unlike existing RGB-D based deep learning approaches [20], [51] that only use small training data sets with less than 1K identities, we train our model on a large RGB-D data set with over 100K identities, where the resulting model achieves more robust performance than RGB based approaches.

Combining the RGB-D approach with the attribute-aware loss function, our new method greatly improves the robustness and accuracy of facial recognition. We tested our method on several datasets, with different identities in diverse facial expressions and lighting conditions. Our method performs consistently better than state-of-the-art approaches that only rely on RGB information and do not consider additional attributes.

To summarize, this paper makes the following major contributions:

- We propose an attribute-aware loss function for CNN-based face recognition, which regularizes the distribution of learned feature vectors with respect to additional attributes and improves the accuracy of recognition results. To the best of our knowledge, this is the first method that utilizes non-facial attributes to improve CNN-based face recognition feature training.
- For neural network training and testing, we construct a large-scale RGB-D face dataset including more than 100k identities mainly with the frontal pose, and a relatively small RGB-D dataset with 952 identities with various poses. This is the first result that verifies the effectiveness of CNN-based RGB-D face recognition with large training data sets.

2 RELATED WORK

Face recognition is a classical research topic in pattern recognition and computer vision, with applications in many areas like biometrics, surveillance system, and information security. For a comprehensive review of 2D face recognition and 3D face recognition methods, one may refer to [1], [27]. This section briefly reviews those techniques that are closely related to our work.

2.1 Deep Learning based Face Recognition

In the past few years, deep learning based face recognition is one of the most active research areas. In this part, we mainly discuss the loss functions used in these methods.

Metric Learning. Metric learning [44], [45], [49] attempts to optimize a parametric notion of distance in a fully/weakly/semi-supervised way such that the similar objects are nearby and dissimilar objects are far apart on a target space. In [49], the learning is done by finding a Mahalanobis distance with a matrix parameter when given some similar pairs of samples. In order to handle more challenging problems, kernel tricks [13], [44] had

been introduced in metric learning to extract nonlinear embeddings. In recent years, more discriminative features can be learned with advanced network architectures that minimize some loss functions based on Euclidean distance, such as contrastive loss [7] and triplet loss [35]. Moreover, these loss functions can be improved by allowing joint comparison among more than one negative example [37] or minimizing the overall classification error [18].

Classification Losses. The most commonly used classification loss is the softmax loss that maps images to deep features and then to predicted labels. Krizhevsky et al. [17] first observed that CNNs trained with softmax loss can produce discriminative feature vectors, which has also been confirmed by other works [36]. However, softmax loss mainly encourages inter-class dispersion, and thus cannot induce strong discriminative features. To enhance the discrimination power of deep features, Wen et al. [46] proposed center loss to enforce intra-class aggregation as well as inter-class dispersion. Meanwhile, Ranjan et al. [29] observed that the softmax loss is biased to the sample distribution, i.e., fitting well to high-quality faces but ignoring the low-quality faces. Adding ℓ_2 -constraints on features to the softmax loss can make the resulting features as discriminative as those trained with center loss. Afterward, Liu and colleagues [23], [24] further improved the features by incorporating an angular margin instead of the Euclidean margin into softmax loss to enhance the inter-class margin and compressing the intra-class angular distribution simultaneously.

2.2 Face Recognition with Attributes

Besides the feature vectors extracted from CNN, other attributes can also be utilized in face recognition tasks. An early study [19] trained 65 “attribute” SVM classifiers to recognize the traits of input facial images such as gender, age, race, and hair color, which are then fused with other features for face recognition. In the context of deep learning, attribute-enhanced face recognition does not gain too much attention. One related work [33] is to exploit CNN based attribute features for authentication on mobile devices, and the facial attributes are trained by a multi-task, partly based Deep Convolutional Neural Network architecture. Hu et al. [11] systematically study the problem of how to fuse face recognition features and facial attribute features to enhance face recognition performance. They reformulate feature fusion as a gated two-stream neural network, which can be efficiently optimized by neural network learning.

Based on the assumption that attributes like gender, age and pose could share low-level features from the representation learning perspective, some studies investigate multi-task learning [30], [32] and show that such attributes could help the face recognition task. In our method, different from the above attribute fusion and multi-task learning methods, the attributes are directly used to guide the face recognition feature learning in the training stage, and they are not needed during the testing stage.

2.3 RGB-D Face Recognition

In recent years, RGB-D based face recognition has attracted increasing attention because of its robustness in an unconstrained environment. Hsu et al. [10] considered a scenario

in which the gallery is a pair of RGB-D images while the probe is a single RGB image captured by a regular camera without the depth channel. They proposed an approach that reconstructs a 3D face from an RGB-D image for each subject in the gallery, aligns the reconstructed 3D model to a probe using facial landmarks, and recognizes the probe using sparse representation based classification. Zhang et al. [51] further considered the problem of multi-modality matching (e.g., RGB-D probe vs. RGB-D gallery) and cross-modality matching (e.g., RGB probe vs. RGB-D) in the same framework. They proposed an approach for RGB-D face recognition that is able to learn complementary features from multiple modalities and common features between different modalities. For the RGB-D vs. RGB-D problem, Goswami et al. [5] proposed to compute an RGB-D image descriptor based on entropy and based on the entropy and saliency, as well as geometric facial attributes from the depth map; then the descriptor and the attributes are fused to perform recognition. Li et al. [21] proposed a multi-channel weighted sparse coding method on the hand-crafted features for RGB-D face recognition.

Although it is straightforward to extend deep learning based face recognition methods from RGB images to RGB-D images, currently there are no large-scale public RGB-D data sets that can be used for training, which limits the practical applications of these approaches. For example, the model proposed in [51] is trained on a dataset with less than 1K identities. To handle this problem, Lee et al. [20] proposed to first train the deep network with a color face dataset, and then fine-tune it on depth face images for transfer learning.

3 METHOD

3.1 Revisiting the Variants of Softmax Loss

Given a training data set $\{\mathbf{x}_i\}_{i=1}^N$ with $\mathbf{x}_i \in \mathbb{R}^{m \times n}$, and their corresponding labels $\{y_i\}_{i=1}^N$ with $y_i \in \mathcal{I} \triangleq \{1, \dots, C\}$, the following classical softmax loss function is widely used in face recognition tasks

$$L_s = - \sum_{i=1}^N \log \left(\frac{\exp(\mathbf{w}_{y_i}^T f(\mathbf{x}_i) + b_{y_i})}{\sum_{j=1}^C \exp(\mathbf{w}_j^T f(\mathbf{x}_i) + b_j)} \right), \quad (1)$$

where $f(\cdot) : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{K \times 1}$ is the learned feature mapping by training CNNs, K is the dimension of deep feature $f(\mathbf{x}_i)$. $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_C] \in \mathbb{R}^{K \times C}$ and $\mathbf{b} = [b_1, \dots, b_C] \in \mathbb{R}^{1 \times C}$ are the weights and biases in the last fully connected layer, and $\{\mathbf{x}_i\}$ can be color or depth images of faces. We denote $\mathbf{f}_i = f(\mathbf{x}_i)$ for simplicity. Typically, during the test phase the mapping $f(\cdot)$ is applied on an image pair $(\mathbf{x}_i, \mathbf{x}_j)$ to extract two deep features $(\mathbf{f}_i, \mathbf{f}_j)$, and the Euclidean distance or cosine distance between the features are computed to determine the similarity between the image pair. Separable features can be learned using softmax loss, but they are not discriminative enough for face recognition.

