

Probabilistic Face Embeddings

Yichun Shi and Anil K. Jain

Michigan State University, East Lansing, MI

shiyichu@msu.edu, jain@cse.msu.edu

Abstract

Embedding methods have achieved success in face recognition by comparing facial features in a latent semantic space. However, in a fully unconstrained face setting, the facial features learned by the embedding model could be ambiguous or may not even be present in the input face, leading to noisy representations. We propose Probabilistic Face Embeddings (PFEs), which represent each face image as a Gaussian distribution in the latent space. The mean of the distribution estimates the most likely feature values while the variance shows the uncertainty in the feature values. Probabilistic solutions can then be naturally derived for matching and fusing PFEs using the uncertainty information. Empirical evaluation on different baseline models, training datasets and benchmarks show that the proposed method can improve the face recognition performance of deterministic embeddings by converting them into PFEs. The uncertainties estimated by PFEs also serve as good indicators of the potential matching accuracy, which are important for a risk-controlled recognition system.

1. Introduction

When humans are asked to describe a face image, they not only give the description of the facial attributes, but also the confidence associated with them. For example, if the eyes are blurred in the image, a person will keep the eye size as an uncertain information and focus on other features. Furthermore, if the image is completely corrupted and no attributes can be discerned, the subject may respond that he/her cannot identify this face. This kind of uncertainty (or confidence) estimation is common and important in human decision making.

On the other hand, the representations used in state-of-the-art face recognition systems are generally confidence-agnostic. These methods depend on an embedding model (e.g. Deep Neural Networks) to give a deterministic point representation for each face image in the latent feature space [27, 35, 20, 34, 4]. A point in the latent space represents the model’s estimation of the facial features in the given image. If the error in the estimation is somehow

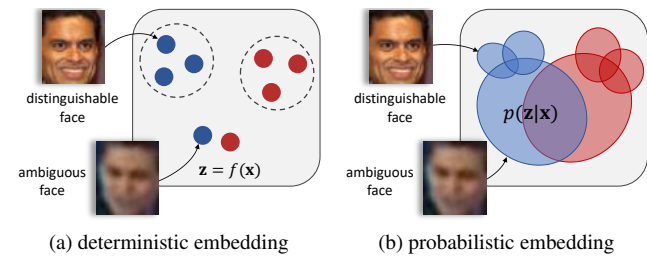


Figure 1: Difference between deterministic face embeddings and probabilistic face embeddings (PFEs). Deterministic embeddings represent every face as a point in the latent space without regards to its feature ambiguity. Probabilistic face embedding (PFE) gives a distributional estimation of features in the latent space instead. **Best viewed in color.**

bounded, the distance between two points can effectively measure the semantic similarity between the corresponding face images. But given a low-quality input, where the expected facial features are ambiguous or absent in the image, a large shift in the embedded points is inevitable, leading to false recognition (Figure 1a).

Given that face recognition systems have already achieved high recognition accuracies on relatively constrained face recognition benchmarks, e.g. LFW [10] and YTF [37], where most facial attributes can be clearly observed, recent face recognition challenges have moved on to more unconstrained scenarios, including surveillance videos [18, 23, 12] (See Figure 2). In these tasks, any type and degree of variation could exist in the face image, where most of the desired facial features learned by the representation model could be absent. Given this lack of information, it is unlikely to find a feature set that could always match these faces accurately. Hence state-of-the-art face recognition systems which obtained over 99% accuracy on LFW have suffered from a large performance drop on IARPA Janus benchmarks [18, 23, 12].

To address the above problems, we propose *Probabilistic Face Embeddings (PFEs)*, which give a distributional estimation instead of a point estimation in the latent space for each input face image (Figure 1b). The mean of the distribution can be interpreted as the most likely latent feature values while the span of the distribution represents the un-




Figure 2: Example images from IJB-A and IJB-S. The first columns show still images, followed by video frames of the respective subjects in the next three columns. These benchmarks present a more unconstrained recognition scenario where there is a large variability in the image quality.

certainty of these estimations. PFE can address the unconstrained face recognition problem in a two-fold way: (1) During matching (face comparison), PFE penalizes uncertain features (dimensions) and pays more attention to more confident features. (2) For low quality inputs, the confidence estimated by PFE can be used to reject the input or actively ask for human assistance to avoid false recognition. Besides, a natural solution can be derived to aggregate the PFE representations of a set of face images into a new distribution with lower uncertainty to increase the recognition performance. The implementation of PFE is open-sourced¹. The contributions of the paper can be summarized as below:

1. An uncertainty-aware probabilistic face embedding (PFE) which represents face images as distributions instead of points.
2. A probabilistic framework that can be naturally derived for face matching and feature fusion using PFE.
3. A simple method that converts existing deterministic embeddings into PFEs without additional training data.
4. Comprehensive experiments showing that the proposed PFE can improve face recognition performance of deterministic embeddings and can effectively filter out low-quality inputs to enhance the robustness of face recognition systems.

2. Related Work

Uncertainty Learning in DNNs To improve the robustness and interpretability of discriminant Deep Neural Networks (DNNs), deep uncertainty learning is getting more attention [14, 5, 15]. There are two main types of uncertainty: *model uncertainty* and *data uncertainty*. Model uncertainty refers to the uncertainty of model parameters given the training data and can be reduced by collecting additional training data [22, 24, 14, 5]. Data uncertainty accounts for the uncertainty in output whose primary source is the inherent noise in input data and hence cannot be eliminated with more training data [15]. The uncertainty studied in our work can be categorized as data uncertainty. Although techniques have been developed for estimating data uncertainty in different tasks, including classification and regression [15],

¹<https://github.com/seasonSH/Probabilistic-Face-Embeddings>

they are not suitable for our task since our target space is not well-defined by given labels². Variational Autoencoders [17] can also be regarded as a method for estimating data uncertainty, but it mainly serves a generation purpose. Specific to face recognition, some studies [6, 16, 45] have leveraged the model uncertainty for analysis and learning of face representations, but to our knowledge, ours is the first work that utilizes data uncertainty³ for recognition tasks.

