

# WebFace260M: A Benchmark for Million-Scale Deep Face Recognition

Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang,  
Xinze Chen, Jiagang Zhu, Tian Yang, Dalong Du, Jiwen Lu, Jie Zhou

**Abstract**—Face benchmarks empower the research community to train and evaluate high-performance face recognition systems. In this paper, we contribute a new million-scale recognition benchmark, containing uncurated 4M identities/260M faces (WebFace260M) and cleaned 2M identities/42M faces (WebFace42M) training data, as well as an elaborately designed time-constrained evaluation protocol. Firstly, we collect 4M name lists and download 260M faces from the Internet. Then, a Cleaning Automatically utilizing Self-Training (CAST) pipeline is devised to purify the tremendous WebFace260M, which is efficient and scalable. To the best of our knowledge, the cleaned WebFace42M is the largest public face recognition training set and we expect to close the data gap between academia and industry. Referring to practical deployments, Face Recognition Under Inference Time conStraint (FRUITS) protocol and a new test set with rich attributes are constructed. Besides, we gather a large-scale masked face sub-set for biometrics assessment under COVID-19. For a comprehensive evaluation of face matchers, three recognition tasks are performed under standard, masked and unbiased settings, respectively. Equipped with this benchmark, we delve into million-scale face recognition problems. A distributed framework is developed to train face recognition models efficiently without tampering with the performance. Enabled by WebFace42M, we reduce 40% failure rate on the challenging IJB-C set and rank 3rd among 430 entries on NIST-FRVT. Even 10% data (WebFace4M) shows superior performance compared with the public training sets. Furthermore, comprehensive baselines are established under the FRUITS-100/500/1000 milliseconds protocols. The proposed benchmark shows enormous potential on standard, masked and unbiased face recognition scenarios. Our WebFace260M website is <https://www.face-benchmark.org>.

**Index Terms**—Large-scale Face Recognition, Masked Face Recognition, Unbiased Face Recognition, Biometric Authentication

## 1 INTRODUCTION

RECOGNIZING faces in the wild has achieved remarkable success due to the boom of neural networks. The key engine of recent face recognition consists of network architecture evolution [37], [38], [47], [70], [74], [77], [98], a variety of loss functions [21], [23], [42], [45], [46], [53], [66], [73], [75], [78], [85], [86], [89], [96], and growing face benchmarks [11], [14], [35], [41], [44], [56], [57], [58], [63], [67], [83], [105], [113], [114]. Even though growing efforts have been devoted to investigating sophisticated networks and losses, academia is restricted by limited training sets and nearly saturated test protocols.

As shown in Table 1 and Figure 1, the largest public training sets in terms of identities and faces are MegaFace2 [58] and MS1M [35], respectively. MegaFace2 contains 4.7M faces of 672K subjects collected from Flickr [82]. MS1M consists of 10M faces of 100K celebrities but the noise rate is around 50% [83]. In contrast, companies from industry can access much larger private data to train face recognition models: Google utilizes 200M images of 8M identities to train FaceNet [66], and Facebook [79] performs training by 500M faces of 10M identities. This data gap hinders researchers

from pushing the frontiers of deep face recognition. The main obstacles for tremendous training data lie in large-scale identity collection, effective and scalable cleaning, and efficient training. For example, IMDB-Face [83] takes 50 annotators to work continuously for one month to obtain 59K identities and 1.7M images, which is labor-intensive and non-scalable.

On the other hand, test sets and protocols play an essential role in analyzing face recognition performance. Popular evaluations including LFW families [41], [113], [114], CFP [67], AgeDB [57], RFW [91], MegaFace [44], and IJB-C [56] mainly target the pursuit of accuracy, which have been almost saturated recently. In real-world application scenarios, face recognition is always restricted by inference time, such as unlocking mobile telephones with a smooth user experience. Lightweight Face Recognition (LFR) Challenge [22] takes a step toward this goal by constraining model size and FLOPs, but actual inference time can vary quite a bit for different networks. Besides, LFR neglects the time cost of face detection and alignment. To the best of our knowledge, NIST-FRVT [2] is the only time-constrained face recognition protocol. However, the strict submission policy (no more than one submission every four calendar months) hinders researchers from freely evaluating their algorithms.

To address the above problems, this paper constructs a new ultra-large-scale face benchmark consisting of 4M identities/260M faces (WebFace260M) as well as a time-constrained assessment protocol. Firstly, a name list of 4M celebrities is gathered and 260M images are downloaded utilizing a search engine. Then, we perform Cleaning Automatically by Self-Training (CAST) pipeline, which is scalable

- Zheng Zhu, Jiwen Lu, and Jie Zhou are with the Beijing National Research Center for Information Science and Technology (BNRist), and the Department of Automation, Tsinghua University, Beijing, 100084, China. E-mail: zhengzhu@tsinghua.edu.cn; lujiwen@tsinghua.edu.cn; jzhou@tsinghua.edu.cn
- Guan Huang, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, and Dalong Du are with XForwardAI, Beijing, China.
- Jiankang Deng is with Imperial College London, London, UK.
- Corresponding author: Jiwen Lu.

TABLE 1

Training sets for deep face recognition. The curated WebFace42M is the largest public training data in terms of both # identities and # images.

Dataset	# Identities	# Images	Images/ID	Cleaning	# Attributes	Availability	Publications
CASIA-WebFace [105]	10 K	0.5 M	47	Automatic	-	Public	Arxiv 2014
CelebFaces [75]	10 K	0.2 M	20	Manual	40	Public	ICCV 2015
VGGFace [63]	2 K	2.6 M	1,000	Semi-automatic	-	Public	BMVC 2015
MS1M [35]	0.1 M	10 M	100	No	-	Public	ECCV 2016
MegaFace2 [58]	0.6 M	4.7 M	7	Automatic	-	Public	CVPR 2017
MS1M-IBUG [23]	85 K	3.8 M	45	Semi-automatic	-	Public	CVPRW 2017
UMDFaces [11]	8 K	0.3 M	45	Semi-automatic	4	Public	IJCB 2017
IMDB-Face [83]	59 K	1.7 M	29	Manual	-	Public	ECCV 2018
VGGFace2 [14]	9 K	3.3 M	363	Semi-automatic	11	Public	FG 2018
MS1MV2 [21]	85 K	5.8 M	68	Semi-automatic	-	Public	CVPR 2019
MS1M-Glint [1]	87 K	3.9 M	44	Semi-automatic	-	Public	-
Facebook [78]	4 K	4.4 M	1,100	-	-	Private	CVPR 2014
Facebook [79]	10 M	500 M	50	-	-	Private	CVPR 2015
Google [66]	8 M	200 M	25	-	-	Private	CVPR 2015
MillionCelebs [109]	0.6 M	18.8 M	30	Automatic	-	Private	CVPR 2020
<b>WebFace260M</b>	<b>4 M</b>	<b>260M</b>	<b>65</b>	No	-	Public	-
<b>WebFace42M</b>	<b>2 M</b>	<b>42M</b>	<b>21</b>	Automatic	7	Public	-

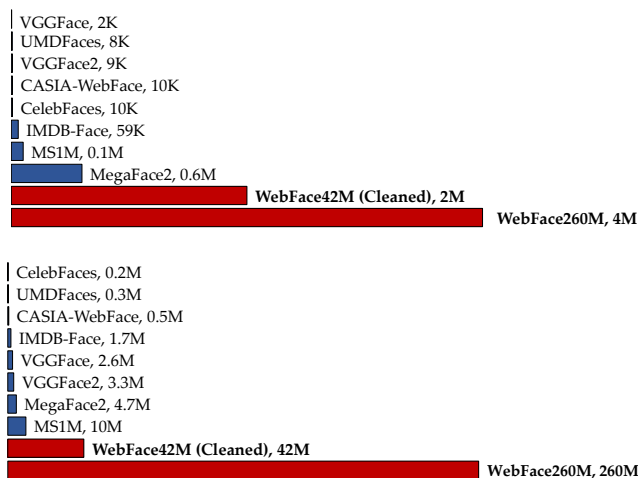


Fig. 1. Comparisons of # identities and # faces for our WebFace data and public training sets.

and does not need any human intervention. The proposed CAST procedure results in high-quality 2M identities and 42M faces (WebFace42M). Meanwhile, rich attributes are provided for further analyzing the statistics of WebFace42M. Referring to various real-world applications, we design the Face Recognition Under Inference Time constraint (FRUITS) protocol, enabling academia to test deep face matchers comprehensively. Specifically, the FRUITS includes three time limit tracks: 100, 500, and 1000 milliseconds, which intend to evaluate deployments on mobile devices, local devices, and clouds, respectively. Since public evaluations are almost saturated [41], [57], [67] and may contain noise [44], [56], we manually construct a new test set with rich attributes to enable FRUITS. Considering the COVID-19 coronavirus epidemic [28], [48] and reported biased face recognition deployments [31], [90], [91], three evaluation tasks are performed: Standard Face Recognition (SFR), Masked Face Recognition (MFR), and Unbiased Face Recognition (UFR). For MFR, we collect a large-scale masked face test sub-set.

Based on the proposed ultra-large-scale benchmark, we delve into million-scale deep face recognition problems.

With such data size, a distributed training framework is developed for efficient optimization, which could perform at a nearly linear acceleration without performance drops. Accuracy on the public and proposed test sets indicates that our training data is indispensable for pushing the frontiers of deep face recognition: WebFace42M achieves 97.70% TAR@FAR=1e-4 on the challenging IJB-C [56] under standard ResNet-100 configurations, reducing near 40% relative error rate compared with state of the art. 10% of our data (WebFace4M) also obtains superior performance than similar-sized MS1M families [1], [21], [23] and MegaFace2 [58]. On the proposed test set, a similar conclusion can be drawn. For SFR, WebFace42M decreases FNMR@FMR=1e-5 from 9.88% (with MS1MV2) to 2.98% under the same settings, which reduces the error rate by 3 times. For MFR, error rates (FNMR@FMR=1e-5) of WebFace42M and MS1MV2 are 42.97% and 69.56%, respectively. For UFR, considering the tremendous scale of the WebFace data, it provides more room for data balancing and fairness exploration. Furthermore, we participate in the NIST-FRVT [2] and rank 3rd among 430 entries based on WebFace42M. At last, we discuss privacy and bias issues in this benchmark.

For baseline comparisons, comprehensive deep face recognition systems are evaluated under FRUITS-100/500/1000 milliseconds protocols. For SFR, different settings of face detection/alignment and feature extraction are explored, covering MobileNet [15], [38], EfficientNet [80], AttentionNet [84], ResNet [37], SENet [39], ResNeXt [100] and RegNet families [64]. For MFR, the influence of mask augmentation in WebFace training data is studied and a strong baseline is established for this difficult problem. In order to investigate the bias in current face recognition systems, we re-sample the WebFace data to obtain an attribute-balanced sub-set and evaluate its influence on recognition fairness. With this new face benchmark, we hope to close the data gap between academia and industry, and facilitate the time-constrained recognition assessment for real-world applications.

The main contributions of this work can be summarized as follows:

- An ultra-large-scale face recognition dataset is constructed for the research community towards closing

the data gap behind the industry. The proposed WebFace260M consists of 4M identities and 260M faces, which provides an excellent resource for million-class deep face cleaning and recognition as shown in Table 1 and Figure 1.

- We contribute the training set WebFace42M which sets new state of the art on challenging IJB-C and ranks 3rd on NIST-FRVT. This cleaned data is automatically purified from WebFace260M by a scalable and effective self-training pipeline. To the best of our knowledge, the resulting WebFace42M is the largest public face recognition training resource.
- The FRUITS protocol as well as a test set with rich attributes are constructed to facilitate the evaluation of real-world applications. Meanwhile, we collect a large-scale masked face set. A series of time-constrained tasks are designed referring to different deployment scenarios, including standard, masked, and unbiased face recognition.
- Based on the new benchmark, we perform extensive million-scale face recognition experiments. Enabled by distributed training framework, comprehensive baselines are established on public and our test sets under the FRUITS protocol. The results indicate substantial insights on three different recognition settings.

This paper is built upon our conference work [118] and significantly extended in several aspects. Firstly, we provide comprehensive visualization results to illustrate the face benchmark. This gives a deeper insight into the diversity and challenges of our training/test data. Secondly, the scale of our test set is significantly increased, making it more challenging for recognition evaluation. Besides, we collect a highly-curated masked sub-set, which contains 862 subjects with real-world masks. Lastly, for evaluation and experiment, standard face recognition is extended to masked and unbiased settings. Corresponding mask augmentation and attribute-balanced baselines are also established. In addition, we present a detailed literature review for deep face recognition and benchmarks. Privacy and bias issues in WebFace260M benchmark are also discussed.