To learn more discriminative features, several variants of softmax loss have been developed by enlarging the inter-class margin and reducing the intra-class variation. Among them, the center loss [46] requires the deep features of each class to gather towards their respective centers $\{\mathbf{c}_i\}_{i=1}^N$:

$$L_c = \frac{1}{2} \sum_i \|\mathbf{f}_i - \mathbf{c}_{y_i}\|_2^2. \quad (2)$$

With the angular softmax loss [23], deep features of each class are compressed using the angular margin τ instead of the Euclidean margin:

$$L_\tau = - \sum_{i=1}^N \log \frac{\exp(\|\mathbf{f}_i\| \cos(\tau \theta_{y_i, i}))}{\exp(\|\mathbf{f}_i\| \cos(\tau \theta_{y_i, i})) + \sum_{j \neq y_i} \exp(\|\mathbf{f}_i\| \cos(\theta_{j, i}))}, \quad (3)$$

where $\theta_{j, i}$ is the angle between vectors $\mathbf{w}_j$ and $\mathbf{f}_i$. There are other variants of softmax loss [3], [42] with a similar form as (3), where the margin and the angle are added instead of being multiplied.

3.2 The Attribute-Aware Loss

To achieve high accuracy for face recognition, it is desirable that the proximity between feature clusters of different classes is consistent with the proximity between the classes (i.e., the underlying faces). Ideally, the more dissimilar two faces are, the further apart their corresponding feature clusters should be. However, this is not guaranteed by the above variants of softmax loss according to our experimental observations (see Fig. 2 and Fig. 7). Since they minimize the intra-class variations and maximize the inter-class margins on the training data, the learned feature mappings tend to produce evenly distributed feature vectors for the training faces. On the other hand, there is no guarantee that the facial images in the training set are evenly distributed in the full face space. As a result, when there exist large variations of sampling density in the training data set, the learned feature mapping may not correctly indicate the proximity of the underlying faces. To address this issue, we can try to introduce a loss function term that regularizes feature proximity with respect to face proximity. However, this is a challenging task as a facial image only reveals the underlying face shape from a certain view direction and can be affected by various factors such as lighting condition and sensor noises. As a result, it is difficult to reliably compute the proximity between two faces by only comparing their scanned images.

Besides the proximity of face shapes, it is also desirable that the learned feature mappings are related to the proximity between other attributes such as gender, ethnicity, and age. For example, if we compare a probe image against a database of facial images to identify Q most likely matches via feature proximity, then it is preferable that all returned images are from persons with the same or similar attributes. The above variants of softmax loss cannot ensure this property either, as they only consider the facial images during the training process.

Motivated by these observations, we propose an attribute-aware loss term that regularizes feature proximity with respect to attribute proximity. Besides the label information, other attributes of the facial images like gender, ethnicity, and age are also given in the training data set. These attributes can be collected during training data construction, and they are independent of the imaging process. We represent the augmented attributes for a facial image $\mathbf{x}_i$ using a vector $\mathbf{p}_i \in \mathbb{R}^{H \times 1}$. Then our attribute-aware loss is formulated as

$$L_a = \frac{1}{2} \sum_{\substack{y_i < y_j \\ d(\mathbf{p}_i, \mathbf{p}_j) < \tau}} \|(\mathbf{f}_i - \mathbf{f}_j) - \mathbf{G}(\mathbf{p}_i - \mathbf{p}_j)\|_2^2, \quad (4)$$

Fig. 1. The face recognition features trained with only the softmax loss tend to be evenly distributed for the training data, which can become less effective when the training data are not evenly sampled from the face space. By augmenting the loss function with the attribute-aware loss term, the learned features have similar distribution as the attributes, which are better suited for recognition.

where $\mathbf{G} \in \mathbb{R}^{K \times H}$ is a parameter matrix to be trained, $d(\mathbf{p}_i, \mathbf{p}_j) = \|\mathbf{p}_i - \mathbf{p}_j\|_2$ is the Euclidean distance between the two attribute vectors, and τ is a user-specified threshold.

Intuitively, this loss term can drive feature clusters with similar attributes towards each other, via a global linear mapping $\mathbf{G}$ that relates the feature difference to the attribute difference. To the best of our knowledge, this is the first work in face recognition that optimizes adjacency of learned features using attribute proximity. As shown in Fig. 1, the learned feature clusters with similar attributes become closer after our adjacency optimization. From another perspective, regularization using additional attributes can help the network pick up other useful cues for face recognition, because attributes such as gender, ethnicity, and age are highly correlated with facial shape and appearance. For example, there can be a notable difference between the facial appearance of two persons with different genders. Therefore, the attribute-aware loss can improve the learned feature mapping by implicitly utilizing the appearance variation related to these attributes.

To better understand the attribute-aware loss, the following proposition clarifies how the difference between $\mathbf{p}_i$ and $\mathbf{p}_j$ influences the difference between $\mathbf{f}_i$ and $\mathbf{f}_j$ in the feature space.

Proposition 1. Let $\{\mathbf{f}_i\}_{i=0}^{N-1}$ be facial features extracted from training samples $\{\mathbf{x}_i\}_{i=0}^{N-1}$ with corresponding labels $\{y_i\}_{i=0}^{N-1}$ and attribute vectors $\{\mathbf{p}_i\}_{i=0}^{N-1}$. Assume that:

- (i) The number of training epochs is large enough to allow all attribute pairs $(\mathbf{p}_i, \mathbf{p}_j)$ that satisfy constraints $y_i \neq y_j$ and $d(\mathbf{p}_i, \mathbf{p}_j) = \|\mathbf{p}_i - \mathbf{p}_j\|_2 < \tau$ to appear sufficiently many times in training phase;
- (ii) For every identity, there exists at least another identity so that

they can meet the above constraints;

- (iii) The training is convergent, the parameter matrix $\mathbf{G}$ is nonsingular, and $\|\mathbf{G}\|_2 \leq M$ where M is a constant.
- (iv) The value

$$\epsilon \triangleq \max_{\substack{y_i < y_j \\ d(\mathbf{p}_i, \mathbf{p}_j) < \tau, \\ i, j \in \{0, \dots, N-1\}}} \|(\mathbf{f}_i - \mathbf{f}_j) - \mathbf{G}(\mathbf{p}_i - \mathbf{p}_j)\|_2 \quad (5)$$

is bounded.

Let $\mathcal{S}_k$ denote the set of all facial features of the k -th identity. Then

- (i) For each k , the Euclidean distance between any pair in $\mathcal{S}_k$ has an upper bound which is linear with threshold τ ;
- (ii) If the Euclidean distance between the attribute vectors of the k_1 -th identity and the k_2 -th identity is smaller than τ , then the Euclidean distance between their average features $\bar{\mathbf{f}}^{k_1} = (\sum_{\mathbf{f}_i \in \mathcal{S}_{k_1}} \mathbf{f}_i) / |\mathcal{S}_{k_1}|$ and $\bar{\mathbf{f}}^{k_2} = (\sum_{\mathbf{f}_j \in \mathcal{S}_{k_2}} \mathbf{f}_j) / |\mathcal{S}_{k_2}|$ also has an upper bound which is linear with threshold τ .

