

PropagationNet: Propagate Points to Curve to Learn Structure Information

Xiehe Huang^{1,2}Weihong Deng¹Haifeng Shen²Xiubao Zhang²Jieping Ye²

¹Pattern Recognition & Intelligent System Laboratory, School of Information and Communication Engineering, Beijing University of Posts and Telecommunications

{xiehe.huang, whdeng}@bupt.edu.cn

²AI Labs, DiDi Chuxing

{huangxiehe.i, shenhaifeng, zhangxiubao, yejieping}@didiglobal.com

Abstract

Deep learning technique has dramatically boosted the performance of face alignment algorithms. However, due to large variability and lack of samples, the alignment problem in unconstrained situations, e.g. large head poses, exaggerated expression, and uneven illumination, is still largely unsolved. In this paper, we explore the instincts and reasons behind our two proposals, i.e. Propagation Module and Focal Wing Loss, to tackle the problem. Concretely, we present a novel structure-infused face alignment algorithm based on heatmap regression via propagating landmark heatmaps to boundary heatmaps, which provide structure information for further attention map generation. Moreover, we propose a Focal Wing Loss for mining and emphasizing the difficult samples under in-the-wild condition. In addition, we adopt methods like CoordConv and Anti-aliased CNN from other fields that address the shift-variance problem of CNN for face alignment. When implementing extensive experiments on different benchmarks, i.e. WFLW, 300W, and COFW, our method outperforms state-of-the-arts by a significant margin. Our proposed approach achieves 4.05% mean error on WFLW, 2.93% mean error on 300W full-set, and 3.71% mean error on COFW.

1. Introduction

Face alignment, aimed at localizing facial landmark, plays an essential role in many face analysis applications, e.g. face verification and recognition [23], face morphing [10], expression recognition [13], and 3D face reconstruction [7]. Recent years have witnessed the constant emergence of fancy face alignment algorithms with considerable performance on various datasets. Yet face alignment in unconstrained situations, e.g. large head pose, exaggerated expression, and uneven illumination, has plagued researchers

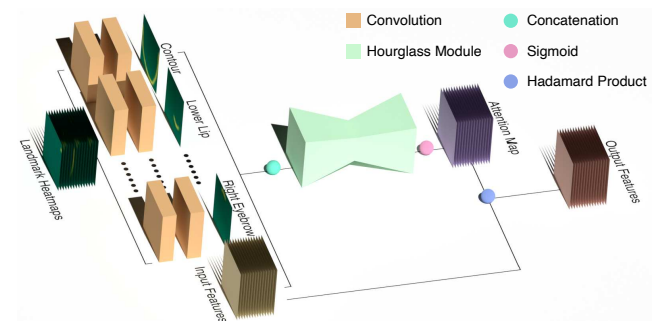


Figure 1. Building blocks of our propagation module. Landmark heatmaps are input to multiple convolutional operations, then concatenated with the output feature maps of last hourglass module, together processed by a two-stage hourglass module, and finally normalized by a sigmoid layer to form an attention map that is imposed on the feature maps.

over the years. Among many other factors, we attribute the mentioned challenges to the disability of CNN to learn facial structure information: if a CNN is enabled to extract the structure of a face in an image, then it can predict facial landmarks more accurately since even the occluded parts of a face, for instance, can be inferred through the shape of the face. This is also the intention of ASM [4]’s designers.

What exactly is structural information? In our work, we deem it to be the statistical mean of landmark coordinates. Perhaps with high variance (such as different head poses), landmark coordinates are still subject to some distribution due to the relative non-transformability of a face shape. In order for a CNN to learn the information, we represent it as facial boundary in this paper (see Fig. 2) following Wu *et al.* [25]. A facial boundary could be the jawline, or the outer contour of a face. Or it could be the edge surrounding a mouth. These boundaries are commonly annotated with a series of points by available datasets due to their difficulty with modeling a line.

In this paper, we propose and implement 3 creative ideas to learn the structural information, *i.e.* Propagation Module, Focal Wing Loss, and Multi-view Hourglass.




Figure 2. Landmarks are connected to make several boundaries.