Probabilistic Face Representation Modeling faces as probabilistic distributions is not a new idea. In the field of face template/video matching, there exists abundant literature on modeling the faces as probabilistic distributions [29, 1], subspace [3] or manifolds [1, 11] in the feature space. However, the input for such methods is a set of face images rather than a single face image, and they use a between-distribution similarity or distance measure, *e.g.* KL-divergence, for comparison, which does not penalize the uncertainty. Meanwhile, some studies [19, 9] have attempted to build a fuzzy model of a given face using the features of face parts. In comparison, the proposed PFE represents each single face image as a distribution in the latent space encoded by DNNs and we use an uncertainty-aware log likelihood score to compare the distributions.

Quality-aware Pooling In contrast to the methods above, recent work on face template/video matching aims to leverage the saliency of deep CNN embeddings by aggregating the deep features of all faces into a single compact vector [41, 21, 39, 7]. In these methods, a separate module learns to predict the quality of each face in the image set, which is then normalized for a weighted pooling of feature vectors. We show that a solution can be naturally derived under our framework, which not only gives a probabilistic explanation for quality-aware pooling methods, but also leads to a more general solution where an image set can also be modeled as a PFE representation.

3. Limitations of Deterministic Embeddings

In this section, we explain the problems of deterministic face embeddings from both theoretical and empirical views. Let \mathcal{X} denote the image space and \mathcal{Z} denote the latent feature space of D dimensions. An ideal latent space \mathcal{Z} should only encode *identity-salient* features and be *disentangled* from identity-irrelevant features. As such, each identity should have a unique intrinsic code $\mathbf{z} \in \mathcal{Z}$ that best represents this person and each face image $\mathbf{x} \in \mathcal{X}$ is an observation sampled from $p(\mathbf{x}|\mathbf{z})$. The process of training face embeddings can be viewed as a joint process of searching for such a latent space \mathcal{Z} and learning the inverse mapping

²Although we are given the identity labels, they cannot directly serve as target vectors in the latent feature space.

³Some in the literature have also used the terminology “data uncertainty” for a different purpose [40].

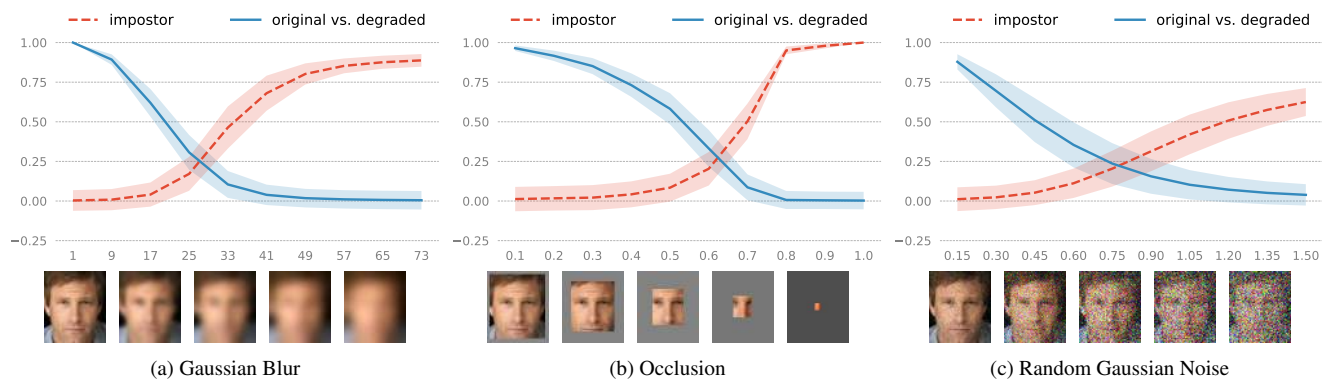


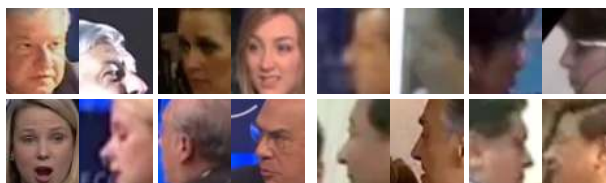
Figure 3: Illustration of *feature ambiguity dilemma*. The plots show the cosine similarity on LFW dataset with different degrees of degradation. Blue lines show the similarity between original images and their respective degraded versions. Red lines show the similarity between impostor pairs of degraded images. The shading indicates the standard deviation. With larger degrees of degradation, the model becomes more confident (very high/low scores) in a wrong way.

$p(\mathbf{z}|\mathbf{x})$. For deterministic embeddings, the inverse mapping is a Dirac delta function $p(\mathbf{z}|\mathbf{x}) = \delta(\mathbf{z} - f(\mathbf{x}))$, where f is the embedding function. Clearly, for any space \mathcal{Z} , given the possibility of noises in \mathbf{x} , it is unrealistic to recover the exact \mathbf{z} and the embedded point of a low-quality input would inevitably shift away from its intrinsic \mathbf{z} (no matter how much training data we have).

The question is whether this shift could be bounded such that we still have smaller intra-class distances compared to inter-class distances. However, this is unrealistic for fully unconstrained face recognition and we conduct an experiment to illustrate this. Let us start with a simple example: given a pair of identical images, a deterministic embedding will always map them to the same point and therefore the distance between them will always be 0, even if these images do not contain a face. This implies that “a pair of images being similar or even the same does not necessarily mean the probability of their belonging to the same person is high”.