Proof. From Eq. (5), we have

$$\|(\mathbf{f}_i - \mathbf{f}_j) - \mathbf{G}(\mathbf{p}_i - \mathbf{p}_j)\|_2 \leq \epsilon, \quad \forall y_i < y_j, d(\mathbf{p}_i, \mathbf{p}_j) < \tau,$$

Thus

$$d(\mathbf{f}_i, \mathbf{f}_j) \leq \|\mathbf{G}(\mathbf{p}_i - \mathbf{p}_j)\|_2 + \epsilon \leq \|\mathbf{G}\|_2 \|\mathbf{p}_i - \mathbf{p}_j\|_2 + \epsilon < M \cdot \tau + \epsilon. \quad (6)$$

Let $\mathbf{f}_\alpha$ and $\mathbf{f}_\beta$ be an arbitrary pair of features in $\mathcal{S}_k$. Then from assumption (ii), we can find another feature $\mathbf{f}_\gamma \in \mathcal{S}_{k'}$ where $k' \neq k$, such that their corresponding attributes $\mathbf{p}_\alpha, \mathbf{p}_\beta, \mathbf{p}_\gamma$ satisfy $d(\mathbf{p}_\alpha, \mathbf{p}_\gamma) < \tau$ and $d(\mathbf{p}_\beta, \mathbf{p}_\gamma) < \tau$. Applying Eq. 6 to the pairs $(\mathbf{p}_\alpha, \mathbf{p}_\gamma)$ and $(\mathbf{p}_\beta, \mathbf{p}_\gamma)$, we obtain

$$d(\mathbf{f}_\alpha, \mathbf{f}_\beta) \leq d(\mathbf{f}_\alpha, \mathbf{f}_\gamma) + d(\mathbf{f}_\beta, \mathbf{f}_\gamma) < 2(M \cdot \tau + \epsilon). \quad (7)$$

Fig. 2. The distribution of deeply learned features under (a) softmax loss and (b) the joint supervision of softmax loss and attribute-aware loss. There are nine identities with three different ages. The points with different colors denote features with different ages.

According to assumptions (iii) and (iv), M is a constant and ϵ is bounded. Therefore, Eq. (7) provides an upper bound for the Euclidean distance between any pair in $\mathcal{S}_k$ which is linear with τ .

For the average features $\bar{\mathbf{f}}^{k_1}$ and $\bar{\mathbf{f}}^{k_2}$ of $\mathcal{S}_{k_1}$ and $\mathcal{S}_{k_2}$, their Euclidean distance satisfies

$$d(\bar{\mathbf{f}}^{k_1}, \bar{\mathbf{f}}^{k_2}) \leq \max_{\substack{\mathbf{f}_i \in \mathcal{S}_{k_1} \\ \mathbf{f}_j \in \mathcal{S}_{k_2}}} d(\mathbf{f}_i, \mathbf{f}_j). \quad (8)$$

If the Euclidean distance between the attribute vectors of the k_1 -th identity and the k_2 -th identity is smaller than τ , then by definition $d(\mathbf{p}_i, \mathbf{p}_j) < \tau$ for any $\mathbf{f}_i \in \mathcal{S}_{k_1}$ and $\mathbf{f}_j \in \mathcal{S}_{k_2}$. This implies $d(\mathbf{f}_i, \mathbf{f}_j) < M \cdot \tau + \epsilon$ according to Eq. (6). Applying this relation to Eq. (8), we obtain

$$d(\bar{\mathbf{f}}^{k_1}, \bar{\mathbf{f}}^{k_2}) < M \cdot \tau + \epsilon,$$

which provides an upper bound that is linear with τ . $\square$

Proposition. 1 implies two properties. First, the attribute-aware loss layer can make intra-class features more compact than using the softmax loss only, similar to the center loss [46]; the smaller the threshold τ is, the more compact the intra-class features may become. Sec. 4.2 evaluates the effects from different values of τ . Second, for two identities with similar attributes, their corresponding feature clusters will not be far away from each other. This can be demonstrated by our experiment in Sec. 4.3.

To showcase the effectiveness of the attribute-aware loss, we present a toy example on a very small RGB face dataset. This dataset is selected from our large-scale RGB-D face dataset presented in Section 4.1, and contains only nine identities with the same gender and ethnicity but different ages. Three of the identities are aged 28, another three aged 50, and the remaining aged 70. We use ResNet-10 [8] to train two models, one with softmax loss only, the other with both the softmax loss and the attribute-aware loss. Details of training with the combined softmax and attribute-aware loss are presented in Sec. 3.3. We reduce the output dimension of the penultimate fully connected layer to two, allowing us to directly plot the learned features in Fig. 2(a) and Fig. 2(b). We can see that the coordinates of features in Fig. 2(a) span

a much larger range than those in Fig. 2(b). It indicates that the two-dimensional features of each identity become more compact through the regularization using attribute-aware loss. We can also observe that features for identities of the same age in Fig. 2(b) are closer to each other than in Fig. 2(a), verifying the second property of the attribute-aware loss.

3.3 Training with the Attribute-Aware Loss

Similar to [46], our attribute-aware loss in Eq. (4) is an auxiliary supervision signal, which can be combined with any variant of softmax loss. For example, it can be combined with the classical softmax to derive a loss function

$$L = L_s + \lambda L_a, \quad (9)$$

where λ is a user-specified weight to balance the two loss terms. In the following, we provide the details of mini-batch training with the loss function L . Each input mini-batch $\mathcal{B}$ consists of M facial data $\{\mathbf{x}_{b_k}\}_{k=1}^M$ where $1 \leq b_k \leq N$, as well as their identity labels $\{y_{b_k}\}_{k=1}^M$ and attributes $\{\mathbf{p}_{b_k}\}_{k=1}^M$. These data are fed to the CNN, the softmax loss layer, and the attribute-aware loss layer respectively, as illustrated in Fig. 3. The addition of an attribute-aware loss layer only introduces a slight overhead for the model size during the training phase, as it contains only one parameter matrix with $K \times H$ parameters, where K and H are the dimensions of the deep facial feature and the attribute vector, respectively. In our implementation, $K = 512$ and $H = 3$, meaning that we only need 1536 additional parameters. In comparison, the backbone network that extracts deep features, which is a 28-layer ResNet [8], requires about 0.3M parameters. Thus, the overhead from the attribute-aware loss layer is almost negligible. Moreover, during the testing phase, the attribute-aware loss layer is not needed and induces no overhead.

All parameters in the CNN and two loss layers can be learned using standard stochastic gradient descent. The

Fig. 3. The framework of our approach. Facial images stacked with normalized point clouds and fed into a deep convolutional neural network. The output features of the CNN are directly fed into a fully connected layer and afterwards the softmax loss layer. They are also fed into the attribute-aware loss layer together with their corresponding attribute vectors.

gradients of L_a with respect to $\mathbf{G}$ and $\mathbf{f}_k$ are computed as:

$$\begin{aligned} \frac{\partial L_a}{\partial \mathbf{G}} &= \sum_{\substack{b_i < b_j \\ d(\mathbf{p}_{b_i}, \mathbf{p}_{b_j}) < \tau}} [\mathbf{G}(\mathbf{p}_{b_i} - \mathbf{p}_{b_j}) - (\mathbf{f}_{b_i} - \mathbf{f}_{b_j})] (\mathbf{p}_{b_i} - \mathbf{p}_{b_j})^T, \\ \frac{\partial L_a}{\partial \mathbf{f}_{b_k}} &= \sum_{\substack{j \neq k \\ d(\mathbf{p}_{b_k}, \mathbf{p}_{b_j}) < \tau}} [(\mathbf{f}_{b_k} - \mathbf{f}_{b_j}) - \mathbf{G}(\mathbf{p}_{b_k} - \mathbf{p}_{b_j})]. \end{aligned} \quad (10)$$

Algorithm 1 summarizes the learning details in the CNNs with joint supervision.

Algorithm 1 The jointly supervised learning algorithm.

Input: Training facial data $\{\mathbf{x}_i\}$ with identity labels $\{y_i\}$ and attributes $\{\mathbf{p}_i\}$; initial parameters Θ for the CNN; parameters $\mathbf{W}$, $\mathbf{b}$ for the softmax loss layer and $\mathbf{G}$ for the attribute-aware loss layer; balancing weight λ ; learning rate α ; iteration count $t = 0$.

Output: Trained parameters Θ for the CNN.