Wu *et al.* produce facial boundaries out of a separate GAN (Generative Adversarial Network) generator. Specifically, Wu *et al.* connected the landmarks to make a blurred line and specified it as ground truth for future training. Unlike their fashion of using an independent CNN to generate facial boundaries, we devise a **Propagation Module** to do this job and incorporate it within our network architecture, in an attempt to substitute a deeper and larger CNN with a computation-efficient Propagation Module. In addition to the module’s embeddability in a CNN, what is more important is that boundary heatmaps are naturally connected with landmark heatmaps. For this reason, it is intuitive for us to use a series of convolution operations to model the connection and propagate a certain number of landmarks (points) to a boundary (a curve). Hence we term the module as Propagation Module.




Figure 3. Fractions of images under extreme conditions for different sets. Extreme conditions include large head pose, exaggerated expression, non-uniform illumination, unrecognizable make-up, object occlusion, and blurry shot.

Data imbalance is a common issue in many fields of AI and so is the case in face alignment community. The structures of a face varies under in-the-wild conditions. For example, the jawline is less widely open when a face is in profile position than when the face is shown frontal. How-

ever, the ratio of data under these two conditions is not actually anywhere near 1 : 1, where the number of frontal images is the same as that of profile ones. As illustrated in Fig. 3, the fractions of images under extreme conditions are rather low, all under 30% across both train set and test set. On the other hand, the fractions on train set deviates largely from those on test set, which means a learned feature adapted to train set can misguide the neural network to make a wrong prediction. This potential non-universal feature, therefore, necessitates a better design of loss function. Based on the primitive AWing [24], we propose a **Focal Wing Loss**, which dynamically adjusts the penalty for incorrect prediction and gears the loss weight (thus the learning rate) for each sample in every batch during training. This indicates that our training process pay attention evenly to both hard-to-learn structures and easy-to-learn ones, so we refer to the loss function as Focal Wing Loss.

Modern-day convolutional neural networks are usually believed to be shift invariant, and so is the stacked hourglass used in our work. Nevertheless, researchers have come to realize the underlying shift variance brought by the introduction of pooling layer, *e.g.* max pooling and average pooling. To resolve this translation-variance, Zhang [31] provided the solution of Anti-aliased CNN, which emulates the traditional signal processing method of anti-aliasing and apply it before every downsampling operation, such as pooling and strided convolution. In our task, we cannot afford to lose structural information when applying pooling layer, so we incorporate Anti-aliased CNN in a special hourglass and name it **Multi-view Hourglass**.

In conclusion, our main contribution include:

- creating a novel propagation module to seamlessly connect landmark heatmaps with boundary heatmaps, a module that can be naturally built into stacked hourglass model.
- devising a loss function termed as Focal Wing Loss to dynamically assign loss weight to a particular sample and tackle data imbalance.
- introducing Anti-aliased CNN from other fields and integrating them within our Multi-View Hourglass Module to add shift equivariance and coordinate information to our network.
- implementing extensive experiments on various datasets as well as ablation studies about the mentioned methods.

2. Related Work

Recently, interest of face alignment community has been largely centering on two mainstream approaches, *i.e.* coordinate regression and heatmap regression, with various



Figure 4. Sample results on WFLW testset. Each column comes from a subset of WFLW, including large head pose, expression, illumination, make-up, occlusion, and blur.

model designs. **Heatmap regression models**, based on fully convolution network (FCN), output a heatmap for each landmark and try to maintain structure information throughout the whole network, therefore, to some extent, dwarfing coordinate regression models in their state-of-the-art performance. MHM [5], one of those heatmap regression models, implements face detection and face alignment consecutively and leverages stacked hourglass model to predict landmark heatmaps. AWing [24], yet another heatmap regression model, modifies L1 loss to derive so-called adaptive wing loss and proves its superiority in CNN-based facial landmark localization. What is common among these 2 models is their adoption of stacked hourglass network. Stacked hourglass model stands out among all FCNs in the field of landmark detection since its debut in [17] for human pose estimation. Its prevalence can be attributed to its repeated bottom-up, top-down processing that allows for capturing information across all scales of an input image.

First raised by Wu *et al.* [25] and later popularized by such researchers as Wang *et al.* [24], **facial boundary** identifies geometry structure of human face and therefore can infuse a network with prior knowledge, be it used for attention mechanism (as in the case of LAB [25]) or for generation of boundary coordinate map (as in the case of AWing [24]). In the former scenario, LAB first utilizes a stacked hourglass model to generate facial boundary map and then incorporates the boundary map to a regression network via feature map fusion. In the latter scenario, AWing encodes boundary prediction as a mask on x-y coordinates and consequently produces two additional feature maps for follow-on convolution. Different from both of them, we generate the boundary heatmap with only several convolution operations instead of a complicated CNN.

