

DeepFace-EMD: Re-ranking Using Patch-wise Earth Mover’s Distance Improves Out-Of-Distribution Face Identification

Hai Phan^{1,2}

pthai1204@gmail.com

¹Auburn University

Anh Nguyen¹

anh.ng8@gmail.com

²Carnegie Mellon University

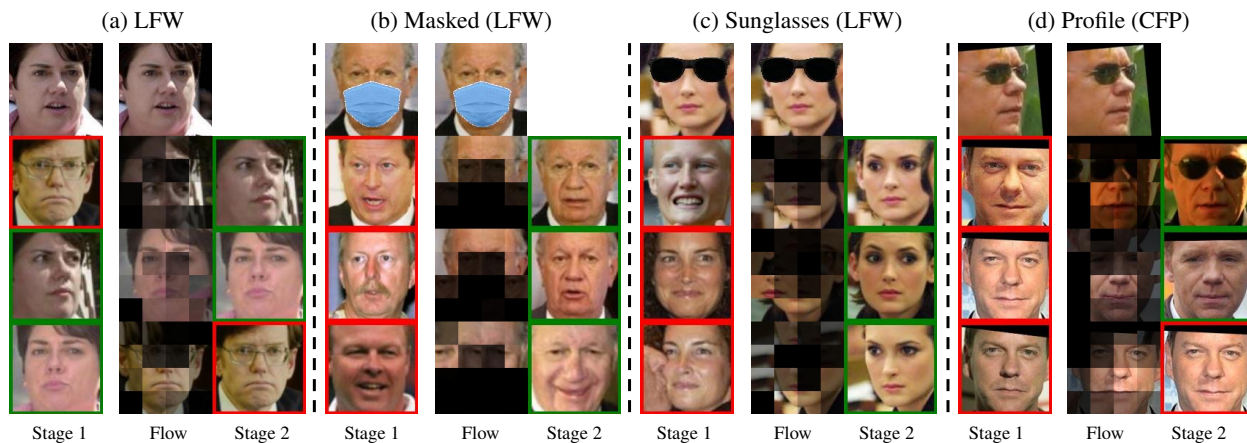


Figure 1. Traditional face identification ranks gallery images based on their cosine distance with the query (top row) at the image-level embedding, which yields large errors upon out-of-distribution changes in the input (e.g. masks or sunglasses; b–d). We find that re-ranking the top- k shortlisted faces from Stage 1 (leftmost column) using their patch-wise EMD similarity w.r.t. the query substantially improves the precision (Stage 2) on challenging cases (b–d). The “Flow” visualization intuitively shows the patch-wise reconstruction of the query face using the most similar patches (i.e. highest flow) from the retrieved face. See Fig. S4 for a full figure with the top-5 candidates.

Abstract

Face identification (FI) is ubiquitous and drives many high-stake decisions made by the law enforcement. A common FI approach compares two images by taking the cosine similarity between their image embeddings. Yet, such approach suffers from poor out-of-distribution (OOD) generalization to new types of images (e.g., when a query face is masked, cropped or rotated) not included in the training set or the gallery. Here, we propose a re-ranking approach that compares two faces using the Earth Mover’s Distance on the deep, spatial features of image patches. Our extra comparison stage explicitly examines image similarity at a fine-grained level (e.g., eyes to eyes) and is more robust to OOD perturbations and occlusions than traditional FI. Interestingly, without finetuning feature extractors, our method consistently improves the accuracy on all tested OOD queries: masked, cropped, rotated, and adversarial

while obtaining similar results on in-distribution images.

1. Introduction

Who shoplifted from the Shinola luxury store in Detroit [4]? Who are you to receive unemployment benefits [5] or board an airplane [1]? Face identification (FI) today is behind the answers to such life-critical questions. Yet, the technology can make errors, leading to severe consequences, e.g. people wrongly denied of unemployment benefits [5] or falsely arrested [2–4,7]. Identifying the person in a single photo remains challenging because, in many cases, the problem is a zero-shot and ill-posed image retrieval task. First, a deep feature extractor may not have seen a normal, non-celebrity person before during training. Second, there may be too few photos of a person in the database for FI systems to make reliable decisions. Third, it is harder to identify when a face in the wild (e.g. from surveillance cameras) is occluded [44, 54] (e.g. wearing masks), distant or

cropped, yielding a new type of photo not in both the training set of deep networks and the retrieval database—i.e., out-of-distribution (OOD) data. For example, face verification accuracy may notoriously drop significantly (from 99.38% to 81.12% on LFW) given an occluded query face (Fig. 1b–d) [44] or adversarial queries [8, 73].

In this paper, we propose to evaluate the performance of state-of-the-art facial feature extractors (ArcFace [19], CosFace [61], and FaceNet [47]) on OOD face *identification* tests. That is, our main task is to recognize the person in a query image given a gallery of known faces. Besides in-distribution (ID) query images, we also test FI models on OOD queries that contain (1) common occlusions, *i.e.* random crops, faces with masks or sunglasses; and (2) adversarial perturbations [73]. Our main findings are:¹

- Interestingly, the OOD accuracy can be substantially improved via a 2-stage approach (see Fig. 2): First, identify a set of the most globally-similar faces from the gallery using cosine distance and then, re-rank these shortlisted candidates by comparing them with the query at the patch-embedding level using the Earth Mover’s Distance (EMD) [45] (Sec. 3 & Sec. 4).
- Across three different models (ArcFace, CosFace, and FaceNet), our re-ranking approach consistently improves the original precision (under all metrics: P@1, R-Precision, and MAP@R) *without finetuning* (Sec. 4). That is, interestingly, the spatial features extracted from these models can be leveraged to compare images patch-wise (in addition to image-wise) and further improve FI accuracy.
- On masked images [59], our re-ranking method (no training) rivals the ArcFace models finetuned directly on masked images (Sec. 4.3).

To our knowledge, our work is the first to demonstrate the remarkable effectiveness of EMD for comparing OOD, occluded and adversarial images at the deep feature level.

2. Methods

2.1. Problem formulation

To demonstrate the generality of our method, we adopt the following simple FI formulation as in [19, 33, 71]: Identify the person in a query image by ranking all gallery images based on their pair-wise similarity with the query. After ranking (Stage 1) or re-ranking (Stage 2), we take the identity of the top-1 nearest image as the predicted identity. **Evaluation** Following [38, 71], we use three common evaluation metrics: Precision@1 (P@1), R-Precision (RP), and MAP@R (M@R). See their definitions in Sec. B1 in [71].

¹Code, demo and data are available at <https://github.com/anguyen8/deepface-emd>

2.2. Networks

Pre-trained models We use three state-of-the-art PyTorch models of ArcFace, FaceNet, and CosFace pre-trained on CASIA [65], VGGFace2 [14], and CASIA, respectively. Their architectures are ResNet-18 [24], Inception-ResNet-v1 [56], and 20-layer SphereFace [33], respectively. See Sec. S1 for more details on network architectures and implementation in PyTorch.