Since the preliminary version of this work was published, we have received dataset access applications<sup>1</sup> from near 400 research groups. Based on the WebFace260M benchmark, we organize the Face Bio-metrics under COVID Workshop and Masked Face Recognition Challenge [18], [117] in ICCV 2021. More than 80 teams have participated in WebFace260M Track under FRUITS protocol and submitted more than 1,000 solutions<sup>2</sup>. In InsightFace Unconstrained Track, all top-3 teams from academia and industry adopt WebFace260M database as their training sets. These results suggest that our WebFace260M is not only an effective benchmark to pursue high-performance face recognition systems but also a meaningful step to reduce the data gap between research laboratories and companies.

## 2 RELATED WORK

Face recognition has been extensively studied in the computer vision literature. Recent years have witnessed a

significant advance in both benchmarks and algorithms, including growing training data, evolutionary evaluation sets, and loss function designs. This section reviews representative progresses in academia and industry.

### 2.1 Face Recognition Training Data

A key aspect in developing face recognition systems is the training data used to learn discriminative face representations. Data collections are extremely important but usually overlooked. Even though some companies have internally labeled private face sets that scale to millions of images [78] or even millions of subjects (Google [66] and Facebook [79]), the situation is quite different for publicly available collections. As shown in Table 1, we give the detailed statistics of widely used training sets in the community, such as CASIA-WebFace [105], VGGFace2 [14], UMDFaces [11], MS1M [35], MegaFace2 [58], and IMDB-Face [83].

CASIA-WebFace [105], VGGFace2 [14] and UMDFaces [11] consist of around 10K identities. CASIA-WebFace [105] is collected by a semi-automated method, which searches face images of celebrities from the Internet. VGGFace2 [14] is an improved version of VGGFace [63] created in order to mitigate the deficiency of its predecessor. The subjects in VGGFace2 are collected from celebrities and famous people such as professors and politicians. Compared to its predecessor, VGGFace2 contains fewer images for each identity but covers a large range of poses, ages and races. To reduce label noise as much as possible, manual and automatic processes are employed. UMDFaces [11] utilizes a mix of human annotators via Amazon Mechanical Turk (AMT) and pre-trained deep-based face analysis tools to build a face dataset that is much tougher than already available sets.

MS1M [35], MegaFace2 [58], and IMDB-Face [83] include more identities than above-mentioned datasets. MS1M [35] retrieves around 100 images for each identity by the Bing search engine [3] using the celebrity's name without any filtering. Therefore, the quality of MS1M is severely biased by label noises, duplicated images, and non-face images present in the set. All of these factors make MS1M hard to be used directly. MegaFace2 [58] contains 672K subjects cleaned from Flickr. However, this dataset only collects 4.7M faces, which results in around 7 images per identity. IMDB-Face [83] claims to be the largest noise-controlled face collection, which contains 1.7M images of 59K celebrities by manual annotation. However, it took 50 annotators to work continuously for one month to clean the dataset, which demonstrates the difficulty of obtaining a large-scale clean dataset for face recognition.

### 2.2 Face Recognition Evaluation

Most popular evaluation sets for face recognition target the pursuit of accuracy. CFP [67], AgeDB [57], CALFW [114] and CPLFW [113] evaluate the verification accuracy under different intra-class variations (such as pose and age). MegaFace [44] and IJB-C [56] serve for both accuracies of large-scale face verification and identification. YTF [97] and IQIYI-Video [54] compare the accuracy of video-based verification. Different model-ensemble and post-processing [68] could be adopted for higher performance under these

1. <https://www.face-benchmark.org/download.html>

2. <https://competitions.codalab.org/competitions/32478#results>

protocols. However, face recognition in real-world application scenarios is always restricted by inference time, such as unlocking mobile phones with a smooth experience or processing multiple channels of surveillance videos on clouds.

Recently, the LFR Challenge [22] takes a step toward this goal by constraining the FLOPs and model size of submissions. Since different neural network architectures can be quite different in terms of real inference time, this protocol is not a straightforward solution. Furthermore, it does not consider face detection and alignment, which are prerequisite components in most modern face recognition systems. To the best of our knowledge, NIST-FRVT [2] is the only benchmark employing the time-constrained protocol. However, the strict submission policy (participants can only send one submission every four calendar months) hinders researchers from freely evaluating their algorithms.

With global COVID-19 pandemic and reported biased systems deployment, up-to-date face recognition datasets focus on comparisons with masks and fairness consideration. RFW [91] aims to evaluate the bias among 4 racial distributions. NIST-FRVT [34] regularly evaluates the bias level of submitted algorithms caused by demographic effects. Due to the sudden outbreak of the coronavirus epidemic, there is yet no comprehensive real-world masked face recognition benchmark. Evaluations on simulated masked images [59], [61] may result in questioned conclusions, and small-scale masked sets [10], [12], [17] can not comprehensively reflect the performance of algorithms. Real-world masked test set RMFRD [95] consists of 525 identities and 5K masked faces, but there exist annotation noises.

### 2.3 Deep Face Recognition

The last decade has witnessed the advance of deep convolutional face recognition techniques. A number of successful face recognition systems, such as DeepFace [78], DeepID [73], [74], [75], [76], FaceNet [66] have achieved impressive performance on face verification and identification. Most of the early works rely on metric-learning based losses [16], [66], [71], and recent researches have switched to margin-based softmax losses due to their efficiency on the large-scale dataset. SphereFace [53], AM-softmax [85], CosFace [89], ArcFace [21] progressively improve the performance on various benchmarks to the newer level. To further improve the margin-based softmax loss, recent works focus on the exploration of adaptive parameters [51], [52], [107], [108], inter-class regularization [27], [112], sample mining [42], [94], learning acceleration [9], [46], [49], [50], [106], etc. There are also many complementary methods proposed to build better face recognition models by promoting desired properties of the produced face representations, such as robustness to noisy labels [19], [40], [93], [109], [116], occlusions [72], [92], [115] and low quality [68], [69], invariance to age [43], [88], [111] and pose [87], [110], ability to mitigate racial bias [32], [90], [91] and domain imbalance [13], [25], [45], [49], to improve the fairness of representations [51], [101].

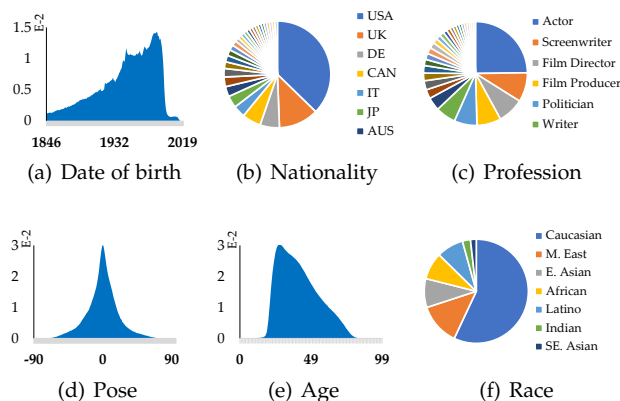


Fig. 2. Birth date, nationality, profession distributions of WebFace260M, and pose (yaw), age, race distributions of WebFace42M.

## 3 WEBFACE260M AND WEBFACE42M

### 3.1 Data Collection and Pre-processing

Knowledge graphs website Freebase [4] and well-curated website IMDB [5] provide excellent resources for collecting celebrity names. Furthermore, commercial search engines such as Google [6] and Bing [3] make it possible to collect images of a specific identity with ranked correlation. Our celebrity name list consists of two parts: the first one is borrowed from MS1M (1M, constructed from the Freebase) and the second one is collected from the IMDB database. There are near 4M celebrity names on the IMDB website, while we find some subjects have no public image from search engines. Therefore, only 3M celebrity names in IMDB are chosen for our benchmark. Based on the name list, celebrity faces are searched and downloaded via Google image search engine [6]. 200 images per identity are downloaded for the top 10% subjects, while 100, 50, 25 images are reserved for the remaining 20%, 30%, 40% subjects, respectively. Finally, we collect 4M identities and 265M images.

In data pre-processing, faces are detected and aligned through five landmarks predicted by RetinaFace [20]. Specifically, the threshold of detection score is set as 0.7 to filter the low confident faces. After pre-processing, there are 4M identities/260M faces (WebFace260M) shown as Table 1. The statistics of WebFace260M are illustrated in Figure 2 including date of birth, nationality and profession. Persons in WebFace260M come from more than 200 distinct countries/regions and more than 500 different professions with the date of birth back to 1846, which guarantees a great diversity in our training data. *During the construction of WebFace260M dataset, privacy and bias problems are our first concerns. Detailed discussion is available in Section 7.*

### 3.2 Cleaned WebFace42M

We perform a CAST pipeline (Section 4) to automatically clean the noisy WebFace260M and obtain a curated training set named WebFace42M, consisting of 42M faces of 2M subjects. Face number in each identity varies from 3 to more than 300, and the average face number is 21 per identity. As shown in Table 1 and Figure 1, WebFace42M offers the largest cleaned training data for face recognition. Compared

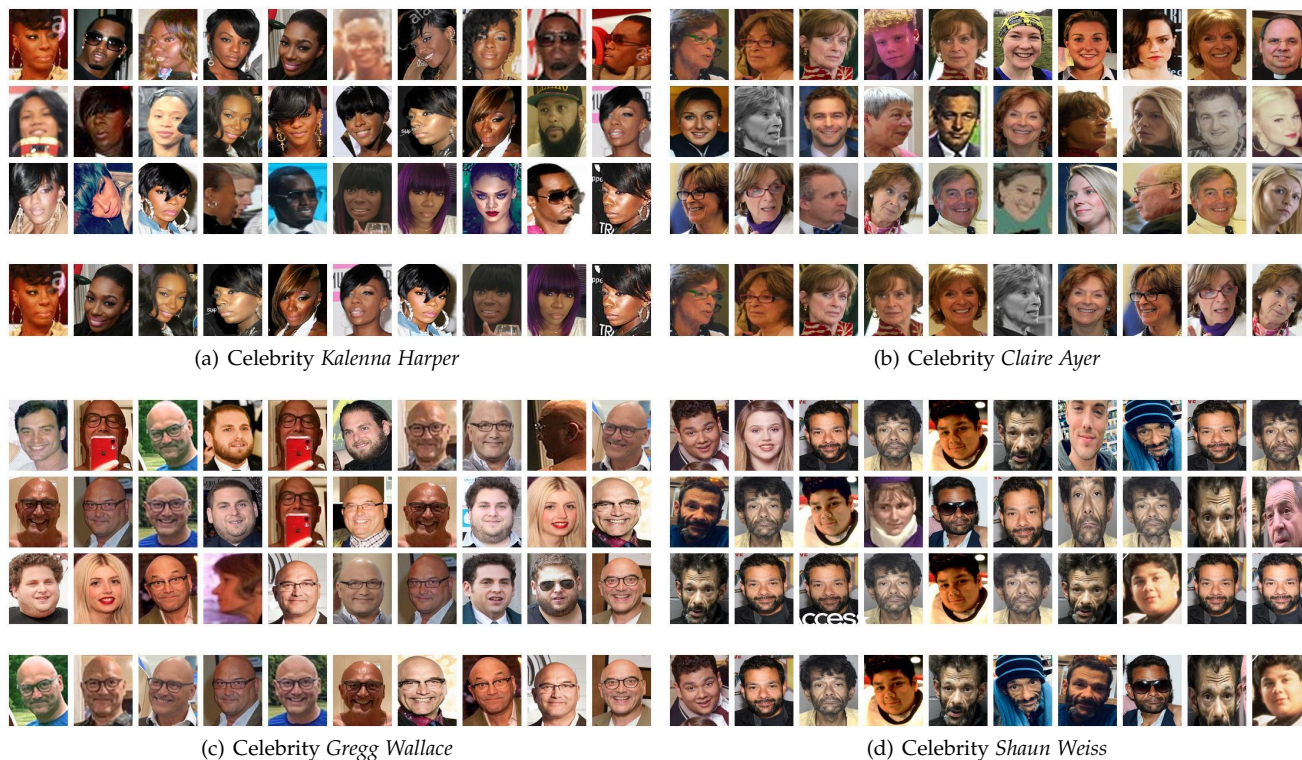


Fig. 3. Visualization of the WebFace data. For each sub-figure, the top part is the randomly selected faces from WebFace260M, while the bottom part shows cleaned faces (also randomly selected) from WebFace42M. Loose cropped faces are shown.

with the MegaFace2 [58] dataset, the proposed WebFace42M includes 3 times more identities (2M vs. 672K), and near 10 times more images (42M vs. 4.7M). Compared with the widely used MS1M [35], our training set is 20 times (2M vs. 100K) and 4 times (42M vs. 10M) more in terms of # identities and # photos. According to [83], there are more than 30% and 50% noises in MegaFace2 and MS1M, while the noise ratio of WebFace42M is lower than 10% (similar to CASIA-WebFace [105]) based on our sampling estimation. With such a large data size, we take a significant step towards closing the data gap between academia and industry.

