

ProS: Facial Omni-Representation Learning via Prototype-based Self-Distillation

Xing Di¹, Yiyu Zheng¹, Xiaoming Liu², and Yu Cheng³

¹ProtagoLabs Inc.

²Michigan State University

³The Chinese University of Hong Kong

{xing.di,yiyu.zheng}@protagolabs.com, liuxm@cse.msu.edu, chengyu@cse.cuhk.edu.hk

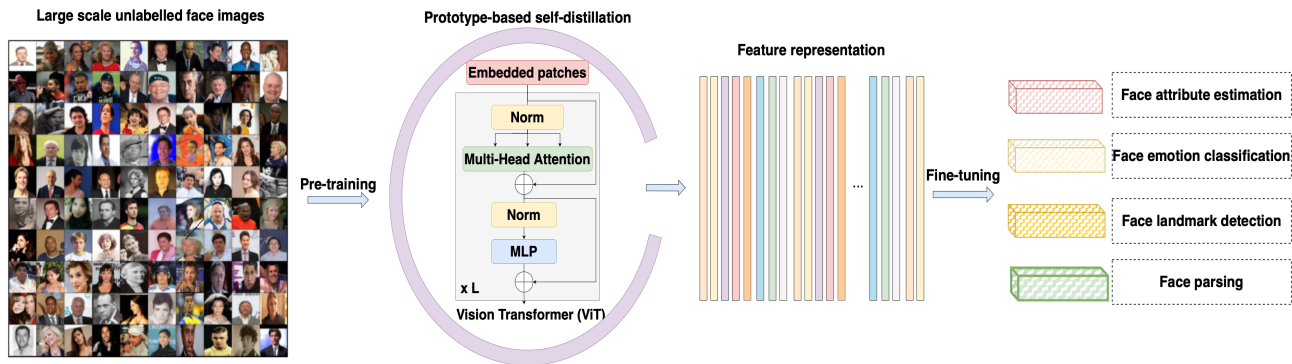


Figure 1. This paper presents a pre-training model that learns facial omni-representations via a **prototype-based self-distillation (ProS)**. For pre-training, ProS learns a general face representation from given large-scale **unlabeled** face images. Afterwards, the learned omni-representation can be conveniently utilized in multiple downstream tasks by simple fine-tuning.

Abstract

This paper presents a novel approach, called *Prototype-based Self-Distillation (ProS)*, for unsupervised face representation learning. The existing supervised methods heavily rely on a large amount of annotated training facial data, which poses challenges in terms of data collection and privacy concerns. To address these issues, we propose ProS, which leverages a vast collection of unlabeled face images to learn a comprehensive facial omni-representation. In particular, ProS consists of two vision-transformers (teacher and student models) that are trained with different augmented images (cropping, blurring, coloring, etc.). Besides, we build a face-aware retrieval system along with augmentations to obtain the curated images comprising predominantly facial areas. To enhance the discrimination of learned features, we introduce a prototype-based matching loss that aligns the similarity distributions between features (teacher or student) and a set of learnable prototypes. After pre-training, the teacher vision transformer serves as a backbone for downstream tasks, including attribute estimation, expression recognition, and landmark alignment, achieved through simple fine-tuning with additional layers. Extensive experiments demonstrate that

our method achieves state-of-the-art performance on various tasks, both in full and few-shot settings. Further, we investigate pre-training with synthetic face images, and ProS exhibits promising performance in this scenario as well.

1. Introduction

Learning good face representation is crucial for face analysis tasks such as face recognition [14, 44, 45, 52, 61, 66, 78–80], attribute estimation [8, 48, 70], expression classification [85, 93], landmark localization [32, 40, 82]. Among these tasks, existing state-of-the-art (SoTA) methods owe their success not only to the sophisticated network design but also large-scale training datasets. However, acquiring manually-annotated large-scale facial images is expensive and difficult for large labor-work and privacy issues [25]. For instance, it is hard to obtain the consent of all involved identities for face recognition datasets.

Recently, self-supervised learning has gained intensive interest due to the remarkable success of training generalizable models in both natural language processing [5, 16, 59, 60] and computer vision [1, 10, 12, 24, 26, 27, 54, 83, 90]. Such a pre-trained model has shown the following advantages: (a) the learned feature shows superiority on transfer-

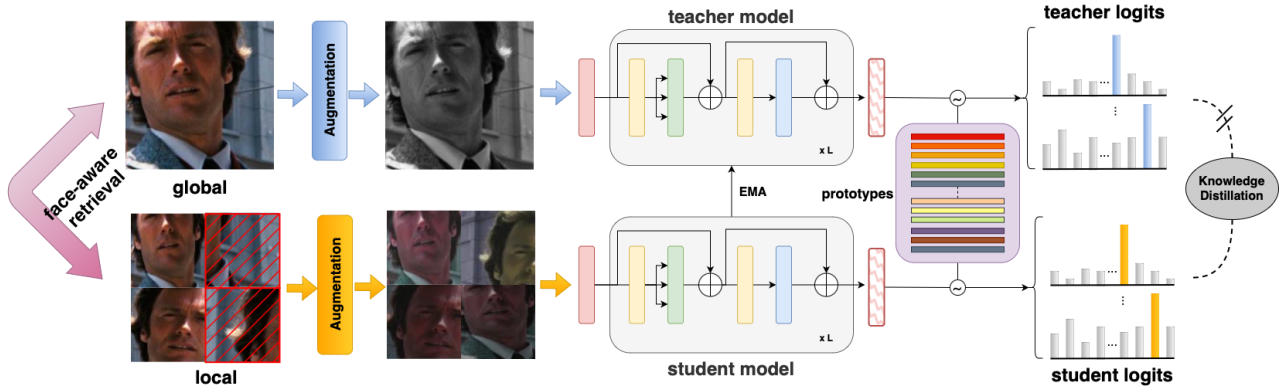


Figure 2. The proposed prototype-based self-distillation method. There are two branches: global and local. The global and local images are obtained through the multi-crops [9, 10]. To obtain the curated face images in local branch, we propose a face-aware retrieval system followed by the augmentations. The teacher and student models have the same vision-transformer architectures (ViT-S/16) but with different parameters. The self-knowledge-distillation between the student and teacher features is penalized via the similarity distribution between features and the learnable prototypes. By this, the networks are forced to leverage the mutual semantics between local and global views.

learning especially in few-shot settings, where it achieves a promising performance when data acquisition is limited. (b) the learned model is scalable for further development on diverse downstream tasks. To our acknowledgment, only few works [6, 95] explored the semi/self-supervised learning on face model. FaRL [95] explored a contrastive loss and masked image modeling for learning features from image-text pairs. FRL [6] learned the face representation based on ResNet [28] via SwAV [9].

