

# STAR Loss: Reducing Semantic Ambiguity in Facial Landmark Detection

Zhenglin Zhou<sup>1\*</sup>, Huaxia Li<sup>2\*</sup>, Hong Liu<sup>3\*</sup>, Nanyang Wang<sup>2</sup>, Gang Yu<sup>2</sup>, Rongrong Ji<sup>1†</sup>

<sup>1</sup> Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University.

<sup>2</sup> Tencent PCG. <sup>3</sup> National Institute of Information.

zhenglinzhou@stu.xmu.edu.cn, {haroldzcli, cagewang, skicyyu}@tencent.com,  
 hliu@nii.ac.jp, rrji@xmu.edu.cn

## Abstract

Recently, deep learning-based facial landmark detection has achieved significant improvement. However, the semantic ambiguity problem degrades detection performance. Specifically, the semantic ambiguity causes inconsistent annotation and negatively affects the model's convergence, leading to worse accuracy and instability prediction. To solve this problem, we propose a **Self-adapTive Ambiguity Reduction (STAR)** loss by exploiting the properties of semantic ambiguity. We find that semantic ambiguity results in the anisotropic predicted distribution, which inspires us to use predicted distribution to represent semantic ambiguity. Based on this, we design the STAR loss that measures the anisotropy of the predicted distribution. Compared with the standard regression loss, STAR loss is encouraged to be small when the predicted distribution is anisotropic and thus adaptively mitigates the impact of semantic ambiguity. Moreover, we propose two kinds of eigenvalue restriction methods that could avoid both distribution's abnormal change and the model's premature convergence. Finally, the comprehensive experiments demonstrate that STAR loss outperforms the state-of-the-art methods on three benchmarks, i.e., COFW, 300W, and WFLW, with negligible computation overhead. Code is at <https://github.com/ZhenglinZhou/STAR>

## 1. Introduction

Facial landmark detection, which aims to locate a group of pre-defined facial landmarks from images [48, 51, 56], is a fundamental problem for many downstream tasks, including face verification [12], face synthetic [1], and 3D face reconstruction [10, 13, 17, 46].

Thanks to the development of Convolutional Neural Networks (CNNs) [20, 36, 39], facial landmark detection has

<sup>†</sup>Corresponding author.

\*Equal contribution. This work was done when Zhenglin Zhou was an intern at Tencent PCG.

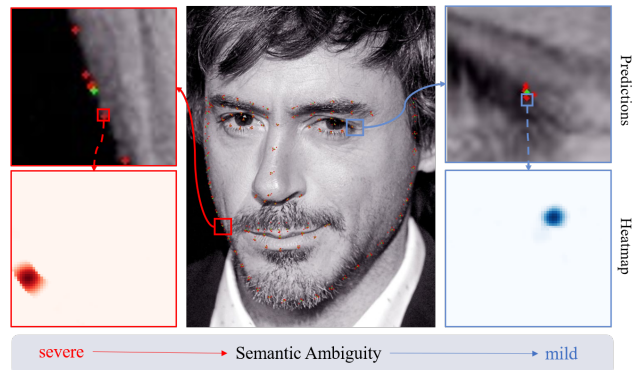


Figure 1. The impact of semantic ambiguity. We visualize the outputs of five models trained with the same architecture under the same experimental setting. (1) The first row shows the predicted facial landmarks for Mr. Tony Stark, marked as red points. And the green point refers to the corresponding mean value. (2) The second row shows the results of predicted probability distribution (i.e., heatmap) from one of the trained models.

improved significantly. At first, coordinate regression methods [3, 16, 33, 47] are proposed to learn the transformation between CNN features and landmark locations via fully connected layers. Recently, the research focus has been the heatmap regression methods, which have shown superiority over coordinate regression methods. The heatmap regression methods [21, 24, 48] predict an intermediate heatmap for each landmark and decode the coordinates from the heatmap. But, the commonly used decoder, Argmax [55], is not differentiable and suffers from quantization error. Recently, some solutions have been proposed [38, 53], and the focus is on the differentiable expectation decoder: soft-Argmax [32]. With the help of soft-Argmax, the heatmap regression method has the advantage of end-to-end training. So the training loss is mainly composed of a regression loss (such as  $\mathcal{L}_2$ ), which makes the model prediction fit the manual annotation.

However, the manual annotation suffers from the semantic ambiguity problem [16, 21, 30]. Specifically, some fa-

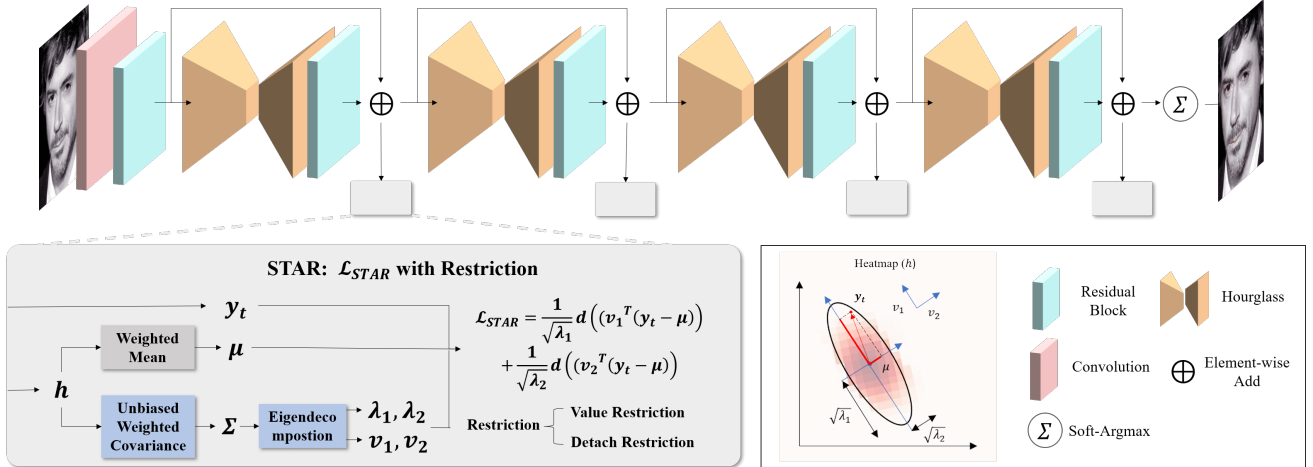


Figure 2. The overview of our framework. We use a four stacked Hourglasses (HG) Network. To mitigate the impact of semantic ambiguity, the STAR loss is applied to each HG module. (Best view in color.)

cial landmarks, especially landmarks located on face contour, do not have a clear and accurate definition. For example, the contour landmarks are defined to evenly distribute around the face contour without a clear definition of the positions [30]. It makes human annotators confused about the position, and it is inevitable to induce inconsistent and imprecise annotations. Thus, we argue that the regression loss will be misled by ambiguous annotations and degrade the model’s convergence and performance. As shown in Figure 1, training the neural network with ambiguous annotations makes the predictions for facial contour landmarks unstable and inaccurate, which will hurt the downstream task [15]. The key problem is to design a new regression loss that mitigates the impact of semantic ambiguity.

