

# Stating Comparison Score Uncertainty and Verification Decision Confidence Towards Transparent Face Recognition

Marco Huber<sup>1,2</sup>

marco.huber@igd.fraunhofer.de

Philipp Terhörst<sup>1,3,4</sup>

philipp.terhoerst@uni-paderborn.de

Florian Kirchbuchner<sup>1</sup>

florian.kirchbuchner@igd.fraunhofer.de

Naser Damer<sup>1,2</sup>

naser.damer@igd.fraunhofer.de

Arjan Kuijper<sup>1,2</sup>

arjan.kuijper@igd.fraunhofer.de

<sup>1</sup> Fraunhofer Institute for Computer Graphics Research IGD  
Darmstadt, Germany

<sup>2</sup> Technical University of Darmstadt  
Darmstadt, Germany

<sup>3</sup> Norwegian University of Science and Technology  
Gjøvik, Norway

<sup>4</sup> Paderborn University  
Paderborn, Germany

---

## Abstract

Face Recognition (FR) is increasingly used in critical verification decisions and thus, there is a need for assessing the trustworthiness of such decisions. The confidence of a decision is often based on the overall performance of the model or on the image quality. We propose to propagate model uncertainties to scores and decisions in an effort to increase the transparency of verification decisions. This work presents two contributions. First, we propose an approach to estimate the uncertainty of face comparison scores. Second, we introduce a confidence measure of the system's decision to provide insights into the verification decision. The suitability of the comparison scores uncertainties and the verification decision confidences have been experimentally proven on three face recognition models on two datasets.

## 1 Introduction

When human operators are asked if two face images show the same individual, they can intuitively state how sure they are about their decisions [22]. They may even conclude that they cannot make a meaningful, justifiable decision. The estimation of confidence and uncertainty regarding decision-making is important and common for humans [24, 26]. However, state-of-the-art face recognition models (FR) [0, 13, 21] do not offer these confidence or uncertainty estimates, even as they outperform humans in some cases [24]. In state-of-the-art FR systems, the similarity of two faces is determined by the similarity of their feature representations generated by the model. This similarity is known as a comparison score and is used to conclude a matching decision depending on a given threshold. Intuitively, the

distance of this score to the decision threshold can be interpreted as the confidence in the decision. This similarity, or rather its deviation from the decision threshold in either direction, can be perceived as an intuitive confidence in the decision. However, this intuitive confidence measure might not be optimal in all cases, especially when the decision threshold is optimized to low false match rates.

The importance of transparent decisions has a special significance in the field of FR. The importance of FR is growing in society and it is becoming ubiquitous and increasingly important in critical systems such as security, law enforcement, and access control [50]. In current automated FR systems, the decision of a match or non-match between two face images is commonly based on face templates extracted with deep learning networks, which makes creating a concrete notion of decision confidence rather a challenging one. This is further stressed by the fact that such verification decisions were shown to be biased towards a large set of factors [29], and thus their lack of confidence transparency might lead to subsequent ethically compromised actions.

This differentiation follows the definition of Peterson and Pitz [23] and refers to the *uncertainty* in the sense of *the belief about the variability of possible outcomes* and to the *confidence* in the sense of *the belief that a given prediction is correct*. Both terms are therefore not interchangeable. Consequently, for face verification, we can only determine the confidence for a verification decision or the uncertainty of a comparison score. The uncertainty of a comparison score describes the uncertainty in the score depending on the uncertainty of the data and the model. In contrast, the confidence of a model’s decisions describes the confidence in the made decision, with lower confidence the model is more unsure regarding its made decision. An example to illustrate the distinction between confidence and uncertainty is shown in Figure 1.



Figure 1: **Transparent Face Verification:** In addition to the usual information, such as the score and the decisions made, we propose the following to increase the transparency of the FR system: the score uncertainty (orange) and the decision confidence (green).

The first work that included uncertainty information in the FR process was done by Shi and Jain [26]. They proposed probabilistic face embeddings (PFEs), that estimate the uncertainty information of a single face image by a stage-wise learning process on top of a pre-trained FR model to incorporate the uncertainty of the facial features into the embedding. Chang et al. [3] enhanced the idea of PFEs and propose to simultaneously learn the feature and the uncertainty of the facial features. Related to the determination of the uncertainty and confidence of an FR model are the works on Face Image Quality (FIQ). The idea of FIQ is to assign each face a quality estimate in the sense of the utility of a face for FR purpose [25]. However, the interaction of two images during the comparison process is disregarded and FIQ can only be considered as an uncertainty or confidence estimate to a limited extent.