To demonstrate this, we conduct an experiment by manually degrading the high-quality images and visualizing their similarity scores. We randomly select a high-quality image of each subject from the LFW dataset [10] and manually insert Gaussian blur, occlusion, and random Gaussian noise to the faces. In particular, we linearly increase the size of Gaussian kernel, occlusion ratio and the standard deviation of the noise to control the degradation degree. At each degradation level, we extract the feature vectors with a 64-layer CNN⁴, which is comparable to state-of-the-art face recognition systems. The features are normalized to a hyper-spherical embedding space. Then, two types of cosine similarities are reported: (1) similarity between pairs of original image and its respective degraded image, and (2) similarity between degraded images of different identities. As shown in Figure 3, for all the three types of degradation, the genuine similarity scores decrease to 0 while the

⁴trained on Ms-Celeb-1M [8] with AM-Softmax [33]



(a) Low-similarity Genuine Pairs (b) High-similarity Impostor Pairs

Figure 4: Example genuine pairs from IJB-A dataset estimated with the lowest similarity scores and impostor pairs with the highest similarity scores (among all possible pairs) by a 64-layer CNN model. The genuine pairs mostly consist of one high-quality and one low-quality image while the impostor pairs are all low-quality images. Note that these pairs are not templates in the verification protocol.

impostor similarity scores converge to 1.0! These indicate two types of errors that can be expected in a fully unconstrained scenario even when the model is very confident (very high/low similarity scores):

- (1) false accept of impostor low-quality pairs and
- (2) false reject of genuine cross-quality pairs.

To confirm this, we test the model on the IJB-A dataset by finding impostor/genuine image pairs with the highest/lowest scores, respectively. The situation is exactly as we hypothesized (See Figure 4). We call this *Feature Ambiguity Dilemma* which is observed when the deterministic embeddings are forced to estimate the features of ambiguous faces. The experiment also implies that there exist a *dark space* where the ambiguous inputs are mapped to and the distance metric is distorted.

4. Probabilistic Face Embeddings

To address the aforementioned problem caused by data uncertainty, we propose to encode the uncertainty into the face representation and take it into account during matching. Specifically, instead of building a model that gives a point estimation in the latent space, we estimate a distribution $p(\mathbf{z}|\mathbf{x})$ in the latent space to represent the potential

appearance of a person’s face⁵. In particular, we use a multivariate Gaussian distribution:

$$p(\mathbf{z}|\mathbf{x}_i) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_i, \boldsymbol{\sigma}_i^2 \mathbf{I}) \quad (1)$$

where $\boldsymbol{\mu}_i$ and $\boldsymbol{\sigma}_i$ are both a D -dimensional vector predicted by the network from the i^{th} input image \mathbf{x}_i . Here we only consider a diagonal covariance matrix to reduce the complexity of the face representation. This representation should have the following properties:

1. The center $\boldsymbol{\mu}$ should encode the most likely facial features of the input image.
2. The uncertainty $\boldsymbol{\sigma}$ should encode the model’s confidence along each feature dimension.

In addition, we wish to use a single network to predict the distribution. Considering that new approaches for training face embeddings are still being developed, we aim to develop a method that could convert existing deterministic face embedding networks to PFEs in an easy manner. In the followings, we first show how to compare and fuse the PFE representations to demonstrate their strength and then propose our method for learning PFEs.

Note 1. *Because of the space limit, we provide the proofs of all the propositions below in the Supplementary Material.*

4.1. Matching with PFEs

Given the PFE representations of a pair of images $(\mathbf{x}_i, \mathbf{x}_j)$, we can directly measure the “likelihood” of them belonging to the same person (sharing the same latent code): $p(\mathbf{z}_i = \mathbf{z}_j)$, where $\mathbf{z}_i \sim p(\mathbf{z}|\mathbf{x}_i)$ and $\mathbf{z}_j \sim p(\mathbf{z}|\mathbf{x}_j)$. Specifically,

$$p(\mathbf{z}_i = \mathbf{z}_j) = \int p(\mathbf{z}_i|\mathbf{x}_i)p(\mathbf{z}_j|\mathbf{x}_j)\delta(\mathbf{z}_i - \mathbf{z}_j)d\mathbf{z}_i d\mathbf{z}_j. \quad (2)$$

In practice, we would like to use the log likelihood instead, whose solution is given by:

$$\begin{aligned} s(\mathbf{x}_i, \mathbf{x}_j) &= \log p(\mathbf{z}_i = \mathbf{z}_j) \\ &= -\frac{1}{2} \sum_{l=1}^D \left(\frac{(\mu_i^{(l)} - \mu_j^{(l)})^2}{\sigma_i^{2(l)} + \sigma_j^{2(l)}} + \log(\sigma_i^{2(l)} + \sigma_j^{2(l)}) \right) \\ &\quad - \text{const}, \end{aligned} \quad (3)$$

where $\text{const} = \frac{D}{2} \log 2\pi$, $\mu_i^{(l)}$ refers to the l^{th} dimension of $\boldsymbol{\mu}_i$ and similarly for $\sigma_i^{(l)}$.

Note that this symmetric measure can be viewed as the expectation of likelihood of one input’s latent code conditioned on the other, that is

$$\begin{aligned} s(\mathbf{x}_i, \mathbf{x}_j) &= \log \int p(\mathbf{z}|\mathbf{x}_i)p(\mathbf{z}|\mathbf{x}_j)d\mathbf{z} \\ &= \log \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\mathbf{x}_i)} [p(\mathbf{z}|\mathbf{x}_j)] \\ &= \log \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\mathbf{x}_j)} [p(\mathbf{z}|\mathbf{x}_i)]. \end{aligned} \quad (4)$$

⁵following the notations in Section 3.

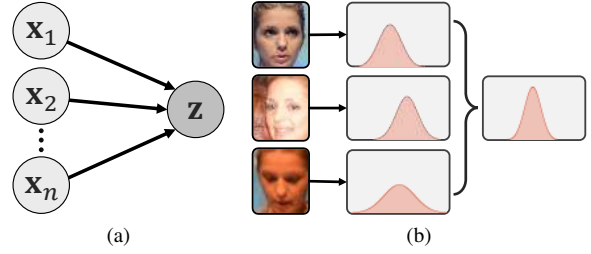


Figure 5: Fusion with PFEs. (a) Illustration of the fusion process as a directed graphical model. (b) Given the Gaussian representations of faces (from the same identity), the fusion process outputs a new Gaussian distribution in the latent space with a more precise mean and lower uncertainty.