In this paper, we propose a novel self-adaptive ambiguity reduction method, STAR loss, by fully exploiting semantic ambiguity. To this end, we explore the impact of semantic ambiguity on the heatmap. Typically, the distribution normalization loss forces the predicted probability distribution to resemble an isotropic Gaussian distribution [32]. However, as shown in Figure 1, compared with the isotropic distribution of the eye corner point, the predicted distribution of the facial contour point is anisotropic. The main difference between the two landmarks is semantic ambiguity, which is more severe in the contour point. We infer that the semantic ambiguity is related to the anisotropic distribution. When the predicted distribution of one facial landmark is anisotropic, this facial landmark has severe semantic ambiguity (leading to model unconvergence), so it is necessary to reduce its impact.

To this end, we begin our story by introducing a customized principal component analysis (PCA) that can process the discrete probability distribution, which contains three steps: weighted mean estimation, unbiased weighted covariance estimation, and eigen-decomposition. We de-

compose a group of predicted distributions and visualize the corresponding principal components. The visualization results show that the first principal component is along with the face contour. Meanwhile, the ambiguous direction for contour landmarks is aligned with the face contour. Therefore, we infer that their first principal component direction is highly consistent with their ambiguity direction.

According to this new observation, we design our STAR loss, which decomposes the prediction error into two principal component directions and divides it by the corresponding energy value. For a facial landmark with anisotropic predicted distribution, the energy of the first principal component is higher than the second. In this way, the error in the first principal component direction can be adaptively suppressed, thereby alleviating the impact of ambiguity annotation on training. However, we find this initial version of STAR loss suffers an abnormal energy increase, leading to premature convergence. To solve this problem, we find that the anomaly results from that the model tends to increase the energy to minimize STAR loss. As a result, we propose two kinds of eigenvalue restriction methods to avoid STAR loss decreasing abnormally.

We evaluate the STAR loss on three widely-used benchmarks, *i.e.*, COFW [5], 300W [35] and WFLW [48]. Experiments show that STAR loss indeed helps deep models achieve competitive performance compared to state-of-the-art methods. Code will be released for reproduction.

## 2. Related Work

At early stages, facial landmark detection is based on statistic models such as Active Appearance Models [8, 31], Constrained Local Models [9], and 3D Morphable Models [2]. Recently, with the development of CNN, deep learning methods have achieved state-of-the-art performance in

facial landmark detection [21, 29, 51, 53], which mainly includes two main branches: coordinate regression method and heatmap regression method.

**Coordinate Regression Method.** Coordinate regression methods [18, 56, 58] utilize fully connected (FC) layers to learn the transformation between CNN features and landmark coordinates. However, coordinate regression methods commonly require a huge number of samples for training, which constrains performance improvements. To this end, many works address this issue from different perspectives – *e.g.*, style transfer [14, 34], semi-supervised learning [3, 15], and virtual data [47]. To further improve the performance, both MDM [41] and RAR [52] use cascaded recurrent networks to sequentially refine the landmark estimation. SDFL [29] and SDL [27] introduce graph constraint through graph convolutional network (GCN). SLPT [50] and RePFormer [26] learn an adaptive inherent relation based on the attention mechanism.

**Heatmap Regression Method.** Heatmap regression methods output an intermediate heatmap for each landmark and take the landmark with the highest intensity as the optimal output. The stacked hourglass network [55], U-Net [11], and HRNet [44] are commonly used to generate the high-quality heatmap. However, heatmap regression methods usually suffer from quantization error since the heatmap is commonly much smaller than the input image. Many works are proposed to mitigate the error, including better heatmap generation [4, 16], mapping construction from heatmap to coordinates [24, 32], coordinates further refinement [4, 38, 53]. Furthermore, many works [21, 45, 48] use facial boundary as the structure constraint to further improve the performance. HSLE [60] automatically constructs a hierarchical structure for learning robust facial landmark detection. LUVLi [24] estimates both the uncertainty and visibility of predicted landmarks. And FaRL [57] proposes the general facial representation learning.

**Semantic Ambiguity in Facial Landmark Detection.** There exist several works [15, 16, 21, 24, 30, 48] addressing semantic ambiguity problems on facial landmark detection. SBR [16] uses the coherency of optical flow between adjacent frames as supervision but is sensitive to illumination and occlusion. LAB [48] uses facial boundary lines as the structure constraint, which is practical but computationally expensive. Moreover, SA [30] proposes a latent variable optimization strategy to find semantically consistent annotations and alleviate random noise during the training stage. However, the complex training strategy limits its application. Closely related to our work is ADNet [21], which presents two key modules, *i.e.*, Anisotropic Direction Loss (ADL) and Anisotropic Attention Module (AAM), to handle the ambiguous annotations problem. Among them, ADL imposes more constraints in normal direction for landmarks on facial boundaries, but the direction and constraint

weight is hand-crafted so that such coarse-grained design degrades its performance.

### 3. Preliminary

We briefly introduce the pipeline of current widely-used regression methods below.

As shown at the top of Figure 2, the basic model of the regression method is stacked of four Hourglasses Networks (HGs) [56]. Each HG generates  $N$  heatmaps, where  $N$  is the number of pre-defined facial landmarks. The normalized heatmap can be viewed as the probability distribution over the predicted facial landmarks. The predicted coordinates are decoded from the heatmap by a soft-Argmax [32].