- 1: **while** not converge **do**
- 2: Compute forward by $L^t = L_s^t + L_a^t$.
- 3: Randomly choose a mini-batch $\mathcal{B}_t$.
- 4: **for** sample $i \in \mathcal{B}_t$ **do**
- 5: Compute backward by $\frac{\partial L^t}{\partial \mathbf{f}_i^t} = \frac{\partial L_s^t}{\partial \mathbf{f}_i^t} + \lambda \frac{\partial L_a^t}{\partial \mathbf{f}_i^t}$
- 6: **end for**
- 7: Update $\mathbf{W}$: $\mathbf{W}^{t+1} = \mathbf{W}^t - \frac{\alpha^t}{M} \frac{\partial L_s^t}{\partial \mathbf{W}^t}$.
- 8: Update $\mathbf{b}$: $\mathbf{b}^{t+1} = \mathbf{b}^t - \frac{\alpha^t}{M} \frac{\partial L_s^t}{\partial \mathbf{b}^t}$.
- 9: Update $\mathbf{G}$: $\mathbf{G}^{t+1} = \mathbf{G}^t - \frac{\alpha^t}{M} \frac{\partial L_a^t}{\partial \mathbf{G}^t}$.
- 10: Update Θ : $\Theta^{t+1} = \Theta^t - \frac{\alpha^t}{M} \sum_{i=1}^M \frac{\partial L^t}{\partial \mathbf{f}_i^t} \cdot \frac{\partial \mathbf{f}_i^t}{\partial \Theta^t}$.
- 11: Increment iteration count t .
- 12: **end while**

3.4 Training with RGB-D facial data

In this paper, we use RGB-D facial images as the training data, to improve robustness to illumination conditions compared with RGB facial images. The RGB-D data are collected using low-cost sensors such as PrimeSense. Using the RGB part of a facial image, we first detect the face region and five landmarks (the eyes, the nose, and two corners of the mouth) using MTCNN [52]. The face is then cropped to 112×96 by similarity transformation, and each RGB color component is normalized from the range $[0, 255]$ into $[-1, 1]$. Afterward, we extract a face region from the corresponding depth image by transferring the RGB face region. Similar to [16], [53], we find the nose tip and crop the point cloud in the face region within an empirically set radius of 90mm. Then we move the center of the cropped facial scan to $(0, 0, z_{opt})$ and reproject it onto a 2D image plane to generate a new depth map of size 112×96 . The value z_{opt} is chosen to enlarge the projection of facial scans onto the image plane as much as possible. Following [9], we compute the depth of each pixel with bilinear interpolation. Using this depth map, we generate a new point cloud under the camera coordinate system. Each point (v_x, v_y, v_z) is further normalized as:

$$\begin{aligned} \bar{v}_x &= (2v_x - x_{\max} - x_{\min}) / (x_{\max} - x_{\min}), \\ \bar{v}_y &= (2v_y - y_{\max} - y_{\min}) / (y_{\max} - y_{\min}), \\ \bar{v}_z &= (2v_z - z_{\max} - z_{\min}) / (z_{\max} - z_{\min}), \end{aligned} \quad (11)$$

where $(x_{\min}, y_{\min}, z_{\min})$ and $(x_{\max}, y_{\max}, z_{\max})$ are the minimum and maximum x -, y - and z -coordinate values among all points, respectively. Augmenting the RGB face region with its normalized point cloud, we obtain a six-channel image with values in $[-1, 1]^6$, which is fed into the deep neural network. Some RGB facial images and their normalized point clouds are shown in Fig. 3.

4 EXPERIMENTAL RESULTS

We conduct extensive experiments to evaluate the effectiveness of our approach. We first test our RGB-D face recognition

Fig. 4. Samples from our large-scale RGB-D dataset. In each row, we show RGB-D data of one individual captured by PrimeSense camera at different locations and times and under different lighting conditions.

TABLE 1
The age and gender distribution of the two training data sets we construct, each including 60K individuals.

Age Group	TRAINING-SET-I		TRAINING-SET-II	
	#male	#female	#male	#female
(0,25]	6613	6613	4000	2231
[26,35]	6613	6613	16504	6245
[36,45]	6613	6613	12558	4994
[46,55]	6612	6612	6612	3688
[56,65]	3530	1889	1648	108
[66,100]	1420	259	1325	87

approach on a large-scale private dataset (Secs. 4.2, 4.3 and 4.6) as well as some public datasets (Sec 4.8). Then we compare our attribute-aware loss with other methods that utilize attributes for face recognition, using some public RGB datasets (Sec 4.9).

4.1 Implementation

Our RGB-D dataset. We construct an RGB-D facial dataset that is captured by PrimeSense camera and contains more than 1.3M RGB-D images of 110K identities, where each identity has at least seven RGB images and their corresponding depth images. Most subjects are captured in the front of the camera with a neutral expression, and the multiple images of each subject are captured at different times and under different lighting conditions. Some samples from this RGB-D facial dataset are shown in Fig. 4. We also record their attributes including age, gender, and ethnicity. Compared with the datasets used for RGB-D face recognition in previous work [20], [51], our dataset contains a much larger number of identities, enabling us to evaluate the effectiveness of our approach in a real-world setting.

Implementation details. All CNN models are implemented using the Caffe library [14] with our modifications. Our CNN models are based on the same architecture as [46], using a 28-layer ResNet [8]. We train the models using

stochastic gradient descent with different loss functions on RGB data, depth data, and their combination, respectively. All CNN models are trained with a batch size of 200 on two GPUs (TITAN XP). The learning rate begins at 0.1, and is divided by ten after 40K and 60K iterations, respectively. The training ends at 70K iterations. The facial data are horizontally flipped for data augmentation. During testing, we extract 512-dimensional deep features from the output of the first fully connected layer. For each test data, we concatenate its 512-dimensional features and its horizontally flipped 512-dimensional features as the final 1024-dimensional representation. In face verification and identification, the similarity between two features is computed using their cosine distance.

4.2 Experiments on the Parameters λ and τ

The parameter λ controls the importance of attribute-aware loss L_a , while the parameter τ decides whether a pair of attribute vectors are close enough to be considered in the attribute-aware loss. Since both of them are important for our loss function, we conduct two experiments to illustrate how λ and τ influence the face recognition performance. We first construct a training set and a test set by sampling the whole dataset. This training set (TRAINING-SET-I) includes about 0.88M RGB-D images of 60K identities, with 91% Caucasians and 9% Asians. Within the training set there are balanced distributions of age and gender, as shown in Tab. 1. The test set includes about 0.22M RGB-D images of 20K identities. The first available neutral image of each identity in the test set is placed in the gallery, and the remaining images are used as probes. We select gender, ethnicity, and age as the attributes for training the model. For the gender attribute, we use 1 to indicate male and -1 for females. For ethnicity, since our dataset only contains Asians and Caucasians, we use 1 for Asians and -1 represents Caucasians. For age, we first truncate the age value at 100, and then linearly map it from the range [0, 100] into [-1, 1]. In this way, we represent the attributes as 3-dimensional vector $\mathbf{p}_i = (p_i^g, p_i^e, p_i^a)$ in $[-1, 1]^3$, where the superscripts g , e and a indicate gender, ethnicity and age, respectively.

Fig. 5. Rank-1 face identification rates of our approach on the RGB part of our constructed test data. Left: results achieved by fixing $\tau = 0.02$ in Eq. (4) and varying λ in Eq. (9). Right: results by fixing $\lambda = 0.001$ and varying τ .

TABLE 2
Comparison of face identification rates (%) using two training datasets with four different loss functions.

