

# SCE-MAE: Selective Correspondence Enhancement with Masked Autoencoder for Self-Supervised Landmark Estimation

Kejia Yin<sup>1,2†\*</sup> Varshanth Rao<sup>2\*</sup> Ruowei Jiang<sup>2</sup> Xudong Liu<sup>1,2†</sup> Parham Aarabi<sup>1</sup> David B. Lindell<sup>1</sup>  
<sup>1</sup>University of Toronto <sup>2</sup>ModiFace

## Abstract

Self-supervised landmark estimation is a challenging task that demands the formation of locally distinct feature representations to identify sparse facial landmarks in the absence of annotated data. To tackle this task, existing state-of-the-art (SOTA) methods (1) extract coarse features from backbones that are trained with instance-level self-supervised learning (SSL) paradigms, which neglect the dense prediction nature of the task, (2) aggregate them into memory-intensive hypercolumn formations, and (3) supervise lightweight projector networks to naïvely establish full local correspondences among all pairs of spatial features. In this paper, we introduce SCE-MAE, a framework that (1) leverages the MAE [14], a region-level SSL method that naturally better suits the landmark prediction task, (2) operates on the vanilla feature map instead of on expensive hypercolumns, and (3) employs a Correspondence Approximation and Refinement Block (CARB) that utilizes a simple density peak clustering algorithm and our proposed Locality-Constrained Repellence Loss to directly hone only select local correspondences. We demonstrate through extensive experiments that SCE-MAE is highly effective and robust, outperforming existing SOTA methods by large margins of  $\sim 20\%$ - $44\%$  on the landmark matching and  $\sim 9\%$ - $15\%$  on the landmark detection tasks.

## 1. Introduction

Facial landmark detection is a computer vision task involving the identification and localization of specific keypoints corresponding to particular positions on a human face. Facial landmarks form the crux for many classical downstream tasks such as 3D face reconstruction [1, 48], face recognition [18, 41], face emotion/expression recognition [32, 33], and more contemporary applications such as facial beauty prediction [4, 16] and face make-up try on [20, 24, 31].

Albeit extremely useful, training facial landmark de-

\*Equal contribution.

†This work was done during the internship at ModiFace.

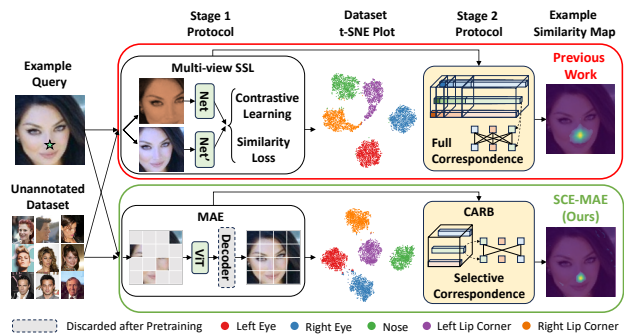


Figure 1. SCE-MAE vs prior self-supervised facial landmark detection methods. **Stage 1:** Prior works (top) use instance-level multi-view SSL paradigms that output less distinct initial local features. Our framework (bottom) leverages MAE to naturally form better initial features that result in *well-defined boundaries* between facial landmarks (see t-SNE plots). **Stage 2:** Prior works operate on *memory-intensive hypercolumns* and supervise *each feature pair* to achieve correspondence. Our framework employs a Correspondence Approximation and Refinement Block (CARB) that operates on the original MAE output and directly hones *only the selected correspondence pairs*. For the example query, SCE-MAE outputs a more-focused and sharper similarity map, demonstrating the superiority of the final features.

tectors requires numerous precise annotations per sample, making it a laborious and expensive ordeal. Furthermore, landmarks are not always semantically well-defined, making their annotations prone to inconsistencies and biases [15, 23, 58], which can severely limit the development of accurate landmark models. Motivated to avoid these demerits, recent works [9, 19, 46] have incorporated the unsupervised [46] and self-supervised learning (SSL) paradigms [12, 13, 47] into their methods. SSL-pretrained models have shown to yield highly effective feature representations without the use of labeled data and, at many times, outperform their supervised counterparts on the target tasks [2, 5, 14].

Facial landmark detection and matching tasks rely on the formation of locally distinct features to differentiate between (1) the facial regions (e.g., eye vs. lip), (2) the components of face parts (e.g., left vs. right corners of the lip),

and finally, (3) the specific pixels of each landmark. In the setting where annotations are severely limited, the recent methods [9, 19] follow a two-stage training protocol. During the first stage, the backbone is trained with a typical SSL objective. In the second stage, the backbone is frozen and a separate light-weight projector network is trained to encode *local correspondences*, i.e., the relationships between the different regions within the same image.

Prior work adopted multi-view SSL protocols [12, 13], which may be less effective on the landmark estimation tasks due to several factors. Firstly, these augment-and-compare pretext tasks prompt the network to learn category-specific signals, but we operate only on a single category, i.e., the human face. Secondly, contrastive learning requires a large and diverse set of negative samples to avoid collapse [52, 53]. Lastly, the training objectives might not directly encourage the model to learn the intricate facial cues within the positive face samples to differentiate between facial regions, which are required for dense tasks [38, 47] such as landmark detection and matching.

On the other hand, the Masked Image Modeling (MIM) protocol [3, 14, 50, 57], which requires the network to reconstruct the masked regions from limited context, *intrinsically* suits our downstream task objective. Based on the observation that the non-landmark regions (e.g., cheeks and foreheads) are larger and more uniform than the sparse and distinctive landmark regions (e.g., the eyes and lip corners), we hypothesize that the reconstruction of the masked landmark regions leads to the formation of effective representations of the facial landmarks. Hence, we choose to adopt the Masked Autoencoder (MAE) [14] as our backbone in the first stage of our framework.