Image pre-processing For all networks, we align and crop input images following the 3D facial alignment in [11] (which uses 5 reference points, 0.7 and 0.6 crop ratios for width and height, and Similarity transformation). All images shown in this paper (*e.g.* Fig. 1) are pre-processed. Using MTCNN, the default pre-processing of all three networks, does not change the results substantially (Sec. S5).

2.3. 2-stage hierarchical face identification

Stage-1: Ranking A common 1-stage face identification [33, 47, 61] ranks gallery images based on their pair-wise cosine similarity with a given query in the last-linear-layer feature space of a pre-trained feature extractor (Fig. 2). Here, our image embeddings are extracted from the *last linear* layer of all three models and are all $\in \mathbb{R}^{512}$.

Stage-2: Re-ranking We re-rank the top- k (where the optimal $k = 100$) candidates from Stage 1 by computing the patch-wise similarity for an image pair using EMD. Overall, we compare faces in **two hierarchical stages** (Fig. 2), first at a coarse, image level and then a fine-grained, patch level.

Via an ablation study (Sec. 3), we find our 2-stage approach (a.k.a. DeepFace-EMD) more accurate than Stage 1 alone (*i.e.* no patch-wise re-ranking) and also Stage 2 alone (*i.e.* sorting the entire gallery using patch-wise similarity).

2.4. Earth Mover’s Distance (EMD)

EMD is an edit distance between two set of weighted objects or distributions [45]. Its effectiveness was first demonstrated in measuring pair-wise image similarity based on color histograms and texture frequencies [45] for image retrieval. Yet, EMD is also an effective distance between two text documents [29], probability distributions (where EMD is equivalent to Wasserstein, *i.e.* Mallows distance) [30], and distributions in many other domains [27, 34, 43]. Here, we propose to harness EMD as a distance between two faces, *i.e.* two sets of weighted facial features.

Let $\mathcal{Q} = \{(q_1, w_{q_1}), \dots, (q_N, w_{q_N})\}$ be a set of N (facial feature, weight) pairs describing a query face where q_i is a feature (*e.g.* left eye or nose) and the corresponding w_{q_i} indicates how important the feature q_i is in FI. The *flow* between \mathcal{Q} and the set of weighted features of a gallery face $\mathcal{G} = \{(g_1, w_{g_1}), \dots, (g_N, w_{g_N})\}$ is any matrix $\mathbf{F} = (f_{ij}) \in \mathbb{R}^{N \times N}$. Intuitively, f_{ij} is the amount of importance weight at q_i that is matched to the weight at g_j . Let d_{ij} be a ground

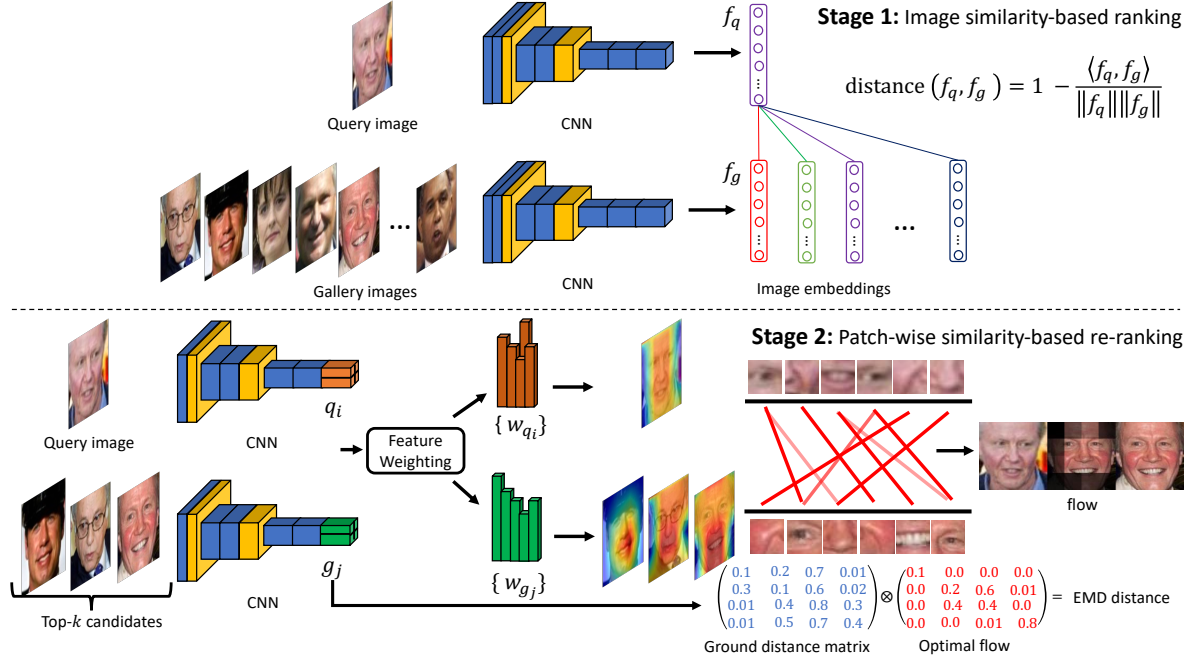


Figure 2. Our 2-stage face identification pipeline. Stage 1 ranks gallery images based on their cosine distance with the query face at the image-embedding level. Stage 2 then re-ranks the top- k shortlisted candidates from Stage 1 using EMD at the patch-embedding level.

distance between (q_i, g_j) and $\mathbf{D} = (d_{ij}) \in \mathbb{R}^{N \times N}$ be the ground distance matrix of all pair-wise distances.

We want to find an optimal flow \mathbf{F} that minimizes the following cost function, *i.e.* the sum of weighted pair-wise distances across the two sets of facial features:

$$\text{COST}(\mathcal{Q}, \mathcal{G}, \mathbf{F}) = \sum_{i=1}^N \sum_{j=1}^N d_{ij} f_{ij} \quad (1)$$

$$\text{s.t.} \quad f_{ij} \geq 0 \quad (2)$$

$$\sum_{j=1}^N f_{ij} \leq w_{q_i}, \text{ and } \sum_{i=1}^N f_{ij} \leq w_{g_j}, i, j \in [1, N] \quad (3)$$

$$\sum_{j=1}^N \sum_{i=1}^N f_{ij} = \min \left(\sum_{j=1}^N w_{g_j}, \sum_{i=1}^N w_{q_i} \right). \quad (4)$$

As in [66, 71], we normalize the weights of a face such that the total weights of features is 1 *i.e.* $\sum_{i=1}^N w_{q_i} = \sum_{j=1}^N w_{g_j} = 1$, which is also the total flow in Eq. (4). Note that EMD is a metric iff two distributions have an equal total weight and the ground distance function is a metric [16].

We use the iterative Sinkhorn algorithm [18] to efficiently solve the linear programming problem in Eq. (1), which yields the final EMD between two faces \mathcal{Q} and \mathcal{G} .