Loss functions	TRAINING-SET-I			TRAINING-SET-II		
	RGB	Depth	RGB + Depth	RGB	Depth	RGB + Depth
(a) softmax	83.12	83.83	88.93	82.46	82.67	86.59
(b) softmax + attribute-aware	92.24	90.70	96.26	92.00	90.32	96.16
(c) softmax + center	96.56	88.49	98.69	95.89	87.85	98.47
(d) softmax + center + attribute-aware	96.78	93.29	98.99	96.28	92.86	98.75

To demonstrate the effectiveness and sensitivity of the two parameters, we train our models jointly supervised with the softmax loss and the attribute-aware loss on the RGB part of the constructed dataset. In the first experiment, we fix $\tau = 0.02$ and vary λ from 0 to 0.003 to learn different models. Performance on the closed-set identification task is the classical evaluation criteria for face recognition. We show the rank-1 identification rates of these models on our test set in Fig. 5(left). We can see that our attribute-aware loss can greatly improve the face recognition performance, especially when λ is in the range $[10^{-5}, 10^{-3}]$. In the second experiment, we fix $\lambda = 0.001$ and vary τ from 0 to 0.04. The corresponding rank-1 identification rates on our test set are shown in Fig. 5(right). It can be observed that the identification rates remain stable for $\tau \in [0.01, 0.04]$. Within this range, there are between 150 and 630 pairs of similar attribute vectors with different identities in one batch. In practice, we prefer to select a small value for τ due to its lower computational cost.

4.3 Experiments on Our RGB-D Dataset

Training & test sets. To better verify whether the attributes and the depth data can improve the face recognition performance, we construct another training set (TRAINING-SET-II) from the whole dataset. This training set also includes about 0.88M RGB-D images of 60K identities, with the same distribution of ethnicity but less balanced distributions of age and gender compared with TRAINING-SET-I (see Tab. 1). We use the probe set and gallery set in Sec 4.2 as our test set. In this experiment, we train and test our models on RGB, depth and their combination, respectively.

Training CNNs. For a fair comparison, we train the CNN model with four loss functions respectively: (a) softmax loss; (b) softmax loss combined with attribute-aware loss; (c) softmax loss combined with center loss [46]; (d) softmax loss combined with center loss and attribute-aware loss. We use models (a) and (c) as baselines, to demonstrate the effect of adding attribute-aware loss into the overall loss function. The weights of center loss and attribute-aware loss are set to 0.001 and 0.0001, respectively. The margin τ used in Eq. (4) is set to 0.01 based on the analysis in Sec. 4.2. For model (d), we use the center loss to fine-tune the CNN model trained by the attribute-aware loss. For fine-tuning, the learning rate begins at 0.01, and is divided by 10 after 20K and 30K iterations, respectively. The fine-tuning ends at 35K iterations.

Results & Analysis. The rank-1 identification rates are shown in Tab. 2. It shows that training with both RGB and depth can achieve better face recognition performance than with RGB or depth alone. As the depth image captured by the PrimeSense camera is of lower quality than the corresponding RGB image, recognition features trained with depth images are less discriminative than those with RGB images. On the other hand, depth information can be a helpful complement to RGB for improving face recognition accuracy.

The benefit of incorporating depth information into face recognition is further shown in Fig. 6, where we show some probes that cannot be identified correctly at rank one using the model trained with softmax loss on RGB data, but can be identified correctly when the model is trained with the same loss on RGB-D data. This is due to the very different lighting conditions between the probes and their corresponding sample in the gallery, which makes it difficult

Fig. 6. RGB-D data of two individuals in our gallery and their corresponding three RGB-D images in our probe set. These probes cannot be identified correctly at rank one without depth data. On top of each RGB probe, we show its ranking (r) in RGB-only probing and the cosine distance (d) between the features extracted from the RGB probe and its gallery sample. Below each depth probe, we show the ranking in RGB-D probing and the cosine distance between features extracted from the RGB-D probe and its gallery sample.

to perform face recognition using only RGB data. On the other hand, depth data is more robust to lighting than RGB data, which helps to refine the learned recognition features. Indeed, the use of depth data improves the ranking of the probes and reduces the cosine distance between the features of each probe and its gallery sample.

In Fig. 7, we show ten clusters of features trained on RGB data with softmax loss (see Fig. 7(a)) and with the combined softmax and attribute-aware loss (see Fig. 7(b)), respectively. The results are visualized in 2D using dimensionality reduction. Without the attribute-aware loss, the five male feature clusters are interlaced with the female feature clusters. In comparison, with an attribute-aware loss there is a more clear separation between male and female feature clusters, and the features within each cluster are closer to each other. In particular, there is one photo of a female subject with large head pose, which looks significantly different from other photos of the same subject. Without attribute-aware loss, this photo is mapped to a feature far away from the feature cluster of the subject (shown as red circles) and becomes an outlier. Using attribute-aware loss, all photo of this female subject is mapped to a tight cluster.

TABLE 3

Comparison of face identification rates (%) using TRAINING-SET-I and TRAINING-SET-II with detected attribute features.

Training set	RGB	Depth	RGB + Depth
TRAINING-SET-I	87.40	88.14	92.84
TRAINING-SET-II	87.04	89.97	92.42

TABLE 4

Comparison of face identification rates (%) by training on the RGB-D data of TRAINING-SET-I with three controlled attributes. Here “a”, “e”, and “g” indicate age, ethnicity, and gender, respectively.

Loss function	e + a	g + a	g + e
softmax + attribute-aware	95.52	96.05	94.16

Tab. 2 also shows that the addition of attribute-aware loss into the overall loss function results in more accurate face identification results compared with the two baseline models without attribute-aware loss. Especially, the model with combined softmax and attribute-aware loss outperforms the model with only softmax loss by a significant margin. Moreover, by comparing the results on TRAINING-SET-I and TRAINING-SET-II, we can observe that the attribute-aware loss improves the robustness of the results when the training datasets are not evenly sampled, because it uses the difference of attributes as a proxy to take the sampling bias into account.

4.4 Experiments on Our RGB-D Dataset with Detected Attribute Features

Since our RGB-D facial dataset records three non-facial attributes for each subject, we can use this dataset to train a network that detects attribute features from facial images. We also evaluate the effectiveness of our approach with such detected attributes, by replacing the ground-truth attributes with detected attribute features in the attribute-aware loss layer. Following [11], we adopt the architecture of Lighten CNN [48] as the attribute detection network. We train this CNN model with a multi-task loss function MOON [32] on the RGB data of TRAINING-SET-I. The learning rate begins at 0.01, and is divided by ten after 40K and 60K iterations, respectively. The training ends at 70K iterations. We then extract 256-dimensional deep attribute features from the penultimate fully connected layer of the attribute detection network, to feed into the attribute-aware loss layer. Using these detected attribute features, we retrain our face recognition network jointly supervised with softmax loss and attribute-aware loss on both TRAINING-SET-I and TRAINING-SET-II. The training details are the same as in Sec. 4.3. The rank-1 identification rates are shown in Tab 3. We can see that using the attribute-aware loss layer with detected attribute features also improve the face recognition accuracy, although not as much as the ground-truth attributes. It shows that our approach can also be applied in scenarios where ground-truth attributes are not available in the training data.