For the second stage, both CL [9] and LEAD [19] utilize objectives to establish correspondences between *each pair* of feature descriptors within the same image. Based on the earlier observation that non-landmark regions are larger and more uniform, we ask the question: *is it necessary to establish correspondences between all feature descriptor pairs?* We hypothesize that the selective refinement of the *important* correspondences utilizes the network’s parameters more effectively. To this end, we employ a novel Correspondence Approximation and Refinement Block. Here, we first differentiate the MAE’s output into attentive (landmark and important facial regions) and inattentive (insignificant facial regions or background) tokens using the first-stage correspondence signals. Next, a clustering algorithm operates on the inattentive tokens and approximates the member tokens using the cluster center. Finally, we supervise a light-weight projector network using a novel Locality-Constrained Repellence Loss that penalizes the erroneous strong correspondences between the different token types weighted by spatial proximity. Here, only the *select correspondences* are directly refined since the loss operates only

on the attentive tokens and inattentive cluster center proxies.

In order to highlight the above stage-wise merits of our approach, we visually compare, at a high level, our framework, which we term Selective Correspondence Enhancement with MAE (SCE-MAE), with prior works in Fig. 1. Our approach not only produces more distinguishable first-stage features but also outputs sharper similarity maps corresponding to the example query, testifying to superior final landmark representations.

In this paper, we show that by leveraging MAE [14] during the first stage and systematically eliminating redundant correspondence learning during the second stage, SCE-MAE can output locally distinct facial landmark representations without the use of labeled data. As a result, it outperforms the previous SOTA methods by large margins. We summarize our contributions below:

1. We are the first to adopt an MIM-trained SSL backbone for the first-stage training of self-supervised facial landmark detection and matching methods. We demonstrate using MAE [14] that the mask-and-predict pretext task more naturally suits the downstream objective and delivers highly potent initial landmark representations.
2. We introduce the Correspondence Approximation and Refinement Block (CARB) during the second-stage to identify and approximate the features of unimportant non-landmark regions, and subsequently operate a novel Locality-Constrained Repellence (LCR) Loss to directly hone only the salient correspondences.
3. We demonstrate the effectiveness and robustness of our framework, SCE-MAE, as it surpasses existing SOTA methods on the landmark matching ( $\sim 20\%$ - $44\%$ ) and detection ( $\sim 9\%$ - $15\%$ ) tasks under various annotation conditions for several challenging datasets.

## 2. Related Works

**Self-Supervised Learning (SSL).** By solving unique pretext tasks, SSL methods are able to learn discriminative feature representations from unlabeled data. Early works explored pretext tasks such as predicting the rotation angle [11] and recovering the original image from random permuted patches [34, 35]. Recently, invariant and contrastive learning based SSL methods [5–8, 12, 13, 37] have gained popularity due to their ability to capture high-level semantic concepts from the data. Invariant learning aims to learn transformation invariant features by forcing the representations of two randomly augmented views of the same image to be similar. Contrastive learning defines different views of an anchor image as positives and views of different images as negatives. Here, the objective is to pull the representations of the anchor and positives together while pushing apart those of the anchor and negatives. To achieve this, MOCO [13] and SimCLR [6] adopted the InfoNCE [36] and NT-Xent [43] losses respectively. These methods operate at

the encoded image or *instance-level* and can be categorized as augment-and-compare SSL methods [25].

Recently, the Masked Image Modeling (MIM) protocol has gained significant momentum [3, 14, 26, 50, 57]. These methods operate at the *region-level* and learn to recover the masked regions from the contextual information contained in the unmasked patches. It has been empirically shown that by using non-extreme masking ratios or patch sizes in Masked Autoencoders (MAE) [14], the representation abstractions capture robust high-level information, while extreme masking ratios capture more low-level information [22]. With higher masking ratios as the norm, MAE executes dense reconstruction, making them intrinsically suitable for dense prediction tasks [14, 26].

For the first stage of self-supervised face landmark detectors, ContrastLandmark (CL) [9] and LEAD [19] utilize MOCO [13] and BYOL [12] pretrained backbones respectively. Since neither MOCO nor BYOL operate explicitly at the sub-image (region/pixel) level, the representations used for the second stage of CL and LEAD are potentially sub-optimal. On the other hand, the sparse nature of facial landmarks perfectly matches the MIM objective to reconstruct the whole view from unmasked patches, resulting in higher fidelity coarse local features. Hence, in our framework, we adopt the MAE [14] as our first stage SSL protocol.

**Unsupervised Landmark Prediction.** To tackle landmark prediction without annotated data, there have been several approaches. Equivalence learning leverages transformation equivalence as a free supervision signal to learn landmark embeddings [44]. Since an undesirable constant vector output would satisfy the objective, adding a diversity loss or enabling similarity enforcement through intermediate auxiliary images are proposed to tackle the issue [45, 46]. Another approach is through generative modeling where landmarks are discovered by training networks with a reconstruction objective [17, 29, 30, 42, 51, 54] such as reconstructing the human image with a different pose [17].

Recent works such as ContrastLandmark (CL) [9] and LEAD [19] have adopted SSL methods to extract coarse features that capture the broad semantic concept and further process them to establish regional/local correspondences. CL and LEAD construct hypercolumns and compact them using proximity-guided and correspondence guided reduction objectives respectively. While both methods reduce the final representation size, hypercolumns are memory-wise enormous structures and operating on them is a computationally intensive process. Furthermore, each spatial feature pair is subject to the optimization objective, neglecting the possibility that some local correspondences do not contribute as much to the downstream task. On the contrary, using our SCE-MAE framework, we do not operate on expensive hypercolumns, and we identify and directly process only salient local correspondences.

### 3. Method

We depict our proposed Selective Correspondence Enhancement with MAE (SCE-MAE) framework in Fig. 2 and detail each component in the following subsections. In Sec. 3.1 we revisit Masked Image Modeling to introduce the MAE [14] as a more suitable and potent first stage protocol. In Sec. 3.2, we elaborate on the setup to execute selective correspondence through the process of reducing the effective number of final correspondence pairs. In Sec. 3.3, we introduce our Correspondence Approximation and Refinement Block, wherein we explain the components of our novel Locality-Constrained Repulsion Loss and how it directly hones only the selected correspondences.