To the best of our knowledge, no work has investigated the uncertainty of a comparison score and the decision confidence in FR yet, besides interpreting the degree of similarity as

the intuitive confidence. The degree of uncertainty regarding a comparison score and the statement of decision confidence deepens the transparency of the used FR model regarding its decision and might prevent incorrect decisions e.g. in front of a court.

In this work, we propose two main contributions. First, we exploit the approximation of model uncertainty through dropout and propagate this uncertainty estimate to the scoring function to obtain a score uncertainty measure. Second, we use the obtained score uncertainty to formulate a function that allows us to determine the confidence of the model regarding its decision. In experiments, we evaluate the sanity of our proposed score uncertainty and compare our proposed decision confidence with the intuitive decision confidence. The comparison score uncertainties and the uncertainty-aware decision confidence can increase the transparency without the need for any further training or adjustment.

## 2 Related Work

In this section, we will look in more detail at PFEs that include uncertainty information into the stored face image embedding and then look at the differences and related work in the area of FIQ since no work yet focused on the exploitation of the score uncertainty. Although we take the deterministic embedding as a basis, our approach can be considered as an expression of PFEs, since we use the uncertainty of the embeddings to derive the score uncertainty. The topic of FIQ is related to the topic of FR uncertainty, as many FIQ approaches relate the utility of face images to model uncertainty.

### 2.1 Probabilistic Face Embeddings

PFEs for a single face image were first introduced by Shi and Jain [26]. They represent each face image as Gaussian distribution with the feature as the mean and the uncertainty of the features as the variance. The variances were obtained by training a model on top of an existing FR model by utilizing the proposed mutual likelihood score. Chang et al. [9] extended the idea and proposed to train features and uncertainty simultaneously. Li et al. [18] proposed a new optimization objective to counter some drawbacks of the approach proposed by Shi and Jain [26]. Li et al. [19] proposed to learn probabilistic ordinal embeddings to be able to perform uncertainty-aware regression. Debnath et al. [5] investigated how uncertainty in face embeddings can be used to produce face clusters of higher quality and proposed an uncertainty-aware face clustering algorithm. Chen et al. [4] proposed to improve the robustness and speed of PFEs by simplifying the matching function and using an output-constraint loss. None of the previous works exploited the calculated uncertainty to gain insights about the uncertainty of the scores or the confidence of the decision, which we propose in this work.

### 2.2 Uncertainty and Face Image Quality

The topic of uncertainty and FIQ are connected and not always clearly distinguishable as authors use the terms interchangeably [18, 51, 53]. Outside the face recognition community, Kendall and Gal [15] differentiate between two types of uncertainty: data uncertainty (aleatoric) that captures the noise in the input data (e.g. blurry face images) and model uncertainty (epistemic) that captures the uncertainty of the model. While FIQ assessment might be interpreted as estimating the data uncertainty independent from the model, approaches

exist that combine both uncertainties to get model-specific quality estimates. SER-FIQ [27] utilized the Bayesian approximation of model uncertainty via dropout to calculate model-specific FIQ scores. MagFace [24] even incorporates the quality as the magnitude into the created face embeddings. QMagFace [28] goes even further and incorporates FIQ into the comparison process. In contrast, EQFace [20] attached a quality-prediction network on an FR model to include FIQ scores in the training process.

It is important to note that FIQ always refers to a single image and does not consider the comparison itself. FIQ values estimate the overall utility of an image to be used in FR and not for a specific comparison pair which is in contrast to our approach. Our approach takes both images, reference and probe, into account to estimate the uncertainty. Some works utilized the comparison in different ways. In PCNet [61], the authors proposed to start from the confidence of mated pairs to derive an FIQ score for both involved images. Ou et al. [22] utilize the similarities between intra- and inter-class samples to derive a quality pseudo-label and therefore also incorporate both, model and data uncertainty and Bhattacharya et al. [11] interpreted the mean of the FIQ scores of a pair as a Face Quality Confidence Score.