We further provide face attribute statistics for WebFace42M. Figure 2 presents the distribution of our cleaned training data in different aspects. WebFace42M covers a large range of poses (Figure 2(d)), ages (Figure 2(e)) and most major races in the world (Figure 2(f)).

### 3.3 Visualization

Figure 3 illustrates four random celebrities from our WebFace data. The original WebFace260M folders downloaded/detected from Internet images are very noisy, containing various wrong detections, unrelated persons, name repetitions, etc. The proposed CAST in the next section can automatically purify each folder to obtain cleaned faces for a certain identity. Specifically, there are large hairstyle and pose variations for celebrities *Kalenna Harper* in WebFace42M. Faces of *Claire Ayer* and *Gregg Wallace* show different expressions and yaw angles. Finally, the proposed CAST cleaning pipeline covers a broad age range for celebrities *Shaun Weiss*. The great diversity of WebFace42M guarantees its quality for training high-performance face recognition models.

## 4 CLEANING AUTOMATICALLY BY SELF-TRAINING

Since the images downloaded from the web are considerably noisy, it is necessary to perform a cleaning step to obtain high-quality training data. Original MS1M [35] does not perform any dataset cleaning, resulting in a near 50% noise ratio, and significantly degrades the performance of the trained models. VGGFace [63], VGGFace2 [14] and IMDB-Face [83] adopt semi-automatic or manual cleaning pipelines, which require expensive labor efforts. It becomes difficult to scale up the current annotation size to even more identities. Although the purification in MegaFace2 [58] is automatic, its procedure is complicated and there are considerably more than 30% noises [83]. Another relevant exploration is to cluster faces via unsupervised approaches [55], [62] or supervised graph-based algorithms [103], [104]. However, these methods assume the whole dataset is curated, which is not suitable for the extremely noisy WebFace260M.

### 4.1 CAST Framework

Recently, self-training [99], [102], a standard approach in semi-supervised learning, is explored to significantly boost the performance of image classification. Different from close-set ImageNet classification [65], directly generating pseudo labels on open-set face recognition is impractical. Considering this inherent limitation, we carefully design the pipeline of Cleaning Automatically by Self-Training (CAST). Our first insight is performing self-training on open-set face recognition data, which is a scalable and efficient cleaning approach. Secondly, we find embedding feature matters in cleaning ultra-large-scale noisy face data.

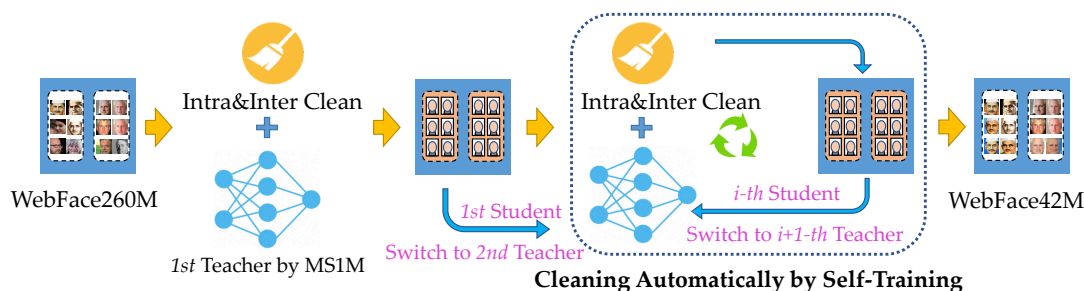


Fig. 4. The proposed Cleaning Automatically by Self-Training (CAST). Firstly, an initial Teacher trained with MS1MV2 is utilized to clean WebFace260M. Then a Student model is trained on the cleaned WebFace data. The CAST is performed by switching the Student as the Teacher until high-quality 42M faces are obtained. Every intra-class/inter-class cleaning is conducted on the initial WebFace260M utilizing different Teacher models.

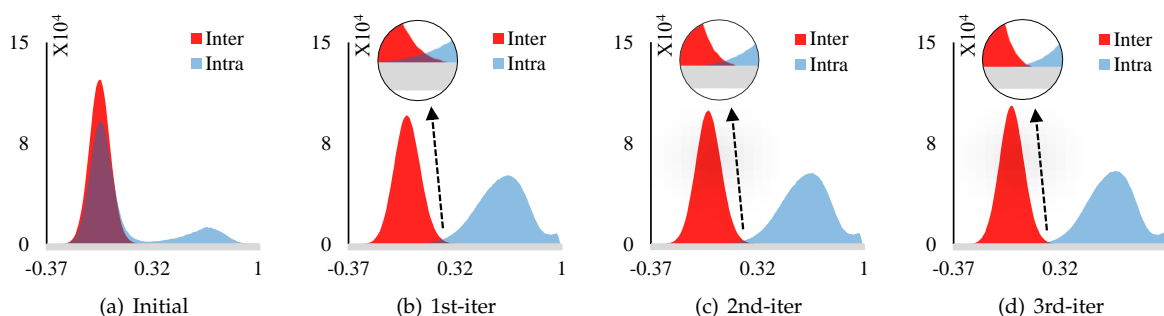


Fig. 5. Inter and intra class similarity distributions during different stages of CAST. Since initial folders are very noisy, score distributions are severely overlapped. Cleaner training set is obtained after more iterations. 100K folders are randomly selected here for showing the statistic changes during iterations.

The overall CAST framework is shown in Figure 4. Following the self-training pipeline, (1) A Teacher model is trained with the public dataset (MS1MV2 [21]) to clean the original 260M images, which mainly consists of intra-class and inter-class cleaning. (2) A Student model is trained on cleaned images from (1). Since the data size is much larger, this Student generalizes better than the Teacher. (3) We iterate this process by switching the Student as the Teacher until high-quality 42M faces are obtained. It is worth noting that each intra-class/inter-class cleaning is conducted on the initial WebFace260M by different Teacher models. All Teacher/Student models adopt ResNet-100 backbone and ArcFace loss function, and other configurations are equivalent to those in WebFace42M training setting (Section 6.1). The CAST pipeline is summarized in Algorithm 1.

## 4.2 Intra-class and Inter-class Cleaning

Since WebFace260M contains various noises such as outliers in a folder and identity overlaps between folders, it is impractical to perform unsupervised or supervised clustering on the whole dataset. Based on the observation that the image search results from Google are sorted by relevance and there is always a dominant subject in each search, the initial folder structure provides strong priors to guide the cleaning strategy: one folder always contains a dominant subject and different folders may contain considerable overlapped identities.

Following these priors, we perform dataset cleaning by a two-step procedure: Firstly, face clustering is parallelly conducted in 4M folders (subjects) to select each dominant identity. Specifically, for each face in a folder, 512-dimensional

embedding feature is extracted by the Teacher model, and then DBSCAN [29] is utilized to cluster faces in this folder. Only the largest cluster (more than 2 faces) in each fold is reserved.  $\epsilon$  and  $n$  of DBSCAN indicate the maximum distance for the radius of a neighborhood, and the minimum number of points required within this distance, respectively. We use  $(1 - \epsilon)$  to denote the similarity of face embeddings in our paper. With more iterations of CAST, the model learns stronger face embeddings. A higher similarity  $(1 - \epsilon)$  has a trend to filter out more number of noisy faces, which is beneficial for creating cleaner datasets. So we empirically set larger values for similarity  $(1 - \epsilon)$  in later iterations of CAST and keep  $n$  fixed. We also investigate other different designs of CAST in Section 6.4. Secondly, we compute the feature center of each subject to perform inter-class cleaning. Two folders are merged if their cosine similarity is higher than 0.7, and the folder containing fewer faces would be deleted when the cosine similarity is between 0.5 and 0.7. As shown in Algorithm 1, lines 2-5 and lines 6-13 are intra-class and inter-class cleaning processes, respectively.

The effectiveness of the above intra-class and inter-class cleaning heavily depends on the quality of the embedding feature, which is guaranteed by the proposed self-training pipeline. The ArcFace model trained on MS1MV2 with ResNet-100 provides a good initial embedding feature to perform first-round cleaning for WebFace260M. Then, this feature is significantly enhanced with more training data in later iterations. Figure 5 illustrates the score distribution during different stages of CAST, which indicates a cleaner training set after more iterations. Furthermore, the ablation study in Table 7 also validates the effectiveness of the CAST

**Algorithm 1** Cleaning Automatically by Self-Training

```

1st Model: The ArcFace model trained on MS1MV2
Input Data: (Noisy) WebFace260M
Output Data: (Cleaned) WebFace42M
1: for  $i = 1$  to  $M$  do
2:   for Each folder in WebFace260M do
3:     Utilizing  $i$ -th Teacher model to extract 512- $d$  feature
       of each face
4:     Performing DBSCAN and reserving the largest
       cluster to obtain intra-class cleaning results
5:   end for
6:   for Every two subjects do
7:     Computing the cosine similarity of feature center
8:     if cosine similarity  $> 0.7$  then
9:       Merging two folders
10:    else if cosine similarity  $> 0.5$  then
11:      Deleting the folder with fewer faces
12:    end if
13:  end for
14:  Training  $i$ -th Student model on  $i$ -th cleaned WebFace
   (not for last iteration)
15:  Converting  $i$ -th Student into  $i+1$ -th Teacher model
   (not for last iteration)
16: end for
17: return Cleaned WebFace data

```

TABLE 2

The identities and images statistics during different cleaning stages.

Stages		# Identities	# Faces
Collect name list and images		4,073,509	265,777,598
Face pre-processing		<b>4,008,130</b>	<b>260,890,076</b>
First iteration	Intra-class	3,341,761	61,792,387
	Inter-class	2,437,140	50,672,354
Second iteration	Intra-class	3,027,814	60,274,892
	Inter-class	2,176,427	47,352,741
Third iteration	Intra-class	2,878,886	58,155,345
	Inter-class	2,070,870	46,220,417
Remove duplicates		2,070,870	43,977,802
Remove test set overlaps		<b>2,059,906</b>	<b>42,474,558</b>

**4.3 Statistics**

The statistics of identities and faces during different cleaning stages are shown in Table 2. After face pre-processing for downloaded images, there are 4,008,130 identities and 260,890,076 faces (WebFace260M). The face set becomes cleaner under more CAST iterations, which results in fewer data. Finally, we obtain 2,059,906 identities and 42,474,558 faces (WebFace42M) after removing duplicates and test set overlaps.

**5 FACE RECOGNITION EVALUATION**

In this section, we firstly introduce the time-constrained face recognition evaluation protocol, which covers various practical applications. Then, three benchmarking tasks (standard, masked, unbiased face recognition) are detailed, including corresponding background, test sets, and metrics.

**5.1 FRUITS Protocol**

As discussed in Section 2.2, most existing face recognition evaluations [41], [44], [56], [67], [97] only focus on accuracy. LFR Challenge [22] adopts FLOPs and model size constraints, which may result in unfair inference time comparisons. Moreover, it ignores the computation module of face detection and alignment. Four months submission interval of NIST-FRVT [2] makes it difficult for the research community to freely perform experiments. In this paper, we design the Face Recognition Under Inference Time constraint (FRUITS) protocol, which enables academia to comprehensively evaluate their face matchers. Referring to [2], inference time is measured on a single core of an Intel Xeon CPU E5-2630-v4@2.20GHz processor. For face recognition evaluation with flip, the batch size is set as 2. Considering different application scenarios, the FRUITS protocol sets a series of tracks:

**FRUITS-100:** The whole face recognition system must distinguish image pairs within 100 milliseconds, including pre-processing (face detection and alignment), feature embedding for recognition, and matching. FRUITS-100 track targets on evaluating lightweight face recognition systems which can be deployed on mobile devices.

**FRUITS-500:** This track follows FRUITS-100 setting, except that the time constraint is increased to 500 milliseconds. This track aims to evaluate modern and popular networks deployed in the local surveillance system.

**FRUITS-1000:** Following NIST-FRVT, FRUITS-1000 constrains the inference time in 1000 milliseconds and aims




Fig. 6. Illustration of the reserved and the rejected face samples. Green and red boxes denote reserved faces and faces rejected by our CAST, respectively.

pipeline. It is worth noting that the proposed CAST pipeline is compatible with any intra-class and inter-class strategies.