#### 3.1. A Revisit of Masked Image Modeling

Masked Image Modeling (MIM) [3, 14, 50, 57] is an SSL paradigm that involves the reconstruction of the original image from the unmasked patches. Taking MAE [14] as an example, given an input image  $x$ , the encoder first divides the image into non-overlapping patches  $x^p$  with positional embedding added to them. A class token is appended to the patch tokens but will not be affected by the following masking procedure. A binary mask  $M$  is randomly sampled to determine the masked out regions. The unmasked patches are denoted by  $\hat{x}^p = x^p \circ M$ , where  $\circ$  symbolizes the Hadamard product, and are processed by the encoder to output the patch embeddings  $\hat{f}^p$ . Finally, MAE uses a special embedding [MASK] to fill in the masked positions,  $f^p = \hat{f}^p + [\text{MASK}] \circ (1 - M)$ , and reconstruct  $x$  from  $f^p$  by minimizing the pixel-level mean squared error via a light-weight decoder. The reconstruction task requires the network to capitalize on the limited semantic context provided by the unmasked patches and the supplied positional information. This encourages the network to forge discriminative features that are optimal for differentiating and localizing the important landmark regions.

#### 3.2. Setup for Selective Correspondence

**Attentive-Inattentive Separation.** The second stage of the framework aims to establish local correspondences effectively to ensure that the representations reflect the extent of similarity and dissimilarity between the different facial regions. To achieve this, we propose to execute *selective correspondence*, i.e., the elimination of the direct refinement of unimportant non-landmark correspondences, and focus on optimizing those that are critical for landmark disambiguation. The first step in this endeavor is to identify potential landmark and non-landmark regions. Due to the observable opposing nature of facial landmarks (sparse and distinct) and non-landmark regions (dense and uniform), we hypothesize that they are coarsely distinguishable using the first stage backbone features.

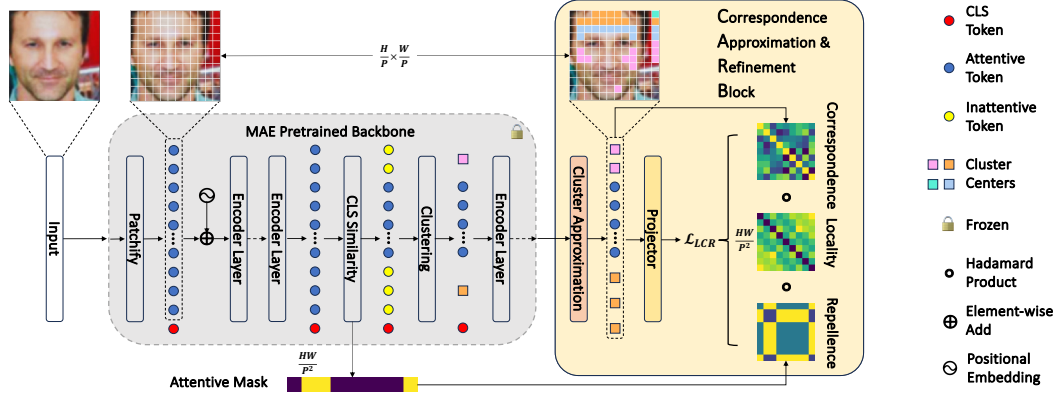


Figure 2. An overview of the second stage of our proposed SCE-MAE. We first split the MAE patch tokens into attentive (blue) and inattentive (yellow) tokens based on CLS token similarity. The inattentive tokens are clustered into  $K$  cluster centers. In the Correspondence Approximation and Refinement Block (CARB), we first substitute the inattentive tokens using the cluster centers (square symbols) and then refine the local features using our novel Locality-Constrained Repellence (LCR) Loss. The LCR loss weakens existing erroneous correspondences in a weighted manner by considering the token-pair proximity (locality) and correspondence type (repellence) constraints.

Following MAE [14], we adopt the ViT [10] as our backbone architecture. The class (CLS) token represents the image and is obtained by aggregating information from the other patch tokens over several layers. Since landmarks are sparse and have more distinct texture, we expect the corresponding tokens to have a large influence on the CLS token representation. After pretraining, we compute a similarity vector between the CLS token and all patch tokens as:

$$Sim_{cls} = \text{Softmax}\left(\frac{K \cdot q_{cls}}{\sqrt{d}}\right) \in \mathbb{R}^N, \quad (1)$$

where  $q_{cls}$ ,  $K$ ,  $d$ , and  $N$  denote the CLS token query vector, the patch token key matrix, latent dimension, and number of patch tokens respectively. Here,  $q_{cls} \in \mathbb{R}^d$  and  $K \in \mathbb{R}^{N \times d}$ . We then split the  $N$  patch tokens into two groups: (1) attentive group, consisting of the  $\eta \cdot N$  tokens that have the highest similarity score with the CLS token, and (2) inattentive group, consisting of the remaining  $(1 - \eta) \cdot N$  tokens. Here,  $\eta$  is a hyperparameter between 0 and 1. We observe that the inattentive tokens mostly cover non-landmark face regions (See Fig. 2), such as cheeks and forehead, as well as background. Henceforth, we presume that attentive tokens cover the landmark and important facial regions, while inattentive tokens correspond to unimportant non-landmark regions.

**Inattentive Token Clustering.** Since several inattentive tokens often correspond to the same facial region (e.g., cheek, forehead, etc.), the downstream correspondence objectives associated with them would likely be redundant. By applying a clustering algorithm on the inattentive tokens, we can represent numerous non-landmark regions with only a handful of cluster centers. Selective correspondence can then be set up by discarding all non-cluster center tokens, ensuring that no correspondence is established with them.