Attention mechanism has enjoyed great popularity in computer vision because the extra “attention” brought by it can guide a CNN to learn worthwhile features and focus on them. In our work, we want our model to focus more on the area of boundary so it can inference landmark more



Figure 5. Sample results on 300W and COFW testsets. Each row demonstrates samples from each dataset.

accurately based on the position of boundary. Unlike LAB [25]’ way of using a ResNet-block-based hourglass to generate attention map, we adopt a multi-view hourglass which can maintain structural information throughout the whole process. Specifically, we incorporate Hierarchical, Parallel & Multi-Scale block [1] to add more sizes of receptive fields and Anti-aliased CNN [31] to improve shift invariance. Larger size of receptive field means that our model can “behold” the whole structure of a face, whereas shift invariance means that our model can still predict boundary heatmaps correctly even if the correspondent face image is shifted a little bit. Moreover, we do not have to downsample boundary heatmaps every time they are fed into the next hourglass whereas LAB does. This is because we do not want to lose boundary information via downsampling.

CNN-based localization models have long been trained with take-away **loss functions**, *e.g.* L1, L2, and smooth L1. These loss functions are indeed useful in common scenarios. Feng *et al.* [8], however, contends that L2 is sensitive to outliers and therefore dwarfed by L1. In order to make their model pay more attention to small and medium range errors, they modify the L1 loss to create Wing Loss which is more effective in landmark coordinates regression models. Based on Wing Loss, Wang *et al.* [24] further bring in adaptiveness to the loss function because they believe “influence” (a concept lent from robust statistics) should be proportional to gradient and balance all errors. The Adaptive Wing Loss they created is proven to be more effective in heatmap regression models.

3. Approach

Based on stacked HG design from Bulat *et al.* [1], our model further integrates it with Propagation Module, anti-aliased block, and CoordConv. Each hourglass outputs feature maps for the following hourglass and landmark heatmaps supervised with ground truth labels. Next in line is a Propagation Module, which generates boundary heatmaps and outputs the feature maps for the follow-on hourglass. This overall process is visualized in Fig. 6.




Figure 6. Overview of our PropagationNet architecture. RGB images are first processed through a series of basic feature extractors, and then fed into several hourglass modules followed by a relation block that outputs the boundary heatmaps.

Table 1. Evaluation of PropagationNet and other state-of-the-arts on COFW testset.

Method	NME	FR _{10%}
Human [2]	5.60	-
RCPR [2]	8.50	20.00
TCDCN [32]	8.05	-
DAC-CSR [9]	6.03	4.73
Wing [8]	5.44	3.75
Awing [24]	4.94	0.99
LAB [25]	3.92	0.39
PropNet(Ours)	3.71	0.20

3.1. Landmark-Boundary Propagation Module

Inspired by attention mechanism, the Landmark-Boundary Propagation Module aims to force the network to pay more “attention” to the boundary area in order to make more accurate prediction of landmark heatmap. To achieve this goal, it first employs an array of convolution operations to transform landmark heatmaps to boundary heatmaps. These operations basically attempt to learn how to translate landmark heatmaps and combine the trajectories to make boundary heatmaps. Every boundary heatmap is generated via a set of several 7×7 convolution operations. Then it concatenates boundary heatmaps and the feature maps from its anterior hourglass module and feeds them into a two-stage hourglass module in order to yield the attention map. Finally, it enhances the feature maps with

attention map and transports those feature maps to its posterior hourglass. This process is visualized in Fig. 1.

During training, the generation of boundary heatmaps is supervised by ground truth heatmaps. As for how to produce ground truth heatmaps, we simply link landmarks together with straight lines and apply Gaussian blurring filter. Each boundary has its semantic meanings. As depicted in Fig. 2, landmarks lying on the jawline are connected to make contour boundary, those denoting the lower lip are connected to make another boundary, and so forth. In total, we obtain $M = 15$ boundary heatmaps.