4.5 Experiments on Our RGB-D Dataset with Controlled Attributes

In Sec. 4.3, we use three attributes in the attribute-aware loss layer to achieve good performance in the test phase. To

Fig. 7. The distribution of learned features supervised by softmax loss (a) and supervised jointly by softmax loss and attribute-aware loss (b) in two-dimensional space after dimension reduction using t-SNE [26]. There are ten feature clusters with different color, five males (pentacle) and five females (circle). All of color images of the first female (red circle) are shown under the learned features.

further understand the importance of each individual attribute, three new kinds of controlled attributes are designed respectively: (a) ethnicity and age (p_i^e, p_i^a); (b) gender and age (p_i^g, p_i^a); (c) gender and ethnicity (p_i^g, p_i^e). We ignore one of the ground-truth attributes, in turn, to observe its corresponding performance change. The dimension of attribute vectors becomes two in the attribute-aware loss layer now. We train the CNN model with the three kinds of controlled attributes on the RGB-D data of TRAINING-SET-I respectively. The training details are the same as in Sec. 4.3. The margin τ in Eq. (4) is still set to 0.01 for the three cases. Since most identities in TRAINING-SET-I are Caucasians, case (a) and case (c) can have more selected pairs in every mini-batch than case (b) when training the CNN models. For the sake of fairness, we also require that the number of selected pairs should not exceed a given threshold in every mini-batch. We set this threshold to 350, which makes sure that the three cases can have a similar number of pairs. The rank-1 identification rates are shown in Tab. 4. We can see the performance of case (c) drops the most, followed by the case (a) and case (b). The result indicates that age is more important than gender in learning strong discriminative features. This may be due to the fact that age is a variable with a larger range of values, which potentially provides more information.

4.6 Experiments on RGB-D Fusion Schemes

In Sec 4.3, we concatenate RGB and depth into a six-channel data as input to the 28-layer ResNet, which is a signal-level fusion scheme for RGB-D face recognition (see Fig. 8(a)).

Fig. 8. The diagrams of the two representative fusion schemes for RGB-D face recognition: (a) signal level fusion, (b) feature level fusion.

It is worth noting that there are other ways to fuse RGB and depth information. Similar to [2], we also explore a feature-level fusion scheme (see Fig. 8(b)), which fuses the features extracted separately from RGB and depth modality networks and then feeds them into the classification layers. We use the same 28-layer ResNet architecture for both the

Fig. 9. Samples from our small-scale RGB-D dataset under different poses. In each row, we show RGB-D data of one individual captured by PrimeSense under the pose with yaw ranging from left 90° to right 90° .

TABLE 5

Comparison of face identification rates (%) and training time (hours) using the RGB-D data of TRAINING-SET-I with different loss functions and fusion schemes. Here “front”, “mid” and “back” indicate fusion operations ahead of the fourth convolution layer, the eighth convolution layer and the first fully connected layer of ResNet-28 [8], respectively.

	signal-level	feature-level		
		front	mid	back
softmax	88.93	88.67	89.44	82.67
softmax + attribute-aware	96.26	95.70	96.47	94.90
training time	~10h	~14h	~14h	~14h

RGB and depth data and test three different types of the feature-level fusion schemes, which are implemented at the front, the middle and the back of ResNet-28 respectively. More concretely, their fusion operation is performed ahead of the fourth convolution layer, the eighth convolution layer and the first fully connected layer of ResNet-28, respectively. The performance of these feature-level fusion schemes is then compared with the signal-level fusion scheme. Since the feature-level fusion schemes involve more parameters, we need to set a smaller batch size for training, which in turn requires more training iterations. In our experiments, we set the batch size to 100, and train for 180K iterations. Each feature-level fusion scheme takes about 14 hours to train, whereas the signal-level fusion scheme takes only about 10 hours. Tab. 5 compares the performance of different fusion schemes. It can be observed that only the feature-level scheme implemented at the middle outperforms the signal-level scheme by a small margin, although with a 40% increase of training time. The feature-level scheme implemented at the back performs much worse than all other schemes, especially for the softmax loss. One possible reason is that the features are concatenated using a fully connected layer followed by a loss function, which is not enough for feature fusion. As pointed out in [2], the fusion of multiple feature descriptors can improve face recognition rates. However, we need to carefully design fusion schemes to achieve better performance.

4.7 Experiments on a Small-Scale RGB-D Dataset Under Pose Variation

Most identities of the RGB-D dataset in Sec. 4.1 are with a frontal pose. To demonstrate the robustness of our method for larger poses, we construct another small-scale RGB-D facial dataset that is captured by PrimeSense camera

in the same indoor scene and contains about 25k RGB-D images of 952 identities, where each identity also has corresponding attributes including gender, ethnicity, and age. The new dataset contains only young Asians with 30% male and 70% female, and covers pose with yaw ranging from left 90° to right 90° . Some samples from this new dataset are shown in Fig. 9. This dataset is available at http://staff.ustc.edu.cn/~juyong/RGBD_dataset.html. We use the first 852 identities to construct a training set and the remaining 100 identities for testing. Similar to [11], the first frontal image of each identity is placed in the gallery, with the remaining images used as probes.

We train the CNN model jointly supervised with softmax loss and attribute-aware loss on the RGB-D data of the training set. The architecture of the CNN model is the same as in Sec. 4.3. The weight of the attribute-aware loss is set to 0.0001 and the distance margin τ is set to 0.001. The learning rate begins at 0.1, and is divided by ten after 1.1k and 1.7k iterations. The training ends at 2k iterations. The rank-1 identification rates are shown in Tab 6. We can see that the model with the combined softmax and attribute-aware loss outperforms the model with only the softmax loss across the whole range of pose variation, with greater improvements for more extreme poses. This example shows that attribute-aware loss can enhance the robustness of facial features with respect to pose variation.

4.8 Experiments on Public RGB-D Datasets

Although our RGB-D face recognition model is trained with data captured by PrimeSense cameras, it can also be applied to other RGB-D data with good generalization ability. In the following, we test our RGB-D face recognition model trained using TRAINING-SET-I on some public RGB-D datasets, and compare it with other RGB-D face recognition methods.

Performance on FRGCv2. The FRGCv2 dataset [28] contains 4,950 scans and texture images of 557 identities captured across multiple sessions and with two kinds of expressions (e.g., neutral and smile). To test face identification performance, we place the first available neutral scan and the texture image of each identity in the gallery, and use the remaining RGB-D images as probes. Tab. 7 reports the rank-1 identification rates of different methods, while Fig. 10(a) shows the cumulative match characteristic (CMC) curves. Our method performs the best under all three data modalities (RGB, Depth, and RGB + Depth).

Performance on Bosphorus. The Bosphorus dataset [34] contains 4,666 scans including color and depth images from 105 identities with rich expression variations, poses, and

TABLE 6
Comparison of face identification rates (%) using the RGB-D data of the small-scale training set under different poses.

Loss functions	-90°	$(-90^\circ, -45^\circ)$	$(-45^\circ, +45^\circ)$	$(+45^\circ, +90^\circ)$	$+90^\circ$
softmax	62.65	88.36	98.59	88.38	50.12
softmax + attribute-aware	81.15	93.97	99.65	92.96	65.88

Fig. 10. Evaluation of our method on three public data sets: FRGC v2 [28], Bosphorus [34] and 3D-TEC [41]. For each test data set, we show cumulative match characteristic (CMC) curves of three models trained using RGB, depth, and RGB-D of TRAINING-SET-I respectively.

occlusions. There are 743 scans that only include expression variations. In our experiment, 105 first neutral depth and color images from each identity are used as the gallery set, and 638 non-neutral scans are used as the probe set. We report the rank-1 identification rate in Tab. 7 and show the CMC curves in Fig. 10(b). Although our training data does not include the public data sets and our trained model is not fine-tuned on any gallery of public test data sets, our model can still achieve the best results on Bosphorus dataset.

Performance on 3D-TEC. The 3D-TEC dataset [41] contains 107 pairs of twins acquired using a Minolta VIVID 910 3D scanner under controlled illumination and background. We use the standard protocol for 3D-TEC (four scenarios: Case I, II, III, and IV) described in [41]. We show the rank-1 identification rate in Tab. 7 and the CMC curves in Figs. 10(c)-(f). Cases III and IV are challenging scenarios where the system does not control the expressions of the subject in a gallery set of twins. Since the depth image of the training data for our models is captured by PrimeSense cameras and of low quality, our models do not work well for these scenarios that require a distinction between subtle differences in the depth data. This also explains why our RGB model performs even better than our RGB-D model.