Different from those previous methods, we propose a vision-transformer framework for learning face representation via **prototype-based self-distillation (ProS)**. As shown in Figure 1, ProS aims to learn the advanced feature representations given large-scale face images **without** labeling. In particular, ProS is trained in self-knowledge-distilling via a local-to-global manner. Our work is inspired by DINO [10] but with (i) a modified sample-to-prototype matching loss and (ii) a proposed face-aware retrieval system for curated data augmentation.

As shown in Figure 2, the global and local images are obtained from the same input image via multi-crop [9, 10] followed by the proposed face-aware retrieval system. The face-aware retrieval system aims to filter out most non-face images. Afterward, the curated images are fed into two sets of separate augmentations. The augmented images in global and local views are given to the teacher and student models respectively. By matching the features extracted from the teacher and student models, the loss gradient is back-propagated to the student model only for updating parameters. The parameters of the teacher model are updated through the exponential moving average [24].

During training, we observed that there are certain amounts of non-face images obtained in the local view as




Figure 3. The multi-cropped samples from MS1M [25]. We compare the (a) uncurated and (b) curated local samples with/without the facial-retrieval system. The global images are shown on top for reference. The non-face images (with red bounding boxes) are deduplicated.

highlighted by the red bounding boxes in Figure 3(a). Recent studies show those unidentifiable images are detrimental to the training procedure [37]. To eliminate those outliers, we build a face-aware retrieval system. In particular, we compute the face embeddings using the pre-trained ArcFace [14] for both the global and local images. The cosine similarity is used as a distance measure between local and global images. We displace those local face images with similarities lower than a threshold. We demonstrate two curated samples for the same input images in Figure 3(b).

After the face-aware retrieval system, we find there are lots of variations in the curated images. To boost the feature

discrimination, we introduce a set of learnable prototypes during training. Instead of directly computing the similarity of the features between the teacher and student samples, we compute two sample-to-prototype distributions: one between teacher samples to the prototypes, and the other one between student samples to the prototypes. We use the differences between these two distributions to penalize the model training. In this manner, both the prototypes and teacher-student models are optimized in every iteration where the features try to get close to positive prototypes and keep away from negative prototypes [15]. In addition, to mitigate the privacy issue of using real face data, we also explore synthesized face images for learning face representations. In particular, we simply train a StyleGAN2 [33, 36] from scratch on MS1M [25]. The synthetic data is generated by randomly-sampled noise from a normal distribution.

To this end, our contributions in this paper can be summarized as follows:

- We propose a novel pre-training framework (ProS) for learning facial omni-representation from large-scale face images without labeling. A learnable prototype-based matching loss and a face-aware retrieval system are introduced along with ProS.
- We conduct extensive experiments for evaluating ProS on various face analysis tasks. Our proposed ProS can achieve state-of-the-art results over different baselines on all the tasks in few-shot settings.
- We explore the capability of ProS on synthetic face images. To our best knowledge, ProS is the first to work on self-supervised pre-training on large-scale synthetic face images. We show that our method still obtains promising performances.

2. Related work

We walk through the related literature on self-supervised training, facial representation learning, and face synthesis in this section.

2.1. Self-supervised training

Self-supervised learning methods [3, 9, 10, 12, 13, 26, 54, 90, 96] have gained remarkable attention as an effective unsupervised learning strategy for learning robust image representations and eliminating the need to annotate vast quantities of data manually. For instance, SimCLR [12] maps the initial embeddings from two augmented views of an image into another space where the infoNCE loss is applied to encourage similarity between the views. DINO [10] feeds two different views of an image into the teacher and student encoders and maps the student network’s weights to the teacher’s by a moving average. SWAV [9] simultaneously clusters the data while enforcing consistency between

cluster assignments when given different augmented views of an image. In addition, MAE [26] and SimMIM [90] are two concurrent masked image modeling (MIM) that directly reconstruct masked image patches.

2.2. Facial representation learning

Existing face representation learning methods can be categorized into two classes: the proxy-based learning [14, 44, 45, 61, 78–80] and pair-wise learning [25, 52, 66]. As the class labels are known, proxy-based learning aims to optimize the similarity between given samples and a set of proxies representing each class. On the one hand, those methods [11, 14, 31, 37, 44, 73, 78, 80] put a margin penalty into the softmax loss and a global comparison between samples and proxies is conducted. On the other hand, those methods [53, 62, 66, 69, 71, 81, 87] involve different triplet strategies (representation selection, hard-mining) in mini or larger batch to leverage the underline pairwise information.

Recently, a few studies have been done for face analysis on few-shot [4], weakly-supervised [95], and self-supervised learning [6, 7, 86]. For instance, Browatzki *et al.* [4] proposed a few-shot face alignment framework with an image reconstruction by an auto encoder-decoder. Zheng *et al.* [95] proposed the FaRL for pre-training the vision-transformer model by leveraging the semantics between web-text and face image pairs. Wiles *et al.* [86] introduced a self-supervised manner for predicting the motion field between two facial images to learn efficient face representations. Vielzeuf *et al.* [77] introduced a common embedding for multi-source features from different trained models by an auto-encoding framework. Bulat *et al.* introduced the unsupervised training model FRL [6] for pre-training ResNet on the collected large-scale dataset.