Specifically, we adopt a simple density peak clustering algorithm [28], wherein two variables  $\rho$  and  $\delta$  are defined for each inattentive token. Here,  $\rho_i$  measures the density of the  $i$ -th token and  $\delta_i$  computes the minimum distance from the  $i$ -th token to any other inattentive token which has a higher density. Mathematically, they are defined as:

$$\rho_i = \exp\left(\sum_{t_j \in T_{inatt}} \|t_i - t_j\|_2^2\right), \quad (2)$$

$$\delta_i = \begin{cases} \min_{j: \rho_j > \rho_i} \|t_i - t_j\|_2, & \text{if } \exists j \text{ s.t. } \rho_j > \rho_i \\ \max_j \|t_i - t_j\|_2, & \text{otherwise} \end{cases}, \quad (3)$$

where  $t_i, t_j \in T_{inatt}$ , and  $T_{inatt}$  denotes all inattentive tokens. Since the cluster center should have higher density than neighbouring tokens and should also be distant to other cluster centers, the cluster center score of the  $i$ -th token is computed by  $\rho_i \cdot \delta_i$ . We select the top- $K_c$  scoring tokens as cluster centers, where  $K_c$  is a hyperparameter. The remaining inattentive tokens are discarded and the cluster center tokens subsequently act as representative proxies for them.

### 3.3. Selective Correspondence using CARB

In our Correspondence Approximation and Refinement Block (CARB), we first substitute the discarded inattentive tokens with their corresponding cluster centers and aggregate the relevant visual features to obtain a complete 2D feature map as illustrated in Fig. 2. With the backbone frozen, the feature map is passed through a light-weight projector, which is supervised by our novel Locality-Constrained Repellence (LCR) Loss. As the LCR loss operates on the features of attentive tokens and inattentive cluster centers, we directly refine only the most important correspondences, thereby achieving *selective correspondence*.

**Locality-Constrained Repellence (LCR) Loss.** We design and operate the LCR loss to yield high-fidelity fine-grained features by optimally refining local correspondences. Henceforth, we use  $T_{att}$  and  $T_{inatt'}$  to denote the attentive and the approximated inattentive tokens (cluster centers) respectively, and define  $T = T_{att} \cup T_{inatt'}$ , as the set of all considered tokens.

We begin by formally defining correspondence, i.e., the probability that a patch token  $t_j$  corresponds to a patch token  $t_i$  in the image  $x$ , which is expressed as:

$$p(t_j|t_i; \Phi, x) = \frac{\exp(\langle \Phi_{t_i}(x), \Phi_{t_j}(x)/\tau \rangle)}{\sum_{t_k \in T} \exp(\langle \Phi_{t_i}(x), \Phi_{t_k}(x)/\tau \rangle)}, \quad (4)$$

where  $\Phi_{t_i}(x)$  is the final projected feature representation of patch  $t_i$ , and  $\tau$  is the temperature parameter.

We observe that image patches that are spatially distant from each other often correspond to different facial regions. Hence, it should follow that strong correspondences between distant patches are likely to be erroneous and should be discouraged. We compute a locality constraint to formalize this idea using the following function:

$$f_{loc}(t_i, t_j) = \log(\|t_i - t_j\| + 1), \quad (5)$$

where  $t_i, t_j \in T$ , and  $\|\cdot\|$  computes the spatial distance. The log function saturates the coefficient in order to discourage the network from excessively focusing on separating very distant correspondences. Although a similar constraint was introduced in [44], the primary motive was to avoid collapse during equivalence learning.

Considering the attentive ( $T_{att}$ ) and the approximated inattentive ( $T_{inatt'}$ ) token sets, there are three types of correspondences: attentive-attentive ( $att - att$ ), attentive-inattentive ( $att - inatt$ ), and inattentive-inattentive ( $inatt - inatt$ ). We introduce a repellence coefficient to quantify the importance of each correspondence type:

$$\lambda_{rep}(t_i, t_j) = \begin{cases} r_{att-att}, & \text{if } t_i, t_j \in T_{att} \\ r_{inatt-inatt}, & \text{if } t_i, t_j \in T_{inatt'} \\ r_{att-inatt}, & \text{otherwise} \end{cases}, \quad (6)$$

where each coefficient  $r$  is a hyperparameter. In practice, we set  $r_{att-att}$  and  $r_{att-inatt}$  to be higher than  $r_{inatt-inatt}$  since we aim to prioritize facial landmark differentiation and landmark vs non-landmark disambiguation over non-landmark differentiation respectively.

Combining all of the above defined components, we mathematically express the LCR loss as:

$$\mathcal{L}_{LCR} = \sum_{t_i \in T} \sum_{t_j \in T} f_{loc}(t_i, t_j) \cdot \lambda_{rep}(t_i, t_j) \cdot p(t_j|t_i; \Phi, x), \quad (7)$$

The LCR loss aims to forge effective local features by systematically weakening erroneous, spatially distant correspondences among and between the important and unimportant patch tokens.



Figure 3. Comparison between the original (red) and re-annotated (green) landmarks in  $AFLW_R$  test set. We denote the original and corrected test sets as  $AFLW_{RO}$  and  $AFLW_{RC}$  respectively.

**Inference.** After training, we obtain optimized representations for all image regions. Hence, during inference, we bypass the clustering and inattentive token approximation procedures and only utilize the original features for the downstream tasks. Additionally, we require to spatially expand the size of the feature input to the projector for fair comparison against prior works. Naïvely decreasing the patch size to expand the output not only quadratically increases the computation and memory costs but also may lead to the formation of inferior feature representations [22, 50]. Instead, we adopt the cover-and-stride technique [49] to produce more fine-grained and rich expanded representations.

## 4. Experiments

**Datasets.** Following prior works [9, 19, 46], we pretrain our backbone on the CelebA [27] dataset, which contains 162,770 images. Face landmark detection is evaluated on four datasets: MAFL [56], 300W [40] and two variants of AFLW [21]. MAFL consists of 19,000 training images and 1,000 test images. 300W has 3148 training images and 689 test images.  $AFLW_M$  contains 10,122 training images and 2995 testing images, which are crops from MTFLL[55].  $AFLW_R$  contains tighter crops of face images where the training and test set has 10,122 and 2,991 images respectively. Note that 300W provides 68 annotations per image while the other three datasets only provides 5 annotations.