**Remove duplicates and test set overlaps.** After CAST, duplicated faces are removed when their cosine similarity is higher than 0.95. Furthermore, the feature center of each subject is compared with popular benchmarks (such as LFW families [41], [113], [114], FaceScrub [60], IJB-C [56], the proposed test set, etc.), and overlaps are removed if the cosine similarity is higher than 0.7.

Figure 6 illustrates the reserved and the rejected face samples in our cleaning process. One can find that noisy and mislabeled faces are successfully rejected by the proposed CAST strategy, and most true positives are reserved in the cleaned WebFace42M set. Significantly, there are diverse expressions and poses among the remaining faces, which clearly show the effectiveness of our CAST and the high quality of our training data.




Fig. 7. Visualization of our test set. Random faces of a certain identity are shown in each row. 1-6 columns: Controlled, 7-12 columns: Wild, 13-18 columns: Masked.

to compare capable recognition models performed on clouds. In our organized Face Bio-metrics under COVID Workshop, FRUITS-1000 is adopted for both masked and standard face recognition challenges.

## 5.2 Standard Face Recognition

### 5.2.1 Test set

Since public evaluations are most saturated and may contain noises, we manually construct an elaborated test set for SFR, MFR and UFR. It is well known that recognizing strangers, especially when they are similar-looking, is a difficult task even for experienced vision researchers. Therefore, our multi-ethnic annotators only select their familiar celebrities, which ensures the high quality of the test data. Besides, annotators are encouraged to gather attribute-balanced faces, and **recognition models are introduced to guide hard sample collection**. The statistics of the final test set are listed in Table 3. In total, there are 60,926 faces of 2,478 identities. Rich attributes (such as age, race, gender, scenario) are accurately annotated. Among all collected data, 57,715 faces are utilized for SFR evaluation. In the future, we will actively maintain and update this test set. In Figure 7, we show the samples of our SFR test set, including *Controlled* and *Wild* faces. One can find that there is great diversity in ages, poses, scenarios, etc.

### 5.2.2 Metrics

Based on the proposed FRUITS protocol and test set, we perform standard 1:1 face verification across various attributes. Table 3 shows numbers of imposter and genuine in different verification settings. *All* means impostors are paired without attention to any attribute, while later comparisons are conducted on age and scenario sub-sets. *Cross-age* refers to cross-age (more than 10 or 20 years span) verification, while *Cross-scene* means pairs are compared between *Controlled* and *Wild* settings. Different algorithms are measured on False Non-Match Rate (FNMR) [2], which is defined as the proportion of mated comparisons below a threshold set to achieve the False Match Rate (FMR) specified. FMR is the proportion of imposter comparisons at or above that threshold. **Lower FNMR at the same FMR is better.**

## 5.3 Masked Face Recognition

### 5.3.1 Face Bio-metrics under COVID

SFR systems usually work with mostly non-occluded faces, which include primary facial features such as eyes, nose, and mouth. However, there are a number of circumstances in which faces are occluded by masks such as in pandemics, medical settings, excessive pollution, or laboratories. According to WHO statistics, there are more than 235,408,082 confirmed COVID-19 cases including 4,809,149 deaths worldwide till October 6, 2021. During the coronavirus epidemic, almost everyone wears a facial mask, which poses a huge challenge to face recognition. Traditional SFR may not effectively recognize the masked faces, but removing the mask for authentication would increase the virus infection risk.

To cope with the above-mentioned challenging scenarios arising from wearing masks, it is crucial to improve the existing face recognition approaches. Recently, some commercial vendors [61] have developed face recognition algorithms capable of handling face masks, and an increasing number of research publications [10], [24], [26], [30], [36] have surfaced on this topic. However, due to the sudden outbreak of the epidemic, there is yet no publicly available large-scale MFR benchmark.

### 5.3.2 Test Set with Real-world Mask

In contrast with simulated [59], [61] or relatively small [10], [12], [17], [95] masked test sets, a real-world comprehensive benchmark for evaluating MFR is developed in this work. Based on the SFR identities, we further collect masked faces for these celebrities. Specifically, as shown in Table 3, there are carefully selected 3,211 masked faces among 862 identities. Subjects with real-world masks are illustrated in Figure 7. Wearing masks causes severe occlusion, resulting in just the periocular area and above visible. Besides, there are changeable mask types, colors, wearing ways, and head poses in real-world applications, which are more practical and challenging than simulated ones.

For MFR, assessment is performed with *Mask-Nonmask* comparisons. Specifically, there are one masked face and



TABLE 3

The statistics of our test set, including SFR, MFR, and UFR evaluations. - means corresponding statistics or comparisons are omitted.

Evaluation	Attributes		# Identities	# Faces	# Impostor	# Genuine
Standard Face Recognition (SFR)	All		2,478	57,715	1,664,475,460	1,006,295
	Age	Cross-age-10	-	-	833,518,127	276,587
		Cross-age-20	-	-	8,421,183,318	61,780
	Scenarios	Controlled	-	22,135	244,828,795	139,250
		Wild	-	35,580	632,449,086	501,324
Cross-scene		-	-	787,197,579	365,721	
Masked Face Recognition (MFR)	All		2,478	60,926	-	-
	Masked		862	3,211	-	-
	Nonmasked		2,478	57,715	-	-
	Mask	Controlled-Masked	-	-	71,042,982	32,503
		Wild-Masked	-	-	114,193,476	53,904
All-Masked		-	-	185,236,458	86,407	
Unbiased Face Recognition (UFR)	Race	Caucasian	990	22,940	262,676,189	434,141
		East Asian	874	22,970	263,392,669	406,296
		African	447	8,528	36,234,242	124,886
		Others	167	3,277	-	-
	Gender	Male	1527	33,872	573,108,363	530,893
		Female	951	23,843	283,757,001	475,402

another face from standard face sets for pair verification. According to the attributes of faces without masks, we evaluate the performance of algorithms under *Controlled-Masked*, *Wild-Masked*, and *All-Masked* settings in Table 3.

## 5.4 Unbiased Face Recognition

### 5.4.1 Bias and Fairness in Face Biometrics

Bias in face recognition means it provides higher accuracy within certain demographic groups and lower performance for other demographics. According to the NIST-FRVT report [34], most submitted recognition algorithms from academia and industry exhibit different levels of biased performances. Deploying such systems may cause significant consequences such as racism. Recent UFR researches [31], [81], [90], [91] mainly focus on balanced data collection/sampling and de-biased algorithm design. For evaluation, most of the tests are performed on RFW set [91], which adopts LFW-like pair comparisons. In this paper, based on the proposed test data and FRUITS protocol, we provide a more challenging and practical UFR evaluation.

### 5.4.2 Test Set and Metrics

To enable a trustworthy face recognition system, it is of importance to investigate the performance on different facial attributes. As shown in Table 3, the test set of UFR is the same as SFR ones. We manually label race and gender attributes for unbiased evaluation. Fairness assessment of ethnicity is reported on Caucasian, East Asian, African, and so does gender. Following common practices [90] in the community, we adopt skewed error ratio (SER) and standard deviation (STD) as the fairness metrics. Specifically, the error ratio of each race and gender attribute is calculated according to  $FNMR@FMR=1e-5$ . Then, SER can be computed by the ratio of the highest to the lowest error rate among race and gender groups. STD is the number of error dispersion among different races and genders.

## 6 MILLION-LEVEL FACE RECOGNITION EXPERIMENTS

Based on the constructed WebFace260M benchmark, we dive deep into the million-scale face recognition in this section.

Firstly, implementations are detailed containing parameter and environment configurations. Then we analyze the speed and performance of distributed framework, which enables large-scale face recognition training. Thirdly, the WebFace data is compared with public counterparts, covering different losses and test sets. Furthermore, the proposed CAST strategy and its key procedures are studied. Lastly, we establish comprehensive baselines for SFR, MFR, UFR, and report the performance on NIST-FRVT.

## 6.1 Implementation Details

**Hyper-parameters.** In order to fairly evaluate the performance of different face recognition models, we reproduce representative algorithms (CosFace [89], ArcFace [21] and CurricularFace [42]) in one Gluon codebase. Margin values in CosFace [89], ArcFace [21] and CurricularFace [42] are set as 0.35, 0.5 and 0.5, respectively. The Stochastic Gradient Descent (SGD) with momentum 0.9 and weight decay 0.0005 is utilized for the network optimization. For large-batch training with cluster, we employ distributed synchronous SGD, which parallelizes the tasks across machines. The default batch size per GPU is set as 64 unless otherwise indicated. The learning rate is set as 0.05 for a single node (8 GPUs), and follows the linear scaling rule [33] for the training on multiple nodes ( $0.05 \times \#$  machines). We decrease the learning rate by  $0.1 \times$  at 8, 12, and 16 epochs, and stop at 20 epochs for all models. Gradual warmup [33] is adopted during the initial 1 epoch and 5 epochs for a single node and multiple nodes training, respectively. During training, we only adopt horizontal flip data augmentation. In DBSCAN, similarity ( $1 - \epsilon$ ) is set as 0.5, 0.55, and 0.6 for 1st, 2nd, and 3rd iterations, respectively.  $n$  is set as 3 for all iterations.

**Hardware and software.** All experiments are performed on a cluster containing 35 nodes, and each node contains 8 GPUs, 2 CPU processors, and 384GB RAM memory. Machines are inter-connected with a 30 Gbps TCP/IP network. For storage, the cluster has 32TB distributed SSDs. For face recognition training, CUDA version 10.0 with cuDNN version 7.6.0 is employed. We use NCCL version 2.4.7 as the communication library.

TABLE 4

Speed and performance comparisons of distributed training. ResNet-100 backbone with ArcFace loss is adopted. B, G and M refer to batch size per GPU, # GPUs per machine, and # machines.  $X$  and  $W$  mean feature and center, and numbers in bracket are the GPU memory usage (MB). Performance is reported on IJB-C (TAR@FAR=1e-4).

Data	B×G×M	FP32/16	Parallel	Speed	Time	IJB-C
10%	32×8×1	FP32	$X$ (7913)	0.6K	39h	96.67
	64×8×1	FP32	$X$ $W$ (7521)	0.9K	26h	96.83
	64×8×1	FP16	$X$ (7551)	1K	23h	96.80
	64×8×1	FP16	$X$ $W$ (7182)	1.8K	13h	96.78
	64×8×4	FP16	$X$ $W$ (7125)	6.3K	4h	96.73
30%	64×8×8	FP16	$X$ $W$ (7119)	12.4K	2h	96.77
	64×8×1	FP16	$X$ $W$ (8901)	1.7K	41h	97.41
	64×8×4	FP16	$X$ $W$ (8519)	5.5K	13h	97.50
100%	64×8×8	FP16	$X$ $W$ (8455)	11.3K	6h	97.47
	32×8×1	FP16	$X$ $W$ (10503)	1K	233h	97.71
	32×8×8	FP16	$X$ $W$ (8359)	6.8K	34h	97.65
	32×8×16	FP16	$X$ $W$ (8297)	12.9K	18h	97.74
	32×8×32	FP16	$X$ $W$ (8221)	25.3K	<b>9h</b>	97.70




Fig. 8. Speed and performance of our distributed training system, which achieves almost linear acceleration with comparable performance. 100% data (WebFace42M) is used in this experiment.

## 6.2 Distributed Training

When using the ultra-large-scale WebFace42M as the training data and computationally demanding backbones as the embedding networks, the model optimization can take several weeks on a single machine. Such a long training time makes it difficult to efficiently perform experiments. Inspired by the distributed optimization on ImageNet [33], we apportion the workload of model training to clusters. To this end, parallelization on both feature  $X$  and center  $W$ , mixed-precision (FP16) and large-batch training are adopted in this work.

The speed and performance of our distributed training system are illustrated in Table 4 and Figure 8. Parallelization on both feature  $X$  and center  $W$  as well as mixed-precision (FP16) significantly reduce the consumption of GPU memory and speed up the training process, while similar performance can be achieved. Equipped with 8 nodes (64 GPUs), the training speed is scaled to 12K samples/s and 11K samples/s on WebFace4M (10% data) and WebFace12M (30% data), respectively. The corresponding training time is only 2 hours and 6 hours. Furthermore, the scaling efficiency of our training system is above 80% when applied to ultra-large-scale WebFace42M on 32 nodes (256 GPUs). Figure 8 shows the ideal speed, actual speed, and corresponding performance for increasing GPUs resource. We can reduce the training time of the ResNet-100 model from 233 hours (1 node) to 9 hours (32 nodes) with comparable performance.

TABLE 6

Performance (%) of ArcFace models trained with ResNet-14 on different portions of WebFace data. TAR@FAR=1e-4 on IJB-C is reported.

Training data	WebFace4M	WebFace12M	WebFace42M
IJB-C	93.13	93.92	94.22




Fig. 9. Performance of ArcFace models (ResNet-100) trained on the WebFace envelopes counterparts trained on the public training data.

## 6.3 Comparisons of Training Data

For comprehensively benchmarking the influence of training data, the proposed WebFace42M is compared with public counterparts including MS1M families [1], [21], [23], [35], MegaFace2 [58] and IMDB-Face [83]. 10% (WebFace4M) and 30% (WebFace12M) random selection of our full data are also employed for further analyses. The statistics of different training sets are illustrated in Table 1. Evaluation sets used in this experiment include popular verification sets (LFW [41], CALFW [114], CPLFW [113], AgeDB [57], CFP-FP [67]), RFW [91], MegaFace [58], IJB-C [56] and our test set.

As we can see from Table 5 and Figure 9, the proposed WebFace42M breaks the bottleneck of training data for deep face recognition across various loss functions and test sets. In summary, almost all best accuracies on each test set are achieved by the proposed WebFace data shown in Table 5. Specifically, WebFace42M reduces 40% error rate on the challenging IJB-C dataset compared with MS1MV2, boosting TAR@FAR=1e-4 from 96.03% to 97.70% with ResNet-100 and ArcFace. Along with the increment of data scale (10%, 30%, and 100%), there exists a consistent improvement in performance as observed in Figure 9, enveloping its public counterparts. On MegaFace identification task, WebFace42M sets a new state of the art by 99.11% rank-1 score. On our test set, for *All* pairs comparison of SFR, WebFace42M decreases the FNMR@FMR=1e-5 from 9.88% to 2.98%, reducing the error rate by more than 3 times. For *All-Masked* pairs comparison of MFR, FNMR@FMR=1e-5 metric is boosted from 69.56% to 42.97% by the proposed training data. Meanwhile, due to the tremendous scale and diversity of WebFace, it benefits UFR according to SER metric. Furthermore, the models trained on 10% data, WebFace4M, impressively achieve superior performance compared to ones trained on MS1M families and MegaFace2, which include even more faces. Undisputedly, this comparison confirms the effectiveness and necessity of our WebFace42M in leveling the playing field for million-scale face recognition.

Besides reporting the face recognition results of ResNet-

TABLE 5

Performance comparisons of our WebFace and public training data. ResNet-100 backbone **without flip test** is adopted. RFW refers to average accuracy on [91], MegaFace refers to rank-1 identification and verification scores on [44], IJB-C is TAR@FAR=1e-5 and 1e-4 on [56]. On our test set, FNMR@FMR=1e-5 on *All* pairs of SFR, FNMR@FMR=1e-5 on *All-Masked* pairs of MFR, race SER metric of UFR are reported (**lower is better**). The best results are marked in bold.

Data	Loss	LFW	CALFW	CPLFW	AgeDB30	CFP-FP	RFW	MegaFace		IJB-C		Our test set ↓		
								Id.	Veri.	10-5	10-4	SFR	MFR	UFR
MS1M	CosFace	98.32	95.19	91.84	96.01	97.09	98.09	96.21	97.73	86.48	92.96	62.91	89.36	2.37
	ArcFace	98.22	94.05	91.59	96.73	97.07	97.64	97.67	97.94	89.03	93.45	61.73	86.42	1.84
	CurricularFace	98.89	95.46	91.18	96.82	96.21	98.12	96.86	97.26	85.47	92.99	64.05	87.78	1.65
MS1M-IBUG	CosFace	99.76	95.83	89.98	94.95	97.83	97.62	97.33	97.60	89.79	94.35	12.22	71.41	1.32
	ArcFace	99.76	95.91	89.41	94.40	97.95	97.27	97.27	97.35	91.50	94.57	13.46	74.61	1.49
	CurricularFace	99.78	95.95	90.37	94.51	97.94	97.86	97.19	97.24	89.85	94.72	13.35	75.05	1.43
MS1MV2	CosFace	99.81	96.18	92.76	98.34	98.18	98.85	98.30	98.33	94.62	96.01	9.04	68.25	1.34
	ArcFace	99.78	96.05	92.93	98.21	98.54	98.98	98.40	98.24	94.05	96.03	9.88	71.81	1.40
	CurricularFace	99.83	96.28	93.05	98.32	98.67	99.02	98.46	98.47	94.01	96.21	9.52	69.56	1.38
MS1M-Glint	CosFace	99.83	96.00	91.03	94.95	98.16	99.59	98.60	98.83	94.00	96.15	12.65	74.33	2.00
	ArcFace	99.71	95.96	90.76	94.20	98.43	99.60	98.48	98.31	93.35	96.24	13.98	78.09	2.19
	CurricularFace	99.80	<b>96.31</b>	91.65	95.47	98.82	<b>99.65</b>	98.57	98.60	93.20	96.31	14.27	74.27	2.12
MegaFace2	CosFace	99.46	93.25	92.46	87.67	89.76	88.90	86.62	89.13	79.75	87.75	58.77	92.42	1.55
	ArcFace	99.50	92.86	88.60	91.06	93.88	89.45	88.28	90.68	81.75	89.35	54.73	91.54	1.57
	CurricularFace	99.48	93.02	90.39	91.07	93.04	90.06	88.32	90.00	82.27	90.11	55.73	91.65	1.56
IMDB-Face	CosFace	99.61	95.08	92.15	96.75	98.48	93.80	94.03	95.12	89.92	93.96	27.88	87.56	2.92
	ArcFace	99.58	94.85	91.83	97.29	98.45	93.08	93.48	94.62	89.09	93.37	30.45	90.25	3.17
	CurricularFace	99.73	94.81	92.44	97.46	98.66	94.11	93.63	94.68	88.78	94.12	30.94	88.54	2.95
WebFace4M	CosFace	99.80	95.95	94.40	97.45	99.25	98.16	97.59	98.32	94.93	96.86	8.98	68.77	1.49
	ArcFace	<b>99.85</b>	95.93	94.31	97.82	99.04	98.14	97.60	98.00	94.72	96.77	10.39	70.97	1.45
	CurricularFace	99.83	96.03	94.21	97.83	99.11	98.14	97.94	98.19	95.15	97.02	9.09	66.58	1.51
WebFace12M	CosFace	99.81	96.18	94.80	97.88	<b>99.38</b>	99.15	98.66	98.75	96.00	97.41	4.34	54.08	1.41
	ArcFace	99.81	96.11	94.68	98.33	99.37	99.08	98.82	98.63	95.96	97.47	4.77	56.47	1.39
	CurricularFace	99.83	96.13	94.70	98.37	99.37	99.18	98.75	<b>98.85</b>	95.99	97.51	4.81	53.05	1.46
WebFace42M	CosFace	99.83	96.11	94.90	<b>98.58</b>	<b>99.38</b>	99.41	99.02	98.57	<b>96.44</b>	97.68	3.01	45.28	1.31
	ArcFace	99.83	96.19	<b>94.93</b>	98.02	99.28	99.33	99.02	98.61	96.23	97.70	<b>2.98</b>	47.25	<b>1.30</b>
	CurricularFace	99.83	96.10	94.85	98.24	<b>99.38</b>	99.39	<b>99.11</b>	98.55	96.38	<b>97.76</b>	3.05	<b>42.97</b>	1.33

100, we also train ArcFace models by using a ResNet-14 network on different portions of our data (10%, 30% and 100%). As given in Table 6, there is also a consistent performance gain for ResNet-14 when more training data are progressively employed. Therefore, the proposed WebFace42M is not only beneficial to the large model (ResNet-100) but also valuable for the lightweight model. This is of significance for mobile devices like cellphones, and we explore smaller models in Section 6.5.

### 6.4 Comparisons of Data Cleaning

As shown in Table 7, the CAST pipeline is compared with other cleaning strategies on the original MS1M [35] and WebFace260M. Specifically, for MS1M results, the initial teacher model is trained on IMDB-Face [83] by using ResNet-100 and ArcFace. Then, CAST is conducted on the noisy MS1M following Section 4. After steps of iteration, our fully automatic cleaning strategy provides highly-curated data for model training, outperforming semi-automatic methods used in [1], [21], [23]. Compared with the most recent GCN-based cleaning [109], the data cleaned by the CAST also achieves higher performance.

**Iterations of CAST.** Table 7 also shows the increasing data purity after more iterations in MS1M and WebFace260M. The accuracy gradually increases from 1st to 3rd iteration, while the 4th iteration shows saturated performance. Therefore, we set the iteration number as 3 for CAST.

**Performance Improvement Reporting of CAST.** From Table 7, one can find the recognition performance improvement as the cleaning iteration continues. Specifically, after the first-round iteration, WebFace data results in 97.28% TAR@FAR=1e-4 on IJB-C, which is already much higher

TABLE 7

Comparisons of CAST and other data cleaning pipelines. Pairs refers to average accuracy on [41], [57], [67], [113], [114]. MegaFace and IJB-C refer to rank-1 identification and TAR@FAR=1e-4 respectively. For MS1M and WebFace by CAST, different iterations are compared. CAST-1 means the first-round iteration.

Data	# Id	# Face	Pairs	MegaFace	IJB-C
MS1M	100K	10M	95.53	97.67	93.45
MS1M-IBUG	85K	3.8M	95.49	97.27	94.57
MS1MV2	85K	5.8M	97.10	98.40	96.03
MS1M-Glint	87K	3.9M	95.81	98.48	96.24
MS1M-GCN [109]	-	-	96.51	-	-
MS1M by CAST-1	94K	6.3M	95.37	97.93	94.31
MS1M by CAST-2	92K	5.5M	97.08	98.47	95.90
MS1M by CAST-3	91K	4.9M	97.42	98.61	96.55
MS1M by CAST-4	91K	4.9M	97.49	98.57	96.52
WebFace by CAST-1	2.4M	46M	97.42	98.64	97.28
WebFace by CAST-2	2.1M	43M	97.53	98.98	97.51
WebFace by CAST-3	2M	42M	97.65	99.02	97.70
WebFace by CAST-4	2M	42M	97.69	99.08	97.66

than MS1MV2 ones (96.03%). There are still some mislabeled identities which cause performance degradation. As the cleaning iteration continues, the number of identities decreases, while resulting WebFace datasets further obtain 97.51% and 97.70% on TAR@FAR=1e-4 of IJB-C. Results on Pairs and MegaFace also lead to similar conclusions. This performance improvement reporting clearly shows the effectiveness of the proposed cleaning method.

**Different Settings of Teacher/Student Model.** In Table 8, we further perform CAST using different settings, including ResNet-200 and MobileFaceNet backbones for Teacher/Student model, and Glint360K [9] dataset for training initial Teacher model. The resulting cleaned WebFace data is then used to train face recognition models,

TABLE 8

Performance (%) of different Teacher/Student models and various data for initial Teacher model. The resulting cleaned WebFace data is then used to train face recognition models, which follows the same ResNet-100 and ArcFace setting for fair comparisons.

Teacher/Student Model	Data for Initial Teacher Model	Pairs	MegaFace	IJB-C
ResNet-100	MS1MV2	97.65	99.02	97.70
ResNet-200	MS1MV2	97.63	99.07	97.71
ResNet-100	Glnt360K	97.66	99.06	97.68
MobileFaceNet	MS1MV2	97.02	98.24	95.82

which adopt ResNet-100 and ArcFace for fair comparisons. The experimental results are summarized as follows: (1) Replacing ResNet-100 with ResNet-200 for Teacher/Student model results in WebFace data with similar performance. We argue that ResNet-100 is competent to perform the cleaning process, and larger backbones offer marginal improvements. It is worth noting that training ResNet-200 costs several times more GPU hours than ResNet-100 ones. (2) Replacing MS1MV2 with Glnt360K for initial Teacher model also results in WebFace data with similar performance. Although Glnt360K is more powerful, it is only utilized in 1st iteration of CAST, while later Teacher/Student models are all trained on WebFace data. The face recognition model trained with MS1MV2 provides a good initial embedding feature to perform 1st round cleaning, which is effective for CAST process. (3) Replacing ResNet-100 with MobileFaceNet for Teacher/Student model shows dramatically degraded cleaning results. Due to the limited capacity, such inferior face recognition models like MobileFaceNet can not handle the ultra-large-scale WebFace training. Figure 10(a) also illustrates the MobileFaceNet performance trained with WebFace42M, which shows very limited recognition accuracy. Considering the trade-off between cleaning cost and quality, we choose ResNet-100 for Teacher/Student model and MS1MV2 for initial Teacher model in CAST.

Furthermore, we investigate the influence of identities duplicates between MS1M and WebFace datasets. As described in Section 3.1, our WebFace celebrity name list consists of two parts: the first one is borrowed from MS1M and the second one is collected from the IMDB database. Since the MS1M identities are most a subset of the WebFace ones, there are few remaining identities in MS1M if removing the duplicates. Alternatively, we remove the MS1M identities from WebFace, which is denoted as WebFace-no-MS1M. Then, WebFace and WebFace-no-MS1M are cleaned following the same settings (the Teacher model initially uses the recognition model trained by MS1M). We find that these 2 datasets obtain similar performance. Specifically, for IJB-C TAR@FAR=1e-4 metric, training sets cleaned from WebFace and WebFace-no-MS1M result in 97.70% and 97.67%, respectively. It is worth noting that MS1M name list has only 100K identities, while WebFace name list has 4M identities. This comparison shows that duplicated identities between MS1M and WebFace do not influence the effectiveness of CAST.

**Intra-class Cleaning.** In this ablation study, we compare different intra-class cleaning modules under the framework of CAST. Both unsupervised (such as K-means [55] and DBSCAN [29]) and supervised styles (such as GCN-D [104]

TABLE 9

Comparisons of different intra-class cleaning strategies for MS1M. ResNet-100 backbone with ArcFace loss is adopted here.

Data	# Id	# Face	Pairs	MegaFace	IJB-C
K-Means	93K	5.2M	95.17	97.31	96.03
DBSCAN	91K	4.9M	<b>97.42</b>	<b>98.61</b>	<b>96.55</b>
GCN-D	86K	4.4M	96.56	98.55	96.48
GCN-V	82K	4.5M	96.93	98.29	96.42

TABLE 10

Configuration and inference time of SFR baselines. Loose cropped test images are resized to  $224 \times 224$  for joint detection and alignment. M-0.25 and R-50 refer to RetinaFace using MobileNet-0.25 (23ms) and ResNet-50 (272ms) as the backbones. FLOPs and Params mean computational complexity and parameter number of recognition module, respectively. Time refers to the duration of the whole system. For evaluation with flip setting, the batch size is set as 2.

Protocol	Det&Align	Embedding	FLOPs	Params	Time
FRUITS -100	M-0.25	ResNet-14	2.1G	19.2M	97ms
	M-0.25	MobileFaceNet (Flip)	230.3M	1.2M	65ms
	M-0.25	EfficientNet-B0	394.2M	11.6M	94ms
	M-0.25	RegNet-800MF	831.0M	23.4M	89ms
FRUITS -500	R-50	ResNet-100	12.1G	65.2M	481ms
	R-50	ResNet-50 (Flip)	6.3G	43.6M	492ms
	R-50	SENet-50	6.3G	43.8M	374ms
	R-50	ResNeXt-100	8.2G	56.2M	411ms
	R-50	RegNet-8GF	8.0G	82.7M	429ms
FRUITS -1000	R-50	ResNet-100 (Flip)	12.1G	65.2M	826ms
	R-50	ResNet-200	23.9G	109.3M	892ms
	R-50	SENet-152	18.1G	101.0M	792ms
	R-50	AttentionNet-152	14.8G	61.3M	785ms
	R-50	RegNet-16GF	16.0G	103.7M	772ms

and GCN-V [103]) are explored to find the dominant subject in each noisy folder. As shown in Table 9, DBSCAN achieves 96.55% TAR@FAR=1e-4 on IJB-C, significantly outperforming K-Means (96.03%) and slightly surpassing the supervised GCN-based ones (96.48% for GCN-D and 96.42% for GCN-V). As the GCN-based strategies may be sub-optimal for the extremely noisy folders, we finally select DBSCAN [29] as our intra-class cleaning module.

## 6.5 Baselines under FRUITS Protocol

### 6.5.1 Standard Face Recognition

In this section, we set up a series of SFR baselines under the proposed FRUITS protocol. Table 10 illustrates various face recognition systems (including different settings of detection, alignment, feature embedding) and their inference time. In our SFR baselines, representative network architectures are explored, covering MobileNet [15], [38], EfficientNet [80], AttentionNet [84], ResNet [37], SENet [39], ResNeXt [100] and RegNet [64] families. All the models are trained on WebFace42M with ArcFace.

Due to the strict time limitation, FRUITS-100 track can only adopt lightweight architectures, including RetinaFace-MobileNet-0.25 [20] for face detection and alignment, ResNet-14, MobileFaceNet (Flip), EfficientNet-B0 and RegNet-800MF for face feature extraction. FMR-FNMR plots on *All* pairs and analyses of attributes are shown in Figure 10(a) and Figure 11(a). Because of the weak detection and recognition modules, the best baseline (RegNet-800MF) only obtains 12.41% FNMR@FMR=1e-5 (lower is better). Therefore, there leaves substantial room for future improvement under the FRUITS-100 protocol.

TABLE 11

Performance (%) of different training data for MFR, including *Controlled-Masked*, *Wild-Masked*, and *All-Masked* comparisons. *Mask* suffix of training data refers to augmentation with simulated masks. FNMR@FMR=10-5 of ResNet-100 and ArcFace is reported.

Data	Controlled-Masked	Wild-Masked	All-Masked
MS1MV2	64.40	76.07	71.81
WebFace4M	61.48	75.43	70.97
WebFace4M-Mask	49.58	65.53	59.89
WebFace12M	45.46	61.57	56.47
WebFace12M-Mask	31.61	46.98	41.45
WebFace42M	36.81	52.83	47.25
WebFace42M-Mask	<b>24.60</b>	<b>39.42</b>	<b>33.87</b>

For the FRUITS-500 protocol, we can employ more capable modern networks, such as RetinaFace-ResNet-50 [20] for pre-processing, and ResNet-100, ResNet-50 (Flip), SENet-50, ResNeXt-100, RegNet-8GF for feature embedding extraction. As shown in Figure 10(b) and Figure 11(b), ResNet-100 exhibits the best overall performance, scoring 2.98% FNMR@FMR=1e-5. For attributes evaluation, ResNet-100 also achieves the lowest FNMR according to the indicators of age, scenarios, race and gender.

Recognition models under the FRUITS-1000 protocol can be more complicated and powerful, therefore we explore ResNet-100 (Flip), ResNet-200, SENet-152, AttentionNet-152 and RegNet-16GF for face feature representations. As shown in Figure 10(c) and Figure 11(c), ResNet-200 performs best in face verification and wins all attribute comparisons. Compared with lightweight FRUITS-100 track, performances of different large models are much closer. This result implies that new designs need to be explored for the heavyweight FRUITS track.

### 6.5.2 Masked Face Recognition

Recognizing identities with masks may be the most challenging face recognition problem, which is essential for biometric authentication during COVID-19. Based on the proposed WebFace benchmark, we perform MFR in this part, and establish a series of baselines for different training settings. ResNet-100 in FRUITS-500 protocol and ArcFace are adopted, while similar conclusions can be drawn for other backbones or losses.

As shown in Table 11, the best-performed public training set, MS1MV2, only scores 71.81% FNMR@FMR=10-5 for *All-Masked* comparisons on our test set (with ArcFace loss). Different proportions of the WebFace data reduce FNMR to 70.97%, 56.47% and 47.25% respectively, which shows the superiority of the proposed dataset again. Moreover, we augment the training data with simulated masks to further investigate this difficult recognition scenario. Specifically, mask renderer in [7] is applied on each face, where wearing height and mask types are randomly chosen. Both simulated and original faces are trained together with 20 epochs. In Table 11, one can find that this simple augmentation strategy effectively boosts the MFR accuracy. For all data proportions, the FNMR improvements are near 10%, which build strong baselines for future MFR researches.

### 6.5.3 Unbiased Face Recognition

Generalized deployments call for robust and fair face recognition systems. In this part, we perform data sampling on

WebFace data and investigate its fairness influence on our test set. Table 12 indicates that MS1MV2 and WebFace4M show considerable bias among race and gender. The scores of SER metric are 1.40, 1.45, 1.88, and 1.92 respectively. The relative rank of different models (trained with WebFace42M) on race and gender attributes is also illustrated in Figure 11. Thanks to the ultra-large-scale of the proposed benchmark, we can sample a balanced race/gender sub-set denoted as WebFace4M-Balanced. According to STD and SER scores, this sampled training data reduces the recognition bias to some extent, surpassing the MS1MV2 and WebFace4M. It is worth noting that there is still demographic bias on the test set even with WebFace4M-Balanced training, which is consistent with the observation in previous studies [90], [91]. More bias-mitigating solutions need to be developed such as loss design, augmentation, and adversarial learning.

In summary, the results show the great potential of WebFace data (including training and test set) for more fair and robust face recognition systems. Based on the proposed WebFace benchmark, we hope to spark UFR researches in the future.

## 6.6 Results on NIST-FRVT

Finally, we report the submission to NIST-FRVT. Following the settings of FRUITS-1000, our system is built based on RetinaFace-ResNet-50 for detection and alignment, and ArcFace-ResNet-200 trained on WebFace42M for feature embedding extraction. The network is accelerated by OpenVINO [8] and the flip test is adopted. The final inference time is near 1300 milliseconds according to the NIST-FRVT report, meeting the latest 1500 milliseconds limitation. Table 13 illustrates top-ranking entries measured by FNMR across six tracks<sup>3</sup>. **Our model trained on the WebFace42M obtains overall 3rd among 430 submissions**, showing impressive performance across different tracks. Specifically, the proposed solution based on WebFace ranks 5th and 3rd on controlled *Visa* and *Mugshot* scenarios, respectively. On more challenging cross-age comparisons (*Mugshot DT ≥ 12Years*), we get 2nd place. For less-controlled *VisaBorder*, *Border* and *Wild* tracks, state-of-the-art performances are also achieved. Considering hundreds of company entries to NIST-FRVT, the WebFace42M takes a significant step towards closing the data gap between academia and industry.

## 7 DISCUSSION AND CONCLUSION

**Discussion** During WebFace260M dataset construction, privacy and bias issues are our primary concerns. For privacy protection, all face images including training and test data are collected from public Internet resources. For data download, we provide strict access for qualified research groups that sign the license, and try our best to guarantee WebFace260M for research purposes only. For bias problem, our dataset has diverse birth dates, poses and ages, while gender and race are inevitably biased due to complex nationality and profession distributions. In evaluations of this work, we especially design the Unbiased Face Recognition (UFR), studying the influence of balanced training data. We argue that the

3. According to report of October 9, 2020




Fig. 10. Performance comparisons of SFR under the FRUITS. The FMR-FNMR plots for *All* pairs verification are drawn, and models are ranked in legend according to FNMR@FMR=1e-5 (**lower FNMR is better**).




Fig. 11. Attributes ranking of different models under the FRUITS. We show the attribute plots under FNMR@FMR=1e-5, which is normalized to 0.5-1.0 for better visualization (**outer is better**).

TABLE 12

Performance of different training data for UFR. For race and gender groups, FNMR@FMR=10-5 of ResNet-100 and ArcFace is reported. STD and SER are main fairness evaluation metrics.

Data	Race						Gender				
	Caucasian	East Asian	African	Avg	STD	SER	Male	Famale	Avg	STD	SER
MS1MV2	0.1050	0.1474	0.1053	0.1192	0.0199	1.40	0.0850	0.1597	0.1224	0.0374	1.88
WebFace4M	0.0942	0.1368	0.1071	0.1127	0.0178	1.45	0.0890	0.1713	0.1302	0.0412	1.92
WebFace4M-Balanced	0.0976	0.1253	0.1025	0.1085	<b>0.0121</b>	<b>1.28</b>	0.0825	0.1332	0.1079	<b>0.0254</b>	<b>1.61</b>

TABLE 13

Results on NIST-FRVT. Numbers in bracket are rankings of different tracks. Our ResNet-200 model using Arcface is trained on WebFace42M. FNMR at corresponding FMR is reported.

Rank	entries	Visa	Mugshot	Mugshot DT <sub>≥12</sub> Years	VisaBorder	Border	Wild
		FNMR@FMR=10-6	FNMR@FMR=10-5	FNMR@FMR=10-5	FNMR@FMR=10-6	FNMR@FMR=10-6	FNMR@FMR=10-5
1	deepglint	0.0027 (3)	0.0032 (8)	0.0033 (4)	0.0043 (3)	0.0084 (4)	0.0301 (4)
2	visionlabs	0.0025 (1)	0.0026 (1)	0.0029 (3)	0.0035 (1)	0.0064 (1)	0.0306 (13)
3	<b>ours</b>	<b>0.0034 (5)</b>	<b>0.0027 (3)</b>	<b>0.0028 (2)</b>	<b>0.0046 (6)</b>	<b>0.0088 (6)</b>	<b>0.0303 (10)</b>
4	dahua	0.0046 (13)	0.0035 (12)	0.0049 (13)	0.0046 (5)	0.0076 (2)	0.0300 (3)
5	cib	0.0061 (20)	0.0030 (6)	0.0041 (7)	0.0048 (8)	0.0578 (76)	0.0302 (25)
6	nazhi	0.0059 (16)	0.0036 (13)	0.0048 (12)	0.0057 (12)	0.0125 (19)	0.0300 (2)
7	vocord	0.0038 (6)	0.0042 (18)	0.0055 (15)	0.0045 (4)	0.0086 (5)	0.0310 (24)
8	ercacat	0.0044 (10)	0.0033 (9)	0.0047 (10)	0.0106 (43)	0.0202 (36)	0.0293 (1)
9	everai-paravision	0.0050 (14)	0.0036 (14)	0.0052 (14)	0.0092 (35)	0.0193 (35)	0.0302 (7)
10	aimall	0.0041 (8)	0.0033 (10)	0.0035 (5)	0.0056 (11)	0.0109 (15)	0.0335 (49)

community could develop more bias-mitigating solutions based on ultra-large-scale WebFace260M benchmark.

**Conclusion** In this paper, we have dived into the million-scale face recognition problem, contributing a tremendous noisy dataset with 260M faces, a high-quality training dataset

with 42M images of 2M identities by using automatic cleaning, a test set containing rich attributes and large-scale masked face sub-set, a time-constrained evaluation protocol, a distributed framework at linear acceleration, a succession of baselines on various scenarios, as well as a final state-of-

the-art model. Equipped with this publicly available face dataset, our model significantly reduces 40% failure rate on IJB-C and ranks 3rd among 430 entries on NIST-FRVT. Besides, baselines built on the proposed WebFace show great potential for masked and unbiased recognition tasks. We hope this benchmark could close the data gap behind the industry, and facilitate future researches of ultra-large-scale face recognition.

### ACKNOWLEDGEMENT

This work was supported in part by the National Key Research and Development Program of China under Grant 2017YFA0700802, in part by the National Natural Science Foundation of China under Grant 62125603 and Grant U1813218, and in part by a grant from the Beijing Academy of Artificial Intelligence (BAAI).

### REFERENCES

[1] <http://trillionpairs.deeplint.com/overview>.

[2] <https://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt-ongoing>.

[3] <https://www.bing.com/image>.

[4] <https://developers.google.com/freebase/>.

[5] <https://www.imdb.com/>.

[6] <https://images.google.com>.

[7] [https://github.com/deepinsight/insightface/tree/master/recognition/\\_tools\\_/](https://github.com/deepinsight/insightface/tree/master/recognition/_tools_/).

[8] <https://github.com/openvinotoolkit/openvino>.

[9] X. An, X. Zhu, Y. Xiao, L. Wu, M. Zhang, Y. Gao, B. Qin, D. Zhang, and Y. Fu. Partial FC: Training 10 million identities on a single machine. *arXiv:2010.05222*, 2020.

[10] A. Anwar and A. Raychowdhury. Masked face recognition for secure authentication. *arXiv:2008.11104*, 2020.

[11] A. Bansal, A. Nanduri, C. D. Castillo, R. Ranjan, and R. Chellappa. UMDFaces: An annotated face dataset for training deep networks. *arXiv:1611.01484v2*, 2016.

[12] F. Boutros, N. Damer, J. N. Kolf, K. Raja, F. Kirchbuchner, R. Ramachandra, A. Kuijper, P. Fang, C. Zhang, F. Wang, et al. MFR 2021: Masked face recognition competition. *IJCB Workshop*, 2021.

[13] D. Cao, X. Zhu, X. Huang, J. Guo, and Z. Lei. Domain balancing: Face recognition on long-tailed domains. In *CVPR*, 2020.

[14] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. VGGFace2: A dataset for recognising faces across pose and age. In *FG*, 2018.

[15] S. Chen, Y. Liu, X. Gao, and Z. Han. MobileFacenets: Efficient CNNs for accurate real-time face verification on mobile devices. In *CCBR*, 2018.

[16] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005.

[17] N. Damer, F. Boutros, M. Süßmilch, F. Kirchbuchner, and A. Kuijper. Extended evaluation of the effect of real and simulated masks on face recognition performance. *IET Biometrics*, 2021.

[18] J. Deng, J. Guo, X. An, Z. Zhu, and S. Zafeiriou. Masked face recognition challenge: The InsightFace track report. *arXiv preprint arXiv:2108.08191*, 2021.

[19] J. Deng, J. Guo, T. Liu, M. Gong, and S. Zafeiriou. Sub-center ArcFace: Boosting face recognition by large-scale noisy web faces. In *ECCV*, 2020.

[20] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou. RetinaFace: Single-shot multi-level face localisation in the wild. In *CVPR*, 2020.

[21] J. Deng, J. Guo, and S. Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.

[22] J. Deng, J. Guo, D. Zhang, Y. Deng, X. Lu, and S. Shi. Lightweight face recognition challenge. In *ICCV Workshop*, 2019.

[23] J. Deng, Y. Zhou, and S. Zafeiriou. Marginal loss for deep face recognition. In *CVPR Workshop*, 2017.

[24] F. Ding, P. Peng, Y. Huang, M. Geng, and Y. Tian. Masked face recognition with latent part detection. In *ACM MM*, 2020.

[25] H. Du, H. Shi, Y. Liu, J. Wang, Z. Lei, D. Zeng, and T. Mei. Semi-siamese training for shallow face learning. In *ECCV*, 2020.

[26] H. Du, H. Shi, Y. Liu, D. Zeng, and T. Mei. Towards nir-vis masked face recognition. *IEEE SPL*, 2021.

[27] Y. Duan, J. Lu, and J. Zhou. UniformFace: Learning deep equidistributed representation for face recognition. In *CVPR*, 2019.

[28] S. E. Eikenberry, M. Mancuso, E. Iboi, T. Phan, K. Eikenberry, Y. Kuang, E. Kostelich, and A. B. Gumel. To mask or not to mask: Modeling the potential for face mask use by the general public to curtail the COVID-19 pandemic. *Infectious Disease Modelling*, 2020.

[29] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *ACM KDD*, 1996.

[30] M. Geng, P. Peng, Y. Huang, and Y. Tian. Masked face recognition with generative data augmentation and domain constrained ranking. In *ACM MM*, 2020.

[31] S. Gong, X. Liu, and A. K. Jain. Jointly de-biasing face recognition and demographic attribute estimation. In *ECCV*, 2020.

[32] S. Gong, X. Liu, and A. K. Jain. Mitigating face recognition bias via group adaptive classifier. In *CVPR*, 2021.

[33] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He. Accurate, large minibatch SGD: Training imagenet in 1 hour. *arXiv:1706.02677*, 2017.

[34] P. Grother, M. Ngan, and K. Hanaoka. *Face Recognition Vendor Test (FRVT): Part 3, Demographic Effects*. National Institute of Standards and Technology, 2019.

[35] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao. MS-Celeb-1M: A dataset and benchmark for large-scale face recognition. In *ECCV*, 2016.

[36] W. Hariri. Efficient masked face recognition method during the COVID-19 pandemic. *arXiv:2105.03026*, 2021.

[37] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[38] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv:1704.04861*, 2017.

[39] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.

[40] W. Hu, Y. Huang, F. Zhang, and R. Li. Noise-tolerant paradigm for training face recognition cnns. In *CVPR*, 2019.

[41] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical report, 2007.

[42] Y. Huang, Y. Wang, Y. Tai, X. Liu, P. Shen, S. Li, J. Li, and F. Huang. CurricularFace: adaptive curriculum learning loss for deep face recognition. In *CVPR*, 2020.

[43] Z. Huang, J. Zhang, and H. Shan. When age-invariant face recognition meets face age synthesis: A multi-task learning framework. In *CVPR*, 2021.

[44] I. Kemelmacher-Shlizerman, S. M. Seitz, D. Miller, and E. Brossard. The MegaFace benchmark: 1 million faces for recognition at scale. In *CVPR*, 2016.

[45] Y. Kim, W. Park, M.-C. Roh, and J. Shin. GroupFace: Learning latent groups and constructing group-based representations for face recognition. In *CVPR*, 2020.

[46] Y. Kim, W. Park, and J. Shin. BroadFace: Looking at tens of thousands of people at once for face recognition. In *ECCV*, 2020.

[47] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *NeurIPS*, 2012.

[48] S. Kwon, A. D. Joshi, C.-H. Lo, D. A. Drew, L. H. Nguyen, C.-G. Guo, W. Ma, R. S. Mehta, F. M. Shebl, E. T. Warner, et al. Association of social distancing and face mask use with risk of covid-19. *Nature Communications*, 2021.

[49] B. Li, T. Xi, G. Zhang, H. Feng, J. Han, J. Liu, E. Ding, and W. Liu. Dynamic class queue for large scale face recognition in the wild. In *CVPR*, 2021.

[50] P. Li, B. Wang, and L. Zhang. Virtual fully-connected layer: Training a large-scale face recognition dataset with limited computational resources. In *CVPR*, 2021.

[51] B. Liu, W. Deng, Y. Zhong, M. Wang, J. Hu, X. Tao, and Y. Huang. Fair loss: Margin-aware reinforcement learning for deep face recognition. In *ICCV*, 2019.

[52] H. Liu, X. Zhu, Z. Lei, and S. Z. Li. Adaptiveface: Adaptive margin and sampling for face recognition. In *CVPR*, 2019.

[53] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphreface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017.

[54] Y. Liu, P. Shi, B. Peng, H. Yan, Y. Zhou, B. Han, Y. Zheng, C. Lin, J. Jiang, and Y. Fan. IQIYI-VID: A large dataset for multi-modal person identification. *arXiv:1811.07548*, 2018.

[55] S. Lloyd. Least squares quantization in PCM. *TIT*, 1982.

- [56] B. Maze, J. Adams, J. A. Duncan, N. Kalka, T. Miller, C. Otto, A. K. Jain, W. T. Niggel, J. Anderson, and J. Cheney. IARPA Janus Benchmark C: Face dataset and protocol. In *ICB*, 2018.
- [57] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotsia, and S. Zafeiriou. AgeDB: The first manually collected in-the-wild age database. In *CVPR Workshop*, 2017.
- [58] A. Nech and I. Kemelmacher-Shlizerman. Level playing field for million scale face recognition. In *CVPR*, 2017.
- [59] A. Negi, K. Kumar, P. Chauhan, and R. Rajput. Deep neural architecture for face mask detection on simulated masked face dataset against COVID-19 pandemic. In *ICCCIS*, 2021.
- [60] H.-W. Ng and S. Winkler. A data-driven approach to cleaning large face datasets. In *ICIP*, 2014.
- [61] M. L. Ngan, P. J. Grother, K. K. Hanaoka, et al. Ongoing Face Recognition Vendor Test (FRVT) Part 6B: Face recognition accuracy with face masks using post-COVID-19 algorithms. 2021.
- [62] C. Otto, D. Wang, and A. K. Jain. Clustering millions of faces by identity. *TPAMI*, 2017.
- [63] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *BMVC*, 2015.
- [64] I. Radosavovic, R. P. Kosaraju, R. Girshick, K. He, and P. Dollár. Designing network design spaces. In *CVPR*, 2020.
- [65] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. In *IJCV*, 2015.
- [66] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015.
- [67] S. Sengupta, J.-C. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs. Frontal to profile face verification in the wild. In *WACV*, 2016.
- [68] Y. Shi and A. K. Jain. Probabilistic face embeddings. In *ICCV*, 2019.
- [69] Y. Shi, X. Yu, K. Sohn, M. Chandraker, and A. K. Jain. Towards universal representation learning for deep face recognition. In *CVPR*, 2020.
- [70] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.
- [71] K. Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NeurIPS*, 2016.
- [72] L. Song, D. Gong, Z. Li, C. Liu, and W. Liu. Occlusion robust face recognition based on mask learning with pairwise differential siamese network. In *ICCV*, 2019.
- [73] Y. Sun, Y. Chen, X. Wang, and X. Tang. Deep learning face representation by joint identification-verification. In *NeurIPS*, 2014.
- [74] Y. Sun, D. Liang, X. Wang, and X. Tang. DeepID3: Face recognition with very deep neural networks. *arXiv:1502.00873*, 2015.
- [75] Y. Sun, X. Wang, and X. Tang. Deep learning face representation from predicting 10,000 classes. In *CVPR*, 2014.
- [76] Y. Sun, X. Wang, and X. Tang. Deeply learned face representations are sparse, selective, and robust. In *CVPR*, 2015.
- [77] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [78] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *CVPR*, 2014.
- [79] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Web-scale training for face identification. In *CVPR*, 2015.
- [80] M. Tan and Q. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019.
- [81] P. Terhörst, J. N. Kolf, M. Huber, F. Kirchbuchner, N. Damer, A. Morales, J. Fierrez, and A. Kuijper. A comprehensive study on face recognition biases beyond demographics. *arXiv:2103.01592*, 2021.
- [82] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li. YFCC100M: The new data in multimedia research. *Communications of the ACM*, 2016.
- [83] F. Wang, L. Chen, C. Li, S. Huang, Y. Chen, C. Qian, and C. Change Loy. The devil of face recognition is in the noise. In *ECCV*, 2018.
- [84] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang. Residual attention network for image classification. In *CVPR*, 2017.
- [85] F. Wang, W. Liu, H. Liu, and J. Cheng. Additive margin softmax for face verification. *IEEE SPL*, 2018.
- [86] F. Wang, X. Xiang, J. Cheng, and A. L. Yuille. Normface: L2 hypersphere embedding for face verification. In *ACM MM*, 2017.
- [87] G. Wang, J. Ma, Q. Zhang, J. Lu, and J. Zhou. Pseudo facial generation with extreme poses for face recognition. In *CVPR*, pages 1994–2003, 2021.
- [88] H. Wang, D. Gong, Z. Li, and W. Liu. Decorrelated adversarial learning for age-invariant face recognition. In *CVPR*, 2019.
- [89] H. Wang, Y. Wang, Z. Zhou, X. Ji, Z. Li, D. Gong, J. Zhou, and W. Liu. CosFace: Large margin cosine loss for deep face recognition. In *CVPR*, 2018.
- [90] M. Wang and W. Deng. Mitigate bias in face recognition using skewness-aware reinforcement learning. *arXiv:1911.10692*, 2019.
- [91] M. Wang, W. Deng, J. Hu, X. Tao, and Y. Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *CVPR*, 2019.
- [92] Q. Wang, T. Wu, H. Zheng, and G. Guo. Hierarchical pyramid diverse attention networks for face recognition. In *CVPR*, 2020.
- [93] X. Wang, S. Wang, J. Wang, H. Shi, and T. Mei. Co-mining: Deep face recognition with noisy labels. In *ICCV*, 2019.
- [94] X. Wang, S. Zhang, S. Wang, T. Fu, H. Shi, and T. Mei. Misclassified vector guided softmax loss for face recognition. In *AAAI*, 2020.
- [95] Z. Wang, G. Wang, B. Huang, Z. Xiong, Q. Hong, H. Wu, P. Yi, K. Jiang, N. Wang, Y. Pei, et al. Masked face recognition dataset and application. *arXiv:2003.09093*, 2020.
- [96] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, 2016.
- [97] L. Wolf, T. Hassner, and I. Maoz. Face recognition in unconstrained videos with matched background similarity. In *CVPR*, 2011.
- [98] X. Wu, R. He, Z. Sun, and T. Tan. A light CNN for deep face representation with noisy labels. *TIFS*, 2018.
- [99] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le. Self-training with noisy student improves Imagenet classification. In *CVPR*, 2020.
- [100] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017.
- [101] X. Xu, Y. Huang, P. Shen, S. Li, J. Li, F. Huang, Y. Li, and Z. Cui. Consistent instance false positive improves fairness in face recognition. In *CVPR*, 2021.
- [102] I. Z. Yalniz, H. Jégou, K. Chen, M. Paluri, and D. Mahajan. Billion-scale semi-supervised learning for image classification. *arXiv:1905.00546*, 2019.
- [103] L. Yang, D. Chen, X. Zhan, R. Zhao, C. C. Loy, and D. Lin. Learning to cluster faces via confidence and connectivity estimation. In *CVPR*, 2020.
- [104] L. Yang, X. Zhan, D. Chen, J. Yan, C. C. Loy, and D. Lin. Learning to cluster faces on an affinity graph. In *CVPR*, 2019.
- [105] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Learning face representation from scratch. *arXiv:1411.7923*, 2014.
- [106] X. Zhang, L. Yang, J. Yan, and D. Lin. Accelerated training for massive classification via dynamic class selection. In *AAAI*, 2018.
- [107] X. Zhang, R. Zhao, Y. Qiao, X. Wang, and H. Li. Adacos: Adaptively scaling cosine logits for effectively learning deep face representations. In *CVPR*, 2019.
- [108] X. Zhang, R. Zhao, J. Yan, M. Gao, Y. Qiao, X. Wang, and H. Li. P2sgrad: Refined gradients for optimizing deep face models. In *CVPR*, 2019.
- [109] Y. Zhang, W. Deng, M. Wang, J. Hu, X. Li, D. Zhao, and D. Wen. Global-local GCN: Large-scale label noise cleansing for face recognition. In *CVPR*, 2020.
- [110] J. Zhao, Y. Cheng, Y. Xu, L. Xiong, J. Li, F. Zhao, K. Jayashree, S. Pranata, S. Shen, J. Xing, et al. Towards pose invariant face recognition in the wild. In *CVPR*, 2018.
- [111] J. Zhao, S. Yan, and J. Feng. Towards age-invariant face recognition. *TPAMI*, 2020.
- [112] K. Zhao, J. Xu, and M.-M. Cheng. RegularFace: Deep face recognition via exclusive regularization. In *CVPR*, 2019.
- [113] T. Zheng and W. Deng. Cross-pose LFW: A database for studying cross-pose face recognition in unconstrained environments. Technical report, 2018.
- [114] T. Zheng, W. Deng, and J. Hu. Cross-age LFW: A database for studying cross-age face recognition in unconstrained environments. *arXiv:1708.08197*, 2017.
- [115] W. Zheng, L. Yan, F.-Y. Wang, and C. Gou. Learning from the web: Webly supervised meta-learning for masked face recognition. In *CVPR Workshop*, 2021.
- [116] Y. Zhong, W. Deng, M. Wang, J. Hu, J. Peng, X. Tao, and Y. Huang. Unequal-training for deep face recognition with long-tailed noisy data. In *CVPR*, 2019.
- [117] Z. Zhu, G. Huang, J. Deng, Y. Ye, J. Huang, X. Chen, J. Zhu, T. Yang, J. Guo, J. Lu, et al. Masked face recognition challenge: The WebFace260M track report. *arXiv preprint arXiv:2108.07189*, 2021.
- [118] Z. Zhu, G. Huang, J. Deng, Y. Ye, J. Huang, X. Chen, J. Zhu, T. Yang, J. Lu, D. Du, and J. Zhou. WebFace260M: A benchmark unveiling the power of million-scale deep face recognition. In *CVPR*, 2021.





**Zheng Zhu** is currently a postdoctoral fellow at Tsinghua University. Before that, he received Ph.D. degree from Institute of Automation, Chinese Academy of Sciences in 2019. He served reviewers in various journals and conferences including TPAMI, TIP, TMM, TCSVT, CVPR, ICCV, ECCV, ICLR. He has co-authored more than 40 journal and conference papers mainly on computer vision and robotics problems, such as face recognition, visual tracking, human pose estimation, and servo control. He has more than 3,000 Google Scholar citations to his work. He organized the Masked Face Recognition Challenge & Workshop in ICCV 2021. He ranked 1st on NIST-FRVT Masked Face Recognition, won the COCO Keypoint Detection Challenge in ECCV 2020 and Visual Object Tracking (VOT) Real-Time Challenge in ECCV 2018. He is a member of the IEEE.



**Junjie Huang** is a researcher in XForwardAI Technology Co., LTD. Before that, he received master degree from Institute of Automation, Chinese Academy of Sciences in 2019. He served reviewers in journals and conferences including TIP, AAAI. He has co-authored up to 17 journal and conference papers mainly on computer vision and robotics problems, such as human pose estimation, face recognition, object detection, panoptic segmentation, and servo control. He is a member of the IEEE.



**Guan Huang** received the B.E. degree from Huazhong University of Science and Technology in 2013 and master degree from Institute of Automation, Chinese Academy of Sciences (CASIA) in 2016. He is currently the director of algorithm department of XForwardAI Technology Co., LTD. He has published research papers in the areas of computer vision and pattern recognition at international journals and conferences such as CVPR, ICCV, AAAI and TIP. His current research interests include computer vision, deep

learning, object recognition and face recognition.



**Xinze Chen** is currently a senior algorithm engineer at XForwardAI Technology Co., LTD since 2019. Before that he was with Horizon Robotics which joined in 2017. He received the M.E. degree from Institute of Automation, Chinese Academy of Sciences in 2017 and the B.E. degree from Huazhong University of Science and Technology (HUST) in 2014. He has co-authored conference papers in CVPR. His research interests focus on computer vision and deep learning, specifically on the topic of face

recognition, person re-identification, 2D/3D object detection, segmentation, and human pose estimation.




**Jiankang Deng** is a Ph.D. candidate in the Intelligent Behaviour Understanding Group (IBUG) at Imperial College London (ICL), supervised by Stefanos Zafeiriou and funded by the Imperial President's PhD Scholarships. He is in the project of EPSRC FACER2VM (Face Matching for Automatic Identity Retrieval, Recognition, Verification and Management). His Ph.D. research topic is face analysis (face detection, face alignment, face recognition and face generation). He is a

reviewer in prestigious computer vision journals and conferences including TPAMI, IJCV, CVPR, ICCV and ECCV. He is one of the main contributors to the widely used open-source platform Insightface. He is a student member of the IEEE.



**Jiagang Zhu** received his B.E. degree from China University of Petroleum, China in 2015, the Ph.D. degree from Institute of Automation, Chinese Academy of Sciences, China in 2020. Currently, he is an algorithm engineer in XForwardAI Technology Co., LTD. His research interests include deep learning and computer vision.



**Tian Yang** is an algorithm engineer at XForwardAI Technology Co., LTD. Before that, he received a Master of Science and a Master of Professional Studies from the University of Leicester and the University of Auckland sequentially in 2017 and 2019. His work is mainly in the area of computer vision, such as face recognition and gait recognition. He has co-authored publications on CVPR and ICCV.



**Yun Ye** is a research scientist at XForwardAI Technology Co., LTD. He received master degree from Institute of Automation, Chinese Academy of Sciences in 2017. He has co-authored conference papers on CVPR and AAAI. His research interests are in computer vision and deep learning, such as face recognition, object detection, and multi-tasks learning.



**Dalong Du** received the master degree from Institute of Computing Technology, Chinese Academy of Sciences in 2011. He is the vice president of XForwardAI Technology Co., LTD. He has published research papers at international journals and conferences such as CVPR and ICCV. His current research interests include machine learning, robotics, computer vision, and deep learning.



**Jiwen Lu** (M'11-SM'15) received the B.Eng. degree in mechanical engineering and the M.Eng. degree in electrical engineering from the Xian University of Technology, Xi'an, China, in 2003 and 2006, respectively, and the Ph.D. degree in electrical engineering from Nanyang Technological University, Singapore, in 2012. He is currently an Associate Professor with the Department of Automation, Tsinghua University, Beijing, China. His current research interests include computer vision and pattern recognition. He was/is a

member of the Image, Video and Multidimensional Signal Processing Technical Committee, Multimedia Signal Processing Technical Committee and the Information Forensics and Security Technical Committee of the IEEE Signal Processing Society, a member of the Multimedia Systems and Applications Technical Committee and the Visual Signal Processing and Communications Technical Committee of the IEEE Circuits and Systems Society, respectively. He serves as the General Co-Chair for the International Conference on Multimedia and Expo (ICME) 2022, the Program Co-Chair for the International Conference on Multimedia and Expo 2020, the International Conference on Automatic Face and Gesture Recognition (FG) 2023, and the International Conference on Visual Communication and Image Processing (VCIP) 2022. He serves as the Co-Editor-in-Chief for Pattern Recognition Letters, an Associate Editor for the IEEE Transactions on Image Processing, the IEEE Transactions

on Circuits and Systems for Video Technology, the IEEE Transactions on Biometrics, Behavior, and Identity Sciences, and Pattern Recognition. He was a recipient of the National Natural Science Funds for Distinguished Young Scholar. He is a Fellow of IAPR.



**Jie Zhou** (M'01-SM'04) received the BS and MS degrees both from the Department of Mathematics, Nankai University, China, in 1990 and 1992, respectively, and the PhD degree from the Institute of Pattern Recognition and Artificial Intelligence, Huazhong University of Science and Technology (HUST), China, in 1995. From then to 1997, he served as a postdoctoral fellow in the Department of Automation, Tsinghua University, China. Since 2003, he has been a full professor in the Department of Automation, Tsinghua University.

His research interests include computer vision and pattern recognition. In recent years, he has authored more than 300 papers have been published in TPAMI, TIP and CVPR. He is an associate editor for TPAMI and two other journals. He received the National Outstanding Youth Foundation of China Award. He is a senior member of the IEEE and Fellow of the IAPR.