2.3. Face synthesis

With the remarkable ability of GANs [22, 23], face synthesis has seen rapid developments, such as StyleGAN [35] and its variations [33, 34, 36] which can generate high-fidelity face images from random noises. Synthetic face data has shown significant improvements on various tasks such as learning pose-invariant models [17, 75, 91, 94], cross-spectrum models [18, 19, 56, 92], bridging domain gaps [43], serving as a training set [38], as well as reducing data bias [39, 58, 63, 69]. Unlike previous methods, we explore the possibility of using synthesized face images for self-supervised pre-training. We hope the synthetic data could be used as an alternative to real face images to avoid privacy issues when collecting data.

3. Proposed method

The proposed prototype-based self-distillation pre-training method is illustrated in Figure 2. In particular, it contains two branches: the global and the local. The global

and local images are multi-cropped from the same original image, resized at different scales and fed into the teacher and student models separately to obtain the corresponding features.

During the experiments, we find there are some non-face images cropped from the local branch which could diminish the discrimination of the learned features. Therefore, a face-aware retrieval system has been built for eliminating non-face images. Specifically, we utilize the pre-trained Arcface [14] to extract the features from both the global and local images. We filter out the local images that have a lower cosine distance between local-global features. In this paper, we set the distance threshold $\theta = -0.5 \in [-1.0, 1.0]$ by visual inspection of the retrieval results.

Inspired by recent findings [37, 68] that a "high" similarity score would be obtained from features of low-quality face images, a set of prototypes is utilized for penalizing the knowledge distillation. During each training iteration, the similarity scores are calculated between the teacher/student image features and these memorized prototype features instead of directly between the teacher/student features themselves. The cross-entropy loss is calculated between the similarity vectors. The student network parameters are optimized by back-propagating the loss gradient while the teacher network parameters are updated via an exponential moving average (EMA) of the student parameters [24]. Additionally, the prototypes are also optimized along with the network parameters by backpropagation.

3.1. Prototype-based self-distillation

Given a large-scale collection of unlabeled face images, the knowledge distillation aims to train the teacher and student models, parameterized as θ_t, θ_s , for matching the output features $f_{\theta_t}(x)$ and $f_{\theta_s}(x)$ given an input image x .

In each iteration, we sample a mini-batch of B images $\{\mathbf{x}_i\}_{i=1}^B$. The global images $\{\mathbf{x}_{m \rightarrow i}^g\}_{m=1}^M$ are obtained by a random crop of i -th original images followed by a set of global augmentations. Similarly, the local images $\{\mathbf{x}_{n \rightarrow i}^l\}_{n=1}^N$ are based on another random crop of i -th original image followed by a set of local augmentations. For brevity, we omit the captions as \mathbf{x}_m^g and \mathbf{x}_n^l . Following prior work [21], we "patchify" the input images into a set of sequential patches without overlapping. After, the global and corresponding local patches are fed into the teacher and student models respectively.

Let $f_{\theta_t}(\mathbf{x}_m^g) \in R^d$ and $f_{\theta_s}(\mathbf{x}_n^l) \in R^d$ denote the d -dim feature vectors obtained from teacher network and student network respectively. Additionally, a set of the learnable prototypes is denoted as $\mathbf{p} \in R^{K \times d}$. Instead of directly utilizing the teacher and student features, we use the cosine similarity between the student/teacher feature and these prototypes as the features. The prediction is calculated as fol-

lows:

$$s_n^l = \text{softmax}\left(\frac{\mathbf{p} \cdot f_{\theta_s}(\mathbf{x}_n^l)}{\tau_l}\right); s_m^g = \text{softmax}\left(\frac{\mathbf{p} \cdot f_{\theta_t}(\mathbf{x}_m^g)}{\tau_g}\right), \quad (1)$$

where \cdot denotes the dot product and $\tau_g \in (0, 1), \tau_l \in (0, 1)$. All the output features are L_2 normalized to mitigate the scale influence. To prevent the model collapse, we choose τ_g to be smaller than τ_l , and the global sharpening is utilized by softmax during training. The training objective is defined as follows:

$$\mathcal{L} = \frac{1}{B} \sum_{i=1}^B \mathcal{H}_i(s_m^g, s_n^l) - \mathcal{H}(\tilde{s}^l), \quad (2)$$

where $\mathcal{H}_i(s_m^g, s_n^l)$ is the i -th cross-entropy between s_m^g, s_n^l as Eq (3):

$$\mathcal{H}_i(s_m^g, s_n^l) = \frac{1}{M \cdot N} \sum_{m=1}^M \sum_{n=1}^N \mathcal{H}(s_m^g, s_n^l), \quad (3)$$

while $\mathcal{H}(\tilde{s}^l)$ is the entropy regularization [2] as Eq (4):

$$\mathcal{H}(\tilde{s}^l) = \frac{1}{B \cdot D} \sum_{i=1}^B \sum_{n=1}^N s_{n \rightarrow i}^l. \quad (4)$$

3.2. Model architecture

The teacher and student models are the vision-transformer encoders [21, 74]. For a fair comparison with other face analysis tasks, ViT-S/16 [74] is chosen, whose number of parameters is similar to the common ResNet-50 [28] (21M vs 23M). Specifically, ViT-S/16 is a 12-layer 384-width visual transformer with 224×224 resolution input. In our work, the global input images are 224×224 while the local images are 96×96 . These global/local images are firstly split into 14×14 and 6×6 patches respectively. Thus, one learnable *cls* token is prepended to the 196/36 embedding. In pre-training, additional 3 fully-connected layers are added as the projector to the output transformer for further optimization by Eq. (2). The prototypes are a set of learnable variables with random initialization. We set the output feature dimension as 256 as the prototypes.

3.3. Pre-training details

The teacher and student models are trained from scratch with randomly initialized parameters. The total training runs 20 epochs with a total batch size 64×4 on 4 Nvidia 3090 GPUs. The AdamW optimizer is utilized with weight decay as 0.04. The initial learning rate is 0.0002 with 2 warmed-up epochs to 0.001 and then cosine decay to $1e-6$ in the next 18 epochs. The teacher/student temperatures are set as 0.025 and 0.1. The number of prototypes is set



Figure 4. Samples of synthetic face images.

to 1,024. The learnable prototypes are randomly initialized with a uniform distribution between $[-1/\sqrt{d}, 1/\sqrt{d}]$, where d is the output dimension here. The prototypes are updated iteratively along with the network parameters by back-propagation. To make use of all the unlabeled images, we directly input the raw face images for pre-training without any further preprocessing like face detection, cropping, or alignment.