**Re-annotation.** Although the  $AFLW_R$  has been used in prior works, the fidelity of the annotations are questionable. In Fig. 3, we visualize several annotation errors in the  $AFLW_R$  test set using red dots. These include errors arising due to semantic mismatches, translations, and random shifts. For a more consistent and reliable evaluation, we re-annotate the  $AFLW_R$  test set and illustrate a few of these corrections using green dots in Fig. 3. In the following sections, we use  $AFLW_{RO}$  and  $AFLW_{RC}$  to denote the original and corrected dataset respectively.

**Implementation Details.** We pretrain our models on the CelebA dataset using MAE with three backbones: DeiT-T, DeiT-S and DeiT-B. All models were trained for 400 epochs




Figure 4. **Qualitative results on landmark matching.** The reference/ground-truth are shown in the top/bottom row. The middle rows show the matching results of our method and prior works, grouped column-wise by errors occurring with the eyes, nose and lip corner landmarks respectively. Our method outputs consistently more accurate matching resulting from leveraging higher fidelity projected features.

with a batch size of 512, a learning rate of  $3e-4$  and patch size of 8. Following DVE [46], we resize the image to  $136 \times 136$  and crop the center  $96 \times 96$  as input for both landmark matching and regression. We set the attentive rate  $\eta$  to 0.25 for DeiT-B and 0.1 for DeiT-T and DeiT-S. We apply clustering after the third encoder layer and set the number of clusters  $K_c$  to 4. For the LCR loss, the three repulsion hyperparameters are set to  $r_{attn-attn} = 5, r_{attn-inattn} = 5, r_{inattn-inattn} = 2$ . Ablation studies on the various hyperparameters are included in the *Supplementary Material*.

#### 4.1. Landmark Matching

**Evaluation Protocol.** Following [46], 1000 reference-and-test image pairs are generated from MAFL test set for evaluation. The first 500 pairs serve as the benchmark for landmark matching between same identities, which contains the original image and its thin-plate-spline (TPS) deformed counterpart. The other 500 pairs are of different identities. During evaluation, all feature maps are bi-linearly up-sampled to the image resolution. Landmark representations of the reference image are used to query the test image. The location with the highest cosine similarity is considered as the matched prediction. Finally, we compute the Mean Pixel Error between the prediction and ground-truth.

**Quantitative Results.** We compare our method with existing SOTA methods in Table 1 by grouping the results based on the final feature size. We use three different backbones to control the number of parameters for a fair comparison. In the first group, our model with DeiT-T, being a *fraction*

Table 1. **Quantitative evaluations on landmark matching.** We report the mean pixel error between the prediction and ground-truth on 1000 image pairs sampled from MAFL. The best and second best results are shown in **bold** and underline respectively. We group the results by the projected feature dimension. Our method outperforms all prior works by large margins within each group for both the same and different identity settings.

Method	#Param. Millions	Feat. Dim.	Same Mean Pixel Error↓	Diff. Mean Pixel Error↓
DVE[46]	12.4	64	0.92	<u>2.38</u>
CL[9]	23.8	64	0.92	2.62
LEAD[19]	23.8	64	<u>0.51</u>	2.60
Ours DeiT-T	5.4	64	<b>0.47</b>	<b>1.99</b>
CL[9]	23.8	128	<u>0.82</u>	<u>2.19</u>
Ours DeiT-S	21.4	128	<b>0.31</b>	<b>1.69</b>
CL[9]	23.8	256	0.71	2.06
LEAD[19]	23.8	256	0.48	2.50
Ours DeiT-S	21.4	256	<u>0.33</u>	<u>1.72</u>
Ours DeiT-B	85.3	256	<b>0.27</b>	<b>1.61</b>

*of the size of prior works*, already outperforms the SOTA. In the second and third group, our method visibly outperforms prior works by large margins of  $\sim 20\%$  and  $\sim 44\%$  for the same and different identities respectively. We attribute this to the highly potent initial features from the MAE pretraining, which, when strategically refined through selective correspondence using CARB, yields distinctive final features that were vital for successful landmark matching.

Table 2. **Quantitative evaluations on landmark detection with all annotated samples.** We compare our method with existing SOTA and report the error as the percentage of inter-ocular distance on four human face datasets: MAFL, AFLW<sub>M</sub>, AFLW<sub>R</sub> and 300W. For AFLW<sub>R</sub>, we report the results on both the original (AFLW<sub>RO</sub>) and corrected (AFLW<sub>RC</sub>) datasets. Our method, despite using significantly smaller features by avoiding expensive hypercolumns, outperforms prior works on all four datasets, even with our smallest backbone, DeiT-T.

Method	#Params. Millions	Feature Dim.	Hypercol. Used	MAFL	AFLW <sub>M</sub>	AFLW <sub>RO</sub>	AFLW <sub>RC</sub>	300W
Inter-ocular Distance (%)↓								
DVE[46]	12.6	64	✗	2.76	6.96	6.33	5.58	4.58
CL[9]	23.8	3840	✓	2.76	6.17	5.69	5.06	4.84
LEAD[19]	23.8	3840	✓	2.44	6.05	5.71	5.11	4.87
Ours DeiT-T	5.4	256	✗	2.20	5.89	5.54	4.86	4.22
Ours DeiT-S	21.4	512	✗	<u>2.08</u>	<u>5.33</u>	<u>5.40</u>	<u>4.69</u>	<b>3.94</b>
Ours DeiT-B	85.3	1024	✗	<b>2.07</b>	<b>5.23</b>	<b>5.33</b>	<b>4.60</b>	<u>3.95</u>

**Qualitative Results.** We visualize our landmark matching results between different identities and compare our results with existing SOTA methods in Fig. 4. The mismatches on different landmarks when using CL [9] and LEAD [19] are shown in different columnar groups, e.g., the first group of three columns contain eye-related mismatches. Our method clearly achieves a more accurate matching performance across all landmarks even on difficult examples such as those wearing eye-glasses. Admittedly, our method experiences some failures when the poses of reference and test samples are vastly dissimilar or when landmark regions are severely occluded. Some of the failure cases are shown in the *Supplementary Material*.