As such, we call it *mutual likelihood score (MLS)*. Different from KL-divergence, this score is unbounded and cannot be seen as a distance metric. It can be shown that the squared Euclidean distance is equivalent to a special case of MLS when all the uncertainties are assumed to be the same:

Property 1. *If $\sigma_i^{(l)}$ is a fixed number for all data \mathbf{x}_i and dimensions l , MLS is equivalent to a scaled and shifted negative squared Euclidean distance.*

Further, when the uncertainties are allowed to be different, we note that MLS has some interesting properties that make it different from a distance metric:

1. *Attention mechanism:* the first term in the bracket in Equation (3) can be seen as a weighted distance which assigns larger weights to less uncertain dimensions.
2. *Penalty mechanism:* the second term in the bracket in Equation (3) can be seen as a penalty term which penalizes dimensions that have high uncertainties.
3. If either input \mathbf{x}_i or \mathbf{x}_j has large uncertainties, MLS will be low (because of penalty) irrespective of the distance between their mean.
4. Only if both inputs have small uncertainties and their means are close to each other, MLS could be very high.

The last two properties imply that PFE could solve the feature ambiguity dilemma if the network can effectively estimate $\boldsymbol{\sigma}_i$.

4.2. Fusion with PFEs

In many cases we have a template (set) of face images, for which we need to build a compact representation for matching. With PFEs, a conjugate formula can be derived for representation fusion (Figure 5). Let $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ be a series of observations (face images) from the same identity and $p(\mathbf{z}|\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ be the posterior distribution after the n^{th} observation. Then, assuming all the observations are conditionally independent (given the latent code \mathbf{z}). It can be shown that:

$$p(\mathbf{z}|\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n+1}) = \alpha \frac{p(\mathbf{z}|\mathbf{x}_{n+1})}{p(\mathbf{z})} p(\mathbf{z}|\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n), \quad (5)$$

where α is a normalization factor. To simplify the notations, let us only consider a one-dimensional case below; the solution can be easily extended to the multivariate case.

If $p(\mathbf{z})$ is assumed to be a noninformative prior, *i.e.* $p(\mathbf{z})$ is a Gaussian distribution whose variance approaches ∞ , the posterior distribution in Equation (5) is a new Gaussian distribution with lower uncertainty (See supplementary material). Further, given a set of face images $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, the parameters of the fused representation can be directly given by:

$$\hat{\mu}_n = \sum_{i=1}^n \frac{\hat{\sigma}_n^2}{\sigma_i^2} \mu_i, \quad (6)$$

$$\frac{1}{\hat{\sigma}_n^2} = \sum_{i=1}^n \frac{1}{\sigma_i^2}. \quad (7)$$

In practice, because the conditional independence assumption is usually not true, *e.g.* video frames include a large amount of redundancy, Equation (7) will be biased by the number of images in the set. Therefore, we take dimension-wise minimum to obtain the new uncertainty.

Relationship to Quality-aware Pooling If we consider a case where all the dimensions share the same uncertainty σ_i for i^{th} input and let the quality value $q_i = \frac{1}{\sigma_i^2}$ be the output of the network. Then Equation (6) can be written as

$$\hat{\mu}_n = \frac{\sum_{i=1}^n q_i \mu_i}{\sum_j q_j}. \quad (8)$$

If we do not use the uncertainty after fusion, the algorithm will be the same as recent quality-aware aggregation methods for set-to-set face recognition [41, 21, 39].

4.3. Learning

Note that any deterministic embedding f , if properly optimized, can indeed satisfy the properties of PFEs: (1) the embedding space is a disentangled identity-salient latent space and (2) $f(\mathbf{x})$ represents the most likely features of the given input in the latent space. As such, in this work we consider a stage-wise training strategy: given a pre-trained embedding model f , we fix its parameters, take $\mu(\mathbf{x}) = f(\mathbf{x})$, and optimize an additional uncertainty module to estimate $\sigma(\mathbf{x})$. When the uncertainty module is trained on the same dataset of the embedding model, this stage-wise training strategy allows us to have a more fair comparison between PFE and the original embedding $f(\mathbf{x})$ than an end-to-end learning strategy.

The uncertainty module is a network with two fully-connected layers which shares the same input as of the bottleneck layer⁶. The optimization criteria is to maximize the

⁶Bottleneck layer refers to the layer which outputs the original face embedding.

mutual likelihood score of all genuine pairs $(\mathbf{x}_i, \mathbf{x}_j)$. Formally, the loss function to minimize is

$$\mathcal{L} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} -s(\mathbf{x}_i, \mathbf{x}_j) \quad (9)$$

where \mathcal{P} is the set of all genuine pairs and s is defined in Equation (3). In practice, the loss function is optimized within each mini-batch. Intuitively, this loss function can be understood as an alternative to maximizing $p(\mathbf{z}|\mathbf{x})$: if the latent distributions of all possible genuine pairs have a large overlap, the latent target \mathbf{z} should have a large likelihood $p(\mathbf{z}|\mathbf{x})$ for any corresponding \mathbf{x} . Notice that because $\mu(\mathbf{x})$ is fixed, the optimization wouldn't lead to the collapse of all the $\mu(\mathbf{x})$ to a single point.

5. Experiments

In this section, we first test the proposed PFE method on standard face recognition protocols to compare with deterministic embeddings. Then we conduct qualitative analysis to gain more insight into how PFE behaves. *Due to the space limit, we provide the implementation details in the supplementary material.*

To comprehensively evaluate the efficacy of PFEs, we conduct the experiments on 7 benchmarks, including the well known **LFW** [10], **YTF** [37], **MegaFace** [13] and four other more unconstrained benchmarks:

CFP [28] contains 7,000 frontal/profile face photos of 500 subjects. We only test on the frontal-profile (FP) protocol, which includes 7,000 pairs of frontal-profile faces.

IJB-A [18] is a template-based benchmark, containing 25,813 faces images of 500 subjects. Each template includes a set of still photos or video frames. Compared with previous benchmarks, the faces in IJB-A have larger variations and present a more unconstrained scenario.

