

LAFS: Landmark-based Facial Self-supervised Learning for Face Recognition

Zhonglin Sun, Chen Feng*, Ioannis Patras, Georgios Tzimiropoulos
 Queen Mary University of London
 London, UK

{zhonglin.sun, chen.feng, i.patras, g.tzimiropoulos}@qmul.ac.uk

Abstract

In this work we focus on learning facial representations that can be adapted to train effective face recognition models, particularly in the absence of labels. Firstly, compared with existing labelled face datasets, a vastly larger magnitude of unlabeled faces exists in the real world. We explore the learning strategy of these unlabeled facial images through self-supervised pretraining to transfer generalized face recognition performance. Moreover, motivated by one recent finding, that is, the face saliency area is critical for face recognition, in contrast to utilizing random cropped blocks of images for constructing augmentations in pretraining, we utilize patches localized by extracted facial landmarks. This enables our method - namely **LAND**mark-based **F**acial **S**elf-supervised learning (**LAFS**), to learn key representation that is more critical for face recognition. We also incorporate two landmark-specific augmentations which introduce more diversity of landmark information to further regularize the learning. With learned landmark-based facial representations, we further adapt the representation for face recognition with regularization mitigating variations in landmark positions. Our method achieves significant improvement over the state-of-the-art on multiple face recognition benchmarks, especially on more challenging few-shot scenarios. The code is available at https://github.com/szlbibiubiu/LAFS_CVPR2024.

1. Introduction

In recent years, face recognition has witnessed significant advancements, owing to the emerging techniques such as advanced loss functions [10, 16, 42, 50, 58, 61, 62] and specialized network structures [7, 9, 17, 32, 33, 40, 55, 60, 64, 66, 70], and the large-scale annotated datasets such as Webface42M [71] and MS1M [24]. However, most of the works have overlooked the impact of initial parameters for supervised training, noted as facial representation, which

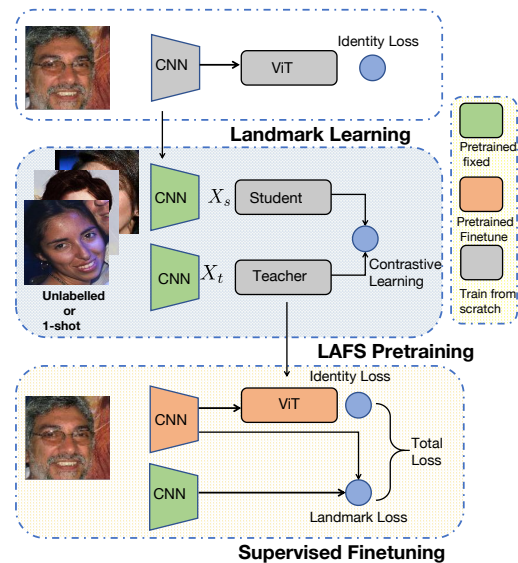


Figure 1. Illustration of our *pretraining* and *finetuning* pipeline for face recognition. First, a landmark CNN is learnt using Part fViT framework [55]. We adopt the landmark CNN to provide facial landmarks for constructing our **LAFS pretraining**. In this framework, the ‘Teacher’ processes the entire set of provided landmarks, while the ‘Student’ operates on subsets of these landmarks. Then, we transfer the ‘Teacher’ for *finetuning* with an additional regularization that penalizes landmark predictions from huge variations.

is commonly proven to be effective. For example, transferring from a good pretrained model derived by either supervised or self-supervised pretraining can yield better results on tasks such as Face Reconstruction [4], Face Anti-spoofing [57] and ImageNet Classification [6, 65]. A recent empirical facial comparison research [4] further reveals that self-supervised learning is able to bring much improvement compared to supervised learning, particularly, in data-limited scenarios. However, self-supervised learning with ResNet yields suboptimal results compared to supervised training under large-scale settings [4]. This observation suggests that previous self-supervised pre-training methods may not effectively scale from limited datasets to scenarios

*Corresponding Author.

with abundant identity information.

Moreover, another challenge for real-world face recognition is that large-scale datasets are not always available, due to issues such as non-commercial licenses and privacy concerns. Collecting new datasets often ensues in a challenging situation where only a few images are annotated for each novel identity [67]. Hence researches under insufficient labelling situation (i.e. few-shot face recognition, especially 1-shot) are worth consideration. Current research on few-shot face recognition focuses primarily on improving accuracy with meta-learning [28, 67, 69], ignoring that a good initial model (facial representation) is also crucial for few-shot face recognition. A potential solution towards maintaining generalized few-shot accuracy is self-supervised learning [44].

Motivated by the above concerns, we are interested in three fundamental questions:

- *With a vast number of unlabelled facial images in the real world, how can we take advantage of those data to train a face recognition model?*
- *With only a limited number of samples for each identity (few-shot learning), how well can a face recognition model perform and to what extent self-supervised learning can be of improvement over straightforward supervised training?*
- *What causes the failure of self-supervised learning when scaling from limited data to large-scale data?*

Our paper aims to address the challenge of learning facial representation that can be effectively adapted to face recognition in few-shot and large-scale scenarios. To solve the problems outlined above, we offer 3 contributions:

- We conduct a series of experiments to answer the question of how well self-supervised learning could behave on few-shot face recognition. We propose a pipeline designed explicitly for few-shot evaluation where we adopt self-supervised methods on an unlabelled(1-shot) pre-training dataset with 1M images which simulate the situation of unlabelled data in real-world applications, then we finetune on a few-shot(including 1-shot) datasets and evaluate on the classical test sets(e.g. IJB-C).
- Part fViT [55] has shown superior localization performance on facial images, indicating that it is feasible to combine a part-based model with self-supervised learning. Inspired by this, We propose a novel landmark-based self-supervised framework for face recognition that pertains entirely to facial parts, where dense landmarks are trained to produce similar representations of sparse landmarks. This facilitates a shift from standard grid learning to landmark-based learning. Fig 1 shows the pipeline for our LAFS pretraining and fine-tuning. Furthermore, we investigate the properties of landmarks and propose two landmark-related augmentations to enhance the representation of self-supervised learning.

- A series of studies are conducted for LAFS, we demonstrate the effectiveness of standard evaluation on large-scale datasets as well as few-shot evaluations, which have never been explored in previous works.

Our main findings are: (a) Without explicit label information, our pipeline, which comprises unlabeled(1-shot) pre-training followed by 1-shot fine-tuning, can deliver accurate face recognition performance. We demonstrate that intra-sample pretraining (self-supervised learning loss) and inter-sample learning (1-shot supervised loss, e.g. CosFace [58]) are both necessary for learning generalized 1-shot face recognition in our pipeline. (b) We find that distinction from the facial research [4], self-supervised learning with Vision Transformer can break the limits for self-supervised learning with ResNet, exceeding supervised training by a large margin. (c) Our proposed LAFS, which works entirely on facial landmarks, is capable of transferring to highly accurate face recognition.

2. Related Work

It is out of scope to review the bulk of papers for face recognition, we introduce the literature for Face Recognition in few-shot and general situations, then Self-supervised Learning.

Face Recognition The goal of face recognition is to learn discriminative feature embedding based on the combination of backbone and losses. To this end, some works [11, 36, 39, 41, 58, 62] concentrate on modifying the softmax loss function which represents separability to counter the discrimination as softmax loss could not handle the global feature distribution especially when we using mini-batch. CosFace [58] adds a margin on the $\cos(\theta)$ to control the distance between the normalized weight and feature embedding, resulting in a more discriminative feature learning scheme. Following that, ArcFace [11] proposes that defining the margin on the angle θ would provide more discriminative results than defining on $\cos(\theta)$. In addition, AdaFace [36] proposes to learn to distinguish low-quality samples and employs adaptive margin to learn more discriminative features. To better extract discriminative representation with loss function, the choice of a suitable backbone plays a pivotal role in computing face images to generate feature embedding. The to-go backbone during the past few years is ResNet [7, 25, 37, 64, 66], recently Vision Transformer are investigated in [1, 9, 55, 70] for face recognition while training ViT still requires a large amount of data or heavy data augmentation. Part fViT [55] is the pioneering study to integrate Vision Transformer with facial parts to learn discriminative facial area which benefits the accuracy. As for the limited data volume, methods for few-shot face recognition [67, 68] mostly focus on exploring dis-

criminative features combined with meta-learning training.

Self-Supervised Learning Self-supervised learning has achieved great performance as an off-the-shelf representation techniques in various computer vision tasks [5, 6, 8, 21–23, 26, 27, 65]. Recently, contrastive learning [2, 5, 6, 8, 19, 20, 23, 26] has made significant strides in exploring image representations based on the instance discrimination task, by considering each sample itself as a unique class. More specifically, two different views of a specific sample are treated as positive to each other while the rest of the dataset is regarded as negative, making the batch size a crucial factor. Large batch sizes are essential for effective training in SimCLR [8]. Instead of contrastive learning, some other works try to avoid instance discrimination tasks. Involving Vision Transformer into consideration, DINO[6] proposes that global augmentation should have similar representations to local augmentations and introduces a centring method to provide stable results. For face recognition applications, Face-MAE[59] utilizes a masked strategy to reconstruct the face and then minimize the embedding distance between the reconstructed face and the original face. Sevastopolskiy et al. [52] proposes StyleGAN [34] is able to boost FR by reconstructing the input image via taking StyleGAN as a self-supervised learning method.

3. Methodology

In Section 3.1, applying DINO [6] for face recognition pre-training is introduced. Then, a novel self-supervised learning method(LAFS) for face recognition which focuses on the representation between full landmarks and the subset of landmarks is explained in Section 3.2. Furthermore, in Section 3.3 we present two types of augmentations. Section 3.4 describes how we simulate using unlabelled data for pre-training. And finally, we explain the finetune strategy for the LAFS in Section 3.5. Details for baseline Part fViT [55] is described in SUPPLEMENTARY MATERIAL 7.1

3.1. DINO for face recognition

We start by adopting the DINO [6] which provides strong feature embedding through contrastive learning between global views and local views. A facial image X is firstly converted into a pair of views, denoted as (X_t, X_s) . The objective of the DINO is as follows:

$$\min_{\theta_s} \mathbf{H}(\mathbf{P}(\text{Cen}(\vartheta_t(X_t)), \vartheta_s(X_s))) \quad (1)$$

Where $\vartheta(*)$ is the backbone, t and s represent the teacher branch and student branch, $\mathbf{P}(*)$ is a projection head with softmax probability, $\text{Cen}(*)$ is centring operation which minus the mean computed on the output and $\mathbf{H}(\mathbf{a}, \mathbf{b}) = -\mathbf{a} \log(\mathbf{b})$ is the cross-entropy. DINO endeavours to investigate the interdependency between global and local features in a given image, then the augmentation for the teacher

branch is treated as a global representation and the student branch is expected to output a representation of local views. We adopt DINO [6] as the baseline SSL method for our pipeline.

3.2. Full-to-Subset Landmark Consistency Learning

In contrast to DINO’s approach which focuses on the global-local relationship of images, our approach aims to leverage the aggregation and structural characteristics of faces by seeking similar representations between a full set and a subset of facial parts (landmarks).

Algorithm 1 Pipeline of the proposed Landmark-based Facial self-supervised pretraining

- 1: Initialization: Input image X , teacher network ϑ_t , student network ϑ_s , pretrained landmark CNN $\delta()$, Landmark Augmentations $\sigma_1()$, DINO augmentations $\sigma_d()$, moving mean \mathbf{l} , $n \leftarrow 0$, maximum iterations τ
- 2: **repeat**
- 3: $n \leftarrow n + 1$
- 4: Sample a batch X^n
- 5: Generate views $X_t^n, X_s^n \leftarrow \sigma_d(X^n), \sigma_d(X^n)$
- 6: $L_t^n, L_s^n \leftarrow \delta(X_t^n), \text{Sample}(\delta(X_s^n))$
- 7: $\hat{L}_t^n, \hat{L}_s^n \leftarrow \sigma_1(L_t^n), \sigma_1(L_s^n)$
- 8: $\hat{L}_t^n \leftarrow \hat{L}_t^n.\text{detach}()$
- 9: Update $\vartheta_s^{n+1} \leftarrow \vartheta_s^n - \nabla_{\vartheta_s^n} \text{Eq 5}$
- 10: Update $\vartheta_t^{n+1} \leftarrow \mathbf{l} * \vartheta_t^n + (\mathbf{1} - \mathbf{l}) \vartheta_s^n$
- 11: **until** converges or $n = \tau$

Output: output model ϑ_t^τ and ϑ_s^τ

Previous landmark localization methods, such as FAN [3], are unable to provide landmarks that are specifically relevant for face recognition, thereby negatively affecting the face recognition performance as shown in Part fViT [55]. To address this issue, we opt for Part fViT [55] for constructing our landmark-based self-supervised learning pipeline. Specifically, as illustrated in (A) part of Figure 2 and Algorithm 1, we utilize a landmark CNN $\delta()$, which is fixed and pre-trained on face recognition task without any positional supervision, to generate a series of facial landmarks before feeding them into the backbone.

$$L_t, L_s = \delta(X_t), \delta(X_s) \quad (2)$$

This landmark CNN can provide stable and highly aggregated image patches for face recognition. It is crucial to note that if the parameters of this landmark CNN are not fixed, it may converge to the standard ViT grids and negatively impact recognition performance. To establish correspondence between all landmarks and a random subset of landmarks, we randomly sample a portion of landmarks on the student network branch. In our case, we randomly sample 36 out of R=196 global landmarks, which corresponds

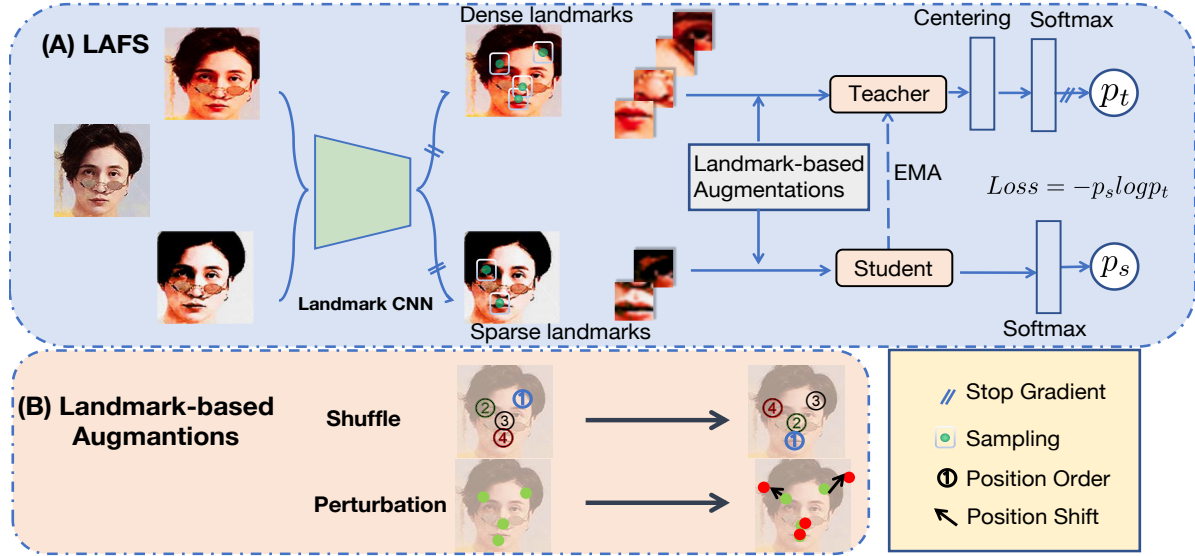


Figure 2. (A) The pipeline of our proposed LAFS framework. Two views of a facial image are first processed by the landmark CNN to provide landmark localization. Then we sample a certain subset of landmarks on the student branch. Following that, landmarks-based augmentations are added before converting into embedding for processing by teacher and student backbones. The representations of the two views are compared by the output of the backbones without label information. Gradients are backpropagated to the student network and the teacher network is updated by the exponential moving average of student parameters. (B) Landmark Augmentations. The upper part is shuffling where the order for sending to fViT is ①, ②, ③, ④, after shuffling the order changed. The bottom part explains the coordinates variation given the perturbation, each position of the green point shifts to the red point after the perturbation.

to the resolution of local views in DINO.

$$\mathbf{L}_t, \mathbf{L}_s = \mathbf{L}_t, \text{Sample}(\mathbf{L}_s) \quad (3)$$

Furthermore, the scale of random crops on the student branch is the same as that of the teacher branch. We append two landmark-based augmentations $\delta_1(\cdot)$ on the learned landmark to enhance the difficulties of the self-learning.

$$\hat{\mathbf{L}}_t, \hat{\mathbf{L}}_s = \delta_1(\mathbf{L}_t), \delta_1(\mathbf{L}_s) \quad (4)$$

Where $\hat{\mathbf{L}}_t, \hat{\mathbf{L}}_s$ are the full landmark views and subsets of landmark views respectively which can be considered as proprietary versions of the global and local views for face recognition. Our proposed method can effectively explore the representation of global facial landmarks and random subset landmarks. The learning objective for LAFS is the same as that of DINO [6]:

$$\min_{\theta_s} \sum_{\hat{\mathbf{L}}_t \in \{\hat{\mathbf{L}}_{t1}^g, \hat{\mathbf{L}}_{t2}^g\}} \sum_{\hat{\mathbf{L}}_s \neq \hat{\mathbf{L}}_t} \mathbf{H}(\mathbf{Q}_t(\hat{\mathbf{L}}_t), \mathbf{Q}_s(\hat{\mathbf{L}}_s)) \quad (5)$$

Where \mathbf{Q}_t and \mathbf{Q}_s are probability produced by the teacher and student backbones, $\hat{\mathbf{L}}_{t1}^g$ and $\hat{\mathbf{L}}_{t2}^g$ are two global landmark views. θ_s is the student network.

3.3. Landmark-based data augmentations

To further regularize the model, we design two landmark-based data augmentations as shown in (B) part of Figure 2,

namely Landmark Coordinate Perturbation and Landmark Shuffle.

Landmark Shuffle The landmark CNN, learned from the face recognition task, is capable of finding the corresponding positions of different faces. As a result, the positions of the patches in the part-based model differ from those in the standard ViT, which assumes sequential and non-overlapping positions for patches. The landmarks produced by the landmark CNN do not impose any order constraints, particularly with regard to position, and adjacent landmarks are overlapping as is described in [55]. However in vision tasks, maintaining the order of patches preserves the overall structure and global information [48, 56]. In light of this, we hypothesize that strict order is not necessary for the part-based model and have adopted a shuffle operation to perturb the model. Surprisingly, this perturbation significantly improves recognition performance, which contradicts the negative impact of shuffle as observed in [48]. We posit that this efficacy of shuffle on landmarks may be attributed to the differences in position encoding between the Part fViT and standard fViT.

Landmark Coordinate Perturbation Despite obtaining relatively stable landmarks with the pre-trained and fixed landmark CNN, we suspect that the landmark CNN may not have the same localization performance for all faces,

and that the adjacent image regions beyond the landmarks are also worth exploring. Thus we design data augmentation to enable the model to explore the area surrounding the landmarks. Specifically, for given landmarks r , we add a coordinate perturbation vector to its coordinate according to the following formula:

$$\mathbf{r} = \mathbf{r} + \alpha * \mathbf{u} \quad (6)$$

Here, \mathbf{u} is subject to standard normal distribution $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$, and α controls the magnitude of the disturbance, resulting in a normal distribution with $\mathcal{N}(\mathbf{0}, \alpha^2)$.

It is important to note that these two augmentations only show effectiveness in our proposed self-supervised learning method. Therefore, applying these augmentations to the standard backbone, specifically perturbing the DINO+fViT architecture by shuffling, results in a significant decrease in performance. Further details such as appropriate perturbation amplitude and results can be found in Section 4.2.2 and Section 4.2.3.

3.4. Unlabeled Face Pretraining

We are curious about how well a face recognition model can go with the situation in real-world where images are available without identity information. To address this issue, we simulate by regarding each unlabelled sample as a novel class, resulting in 1-shot face recognition. As a traditional solution, straightforwardly applying classification-based loss such as CosFace [58] will provide disastrous performance as shown in Section 4.2.1. Inspired by Lu et al. [44] that self-supervised can be a good few-shot learner, we construct a pipeline for few-shot face recognition, where we pretrain on a large-scale unlabelled facial dataset then finetune on few-shot face datasets and finally evaluate the performance on traditional testset. Our experiments make use of the WebFace260M dataset [71], a clean dataset consisting of 42 million images and 2 million identities, we are able to sample a 1-shot dataset of 1 million classes for our task to simulate the case of real-world application where each image is treated as a new identity by Self-supervised Contrastive Learning. Discussion regarding the difference between 1-shot and unlabelled is available at SUPPLEMENTARY MATERIAL 8.4.1

3.5. Finetune

Inter-sample Separability or Intra-sample Similarity Self-supervised learning methods primarily focus on exploring the representation of different views of an input image. This raises a pertinent question: When finetuning in a 1-shot setting, which loss function, self-supervised or classification-based, performs better? We address this query through our experiments, wherein we examine applying DINO [6] loss (Intra-sample Similarity) and Cosface [58] (Inter-sample Separability) for finetuning on the

1-shot dataset in experiment 4.2.1. DINO loss could not provide superior results compared with Cosface, and it slightly benefits the model without fine-tuning. This indicates that for 1-shot face recognition, simply applying classification loss or self-supervised loss would not benefit the recognition performance, instead, the pipeline unlabeled(1-shot) pretraining with intra-sample similarity (self-supervised learning) and fine-tuning with inter-sample separability (Supervised classification) is crucial for addressing the 1-shot problem.

Landmark supervision We have 3 finetuning options for LAFS finetuning: (a) Fix the landmark CNN and train fViT only; (b) Finetune the whole backbone (landmark CNN and fViT); (c) As shown in the bottom part of Figure 1, we use an extra pre-trained and fixed landmark CNN to provide landmark localization coordinates as soft-label, and the trainable landmark CNN could be either from the supervised pretraining or self-supervised(e.g. DINO) pretraining. The overall finetuning loss is:

$$L_{total} = L_{id} + \beta ||\hat{r}, r||_2 \quad (7)$$

Where \hat{r} is the landmarks predicted by the fixed landmark CNN, r is the predicted coordinates, and β controls the variation of landmarks. L_{id} is CosFace [58] to learn discriminative embedding. We ablate the choice of the landmark finetuning options in SUPPLEMENTARY MATERIAL 8.3.2.

4. Experiments

In this section, we examine the accuracy of models transferring from the proposed LAFS on several well-known datasets as well as the few-shot evaluation, and conduct a series of ablation studies to illustrate the effectiveness of our proposed methods. Text in bold means the best results obtained. We include pretraining, finetuning details as well as the models opted in SUPPLEMENTARY MATERIAL 8.

4.1. Implementation details

To fairly compare with our other methods, for the self-supervised pipeline, we randomly sample 1M images with 1 image per identity from the first 50% part of the Webface42M dataset for self-supervised pretraining, namely Webface-1shot. As for supervised finetuning, we opt for MS1MV3 [12] containing 93,431 identities and 10% of Webface42M [71] dataset, coined as Webface4M, comprising 4M images with 200K identities. We also test how the unlabelled in-the-wild dataset behaves in SUPPLEMENTARY MATERIAL 8.1.1.

4.2. Ablation Studies

A number of studies are conducted to demonstrate the effectiveness of different parts of methods. For self-supervised

pretraining, we adopt Part fViT-Tiny with patch number 196, size of 15.28M and 2.48G Flops. We put ablation studies for the number of shots, finetuning landmark supervision, comparison of global landmark view between global view and impact of the in-the-wild dataset in SUPPLEMENTARY MATERIAL 8.2.

4.2.1 Finetune with Similarity or Separability

We consider different loss options when transferring from LAFS, specifically the Inter-sample Separability (CosFace [58] loss) and Intra-sample Similarity (DINO [6] loss). We report the results of simply adopting the pretrain model for evaluation, then CosFace and DINO in Table 1. The accuracy for directly training CosFace [58] on the Webface1M without pretraining is also included. We can conclude that DINO improve the model without fine-tuning, and CosFace can bring large benefits in terms of accuracy. Then training 1-shot data from scratch leads to invalid results.

Finetune Loss	LFW	CFP-FP	AgeDB
w/o Finetune	75.20	70.82	55.90
DINO	76.96	73.44	56.65
CosFace	88.58	76.04	65.84
CosFace, 1M	50.00	50.00	50.00

Table 1. Ablation studies for different types of finetune loss. W/o Finetune means directly adopting pretrained model for evaluation. CosFace, 1M means the model is trained from scratch on the 1M 1-shot training set. The backbone adopted is Part fViT-B

4.2.2 Effect of Landmark Shuffle

Herein we challenge the property of the negative impact of shuffle for Vision Transformer which can be found in Table 2 (middle part), as well as the performance of standard fViT with shuffle operations. We can conclude that the shuffling operation indeed disturbs the global information for fViT, but it leads to better recognition performance for Part fViT.

4.2.3 Effect of Landmark Coordinates Perturbation

We disturb the learned landmark by adding a noise subjected to normal distribution. As listed in Table 2 (bottom part), the magnitude with $\alpha = 2$ performs the best towards finetuning results. While too large perturbation would even perform worse than without perturbation.

We also conduct experiments to test the impact of these two augmentations on the standard grid (i.e. fViT). From the results, the following conclusions are observed: (1) Shuffling operation leads to worse results for the standard fViT, which is consistent with previous findings [48]. (2)

Experiment	Backbone	Content	1% data	10% data
Baseline	Part fViT	w/o ssl	6.91	90.87
		LAFS w/o Aug	38.05	91.06
	fViT	w/o ssl	8.54	86.9
		DINO	36.27	87.15
Landmark Shuffle	Part fViT (LAFS)	w/o shuffle	38.05	91.06
		shuffle	40.01	91.53
	fViT (DINO)	w/o shuffle	36.27	87.15
		shuffle	30.96	78.95
Landmark Perturbation	Part fViT (LAFS)	$\alpha = 0$	38.05	91.06
		$\alpha = 2$	39.44	91.45
		$\alpha = 5$	39.16	91.37
	fViT (DINO)	$\alpha = 0$	36.27	87.15
$\alpha = 2$		36.30	87.23	

Table 2. Ablation studies for the effectiveness of Landmark augmentations, the top parts are the results for without landmark augmentations. We report IJB-B when TAR@FAR=1e-4. The $x\%$ in Table 2 means using only $x\%$ labelled data from the whole WebFace4M dataset.

Perturbation augmentation could bring small benefits in increasing the performance of the standard fViT.

4.3. Comparison with the State-of-the-Art

We experiment with ResNet-100, fViT-B and Part fViT-B models to evaluate the performance transferring from the self-supervised pipeline on well-known datasets to compare with the recent state-of-the-art methods as well as the results fine-tuned with few-shot samples.

The training (or fine-tuning) datasets we used for quantitative results are Webface4M [71], few-shot subsets of Webface4M and MS1MV3 [24]. We achieve SOTA or near SOTA results on the Webface4M and MS1MV3 dataset, and our proposed LAFS pretraining outperforms other methods for few-shot evaluation. Note that those results are fine-tuned with (c) mentioned in Section 3.5.

4.3.1 Few-shot Evaluation

We report the performance few-shot evaluation on a few subsets of Webface4M [71] in Table 3, specifically 1% labels with randomly 1, 2, 4, 10 shots and all images per label, and 10% labels with 1, 2, 4, 10 shots and 10% labels with their available images. The evaluation is performed on three backbones: ResNet-100, fViT-B, and Part fViT-B, considering both the supervised training and the self-supervised pretraining settings. We also test the impact of the in-the-wild dataset Flickr which is collected by us in SUPPLEMENTARY MATERIAL 8.3. Results for Part fViT with fixed landmark CNN under DINO training are also provided. Based on the evaluation results, we can draw the following conclusions:

- Without self-supervised pretraining, we consistently observe that ResNet outperforms both fViT and Part fViT in

Data Amount	Pretrain Method	Backbone	IJB-B					IJB-C					
			1 shot	2 shot	4 shot	10 shot	all	1 shot	2 shot	4 shot	10 shot	all	
1%	Scratch	ResNet	14.13	19.69	30.58	22.19	58.66	15.39	22.15	33.56	25.01	63.16	
		fViT	1.67	3.79	6.6	16.91	27.22	1.84	4.03	7.36	19.03	29.19	
		Part fViT	0.64	0.89	1.46	3.47	6.22	0.58	0.97	1.59	3.73	5.99	
	DINO	ResNet	13.98	22.70	36.85	52.72	72.38	15.96	26.79	41.87	57.19	77.3	
		fViT	33.48	41.67	56.10	66.20	74.42	37.49	46.93	60.13	70.44	78.95	
		Part fViT	18.51	23.41	35.19	49.19	69.37	21.40	27.23	39.38	53.70	75.82	
	LAFS	Part fViT	33.97	42.80	57.32	68.76	75.67	37.89	47.03	61.09	72.57	79.66	
	10%	Scratch	ResNet	27.23	41.96	59.27	68.25	89.33	28.84	45.25	63.26	72.40	92.28
			fViTB	7.53	13.32	19.85	45.45	81.37	7.63	14.21	22.11	52.80	84.62
Part fViT			3.11	3.97	14.66	54.47	84.47	3.32	4.25	17.21	58.75	88.09	
DINO		ResNet	27.63	48.97	76.83	87.58	90.99	27.91	59.58	81.29	90.26	92.28	
		fViT	46.99	66.61	81.56	88.35	91.35	51.35	70.73	85.02	91.42	93.85	
		Part fViT	29.82	47.48	75.92	85.39	90.37	33.87	52.45	74.60	89.07	91.44	
LAFS		Part fViT	48.56	66.71	81.67	88.70	92.16	51.97	71.10	85.09	91.56	94.32	

Table 3. The comparison of the proposed methods on few-shot evaluation

few-shot evaluation scenarios. This gap can be attributed to the overfitting property of Vision Transformer when trained on limited amounts of data [54, 70].

- Part fViT can exceed fViT when a sufficient volume of data is available under the supervised training setting, specifically when the data amount is equal to or larger than 10% label with 10 shots.
- Self-supervised learning consistently yields superior results compared to training without self-supervised training, except for the ResNet under 1% label with 1-shot setting. Notably, the benefits of self-supervised learning become more prominent when the volume of available data decreases, which aligns with the observations from the facial research [4].
- Consistent with the findings of DINO [6], we observe that the improvement achieved through self-supervised learning for ResNet is not as significant as that observed for fViT in most cases.
- Our proposed LAFS pretraining could always obtain effective results compared with models pre-trained by DINO. Particularly, when the number of shots reaches the maximum available data, we observe significant improvements over DINO which aligns with supervised finding, implying that more diversity in each identity will help with finding effective landmarks. However, DINO with Part fViT would provide worse training results compared to fViT with DINO, which illustrates the effectiveness of our LAFS pipeline.

4.3.2 Webface4M

We present the results of the proposed self-supervised pre-training on the 1-shot image set, then finetuning on Webface4M in Table 4. We also test the results of ResNet-100 with augmentations as well as ResNet-100 pretrained from DINO. We also self-implement Part fViT on Webface4M

Method	SSL	LFW	CFP-FP	AgeDB	IJB-B	IJB-C
ArcFace w/o Aug [71]	-	99.85	99.04	97.82	-	96.77
CosFace w/o Aug [71]	-	99.80	99.25	97.45	-	96.86
QAFace [49]	-	99.85	99.21	97.91	95.67	97.20
AdaFace [36]	-	99.80	99.17	97.90	96.03	97.39
ArcFace PFC-0.3 [1]	-	99.85	99.23	98.01	95.64	97.22
ResNet, w/ Aug	-	99.83	99.16	97.89	95.94	97.34
ResNet, w/ Aug	DINO	99.82	99.14	97.9	95.83	97.29
fViT [55]	-	99.83	99.04	97.48	95.72	97.15
fViT	DINO	99.83	99.10	97.77	95.91	97.33
Part fViT [55]	-	99.83	99.13	97.73	96.05	97.4
Part fViT	LAFS	99.83	99.18	97.90	96.28	97.55
fViT, w/o Aug	DINO	99.85	99.05	97.70	95.64	97.12
Part fViT, w/o Aug	LAFS	99.81	99.14	97.83	96.07	97.44

Table 4. Comparison with the state-of-the-art on Webface4m. LAFS outperform other methods on several testsets.

with the same hyper-parameters of [55] for comparison as it is not reported.

It can be observed on large pose evaluation provided by the CFP-FP protocol, Part fViT boosts the fViT by 0.09 while producing similar results with ResNet-100 under the same training augmentations. Self-supervised methods bring small improvement, such that LAFS outperforms Part fViT by 0.05 and DINO with fViT improves supervised fViT by 0.06. However, for age variation evaluation provided by the AgeDB-30 dataset, the self-supervised learning proves advantageous for all considered backbones, our LAFS improves Part fViT by 0.17, fViT with DINO improves fViT by 0.29, ResNet with DINO improves ResNet by 0.01. And LAFS still surpasses baseline fViT with DINO which is 0.13. On **IJB-B** and **IJB-C** datasets, baseline fViT could produce competitive results compared with AdaFace [36], and the proposed LAFS is able to achieve much higher accuracy by 0.16 on IJB-C and 0.25 on IJB-B than AdaFace. In terms of improvement of SSL, SSL could generate promising improvement (i.e. DINO improves fViT by 0.18 and LAFS boosts Part fViT by 0.15 on IJB-C).

Notably, our method achieves better performance even

Method	SSL	Train Data	LFW	CFP-FP	AgeDB	IJB-B	IJB-C	MegaFace/id	MegaFace/ver
CosFace[58]	-	MS1MV2	99.81	98.12	98.11	94.80	96.37	97.91	97.91
ArcFace[11]	-	MS1MV2	99.83	92.27	92.28	94.25	96.03	98.35	98.48
Sub-center ArcFace[13]	-	MS1MV2	99.80	98.80	98.31	94.94	96.28	98.16	98.36
FAN-Face[66]	-	MS1MV2	99.85	98.63	98.38	94.97	96.38	98.70	98.95
VirFace[38]	-	MS1MV2	99.56	97.15	-	88.90	90.54	-	-
MagFace[46]	-	MS1MV2	99.83	98.46	96.15	94.51	95.97	-	-
Face Transformer[70]	-	MS1MV2	99.83	96.19	97.82	-	95.96	-	-
ArcFace-challenge[15]	-	MS1MV3	99.85	99.06	98.48	-	96.81	-	-
VPL[16]	-	MS1MV3	99.83	99.11	98.60	95.56	96.76	98.80	98.97
AdaFace[36]	-	MS1MV3	99.83	99.03	98.17	95.84	97.09	-	-
ResNet, w/ Aug	-	MS1MV3	99.85	99.14	98.16	96.03	97.22	98.78	98.67
ResNet, w/ Aug	DINO	MS1MV3	99.85	99.12	98.14	95.86	97.01	98.50	98.53
fViT [55]	-	MS1MV3	99.85	99.01	98.13	95.97	97.21	98.69	98.91
fViT, ours	DINO	MS1MV3	99.85	99.08	98.10	96.13	97.27	98.70	98.65
Part fViT [55]	-	MS1MV3	99.83	99.21	98.29	96.11	97.29	98.96	98.78
Part fViT, ours	LAFS	MS1MV3	99.86	99.15	98.33	96.24	97.48	98.66	98.95
fViT, w/o Aug, ours	DINO	MS1MV3	99.81	99.02	98.10	95.76	97.14	98.15	98.31
Part fViT, w/o Aug, ours	LAFS	MS1MV3	99.83	99.11	98.33	95.98	97.20	98.73	98.62

Table 5. Comparison with the state-of-the-art on MS1MV3 datasets. Our proposed LAFS achieve state-of-the-art results on most datasets.

without data augmentation compared to **methods applying strong augmentations (AdaFace, QAFace and Part fViT)**, further validating its superior performance. LAFS without augmentations outperforms the SOTA method AdaFace [36] by 0.04 on IJB-B and 0.05 on IJB-C. fViT with DINO behaves slightly worse than ArcFace PFC-0.3 [1] and QAFace [49] which is 95.64 on IJB-B and 97.12 on IJB-C, but still a large increase compared to supervised CosFace [58]. Another observation is that SSL could only provide marginal improvement or even worse for ResNet, which is consistent with the facial research [4].

4.3.3 MS1MV3

As the original version of MS1MV*, MS-Celeb [24], is abandoned by the creator, AdaFace [36] advocates for researchers to transition towards using the Webface4M [71], hence for the sake of reference, we report the results of the models trained on MS1MV3, and test on various benchmarks. The results are shown in Table 5.

As observed, fViT with DINO pretraining is able to boost the accuracy, especially on IJB-B and IJB-C, by a large margin over supervised fViT, e.g. 0.16 on IJB-B and 0.06 on IJB-C. Remarkably, the performance achieved by LAFS for Part fViT considerably outweighs that of DINO, achieving the best performance on MS1MV3 which improves Part fViT-B without self-supervised learning by 0.13 on IJB-B and 0.19 on IJB-C. Also LAFS outperforms the state-of-the-art method AdaFace [36] by 0.4 and 0.39 respectively.

In addition, regarding results without strong data augmentation, our LAFS outperforms the SOTA method AdaFace [36] by 0.14 on IJB-B and 0.11 on IJB-C. And for ResNet-100 the same phenomenon is observed that with 100% of the dataset, SSL is not capable of bringing im-

provement when the data volume is large, it even results in a reduction in accuracy, for instance, dropping from 97.22 with supervised training to 97.01 with SSL pretraining.

5. Conclusions

We propose landmark-based self-supervised learning (LAFS) to learn generalized representations that can be adapted to highly accurate face recognition. We offer several key contributions: (A) A self-supervised pipeline utilizing un-labelled data for training generalized representation, which can be adopted to evaluate the effectiveness of self-supervised representation. (B) LAFS, a novel self-supervised learning framework called landmark-based self-supervised learning for face recognition, minimises the representation of all landmarks and fewer landmarks to provide promising improvement when transferring to face recognition. (C) Two augmentations, namely Landmark Shuffle and Landmark Coordinates Perturbation, are effective and robust for learning more generalized representations for LAFS. Overall, our approach achieves state-of-the-art or near state-of-the-art performance on several face recognition benchmarks including few-shot evaluation, demonstrating the effectiveness of our proposed methods. However, the landmarks in LAFS are fixed and don't benefit from self-supervised training. In the future, we will explore more techniques to refine the landmarks during self-supervised training stages. In addition, we will investigate to what extent the LAFS can behave on other facial tasks beyond face recognition.

Acknowledgments: This research utilised QMUL's Apocrita HPC facility, supported by QMUL Research-IT and ITS Research.

References

- [1] Xiang An, Jiankang Deng, Jia Guo, Ziyong Feng, XuHan Zhu, Jing Yang, and Tongliang Liu. Killing two birds with one stone: Efficient and robust training of face recognition cnns by partial fc. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4042–4051, 2022. [2](#), [7](#), [8](#)
- [2] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *European Conference on Computer Vision*, pages 456–473. Springer, 2022. [3](#)
- [3] Adrian Bulat and Georgios Tzimiropoulos. Human pose estimation via convolutional part heatmap regression. In *ECCV*, 2016. [3](#)
- [4] Adrian Bulat, Shiyang Cheng, Jing Yang, Andrew Garbett, Enrique Sanchez, and Georgios Tzimiropoulos. Pre-training strategies and datasets for facial representation learning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIII*, pages 107–125. Springer, 2022. [1](#), [2](#), [7](#), [8](#)
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020. [3](#)
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. [1](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [7] Dong Chen, Xudong Cao, Fang Wen, and Jian Sun. Blessing of dimensionality: High-dimensional feature and its efficient compression for face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3025–3032, 2013. [1](#), [2](#), [3](#)
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [3](#)
- [9] Jun Dan, Yang Liu, Haoyu Xie, Jiankang Deng, Haoran Xie, Xuansong Xie, and Baigui Sun. Transface: Calibrating transformer training for face recognition from a data-centric perspective. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20642–20653, 2023. [1](#), [2](#)
- [10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. [1](#)
- [11] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. [2](#), [8](#), [1](#)
- [12] Jiankang Deng, Jia Guo, Debing Zhang, Yafeng Deng, Xiangu Lu, and Song Shi. Lightweight face recognition challenge. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. [5](#)
- [13] Jiankang Deng, Jia Guo, Tongliang Liu, Mingming Gong, and Stefanos Zafeiriou. Sub-center arcface: Boosting face recognition by large-scale noisy web faces. In *European Conference on Computer Vision*, pages 741–757. Springer, 2020. [8](#)
- [14] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotzia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5203–5212, 2020. [1](#)
- [15] Jiankang Deng, Jia Guo, Xiang An, Zheng Zhu, and Stefanos Zafeiriou. Masked face recognition challenge: The insight-face track report. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1437–1444, 2021. [8](#)
- [16] Jiankang Deng, Jia Guo, Jing Yang, Alexandros Lattas, and Stefanos Zafeiriou. Variational prototype learning for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11906–11915, 2021. [1](#), [8](#)
- [17] Changxing Ding and Dacheng Tao. Trunk-branch ensemble convolutional neural networks for video-based face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):1002–1014, 2017. [1](#), [3](#)
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [1](#)
- [19] Chen Feng and Ioannis Patras. Adaptive soft contrastive learning. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 2721–2727. IEEE, 2022. [3](#)
- [20] Chen Feng and Ioannis Patras. Maskcon: Masked contrastive learning for coarse-labelled dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19913–19922, 2023. [3](#)
- [21] Chen Feng, Georgios Tzimiropoulos, and Ioannis Patras. SSR: An efficient and robust framework for learning with unknown label noise. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022. [3](#)
- [22] Zheng Gao, Chen Feng, and Ioannis Patras. Self-supervised representation learning with cross-context learning between global and hypercolumn features. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1773–1783, 2024.
- [23] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. [3](#)
- [24] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for

- large-scale face recognition. In *European conference on computer vision*, pages 87–102. Springer, 2016. 1, 6, 8
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 1
- [26] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 3
- [27] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 3
- [28] Ashwamegha Holkar, Rahee Walambe, and Ketan Kotecha. Few-shot learning for face recognition in the presence of image discrepancies for limited multi-class datasets. *Image and Vision Computing*, 120:104420, 2022. 2
- [29] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1314–1324, 2019. 1
- [30] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008. 1
- [31] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Advances in neural information processing systems*, 28:2017–2025, 2015. 1
- [32] Bong-Nam Kang, Yonghyun Kim, and Daijin Kim. Pairwise relational networks for face recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 628–645, 2018. 1
- [33] Bong-Nam Kang, Yonghyun Kim, Bongjin Jun, and Daijin Kim. Hierarchical feature-pair relation networks for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 1
- [34] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33:12104–12114, 2020. 3
- [35] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4873–4882, 2016. 1
- [36] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18750–18759, 2022. 2, 7, 8
- [37] Pengyu Li. BioNet: A biologically-inspired network for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10344–10354, 2023. 2
- [38] Pengyu Li, Biao Wang, and Lei Zhang. Virtual fully-connected layer: Training a large-scale face recognition dataset with limited computational resources. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13315–13324, 2021. 8
- [39] Qiufu Li, Xi Jia, Jiancan Zhou, Linlin Shen, and Jinming Duan. Unitsface: Unified threshold integrated sample-to-sample loss for face recognition. *arXiv preprint arXiv:2311.02523*, 2023. 2
- [40] Jingtuo Liu, Yafeng Deng, Tao Bai, Zhengping Wei, and Chang Huang. Targeting ultimate accuracy: Face recognition via deep embedding. *arXiv preprint arXiv:1506.07310*, 2015. 1
- [41] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *ICML*, page 7, 2016. 2
- [42] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. 1
- [43] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 1, 2
- [44] Yuning Lu, Liangjian Wen, Jianzhuang Liu, Yajing Liu, and Xinmei Tian. Self-supervision can be a good few-shot learner. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIX*, pages 740–758. Springer, 2022. 2, 5
- [45] Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al. Iarpa janus benchmark-c: Face dataset and protocol. In *2018 International Conference on Biometrics (ICB)*, pages 158–165. IEEE, 2018. 1
- [46] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14225–14234, 2021. 8
- [47] 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 Workshops*, pages 51–59, 2017. 1
- [48] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *Advances in Neural Information Processing Systems*, 34: 23296–23308, 2021. 4, 6
- [49] Mohammad Saeed Ebrahimi Saadabadi, Sahar Rahimi Malakshan, Ali Zafari, Moktari Mostofa, and Nasser M Nasrabadi. A quality aware sample-to-sample comparison for face recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6129–6138, 2023. 7, 8

- [50] 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. [1](#)
- [51] Soumyadip Sengupta, Jun-Cheng Chen, Carlos Castillo, Vishal M Patel, Rama Chellappa, and David W Jacobs. Frontal to profile face verification in the wild. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016. [1](#)
- [52] Artem Sevastopolskiy, Yury Malkov, Nikita Durasov, Luisa Verdoliva, and Matthias Nießner. How to boost face recognition with stylegan? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20924–20934, 2023. [3](#)
- [53] Bharat Singh, Soham De, Yangmuzi Zhang, Thomas Goldstein, and Gavin Taylor. Layer-specific adaptive learning rates for deep networks. In *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*, pages 364–368. IEEE, 2015. [2](#)
- [54] Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *arXiv preprint arXiv:2106.10270*, 2021. [7](#)
- [55] Zhonglin Sun and Georgios Tzimiropoulos. Part-based face recognition with vision transformers. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022. [1](#), [2](#), [3](#), [4](#), [7](#), [8](#)
- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [4](#)
- [57] Chien-Yi Wang, Yu-Ding Lu, Shang-Ta Yang, and Shang-Hong Lai. Patchnet: A simple face anti-spoofing framework via fine-grained patch recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20281–20290, 2022. [1](#)
- [58] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5265–5274, 2018. [1](#), [2](#), [5](#), [6](#), [8](#)
- [59] Kai Wang, Bo Zhao, Xiangyu Peng, Zheng Zhu, Jiankang Deng, Xinchao Wang, Hakan Bilen, and Yang You. FaceMAE: Privacy-preserving face recognition via masked autoencoders. *arXiv preprint arXiv:2205.11090*, 2022. [3](#)
- [60] Qiangchang Wang, Tianyi Wu, He Zheng, and Guodong Guo. Hierarchical pyramid diverse attention networks for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8326–8335, 2020. [1](#)
- [61] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European conference on computer vision*, pages 499–515. Springer, 2016. [1](#)
- [62] Yandong Wen, Weiyang Liu, Adrian Weller, Bhiksha Raj, and Rita Singh. SphereFace2: Binary classification is all you need for deep face recognition. *arXiv preprint arXiv:2108.01513*, 2021. [1](#), [2](#)
- [63] Cameron Whitelam, Emma Taborsky, Austin Blanton, Brianna Maze, Jocelyn Adams, Tim Miller, Nathan Kalka, Anil K Jain, James A Duncan, Kristen Allen, et al. Iarpa janus benchmark-b face dataset. In *proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 90–98, 2017. [1](#)
- [64] Weidi Xie, Li Shen, and Andrew Zisserman. Comparator networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 782–797, 2018. [1](#), [2](#)
- [65] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. SimMIM: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. [1](#), [3](#), [2](#)
- [66] Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. Fanface: a simple orthogonal improvement to deep face recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 12621–12628, 2020. [1](#), [2](#), [8](#), [3](#)
- [67] Xingyu Yang, Mengya Han, Yong Luo, Han Hu, and Yonggang Wen. Two-stream prototype learning network for few-shot face recognition under occlusions. *IEEE Transactions on Multimedia*, 25:1555–1563, 2023. [2](#)
- [68] Yu-Feng Yu, Dao-Qing Dai, Chuan-Xian Ren, and Ke-Kun Huang. Discriminative multi-scale sparse coding for single-sample face recognition with occlusion. *Pattern Recognition*, 66:302–312, 2017. [2](#)
- [69] Wenbo Zheng, Chao Gou, and Fei-Yue Wang. A novel approach inspired by optic nerve characteristics for few-shot occluded face recognition. *Neurocomputing*, 376:25–41, 2020. [2](#)
- [70] Yaoyao Zhong and Weihong Deng. Face transformer for recognition. *arXiv preprint arXiv:2103.14803*, 2021. [1](#), [2](#), [7](#), [8](#)
- [71] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Jiwen Lu, Dalong Du, et al. Webface260m: A benchmark unveiling the power of million-scale deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10492–10502, 2021. [1](#), [5](#), [6](#), [7](#), [8](#)