## 3 Methodology

We propose two main contributions: 1) the score uncertainty methodology that allows the interpretation of the propagated error of the comparison score, 2) a decision confidence approach that allows the interpretation of the reliability of a decision. Since both of these contributions build on the model uncertainties on the feature-level, we will first introduce how these are estimated.

### 3.1 Embedding Uncertainty Estimation

Given a face image  $I$  and an FR model  $M$  that provides a face embedding function  $\theta$ , we use  $\theta$  to create a deterministic feature representation  $\theta(I)$ . To estimate the model uncertainty, we apply  $t = 100$  stochastic forward passes (as recommended by Terhörst et al. [27]) with different dropout patterns being applied to  $M$ . This was mathematically proved by Gal and Ghahramani [9] to estimate the model’s uncertainty. Doing this, we create a set of stochastic embeddings and calculate the uncertainty  $\Delta_I$  of the feature representation for each feature dimension of the image  $I$  as the standard deviation of this set. We obtain a probabilistic embedding consisting of the deterministic embedding and its uncertainty in each feature dimension. A visualization of the uncertainty estimation process of the embedding can be seen in Figure 2. The computational burden can be reduced by limiting the stochastic forward passes to the last layer.

### 3.2 Score Uncertainty

After the calculation of the uncertainty  $\Delta_I$  for an embedding  $\theta(I)$  of an image  $I$ , we obtain a probabilistic face embedding consisting of the actual deterministic embedding and the uncertainty in each feature dimension. To decide whether two embeddings represent the same identity or not, we need a decision measure. Typically, the cosine similarity is calculated to determine the similarity of two vectors that represent two face images [6, 13], but theoretically the method can be applied to any differentiable comparison function. The cosine




Figure 2: **Visualization of the uncertainty estimation of the embedding:** The input image  $I$  is used to create the deterministic representation  $\theta(I)$ .  $I$  is also propagated  $t$  times through different subnetworks of  $M$  to create a set of stochastic embeddings  $\theta_t(I)$ . These stochastic embeddings are used to estimate the uncertainty of the embedding in form of the standard deviation  $\Delta_f = \sigma(\theta_t(I))$ .

similarity  $S_C$  is defined as

$$S_C(X, Y) = \frac{X \cdot Y}{\|X\| \|Y\|} \quad (1)$$

where  $X$  and  $Y$  are face embeddings. Since the embeddings are often normalized before the comparisons, and thus mapped to a unit-sphere, the formula simplifies to:

$$S_C(X, Y) = \sum_{n=1}^N x_n \cdot y_n \quad (2)$$

as  $\|X\| \|Y\| = 1$ .  $N$  denotes the number of dimensions of the embeddings  $X, Y$ . To obtain the uncertainty of the comparison score, we make use of the propagation of uncertainty formula, presented in [1], while neglecting cross-feature correlations and assuming independence of the dimensions:

$$\Delta_E = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 \cdot \Delta^2} \quad (3)$$

where  $\Delta_E$  denotes the propagated uncertainty. The assumption of independence and the neglecting of the correlations was for computational simplification. An investigation of the features and their correlations leads to the insight, that the dimensions are to some degree disentangled and ways to incorporate or estimate the cross-correlations are open for future research. We set  $f$  to the cosine similarity function and combine both embeddings,  $X$  and  $Y$  to  $Z = [X, Y]$ . By this, we achieve:

$$\begin{aligned} \Delta_{S_C} &= \sqrt{\sum_{i=1}^{2n} \left(\frac{\partial f}{\partial z_i}\right)^2 \cdot \Delta_i^2} \\ &= \sqrt{\sum_{i=1}^n \left(\frac{\partial f}{\partial x_i}\right)^2 \cdot \Delta_{x_i}^2} + \sqrt{\sum_{i=1}^n \left(\frac{\partial f}{\partial y_i}\right)^2 \cdot \Delta_{y_i}^2} \\ &= \sqrt{\sum_{i=1}^n y_i^2 \cdot \Delta_{x_i}^2 + \sum_{i=1}^n x_i^2 \cdot \Delta_{y_i}^2} \end{aligned} \quad (4)$$

$\Delta_{S_C}$ , therefore, estimates the uncertainty of the comparison score  $S_C$  based on the individual uncertainties of the feature of both embeddings of the comparison  $X, Y$ . This estimation of the score uncertainty allows a deeper insight into the calculated similarity score as it additionally takes the uncertainty of the representation into account. If the uncertainty in the face image representation is larger, it also represents a higher uncertainty in the comparison score.