3.2. Focal Wing Loss

Adaptive wing loss [24] is derived from wing loss [8] and is basically a variant of smooth L1 loss except that the smooth quadratic curve is replaced by a logarithmic curve. It is piecewise-defined as Eq. (1), where $A = \omega(\alpha - y)(\theta/\epsilon)^{\alpha-y-1}/(1 + (\theta/\epsilon)^{\alpha-y})/\epsilon$ and $\Omega = \theta A - \omega \ln(1 + (\theta/\epsilon)^{\alpha-y})$ are defined to make the loss function continuous and smooth at $|y - \hat{y}| = \theta$ and ω , θ , α , and ϵ are hyper-parameters that affect the none-L1 range and the gradient between it.

$$AWing(x) = \begin{cases} \omega \ln(1 + |\frac{y-\hat{y}}{\epsilon}|^{\alpha-y}), & |y - \hat{y}| < \theta \\ A|y - \hat{y}| - \Omega, & |y - \hat{y}| \geq \theta \end{cases} \quad (1)$$

Table 2. Evaluation of PropagationNet and other state-of-the-arts on 300W testset.

Method	Common Subset	Challenging Subset	Fullset
Inter-pupil Normalization			
CFAN [30]	5.50	16.78	7.69
SDM [29]	5.57	15.40	7.50
LBF [18]	4.95	11.98	6.32
CFSS [34]	4.73	9.98	5.76
TCDCN [33]	4.80	8.60	5.54
MDM [21]	4.83	10.14	5.88
RAR [28]	4.12	8.35	4.94
DVLN [27]	3.94	7.62	4.66
TSR [15]	4.36	7.56	4.99
DSRN [16]	4.12	9.68	5.21
LAB [25]	4.20	7.41	4.92
RCN ⁺ (L+ELT) [11]	4.20	7.78	4.90
DCFE [22]	3.83	7.54	4.55
Wing [8]	3.27	7.18	4.04
AWing [24]	3.77	6.52	4.31
PropNet(Ours)	3.70	5.75	4.10
Inter-ocular Normalization			
PCD-CNN [12]	3.67	7.62	4.44
CPM+SBR [6]	3.28	7.58	4.10
SAN [6]	3.34	6.60	3.98
LAB [25]	2.98	5.19	3.49
DU-Net [20]	2.90	5.15	3.35
AWing [24]	2.72	4.52	3.07
PropNet(Ours)	2.67	3.99	2.93

In order to address data imbalance, we introduce a factor named **Focal Factor**. For a class c and a sample n , it is mathematically defined as:

$$\sigma_n^{(c)} = \begin{cases} 1, & \text{if } \sum_{n=1}^N s_n^{(c)} = 0 \\ \frac{N}{\sum_{n=1}^N s_n^{(c)}}, & \text{otherwise} \end{cases} \quad (2)$$

where $s_n^{(c)}$ is binary number: when $s_n^{(c)} = 0$, the sample n does not belong to class c ; when $s_n^{(c)} = 1$, the sample n belongs to class c . In this paper, a sample that belongs to a certain class means the sample has the c th attribute, such as large head pose, exaggerated expression, *etc.* For WFLW dataset, these attributes are labeled in annotation file, while for COFW and 300W we hand-label these attributes by ourselves and use them when training. Also note that the Focal Factor is defined batch-wise, which means it is fluctuating during the training process and again it dynamically adjusts loss weight on every sample in a batch. Furthermore, the weight loss is a sum of all focal factors from different class, as can be seen in the following definition (3). This indicates that we intend to balance the data across all classes because a facial image can be subjected to multiple extreme conditions, *e.g.* a blurry facial image with large head pose.

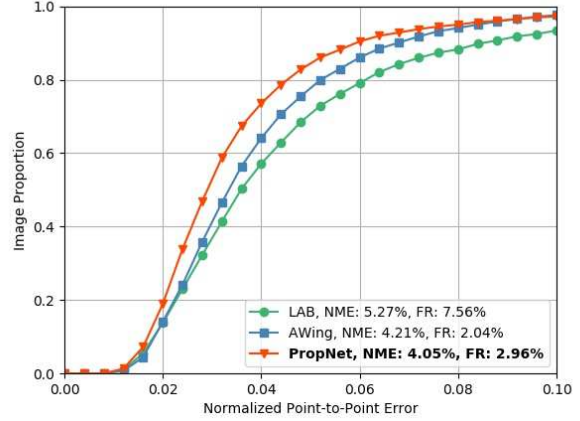


Figure 7. CED for WFLW testset. NME and FR_{10%} are reported at the legend for comparison. We compare our methods with other state-of-the-arts with source codes available, including LAB [25] and AWing [24].