IJB-C [23] is an extension of IJB-A with 140,740 faces images of 3,531 subjects. The verification protocol of IJB-C includes more impostor pairs so we can compute True Accept Rates (TAR) at lower False Accept Rates (FAR).

IJB-S [12] is a surveillance video benchmark containing 350 surveillance videos spanning 30 hours in total, 5,656 enrollment images, and 202 enrollment videos of 202 subjects. Many faces in this dataset are of extreme pose or low-quality, making it one of the most challenging face recognition benchmarks (See Figure 2 for example images).

We use the CASIA-WebFace [42] and a cleaned version⁷ of MS-Celeb-1M [8] as training data, from which we remove the subjects that are also included in the test datasets⁸.

⁷https://github.com/inlmouse/MS-Celeb-1M_WashList.

⁸84 and 4,182 subjects were removed from CASIA-WebFace and MS-Celeb-1M, respectively.

Base Model	Representation	LFW	YTF	CFP-FP	IJB-A
Softmax + Center Loss [35]	Original	98.93	94.74	93.84	78.16
	PFE	99.27	95.42	94.51	80.83
Triplet [27]	Original	97.65	93.36	89.76	60.82
	PFE	98.45	93.96	90.04	61.00
A-Softmax [20]	Original	99.15	94.80	92.41	78.54
	PFE	99.32	94.94	93.37	82.58
AM-Softmax [33]	Original	99.28	95.64	94.77	84.69
	PFE	99.55	95.92	95.06	87.58

Table 1: Results of models trained on CASIA-WebFace. “Original” refers to the deterministic embeddings. The better performance among each base model are shown in bold numbers. “PFE” uses mutual likelihood score for matching. IJB-A results are verification rates at FAR=0.1%.

Method	Training Data	LFW	YTF	MF1 Rank1	MF1 Veri.
DeepFace+ [31]	4M	97.35	91.4	-	-
FaceNet [27]	200M	99.63	95.1	-	-
DeepID2+ [30]	300K	99.47	93.2	-	-
CenterFace [35]	0.7M	99.28	94.9	65.23	76.52
SphereFace [20]	0.5M	99.42	95.0	75.77	89.14
ArcFace [4]	5.8M	99.83	98.02	81.03	96.98
CosFace [34]	5M	99.73	97.6	77.11	89.88
L2-Face [25]	3.7M	99.78	96.08	-	-
Baseline	4.4M	99.70	97.18	79.43	92.93
PFE _{fuse}	4.4M	-	97.32	-	-
PFE _{fuse+match}	4.4M	99.82	97.36	78.95	92.51

Table 2: Results of our models (last three rows) trained on MS-Celeb-1M and state-of-the-art methods on LFW, YTF and MegaFace. The MegaFace verification rates are computed at FAR=0.0001%. “-” indicates that the author did not report the performance on the corresponding protocol.

5.1. Experiments on Different Base Embeddings

Since our method works by converting existing deterministic embeddings, we want to evaluate how it works with different base embeddings, *i.e.* face representations trained with different loss functions. In particular, we implement the following state-of-the-art loss functions: Softmax+Center Loss [35], Triplet Loss [27], A-Softmax [20] and AM-Softmax [33]⁹. To be aligned with previous work [20, 34], we train a 64-layer residual network [20] with each of these loss functions on the CASIA-WebFace dataset as base models. All the features are ℓ_2 -normalized to a hyper-spherical embedding space. Then we train the uncertainty module for each base model on the CASIA-WebFace again for 3,000 steps. We evaluate the performance on four benchmarks: LFW [10], YTF [37], CFP-FP [28] and IJB-A [18], which present different challenges in face recognition. The results are shown in Table 1. The PFE improves over the original representation in all cases, indicating the proposed method is robust with different embeddings and testing scenarios.

5.2. Comparison with State-Of-The-Art

To compare with state-of-the-art face recognition methods, we use a different base model, which is a 64-layer

⁹We also tried implementing ArcFace [4] but it does not converge well in our case. So we did not use it.

Method	Training Data	IJB-A (TAR@FAR)		CFP-FP
		0.1%	1.0%	
DR-GAN [32]	1M	53.9 ± 4.3	77.4 ± 2.7	93.41
Yin <i>et al.</i> [44]	0.5M	73.9 ± 4.2	77.5 ± 2.5	94.39
TPE [26]	0.5M	90.0 ± 1.0	93.4 ± 0.5	89.17
NAN [41]	3M	88.1 ± 1.1	94.1 ± 0.8	-
QAN [21]	5M	89.31 ± 3.92	94.20 ± 1.53	-
Cao <i>et al.</i> [2]	3.3M	90.4 ± 1.4	95.8 ± 0.6	-
Multicolumn [39]	3.3M	92.0 ± 1.3	96.2 ± 0.5	-
L2-Face [25]	3.7M	94.3 ± 0.5	97.00 ± 0.4	-
Baseline	4.4M	93.30 ± 1.29	96.15 ± 0.71	92.78
PFE _{fuse}	4.4M	94.59 ± 0.72	95.92 ± 0.73	-
PFE _{fuse+match}	4.4M	95.25 ± 0.89	97.50 ± 0.43	93.34

Table 3: Results of our models (last three rows) trained on MS-Celeb-1M and state-of-the-art methods on CFP (frontal-profile protocol) and IJB-A.