4.9 Comparison with Other Methods that Utilize Attributes

To demonstrate the effectiveness of our attribute-aware loss, we also compare our method with GTNN [11], a state-of-the-art face recognition method utilizing attributes. GTNN fuses the facial recognition features (FRFs) and facial attribute features (FAFs) together via a nonlinear tensor-based fusion method, and the fused features are used for the final face recognition. Different from GTNN, our method needs attribute information only for the training stage but not the testing stage.

Training & test sets. We train GTNN and our model on publicly available web-collected RGB datasets, including CASIA-Webface [50] and CelebA [25]. CASIA-Webface includes about 494K facial RGB images of 10,575 identities without facial attributes. CelebA is a large-scale facial attributes dataset with more than 200K celebrity RGB images of 10,177 identities, each with 40 attribute annotations. Same as [11], we remove attributes that do not characterize a specific person and retain 13 attributes. The results are tested on the LFW dataset [12], which contains 13,233 RGB images of 5,749 identities. During testing, LFW is divided into ten predefined subsets for cross-validation. We follow the standard *Unrestricted, Labeled Outside Data Results* protocol

TABLE 7

Comparison of Rank-1 identification rates (%) with state-of-the-art methods on public data sets. Except for our method, the result data are taken from the respective papers [4], [5], [16], [21], [22], [53]. The best result for each data set is highlighted in bold font.

Method	Training data	Modality	FRGC	Bosphorus	3DTEC			
					Case I	Case II	Case III	Case IV
GoogleNet [39]	Private [53]	RGB	21.51	63.44	-	-	-	-
VGG-Face [27]	Private [53]	RGB	87.92	96.39	-	-	-	-
Ours	TRAINING-SET-I	RGB	95.69	96.08	95.79	98.13	84.11	84.11
Faltemier et al. [4]	FRGC [28]	Depth	-	-	94.40	93.50	72.40	72.90
Li et al. [22]	Part of Bosphorus [34]	Depth	96.30	95.40	93.90	96.30	90.70	91.60
Kim et al. [16]	FRGC [28]	Depth	-	99.20	94.80	94.80	81.30	79.90
FR3DNet [53]	Private [53]	Depth	97.06	96.18	-	-	-	-
Ours	TRAINING-SET-I	Depth	97.45	99.37	93.46	95.33	73.83	75.23
Goswami et al. [5]	None	RGB+Depth	-	-	95.80	94.30	90.10	92.40
Li et al. [21]	Part of FRGC [28]	RGB+Depth	95.20	99.40	-	-	-	-
Ours	TRAINING-SET-I	RGB+Depth	98.52	99.52	96.73	97.66	84.58	83.18

for testing.

Training CNNs. The input for GTNN consists of pre-extracted FRFs and FAFs. We train a CNN model (MODEL-A) with softmax loss on CASIA-Webface to extract FRFs. We also train a CNN model (MODEL-B) with the loss function MOON [32] on CelebA to extract FAFs. Then we use MODEL-A and MODEL-B to extract FRFs and FAFs on CASIA-Webface respectively. Therefore, we can train the GTNN (MODEL-C) with the softmax loss on those FRFs and FAFs. In order to compare with GTNN, we regard those FAFs as the attributes and train a CNN model (MODEL-D) with the combined softmax and attribute-aware loss on CASIA-Webface. MODEL-A, MODEL-B and MODEL-C are trained according to the details given in [11], [32], [46]. For training MODEL-D, the learning rate begins at 0.1, and is divided by ten after 16K and 24K iterations, respectively. The training ends at 28K iterations. The weight of the attribute-aware loss is 0.001 and the margin τ is set to 0.58, such that there are about 300 pairs of similar attribute vectors with different identities in one batch.

Results & Analysis. The face verification rates of four models on the LFW dataset are reported in Tab. 8. The GTNN [11] fusion method improves the verification rate compared with training with FRFs or FAFs alone, while our attribute-aware achieves the best accuracy among the four approaches. Another thing to note is the greater efficiency of our approach. Training with our attribute-aware loss layer improves the accuracy of the resulting FRFs, with an almost negligible increase in the number of training parameters. In the testing phase, the attribute-aware loss layer is not needed and incurs no overhead. In comparison, GTNN fuses FRFs and FAFs into high-dimensional features via a tensor-based approach that involves a large number of training parameters. Although the number of parameters can be reduced by Tucker decomposition as shown in [11], the rank of the tensor must be carefully chosen to achieve a good balance between accuracy and efficiency. Moreover, GTNN requires training the FRF model, the FAF model and the fusion model, whereas our approach only requires training a single FRF model.

TABLE 8

Comparison with fusion method GTNN [11] on LFW data set.

Models	Descriptions	Acc.(%)
MODEL-A	FRFs	97.43
MODEL-B	FAFs	73.92
MODEL-C	GTNN [11] fusion	97.78
MODEL-D	Our method	98.42

5 DISCUSSION & CONCLUSION

We have presented an attribute-aware loss function for CNN-based face recognition, which regularizes the distribution of learned recognition features with respect to additional attributes. The novel attribute-aware loss could help resolve the problem of uneven sampling in the training dataset, and improves the face recognition accuracy. Besides, we train an RGB-D based face recognition model using a large data set with over 100K identities. The experimental results demonstrate the effectiveness of our novel attribute-aware loss, and the good generalization ability of our trained RGB-D face recognition model.

Our work is the first method on using non-facial attributes which are invariant to the capture environments to regularize the face recognition feature mapping. In this work, although we only use gender, age, and ethnicity attribute for regularization, the experimental results still demonstrate big improvements. Other attributes could also be added for further study. Another point that needs to be further investigated is how the transformation formulation between recognition feature and attribute feature influences the final recognition performance. For example, we could consider replacing the linear transformation $\mathbf{G}$ with nonlinear transformation. Currently, we add the attribute-aware loss term to classification loss like softmax. How to combine it with metric learning algorithms like triplet loss used in [35] is an interesting research problem.

ACKNOWLEDGMENTS

This work was supported by National Natural Science Foundation of China (No. 61672481), and Youth Innovation Promotion Association CAS (No. 2018495).