3.4. Synthetic data

We train the original StyleGAN2 [36] on MS1M dataset [25] to obtain the synthetic data. In particular, the images are all resized to 256×256 to train the adversarial generative networks. The “paper256” setting [33] is utilized. When training on different data sizes, images are randomly selected. Once the training is completed, the synthetic face images are obtained through the corresponding generator via input noise vectors, which are sampled from a standard normal distribution. Samples of the synthesized face images are shown in Figure 4.

3.5. Downstream tasks

Face attributes recognition is a multi-class classification task that predicts multiple facial attributes (*e.g.* gender, race, hair color) given one facial image. In this work, we evaluate the pre-trained model on two datasets: CelebA [48] and LFWA [30, 48], containing 202, 599 and 13, 143 images respectively. They both have 40 annotated attributes per image. Following the protocols [48, 70, 95] we use 162, 770 images for training and 19, 962 for testing on CelebA, and 6, 263 images for training and the rest for testing on LFWA. Other than the initial weights from ProS, we follow the same protocols and the average accuracy on all attributes is reported.

Face expression recognition is a single-class classification task that estimates one facial expression (*e.g.* happy, anger, disguising) of a given face image. We evaluate the pre-trained model on two datasets: RAF-DB [41, 42] and AffectNet [51]. RAF-DB contains around 29, 670 face images from real-world databases, of which 15, 339 images for 7 expression classifications. There are 12, 271 images for training and the remaining 3, 068 for testing. AffectNet [51] is a large-scale database for facial expressions. We use the most challenging AffectNet8 (including the additional “contempt” category) data, with 287, 651 training images and 3, 999 testing images. The average accuracy from

all emotion classes is used as the evaluation metric.

Face alignment targets to regress the 2D coordinates of face landmarks on a face image. We evaluate our proposed model on two popular datasets: 300W [64, 65] and WFLW [88]. 300W dataset contains 68 landmarks per face with 3, 837 training images and 600 testing images. The WFLW dataset contains 68 landmarks as well, with 7, 500 training and 2, 500 testing samples. We measure the performance by the normalized mean error (NME).

4. Experimental results

Besides the results in this section, more ablation studies, like “baseline methods pre-trained on face dataset” and “numbers of prototypes”, can be found in *supplementary material*.

4.1. Implementation

After pre-training, the teacher model is used for downstream tasks training with additional head(s), both with end-to-end fine-tuning or head-only fine-tuning. For different tasks, the head designs are slightly varied. We donate the features from h -th head of k -th layer, including the last and intermediate layers, of the visual transformer as $feat_k = \{f_{cls,k}, f_{1,k}, f_{2,k}, \dots, f_{h,k}\}$, where $k = \{1, 2, \dots, 12\}$.

In particular, we use the multi-task classifiers [8] for face attributes classification. The cls-token feature vector from the last layer ($f_{cls,12}$) is layer-normalized and appended with 40 separate linear layers to generate the logits for binary classification on each attribute. The model is trained with the averaged binary-cross-entropy loss on each head and is optimized by AdamW [49]. We set the effective learning rate as $5e-4$, weight decay as 0.05, and layer decay as 0.65. The learning rate decreases to zero in 100 epochs.

For face expression recognition, we use the original ViT-S/16 but change the last linear layer output dimension. Specifically, the output vector dimension for RAF-DB is set to 7, and AffectNet8 is set to 8. When fine-tuning, the learning rate is $5e-4$, the weight decay is 0.5, and the layer decay is 0.65. The total training epochs are 100 for RAF-DB and 10 for AffectNet8. Like previous settings [85], we use the imbalanced data sampler for AffectNet8.

For face alignment, the ground-truth landmarks are rendered as Gaussian heat-map at a size of 128×128 with σ and values $\in [0, 1]$ [32, 82]. The non-cls tokens on layers $\{4, 6, 8, 12\}$ are utilized. To leverage the spatial distribution of these tokens, we reshape each to the 2D feature map of 14×14 . The UperNet [89] is followed to fuse these feature maps from each layers to a final heat-map logits [47, 95]. Following the prior work [95], a simple soft-label cross-entropy loss is utilized for training the model. We use the AdamW with a learning rate of 0.01 and a weight decay of $1e-5$.

Table 1. Performance comparison of pre-trained models on various downstream tasks. We choose ViT-16/B for MAE due to the availability. [Keys: **Best**, **Second best**]

Method	Model - # of params	Pre-train Datasets	Data Scale	Supervision	Downstream Tasks					
					CelebA [48]		RAF-DB [41,42]		300W [64,65]	
					mAcc. ↑		mAcc. ↑		NME _{inter-ocular} ↓	
					100%	1%	100%	10%	100%	10%
DeiT [74]	ViT-S/16 - 21M	ImageNet-1K	1.3M	images, labels	90.79	88.27	87.87	75.36	3.40	4.34
MAE* [26]	ViT-B/16 - 86M	ImageNet-1K	1.3M	images	91.16	90.17	88.33	76.84	3.36	4.13
DINO [10]	ViT-S/16 - 21M	ImageNet-1K	1.3M	images	91.25	89.62	88.23	75.85	3.53	4.57
MSN [1]	ViT-S/16 - 21M	ImageNet-1K	1.3M	images	91.17	89.99	87.81	76.21	3.48	4.26
FRL [6]	ResNet50 - 23M	Large-Scale-Face	5M	face images	91.04	90.04	90.07	80.57	3.85	4.25
ProS-full-real	ViT-S/16 - 21M	MS1M	8.6M	face images	91.88	90.86	91.04	82.10	3.27	3.92

Table 2. Comparison with baseline methods of face attribute estimation on CelebA and LFWA datasets with limited data. The mAcc. ↑ is used as the evaluation metric. [Keys: **SoTA**, **Best**, **Second best**]