Method	Training Data	IJB-C (TAR@FAR)			
		0.001%	0.01%	0.1%	1%
Yin <i>et al.</i> [43]	0.5M	-	-	69.3	83.8
Cao <i>et al.</i> [2]	3.3M	74.7	84.0	91.0	96.0
Multicolumn [39]	3.3M	77.1	86.2	92.7	96.8
DCN [38]	3.3M	-	88.5	94.7	98.3
Baseline	4.4M	70.10	85.37	93.61	96.91
PFE _{fuse}	4.4M	83.14	92.38	95.47	97.36
PFE _{fuse+match}	4.4M	89.64	93.25	95.49	97.17

Table 4: Results of our models (last three rows) trained on MS-Celeb-1M and state-of-the-art methods on IJB-C.

network trained with AM-Softmax on the MS-Celeb-1M dataset. Then we fix the parameters and train the uncertainty module on the same dataset for 12,000 steps. In the following experiments, we compare 3 methods:

- **Baseline** only uses the original features of the 64-layer deterministic embedding along with cosine similarity for matching. Average pooling is used in case of template/video benchmarks.
- **PFE_{fuse}** uses the uncertainty estimation σ in PFE and Equation (6) to aggregate the features of templates but uses cosine similarity for matching. If the uncertainty module could estimate the feature uncertainty effectively, fusion with σ should be able to outperform average pooling by assigning larger weights to confident features.
- **PFE_{fuse+match}** uses σ both for fusion and matching (with mutual likelihood scores). Templates/videos are fused based on Equation (6) and Equation (7).

In Table 2 we show the results on three relatively easier benchmarks: LFW, YTF and MegaFace. Although the accuracy on LFW and YTF are nearly saturated, the proposed PFE still improves the performance of the original representation. Note that MegaFace is a biased dataset: because all the probes are high-quality images from FaceScrub, the positive pairs in MegaFace are both high-quality images while the negative pairs only contain at most one low-quality image¹⁰. Therefore, neither of the two types of error caused

¹⁰The negative pairs of MegaFace in the verification protocol only in-

Method	Training Data	Surveillance-to-Single					Surveillance-to-Booking					Surveillance-to-Surveillance				
		Rank-1	Rank-5	Rank-10	1%	10%	Rank-1	Rank-5	Rank-10	1%	10%	Rank-1	Rank-5	Rank-10	1%	10%
C-FAN [7]	5.0M	50.82	61.16	64.95	16.44	24.19	53.04	62.67	66.35	27.40	29.70	10.05	17.55	21.06	0.11	0.68
Baseline	4.4M	50.00	59.07	62.70	7.22	19.05	47.54	56.14	61.08	14.75	22.99	9.40	17.52	23.04	0.06	0.71
PFE _{fuse}	4.4M	53.44	61.40	65.05	10.53	22.87	55.45	63.17	66.38	16.70	26.20	8.18	14.52	19.31	0.09	0.63
PFE _{fuse+match}	4.4M	50.16	58.33	62.28	31.88	35.33	53.60	61.75	64.97	35.99	39.82	9.20	20.82	27.34	0.84	2.83

Table 5: Performance comparison on three protocols of IJB-S. The performance is reported in terms of rank retrieval (closed-set) and TPIR@FPIR (open-set) instead of the media-normalized version [12]. The numbers “1%” and “10%” in the second row refer to the FPIR.

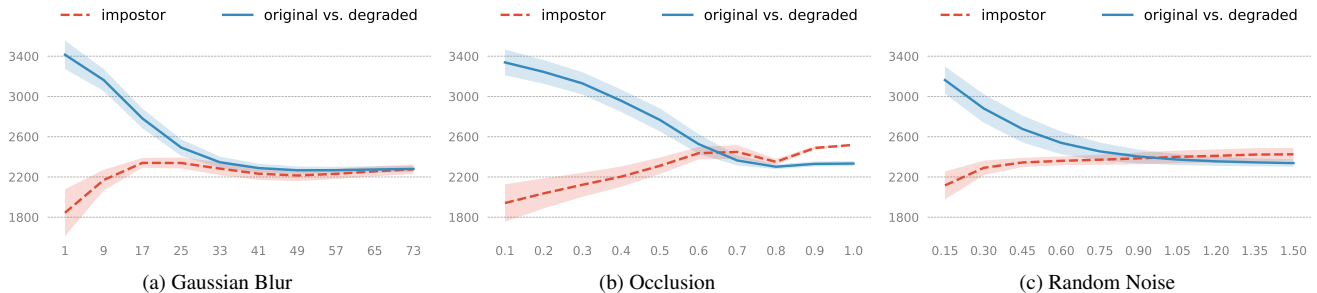


Figure 6: Repeated experiments on feature ambiguity dilemma with the proposed PFE. The same model in Figure 3 is used as the base model and is converted to a PFE by training an uncertainty module. No additional training data nor data augmentation is used for training.




Figure 7: Example genuine pairs from IJB-A dataset estimated with the lowest mutual likelihood scores and impostor pairs with the highest scores by the PFE version of the same 64-layer CNN model in Section 3. In comparison to Figure 4, most images here are high-quality ones with clear features, which can mislead the model to be confident in a wrong way. Note that these pairs are not templates in the verification protocol.

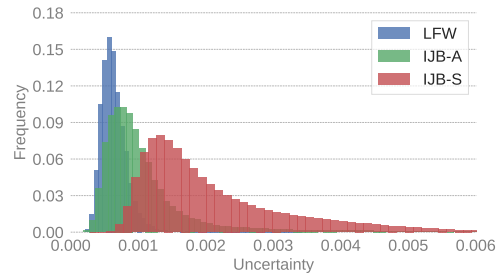


Figure 8: Distribution of estimated uncertainty on different datasets. Here, “Uncertainty” refers to the harmonic mean of σ across all feature dimensions. Note that the estimated uncertainty is proportional to the complexity of the datasets. **Best viewed in color.**

by the feature ambiguity dilemma (Section 3) will show up in MegaFace and it naturally favors deterministic embeddings. However, the PFE still maintains the performance in this case. We also note that such a bias, namely the target gallery images being of higher quality than the rest of gallery, would not exist in real world applications.

In Table 3 and Table 4 we show the results on three more challenging datasets: CFP, IJB-A and IJB-C. The images in these datasets present larger variations in pose, occlusion, etc, and facial features could be more ambiguous. As such, we can see that PFE achieves a more significant improvement on these three benchmarks. In particular on IJB-C at FAR=0.001%, PFE reduces the error rate by 64%. In addition, simply fusing the original features with the learned uncertainty (PFE_{fuse}) also helps the performance.

In Table 5 we report the results on three protocols of the latest benchmark, IJB-S. Again, PFE is able to improve the performance in most cases. Notice that the gallery templates in the “Surveillance-to-still” and “Surveillance-to-booking” all include high-quality frontal mugshots, which include those between probes and distractors.

present little feature ambiguity. Therefore, we only see a slight performance gap in these two protocols. But in the most challenging “surveillance-to-surveillance” protocol, larger improvement can be achieved by using uncertainty for matching. Besides, PFE_{fuse+match} improves the performance significantly on all the open-set protocols, which indicates that MLS has more impact on the absolute pairwise score than the relative ranking.

5.3. Qualitative Analysis

Why and when does PFE improve performance? We first repeat the same experiments in Section 3 using the PFE representation and MLS. The same network is used as the base model here. As one can see in Figure 6, although the scores of low-quality impostor pairs are still increasing, they converge to a point that is lower than the majority of genuine scores. Similarly, the scores of cross-quality genuine pairs converge to a point that is higher than the majority of impostor scores. This means the two types of errors discussed in Section 3 could be solved by PFE. This is further confirmed by the IJB-A results in Figure 7. Figure 8 shows

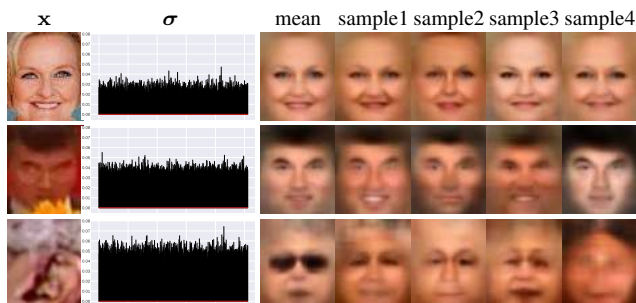


Figure 9: Visualization results on a high-quality, a low-quality and a mis-detected image from IJB-A. For each input, 5 images are reconstructed by a pre-trained decoder using the mean and 4 randomly sampled \mathbf{z} vectors from the estimated distribution $p(\mathbf{z}|\mathbf{x})$.

the distribution of estimated uncertainty on LFW, IJB-A and IJB-S. As one can see, the “variance” of uncertainty increases in the following order: LFW < IJB-A < IJB-S. Comparing with the performance in Section 5.2, we can see that PFE tends to achieve larger performance improvement on datasets with more diverse image quality.

What does DNN see and not see? To answer this question, we train a decoder network on the original embedding, then apply it to PFE by sampling \mathbf{z} from the estimated distribution $p(\mathbf{z}|\mathbf{x})$ of given \mathbf{x} . For a high-quality image (Figure 9 Row 1), the reconstructed images tend to be very consistent without much variation, implying the model is very certain about the facial features in this images. In contrast, for a lower-quality input (Figure 9 Row 2), larger variation can be observed from the reconstructed images. In particular, attributes that can be clearly discerned from the image (*e.g.* thick eye-brow) are still consistent while attributes cannot (*e.g.* eye shape) be discerned have larger variation. As for a mis-detected image (Figure 9 Row 3), significant variation can be observed in the reconstructed images: the model does not see any salient feature in the given image.

6. Risk-controlled Face Recognition

In many scenarios, we may expect a higher performance than our system is able to achieve or we may want to make sure the system’s performance can be controlled when facing complex application scenarios. Therefore, we would expect the model to reject input images if it is not confident. A common solution for this is to filter the images with a quality assessment tool. We show that PFE provides a natural solution for this task. We take all the images from LFW and IJB-A datasets for image-level face verification (We do not follow the original protocols here). The system is allowed to “filter out” a proportion of all images to maintain a better performance. We then report the TAR@FAR= 0.001% against the “Filter Out Rate”. We consider two criteria for filtering: (1) the detection score of MTCNN [36] and (2) a confidence value predicted by our uncertainty module. Here the confidence for i^{th} sample is defined as the inverse of har-




Figure 10: Example images from LFW and IJB-A that are estimated with the highest (H) confidence/quality scores and the lowest (L) scores by our method and MTCNN face detector.

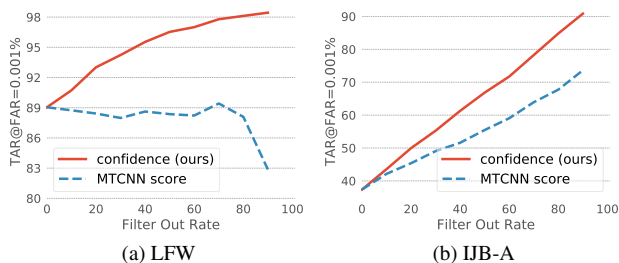


Figure 11: Comparison of verification performance on LFW and IJB-A (not the original protocol) by filtering a proportion of images using different quality criteria.

monic mean of σ_i across all dimensions. For fairness, both methods use the original deterministic embedding representations and cosine similarity for matching. To avoid saturated results, we use the model trained on CASIA-WebFace with AM-Softmax. The results are shown in Figure 11. As one can see, the predicted confidence value is a better indicator of the potential recognition accuracy of the input image. This is an expected result since PFE is trained under supervision for the particular model while an external quality estimator is unaware of the kind of features used for matching by the model. Example images with high/low confidence/quality scores are shown in Figure 10.

7. Conclusion

We have proposed probabilistic face embeddings (PFEs), which represent face images as distributions in the latent space. Probabilistic solutions were derived to compare and aggregate the PFE of face images. Unlike deterministic embeddings, PFEs do not suffer from the feature ambiguity dilemma for unconstrained face recognition. Quantitative and qualitative analysis on different settings showed that PFEs can effectively improve the face recognition performance by converting deterministic embeddings to PFEs. We have also shown that the uncertainty in PFEs is a good indicator for the “discriminative” quality of face images. In the future work we will explore how to learn PFEs in an end-to-end manner and how to address the data dependency within face templates.

References

- [1] Ognjen Arandjelovic, Gregory Shakhnarovich, John Fisher, Roberto Cipolla, and Trevor Darrell. Face recognition with image sets using manifold density divergence. In *CVPR*, 2005. 2
- [2] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *FG*, 2018. 6
- [3] Hakan Cevikalp and Bill Triggs. Face recognition based on image sets. In *CVPR*, 2010. 2
- [4] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 1, 6
- [5] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016. 2
- [6] Sixue Gong, Vishnu Naresh Boddeti, and Anil K Jain. On the capacity of face representation. *arXiv:1709.10433*, 2017. 2
- [7] Sixue Gong, Yichun Shi, and Anil K Jain. Video face recognition: Component-wise feature aggregation network (c-fan). In *ICB*, 2019. 2, 7
- [8] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *ECCV*, 2016. 3, 5
- [9] PS Hiremath, Ajit Danti, and CJ Prabhakar. Modelling uncertainty in representation of facial features for face recognition. In *Face recognition*. 2007. 2
- [10] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical report, University of Massachusetts, Amherst, 2007. 1, 3, 5, 6
- [11] Zhiwu Huang, Ruiping Wang, Shiguang Shan, Xianqiu Li, and Xilin Chen. Log-euclidean metric learning on symmetric positive definite manifold with application to image set classification. In *ICML*, 2015. 2
- [12] Nathan D. Kalka, Brianna Maze, James A. Duncan, Kevin J. OConnor, Stephen Elliott, Kaleb Hebert, Julia Bryan, and Anil K. Jain. IJB-S : IARPA Janus Surveillance Video Benchmark . In *BTAS*, 2018. 1, 2, 5, 7
- [13] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *CVPR*, 2016. 5
- [14] Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla. Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. In *BMVC*, 2015. 2
- [15] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*, 2017. 2
- [16] Salman Khan, Munawar Hayat, Waqas Zamir, Jianbing Shen, and Ling Shao. Striking the right balance with uncertainty. *arXiv:1901.07590*, 2019. 2
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2013. 2
- [18] Brendan F Klare, Ben Klein, Emma Taborsky, Austin Blanton, Jordan Cheney, Kristen Allen, Patrick Grother, Alan Mah, and Anil K Jain. Pushing the frontiers of unconstrained face detection and recognition: IARPA Janus Benchmark A. In *CVPR*, 2015. 1, 2, 5, 6
- [19] Haoxiang Li, Gang Hua, Zhe Lin, Jonathan Brandt, and Jianchao Yang. Probabilistic elastic matching for pose variant face verification. In *CVPR*, 2013. 2
- [20] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017. 1, 6
- [21] Yu Liu, Junjie Yan, and Wanli Ouyang. Quality aware network for set to set recognition. In *CVPR*, 2017. 2, 5, 6
- [22] David JC MacKay. A practical bayesian framework for back-propagation networks. *Neural Computation*, 1992. 2
- [23] Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Otto Charles, Anil K Jain, Niggel Tyler, Janet Anderson, Jordan Cheney, and Patrick Grother. Iarpa janus benchmark-c: Face dataset and protocol. In *ICB*, 2018. 1, 5
- [24] Radford M Neal. *Bayesian learning for neural networks*. PhD thesis, University of Toronto, 1995. 2
- [25] Rajeev Ranjan, Carlos D Castillo, and Rama Chellappa. L2-constrained softmax loss for discriminative face verification. *arXiv:1703.09507*, 2017. 6
- [26] Swami Sankaranarayanan, Azadeh Alavi, Carlos D Castillo, and Rama Chellappa. Triplet probabilistic embedding for face verification and clustering. In *BTAS*, 2016. 6
- [27] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015. 1, 6
- [28] Soumyadip Sengupta, Jun-Cheng Chen, Carlos Castillo, Vishal M Patel, Rama Chellappa, and David W Jacobs. Frontal to profile face verification in the wild. In *WACV*, 2016. 5, 6
- [29] Gregory Shakhnarovich, John W Fisher, and Trevor Darrell. Face recognition from long-term observations. In *ECCV*, 2002. 2
- [30] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deeply learned face representations are sparse, selective, and robust. In *CVPR*, 2015. 6
- [31] Yaniv Taigman, Ming Yang, Marc’ Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *CVPR*, 2014. 6
- [32] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *CVPR*, 2017. 6
- [33] Feng Wang, Jian Cheng, Weiyang Liu, and Haijun Liu. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 2018. 3, 6
- [34] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, 2018. 1, 6
- [35] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, 2016. 1, 6

- [36] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, 2016. 8
- [37] Lior Wolf, Tal Hassner, and Itay Maoz. Face recognition in unconstrained videos with matched background similarity. In *CVPR*, 2011. 1, 5, 6
- [38] Weidi Xie, Li Shen, and Andrew Zisserman. Comparator networks. In *ECCV*, 2018. 6
- [39] Weidi Xie and Andrew Zisserman. Multicolumn networks for face recognition. In *BMVC*, 2018. 2, 5, 6
- [40] Yong Xu, Xiaozhao Fang, Xuelong Li, Jiang Yang, Jane You, Hong Liu, and Shaohua Teng. Data uncertainty in face recognition. *IEEE Trans. on Cybernetics*, 2014. 2
- [41] Jiaolong Yang, Peiran Ren, Dongqing Zhang, Dong Chen, Fang Wen, Hongdong Li, and Gang Hua. Neural aggregation network for video face recognition. In *CVPR*, 2017. 2, 5, 6
- [42] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv:1411.7923*, 2014. 5
- [43] Bangjie Yin, Luan Tran, Haoxiang Li, Xiaohui Shen, and Xiaoming Liu. Towards interpretable face recognition. *arXiv:1805.00611*, 2018. 6
- [44] Xi Yin and Xiaoming Liu. Multi-task convolutional neural network for pose-invariant face recognition. *IEEE Trans. on Image Processing*, 2018. 6
- [45] Umara Zafar, Mubeen Ghafoor, Tehseen Zia, Ghufraan Ahmed, Ahsan Latif, Kaleem Razzaq Malik, and Abdullahi Mohamud Sharif. Face recognition with bayesian convolutional networks for robust surveillance systems. *EURASIP Journal on Image and Video Processing*, 2019. 2