REFERENCES

- [1] A. F. Abate, M. Nappi, D. Riccio, and G. Sabatino. 2d and 3d face recognition: A survey. *Pattern Recognition Letters*, 28(14):1885–1906, 2007.
- [2] J. Cui, H. Han, S. Shan, and X. Chen. Rgb-d face recognition: A comparative study of representative fusion schemes. In *Chinese Conference on Biometric Recognition*, pages 358–366. Springer, 2018.
- [3] J. Deng, J. Guo, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *arXiv preprint arXiv:1801.07698*, 2018.
- [4] T. C. Faltemier, K. W. Bowyer, and P. J. Flynn. A region ensemble for 3-d face recognition. *IEEE Transactions on Information Forensics and Security*, 3(1):62–73, 2008.
- [5] G. Goswami, M. Vatsa, and R. Singh. RGB-D face recognition with texture and attribute features. *IEEE Trans. Information Forensics and Security*, 9(10):1629–1640, 2014.
- [6] Y. Guo, J. Zhang, J. Cai, B. Jiang, and J. Zheng. Cnn-based real-time dense face reconstruction with inverse-rendered photo-realistic face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [7] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *Computer vision and pattern recognition, 2006 IEEE computer society conference on*, volume 2, pages 1735–1742. IEEE, 2006.
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [9] M. Hernandez, J. Choi, and G. Medioni. Near laser-scan quality 3-d face reconstruction from a low-quality depth stream. *Image and Vision Computing*, 36:61–69, 2015.
- [10] G. Hsu, Y. Liu, H. Peng, and P. Wu. Rgb-d-based face reconstruction and recognition. *IEEE Trans. Information Forensics and Security*, 9(12):2110–2118, 2014.
- [11] G. Hu, Y. Hua, Y. Yuan, Z. Zhang, Z. Lu, S. S. Mukherjee, T. M. Hospedales, N. M. Robertson, and Y. Yang. Attribute-enhanced face recognition with neural tensor fusion networks. In *ICCV*, pages 3764–3773, 2017.
- [12] G. B. Huang and E. Learned-Miller. Labeled faces in the wild: Updates and new reporting procedures. *Dept. Comput. Sci., Univ. Massachusetts Amherst, Amherst, MA, USA, Tech. Rep*, pages 14–003, 2014.
- [13] P. Jain, B. Kulis, J. V. Davis, and I. S. Dhillon. Metric and kernel learning using a linear transformation. *Journal of Machine Learning Research*, 13(Mar):519–547, 2012.
- [14] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [15] I. Kemelmacher-Shlizerman, S. M. Seitz, D. Miller, and E. 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.
- [16] D. Kim, M. Hernandez, J. Choi, and G. Medioni. Deep 3d face identification. *arXiv preprint arXiv:1703.10714*, 2017.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [18] B. Kumar, G. Carneiro, I. Reid, et al. Learning local image descriptors with deep siamese and triplet convolutional networks by minimising global loss functions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5385–5394, 2016.
- [19] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar. Attribute and simile classifiers for face verification. In *IEEE 12th International Conference on Computer Vision, ICCV 2009, Kyoto, Japan, September 27 - October 4, 2009*, pages 365–372, 2009.
- [20] Y. Lee, J. Chen, C. W. Tseng, and S. Lai. Accurate and robust face recognition from RGB-D images with a deep learning approach. In *Proceedings of the British Machine Vision Conference 2016, BMVC 2016, York, UK, September 19-22, 2016*, 2016.
- [21] B. Y. L. Li, M. Xue, A. S. Mian, W. Liu, and A. Krishna. Robust RGB-D face recognition using kinect sensor. *Neurocomputing*, 214:93–108, 2016.
- [22] H. Li, D. Huang, J.-M. Morvan, L. Chen, and Y. Wang. Expression-robust 3d face recognition via weighted sparse representation of multi-scale and multi-component local normal patterns. *Neurocomputing*, 133:179–193, 2014.
- [23] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphreface: Deep hypersphere embedding for face recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 1, 2017.
- [24] W. Liu, Y. Wen, Z. Yu, and M. Yang. Large-margin softmax loss for convolutional neural networks. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 507–516, 2016.
- [25] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3730–3738, 2015.
- [26] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [27] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proceedings of the British Machine Vision Conference 2015, BMVC 2015, Swansea, UK, September 7-10, 2015*, pages 41.1–41.12, 2015.
- [28] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *Computer vision and pattern recognition, 2005. CVPR 2005. IEEE computer society conference on*, volume 1, pages 947–954. IEEE, 2005.
- [29] R. Ranjan, C. D. Castillo, and R. Chellappa. L2-constrained softmax loss for discriminative face verification. *arXiv preprint arXiv:1703.09507*, 2017.
- [30] R. Ranjan, S. Sankaranarayanan, C. D. Castillo, and R. Chellappa. An all-in-one convolutional neural network for face analysis. In *12th IEEE International Conference on Automatic Face & Gesture Recognition*, pages 17–24, 2017.
- [31] E. Richardson, M. Sela, R. Or-El, and R. Kimmel. Learning detailed face reconstruction from a single image. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5553–5562, 2017.
- [32] E. M. Rudd, M. Günther, and T. E. Boulton. Moon: A mixed objective optimization network for the recognition of facial attributes. In *European Conference on Computer Vision*, pages 19–35. Springer, 2016.
- [33] P. Samangouei and R. Chellappa. Convolutional neural networks for attribute-based active authentication on mobile devices. In *8th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS 2016, Niagara Falls, NY, USA, September 6-9, 2016*, pages 1–8, 2016.
- [34] A. Savran, N. Alyüz, H. Dibeklioğlu, O. Çeliktutan, B. Gökberk, B. Sankur, and L. Akarun. Bosphorus database for 3d face analysis. In *European Workshop on Biometrics and Identity Management*, pages 47–56. Springer, 2008.
- [35] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015.
- [36] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 806–813, 2014.
- [37] K. Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, pages 1849–1857, 2016.
- [38] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3476–3483, 2013.
- [39] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015.
- [40] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708, 2014.
- [41] V. Vijayan, K. W. Bowyer, P. J. Flynn, D. Huang, L. Chen, M. Hansen, O. Ocegueda, S. K. Shah, and I. A. Kakadiaris. Twins 3d face recognition challenge. In *Biometrics (IJCB), 2011 International Joint Conference on*, pages 1–7. IEEE, 2011.
- [42] F. Wang, J. Cheng, W. Liu, and H. Liu. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25(7):926–930, 2018.
- [43] F. Wang, X. Xiang, J. Cheng, and A. L. Yuille. Normface: L₂ hypersphere embedding for face verification. In *Proceedings of the 2017 ACM on Multimedia Conference, MM 2017, Mountain View, CA, USA, October 23-27, 2017*, pages 1041–1049, 2017.
- [44] J. Wang, H. T. Do, A. Woznica, and A. Kalousis. Metric learning

- with multiple kernels. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 1170–1178. Curran Associates, Inc., 2011.
- [45] K. Q. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10(Feb):207–244, 2009.
- [46] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*, pages 499–515. Springer, 2016.
- [47] L. Wolf, T. Hassner, and I. Maoz. Face recognition in unconstrained videos with matched background similarity. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 529–534. IEEE, 2011.
- [48] Z. S. Xiang Wu, Ran He. A lightened cnn for deep face representation. *arXiv preprint arXiv:1511.02683v1*, 2015.
- [49] E. P. Xing, A. Y. Ng, M. I. Jordan, and S. Russell. Distance metric learning, with application to clustering with side-information. In *Proceedings of the 15th International Conference on Neural Information Processing Systems, NIPS’02*, pages 521–528, 2002.
- [50] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014.
- [51] H. Zhang, H. Han, J. Cui, S. Shan, and X. Chen. RGB-D face recognition via deep complementary and common feature learning. In *13th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2018, Xi’an, China, May 15-19, 2018*, pages 8–15, 2018.
- [52] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016.
- [53] S. Zulqarnain Gilani and A. Mian. Learning from millions of 3d scans for large-scale 3d face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1896–1905, 2018.

Luo Jiang is currently working towards the PhD degree at the University of Science and Technology of China. He obtained his bachelor degree in 2013 from the Huazhong University of Science and Technology, China. His research interests include computer graphics, image processing and deep learning.

Juyong Zhang is an associate professor in the School of Mathematical Sciences at University of Science and Technology of China. He received the BS degree from the University of Science and Technology of China in 2006, and the PhD degree from Nanyang Technological University, Singapore. His research interests include computer graphics, computer vision, and numerical optimization. He is an associate editor of The Visual Computer.

Bailin Deng is a lecturer in the School of Computer Science and Informatics at Cardiff University. He received the BEng degree in computer software (2005) and the MSc degree in computer science (2008) from Tsinghua University (China), and the PhD degree in technical mathematics from Vienna University of Technology (Austria). His research interests include geometry processing, numerical optimization, computational design, and digital fabrication. He is a member of the IEEE.