Dataset	CelebA [48]					LFWA [30,48]				
Portion	0.2%	0.5%	1%	2%	100%	5%	10%	20%	50%	100%
# of training data	325	843	1,627	3,255	162,770	313	626	1,252	3,131	6,263
PS-MCNN [8]	-	-	-	-	92.98	-	-	-	-	87.36
Slim-CNN [67]	79.90	80.20	80.96	82.32	91.24	70.90	71.49	72.12	73.45	76.02
FixMath [72]	80.22	84.19	85.77	86.14	89.78	71.42	72.78	75.10	80.87	83.84
VAT [50]	81.44	84.02	86.30	87.28	91.44	72.19	74.42	76.26	80.55	84.68
SSPL [70]	86.67	88.05	88.84	89.58	91.77	78.68	81.65	83.45	85.43	86.53
FARL [95]	88.51	89.12	90.24	90.55	91.88	82.57	83.58	84.80	85.95	86.69
ProS-1M-syn	88.60	89.78	90.57	90.92	91.57	82.69	83.92	85.50	86.75	86.83
ProS-1M-real	88.70	90.15	90.72	91.08	91.58	82.73	84.57	85.24	86.79	87.06
ProS-full-real	88.76	90.43	90.86	91.17	91.88	83.25	85.13	86.25	86.85	87.08

4.2. Comparing with other pre-training models

We clarify our models under different settings:

- ProS-1M-syn: pre-trained with 1M synthetic images from the generator, which is trained with randomly selected 1M real images.
- ProS-1M-real: pre-trained with randomly selected 1M real images from MS1M dataset.
- ProS-full-real: pre-trained with all real images from MS1M dataset.

We investigate how the pre-training models influence the downstream tasks’ performance in Table 1. We compare the models from different architectures on various datasets. In particular, five different pre-training models are included: (1) DeiT [74]: was the improved ViT trained on ImageNet-1K with distillation under full supervision. (2) MAE [26]: was an auto-encoder learner for images reconstructed from masked input. It was trained on ImageNet-1K with self-supervision. (3) DINO [10] was trained on ImageNet-1K as a form of mean teacher self-distillation under image self-supervision. (4) MSN [1] was a masked Siamese network trained on ImageNet-1K with self-supervised learning. (5) FRL¹ [6] trained ResNet50 [28] on a large-scale

¹<https://github.com/ladrianb/unsupervised-face-representation>

face dataset without labels. For a fair comparison, we use ViT-S/16 as the backbone for DeiT, DINO, and MSN but ViT-B/16 for MAE due to the availability. We compare the proposed method with those baselines in the downstream tasks as illustrated in Table 1. As we can observe, all these models show a reasonable performance. The ProS-full-real shows superior performance over both fully-supervised and self-supervised methods. The ProS-1M-syn shows a competitive performance to the other baselines as well.

4.3. Comparing with state-of-the-art face methods

We compare our proposed model with other SoTA methods in both full-shot and few-shot settings in multiple downstream tasks. All the input images are resized to 224×224 and the official aligned version (if applicable) is used. We conduct all the following experiments five times and report the average performance.

Face attributes recognition We compare our proposed method with baseline methods under both full-shot and few-shot settings. As we can observe in Table 2, our proposed method shows superiority over all the other methods in few-shot and rank the 2-nd under the full-shot with FaRL [95]. Note the PS-MCNN-LC [8] achieved a higher accuracy by using extra identity labels and a fine-grained network design to leverage the attribute relation. Meanwhile, we can observe that our method indeed benefits from a larger data scale (ProS-1M-real vs ProS-full-real). With only 1M syn-

Table 3. Comparison with SoTA results of facial expression recognition on AffectNet8 and RAF-DB datasets. The mAcc. \uparrow is used here as the evaluation metric. [Keys: SoTA, Best, Second best]

Methods	AffectNet8 [51]			RAF-DB [41,42]			
	2%	10%	Full	1%	2%	10%	Full
EAC [93]	-	-	63.11	57.95	64.05	82.07	89.99
DAN [85]	43.16	52.41	62.09	53.17	58.46	78.05	89.70
ProS-1M-syn	43.46	49.96	62.59	58.74	66.13	80.11	89.06
ProS-1M-real	43.64	50.16	63.44	61.04	67.60	80.32	89.83
ProS-full-real	45.91	50.66	63.64	63.06	70.61	82.10	91.04

thetic data, ProS-1M-syn also outperforms the baselines in all the few-shot settings. It is impressive to see that when trained with only 50% data in LFWA, our proposed methods are still better than the close competitor FaRL [95] in 100% data usage. When comparing results from ProS-1M-syn vs ProS-1M-real, the synthetic face data give competitive results with the one with real data in both settings. In general, the largest real data model ProS-full-real outperforms both ProS-1M-syn and ProS-1M-real, which achieves the best results among all in the few-shot settings.

Face expression recognition We conduct another set of experiments on facial expression recognition. Similar to previous experiments, we evaluate in both full-shot and few-shot settings. Due to the limited models for few-shot face expression recognition, we compare the proposed models with two SoTA baselines: DAN [85] (ResNet-18) and EAC [93] (ResNet-50). For the full-shot results, the results are copied from the paper report. For the few-shot results, we run the experiments with their published codes²³. For a fair comparison, both DAN and EAC models were initialized with trained weights from fully-supervised training on MS1M dataset. As shown in Table 3, the ProS-full-real model achieves a better performance than the EAC and DAN methods in both RAF-DB and AffectNet8 datasets under both full and limited data settings. Similarly, we can observe the larger pre-training data gives better performance (ProS-full-real vs ProS-1M-real). When comparing the AffectNet8 results of DAN using 100% and 10% training data to the ones from our ProS, one reason for the degradation on 100% could be the long-tail bias [45, 55] from the MS1M weights in fully-supervised training.

Face alignment We also evaluate the ProS models on face alignment tasks using WFLW and 300W test-set. As shown in Table 4, our ProS-full-real model surpasses all baselines under the full settings and most under the few-shot settings. We also include the SoTA (FaRL [95]) method in this table, which was trained in semi-supervision on 20M web-text and image pairs while we are self-supervision on fewer face images only. For the close competitor FRL [6], it only surpasses ours under the 0.7% few-shot setting on WFLW.

²<https://github.com/yaoing/DAN>

³<https://github.com/zyh-uiaaaa/Erasing-Attention-Consistency>

When comparing ProS-1M-real and ProS-1M-syn, we can see using the real face images leads to a better result.

4.4. Visualization of learned prototypes

In order to exhibit the variances captured by the learned prototypes, we perform a t-SNE visualization [76] on the set of 1024 prototypes derived from the ProS-1M-real and ProS-full-real. Specifically, we employ nearest-neighbor retrieval to establish connections between the learned prototypes and the training images. Figure 5 illustrates the results, indicating that the learned prototypes are predominantly dispersed in a sparse distribution. Comparing with Figure 5a and Figure 5b, we can observe that the learned prototypes show a better span coverage when using the full size of the training dataset.

4.5. Visualization of pre-training model

In order to investigate the knowledge acquired through the pre-training and fine-tuning stages of our proposed ProS model, we visualize the attention mechanism using heatmaps, as depicted in Figure 6. Specifically, we use the teacher model of ProS-full-real and analyze the heatmap values obtained from the attention heads in the final self-attention layer. During the pre-training phase, our ProS model successfully localizes the facial region in input images, effectively capturing various variations such as pose and scale. These heatmaps exhibit a similar pattern to human attention patterns. Comparing the heatmaps generated by the pre-trained model and the fine-tuned models, we observe that the features learned through fine-tuning become more task-specific, while the features derived from the pre-trained model maintain a more generalized representation.

5. Discussion

Ethical considerations For training, ProS requires large-scale face images. We use the MS1M data, one of the largest datasets in the world, collected by Microsoft in 2016. The images in this dataset are scraped off the web under the terms of the Creative Commons license and are limited to the academic usage of the photos. Studies [97] show that the MS1M dataset is subject to biases due to label noise, duplicate images, and non-face images. It is essential to acknowledge and address these biases to prevent the propagation of unfair or discriminatory practices in the development and deployment of face recognition systems.

Social Impact and Limitation We have identified certain limitations in the pre-training phase of our ProS model. Specifically, we found that the bias issue present in the MS1M dataset is still present in the learned representations of ProS. Moreover, our model does not perform exceptionally well in face parsing tasks on the LaPa dataset [46], as evidenced in the supplementary documents. One possible

Table 4. Comparison of facial alignment on WFLW and 300W (test-set) dataset. The results of each column in 300W stand for the Common, Challenge, and Full subsets respectively. The $NME_{inter-ocular} \downarrow$ is used here as the evaluation metric. [Keys: **SoTA**, **Best**, **Second best**]

Methods	WFLW [88]					300W [64,65]								
	0.7%	5%	10%	20%	100%	1.5%			10%			100%		
FaRL [95]	6.02	4.83	4.55	4.33	4.03	5.87	3.24	3.76	2.81	4.83	3.21	2.56	4.45	2.93
RCN+ [29]	-	-	-	-	-	-----	-----	-----	-----	6.63	4.47	3.00	4.98	3.46
SA [57]	-	-	7.20	6.00	4.39	-----	-----	-----	-----	4.27	3.21	6.49	3.86	
TS ³ [20]	-	-	-	-	-	-----	-----	-----	4.67	9.26	5.64	2.91	5.90	3.49
3FabRec [4]	8.39	7.68	6.73	6.51	5.62	4.55	7.39	5.10	3.88	6.88	4.47	3.36	5.74	3.82
FSMA [84]	-	-	-	-	-	-----	-----	-----	3.59	7.01	4.45	3.12	6.14	3.88
FRL [6]	7.11	-	5.44	-	4.57	-----	-----	-----	-----	4.25	-----	-----	-----	3.85
ProS-1M-syn	8.63	6.25	5.70	5.18	4.55	4.65	8.89	5.49	3.70	6.57	4.25	2.95	5.06	3.36
ProS-1M-real	8.13	6.08	5.56	5.11	4.51	4.56	8.55	5.35	3.54	6.20	4.06	2.90	5.05	3.32
ProS-full-real	7.73	5.75	5.30	4.90	4.37	4.38	7.88	5.08	3.41	6.01	3.92	2.86	4.92	3.27




Figure 5. t-SNE visualization [76] of learned prototypes by finding the nearest neighbor in synthesized training face images.

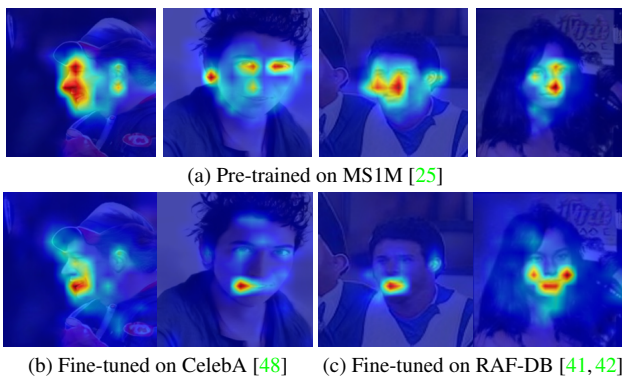


Figure 6. Comparison of attention heatmaps of teacher models from pre-trained (a) on MS1M [25]; and fine-tuned (b) on CelebA [48] and (c) on RAF-DB [41, 42] respectively.

explanation for this performance gap is that the learned features in ProS tend to be more semantic-specific rather than

spatial-specific. This observation is further supported by Figure 6, where we can observe that the attention mechanism does not adequately attend to the hair region.

6. Conclusion

In this paper, we present a self-supervised pre-training method (ProS) for learning face representation from unlabeled large-scale images only. One modified prototype-based matching loss and a face-aware retrieval system are introduced along with the ProS. We explore the ProS on both real and synthetic face images. In addition, we show the face representations learned from ProS can be well-transferred to multiple downstream face analysis tasks including attribute estimation, expression recognition, and face alignment. Compared with SoTA methods, our proposed ProS shows the superiority of performance in the limited data. Moreover, the proposed method surpasses the previous SoTA methods in facial expression recognition.

References

- [1] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Michael Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. *arXiv preprint arXiv:2204.07141*, 2022. 1, 6
- [2] Adrien Bardes, Jean Ponce, and Yann LeCun. Vircreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021. 4
- [3] Adrien Bardes, Jean Ponce, and Yann Lecun. Vircreg: Variance-invariance-covariance regularization for self-supervised learning. In *ICLR 2022-10th International Conference on Learning Representations*, 2022. 3
- [4] Bjorn Browatzki and Christian Wallraven. 3fabrec: Fast few-shot face alignment by reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6110–6120, 2020. 3, 8
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 1
- [6] Adrian Bulat, Shiyang Cheng, Jing Yang, Andrew Garbett, Enrique Sanchez, and Georgios Tzimiropoulos. Pre-training strategies and datasets for facial representation learning. In *ECCV*, 2022. 2, 3, 6, 7, 8
- [7] Zhixi Cai, Shreya Ghosh, Kalin Stefanov, Abhinav Dhall, Jianfei Cai, Hamid Rezatofighi, Reza Haffari, and Munawar Hayat. Marlin: Masked autoencoder for facial video representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1493–1504, June 2023. 3
- [8] Jiajiong Cao, Yingming Li, and Zhongfei Zhang. Partially shared multi-task convolutional neural network with local constraint for face attribute learning. In *Proceedings of the IEEE Conference on computer vision and pattern recognition*, pages 4290–4299, 2018. 1, 5, 6
- [9] 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. 2, 3
- [10] 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, 2, 3, 6
- [11] Binghui Chen, Weihong Deng, and Junping Du. Noisy softmax: Improving the generalization ability of dcnn via postponing the early softmax saturation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5372–5381, 2017. 3
- [12] 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. 1, 3
- [13] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020. 3
- [14] 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. 1, 2, 3, 4
- [15] 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. 3
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1
- [17] Xing Di, Shuowen Hu, and Vishal M Patel. Heterogeneous face frontalization via domain agnostic learning. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pages 01–08. IEEE, 2021. 3
- [18] Xing Di, Benjamin S Riggan, Shuowen Hu, Nathaniel J Short, and Vishal M Patel. Polarimetric thermal to visible face verification via self-attention guided synthesis. In *2019 International Conference on Biometrics (ICB)*, pages 1–8. IEEE, 2019. 3
- [19] Xing Di, He Zhang, and Vishal M Patel. Polarimetric thermal to visible face verification via attribute preserved synthesis. In *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–10. IEEE, 2018. 3
- [20] Xuanyi Dong and Yi Yang. Teacher supervises students how to learn from partially labeled images for facial landmark detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 783–792, 2019. 8
- [21] 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. 4
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. 3
- [23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 3
- [24] 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. 1, 2, 4
- [25] Yandong Guo, Lei Zhang, Yuxiao Hu, X. He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *ECCV*, 2016. 1, 2, 3, 5, 8
- [26] 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. 1, 3, 6
- [27] 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. 1
- [28] 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, 4, 6
- [29] Sina Honari, Pavlo Molchanov, Stephen Tyree, Pascal Vincent, Christopher Pal, and Jan Kautz. Improving landmark localization with semi-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1546–1555, 2018. 8
- [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. 5, 6
- [31] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, and Feiyue Huang Jilin Li. Curricularface: Adaptive curriculum learning loss for deep face recognition. *CVPR*, pages 1–8, 2020. 3
- [32] Yangyu Huang, Hao Yang, Chong Li, Jongyoo Kim, and Fangyun Wei. Adnet: Leveraging error-bias towards normal direction in face alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3080–3090, 2021. 1, 5
- [33] 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, 5
- [34] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. 3
- [35] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 3
- [36] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. 3, 5
- [37] 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, 3, 4
- [38] Minchul Kim, Feng Liu, Anil Jain, and Xiaoming Liu. DC-Face: Synthetic face generation with dual condition diffusion model. In *CVPR*, 2023. 3
- [39] Adam Kortylewski, Bernhard Egger, Andreas Schneider, Thomas Gerig, Andreas Morel-Forster, and Thomas Vetter. Analyzing and reducing the damage of dataset bias to face recognition with synthetic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 3
- [40] Abhinav Kumar, Tim K Marks, Wenxuan Mou, Ye Wang, Michael Jones, Anoop Cherman, Toshiaki Koike-Akino, Xiaoming Liu, and Chen Feng. Luvli face alignment: Estimating landmarks’ location, uncertainty, and visibility likelihood. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8236–8246, 2020. 1
- [41] Shan Li and Weihong Deng. Reliable crowdsourcing and deep locality-preserving learning for unconstrained facial expression recognition. *IEEE Transactions on Image Processing*, 28(1):356–370, 2019. 5, 6, 7, 8
- [42] Shan Li, Weihong Deng, and JunPing Du. Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2584–2593. IEEE, 2017. 5, 6, 7, 8
- [43] Feng Liu, Minchul Kim, Anil Jain, and Xiaoming Liu. Controllable and guided face synthesis for unconstrained face recognition. In *ECCV*, 2022. 3
- [44] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. 1, 3
- [45] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *ICML*, 2016. 1, 3, 7
- [46] Yinglu Liu, Hailin Shi, Hao Shen, Yue Si, Xiaobo Wang, and Tao Mei. A new dataset and boundary-attention semantic segmentation for face parsing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11637–11644, 2020. 7
- [47] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 5
- [48] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. 1, 5, 6, 8
- [49] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5
- [50] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. 6