As a result, we have the loss of landmark as:

$$L_{lm} = \frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C \sigma_n^{(c)} \sum_{k=1}^K AWing(y_n^{(k)}, \hat{y}_n^{(k)}) \quad (3)$$

where N, C, K respectively denote batch size, number of classes (subsets) and number of coordinates. In our case, $C = 6$ for 6 attributes: head pose, expression, illumination, make-up, occlusion, and blurring; $K = 196$ for 98 landmarks that are considered in WFLW dataset. $y_n^{(k)}$ and $\hat{y}_n^{(k)}$ separately stand for the ground truth heatmap of sample n , landmark k and the corresponding predicted heatmap.

Similarly, we define the loss function for boundary heatmap prediction as:

$$L_{bd} = \frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C \sigma_n^{(c)} \sum_{m=1}^M AWing(z_n^{(m)}, \hat{z}_n^{(m)}) \quad (4)$$

where M denotes the total number of boundaries. $z_n^{(m)}$ and $\hat{z}_n^{(m)}$ are respectively the ground truth boundary heatmap of sample n , boundary m and the corresponding predicted boundary heatmap.

Finally we obtain the holistic loss function as:

$$Loss = L_{lm} + \beta \cdot L_{bd} \quad (5)$$

where β is a hyper-parameter for balancing two tasks.

3.3. Multi-view Hourglass Module

Different from traditional hourglass networks using bottleneck block as their building blocks, we adopt the **hierarchical, parallel and multi-scale residual architecture** proposed by Bulat *et al.* [1]. We think the architecture is

Table 3. Evaluation of PropagationNet and other state-of-the-arts on WFLW testset and its subsets.

Metric	Method	Testset	Pose Subset	Expression Subset	Illumination Subset	Make-up Subset	Occlusion Subset	Blur Subset
NME (%)	ESR [3]	11.13	25.88	11.47	10.49	11.05	13.75	12.20
	SDM [29]	10.29	24.10	11.45	9.32	9.38	13.03	11.28
	CFSS [34]	9.07	21.36	10.09	8.30	8.74	11.76	9.96
	DVLN [26]	6.08	11.54	6.78	5.73	5.98	7.33	6.88
	LAB [25]	5.27	10.24	5.51	5.23	5.15	6.79	6.12
	Wing [8]	5.11	8.75	5.36	4.93	5.41	6.37	5.81
	PropNet(Ours)	4.05	6.92	3.87	4.07	3.76	4.58	4.36
FR _{10%} (%)	ESR [3]	35.24	90.18	42.04	30.80	38.84	47.28	41.40
	SDM [29]	29.40	84.36	33.44	26.22	27.67	41.85	35.32
	CFSS [34]	20.56	66.26	23.25	17.34	21.84	32.88	23.67
	DVLN [26]	10.84	46.93	11.15	7.31	11.65	16.30	13.71
	LAB [25]	7.56	28.83	6.37	6.73	7.77	13.72	10.74
	Wing [8]	6.00	22.70	4.78	4.30	7.77	12.50	7.76
	PropNet(Ours)	2.96	12.58	2.55	2.44	1.46	5.16	3.75
AUC _{10%}	ESR [3]	0.2774	0.0177	0.1981	0.2953	0.2485	0.1946	0.2204
	SDM [29]	0.3002	0.0226	0.2293	0.3237	0.3125	0.2060	0.2398
	CFSS [34]	0.3659	0.0632	0.3157	0.3854	0.3691	0.2688	0.3037
	DVLN [26]	0.4551	0.1474	0.3889	0.4743	0.4494	0.3794	0.3973
	LAB [25]	0.5323	0.2345	0.4951	0.5433	0.5394	0.4490	0.4630
	Wing [8]	0.5504	0.3100	0.4959	0.5408	0.5582	0.4885	0.4918
	PropNet(Ours)	0.6158	0.3823	0.6281	0.6164	0.6389	0.5721	0.5836

beneficial to landmark localization due to its multiple receptive fields and the various scale of images those fields can bring. That means we have features describing the larger structure a human face as well as details about each boundary. Hence we name hourglass module as Multi-view Hourglass module and the architecture itself as Multi-view Block, as seen in Fig. 6.

On the other hands, we implement **anti-aliased CNN** in place of pooling layer used in traditional hourglass networks. One reason for this is to maintain shift equality in our network while another reason is that we do not want to lose some detail information caused by pooling layer or strided convolution.