Formally, given a discrete probability distribution  $h$ , we define the value  $h_i$  as the probability of the predicted landmark locating at  $y_i \in \mathbb{R}^2$ . The expectation coordinate  $\mu$  is decoded by soft-Argmax:

$$\mu = \text{soft-Argmax}(h) = \sum_i h_i y_i. \quad (1)$$

Based on these predicted coordinates, we can use the regression loss to learn the parameter of the model. The regression loss can be viewed as a distance  $d(\cdot)$  between predicted coordinates  $\mu$  and manual annotations  $y_t$ . And  $l_1$ -distance,  $l_2$ -distance, smooth- $l_1$  distance, and Wing distance [18] are widely used. So, the regression loss can be formulated as follows:

$$\mathcal{L}_{reg} = d(y_t - \mu) = d(y_t - \sum_i h_i y_i). \quad (2)$$

The manual annotations suffer from the semantic ambiguity problem. As shown in the gray box of Figure 2, the key to this problem lies in the regression loss.

### 4. Method

The semantic ambiguity problem degrades detection performance considerably. To handle this problem, we design a novel STAR loss, which is influenced by predicted distribution. Specifically, STAR loss is encouraged to be small when the predicted distribution is anisotropic. To represent the shape of the predicted distribution, we first introduce a customized PCA applied to the discrete probability distribution. Then, we visualize the results of PCA computed on the predicted distribution and discuss the relevance between its first principal component and the semantic ambiguity. Based on this, we present the STAR loss in detail, which adaptively suppresses the prediction error in the first principal component direction to mitigate the impact of ambiguity annotation during the training phase.

#### 4.1. Analysis of the Semantic Ambiguity

**The PCA of Discrete Probability Distribution.** We first design a customized PCA applied to the discrete

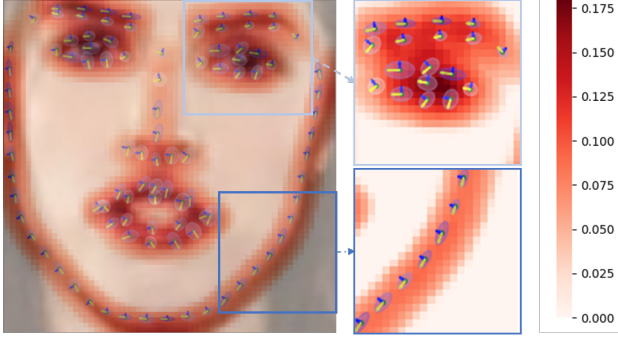


Figure 3. The visualization of the PCA results. The yellow and blue arrows indicate the predicted discrete distribution’s first and second principal components, respectively. The length of the arrow is the corresponding eigenvalues. We use the  $(\lambda_1/\lambda_2)$  to formulate the ambiguity, which is represented by the shading of the blue ellipse. (Best view in color.)

probability distribution, which is composed of two steps: (1) Approximate the means and covariance matrix; (2) Compute the eigen-decomposition of the covariance matrix.

**STEP 1:** similar to the weighted mean in Eq.(1), the weighted covariance matrix is given by:

$$\Sigma_b = \frac{\sum_i h_i (y_i - \mu)^T (y_i - \mu)}{V_1}, \quad (3)$$

where  $V_1 = \sum_i h_i$ . To approximate more precisely, we determine a correction factor to yield an unbiased estimator. The correction factor is  $(1 - V_2/V_1^2)$ , where  $V_2 = \sum_i h_i^2$ . So, the unbiased weighted estimate of the covariance matrix, with Bessel’s correction, is given by Eq.(4). And we commonly use  $\Sigma = \Sigma_{ub}$ .

$$\Sigma_{ub} = \frac{\sum_i h_i (y_i - \mu)^T (y_i - \mu)}{V_1 - (V_2/V_1)}. \quad (4)$$

**STEP 2:** we compute the eigen-decomposition of covariance matrix, which can be formulated as below:

$$\Sigma = V L V^{-1}, \quad (5)$$

where  $V = [v_1, v_2] \in \mathbb{R}^{2 \times 2}$  is the eigenvectors matrix, and  $L = \text{diag}(\lambda_1, \lambda_2) \in \mathbb{R}^{2 \times 2}$  is the corresponding eigenvalues matrix. We refer  $(v_1, \lambda_1)$  and  $(v_2, \lambda_2)$  to represent the first and second principal component, respectively.

**Discussion of Relevance.** In Figure 3, we visualize the PCA results of predicted discrete probability distribution (More visualization can be found in supplementary materials). On the one hand, a significant feature of the contour landmarks is that their first principal component is along with the face contour. Meanwhile, as discussed above, the ambiguity direction for the contour landmarks is also along with the face contour. That is, for landmarks with an

anisotropic predicted distribution, *the first principal component direction is highly consistent with their ambiguity direction*. On the other hand, we visualize the elliptical eccentricity  $(\lambda_1/\lambda_2)$  by the shading of blue. Compared with the near-white eye corner, the color of the contour point is blue. It indicates that eccentricity is higher when the ambiguity is more severe. We further infer that *the corresponding energy can represent the ambiguity intensity*. Because the energy indicates the variance along with the corresponding direction of the principal component. Meanwhile, the high variance in an anisotropic predicted distribution mainly results from inconsistent annotation caused by semantic ambiguity. Based on these two observations, we introduce STAR loss.

## 4.2. STAR Loss

The proposed STAR loss belongs to a self-adaptive ambiguity reduction regression loss. And we realize it by an ambiguity-guided decomposition. Meanwhile, we propose two kinds of eigenvalue restriction methods to avoid STAR loss decrease abnormally.

**Ambiguity-guided Decomposition.** Following the motivations above, STAR loss decomposes the prediction error into the two principal component directions and divides by the corresponding energy value to reduce ambiguity effect self-adaptively. Thus, the STAR loss can be formulated as:

$$\mathcal{L}_{STAR}(y_t, \mu, d) = \frac{1}{\sqrt{\lambda_1}} d(v_1^T (y_t - \mu)) + \frac{1}{\sqrt{\lambda_2}} d(v_2^T (y_t - \mu)), \quad (6)$$

where  $v_1^T$  and  $v_2^T$  project the prediction error  $(y_t - \mu)$  according to ambiguity directions,  $\lambda_1$  and  $\lambda_2$  adaptively imposes constraint, and  $d$  denotes the distance function. Because the decomposition operation does not affect the error metric, STAR loss can be used with an arbitrary distance function and benefits from their improvement, such as smooth-l1, Wing [18], Awing [45], etc.

Take an anisotropic distribution as an example where the  $\lambda_1$  is much larger than  $\lambda_2$ . The  $\mathcal{L}_{STAR}$  results in a small prediction error in the direction of the first principal component  $v_1$ . It makes the model pay less attention to the prediction error in the direction of  $v_1$ , which is mainly caused by inconsistent manual annotation. As a result, with the help of  $\mathcal{L}_{STAR}$ , the detection model can mitigate the impact of semantic ambiguity, leading to better convergence.

**Eigenvalue Restriction.** According to our experiments, we find an abnormal increase of the eigenvalue  $\lambda_i$  that leads to the model’s premature convergence. To solve this problem, we consider the main reason is that the eigenvalue  $\lambda$  is a denominator in  $\mathcal{L}_{STAR}$ . To minimize  $\mathcal{L}_{STAR}$ , the model tends to increase the eigenvalue  $\lambda_i$ . Therefore, we propose two kinds of restriction on the eigenvalue to avoid  $\mathcal{L}_{STAR}$






Figure 4. The NME results of different distance functions with varying input image resolution in WFLW. The blue and red lines indicate the distance function without and with STAR loss, respectively.

decrease abnormally. The first one is a loss restriction term, namely value restriction, which is formulated as below:

$$\mathcal{L}_{value} = \frac{\lambda_1 + \lambda_2}{2}. \quad (7)$$

It directly restricts the abnormal increase of the eigenvalue. On the other hand, we propose to detach the gradient of eigenvalue and eigenvectors, namely detach restriction. It cuts off the backpropagation from eigenvalue and eigenvectors so that they act as constant values in STAR loss. Thus, STAR loss could use them but not directly affect the shape of the predicted distribution. The ablation study shows that both restrictions help mitigate the abnormal increase problem. And we use the value restriction as the default setting unless mentioned.

## 5. Experiments

First, we present the experiment settings, including training details, model setup, dataset, and evaluation metrics. Second, we investigate the benefits of STAR loss with extensive experiments. Third, we discuss the relationship between STAR Loss and related works.

### 5.1. Experiment Settings

**Data Augmentation.** We use the same data augmentation strategy in [21] for all experiments. The input image is generated in two steps: (1) Crop the face regions and resize them into  $256 \times 256$ . (2) Perform augmentations with random rotation ( $18^\circ$ ), random scaling ( $\pm 10\%$ ), random crop ( $\pm 5\%$ ), random gray (20%), random blur (30%), random occlusion (40%) and random horizontal flip (50%).

**The Setup of Model.** We use a four-stacked hourglass model as the backbone. The recursive step in HG is set to 3. Each hourglass module outputs a  $64 \times 64$  feature map. We follow the training strategy introduces in [21]. Specifically, we employed an Adam optimizer and an initial learning rate of  $1 \times 10^{-3}$ . We train the model for 500 epochs and reduce the learning rate by 10 at each epoch of 200, 350, and 450. The model was trained on four GPUs (32GB NVIDIA Tesla V100), where the batch size of each GPU is 16. We

also study some commonly used tricks that can further improve the performance, such as Distribution Regularization (DR) [32] and Anisotropic Attention Module (AAM) [21].

**Datasets.** We evaluate our method on three commonly used public datasets including COFW [5], 300W [35], WFLW [48]. Meanwhile, we use a video-based dataset, 300VW [7,37,42], to make a cross-dataset validation. In detail, **COFW** contains 1,345 training images and 507 testing images with 29 landmarks. **300W** contains 3,148 training images and 689 test images. All images are labeled with 68 landmarks. We use the common setting on 300-W, where the test set is split into the common (554 images) and challenge (135 images) sets. **WFLW** is currently the most used dataset in facial landmark detection, which contains 7,500 training images and 2,500 test images with 98 landmarks. **300VW**: we use the test sets to validate the model trained on 300W. It provides three test sets: Category-A (well-lit, 31 videos with 62,135 frames), Category-B (mild unconstrained, 19 videos with 32,805 videos), and Category-C (challenging, 14 videos with 26,338 frames).

**Evaluation Metrics.** We use three commonly used evaluation metrics: Normalized Mean Error (NME), Failure Rate (FR), and Area Under Curve (AUC), to evaluate the landmark detection performance. Specifically, (1) NME: we use the inter-ocular distance in 300W, WFLW dataset, and the inter-pupils distance in COFW dataset for normalization; (2) FR: we set the thresholds by 10% for WFLW; (3) AUC: we set the thresholds by 10% for WFLW.

### 5.2. Accuracy Evaluation

We compared our STAR loss against 14 state-of-the-art baselines (as reported in Table 1). We design two interesting experiments on 300W, including within-dataset validation and cross-dataset validation. First, the results of within-dataset validation are shown in Table 1. Our method can achieve state-of-the-art performance on all test sets. Primarily, we have a 0.2 improvement in the challenge subset compared with the state-of-the-art methods. Second, for cross-dataset validation, the model is trained on 300W and evaluated on 300VW. Meanwhile, the results of cross-dataset validation are shown in Table 2. We evaluate different dis-

Method	Backbone	Param. (M)	WFLW-Full		COFW (NME)	300W (NME)			
			(NME)	(FR <sub>10%</sub> )		(AUC <sub>10%</sub> )	Full	Comm.	Chal.
LAB [48]	Hourglass	12.26	5.27	7.56	0.532	-	3.49	2.98	5.19
Wing [18]	ResNet-50	25	4.99	6.00	0.550	5.44	-	-	-
DeCaFa [11]	U-Net	10	4.62	4.84	0.563	-	3.39	2.93	5.26
HRNet [44]	HRNet-W18	9.7	4.6	4.64	-	-	3.32	2.87	5.15
Awing [45]	Hourglass	25.1	4.36	2.84	0.572	4.94	3.07	2.72	4.52
AVS + SAN [33]	ITN-CPM	-	4.39	4.08	0.591	-	3.86	3.21	6.46
DAG [28]	HRNet-W18	-	4.21	3.04	0.589	-	3.04	2.62	4.77
LUVLi [24]	DU-Net	-	4.37	3.12	0.577	-	3.23	2.76	5.16
ADNet [21]	Hourglass	13.37	4.14	2.72	<b>0.602</b>	4.68	<b>2.93</b>	<b>2.53</b>	4.58
HIH [53]	Hourglass	22.68	<b>4.08</b>	<b>2.60</b>	<b>0.605</b>	<b>4.63</b>	3.09	2.65	4.89
PIPNet [22]	ResNet-101	45.7	4.31	-	-	-	3.19	2.78	4.89
SLPT [50]	HRNet-W18C-lite	9.98	4.14	2.76	0.595	4.79	3.17	2.75	4.90
DTLD [25]	ResNet-18	13.3	<b>4.08</b>	2.76	-	-	2.96	2.59	<b>4.50</b>
RePFormer [26]	ResNet-101	-	4.11	-	-	-	3.01	-	-
Ours (smooth- $l1$ distance)	Hourglass	13.37	<b>4.02</b>	<b>2.32</b>	<b>0.605</b>	<b>4.62</b>	<b>2.87</b>	<b>2.52</b>	<b>4.32</b>

Table 1. Comparing with state-of-the-art methods on COFW, 300W and WFLW. The **best** and **second best** results are marked in colors of red and blue, respectively. We mainly report NME score on COFW, 300W, and WFLW. On WFLW, we also report the FR and AUC, whose thresholds are both set to 10%.

wo./w.	Category-A	Category-B	Category-C
$l1$	4.08/4.08	<b>3.33/3.37</b>	8.71/ <b>8.55</b>
$l2$	4.16/ <b>3.96</b>	3.48/ <b>3.35</b>	8.44/ <b>8.27</b>
smooth- $l1$	4.00/ <b>3.97</b>	3.40/ <b>3.39</b>	8.69/ <b>8.42</b>

Table 2. The cross-dataset validations with different distance functions. We train all modes on 300W with or without using STAR loss, and we report the NME scores on 300W.

tance functions combined with STAR loss. The results show that our method achieves an average improvement of 0.2 in Category-C. Notably, the subset Category-C is the most challenging one, where the video frames are blurred and with strong occlusions. We observe a significant improvement in the challenge test of both within-dataset and cross-dataset validation, which verifies the effectiveness of our method. Meanwhile, our method achieves state-of-the-art performance in the COFW dataset.

The WFLW is a very challenging benchmark because it contains many images under adverse conditions such as pose, make-up, illumination, blur, and expression. We report the results in Table 1, where our method still achieves the best results. In detail, our STAR loss achieves an improvement of 0.06, 0.28 in NME, and FR<sub>10%</sub> metrics of the full test set, respectively, when compared to the state-of-the-art methods. Compared with the increase in COFW and 300-W datasets, the improvement achieved by STAR loss in the WFLW dataset is sufficient. We conclude that STAR loss has a better performance over a big dataset, *i.e.*, the proposed STAR loss will mitigate the effect of annotation

wo./w.	$l1$	smooth- $l1$	$l2$	Wing
COFW				
basic	5.00/ <b>4.76</b>	5.02/ <b>4.78</b>	5.59/ <b>4.89</b>	4.94/ <b>4.79</b>
+DR	4.78/ <b>4.70</b>	4.76/ <b>4.70</b>	5.06/ <b>4.73</b>	4.76/ <b>4.67</b>
+AAM	5.01/ <b>4.77</b>	4.88/ <b>4.62</b>	5.44/ <b>4.77</b>	4.95/ <b>4.70</b>
300W				
basic	3.13/ <b>2.98</b>	3.12/ <b>3.00</b>	3.45/ <b>3.05</b>	3.12/ <b>3.02</b>
+DR	3.04/ <b>2.94</b>	3.02/ <b>2.96</b>	3.20/ <b>2.99</b>	3.04/ <b>2.96</b>
+AAM	2.97/ <b>2.89</b>	2.93/ <b>2.87</b>	3.08/ <b>2.91</b>	2.93/ <b>2.88</b>
WFLW				
basic	4.35/ <b>4.16</b>	4.34/ <b>4.17</b>	4.79/ <b>4.31</b>	4.32/ <b>4.17</b>
+DR	4.17/ <b>4.13</b>	4.18/ <b>4.10</b>	4.29/ <b>4.21</b>	4.18/ <b>4.12</b>
+AAM	4.19/ <b>4.05</b>	4.14/ <b>4.02</b>	4.37/ <b>4.11</b>	4.15/ <b>4.05</b>

Table 3. The robustness analysis on COFW, 300W, and WFLW. We report the NME score for basic model trained with or without using our STAR loss. We also study the influence of two widely used tricks, like DR and AAM.

error and benefit from the large amounts of noise data.

### 5.3. Ablation Study

**Evaluation on Different Distance Functions.** To investigate the effect of the distance function, we test the STAR loss with the commonly used distance function, such as  $l1$ -distance,  $l2$  distance, smooth- $l1$  distance, and Wing distance [18] in COFW, 300W and WFLW dataset, respectively. Note that the smooth- $l1$  distance used here is formulated as below:

$$\begin{cases} \frac{0.5}{s}(y_t - \mu)^2, & |y_t - \mu| < s \\ |y_t - \mu| - 0.5s, & \text{otherwise} \end{cases}, \quad (8)$$




Figure 5. Qualitative results of different distance functions w.o/w. STAR loss on WFLW dataset. Green and red points represent the predicted and ground-truth points, respectively. The yellow circles indicate the clear failures, which are solved with the help of STAR loss. (Best view in color and zoom in.)

where  $s$  is a threshold that is usually set as  $s = 0.01$ . We refer to the regression loss defined in Eq.(2) as a baseline. The results are shown in Table 3. Compared with the baseline, STAR loss can achieve an average improvement of 0.34, 0.19, and 0.25 on COFW, 300W, and WFLW among different distance functions, respectively, with negligible computational. Meanwhile, we find that STAR loss has a significant effect when combined with  $l_2$ -distance.

Moreover, we visualize the failure cases of different distance functions with and without STAR loss on WFLW dataset. As shown in Figure 5, the results with STAR exhibits better structural constraints. We infer that STAR works as a label regularization, which helps the model avoid over-fitting ambiguity annotation and pay more attention to the relation between landmarks.

**Evaluation on Different Distribution Normalization.** We evaluate our proposed STAR loss with different distribution normalizations, such as Distribution Regularization (DR) [32] and Anisotropic Attention Module (AAM) [21]. DR is a stricter regularization that forces the heatmap to resemble a Gaussian distribution. AAM is an attention module that learns a point-edge heatmap as an attention mask and forces the heatmap to resemble a mixture of Gaussian distribution and adjacent boundary [24] distribution. In DR, we use Jensen-Shannon divergence as the divergence measure and set the Gaussian  $\sigma$  equals to 1.0. Regarding AAM, we follow the setting in paper [21].

In Table 3, both of them have a positive effect: (1) In the WFLW dataset, STAR loss achieve an average improvement of 0.07 and 0.16 on DR and AAM, respectively; (2) In the COFW dataset, the improvement is 0.14 and 0.35, respectively; (3) In the 300W dataset, the improvement is

$w$	1	2	3	5	10	20
NME	4.17	4.20	4.18	4.18	4.18	4.20

Table 4. The analysis of  $w$ . We report the NME scores for STAR loss with smooth- $l_1$  distance with the varying  $w$  on WFLW.

0.11 and 0.09, respectively; We observe that the AAM is better suited to STAR loss because the boundary information in AAM further helps. And the best results among the three datasets are taken with the combination of smooth- $l_1$  distance, AAM, and STAR loss, which refers to the default model setup in accuracy evaluation. Notice that the model used for visualization in Figure 1, and Figure 3 is trained with the combination of  $\mathcal{L}_2$  and DR.

**Evaluation on Different Input Image Resolution.** We investigate the influence of different input image resolutions (64px, 128px, and 256px). The results are shown in Figure 4, STAR loss has a positive effect on all input image resolution. Specifically, STAR loss can achieve an average improvement of 0.20, 0.30, and 0.25 in 64px, 128px, and 256px among different distance functions, respectively. Notice that the baseline with STAR loss in 128px is equal to or better than the baseline in 256px. It means that STAR loss is not very sensitive to the input image size, which encourages using STAR loss in a small model.

**Evaluation on Different Restrictions.** Firstly, we investigate the value restriction and discuss the effect of its weight  $w$  in STAR loss. As shown in Table 4, STAR loss is not sensitive to  $w$ , and we choose  $w = 1$  as the default setting in our experiments. Then, we analyze the detach restriction. As shown in Table 5, the restriction is essential, and the effect of detach restriction is comparable with the value restriction. This observation indicates that the improvement



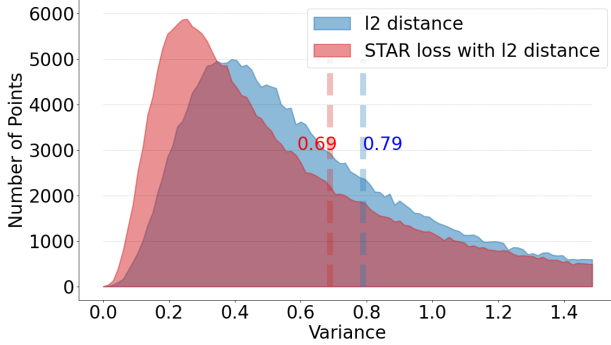


Figure 6. The variance distribution. We visualize the variance statistic of five models trained under the same experimental setting with and without STAR loss. The red and blue lines indicate the corresponding mean value of  $l_2$  distance function with and without STAR loss, respectively.

is mainly contributed by STAR loss but not the value restriction term.

**Evaluation on Stability.** We evaluate the effect of STAR loss on stability. To quantitatively analyze the stability, we design a toy experiment: 1) train  $N = 5$  models under the same experimental setting with and without STAR loss, respectively; 2) count the variance of  $N$  models’ predictions on the WFLW test set. As shown in Figure 6, the predictions of models trained with our STAR loss distribute more intensively within a small variance range. Meanwhile, the mean value of prediction variances drops from 0.79 to 0.69. It indicates that STAR loss does mitigate the impact of ambiguous annotations and makes prediction more stable.

#### 5.4. Discussion with Related Solutions

**ADL v.s. STAR Loss.** In ADNet [21], Huang *et al.* introduce an Anisotropic Direction Loss (ADL), which imposes strong constraints in the normal direction and weak constraints in the tangent direction for landmarks on the facial boundary. The normal direction is calculated by the slope from its adjacent points, and the constraint strength is a hyperparameter. We argue that the hand-crafted design degrades its performance. Compared with it, STAR loss estimates the direction and constraint weight via PCA of heatmaps, which is a fine-grained and self-adaptive scheme. Then, we replaced ADL with STAR in ADNet and adopted the same setup for a fair comparison. As shown in Table 6, STAR loss brings further improvements of 0.12, 0.28, 0.08, 0.07, 0.13, 0.11, 0.10 in NME in full WFLW and six sub-categories, respectively. Besides, STAR is plug-and-play for heatmap-based methods with an arbitrary number of landmarks without extra direction computation design.

**LUVLi Loss v.s. STAR Loss.** LUVLi [24] introduces an uncertainty-based loss (*e.g.*, Gaussian Log-Likelihood (GLL)) to facial landmark detection. The GLL consists of a Mahalanobis distance and a regularization term. As dis-

WFLW	basic	+AAM
baseline	4.34	4.14
wo. restriction	4.60	4.61
w. value restriction	4.17	4.02
w. detach restriction	4.20	4.05

Table 5. The analysis of restriction method. We report the NME scores for STAR loss (smooth- $l_1$  distance) without and with different restriction methods on WFLW.

WFLW	Full	Pose	Exp.	Ill.	Make.	Occl.	Blur
ADL [21]	4.14	6.96	4.38	4.09	4.05	5.06	4.79
STAR	<b>4.02</b>	<b>6.76</b>	<b>4.27</b>	<b>3.97</b>	<b>3.84</b>	<b>4.80</b>	<b>4.58</b>

Table 6. The comparison between ADL and STAR loss.

cussed in Sec. 5.3, the value restriction is unnecessary for STAR loss. So, we mainly compare the Mahalanobis distance with  $\mathcal{L}_{STAR}$  by a toy experiment: replace  $\mathcal{L}_{STAR}$  with Mahalanobis distance and adopt the same setup for a fair comparison. The experiment shows that replacing  $\mathcal{L}_{STAR}$  with Mahalanobis can take 4.32 in NME in WFLW. However, our  $\mathcal{L}_{STAR}$  takes 4.17, which has a further 0.15 NME improvement in WFLW. So, compared with LUVLi, STAR is lighter and better: 1) propose the heatmap-based covariance estimation, simplifying model structure; 2) design the loss function from the perspective of ambiguity reduction, thus achieving significant performance gain.

## 6. Conclusion

We study the semantic ambiguity problem in facial landmark detection and propose a self-adaptive ambiguity reduction method, STAR loss. We first observe that the predicted distribution can represent semantic ambiguity. Then we use PCA to indicate the character of the predicted distribution and indirectly formulate the direction and intensity of semantic ambiguity. Based on this, STAR loss adaptively suppresses the prediction error in the ambiguity direction to mitigate the impact of ambiguity annotation in training. With the help of STAR loss, our method achieves state-of-the-art performance on three benchmarks.

**Acknowledgement.** This work was supported in part by National Key R&D Program of China (No.2022ZD0118202), the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No.U21B2037, No.U22B2051, No.62176222, No.62176223, No.62176226, No.62072386, No.62072387, No.62072389, No.62002305 and No.62272401), the National Science Foundation of Fujian Province of China (No.2021J01002, No.2022J06001) and the Grant-in-Aid for JSPS Fellows Grant (No. 21F20377).



## References

- [1] Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, and Gang Hua. Towards open-set identity preserving face synthesis. In *CVPR*, 2018. 1
- [2] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3d morphable model. *TPAMI*, 2003. 2
- [3] Bjorn Browatzki and Christian Wallraven. 3fabrec: Fast few-shot face alignment by reconstruction. In *CVPR*, 2020. 1, 3
- [4] Adrian Bulat, Enrique Sanchez, and Georgios Tzimiropoulos. Subpixel heatmap regression for facial landmark localization. In *BMVC*, 2021. 3
- [5] Xavier P Burgos-Artizzu, Pietro Perona, and Piotr Dollár. Robust face landmark estimation under occlusion. In *ICCV*, 2013. 2, 5, 12
- [6] Xudong Cao, Yichen Wei, Fang Wen, and Jian Sun. Face alignment by explicit shape regression. *IJCV*, 2014. 13
- [7] Grigoris G Chrysos, Epameinondas Antonakos, Stefanos Zafeiriou, and Patrick Snape. Offline deformable face tracking in arbitrary videos. In *ICCVW*, 2015. 5
- [8] Timothy F. Cootes, Gareth J. Edwards, and Christopher J. Taylor. Active appearance models. *TPAMI*, 2001. 2
- [9] David Cristinacce and Tim Cootes. Automatic feature localisation with constrained local models. *Pattern Recognition*, 2008. 2
- [10] Radek Danecek, Michael J. Black, and Timo Bolkart. EMOCA: Emotion driven monocular face capture and animation. In *CVPR*, 2022. 1
- [11] Arnaud Dapogny, Kevin Bailly, and Matthieu Cord. Decafa: Deep convolutional cascade for face alignment in the wild. In *ICCV*, 2019. 3, 6, 12, 13
- [12] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 1
- [13] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *CVPRW*, 2019. 1
- [14] Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Style aggregated network for facial landmark detection. In *CVPR*, 2018. 3, 12
- [15] Xuanyi Dong, Yi Yang, Shih-En Wei, Xinshuo Weng, Yaser Sheikh, and Shoou-I Yu. Supervision by registration and triangulation for landmark detection. *TPAMI*, 2020. 2, 3
- [16] Xuanyi Dong, Shoou-I Yu, Xinshuo Weng, Shih-En Wei, Yi Yang, and Yaser Sheikh. Supervision-by-registration: An unsupervised approach to improve the precision of facial landmark detectors. In *CVPR*, 2018. 1, 3
- [17] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ToG*, 2021. 1
- [18] Zhen-Hua Feng, Josef Kittler, Muhammad Awais, Patrik Huber, and Xiao-Jun Wu. Wing loss for robust facial landmark localisation with convolutional neural networks. In *CVPR*, 2018. 3, 4, 6, 12, 13
- [19] Zhen-Hua Feng, Josef Kittler, William Christmas, Patrik Huber, and Xiao-Jun Wu. Dynamic attention-controlled cascaded shape regression exploiting training data augmentation and fuzzy-set sample weighting. In *CVPR*, 2017. 12
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1
- [21] 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, 3, 5, 6, 7, 8, 11, 12, 13
- [22] Haibo Jin, Shengcai Liao, and Ling Shao. Pixel-in-pixel net: Towards efficient facial landmark detection in the wild. *IJCV*, 2021. 6, 12, 13
- [23] Amit Kumar and Rama Chellappa. Disentangling 3d pose in a dendritic cnn for unconstrained 2d face alignment. In *CVPR*, 2018. 12
- [24] 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, 3, 6, 7, 8, 12, 13
- [25] Hui Li, Zidong Guo, Seon-Min Rhee, Seungju Han, and Jae-Joon Han. Towards accurate facial landmark detection via cascaded transformers. In *CVPR*, 2022. 6, 12
- [26] Jinpeng Li, Haibo Jin, Shengcai Liao, Ling Shao, and Pheng-Ann Heng. Repformer: Refinement pyramid transformer for robust facial landmark detection. In *IJCAI*, 2022. 3, 6, 13
- [27] Weijian Li, Yuhang Lu, Kang Zheng, Haofu Liao, Chihung Lin, Jiebo Luo, Chi-Tung Cheng, Jing Xiao, Le Lu, Chang-Fu Kuo, et al. Structured landmark detection via topology-adapting deep graph learning. In *ECCV*, 2020. 3, 13
- [28] Weijian Li, Yuhang Lu, Kang Zheng, Haofu Liao, Chihung Lin, Jiebo Luo, Chi-Tung Cheng, Jing Xiao, Le Lu, Chang-Fu Kuo, et al. Structured landmark detection via topology-adapting deep graph learning. In *ECCV*, 2020. 6
- [29] Chunze Lin, Beier Zhu, Quan Wang, Renjie Liao, Chen Qian, Jiwen Lu, and Jie Zhou. Structure-coherent deep feature learning for robust face alignment. *TIP*, 2021. 3, 12, 13
- [30] Zhiwei Liu, Xiangyu Zhu, Guosheng Hu, Haiyun Guo, Ming Tang, Zhen Lei, Neil M Robertson, and Jinqiao Wang. Semantic alignment: Finding semantically consistent ground-truth for facial landmark detection. In *CVPR*, 2019. 1, 2, 3
- [31] Iain Matthews and Simon Baker. Active appearance models revisited. *IJCV*, pages 135–164, 2004. 2
- [32] Aiden Nibali, Zhen He, Stuart Morgan, and Luke Prendergast. Numerical coordinate regression with convolutional neural networks. In *ECCV*, 2018. 1, 2, 3, 5, 7
- [33] Shengju Qian, Keqiang Sun, Wayne Wu, Chen Qian, and Jiaya Jia. Aggregation via separation: Boosting facial landmark detector with semi-supervised style translation. In *ICCV*, 2019. 1, 6
- [34] Shengju Qian, Keqiang Sun, Wayne Wu, Chen Qian, and Jiaya Jia. Aggregation via separation: Boosting facial landmark detector with semi-supervised style translation. In *ICCV*, 2019. 3

- [35] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *ICCVW*, 2013. 2, 5
- [36] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 1
- [37] Jie Shen, Stefanos Zafeiriou, Grigoris G Chrysos, Jean Kossai, Georgios Tzimiropoulos, and Maja Pantic. The first facial landmark tracking in-the-wild challenge: Benchmark and results. In *ICCVW*, 2015. 5
- [38] Ying Tai, Yicong Liang, Xiaoming Liu, Lei Duan, Jilin Li, Chengjie Wang, Feiyue Huang, and Yu Chen. Towards highly accurate and stable face alignment for high-resolution videos. In *AAAI*, 2019. 1, 3
- [39] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 1
- [40] Zhiqiang Tang, Xi Peng, Shijie Geng, Lingfei Wu, Shaoting Zhang, and Dimitris Metaxas. Quantized densely connected u-nets for efficient landmark localization. In *ECCV*, 2018. 12
- [41] George Trigeorgis, Patrick Snape, Mihalis A Nicolaou, Epameinondas Antonakos, and Stefanos Zafeiriou. Mnemonic descent method: A recurrent process applied for end-to-end face alignment. In *CVPR*, 2016. 3, 12
- [42] Georgios Tzimiropoulos. Project-out cascaded regression with an application to face alignment. In *CVPR*, 2015. 5
- [43] Roberto Valle, Jose M Buenaposada, Antonio Valdes, and Luis Baumela. A deeply-initialized coarse-to-fine ensemble of regression trees for face alignment. In *ECCV*, 2018. 12
- [44] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *TPAMI*, 2019. 3, 6, 12
- [45] Xinyao Wang, Liefeng Bo, and Li Fuxin. Adaptive wing loss for robust face alignment via heatmap regression. In *ICCV*, 2019. 3, 4, 6, 12, 13
- [46] Erroll Wood, Tadas Baltrusaitis, Charlie Hewitt, Matthew Johnson, Jingjing Shen, Nikola Milosavljevic, Daniel Wilde, Stephan Garbin, Toby Sharp, Ivan Stojiljkovic, et al. 3d face reconstruction with dense landmarks. *arXiv*, 2022. 1
- [47] Erroll Wood, Tadas Baltrušaitis, Charlie Hewitt, Sebastian Dziadzio, Matthew Johnson, Virginia Estellers, Thomas J. Cashman, and Jamie Shotton. Fake it till you make it: Face analysis in the wild using synthetic data alone, 2021. 1, 3
- [48] Wayne Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, and Qiang Zhou. Look at boundary: A boundary-aware face alignment algorithm. In *CVPR*, 2018. 1, 2, 3, 5, 6, 12, 13
- [49] Wenyao Wu and Shuo Yang. Leveraging intra and inter-dataset variations for robust face alignment. In *CVPRW*, 2017. 12, 13
- [50] Jiahao Xia, Weiwei Qu, Wenjian Huang, Jianguo Zhang, Xi Wang, and Min Xu. Sparse local patch transformer for robust face alignment and landmarks. In *CVPR*, 2022. 3, 6, 12, 13
- [51] Jiahao Xia, Weiwei Qu, Wenjian Huang, Jianguo Zhang, Xi Wang, and Min Xu. Sparse local patch transformer for robust face alignment and landmarks inherent relation learning. In *CVPR*, 2022. 1, 3
- [52] Shengtao Xiao, Jiashi Feng, Junliang Xing, Hanjiang Lai, Shuicheng Yan, and Ashraf Kassim. Robust facial landmark detection via recurrent attentive-refinement networks. In *ECCV*, 2016. 3, 12
- [53] Lan Xing, Hu Qinghao, and Cheng Jian. HIH: towards more accurate face alignment via heatmap in heatmap, 2021. 1, 3, 6, 12, 13
- [54] Xuehan Xiong and Fernando De la Torre. Supervised descent method and its applications to face alignment. In *CVPR*, 2013. 12, 13
- [55] Jing Yang, Qingshan Liu, and Kaihua Zhang. Stacked hour-glass network for robust facial landmark localisation. In *CVPR*, 2017. 1, 3
- [56] Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Facial landmark detection by deep multi-task learning. In *ECCV*, 2014. 1, 3, 12
- [57] 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 *CVPR*, 2021. 3
- [58] Erjin Zhou, Haoqiang Fan, Zhimin Cao, Yuning Jiang, and Qi Yin. Extensive facial landmark localization with coarse-to-fine convolutional network cascade. In *ICCVW*, 2013. 3
- [59] Shizhan Zhu, Cheng Li, Chen Change Loy, and Xiaoou Tang. Face alignment by coarse-to-fine shape searching. In *CVPR*, 2015. 12, 13
- [60] Xu Zou, Sheng Zhong, Luxin Yan, Xiangyun Zhao, Jiahuan Zhou, and Ying Wu. Learning robust facial landmark detection via hierarchical structured ensemble. In *CVPR*, 2019. 3