Facial features In image retrieval using EMD, a set of features $\{q_i\}$ can be a collection of dominant colors [45], spatial frequencies [45], or a histogram-like descriptor based on the local patches of reference identities [60]. Inspired

by [60], we also divide an image into a grid but we take the embeddings of the local patches from the last convolutional layers of each network. That is, in FI, face images are aligned and cropped such that the entire face covers most of the image (see Fig. 1a). Therefore, without facial occlusion, every image patch is supposed to contain useful identity information, which is in contrast to natural photos [66].

Our grid sizes $H \times W$ for ArcFace, FaceNet, and CosFace are respectively, 8×8 , 3×3 , and 6×7 , which are the corresponding spatial dimensions of their last convolutional layers (see definitions of these layers in Sec. S1). That is, each feature q_i is an embedding of size $1 \times 1 \times C$ where C is the number of channels (*i.e.* 512, 1792, and 512 for ArcFace, FaceNet, and CosFace, respectively).

Ground distance Like [66, 71], we use cosine distance as the ground distance d_{ij} between the embeddings (q_i, g_j) of two patches:

$$d_{ij} = 1 - \frac{\langle q_i, g_j \rangle}{\|q_i\| \|g_j\|} \quad (5)$$

where $\langle \cdot \rangle$ is the dot product between two feature vectors.

2.5. Feature weighting

EMD in our FI intuitively is an optimal plan to match all weighted features across two images. Therefore, how to weight features is an important step. Here, we thoroughly explore five different feature-weighting techniques for FI.

Uniform Zhang et al. [66] found that it is beneficial to assign lower weight to less informative regions (*e.g.* background or occlusion) and higher weight to discriminative

areas (e.g. those containing salient objects). Yet, assigning an equal weight to all $N = H \times W$ patches is worth testing given that background noise is often cropped out of the pre-processed face image (Fig. 1):

$$w_{q_i} = w_{g_i} = \frac{1}{N}, \text{ where } 1 \leq k \leq N \quad (6)$$

Average Pooling Correlation (APC) Instead of uniformly weighting all patch embeddings, an alternative from [66] would be to weight a given feature q_i proportional to its correlation to the entire other *image* in consideration. That is, the weight w_{q_i} would be the dot product between the feature q_i and the average pooling output of all embeddings $\{g_j\}_1^N$ of the gallery image:

$$w_{q_i} = \max\left(0, \langle q_i, \frac{\sum_j g_j}{N} \rangle\right), w_{g_j} = \max\left(0, \langle g_j, \frac{\sum_i q_i}{N} \rangle\right) \quad (7)$$

where $\max(\cdot)$ keeps the weights always non-negative. APC tends to assign near-zero weight to occluded regions and, interestingly, also minimizes the weight of eyes and mouth in a non-occluded gallery image (see Fig. 3b; blue shades around both the mask and the non-occluded mouth).

Cross Correlation (CC) APC [66] is different from CC introduced in [71], which is the same as APC except that CC uses the output vector from the last linear layer (see `code`) instead of the global average pooling vector in APC.

Spatial Correlation (SC) While both APC and CC “summarize” an entire other gallery image into a vector first, and then compute its correlation with a given patch q_i in the query. In contrast, an alternative, inspired by [53], is to take the sum of the cosine similarity between the query patch q_i and every patch in each gallery image $\{g_j\}_1^N$:

$$w_{q_i} = \max\left(0, \sum_j \frac{\langle q_i, g_j \rangle}{\|q_i\| \|g_j\|}\right), w_{g_j} = \max\left(0, \sum_i \frac{\langle q_i, g_j \rangle}{\|q_i\| \|g_j\|}\right) \quad (8)$$

We observe that SC often assigns a higher weight to occluded regions e.g., masks and sunglasses (Fig. 3b).

Landmarking (LMK) While the previous three techniques adaptively rely on the image-patch similarity (APC, CC) or patch-wise similarity (SC) to weight a given patch embedding, their considered important points may or may not align with facial landmarks, which are known to be important for many face-related tasks. Here, as a baseline for APC, CC, and SC, we use dlib [26] to predict 68 keypoints in each face image (see Fig. 3c) and weight each patch-embedding by the density of the keypoints inside the patch area. Our LMK weight distribution appears Gaussian-like with the peak often right below the nose (Fig. 3c).




Figure 3. The results of assigning weights to 4×4 patches for ArcFace under three different techniques. Based on per-patch density of detected landmarks (---), LMK (c) often assigns higher weight to the center of a face (regardless of occlusions). In contrast, SC and APC assign higher weight to patches of higher patch-wise and patch-vs-image similarity, respectively. APC tends to down-weight a facial feature (e.g. blue regions around sunglasses or mouth) if its corresponding feature is occluded in the other image (b). In contrast, SC is insensitive to occlusions (a). See Fig. S1 & Fig. S2 for more heatmap examples of feature weighting.

3. Ablation Studies

We perform three ablation studies to rigorously evaluate the key design choices in our 2-stage FI approach: (1) Which feature-weighting techniques to use (Sec. 3.1)? (2) re-ranking using both EMD and cosine distance (Sec. 3.2); and (3) comparing patches or images in Stage 1 (Sec. 3.3).

Experiment For all three experiments, we use ArcFace to perform FI on both LFW [65] and LFW-crop. For **LFW**, we take all 1,680 people who have ≥ 2 images for a total of 9,164 images. When taking each image as a query, we search in a gallery of the remaining 9,163 images. For the experiments with **LFW-crop**, we use all 13,233 original LFW images as the gallery. To create a query set of 13,233 cropped images, we clone the gallery and crop each image randomly to its 70% and upsample it back to the original size of 128×128 (see examples in Fig. 5d). That is, LFW-crop tests identifying a cropped (i.e. close-up, and misaligned) image given the unchanged LFW gallery. LFW and LFW-crop tests offer contrast insights (ID vs. OOD).

In Stage 2, i.e. re-ranking the top- k candidates, we test different values of $k \in \{100, 200, 300\}$ and do not find the performance to change substantially. At $k = 100$, our 2-stage precision is already close to the maximum precision of 99.88 under a perfect re-ranking (see Tab. 1a; Max prec.).

3.1. Comparing feature weighting techniques

Here, we evaluate the precision of our 2-stage FI as we sweep across five different feature-weighting techniques and two grid sizes (8×8 and 4×4). In an 8×8 grid, we observe that some facial features such as the eyes are often split in half across two patches (see Fig. S5), which may impair the patch-wise similarity. Therefore, for each weight-

ing technique, we also test average-pooling the 8×8 grid into 4×4 and performing EMD on the resultant 16 patches.

Results First, we find that, on LFW, our image-similarity-based techniques (APC, SC) outperform the LMK baseline (Tab. 1a) despite not using landmarks in the weighting process, verifying the effectiveness of adaptive, similarity-based weighting schemes.

Second, interestingly, in FI, we find that Uniform, APC, and SC all outperform the CC weighting proposed in [66, 71]. This is in stark contrast to the finding in [71] that CC is better than Uniform (perhaps because face images do not have background noise and are close-up). Furthermore, using the global average-pooling vector from the channel (APC) substantially yields more useful spatial similarity than the last-linear-layer output as in CC implementation (Tab. 1b; 96.16 vs. 91.31 P@1).

Third, surprisingly, despite that a patch in a 8×8 grid does not enclose an entire, fully-visible facial feature (e.g. an eye), all feature-weighting methods are on-par or better on an 8×8 grid than on a 4×4 (e.g. Tab. 1b; APC: 96.16 vs. 95.32). Note that the optimal flow visualized in a 4×4 grid is more interpretable to humans than that on a 8×8 grid (compare Fig. 1 vs. Fig. S5).

Fourth, across all variants of feature weighting, our 2-stage approach consistently and *substantially* outperforms the traditional Stage 1 alone on LFW-crop, suggesting its robust effectiveness in handling OOD queries.

Fifth, under a perfect re-ranking of the top- k candidates (where $k = 100$), there is only 1.4% headroom for improvement upon Stage 1 alone in LFW (Tab. 1a; 98.48 vs. 99.88) while there is a large $\sim 12\%$ headroom in LFW-crop (Tab. 1a; 87.35 vs. 98.71). Interestingly, our re-ranking results approach the upperbound re-ranking precision (e.g. Tab. 1b; 96.26 of Uniform vs. 98.71 Max prec. at $k = 100$).

3.2. Re-ranking using both EMD & cosine distance

We observe that for some images, re-ranking using patch-wise similarity at Stage 2 does not help but instead hurt the accuracy. Here, we test whether linearly combining EMD (at the *patch*-level embeddings as in Stage 2) and cosine distance (at the *image*-level embeddings as in Stage 1) may improve *re-ranking* accuracy further (vs. EMD alone).

Experiment We use the grid size of 8×8 , *i.e.* the better setting from the previous ablation study (Sec. 3.1). For each pair of images, we linearly combine their patch-level EMD (θ_{EMD}) and the image-level cosine distance (θ_{Cosine}) as:

$$\theta = \alpha \times \theta_{EMD} + (1 - \alpha) \times \theta_{Cosine} \quad (9)$$

Sweeping across $\alpha \in \{0, 0.3, 0.5, 0.7, 1\}$, we find that changing α has a marginal effect on the P@1 on LFW. That is, the P@1 changes in [95, 98.5] with the lowest accuracy being 95 when EMD is exclusively used, *i.e.* $\alpha = 1$ (see Fig. 4a). In contrast, for LFW-crop, we find the accuracy to

ArcFace	Method	P@1	RP	M@R
LFW vs. LFW (a)	Stage 1 alone [19]	98.48	78.69	78.29
	Max prec. at $k = 100$	99.88	81.32	-
	CC [71] (8×8)	98.42	78.35	77.91
	CC [71] (4×4)	81.69	76.29	72.47
	APC (8×8)	98.60	78.63	78.23
	APC (4×4)	98.54	78.57	78.16
	Uniform (8×8)	98.66	78.73	78.35
	Uniform (4×4)	98.63	78.72	78.33
	SC (8×8)	98.66	78.74	78.35
	SC (4×4)	98.65	78.72	78.33
LFW-crop vs. LFW (b)	LMK (8×8)	98.35	78.43	77.99
	LMK (4×4)	98.31	78.38	77.90
	Stage 1 alone [19]	87.35	71.38	69.04
	Max prec. at $k = 100$	98.71	89.13	-
	CC [71] (8×8)	91.31	72.33	70.00
	CC [71] (4×4)	63.12	56.03	51.00
	APC (8×8)	96.16	76.60	74.57
	APC (4×4)	95.32	75.37	73.25
	Uniform (8×8)	96.26	78.08	76.25
	Uniform (4×4)	95.53	77.15	75.29
SC (8×8)	96.19	78.05	76.20	
SC (4×4)	95.42	77.12	75.25	

Table 1. Comparison of five feature-weighting techniques for ArcFace [19] patch embeddings on LFW [65] and LFW-crop datasets. Performance is often slightly better on a 8×8 grid than on a 4×4 . Our 2-stage approach consistently outperforms the vanilla Stage 1 alone and approaches closely the maximum re-ranking precision at $k = 100$.

monotonically increases as we increase α (Fig. 4b). That is, **the higher the contribution of patch-wise similarity, the better re-ranking** accuracy on the challenging randomly-cropped queries. We choose $\alpha = 0.7$ as the best and default choice for all subsequent FI experiments. Interestingly, our proposed distance (Eq. 9) also yields a state-of-the-art face verification result on MLFW [59] (Sec. S4).




Figure 4. The P@1 of our 2-stage FI when sweeping across α for linearly combining EMD and cosine distance on LFW (a) and LFW-crop images (b) when using APC feature weighting. Trends are similar for all other feature-weighting methods (see Fig. S3).

3.3. Patch-wise EMD for ranking or re-ranking

Given that re-ranking using EMD at the patch-embedding space substantially improves the precision of FI compared to Stage 1 alone (Tab. 1), here, we test performing such patch-wise EMD sorting at Stage 1 instead of Stage 2.

Experiment That is, we test ranking images using EMD at the patch level instead of the standard cosine distance at the image level. Performing patch-wise EMD at Stage 1 is significantly slower than our 2-stage approach, *e.g.*, **~12 times slower** (729.20s vs. 60.97s, in total, for 13,233 queries). That is, Sinkhorn is a slow, iterative optimization method and the EMD at Stage 2 has to sort only $k = 100$ (instead of 13,233) images. In addition, FI by comparing images patch-wise using EMD at Stage 1 yields consistently worse accuracy than our 2-stage method under all feature-weighting techniques (see Tab. S1 for details).

4. Additional Results

To demonstrate the generality and effectiveness of our 2-stage FI, we take the best hyperparameter settings ($\alpha = 0.7$; APC) from the ablation studies (Sec. 3) and use them for three different models (ArcFace [19], CosFace [61], and FaceNet [47]), which have different grid sizes.

We test the three models on five different OOD query types: (1) faces wearing masks or (2) sunglasses; (3) profile faces; (4) randomly cropped faces; and (5) adversarial faces.

4.1. Identifying occluded faces

Experiment We perform our 2-stage FI on three datasets: CFP [48], CALFW [72], and AgeDB [37]. 12,173-image CALFW and 16,488-image AgeDB have age-varying images of 4,025 and 568 identities, respectively. CFP has 500 people, each having 14 images (10 frontal and 4 profile).

To test our models on challenging OOD queries, in CFP, we use its 2,000 profile faces in CFP as queries and its 5,000 frontal faces as the gallery. To create OOD queries using CFP², CALFW, and AgeDB, we automatically occlude all images with masks and sunglasses by detecting the landmarks of eyes and mouth using dlib and overlaying black sunglasses or a mask on the faces (see examples in Fig. 1). We also take these three datasets and create randomly cropped queries (as for LFW-crop in Sec. 3). For all datasets, we test identifying occluded query faces given the original, unmodified gallery. That is, for every query, there is ≥ 1 matching gallery image.

Results First, for all three models and all occlusion types, *i.e.* due to masks, sunglasses, crop, and self-occlusion (profile queries in CFP), **our method consistently outperforms the traditional Stage 1 alone** approach under all three precision metrics (Tables 2, S8, & S4).

Second, across all three datasets, we find the **largest improvement** that our Stage 2 provides upon the Stage 1 alone is when the queries are **randomly cropped or masked** (Tab. 2). In some cases, the Stage 1 alone using cosine distance is not able to retrieve any relevant examples among

Dataset	Model	Method	P@1	RP	M@R
CALFW (Mask)	ArcFace	ST1	96.81	53.13	51.70
		Ours	99.92	57.27	56.33
	CosFace	ST1	98.54	43.46	41.20
		Ours	99.96	59.85	58.87
	FaceNet	ST1	77.63	39.74	36.93
		Ours	96.67	45.87	44.53
CALFW (Sunglass)	ArcFace	ST1	51.11	29.38	26.73
		Ours	54.95	30.66	27.74
	CosFace	ST1	45.20	25.93	22.78
		Ours	49.67	26.98	24.12
	FaceNet	ST1	21.68	13.70	10.89
		Ours	25.07	15.04	12.16
CALFW (Crop)	ArcFace	ST1	79.13	43.46	41.20
		Ours	92.57	47.17	45.68
	CosFace	ST1	10.99	6.45	5.43
		Ours	25.99	12.35	11.13
	FaceNet	ST1	79.47	44.40	41.99
		Ours	85.71	45.91	43.83
AgeDB (Mask)	ArcFace	ST1	96.15	39.22	30.41
		Ours	99.84	39.22	33.18
	CosFace	ST1	98.31	38.17	31.57
		Ours	99.95	39.70	33.68
	FaceNet	ST1	75.99	22.28	14.95
		Ours	96.53	24.25	17.49
AgeDB (Sunglass)	ArcFace	ST1	84.64	51.16	44.99
		Ours	87.06	50.40	44.27
	CosFace	ST1	68.93	34.90	27.30
		Ours	75.97	35.54	28.12
	FaceNet	ST1	56.77	27.92	20.00
		Ours	61.21	28.98	21.11
AgeDB (Crop)	ArcFace	ST1	79.92	32.66	26.19
		Ours	92.92	32.93	26.60
	CosFace	ST1	10.11	4.23	2.18
		Ours	19.58	4.95	2.76
	FaceNet	ST1	80.80	31.50	24.27
		Ours	86.74	31.51	24.32

Table 2. When the queries (from CALFW [72] and AgeDB [37]) are occluded by masks, sunglasses, or random cropping, our 2-stage method (8×8 grid; APC) is substantially more robust to the Stage 1 alone baseline (ST1) with up to +13% absolute gain (*e.g.* P@1: 79.13 to 92.57). The conclusions are similar for other feature-weighting methods (see Tab. S2 and Tab. S3).

the top-5 but our re-ranking manages to push three relevant faces into the top-5 (Fig. 5d).

Third, we observe that for faces with masks or sunglasses, APC interestingly often excludes the mouth or eye regions from the fully-visible gallery faces when computing the EMD patch-wise similarity with the corresponding occluded query (Fig. 3). The same observation can be seen in the visualizations of the most similar patch pairs, *i.e.* highest flow, for our same 2-stage approach that uses either 4×4 grids (Fig. 5 and Fig. 1) or 8×8 grids (Fig. S5).

4.2. Robustness to adversarial images

Adversarial examples pose a huge challenge and a serious security threat to computer vision systems [28, 39] including FI [50, 73]. Recent research suggests that the

²We only apply masks and sunglasses on the frontal images of CFP.

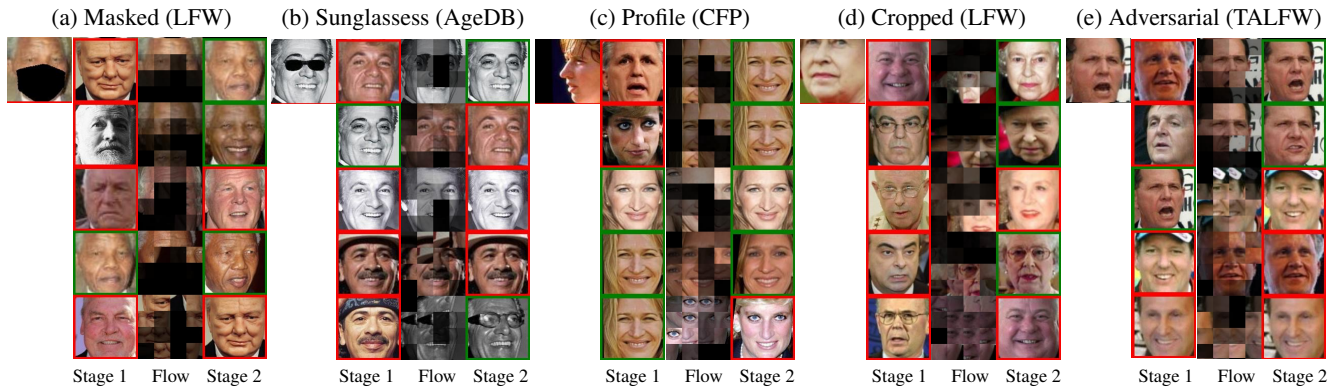


Figure 5. Figure in a similar format to that of Fig. 1. Our re-ranking based on patch-wise similarity using ArcFace (4×4 grid; APC) pushes more relevant gallery images higher up (here, we show top-5 results), improving face identification precision under various types of occlusions. The “Flow” visualization intuitively shows the patch-wise reconstruction of the query (top-left) given the highest-correspondence patches (*i.e.* largest flow) from a gallery face. The darker a patch, the lower the flow. For example, despite being masked out $\sim 50\%$ of the face (a), Nelson Mandela can be correctly retrieved as Stage 2 finds gallery faces with similar forehead patches. See Fig. S5 for a similar figure as the results of running our method with an 8×8 grid (*i.e.* smaller patches), which yields slightly better precision (Tab. 1).

Dataset	Model	Method	P@1	RP	M@R
TALFW [73] vs. LFW [65]	ArcFace	ST1	93.49	81.04	80.35
		Ours	96.64	82.72	82.10
	CosFace	ST1	96.49	83.57	82.99
		Ours	99.07	85.48	85.03
	FaceNet	ST1	95.33	79.24	78.19
		Ours	97.26	80.33	79.39

Table 3. Our re-ranking (8×8 grid; APC) consistently improves the precision over Stage 1 alone (ST1) when identifying adversarial TALFW [73] images given an in-distribution LFW [65] gallery. The conclusions also carry over to other feature-weighting methods (more results in Tab. S5).

patch representation may be the key behind ViT impressive robustness to adversarial images [9, 35, 49]. Motivated by these findings, we test our 2-stage FI on TALFW [73] queries given an original 13,233-image LFW gallery.

Experiment TALFW contains 4,069 LFW images perturbed adversarially to cause face verifiers to mislabel [73].

Results Over the entire TALFW query set, we find our re-ranking to consistently outperform the Stage 1 alone under all three metrics (Tab. 3). Interestingly, the improvement (of ~ 2 to 4 points under P@1 for three models) is larger than when tested on the original LFW queries (around 0.12 in Tab. 1a), verifying our patch-based re-ranking robustness when queries are perturbed with very small noise. That is, our approach can improve FI precision when the perturbation size is either small (adversarial) or large (*e.g.* masks).

4.3. Re-ranking rivals finetuning on masked images

While our approach does not involve re-training, a common technique for improving FI robustness to occlusion is data augmentation, *i.e.* re-train the models on occluded

data in addition to the original data. Here, we compare our method with data augmentation on masked images.

Experiment To generate augmented, masked images, we follow [10] to overlay various types of masks on CASIA images to generate $\sim 415K$ masked images. We add these images to the original CASIA training set, resulting in a total of $\sim 907K$ images (10,575 identities). We finetune ArcFace on this dataset with the same original hyperparameters [6] (see Sec. S2). We train three models and report the mean and standard deviation (Tab. 4).

For a fair comparison, we evaluate the finetuned models and our no-training approach on the MLFW dataset [59], instead of our self-created masked datasets. That is, the query set has 11,959 MLFW masked-face images and the gallery is the entire 13,233-image LFW.

Results First, we find that finetuning ArcFace improves its accuracy in FI under Stage 1 alone (Tab. 4; 39.79 vs. 41.64). Yet, our 2-stage approach still substantially outperforms Stage 1 alone, both when using the original and the finetuned ArcFace (Tab. 4; 48.23 vs. 41.64). Interestingly, we also test using the finetuned model in our DeepFace-EMD framework and finds it to approach closely the best no-training result (46.21 vs. 48.23).

5. Related Work

Face Identification under Occlusion Partial occlusion presents a significant, ill-posed challenge to face identification as the AI has to rely only on incomplete or noisy facial features to make decisions [44]. Most prior methods propose to improve FI robustness by augmenting the training set of deep feature extractors with partially-occluded faces [23, 41, 57, 59, 63, 63]. Training on augmented, occluded data encourages models to rely more on local, dis-

ArcFace	Method	P@1	RP	M@R
Pre-trained	(a) ST1	39.79	35.10	33.32
	(b) Ours	48.23	41.43	39.71
Finetuned	(c) ST1	41.64± 0.16	34.67± 0.24	32.66± 0.25
	(d) Ours	46.21 ± 0.27	38.65 ± 0.26	36.73± 0.26

Table 4. Our 2-stage approach (b) using ArcFace (8×8 grid; APC) substantially outperforms Stage 1 alone (a) on identifying masked images of MLFW given the unmasked gallery of LFW. Interestingly, our method (b) also outperforms Stage 1 alone when ArcFace has been finetuned on masked images (c). In (c), we report the mean and std over three finetuned models.

criminative facial features [41]; however, does not prevent FI models from misbehaving on new OOD occlusion types, especially under adversarial scenarios [50]. In contrast, our approach (1) does not require re-training or data augmentation; and (2) harnesses both image-level features (stage 1) and local, patch-level features (stage 2) for FI.

A common alternative is to learn to generate a spatial feature mask [36, 44, 52, 58] or an attention map [63] to *exclude* the occluded (*i.e.* uninformative or noisy) regions in the input image from the face matching process. Motivated by these works, we tested five methods for inferring the importance of each image patch (Sec. 3) for EMD computation. Early works used hand-crafted features and obtained limited accuracy [32, 36, 40]. In contrast, the latter attempts took advantage of deep architectures but requires a separate occlusion detector [52] or a masking subnetwork in a custom architecture trained end-to-end [44, 58]. In contrast, we leverage directly the pre-trained state-of-the-art image embeddings (of ArcFace, CosFace, & FaceNet) and EMD to exclude the occluded regions from an input image *without any* architectural modifications or re-training.

Another approach is to predict occluded pixels and then perform FI on the recovered images [25, 31, 62, 64, 70, 75]. Yet, how to recover a non-occluded face while preserving true identity remains a challenge to state-of-the-art GAN-based de-occlusion methods [13, 20, 22].

Re-ranking in Face Identification Re-ranking is a popular 2-stage method for refining image retrieval results [69] in many domains, *e.g.* person re-identification [46], localization [51], or web image search [17]. In FI, Zhou et. al. [74] used hand-crafted patch-level features to encode an image for ranking and then used multiple reference images in the database to re-rank each top- k candidate. The social context between two identities has also been found to be useful in re-ranking photo-tagging results [12]. Swearingen et. al. [55] found that harnessing an external “disambiguator” network trained to separate a query from lookalikes is an effective re-ranking method. In contrast to the prior work, we do not use extra images [74] or external knowledge [12]. Compared to face re-ranking [21, 42], our method is the first re-rank candidates based on a pair-wise similarity score

computed from both the image-level and patch-level similarity computed off of state-of-the-art deep facial features.

EMD for Image Retrieval While EMD is a well-known metric in image retrieval [45], its applications on *deep* convolutional features of images have been relatively under-explored. Zhang et al. [66, 67] recently found that classifying fine-grained images (of dogs, birds, and cars) by comparing them patch-wise using EMD in a deep feature space improves few-shot fine-grained classification accuracy. Yet, their success has been limited to *few-shot*, 5-way and 10-way classification with smaller networks (ResNet-12 [24]). In contrast, here, we demonstrate a substantial improvement in FI using EMD without re-training the feature extractors.

Concurrent to our work, Zhao et al. [71] proposes DIML, which exhibits consistent improvement of $\sim 2\text{--}3\%$ in image retrieval on images of birds, cars, and products by using the sum of cosine distance and EMD as a “structural similarity” score for ranking. They found that CC is more effective than assigning uniform weights to image patches [70]. Interestingly, via a rigorous study into different feature-weighting techniques, we find novel insights specific for FI: Uniform weighting is more effective than CC. Unlike prior EMD works [60, 66, 67, 71], ours is **the first** to show the significant effectiveness of EMD on (1) occluded and adversarial OOD images; and (2) on face identification.

6. Discussion and Conclusion

Limitations Solving patch-wise EMD via Sinkhorn is slow, which may prohibit it from being used to sort a much larger image sets (see run-time reports in Tab. S1). Furthermore, here, we used EMD on two distributions of equal weights; however, the algorithm can be used for unequal-weight cases [16, 45], which may be beneficial for handling occlusions. While substantially improving FI accuracy under the four occlusion types (*i.e.*, masks, sunglasses, random crops, and adversarial images), re-ranking is only marginally better than Stage 1 alone on ID and profile faces, which is interesting to understand deeper in future research.