### 3.3 Decision Confidence

Since the score uncertainty  $\Delta_{S_C}$  only indicates how certain the model is about the score and not how confident it is about a decision made, we also propose a decision confidence approach. Following the definition of confidence, which refers to a decision and the belief that the decision made is correct, we incorporate the decision threshold  $d$  into our approach. Since we propagate uncertainties to derive the decision confidence, a derivable decision function is needed. Therefore, we approximate the decision function, usually in the form of a step function, by a modified sigmoid function  $\delta$ :

$$\delta(S_C) = \frac{1}{1 + e^{-\alpha(S_C - d)}} \quad (5)$$

where  $d$  denominates the given decision threshold,  $S_C$  the comparison score, and  $\alpha$  a scaling parameter to scale the range of the decision confidence. The  $\alpha$  value represents how sharp it resembles the decision step function and thus, it weights the trade-off between the importance of the uncertainty and the distance to the decision threshold. We now also apply the propagation of uncertainty formula [17] to gain a confidence estimate of the decision made depending on the decision threshold and on the uncertainty of the embeddings. The confidence estimate is defined as:

$$\Lambda_d(S_C) = 1 - \left[ \sum^N \delta(S_C) \cdot (1 - \delta(S_C)) \cdot \alpha x_n \cdot \Delta_{y_n}^2 + \sum^N \delta(S_C) \cdot (1 - \delta(S_C)) \cdot \alpha y_n \cdot \Delta_{x_n}^2 \right]^{\frac{1}{2}} \quad (6)$$

where  $x_n$  and  $y_n$  refers to features of the embeddings  $X$  and  $Y$  and  $\Delta_{x_n}$ , and  $\Delta_{y_n}$ , to the corresponding feature uncertainties. To obtain the confidence, we subtract the propagated uncertainty from 1 to get the intuitive understanding from increasing and decreasing confidence in a natural sense.

## 4 Experimental Setup

### 4.1 Face Recognition Models

To show the validity of our proposed uncertainty scores and decisions confidences, we use three different well-performing FR models as a basis. All models share the same ResNet-100 [18] architecture and have been trained with their specific loss function. The models were trained on the MS1M-V2 [6] dataset and were officially provided as pre-trained models by the authors. The used models are named ArcFace [6], CurricularFace [13], and MagFace [21]. We use these models to show that our proposed uncertainties and confidences are model-independent and can be transferred to a range of diverse models.

## 4.2 Datasets & Evaluation Benchmarks

To investigate the meaningfulness of our proposed uncertainty estimation approach and the proposed decision confidence, we evaluate on two difficult datasets with a higher variance in the quality of the images. We chose the Adience dataset [10] that consists of 26k images of over 2k different subjects as the images vary in a wide range of image quality, age, and pose. For the same reason, we also chose the XQFW [16] dataset that is based on the LFW [12] dataset and consists of synthetically degraded images as well as unaltered ones. Originally, intended to evaluate performance on cross-quality images, we use it to show that our uncertainty can also handle low-quality pairs as well as high-quality pairs. For this, we do not limit ourselves to the defined pairs for the evaluation of the uncertainty and the confidence and consider additional pairs. Both datasets are used to evaluate the score uncertainty and the decision confidence.