3.4. Anti-aliased CNN and CoordConv

CoordConv [14] is applied in our work to learn either complete translation invariance or ranging degrees of translation dependence. Anti-aliased CNN [31] is also used to replace pooling layer or strided convolution in our work to preserve shift equality. We term it anti-aliased block, seen in Fig. 6.

4. Experiments

4.1. Evaluation Metrics

Normalized Mean Error (NME) is a widely used metric to evaluate the performance of a facial landmark localization algorithm. Pixel-wise absolute distance is normal-

ized over a distance that takes face size into consideration. Error of each keypoint is calculated this way and then averaged to get the final result. See Equation (6).

$$NME(P, \hat{P}) = \frac{1}{L} \sum_{l=1}^L \frac{\|p_l - \hat{p}_l\|_2}{d} \quad (6)$$

where P, \hat{P} are respectively the ground truth coordinates of all points and the predicted ones for a face image, L is the total number of keypoints, and p_l, \hat{p}_l are both 2-dimensional vector presenting the x-y coordinates of the i-th keypoint. Especially, d is the mentioned normalization factor, be it inter-pupil distance or inter-ocular distance. The latter could be distance between the inner corners of eyes (not commonly used) or the outer corner of eyes which we use in our evaluation. For 300W dataset, both factors is applied; for COFW dataset, we use only inter-pupil distance; for WFLW dataset, inter-ocular distance is adopted.

Failure Rate (FR) provides another insight into a face alignment algorithm design. The NME computed on every image is thresholded at, for example, 8% or 10%. If the NME of an image is larger than the threshold, the sample is considered to be a failure. We derive the FR from the rate of failures in a testset.

Area under the Curve (AUC) is yet another metric popular among designers of facial landmark detection algorithm. Basically, we derive it from the CED curve: by plotting the curve from zero to the threshold for FR, we have

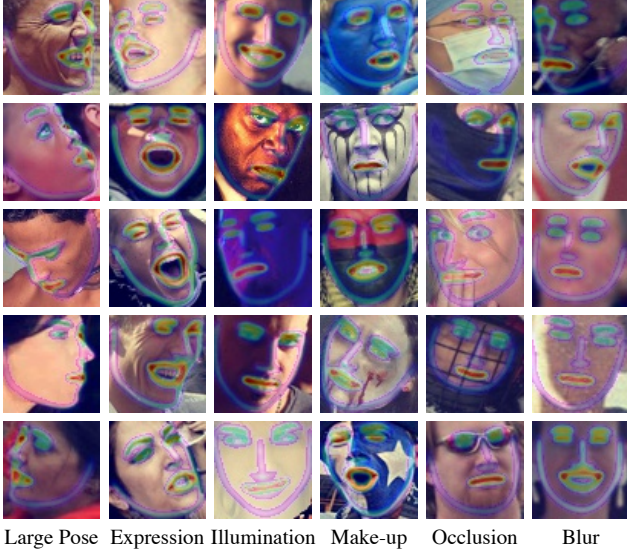


Figure 8. Image samples from WFLW testset imposed with generated boundary heatmap. Each column comes from different subset.

a non-negative curve under which the area is calculated to be AUC. An AUC increment implies that more samples in testset are well predicted.

4.2. Datasets

We perform training and testing of our model on 3 datasets: the challenging dataset WFLW [25] which consists of 10,000 faces (7,500 for training and 2,500 for testing) with 98 fully manual annotated landmarks and is probably hitherto the largest open dataset for face alignment with a large number of keypoints annotation; COFW dataset [2] which contains 1852 face images (1,345 for training and 507 for testing) with 29 annotated landmarks and features heavy occlusions and large shape variations; 300W [19] which is the first facial landmark localization benchmark and for its testset, includes 554 samples for common subset and 135 images for challenging subset.

On **WFLW** dataset, we achieve state-of-the-art performance. See Table 3. Compared with the second leading algorithm, *i.e.* Wing, we improve 3 metrics by about 20% for NME, around 51% for $FR_{10\%}$, and roughly 12% for $AUC_{10\%}$. More importantly, we outperform all the others algorithms on all subsets, which means our model remains robust against different in-the-wild conditions. Special attention should be paid to pose and make-up subsets, where we made a significant improvement. Some sample from the testset can be viewed in Fig. 4. Besides, we also draw Cumulative Error Distribution (CED) curves (see Fig. 7) for algorithms with available released code, including LAB [25] and AWing [24]. From the figure, it is obvious that our PropNet curve is higher than the rest two between 0.02 and

0.08, which means we are able to predict facial landmarks of a larger fraction of images in WFLW testset.