- [51] A. Mollahosseini, B. Hasani, and M. H. Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, PP(99):1–1, 2017. 5, 7
- [52] Aaron Nech and Ira Kemelmacher-Shlizerman. Level playing field for million scale face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7044–7053, 2017. 1, 3
- [53] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4004–4012, 2016. 3
- [54] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023. 1, 3
- [55] Wanli Ouyang, Xiaogang Wang, Cong Zhang, and Xiaokang Yang. Factors in finetuning deep model for object detection with long-tail distribution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 864–873, 2016. 7
- [56] Domenick Poster, Matthew Thielke, Robert Nguyen, Srinivasan Rajaraman, Xing Di, Cedric Nimpa Fondje, Vishal M Patel, Nathaniel J Short, Benjamin S Riggan, Nasser M Nasrabadi, et al. A large-scale, time-synchronized visible and thermal face dataset. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1559–1568, 2021. 3
- [57] Shengju Qian, Keqiang Sun, Wayne Wu, Chen Qian, and Jia-ya Jia. Aggregation via separation: Boosting facial landmark detector with semi-supervised style translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10153–10163, 2019. 8
- [58] Haibo Qiu, Baosheng Yu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. Synface: Face recognition with synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10880–10890, 2021. 3
- [59] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 1
- [60] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 1
- [61] Rajeev Ranjan, Carlos D Castillo, and Rama Chellappa. L2-constrained softmax loss for discriminative face verification. *arXiv preprint arXiv:1703.09507*, 2017. 1, 3
- [62] Oren Rippel, Manohar Paluri, Piotr Dollar, and Lubomir Bourdev. Metric learning with adaptive density discrimination. *ICLR*, 2016. 3
- [63] Nataniel Ruiz, Barry-John Theobald, Anurag Ranjan, Ahmed Hussein Abdelaziz, and Nicholas Apostoloff. Morphgan: One-shot face synthesis gan for detecting recognition bias. *arXiv preprint arXiv:2012.05225*, 2020. 3
- [64] Christos Sagonas, Epameinondas Antonakos, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: Database and results. *Image and vision computing*, 47:3–18, 2016. 5, 6, 8
- [65] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 397–403, 2013. 5, 6, 8
- [66] 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, 3
- [67] Ankit Kumar Sharma and Hassan Foroosh. Slim-cnn: A light-weight cnn for face attribute prediction. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 329–335. IEEE, 2020. 6
- [68] Yichun Shi and Anil K Jain. Probabilistic face embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6902–6911, 2019. 4
- [69] Yichun Shi, Xiang Yu, Kihyuk Sohn, Manmohan Chandraker, and Anil K Jain. Towards universal representation learning for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6817–6826, 2020. 3
- [70] Ying Shu, Yan Yan, Si Chen, Jing-Hao Xue, Chunhua Shen, and Hanzi Wang. Learning spatial-semantic relationship for facial attribute recognition with limited labeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11916–11925, 2021. 1, 5, 6
- [71] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29, 2016. 3
- [72] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 6
- [73] Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6398–6407, 2020. 3
- [74] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. 4, 6
- [75] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1415–1424, 2017. 3
- [76] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 7, 8

- [77] Valentin Vielzeuf, Alexis Lechery, Stéphane Pateux, and Frédéric Jurie. Towards a general model of knowledge for facial analysis by multi-source transfer learning. *arXiv preprint arXiv:1911.03222*, 2019. **3**
- [78] Feng Wang, Jian Cheng, Weiyang Liu, and Haijun Liu. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25(7):926–930, 2018. **1, 3**
- [79] Feng Wang, Xiang Xiang, Jian Cheng, and Alan Loddon Yuille. Normface: L2 hypersphere embedding for face verification. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1041–1049, 2017. **1, 3**
- [80] 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, 3**
- [81] Jing Wang, Yu Cheng, and Rogério Schmidt Feris. Walk and learn: Facial attribute representation learning from ego-centric video and contextual data. *CVPR*, pages 2295–2304, 2016. **3**
- [82] Xinyao Wang, Liefeng Bo, and Li Fuxin. Adaptive wlog loss for robust face alignment via heatmap regression. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6971–6981, 2019. **1, 5**
- [83] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14668–14678, 2022. **1**
- [84] Zhen Wei, Bingkun Liu, Weinong Wang, and Yu-Wing Tai. Few-shot model adaptation for customized facial landmark detection, segmentation, stylization and shadow removal. *arXiv preprint arXiv:2104.09457*, 2021. **8**
- [85] Zhengyao Wen, Wenzhong Lin, Tao Wang, and Ge Xu. Distract your attention: multi-head cross attention network for facial expression recognition. *arXiv preprint arXiv:2109.07270*, 2021. **1, 5, 7**
- [86] Olivia Wiles, A Koepke, and Andrew Zisserman. Self-supervised learning of a facial attribute embedding from video. *arXiv preprint arXiv:1808.06882*, 2018. **3**
- [87] Chao-Yuan Wu, R Manmatha, Alexander J Smola, and Philipp Krahenbuhl. Sampling matters in deep embedding learning. In *Proceedings of the IEEE international conference on computer vision*, pages 2840–2848, 2017. **3**
- [88] Wayne Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, and Qiang Zhou. Look at boundary: A boundary-aware face alignment algorithm. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2129–2138, 2018. **5, 8**
- [89] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the European conference on computer vision (ECCV)*, pages 418–434, 2018. **5**
- [90] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simsim: 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**
- [91] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. Towards large-pose face frontalization in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3990–3999, 2017. **3**
- [92] Aijing Yu, Haoxue Wu, Huaibo Huang, Zhen Lei, and Ran He. Lamp-hq: A large-scale multi-pose high-quality database and benchmark for nir-vis face recognition. *International Journal of Computer Vision*, 129(5):1467–1483, 2021. **3**
- [93] Yuhang Zhang, Chengrui Wang, Xu Ling, and Weihong Deng. Learn from all: Erasing attention consistency for noisy label facial expression recognition. In *Proceedings of the European conference on computer vision (ECCV)*, 2022. **1, 7**
- [94] Jian Zhao, Lin Xiong, Jianshu Li, Junliang Xing, Shuicheng Yan, and Jiashi Feng. 3d-aided dual-agent gans for unconstrained face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 41(10):2380–2394, 2018. **3**
- [95] Yinglin Zheng, Hao Yang, Ting Zhang, Jianmin Bao, Dongdong Chen, Yangyu Huang, Lu Yuan, Dong Chen, Ming Zeng, and Fang Wen. General facial representation learning in a visual-linguistic manner. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18697–18709, 2022. **2, 3, 5, 6, 7, 8**
- [96] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Image bert pre-training with online tokenizer. In *ICLR 2022-10th International Conference on Learning Representations*, 2022. **3**
- [97] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Dalong Du, Jiwen Lu, et al. Webface260m: A benchmark for million-scale deep face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. **7**