## 4.2. Landmark Detection

**Evaluation Protocol.** Following prior works [9, 19, 46], we freeze the pretrained backbone and projector, and only train a light-weight regressor. Our regressor consists of a convolution block (instead of a linear layer) and a linear layer when training on all annotated samples. The convolution block utilizes the spatial context to produce  $I$  intermediate heatmaps for each landmark which are converted to  $I$  pairs of 2D coordinates by a soft-argmax operation and fed to a linear layer that outputs the final landmark prediction. We set  $I = 50$  for all experiments [9, 19, 46]. We leverage the first and second stage concatenated features as a more robust input to the regressor as we expect the backbone to provide rich task-agnostic representations during the first stage and the projector to supplement task-specific cues critical for landmark detection during the second stage. For a fair comparison, the reported results are produced with their official implementation and checkpoints, if available.

**All Annotated Samples.** We compare our method with prior works on landmark detection benchmarks in Table 2. Once again, even with our smallest backbone DeiT-T, being *a fraction of the size of prior works*, outperforms existing SOTA. Considering our best results, our method achieves ~9%-15% performance gain across different benchmarks.

Such a compelling performance is an attestation to the excellent discriminative ability of our features, which provide intricate disambiguation cues to the regressor for locating landmarks. We also highlight that *all methods* achieve lower error with our re-annotated AFLW<sub>RC</sub> test set, hence confirming the higher annotation quality.

**Limited Annotated Samples.** We compare our method with prior works on landmark detection under limited annotation in Table 3. Our method outperforms all existing SSL-based methods with significant performance gain under all annotation and feature dimension settings. Specifically, we achieved a relative gain of 8.6% *on average* and *as high as 20.1%* compared to the existing SOTA. Furthermore, we observe a smaller standard deviation with repeated experiments, indicating that our method produces optimal features more consistently, hence attesting to its robustness.

## 4.3. Ablation Studies

**Importance of Each Component.** To better understand our proposed SCE-MAE framework, we report the component-wise ablation analysis on the landmark matching task in Table 4. The first three rows indicate the usage of only the first-stage backbone features, while the last two rows, respectively, indicate the inclusion of clustering and LCR loss in our framework. CL [9] and LEAD [19] utilize hypercolumns whereas we leverage the vanilla last layer features of the pretrained MAE, which we indicate as Baseline. Using the backbone alone, our baseline outperforms CL and LEAD, validating our motivation that MIM is a more suitable pretext task for landmark representation learning. We observe that the clustering assists the matching between the same identity while the LCR loss boosts the matching performance between different identities. Overall, these trends align with our expectations: initially, the region-level first-stage MAE features capture local intricacies but are too raw to generalize the landmarks across different identities; the clustering disambiguates the landmarks from the unimportant regions, which improves the same identity matching

Table 3. **Quantitative evaluations on landmark detection with limited annotated samples.** We compare our method with existing SOTA under different annotation settings on the AFLW<sub>M</sub> dataset and report the error as the percentage of inter-ocular distance. Our method proves to be more effective by offering notably lower errors and more robust by yielding a lower std of errors than prior works.

Method	Feat. Dim.	Number of Annotated Samples					
		1	5	10	20	50	100
DVE[46]	64	<b>14.23 ± 1.45</b>	12.04 ± 2.03	12.25 ± 2.42	11.46 ± 0.83	12.76 ± 0.53	11.88 ± 0.16
CL[9]	64	24.87 ± 2.67	15.15 ± 0.53	13.52 ± 1.08	11.77 ± 0.68	11.57 ± 0.03	10.06 ± 0.45
LEAD[19]	64	21.80 ± 2.54	13.34 ± 0.43	11.50 ± 0.34	10.13 ± 0.45	9.29 ± 0.41	9.11 ± 0.25
SCE-MAE (Ours)	64	18.41 ± 1.21	<b>11.79 ± 0.44</b>	<b>10.57 ± 0.24</b>	<b>9.65 ± 0.14</b>	<b>8.60 ± 0.17</b>	<b>8.31 ± 0.06</b>
CL[9]	128	27.31 ± 1.39	18.66 ± 4.59	13.39 ± 0.30	11.77 ± 0.85	10.25 ± 0.22	9.46 ± 0.05
LEAD[19]	128	21.20 ± 1.67	13.22 ± 1.43	10.83 ± 0.65	9.69 ± 0.41	8.89 ± 0.20	8.83 ± 0.33
SCE-MAE (Ours)	128	<b>20.14 ± 1.76</b>	<b>11.99 ± 0.71</b>	<b>10.40 ± 0.22</b>	<b>9.25 ± 0.14</b>	<b>8.49 ± 0.19</b>	<b>7.96 ± 0.21</b>
CL[9]	256	28.00 ± 1.39	15.85 ± 0.86	12.98 ± 0.16	11.18 ± 0.19	9.56 ± 0.44	9.30 ± 0.20
LEAD[19]	256	21.39 ± 0.74	12.38 ± 1.28	11.01 ± 0.48	10.06 ± 0.59	8.51 ± 0.09	8.56 ± 0.21
SCE-MAE (Ours)	256	<b>17.08 ± 1.35</b>	<b>11.28 ± 0.54</b>	<b>10.30 ± 0.09</b>	<b>8.95 ± 0.08</b>	<b>8.20 ± 0.20</b>	<b>7.58 ± 0.09</b>

Table 4. **Component-wise ablation on landmark matching.** The first three rows compare the results using backbone features only. Our raw MAE features (Baseline) are compared against the hypercolumns used in CL and LEAD. The last two rows indicate the inclusion of the clustering and the LCR loss in our framework.

Method	Cluster	$\mathcal{L}_{LCR}$	Same Mean Pixel Error↓	Diff.
CL [9]	-	-	0.69	5.37
LEAD [19]	-	-	2.35	6.22
Baseline	<b>X</b>	<b>X</b>	0.55	<u>3.51</u>
SCE-MAE (Ours)	<b>✓</b>	<b>X</b>	<u>0.30</u>	4.04
	<b>✓</b>	<b>✓</b>	<b>0.27</b>	<b>1.61</b>

performance; and finally, the LCR loss forges critical local correspondences between important facial regions, resulting in the best performance for both settings.

**Visualization of Landmark Representations.** We visualize the t-SNE plots of the landmark representations corresponding to 1000 test images in Fig. 5. Since LEAD [19] only performs knowledge distillation in its second stage, we use the first-stage hypercolumn representations as it can be considered as the upper bound of the second-stage objective. For each method, we execute t-SNE 100 times and report the mean and standard deviation of the Silhouette Coefficient [39], a metric (higher is better) indicating the quality of the clustering as a function of the mean intra-cluster (lower is better) and inter-cluster (higher is better) distance. For CL [9], the samples within the cluster are more scattered, resulting in a higher intra-cluster distance. For LEAD, though the clusters are more dense, the clusters of the left/right lip corner and nose are not clearly separated, resulting in a lower inter-cluster distance. Our method results in both well-separated and dense clusters, which is reflected by a high Silhouette Coefficient, thereby corroborating the superior quality of our landmark representations.




Figure 5. **t-SNE plot of the landmark representations.** † denotes the usage of the stage 1 hypercolumn representations. SC denotes the Silhouette Coefficient [39], a score (higher is better) which measures the quality of the clustering. Our method results in both a clear separation between the landmarks and the densest landmark clusters, resulting in the highest Silhouette Coefficient.

## 5. Conclusion

In this work, we present SCE-MAE, a two-stage framework to tackle the self-supervised face landmark estimation tasks. To learn effective and locally distinct representations, we target structured improvements on both stages. For the first stage, we leverage the region-level MAE instead of instance-level SSL methods to derive more potent initial representations. For the second stage, we demonstrate that it is beneficial to identify important facial regions and directly hone only the salient correspondences. Our approach yields discriminative and high-quality landmark representations that result in superior performance over prior SOTA works on both the landmark matching and detection tasks. Due to the nature of facial data, we believe that further research on the sparsification of the correspondence computation through the systematic elimination of insignificant correspondences could allow future self-supervised landmark estimation methods to better exploit inter-landmark dependencies and form higher-caliber landmark representations.



## References

- [1] 3d face reconstruction and dense alignment with a new generated dataset. *Displays*, page 102094, 2021. 1
- [2] Randall Balestriero and Yann LeCun. Contrastive and non-contrastive self-supervised learning recover global and local spectral embedding methods. In *NeurIPS*, 2022. 1
- [3] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEiT: BERT pre-training of image transformers. In *ICLR*, 2021. 2, 3
- [4] F. Bougourzi, F. Dornaika, and A. Taleb-Ahmed. Deep learning based face beauty prediction via dynamic robust losses and ensemble regression. *Knowledge-Based Systems*, 2022. 1
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 1, 2
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [8] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021. 2
- [9] Zezhou Cheng, Jong-Chyi Su, and Subhransu Maji. On equivariant and invariant learning of object landmark representations. In *ICCV*, 2021. 1, 2, 3, 5, 6, 7, 8
- [10] 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 *ICLR*, 2021. 4
- [11] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 2
- [12] 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. In *NeurIPS*, 2020. 1, 2, 3
- [13] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 3
- [14] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 1, 2, 3, 4
- [15] Yangyu Huang, Hao Yang, Chong Li, Jongyoo Kim, and Fangyun Wei. Adnet: Leveraging error-bias towards normal direction in face alignment. In *ICCV*, 2021. 1
- [16] Tharun J. Iyer, Rahul K., Ruban Nersisson, Zheming Zhuang, Alex Noel Joseph Raj, and Imthiaz Refayee. Machine learning-based facial beauty prediction and analysis of frontal facial images using facial landmarks and traditional image descriptors. *Computational Intelligence and Neuroscience*, 2021. 1
- [17] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks through conditional image generation. In *NeurIPS*, 2018. 3
- [18] Aniwat Juhong and C. Pintavirooj. Face recognition based on facial landmark detection. In *2017 10th Biomedical Engineering International Conference (BMEiCON)*, 2017. 1
- [19] Tejan Karmali, Abhinav Atrishi, Sai Sree Harsha, Susmit Agrawal, Varun Jampani, and R Venkatesh Babu. LEAD: Self-supervised landmark estimation by aligning distributions of feature similarity. In *WACV*, 2022. 1, 2, 3, 5, 6, 7, 8
- [20] Robin Kips, Ruowei Jiang, Sileye Ba, Edmund Phung, Parham Aarabi, Pietro Gori, Matthieu Perrot, and Isabelle Bloch. Deep graphics encoder for real-time video makeup synthesis from example. In *CVPRW*, 2021. 1
- [21] Martin Koestinger, Paul Wohlhart, Peter M Roth, and Horst Bischof. Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. In *ICCVW*, 2011. 5
- [22] Lingjing Kong, Martin Q Ma, Guangyi Chen, Eric P Xing, Yuejie Chi, Louis-Philippe Morency, and Kun Zhang. Understanding masked autoencoders via hierarchical latent variable models. In *CVPR*, 2023. 3, 5
- [23] Abhinav Kumar, Tim K. Marks, Wenxuan Mou, Ye Wang, Michael Jones, Anoop Cherian, Toshiaki Koike-Akino, Xiaoming Liu, and Chen Feng. Luvli face alignment: Estimating landmarks' location, uncertainty, and visibility likelihood. In *CVPR*, 2020. 1
- [24] TianXing Li, Zhi Yu, Edmund Phung, Brendan Duke, Irina Kezele, and Parham Aarabi. Lightweight real-time makeup try-on in mobile browsers with tiny cnn models for facial tracking. In *CVPRW*, 2019. 1
- [25] Wei Li, Jiahao Xie, and Chen Change Loy. Correlational image modeling for self-supervised visual pre-training. In *CVPR*, 2023. 3
- [26] Jihao Liu, Xin Huang, Jinliang Zheng, Yu Liu, and Hongsheng Li. MixMAE: Mixed and masked autoencoder for efficient pretraining of hierarchical vision transformers. In *CVPR*, 2023. 3
- [27] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 5
- [28] Sifan Long, Zhen Zhao, Jimin Pi, Shengsheng Wang, and Jingdong Wang. Beyond attentive tokens: Incorporating token importance and diversity for efficient vision transformers. In *CVPR*, 2023. 4
- [29] Dominik Lorenz, Leonard Bereska, Timo Milbich, and Björn Ommer. Unsupervised part-based disentangling of object shape and appearance. In *CVPR*, 2019. 3
- [30] Dimitrios Mallis, Enrique Sanchez, Matthew Bell, and Georgios Tzimiropoulos. Unsupervised learning of object landmarks via self-training correspondence. In *NeurIPS*, 2020. 3
- [31] Davide Marelli, Simone Bianco, and Gianluigi Ciocca. Designing an AI-based virtual try-on web application. *Sensors (Basel)*, 2022. 1