On **COFW** dataset, our algorithm outperforms the other models. See Table 1. As we all know that COFW is well-known for heavy occlusion and wide range of head pose variation, our leading NME and $FR_{10\%}$ prove that our algorithm stays robust again those extreme situations. This also implies that Propagation Module is able to infuse the network with geometrical structure of a human face because only this structure remains in those worst case scenarios. We can see this in Fig. 5.

On **300W** dataset, our model performs excellently on both two subsets and the fullset when compared with other algorithms using inter-ocular normalization factor, as the upper part of Table 2 demonstrates. In terms of metrics vis-à-vis inter-pupil normalization, we have similar metrics with the other leading algorithm on the common set and the fullset, but beat them on the challenging set. This suggests that our algorithm can make plausible predictions even in deplorable situations. This is obviously demonstrated in Fig. 5. A potential reason for relative higher NME with inter-pupil normalization is that 300W annotates some out-of-bounding-box facial parts, *e.g.* chin, with a flat line along the bounding box rather than sticks to the truth. Therefore, this annotation style makes it difficult for our model to learn facial structure.

4.3. Implementation Details

Every input image is cropped and resized to 256×256 and the output feature map of every hourglass module is 64×64 . In our network architecture, we adopt four stacked hourglass modules. During training process, we use Adam to optimize our neural network with the initial learning rate set as 1×10^{-4} . Besides, data augmentation is implemented at training time: random rotation ($\pm 30^\circ$), random scaling ($\pm 15\%$), random crop ($\pm 25px$), and random horizontal flip (50%). At test time, we adopt the same strategy of slightly modifying the predicted result as [17], that is, the coordinate of highest heatmap response is shifted a quarter pixel away to the coordinate of second highest response next to it. Moreover, we empirically set the hyper-parameters in our loss function to be: $\alpha = 2.1, \beta = 0.5, \omega = 14, \epsilon = 1.0, \theta = 0.5$.

4.4. Ablation study

Our algorithm is comprised of several pivotal designs, *i.e.* Propagation Module (PM), Hourglass Module (HM), and Focal Wing Loss. We will delve into the effectiveness of these components in the following paragraph. For comparison, we use stacked hourglass model with ResNet block as our baseline and it is trained with adaptive wing loss.

Propagation Module plays an important role in enhancing our model’s performance. It makes the largest improve-

Table 4. The potential of Propagation Module’s (PM) contribution to our model’s performance.

Method	Without PM	With PM
NME	4.81	4.48
FR _{10%}	3.36	3.12
AUC _{10%}	0.5132	0.5421

ment to our model. We set our baseline as a stacked hourglass network without this module. See Table 4 and compare the relation-block-enhanced model with the baseline model. We can observe -6.86% , -7.14% , 5.63% increase respectively in NME (the lower the better), FR (the lower the better), and AUC (the larger the better). From Fig. 8, we can see the actual results of generated boundary heatmaps. They are consistent with our expectation, and substantiate our presumption that landmark heatmaps can be propagated to boundary heatmaps via a few consecutive convolution operations. Furthermore, note that our algorithm remain robust in extreme conditions, especially when human face is being occluded, which means the structural information has been captured through our propagation module.

Table 5. The potential of Multi-view Hourglass Module’s (MHM) contribution to our model’s performance compared to baseline model with Bottleneck Block (BB).

Method	BB	MHM
NME	4.81	4.67
FR _{10%}	3.36	3.16
AUC _{10%}	0.5132	0.5300

Table 6. Comparison between anti-aliased CNN (AC) with different size of Gaussian kernel.

Method	BL	BL+AC-2	BL+AC-3	BL+AC-5
NME	4.81	4.79	4.67	4.75
FR _{10%}	3.36	3.80	3.16	3.76
AUC _{10%}	0.5132	0.5178	0.5300	0.5200

Hourglass Module is one effective module to improve our network’s performance on WFLW dataset. Take a look at Table 5. In comparison with the baseline model with bottleneck block, it increases all three metrics by about -2.91% , -5.95% , 3.27% . When encountered with the choice of Gaussian kernel size for anti-aliased CNN, we compare different sizes against the baseline model. See Table 6. We use AC- n to indicate the Gaussian kernel of size n . For example, AC-3 stands out and we use the size 3 in the rest experiments.