## 4.3 Evaluation Metrics

To investigate the rationale of our estimated uncertainty and our decision confidence, we use Error-vs-Reject curves (ERC) [10]. ERCs were proposed as a measure for the performance of a biometric quality estimator and they show the verification error rate over the fraction of unconsidered face images, where the images with the lowest quality are removed. Starting from the whole set of pairs, we subsequently remove the pairs with the highest uncertainty/lowest confidence and evaluate the performance again. We adapt the ERC and use them to evaluate the expressiveness of our uncertainty and confidence. For the uncertainty, the accuracy should ideally increase if uncertain scores are removed. For the confidence, the accuracy should increase, if pairs with low decision confidence are removed. For the experiments, we take all genuine pairs and select 30 random imposter images for each image from all possible imposter images. For the investigation of the usefulness of the score uncertainty, we compare our proposed score uncertainty with the MagFace [20] FIQ values. Although FIQ is designed and intended to be used to estimate the quality of single images, some evaluation protocols do suggest rejecting the comparison pairs if one of the images in the pair is of low-quality [21]. Based on that, we provide a minimized evaluation to compare the performance of such a technique (despite its limitations). On all ERC, we report the verification performance in terms of false non-match rate (FNMR) at fixed false match rate (FMR) following the international standard [22]. We also follow the best practice for high-security scenarios, e.g. automatic border control systems proposed by the European Border and Coast Guard Agency (Frontex) [8] and evaluate at the operation point of fixed FMR < 0.1%.

For the investigation of our decision confidence, we compare our proposed decision confidence with an intuitive confidence. This intuitive confidence is derived from the comparison score and its distance to the decision threshold  $D_{th}$ . The decision threshold  $D_{th}$  is also determined at the fixed FMR and does not change as comparisons are removed. The intuitive idea is, that the confidence of a correct decision increases with a higher or lower similarity, starting from the decision threshold that divides the similarity range into match and no match. The intuitive confidence  $\Lambda_{Int}$  can therefore be defined as:

$$\Lambda_{Int}(s) = |s - D_{th}| \quad (7)$$

where  $s$  is the given similarity score and  $D_{th}$  is the decision threshold of the system. A visualization of the intuitive confidence is shown in Figure 6(a). It is important to keep in



Figure 3: **Transparent Face Verification:** For each pair not only the decision is reported, but also the score uncertainty and the decision confidence. In the upper row even though the intuitive confidence is almost the same, the proposed decision confidence shows that the confidence in the false no-match decision on the left is lower in comparison to the correct no-match decision on the right. In the bottom row, the decision confidence is higher if the score uncertainty is smaller and the images are less similar.

mind, that as the intuitive confidence only depends on the similarity score and its distance to the threshold all pairs with the same similarity score share the same confidence.

## 5 Results


In this section, we investigate the performance of our two contributions separately. We start with evaluating the estimated score uncertainty by looking at ERCs to investigate the sanity of the proposed approach. In a second evaluation, we compare our proposed decision confidence against the intuitive confidence on two datasets to show that a deeper interpretation of the system’s confidence in a decision is possible. An overview of all contributions combined can be seen in Figure 3.

### 5.1 Investigating Score Uncertainty


To evaluate the utility of our proposed score uncertainty we plot ERCs on two datasets, Adience [17] and XQLFW [16], at the operation point of FMR=0.001 and report the performance as the FNMR at this point. We compare our proposed score uncertainty in a minimized evaluation with the FIQ provided by the MagFace embeddings. The ERCs are shown in Figure 4. We can observe that performance improves, i.e. fewer errors are made, when those comparisons are removed where the score has the highest uncertainty. This can be observed on all three models on both utilized datasets. Even if the minimum quality is in some cases better than the proposed uncertainty, we emphasize the differences: while the minimum quality approach reduces the problem of the inaccuracy of the comparison to the lowest quality, our approach processes the inaccuracy of the mapping and considers both images.

### 5.2 Analysing Decision Confidence


For the evaluation of the decision confidence, we also investigate the performance by looking at ERCs and compare our proposed decision confidence with the intuitive decision confidence as defined above in Subsection 4.3. It is important to note that the decision confidence



**Figure 4: Score uncertainty evaluation using ERC:** The comparisons are unconsidered either based on their score uncertainty or on their minimum quality based on the MagFace [20] embedding quality score. With a higher uncertainty regarding the score, more wrong decisions arise, accordingly the error decreases when the comparisons with the highest uncertainty are removed. The steadily decrease indicating the strong usefulness of the estimated uncertainties. (AF = ArcFace [16], MF = MagFace [20], CF = CurricularFace [13])



**Figure 5: Evaluating Decision Confidence (FMR@0.001):** The intuitive and the proposed confidence are compared. Overall the proposed decision confidence estimated the reliability of the decision more accurately than the intuitive decision confidence.