Instead of using pre-trained models, it might be interesting to re-train new models explicitly on patch-wise correspondence tasks, which may yield better patch embeddings for our re-ranking. In sum, we propose DeepFace-EMD, a 2-stage approach for comparing images hierarchically: First at the image level and then at the patch level. DeepFace-EMD shows impressive robustness to occluded and adversarial faces and can be easily integrated into existing FI systems in the wild.

Acknowledgement We thank Qi Li, Peijie Chen, and Giang Nguyen for their feedback on manuscript. We also thank Chi Zhang, Wenliang Zhao, Chengrui Wang for releasing their DeepEMD, DIML, and MLFW code, respectively. AN was supported by the NSF Grant No. 1850117 and a donation from NaphCare Foundation.

References

- [1] Facial recognition tech at hartsfield-jackson wins over most international delta customers - atlanta business chronicle. <https://www.bizjournals.com/atlanta/news/2019/06/25/facial-recognition-tech-at-hartsfield-jackson-wins.html>. (Accessed on 11/09/2021). 1
- [2] Flawed facial recognition leads to arrest and jail for new jersey man - the new york times. <https://www.nytimes.com/2020/12/29/technology/facial-recognition-misidentify-jail.html>. (Accessed on 11/09/2021). 1
- [3] Michigan man wrongfully accused with facial recognition urges congress to act. <https://www.detroitnews.com/story/news/politics/2021/07/13/house-panel-hear-michigan-man-wrongfully-accused-facial-recognition/7948908002/>. (Accessed on 11/09/2021). 1
- [4] The new lawsuit that shows facial recognition is officially a civil rights issue — mit technology review. <https://www.technologyreview.com/2021/04/14/1022676/robert-williams-facial-recognition-lawsuit-aclu-detroit-police/>. (Accessed on 04/15/2021). 1
- [5] The pandemic is testing the limits of face recognition — mit technology review. <https://www.technologyreview.com/2021/09/28/1036279/pandemic-unemployment-government-face-recognition/>. (Accessed on 11/09/2021). 1
- [6] ronghuaiyang/arcface-pytorch. <https://github.com/ronghuaiyang/arcface-pytorch>. (Accessed on 11/16/2021). 7
- [7] Wrongfully arrested man sues detroit police following false facial-recognition match - the washington post. <https://www.washingtonpost.com/technology/2021/04/13/facial-recognition-false-arrest-lawsuit/>. (Accessed on 11/09/2021). 1
- [8] Brandon Amos, Bartosz Ludwiczuk, and Mahadev Satyanarayanan. Openface: A general-purpose face recognition library with mobile applications. Technical report, CMU-CS-16-118, CMU School of Computer Science, 2016. 2
- [9] Anonymous. Patches are all you need? In *Submitted to The Tenth International Conference on Learning Representations*, 2022. under review. 7
- [10] Aqeel Anwar and Arijit Raychowdhury. Masked face recognition for secure authentication. *ArXiv*, abs/2008.11104, 2020. 7, 21
- [11] Chandrasekhar Bhagavatula, Chenchen Zhu, Khoa Luu, and Marios Savvides. Faster than real-time facial alignment: A 3d spatial transformer network approach in unconstrained poses. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4000–4009, 2017. 2
- [12] Samarth Bharadwaj, Mayank Vatsa, and Richa Singh. Aiding face recognition with social context association rule based re-ranking. In *IEEE International Joint Conference on Biometrics*, pages 1–8. IEEE, 2014. 8
- [13] Jiancheng Cai, Hu Han, Jiyun Cui, Jie Chen, Li Liu, and S Kevin Zhou. Semi-supervised natural face de-occlusion. *IEEE Transactions on Information Forensics and Security*, 16:1044–1057, 2020. 8
- [14] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE International Conference on Automatic Face Gesture Recognition (FG 2018)*, pages 67–74, 2018. 2
- [15] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 67–74, 2018. 12
- [16] Scott Cohen. *Finding color and shape patterns in images*. stanford university, 1999. 3, 8
- [17] Jingyu Cui, Fang Wen, and Xiaou Tang. Real time google and live image search re-ranking. In *Proceedings of the 16th ACM international conference on Multimedia*, pages 729–732, 2008. 8
- [18] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2013. 3
- [19] Jiankang Deng, Jia Guo, Xue Niannan, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 2, 5, 6, 12, 13
- [20] Jiayuan Dong, Liyan Zhang, Hanwang Zhang, and Weichen Liu. Occlusion-aware gan for face de-occlusion in the wild. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2020. 8
- [21] Xingbo Dong, Soohyung Kim, Zhe Jin, Jung Yeon Hwang, Sangrae Cho, and A.B.J. Teoh. Open-set face identification with index-of-max hashing by learning. *Pattern Recognition*, 103:107277, 2020. 8
- [22] Shiming Ge, Chenyu Li, Shengwei Zhao, and Dan Zeng. Occluded face recognition in the wild by identity-diversity inpainting. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(10):3387–3397, 2020. 8
- [23] Jianzhu Guo, Xiangyu Zhu, Zhen Lei, and Stan Z Li. Face synthesis for eyeglass-robust face recognition. In *Chinese Conference on biometric recognition*, pages 275–284. Springer, 2018. 7
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 2, 8, 12
- [25] Ran He, Wei-Shi Zheng, Bao-Gang Hu, and Xiang-Wei Kong. A regularized correntropy framework for robust pattern recognition. *Neural computation*, 23(8):2074–2100, 2011. 8
- [26] Davis E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009. 4
- [27] Sachin Kumar, Soumen Chakrabarti, and Shourya Roy. Earth mover’s distance pooling over siamese lstms for automatic short answer grading. In *IJCAI*, pages 2046–2052, 2017. 2

- [28] Alexey Kurakin, Ian Goodfellow, Samy Bengio, et al. Adversarial examples in the physical world, 2016. [6](#)
- [29] Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. From word embeddings to document distances. In *International conference on machine learning*, pages 957–966. PMLR, 2015. [2](#)
- [30] Elizaveta Levina and Peter Bickel. The earth mover’s distance is the mallows distance: Some insights from statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 251–256. IEEE, 2001. [2](#)
- [31] Xiao-Xin Li, Dao-Qing Dai, Xiao-Fei Zhang, and Chuan-Xian Ren. Structured sparse error coding for face recognition with occlusion. *IEEE transactions on image processing*, 22(5):1889–1900, 2013. [8](#)
- [32] Zhifeng Li, Wei Liu, Dahua Lin, and Xiaoou Tang. Non-parametric subspace analysis for face recognition. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 2, pages 961–966. IEEE, 2005. [8](#)
- [33] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheraface: Deep hypersphere embedding for face recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [2](#), [12](#), [14](#)
- [34] Noam Lupu, Lucía Selios, and Zach Warner. A new measure of congruence: The earth mover’s distance. *Political Analysis*, 25(1):95–113, 2017. [2](#)
- [35] Kaleel Mahmood, Rigel Mahmood, and Marten Van Dijk. On the robustness of vision transformers to adversarial examples. *arXiv preprint arXiv:2104.02610*, 2021. [7](#)
- [36] Rui Min, Abdenour Hadid, and Jean-Luc Dugelay. Improving the recognition of faces occluded by facial accessories. In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 442–447. IEEE, 2011. [8](#)
- [37] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: the first manually collected, in-the-wild age database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop*, volume 2, page 5, 2017. [6](#), [17](#)
- [38] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. A metric learning reality check. In *ECCV*, pages 681–699. Springer, 2020. [2](#)
- [39] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015. [6](#)
- [40] Hyun Jun Oh, Kyoung Mu Lee, and Sang Uk Lee. Occlusion invariant face recognition using selective local non-negative matrix factorization basis images. *Image and Vision computing*, 26(11):1515–1523, 2008. [8](#)
- [41] Elad Osherov and Michael Lindenbaum. Increasing cnn robustness to occlusions by reducing filter support. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 550–561, 2017. [7](#), [8](#)
- [42] Chunlei Peng, Nannan Wang, Jie Li, and Xinbo Gao. Re-ranking high-dimensional deep local representation for nirvis face recognition. *IEEE Transactions on Image Processing*, 28(9):4553–4565, 2019. [8](#)
- [43] Yuxin Peng, Cuihua Fang, and Xiaoou Chen. Using earth mover’s distance for audio clip retrieval. In *Pacific-Rim Conference on Multimedia*, pages 405–413. Springer, 2006. [2](#)
- [44] Haibo Qiu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. End2end occluded face recognition by masking corrupted features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [1](#), [2](#), [7](#), [8](#)
- [45] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000. [2](#), [3](#), [8](#)
- [46] M Saquib Sarfraz, Arne Schumann, Andreas Eberle, and Rainer Stiefelhagen. A pose-sensitive embedding for person re-identification with expanded cross neighborhood re-ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 420–429, 2018. [8](#)
- [47] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. [2](#), [6](#), [12](#), [14](#)
- [48] S. Sengupta, J.C. Cheng, C.D. Castillo, V.M. Patel, R. Chellappa, and D.W. Jacobs. Frontal to profile face verification in the wild. *IEEE Conference on Applications of Computer Vision*, February 2016. [6](#), [18](#), [21](#)
- [49] Rulin Shao, Zhouxing Shi, Jinfeng Yi, Pin-Yu Chen, and Cho-Jui Hsieh. On the adversarial robustness of visual transformers. *arXiv preprint arXiv:2103.15670*, 2021. [7](#)
- [50] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 acm sigsac conference on computer and communications security*, pages 1528–1540, 2016. [6](#), [8](#)
- [51] Xiaohui Shen, Zhe Lin, Jonathan Brandt, Shai Avidan, and Ying Wu. Object retrieval and localization with spatially-constrained similarity measure and k-nn re-ranking. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3013–3020. IEEE, 2012. [8](#)
- [52] Lingxue Song, Dihong Gong, Zhifeng Li, Changsong Liu, and Wei Liu. Occlusion robust face recognition based on mask learning with pairwise differential siamese network. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 773–782, 2019. [8](#)
- [53] Abby Stylianou, Richard Souvenir, and Robert Pless. Visualizing deep similarity networks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019. [4](#)
- [54] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deeply learned face representations are sparse, selective, and robust. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2892–2900, 2015. [1](#)
- [55] Thomas Swearingen and Arun Ross. Lookalike disambiguation: Improving face identification performance at top ranks.

- In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 10508–10515. IEEE, 2021. 8
- [56] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*, 2017. 2, 12
- [57] Daniel Sáez Trigueros, Li Meng, and Margaret Hartnett. Enhancing convolutional neural networks for face recognition with occlusion maps and batch triplet loss. *Image and Vision Computing*, 79:99–108, 2018. 7
- [58] Weitao Wan and Jiansheng Chen. Occlusion robust face recognition based on mask learning. In *2017 IEEE international conference on image processing (ICIP)*, pages 3795–3799. IEEE, 2017. 8
- [59] Chengrui Wang, Han Fang, Yaoyao Zhong, and Weihong Deng. Mlfw: A database for face recognition on masked faces. *arXiv preprint arXiv:2109.05804*, 2021. 2, 5, 7, 14, 19
- [60] Fan Wang and Leonidas J Guibas. Supervised earth mover’s distance learning and its computer vision applications. In *European Conference on Computer Vision*, pages 442–455. Springer, 2012. 3, 8
- [61] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, pages 5265–5274, 2018. 2, 6, 12
- [62] John Wright, Allen Y Yang, Arvind Ganesh, S Shankar Sastri, and Yi Ma. Robust face recognition via sparse representation. *IEEE transactions on pattern analysis and machine intelligence*, 31(2):210–227, 2008. 8
- [63] Xiang Xu, Nikolaos Sarafianos, and Ioannis A Kakadiaris. On improving the generalization of face recognition in the presence of occlusions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 798–799, 2020. 7, 8
- [64] Meng Yang, Lei Zhang, Jian Yang, and David Zhang. Robust sparse coding for face recognition. In *CVPR 2011*, pages 625–632. IEEE, 2011. 8
- [65] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014. 2, 4, 5, 7, 12, 19
- [66] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. Deepemd: Differentiable earth mover’s distance for few-shot learning, 2020. 3, 4, 5, 8, 14
- [67] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 8
- [68] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016. 15
- [69] Xuanmeng Zhang, Minyue Jiang, Zhedong Zheng, Xiao Tan, Errui Ding, and Yi Yang. Understanding image retrieval re-ranking: A graph neural network perspective. *arXiv preprint arXiv:2012.07620*, 2020. 8
- [70] Fang Zhao, Jiashi Feng, Jian Zhao, Wenhan Yang, and Shuicheng Yan. Robust lstm-autoencoders for face de-occlusion in the wild. *IEEE Transactions on Image Processing*, 27(2):778–790, 2017. 8
- [71] Wenliang Zhao, Yongming Rao, Ziyi Wang, Jiwen Lu, and Jie Zhou. Towards interpretable deep metric learning with structural matching. In *ICCV*, 2021. 2, 3, 4, 5, 8
- [72] Tianyue Zheng, Weihong Deng, and Jiani Hu. Cross-age LFW: A database for studying cross-age face recognition in unconstrained environments. *CoRR*, abs/1708.08197, 2017. 6, 16
- [73] Yaoyao Zhong and Weihong Deng. Towards transferable adversarial attack against deep face recognition. *IEEE Transactions on Information Forensics and Security*, 16:1452–1466, 2020. 2, 6, 7, 19
- [74] Shengyao Zhou, Junfan Luo, Junkun Zhou, and Xiang Ji. Asarcface: Asymmetric additive angular margin loss for fair-face recognition. In *ECCV Workshops*, 2020. 8
- [75] Zihan Zhou, Andrew Wagner, Hossein Mobahi, John Wright, and Yi Ma. Face recognition with contiguous occlusion using markov random fields. In *2009 IEEE 12th international conference on computer vision*, pages 1050–1057. IEEE, 2009. 8