Focal Wing Loss also contributes to the improvement of our model’s performance. As can be seen in Table 7, it gives a rise to three metric increments compared to baseline model trained with AWing by around -3.53% , -1.19% , 1.23% respectively. In addition, we can also see from Table 3 that our model performs better than other state-of-the-arts

Table 7. The potential of Focal Wing Loss’s contribution to overall performance.

Method	AWing	Focal Wing
NME	4.81	4.64
FR _{10%}	3.36	3.32
AUC _{10%}	0.5132	0.5195

on every subsets, which means data imbalance has been effectively addressed and once again it helps our network to sustain its robustness against extreme conditions (see fig. 4).

Table 8. Complexity of PropNet and some other state-of-the-arts.

Method	LAB[25]	AWing[24]	PropNet
#params (M)	12.29	24.15	36.30
FLOPS (G)	18.85	26.79	42.83

See Table 8. We make a comparison of computational complexity with some of the open-source state-of-the-arts. As can be seen in the table, we have greater number of parameters and FLOPS than the other two, which may explain why we achieve better performance than them.

5. Conclusion

In our paper, we pinpoint the long-ignored relation between landmark heatmaps and boundary heatmaps. To this end, we propose a Propagation Module to capture the structure information of human face and bridge the gap between landmark heatmap and boundary heatmap. This module is proven by our extensive experiments on widely recognized datasets to be effective and beneficial to the improvement of our algorithm’s performance.

Then we creatively formulate our method to solve data imbalance by introducing the Focal Factor, a factor attempting to dynamically accommodate the loss weight on each sample in a batch. As our ablation studies show, it makes our algorithm more robust under extreme conditions.

Finally, we also redesign hourglass network by incorporating multi-view blocks and anti-aliased network. The multi-view blocks enables our network to have both macro and micro receptive fields while the anti-aliased architecture make our network shift invariant again. Our ablation studies substantiate its usefulness in the enhancement of our performance.

6. Acknowledgment

This work was sponsored by DiDi GAIA Research Collaboration Initiative.

References

- [1] Adrian Bulat and Georgios Tzimiropoulos. Binarized convolutional landmark localizers for human pose estimation and

- face alignment with limited resources. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3706–3714, 2017. 3, 5
- [2] Xavier P Burgos-Artizzu, Pietro Perona, and Piotr Dollár. Robust face landmark estimation under occlusion. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1513–1520, 2013. 4, 7
- [3] Xudong Cao, Yichen Wei, Fang Wen, and Jian Sun. Face alignment by explicit shape regression. *International Journal of Computer Vision*, 107(2):177–190, 2014. 6
- [4] Timothy F Cootes, Christopher J Taylor, David H Cooper, and Jim Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995. 1
- [5] Jiankang Deng, George Trigeorgis, Yuxiang Zhou, and Stefanos Zafeiriou. Joint multi-view face alignment in the wild. *IEEE Transactions on Image Processing*, 28(7):3636–3648, 2019. 3
- [6] Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Style aggregated network for facial landmark detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 379–388, 2018. 5
- [7] Pengfei Dou, Shishir K Shah, and Ioannis A Kakadiaris. End-to-end 3d face reconstruction with deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5908–5917, 2017. 1
- [8] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2235–2245, 2018. 3, 4, 5, 6
- [9] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2481–2490, 2017. 4
- [10] Tal Hassner, Shai Harel, Eran Paz, and Roei Enbar. Effective face frontalization in unconstrained images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4295–4304, 2015. 1
- [11] Sina Honari, Pavlo Molchanov, Stephen Tyree, Pascal Vincent, Christopher Pal, and Jan Kautz. Improving landmark localization with semi-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1546–1555, 2018. 5
- [12] Amit Kumar and Rama Chellappa. Disentangling 3d pose in a dendritic cnn for unconstrained 2d face alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 430–439, 2018. 5
- [13] Shan Li and Weihong Deng. Deep facial expression recognition: A survey. *CoRR*, abs/1804.08348, 2018. 1
- [14] Rosanne Liu, Joel Lehman, Piero Molino, Felipe Petroski Such, Eric Frank, