

Gaussian Vector: An Efficient Solution for Facial Landmark Detection

Yilin Xiong^{1,2}[0000-0003-1989-916X], Zijian Zhou²[0000-0003-3315-3962], Yuhao Dou²[0000-0002-4407-9609], and Zhizhong Su²[0000-0003-2312-9985]

¹ Central South University, Changsha, Hunan, CN

² Horizon Robotics

yilin.xiong@csu.edu.cn

{zijian.zhou, yuhao.dou, zhizhong.su}@horizon.ai

Abstract. Significant progress has been made in facial landmark detection with the development of Convolutional Neural Networks. The widely-used algorithms can be classified into coordinate regression methods and heatmap based methods. However, the former loses spatial information, resulting in poor performance while the latter suffers from large output size or high post-processing complexity. This paper proposes a new solution, Gaussian Vector, to preserve the spatial information as well as reduce the output size and simplify the post-processing. Our method provides novel vector supervision and introduces Band Pooling Module to convert heatmap into a pair of vectors for each landmark. This is a plug-and-play component which is simple and effective. Moreover, Beyond Box Strategy is proposed to handle the landmarks out of the face bounding box. We evaluate our method on 300W, COFW, WFLW and JD-landmark. That the results significantly surpass previous works demonstrates the effectiveness of our approach.

1 Introduction

Facial landmark detection is a critical step for face-related computer vision applications, e.g. face recognition [1–3], face editing [4] and face 3D reconstruction [5–8]. Researchers have achieved great success in this field, especially after using Convolutional Neural Networks (CNNs). The popular algorithms can be divided into coordinate regression methods and heatmap based methods.

Coordinate regression methods use Fully Connected (FC) layers to predict facial landmark coordinates [9, 10]. However, the compacted feature maps before the final FC layers cause spatial information loss, leading to performance degradation.

As the advent of Fully Convolution Network [11] and encoder-decoder structure [12], a variety of heatmap based methods [13–18] sprung up to preserve the spatial information. Heatmap based methods predict probability response maps, usually called heatmap, where each pixel predicts the probability that it is the landmark location, and find the highest response position by argmax operation. Heatmap based methods boost the performance compared with coordinate regression methods by better utilizing spatial information.




Fig. 1. Comparison of existing methods and our proposed method. (a) coordinate regression methods and (b) heatmap based methods utilize the scalar coordinates and the heatmap to represent the spatial location of landmarks respectively. (c) Gaussian Vector locates a landmark via a pair of vectors.

However, heatmap based methods suffer from following problems. **1)** The large output heatmap tensor and complicated post-processing including argmax operation cause heavy burden of data transmission and computation in embedded systems. **2)** The ground truth heatmap consists of a small patch of foreground pixels with positive values and the remaining is full of background pixels with zero values. The background dominates the training process, which causes slow convergence. **3)** Only very few pixels around the highest response position is exploited to locate the landmark. Spatial information is not fully used, causing detection error. **4)** The input of landmark detection model is the face region of interest (RoI) acquired by a face detector. An imperfect face detector causes face bounding box shifting resulting in some of the facial landmarks out of the bounding box. It is difficult for the heatmap based methods to construct a Gaussian surface to locate the landmarks out of the bounding box.

In this case, we are motivated to propose a new facial landmark detection algorithm, which not only maintains the spatial information but also overcome the problems above.

First of all, inspired by the heatmap based methods, we use Gaussian formula to encode landmark coordinates into vectors as supervision. Vector label raises the proportion of foreground pixels in ground truth, which is beneficial to convergence. **Next**, Band Pooling Module (BPM) is proposed to convert the $h \times w$ output heatmap into $h \times 1$ and $1 \times w$ vectors as prediction. On the one hand, vector prediction greatly eases the data transmission burden and simplifies the post-processing. On the other hand, pooling operation helps the network take advantage of more spatial information, which further improves the performance. **Last but not least**, we propose Beyond Box Strategy to handle landmarks out of the bounding box. The strategy helps our proposed method to reduce the error caused by bounding box shifting.

We compare our approach with both coordinate regression methods and heatmap based methods in Fig. 1. Considering the utilization of Gaussian formula and vector prediction, we call our method Gaussian Vector (GV).

We evaluate our method on four popular facial landmark detection benchmarks including 300W [19], COFW [20], WFLW [16] and the recent Grand Challenge of 106-p Facial Landmark Localization benchmark [21] (JD-landmark for short). Our approach outperforms state-of-the-art methods in 300W, COFW and WFLW. Without bells and whistles, we achieve the second place in JD-landmark by utilizing a single model.

In summary, our main contributions include:

- We propose an efficient method, Gaussian Vector, for facial landmark detection, which boosts the model performance and reduces the system complexity.
- We introduce novel vector label to encode landmark location for supervision, which alleviates the imbalance of foreground and background in heatmap supervision.
- We design a plug-and-play module, Band Pooling Module (BPM), that effectively converts heatmap into vectors. By this means, the model takes more spatial information into account yet outputs smaller tensor.
- We propose the Beyond Box Strategy, that enables our method to predict landmarks out of the face bounding box.
- Our method achieves state-of-the-art results on 300W, COFW, WFLW and JD-landmark.

2 Related Work

Facial landmark detection, or face alignment, has become a hot topic in computer vision for decades. Early classical methods include Active Appearance Models (AAMs) [22–24], Active Shape Models (ASMs) [25], Constrained Local Models (CLMs) [26], Cascaded Regression Models [27, 20, 28, 29] and so on. With the development of deep learning, CNN-based methods have dominated this field. The prevailing approaches fall into two categories.

Coordinate regression methods take the face image as input and output landmark coordinates directly. This kind of methods get better results than the classic without extra post-processing. Sun *et al.* [9] borrowed the idea of Cascaded Regression and combined serial sub-networks to reduce the location error. TCDCN [10] researched the possibility of improving detection robustness through multi-task learning. MDM [30] employed Recurrent Neural Network (RNN) to replace the cascaded structure in Cascaded Regression. Feng *et al.* [31] presented a new loss function to pay more attention to small and medium error. Most of the approaches above are proposed with cascaded structure, which largely increases the system complexity.

Heatmap based methods are popular in recent years. Benefit from the encoder-decoder structure, heatmap based methods preserve spatial information by upsampling and significantly improve state-of-the-art performance on all benchmarks. DAN [15] utilized the heatmap information for each stage in the cascaded network and reduced the localization error. Stacked Hourglass [17]

introduces repeated bottom-up, top-down architecture for points detection and proposes a shift-based post-processing strategy. HRNet [32, 33] presented a novel network architecture for landmark detection, which maintained high-resolution representations through the design of dense fusions on cross-stride layers. LAB [16] proposed a boundary-aware method, which used the stacked hourglass [17] to regress the boundary heatmap for capturing the face geometry structure. Another landmarks regressor fused the structure information in multiple stages and got the final results. AWing [14] followed the work of Wing loss [31] and employed the revised version in the heatmap based methods. Also, boundary information and CoordConv [34] are utilized for promoting landmark detection.

However, recovering landmark coordinates from the heatmap still remains to be simplified. Efforts has been devoted to improve the efficiency of this process recently [35–37]. Specifically, Sun *et al.* [35] proposed a differentiable operation to convert heatmap into numerical coordinates for end-to-end training. Yet the calculation of exponent is not computation-efficient in embedded systems.

We believe that brand new supervision is necessary to address problems above. To this end, we encode the location information into vectors for supervision.

3 Our approach

We first present an overview of our approach in Sec. 3.1. Then Sec. 3.2 introduces the generation of vector label. After that, we illustrate the implementation details of our core component BPM in Sec. 3.3. Finally, we propose a strategy to predict the landmarks out of the face bounding box in Sec. 3.4.

3.1 Overview

An overview of our pipeline is shown in Fig. 2. Similar to heatmap based methods, we use the encoder-decoder structure in our model. The encoder, with fully convolution layers, extracts spatial information on different scales. For the decoder, Feature Pyramid Network (FPN) [38] is adopted. It combines the features of both higher and lower levels through the top-down and bottom-up pathways, which improves the pixel-level detection tasks.

The network input is cropped and normalized face image. The network outputs 1/4 resolution heatmap to the original input, where each channel is corresponding to one landmark. A plug-and-play component Band Pooling Module (BPM) compresses heatmap into vectors for end-to-end training.

During training, the predicted vector is supervised by the vector label which is generated from Gaussian formula. More details about vector label can be found in Sec. 3.2. Mean Square Error (MSE) is used as loss function.

During inference, the maximum of predicted vectors indicates the horizontal and vertical location of the landmark respectively. A shift strategy [17], widely used in heatmap based methods, is applied according to its neighbor pixels to reduce location error.

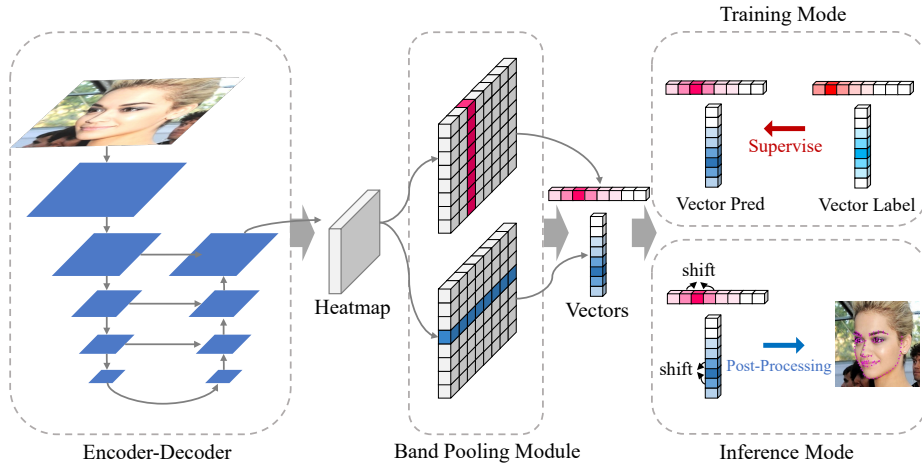


Fig. 2. The overview of the proposed method. Given the face image inputs to the encoder-decoder structure, we can obtain the output heatmap and generate the prediction vectors through Band Pooling Module, match it with vector label for training. Landmark coordinates can be recovered from the vectors by locating the maximum value indexes. Shift strategy [17] is used to reduce quantization error.

3.2 Vector Label

Vector label $\hat{\mathbf{G}} = (\hat{\mathbf{x}}, \hat{\mathbf{y}})$ for each landmark is generated by the Gaussian formula. Take vector $\hat{\mathbf{x}} \in \mathbb{R}^w$ as an example, where w means the width of the output heatmap. For the landmark coordinate $p(x_0, y_0)$, we first calculate the Euclidean distance between each pixel and x_0 in the vector, getting the distance vector $\mathbf{d} = \{d_x^i\}_{i=0}^{w-1}$, each d_x^i as:

$$d_x^i = \sqrt{(x_i - x_0)^2}, i \in [0, w) \quad (1)$$

Then, we use quasi-Gaussian distribution to transform the distance vector \mathbf{d} to vector label $\hat{\mathbf{x}} = \{\hat{x}^i\}_{i=0}^{w-1}$ with standard deviation σ . The elements are defined as:

$$\hat{x}^i = \begin{cases} e^{-\frac{d_x^i{}^2}{2\sigma^2}}, & \text{if } d_x^i < 3\sigma \text{ and } d_x^i \neq 0 \\ 1 + \theta, & \text{if } d_x^i = 0 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

We recommend $\sigma=2$ in most cases. Compared with the normal Gaussian distribution, we reinforce the distribution peak to enhance the supervision by adding a positive constant θ . By this enhancement, we enlarge the value gaps between the horizontal landmark location x_0 and the others to avoid deviation of the highest response position. In the same way, we can get $\hat{\mathbf{y}} \in \mathbb{R}^h$, where h is the heatmap height.

The proposed vector label is beneficial to the model training in two aspects. On the one hand, generating vectors instead of heatmap accelerates the label preparing. On the other hand, extreme foreground background imbalance, as mentioned above, is alleviated in a straightforward way.

3.3 Band Pooling Module

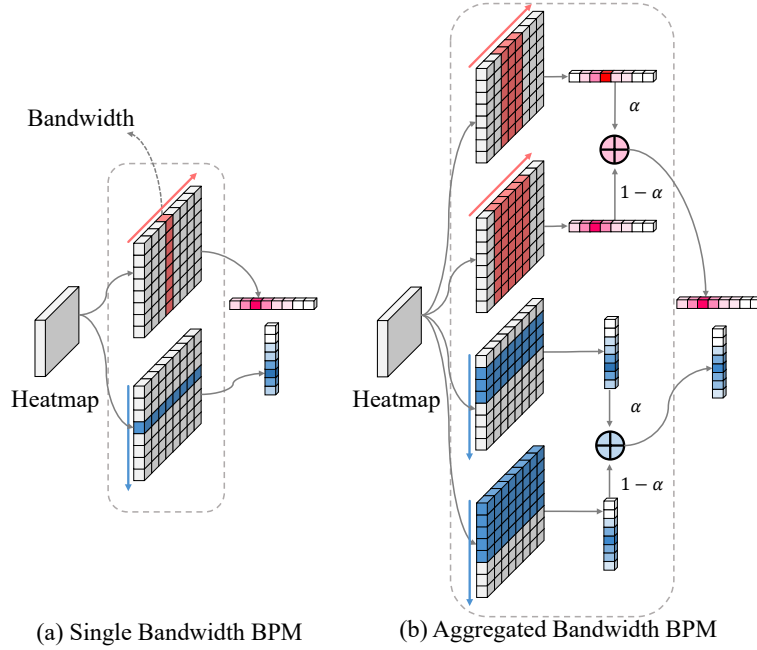


Fig. 3. The illustration of BPM. The vertical band B_v (in red) and the horizontal band B_h (in blue) slide on the output heatmap. We illustrate the vector generated by single bandwidth (a) and aggregated bandwidth (b). Specifically, the feature vectors of bandwidth $l=3$ and $l=5$ are able to be aggregated by weighted elementwise sum.

For end-to-end training, a plug-and-play component Band Pooling Module is presented. The BPM aims to perform the following conversion:

$$\mathbf{H} \rightarrow \mathbf{G} \quad (3)$$

\mathbf{H} is the output feature map and $\mathbf{G} = (\mathbf{x}, \mathbf{y})$ refers to the predicted vectors, matching with vector label $\hat{\mathbf{G}}$ for training.

Band Pooling Module, consisting of horizontal band $B_h \in \mathbb{R}^{l \times w}$ and vertical band $B_v \in \mathbb{R}^{h \times l}$, plays a crucial role in the process of generating predicted vectors. The long side of the band is the same as side length of heatmap (w for

B_h and h for B_v). The bandwidth, denoted as l , is a much smaller odd number. The two long and thin bands slide on the output heatmap and average the values of pixels that fall into the band regions, generating predicted vectors \mathbf{x} and \mathbf{y} . The pooling operation is shown in Fig. 3(a). It can be described as:

$$\mathbf{x}_l = AvgPooling(\mathbf{H}; B_v, l) \quad (4)$$

The bandwidth l is adjustable, which controls the receptive field size of the vector elements. Undersize bandwidth strictly limits the receptive field, resulting in loss of spatial information, which is essential for the landmark detection. Conversely, overextended receptive field brings in redundant information from the background or nearby landmarks, leading to confusion. In practical, bandwidth l is suggested to be chosen from 1 and 7 and we show the effect of different l in Table. 4.5.

In addition, we extend the basic BPM through fusing the vectors generated by different bandwidths. In general considering the performance and complexity, we aggregate the vectors from only two different bandwidth l_1 and l_2 by weights.

$$\mathbf{x} = \alpha \mathbf{x}_{l_1} + (1 - \alpha) \mathbf{x}_{l_2}, 0 < \alpha < 1 \quad (5)$$

Generally, α is set to be 0.5. Similarly, we can get \mathbf{y} . Zero padding is necessary on the edges when $l > 1$. The aggregation is shown in Fig. 3(b).

In this way, the model comprehensively takes spatial information into account via the bigger bandwidth and pay more attention to the adjacent region through the smaller.

By using BPM, we narrow the search regions from 2D maps to 1D vectors when seeking the maximum of the prediction, leading to large reduction on post-processing complexity (from $O(N^2)$ to $O(N)$).

3.4 Beyond the Box

The face bounding box has a significant impact on facial landmark detection. An oversize bounding box brings in redundant information, especially in the face-crowded image. But a tight box may cause facial landmarks to locate out of the face bounding box. To solve this problem, we propose an effective strategy to predict outer landmarks.

In general, the predicted vectors present smoother quasi-Gaussian distribution than vector label, where the maximum implies the landmark location. When the landmark locates inside, the maximum position is in the middle of the vector. On the contrary, when the landmark falls out of the input, the maximum is close to one of the endpoints and the quasi-Gaussian distribution is cut off. An example of the outside landmark is displayed in Fig. 4. Apparently, the maximum value of \mathbf{y} is much smaller than the normal ones.

Fortunately, the maximum value tells us whether the landmark goes beyond and how far it goes. If the distribution is cut off and the maximum locate at endpoint (0 or $w - 1$ for index s), we assume x^s obeys the distribution:

$$x^s = \tau e^{-\frac{d^2}{2\sigma^2}} \quad (6)$$

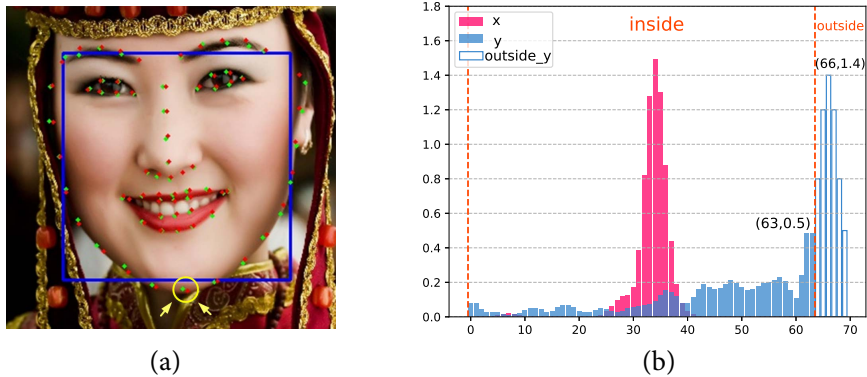


Fig. 4. (a) A visualization of the prediction with the Beyond Box Strategy. The ground truth and prediction landmarks are shown in red and green, respectively. The blue rectangle refers to the face bounding box. (b) We use a histogram to visualize the predicted vector of landmark in the yellow circle of the left image. The red dashed line shows the distribution of predicted vector inside, while the outside shows the estimated distribution from Beyond Box Strategy when the maximum locates at the endpoint of vector (e.g. \mathbf{y}). We calculate the position of ideal distribution peak, i.e. (66, 1.4) in histogram, according to Eq. 6.

where d indicates the distance between s and the distribution peak. τ is a scale parameter and σ is the standard deviation, which control the peak value and smoothness of distribution respectively. The assumption helps the model to estimate the location of the distribution peak, which works only if the vector maximum locates at one of the endpoints, just like the blue vector in Fig. 4(b).

Assume that the maximum locates at the endpoint, we can get the revised landmark location:

$$loc = \begin{cases} -\sqrt{-2\sigma^2 \log(\frac{x^0}{\tau})} & \text{if } \underset{i}{\operatorname{argmax}} x^i = 0 \\ w - 1 + \sqrt{-2\sigma^2 \log(\frac{x^{w-1}}{\tau})} & \text{if } \underset{i}{\operatorname{argmax}} x^i = w - 1 \end{cases} \quad (7)$$

Generally, τ ranges from $0.5 \times (1 + \theta)$ to $0.8 \times (1 + \theta)$ and σ equals to that in Eq. 2.

4 Experiments

To verify the effectiveness of our method, a series of experiments are conducted on public datasets in Sec. 4.4. Ablation study is carried out on 300W in Sec. 4.5 to test the proposed components. In Sec. 4.6, we perform efficiency analysis to validate the points made above. In addition, we give comprehensive analysis of the merits of vector supervision in the supplement material.

4.1 Datasets

300W [19] is a widely-used benchmark for facial landmark detection, which is a compound dataset including HELEN [39], LFPW [40], AFW [41], and IBUG [19] and re-annotated by the 68-point scheme. Same as the previous works [16, 13], we utilize 3148 images for training and 689 images for validation. The full test set can be divided into the common subset and the challenge subset.

COFW [20] is designed to present faces in realistic conditions. Faces in the pictures show large variations in shape and occlusion due to differences in expression, as well as pose. COFW has 1354 images for training and 507 images for testing. Each image is annotated with 29 points consistent with LFPW.

WFLW [16] is a recently proposed large dataset which consists of 7500 training images and 2500 testing images with 98-point annotation. It is one of the most challenging benchmarks for facial landmark detection because of existing extreme disturbance such as occlusion, large pose, illumination, and blur. In order to evaluate the model in variant conditions, the test dataset is categorized as multiple subsets.

JD-landmark [21] is introduced for competition in 2019. Each image is manually annotated with 106-point landmarks. It contains three parts of training, validation, and test sets. The training set provides 11393 images with different poses. The validation and test sets come from web open datasets and cover large variations in pose, expression and occlusion, including 2000 images respectively.

4.2 Evaluation Metrics

To facilitate comparison with previous works, we adopt different evaluation metrics for different benchmarks including Normalized Mean Error (NME), Area Under Curve (AUC) based on Cumulative Error Distribution (CED) curve, and Failure Rate (FR).

NME is widely used in facial landmark detection to evaluate the quality of models. It is calculated by average the Euclidean distance between predicted and ground truth landmarks, then normalized to eliminate the impact caused by the image size inconsistency. NME for each image is defined as:

$$NME = \frac{1}{L} \sum_{k=1}^L \frac{\|p_k - \hat{p}_k\|_2}{d} \quad (8)$$

Where L refers to the number of landmarks. p_k and \hat{p}_k refer to the predicted and ground truth coordinates of the k -th landmark respectively. d is the normalization factor, such as the distance of eye centers (inter-pupil normalization, IPN) or the distance of eye corners (inter-ocular normalization, ION). Specially, For JD-landmark, $d = \sqrt{w_{bbox} \times h_{bbox}}$, where w_{bbox} and h_{bbox} are the width and height of the enclosing rectangle of the ground truth landmarks.

FR is another widely-used metric. It is the percentage of failure samples whose NME is larger than a threshold. As the threshold rises from zero to the target value, we plot the proportion of success (samples with smaller NME than the threshold) as the CED curve. AUC calculates the area under the CED curve.

4.3 Implementation Details

Firstly, for each face bounding box, we extend the short side to the same as the long to avoid *image distortion* since the bounding box is not square in most cases. Then we enlarge the bounding boxes in WFLW and COFW by 25% and 10% respectively, and keep the extended square box size for 300W and JD-landmark. Finally, the images are cropped and resized to 256×256 according to the square boxes from the previous step.

In terms of data augmentation, we use standard methods: 30° max angle of random rotation with the probability of 50%, random scale 75%~125% for all images, random flip horizontally with the probability of 50%. We also try random occlusion in training which improves the performance. We optimize models with Adam [42] for all benchmarks. For bandwidth l , we choose both 3 and 5 for aggregating. For σ in Eq. 2, it is set to 2 on all benchmarks. Our models are trained with MXNet/Gluon 1.5.1 [43] on $4 \times$ TITAN X GPUs. During inference our baseline model costs 28ms per face for the entire pipeline (network forward and post-processing) with batch size = 1 and $1 \times$ TITAN X.

Table 1. NME (%) on 300W Common, Challenging, Full subset

Method	Common	Challenging	Fullset
Inter-pupil Normalization			
RCPR[20]	6.18	17.26	8.35
LBF[44]	4.95	11.98	6.32
TCDCN[10]	4.80	6.80	5.54
RAR[45]	4.12	8.35	4.94
LAB[16]	4.20	7.41	4.92
DCFE[46]	3.83	7.54	4.55
AWing[14]	3.77	6.52	4.31
GV(ResNet50)	3.78	6.98	4.41
GV(HRNet)	3.63	6.51	4.19
Inter-ocular Normalization			
DAN[15]	3.19	5.24	3.59
DCFE[46]	2.76	5.22	3.24
LAB[16]	2.98	5.19	3.49
HRNetV2[32]	2.87	5.15	3.32
AWing[14]	2.72	4.52	3.07
GV(ResNet50)	2.73	4.84	3.14
GV(HRNet)	2.62	4.51	2.99

Table 2. NME (%) and FR (%) on COFW testset

Method	NME	FR _{10%}
Inter-pupil Normalization		
Human[20]	5.60	-
RCPR[20]	8.50	20.0
RAR[45]	6.03	4.14
TCDCN[10]	8.05	20.0
DAC-CSR[47]	6.03	4.73
PCD-CNN[48]	5.77	3.73
DCFE[46]	5.27	-
AWing[14]	4.94	0.99
GV(ResNet50)	4.92	0.99
GV(HRNet)	4.85	0.59
Inter-ocular Normalization		
LAB[47]	3.92	0.39
MobileFAN[13]	3.66	0.59
HRNetV2[32]	3.45	0.19
GV(ResNet50)	3.42	0.59
GV(HRNet)	3.37	0.39

4.4 Evaluation on different benchmarks

We provide two models based on ReNetV2-50 [49] and HRNetV2-W32 [32], represented as GV(ResNet50) and GV(HRNet) respectively, to verify the generality and effectiveness of our approach. Experiments show that our approach achieves state of the art on all benchmarks.³

³ In addition, we analyze the merits of vector supervision in supplemental material.

Table 3. Evaluation on WFLW

Metric	Method	Testset	Pose	Expression	Illumination	Make-up	Occlusion	Blur
			Subset	Subset	Subset	Subset	Subset	Subset
NME(%) (Lower is better)	ESR[27]	11.13	25.88	11.47	10.49	11.05	13.75	12.20
	SDM[50]	10.29	24.10	11.45	9.32	9.38	13.03	11.28
	CFSS[51]	9.07	21.36	10.09	8.30	8.74	11.76	9.96
	DVLN[52]	6.08	11.54	6.78	5.73	5.98	7.33	6.88
	LAB[16]	5.27	10.24	5.51	5.23	5.15	6.79	6.32
	Wing[31]	5.11	8.75	5.36	4.93	5.41	6.37	5.81
	AWing[14]	4.36	7.38	4.58	4.32	4.27	5.19	4.96
	GV(ResNet50)	4.57	7.91	4.85	4.51	4.45	5.50	5.28
	GV(HRNet)	4.33	7.41	4.51	4.24	4.18	5.19	4.93
FR _{10%} (%) (Lower is better)	ESR[27]	35.24	90.18	42.04	30.80	38.84	47.28	41.40
	SDM[50]	29.40	84.36	33.44	26.22	27.67	41.85	35.32
	CFSS[51]	20.56	66.26	23.25	17.34	21.84	32.88	23.67
	DVLN[52]	10.84	46.93	11.15	7.31	11.65	16.30	13.71
	LAB[16]	7.56	28.83	6.37	6.73	7.77	13.72	10.74
	Wing[31]	6.00	22.70	4.78	4.30	7.77	12.50	7.76
	AWing[14]	2.84	13.50	2.23	2.58	2.91	5.98	3.75
	GV(ResNet50)	4.44	19.94	3.18	3.72	3.88	8.83	6.47
	GV(HRNet)	3.52	16.26	2.55	3.30	3.40	6.79	5.05
AUC _{10%} (Higher is better)	ESR[27]	0.2774	0.0177	0.1981	0.2953	0.2485	0.1946	0.2204
	SDM[50]	0.3002	0.0226	0.2293	0.3237	0.3125	0.2060	0.2398
	CFSS[51]	0.3659	0.0632	0.3157	0.3854	0.3691	0.2688	0.3037
	DVLN[52]	0.4551	0.1474	0.3889	0.4743	0.4494	0.3794	0.3973
	LAB[16]	0.5323	0.2345	0.4951	0.5433	0.5394	0.4490	0.4630
	Wing[31]	0.5504	0.3100	0.4959	0.5408	0.5582	0.4885	0.4918
	AWing[14]	0.5719	0.3120	0.5149	0.5777	0.5715	0.5022	0.5120
	GV(ResNet50)	0.5568	0.2769	0.5321	0.5673	0.5640	0.4782	0.4959
	GV(HRNet)	0.5775	0.3166	0.5636	0.5863	0.5881	0.5035	0.5242

Evaluation on 300W. Our proposed method achieves the best performance on 300W as shown in Table 1. GV(ResNet50) achieves 3.14% for ION (4.41% for IPN), and GV(HRNet) achieves 2.99% for ION (4.19% for IPN). Both models get remarkable results.

Evaluation on COFW. Experiments on COFW is displayed in Table 2. Our baseline model GV(ResNet50) provides a better result than the previously best work AWing [14] with 4.92% for IPN. GV(HRNet) achieves newly state-of-the-art performance with 4.85% and reduces FR from 0.99% to 0.39%.

Evaluation on WFLW. It is one of the most challenging benchmarks. We comprehensively compare the results of NME, FR@0.1, and AUC@0.1 with the previous works on the whole test set as well as all subsets in Table 3. Our method achieves evident improvements of NME and AUC in almost all subsets. Notably, ION on the whole test set decreases to 4.33%. We comprehensively compare our method with AWing in the supplement material.

Evaluation on JD-landmark. The results of JD-landmark are exhibited on Table 4. GV(HRNet) achieves comparable results (1.35% for NME with FR 0.10%) to the champion (1.31% for NME with FR 0.10%) with a single model excluding extra competition tricks. It is noteworthy that the champion [21] employed AutoML [53] on architecture search, ensemble model, and a well-designed

data augmentation scheme etc. The results further prove the briefness and effectiveness of our method for facial landmark detection.

Table 4. The leadboard of JD-landmark

Ranking	Team	AUC(@0.08)(%)	NME(%)
No.1	Baidu VIS	84.01	1.31
No.2	USTC-NELSLIP	82.68	1.41
No.3	VIC iron man	82.22	1.42
No.4	CIGIT-ALIBABA	81.53	1.50
No.5	Smiles	81.28	1.50
Baseline	-	73.32	2.16
GV(HRNet)	-	83.34	1.35

Table 5. Ablation study on our method. BBS means Beyond Box Strategy.

Backbone	Original face bounding box		Shrunk face bounding box	
	BPM	NME(%)	BPM	BBS NME(%)
	ResNet50	-	3.42	✓
	✓	3.14	✓	4.10




Fig. 5. CED curves of models in Table 5. SB means Shrinking the bounding box by 10%.

4.5 Ablation study

To demonstrate the effectiveness of Band Pooling Module and Beyond Box strategy, ablation study is implemented on 300W. We choose ResNet50 as the backbone and ION as the test metric.

Effectiveness of BPM. We compare the performance with and without the BPM, that is, our proposed method and the heatmap based method. The result of each model is shown in Table 5, and the BPM greatly improves the model performance from 3.42% to 3.14%. Moreover, we provide CED curves to evaluate the accuracy difference in Fig. 4.4.

Choice of the bandwidth. Experiments show that the choice of bandwidth l has a great impact on the performance of the model in Table 4.5. It is obvious that oversize or undersize bandwidth does harm to the model performance (3.94% when $l=1$ and 3.38% when $l=7$). By aggregating different bandwidth, we get better localization results. Specifically, we obtain the best ION 3.14% when fusing the vectors from bandwidth $l=3$ and $l=5$.

Choice of θ in Eq. 2. We conduct a series of experiments on 300W with θ range from 0 to 4. Note the backbone is ResNet50. Experiments show that the appropriate θ gets better results in Table. 7

Improvements from the Beyond Box Strategy. In Sec. 3.4, we introduce a strategy to predict the landmarks out of the face bounding box. Now a plain experiments is conducted to illustrate the effectiveness of it. We shrink the ground truth bounding box of 300W fullset by 10% in each dimension to make some points fall out of the image. In this case, the prediction results turn worse (ION raised from 3.14% to 4.57%, shown in Table 5) since the certain parts of face are invisible. We compare different choices for τ and ρ , as illustrated in Table 6. When we pick $\tau=2.0$ and $\rho=2.0$, ION decreases from 4.57% to 4.10%.

Table 6. Different choices for τ and σ in Beyond Box Strategy.

τ	σ	1.0	1.5	2.0	2.5	3.0
1.0		4.42	4.35	4.29	4.26	4.24
1.5		4.33	4.23	4.15	4.12	4.14
2.0		4.28	4.17	4.10	4.12	4.22
2.5		4.26	4.15	4.13	4.23	4.43


Table 7. Evaluation on different values of θ .

θ	0	1	2	3	4
ION(%)	3.32	3.23	3.21	3.14	3.25

Table 8. Comparison among different bandwidth. Diagonal elements indicate that vector generated by single bandwidth BPM, as bandwidth $l=1$ or $l=3$. The others attained by aggregating different bandwidth where $\alpha=0.5$ in Eq. 5.

bandwidth	$l=1$	$l=3$	$l=5$	$l=7$
$l=1$	3.94	3.67	3.61	3.74
$l=3$	-	3.31	3.14	3.22
$l=5$	-	-	3.21	3.28
$l=7$	-	-	-	3.38

4.6 Efficiency Analysis

**Fig. 6.** The training convergence plots of different method. Note the logarithm of the y-axis here.

The efficiency of our method is examined in both convergence rate and time cost as follows.

The comparison of convergence rate. We argue that the vector label speeds up model convergence and will experimentally demonstrate that. In addition, a solution for imbalance has been proposed in AWing [14]. By using a weight map, AWing applied different weights on hard examples(1.0) and easy examples(0.1) to alleviate the imbalance of heatmap. Experiments on 300W show the impact of weight map and vector label on convergence. We use the same

learning rate and batch size for a fair comparison. The metric is ION. The results show both weight map and vector label facilitate the convergence and our method is better.

Table 9. Comparison of the time cost.

method	output tensor size	net-forward (ms)	post-processing (ms)	transmission (ms)
Heatmap	(68, 64, 64)	6.129	2.039	0.311
Gaussian Vector	(68, 64, 2)	6.158	0.062	0.048

Time cost analysis of Heatmap based method and Gaussian Vector. In embedded systems, the network runs on neural net accelerators (e.g. TPU) and post-processing on CPU (e.g. ARM). The former accelerates common ops (Conv/Pool excluding argmax) significantly. So BPM increases little computation complexity and time. However, post-processing is the bottleneck since CPU has to process many other tasks synchronously. For the limitation of the practical system (MT6735 with a NN accelerator), we use MobileNetV3 as the backbone to test the time cost of heatmap based method and ours. As Table. 9, our method saves 26% time of the whole process in this case.

5 Conclusion

In this paper, we present a simple yet effective solution for facial landmark detection. We provide novel vector label for supervision, Band Pooling Module to convert output heatmap into vectors, as well as Beyond Box Strategy for predicting landmarks out of the input. By evaluating on different benchmarks, we demonstrate the efficiency of our proposed method. Also, our method and its components are validated via comprehensive ablation study.

References

1. Sun, Y., Wang, X., Tang, X.: Deep learning face representation from predicting 10,000 classes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2014) 1891–1898
2. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Sphereface: Deep hypersphere embedding for face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2017) 212–220
3. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 4690–4699
4. Thies, J., Zollhofer, M., Stamminger, M., Theobalt, C., Nießner, M.: Face2face: Real-time face capture and reenactment of rgb videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2016) 2387–2395

5. Dou, P., Shah, S.K., Kakadiaris, I.A.: End-to-end 3d face reconstruction with deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2017) 5908–5917
6. Roth, J., Tong, Y., Liu, X.: Unconstrained 3d face reconstruction. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2015) 2606–2615
7. Feng, Y., Wu, F., Shao, X., Wang, Y., Zhou, X.: Joint 3d face reconstruction and dense alignment with position map regression network. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 534–551
8. Zhu, X., Liu, X., Lei, Z., Li, S.Z.: Face alignment in full pose range: A 3d total solution. *IEEE transactions on pattern analysis and machine intelligence* **41** (2017) 78–92
9. Sun, Y., Wang, X., Tang, X.: Deep convolutional network cascade for facial point detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2013) 3476–3483
10. Zhang, Z., Luo, P., Loy, C.C., Tang, X.: Facial landmark detection by deep multi-task learning. In: European conference on computer vision, Springer (2014) 94–108
11. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2015) 3431–3440
12. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention, Springer (2015) 234–241
13. Zhao, Y., Liu, Y., Shen, C., Gao, Y., Xiong, S.: Mobilefan: Transferring deep hidden representation for face alignment. *Pattern Recognition* **100** (2020) 107114
14. Wang, X., Bo, L., Fuxin, L.: Adaptive wing loss for robust face alignment via heatmap regression. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 6971–6981
15. Kowalski, M., Naruniec, J., Trzcinski, T.: Deep alignment network: A convolutional neural network for robust face alignment. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. (2017) 88–97
16. Wu, W., Qian, C., Yang, S., Wang, Q., Cai, Y., Zhou, Q.: Look at boundary: A boundary-aware face alignment algorithm. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2018) 2129–2138
17. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: European conference on computer vision, Springer (2016) 483–499
18. Dong, X., Yu, S.I., Weng, X., Wei, S.E., Yang, Y., Sheikh, Y.: Supervision-by-registration: An unsupervised approach to improve the precision of facial landmark detectors. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 360–368
19. Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., Pantic, M.: 300 faces in-the-wild challenge: The first facial landmark localization challenge. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. (2013) 397–403
20. Burgos-Artizzu, X.P., Perona, P., Dollár, P.: Robust face landmark estimation under occlusion. In: Proceedings of the IEEE international conference on computer vision. (2013) 1513–1520
21. Liu, Y., Shen, H., Si, Y., Wang, X., Zhu, X., Shi, H., Hong, Z., Guo, H., Guo, Z., Chen, Y., et al.: Grand challenge of 106-point facial landmark localization. In: 2019 IEEE International Conference on Multimedia & Expo Workshops (ICMEW), IEEE (2019) 613–616

22. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. *IEEE Transactions on pattern analysis and machine intelligence* **23** (2001) 681–685
23. Kahraman, F., Gokmen, M., Darkner, S., Larsen, R.: An active illumination and appearance (aia) model for face alignment. In: 2007 IEEE Conference on Computer Vision and Pattern Recognition, IEEE (2007) 1–7
24. Matthews, I., Baker, S.: Active appearance models revisited. *International journal of computer vision* **60** (2004) 135–164
25. Milborrow, S., Nicolls, F.: Locating facial features with an extended active shape model. In: European conference on computer vision, Springer (2008) 504–513
26. Cristinacce, D., Cootes, T.F.: Feature detection and tracking with constrained local models. In: *Bmvc*. Volume 1., Citeseer (2006) 3
27. Cao, X., Wei, Y., Wen, F., Sun, J.: Face alignment by explicit shape regression. *International Journal of Computer Vision* **107** (2014) 177–190
28. Chen, D., Ren, S., Wei, Y., Cao, X., Sun, J.: Joint cascade face detection and alignment. In: European conference on computer vision, Springer (2014) 109–122
29. Xiong, X., De la Torre, F.: Global supervised descent method. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2015) 2664–2673
30. Trigeorgis, G., Snape, P., Nicolaou, M.A., Antonakos, E., Zafeiriou, S.: Mnemonic descent method: A recurrent process applied for end-to-end face alignment. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2016) 4177–4187
31. Feng, Z.H., Kittler, J., Awais, M., Huber, P., Wu, X.J.: Wing loss for robust facial landmark localisation with convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 2235–2245
32. Sun, K., Zhao, Y., Jiang, B., Cheng, T., Xiao, B., Liu, D., Mu, Y., Wang, X., Liu, W., Wang, J.: High-resolution representations for labeling pixels and regions. arXiv preprint arXiv:1904.04514 (2019)
33. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 5693–5703
34. Liu, R., Lehman, J., Molino, P., Such, F.P., Frank, E., Sergeev, A., Yosinski, J.: An intriguing failing of convolutional neural networks and the coordconv solution. In: *Advances in Neural Information Processing Systems*. (2018) 9605–9616
35. Sun, X., Xiao, B., Wei, F., Liang, S., Wei, Y.: Integral human pose regression. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 529–545
36. Nibali, A., He, Z., Morgan, S., Prendergast, L.: Numerical coordinate regression with convolutional neural networks. arXiv preprint arXiv:1801.07372 (2018)
37. Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research* **17** (2016) 1334–1373
38. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2017) 2117–2125
39. Le, V., Brandt, J., Lin, Z., Bourdev, L., Huang, T.S.: Interactive facial feature localization. In: European conference on computer vision, Springer (2012) 679–692
40. Belhumeur, P.N., Jacobs, D.W., Kriegman, D.J., Kumar, N.: Localizing parts of faces using a consensus of exemplars. *IEEE transactions on pattern analysis and machine intelligence* **35** (2013) 2930–2940

41. Zhu, X., Ramanan, D.: Face detection, pose estimation, and landmark localization in the wild. In: 2012 IEEE conference on computer vision and pattern recognition, IEEE (2012) 2879–2886
42. Kingma, D.P., Ba, J.L.: (Adam: A method for stochastic optimization)
43. Chen, T., Li, M., Li, Y., Lin, M., Wang, N., Wang, M., Xiao, T., Xu, B., Zhang, C., Zhang, Z.: Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. arXiv preprint arXiv:1512.01274 (2015)
44. Ren, S., Cao, X., Wei, Y., Sun, J.: Face alignment at 3000 fps via regressing local binary features. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2014) 1685–1692
45. Xiao, S., Feng, J., Xing, J., Lai, H., Yan, S., Kassim, A.: Robust facial landmark detection via recurrent attentive-refinement networks. In: European conference on computer vision, Springer (2016) 57–72
46. Valle, R., Buenaposada, J.M., Valdes, A., Baumela, L.: A deeply-initialized coarse-to-fine ensemble of regression trees for face alignment. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 585–601
47. Feng, Z.H., Kittler, J., Christmas, W., Huber, P., Wu, X.J.: Dynamic attention-controlled cascaded shape regression exploiting training data augmentation and fuzzy-set sample weighting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2017) 2481–2490
48. Kumar, A., Chellappa, R.: Disentangling 3d pose in a dendritic cnn for unconstrained 2d face alignment. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 430–439
49. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: European conference on computer vision, Springer (2016) 630–645
50. Xiong, X., De la Torre, F.: Supervised descent method and its applications to face alignment. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2013) 532–539
51. Zhu, S., Li, C., Change Loy, C., Tang, X.: Face alignment by coarse-to-fine shape searching. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2015) 4998–5006
52. Wu, W., Yang, S.: Leveraging intra and inter-dataset variations for robust face alignment. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. (2017) 150–159
53. He, Y., Lin, J., Liu, Z., Wang, H., Li, L.J., Han, S.: Amc: Automl for model compression and acceleration on mobile devices. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 784–800