- [32] M. I. N. P. Munasinghe. Facial expression recognition using facial landmarks and random forest classifier. In *2018 IEEE/ACIS 17th International Conference on Computer and Information Science (ICIS)*, 2018. [1](#)
- [33] Quang Tran Ngoc, Seunghyun Lee, and Byung Cheol Song. Facial landmark-based emotion recognition via directed graph neural network. *Electronics*, 2020. [1](#)
- [34] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. [2](#)
- [35] Mehdi Noroozi, Ananth Vinjimoor, Paolo Favaro, and Hamed Pirsiavash. Boosting self-supervised learning via knowledge transfer. In *CVPR*, 2018. [2](#)
- [36] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [2](#)
- [37] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. DINOv2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [2](#)
- [38] Byungseok Roh, Wuhyun Shin, Ildoo Kim, and Sungwoong Kim. Spatially consistent representation learning. In *CVPR*, 2021. [2](#)
- [39] Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 1987. [8](#)
- [40] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *ICCVW*, 2013. [5](#)
- [41] Adil Sarsenov and Konstantin Latuta. Face recognition based on facial landmarks. In *2017 IEEE 11th International Conference on Application of Information and Communication Technologies (AICT)*, 2017. [1](#)
- [42] Zhixin Shu, Mihir Sahasrabudhe, Riza Alp Guler, Dimitris Samaras, Nikos Paragios, and Iasonas Kokkinos. Deforming autoencoders: Unsupervised disentangling of shape and appearance. In *ECCV*, 2018. [3](#)
- [43] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NeurIPS*, 2016. [2](#)
- [44] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object frames by dense equivariant image labelling. In *NeurIPS*, 2017. [3, 5](#)
- [45] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks by factorized spatial embeddings. In *ICCV*, 2017. [3](#)
- [46] James Thewlis, Samuel Albanie, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of landmarks by descriptor vector exchange. In *ICCV*, 2019. [1, 3, 5, 6, 7, 8](#)
- [47] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *CVPR*, 2021. [1, 2](#)
- [48] Erroll Wood, Tadas Baltrušaitis, Charlie Hewitt, Matthew Johnson, Jingjing Shen, Nikola Milosavljević, Daniel Wilde, Stephan Garbin, Toby Sharp, Ivan Stojiljković, Tom Cashman, and Julien Valentin. 3d face reconstruction with dense landmarks. In *ECCV*, 2022. [1](#)
- [49] Ronald Xie, Kuan Pang, Gary D Bader, and Bo Wang. MAESTER: Masked autoencoder guided segmentation at pixel resolution for accurate, self-supervised subcellular structure recognition. In *CVPR*, 2023. [5](#)
- [50] 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 *CVPR*, 2022. [2, 3, 5](#)
- [51] Yinghao Xu, Ceyuan Yang, Ziwei Liu, Bo Dai, and Bolei Zhou. Unsupervised landmark learning from unpaired data. *arXiv preprint arXiv:2007.01053*, 2020. [3](#)
- [52] Chun-Hsiao Yeh, Cheng-Yao Hong, Yen-Chi Hsu, Tyng-Luh Liu, Yubei Chen, and Yann LeCun. Decoupled contrastive learning. In *ECCV*, 2022. [2](#)
- [53] Chaoning Zhang, Kang Zhang, Trung X. Pham, Axi Niu, Zhinan Qiao, Chang D. Yoo, and In So Kweon. Dual temperature helps contrastive learning without many negative samples: Towards understanding and simplifying moco. In *CVPR*, 2022. [2](#)
- [54] Yuting Zhang, Yijie Guo, Yixin Jin, Yijun Luo, Zhiyuan He, and Honglak Lee. Unsupervised discovery of object landmarks as structural representations. In *CVPR*, 2018. [3](#)
- [55] Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Facial landmark detection by deep multi-task learning. In *ECCV*, 2014. [5](#)
- [56] Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Learning deep representation for face alignment with auxiliary attributes. *PAMI*, 2015. [5](#)
- [57] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Image BERT pre-training with online tokenizer. In *ICLR*, 2021. [2, 3](#)
- [58] Zhenglin Zhou, Huaxia Li, Hong Liu, Nanyang Wang, Gang Yu, and Rongrong Ji. STAR Loss: Reducing semantic ambiguity in facial landmark detection. In *CVPR*, 2023. [1](#)