**Figure 6: Comparison of the confidence measures (Heatmap):** On the Adience dataset using the MagFace model. While the intuitive confidence only depends on the comparison score and its distance to the decision threshold (black asymptote), the proposed decision confidence also considers the uncertainty of the score and provides a more natural understanding of confidence.

is threshold dependent as it defines the confidence about a made decision. The results are presented in Figure 5 on the operational point of 0.1% FMR. For  $\alpha$ , we chose for the ArcFace model  $\alpha = 2$  and for the other models, MagFace and CurricularFace  $\alpha = 5$ . Other values for  $\alpha$  are shown in the supplementary material. With the  $\alpha$ -parameter the influence of the threshold can be scaled in relation to the uncertainty. On the Adience dataset, we observe that a minor improvement can be achieved using the proposed decision confidence in comparison to the intuitive confidence. Therefore, it is possible and useful to find more sophisticated ways of interpreting the confidence of a model’s decision besides the similarity of pairs and the distance to an operational threshold. A larger improvement can be observed in the experiments on the harder XQLEW dataset. As we do not limit the pairs to be cross-quality but also take high- and low-quality pairs into consideration, the decision confidence shows a higher increase in reliability of the made decision. For a deeper insight into our proposed decision confidence in comparison to the intuitive confidence, we show in Figure 6 the heatmaps of the confidences depending on the score. While for the intuitive solution the same comparison score leads to the same degree of confidence, in our proposed approach the uncertainty also is taken into consideration. This allows a more natural interpretation of confidence as it not only depends on the similarity but also on the uncertainty of the similarity.

## 6 Conclusion

In this work, we proposed 1) to estimate the uncertainty of the comparison score of FR models, 2) to provide decision confidence on the FR systems’ decision. After showing the meaningfulness of our proposed score uncertainty, we proposed to use the score uncertainty to derive decision confidence that is able to beat the intuitive confidence derived from the similarity of the embedded face images and the distance to the decision threshold. By introducing the new concepts of score uncertainty and decision confidence we highlight the need for transparency in the biometric decision-making process. This is necessary to increase the trust and reliability of FR systems and to draw attention to possible biases in form of lacking decision confidence, especially as FR becomes more ubiquitous and widespread in our society.

## Acknowledgement

This research work has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. This work was also partially carried out during the tenure of an ERCIM 'Alain Bensoussan' Fellowship Programme.

## References

- [1] Shubhobrata Bhattacharya, Chirag Kyal, and Aurobinda Routray. Simplified face quality assessment (SFQA). *Pattern Recognit. Lett.*, 147:108–114, 2021.
- [2] Fadi Boutros, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Elasticface: Elastic margin loss for deep face recognition. In *CVPR Workshops*, pages 1577–1586. IEEE, 2022.
- [3] Jie Chang, Zhonghao Lan, Changmao Cheng, and Yichen Wei. Data uncertainty learning in face recognition. In *CVPR*, pages 5709–5718. Computer Vision Foundation / IEEE, 2020.
- [4] Kai Chen, Taihe Yi, and Qi Lv. Fast and reliable probabilistic face embeddings based on constrained data uncertainty estimation. *Image Vis. Comput.*, 121:104429, 2022.
- [5] Biplob Debnath, Giuseppe Coviello, Yi Yang, and Srimat Chakradhar. Uac: An uncertainty-aware face clustering algorithm. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 3487–3495, October 2021.
- [6] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, pages 4690–4699. Computer Vision Foundation / IEEE, 2019.
- [7] Eran Eidinger, Roe Enbar, and Tal Hassner. Age and gender estimation of unfiltered faces. *IEEE Trans. Inf. Forensics Secur.*, 9(12):2170–2179, 2014.
- [8] Frontex. Best practice technical guidelines for automated border control (abc) systems, 2015.
- [9] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1050–1059. JMLR.org, 2016.
- [10] Patrick Grother and Elham Tabassi. Performance of biometric quality measures. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(4):531–543, 2007.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778. IEEE Computer Society, 2016.
- [12] 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 07-49, University of Massachusetts, Amherst, October 2007.

- [13] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, and Feiyue Huang. Curricularface: Adaptive curriculum learning loss for deep face recognition. In *CVPR*, pages 5900–5909. Computer Vision Foundation / IEEE, 2020.
- [14] ISO/IEC 19795-1:2006). ISO/IEC 19795-1:2006 Information technology — Biometric performance testing and reporting. Standard, International Organization for Standardization, 2016.
- [15] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*, pages 5574–5584, 2017.
- [16] Martin Knoche, Stefan Hörmann, and Gerhard Rigoll. Cross-quality LFW: A database for analyzing cross-resolution image face recognition in unconstrained environments. *CoRR*, abs/2108.10290, 2021.
- [17] H.H. Ku. Notes on the use of propagation of error formulas. *Journal of Research of the National Bureau of Standards. Section C: Engineering and Instrumentation*, 70C, 1966.
- [18] Shen Li, Jianqing Xu, Xiaqing Xu, Pengcheng Shen, Shaoxin Li, and Bryan Hooi. Spherical confidence learning for face recognition. In *CVPR*, pages 15629–15637. Computer Vision Foundation / IEEE, 2021.
- [19] Wanhua Li, Xiaoke Huang, Jiwen Lu, Jianjiang Feng, and Jie Zhou. Learning probabilistic ordinal embeddings for uncertainty-aware regression. In *CVPR*, pages 13896–13905. Computer Vision Foundation / IEEE, 2021.
- [20] Rushuai Liu and Weijun Tan. Eqface: A simple explicit quality network for face recognition. In *CVPR Workshops*, pages 1482–1490. Computer Vision Foundation / IEEE, 2021.
- [21] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *CVPR*, pages 14225–14234. Computer Vision Foundation / IEEE, 2021.
- [22] Fu-Zhao Ou, Xingyu Chen, Ruixin Zhang, Yuge Huang, Shaoxin Li, Jilin Li, Yong Li, Liujuan Cao, and Yuan-Gen Wang. SDD-FIQA: unsupervised face image quality assessment with similarity distribution distance. In *CVPR*, pages 7670–7679. Computer Vision Foundation / IEEE, 2021.
- [23] Dane K Peterson and Gordon F Pitz. Confidence, uncertainty, and the use of information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(1): 85, 1988.
- [24] P. Jonathon Phillips and Alice J. O’Toole. Comparison of human and computer performance across face recognition experiments. *Image Vis. Comput.*, 32(1):74–85, 2014.
- [25] Torsten Schlett, Christian Rathgeb, Olaf Henniger, Javier Galbally, Julian Fierrez, and Christoph Busch. Face image quality assessment: A literature survey. *CoRR*, abs/2009.01103, 2020.
- [26] Yichun Shi and Anil K. Jain. Probabilistic face embeddings. In *ICCV*. IEEE, 2019.

- [27] Philipp Terhörst, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. SER-FIQ: unsupervised estimation of face image quality based on stochastic embedding robustness. In *CVPR*, pages 5650–5659. Computer Vision Foundation / IEEE, 2020.
- [28] Philipp Terhörst, Malte Ihlefeld, Marco Huber, Naser Damer, Florian Kirchbuchner, Kiran B. Raja, and Arjan Kuijper. Qmagface: Simple and accurate quality-aware face recognition. *CoRR*, abs/2111.13475, 2021.
- [29] Philipp Terhörst, Jan Niklas Kolf, Marco Huber, Florian Kirchbuchner, Naser Damer, Aythami Morales Moreno, Julian Fierrez, and Arjan Kuijper. A comprehensive study on face recognition biases beyond demographics. *IEEE Transactions on Technology and Society*, 3(1):16–30, 2021.
- [30] Mei Wang and Weihong Deng. Deep face recognition: A survey. *Neurocomputing*, 429:215–244, 2021.
- [31] Weidi Xie, Jeffrey Byrne, and Andrew Zisserman. Inducing predictive uncertainty estimation for face verification. In *BMVC*. BMVA Press, 2020.
- [32] Nick Yeung and Christopher Summerfield. Metacognition in human decision-making: confidence and error monitoring. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594):1310–1321, 2012.
- [33] Kecheng Zheng, Cuiling Lan, Wenjun Zeng, Zhizheng Zhang, and Zheng-Jun Zha. Exploiting sample uncertainty for domain adaptive person re-identification. In *AAAI*, pages 3538–3546. AAAI Press, 2021.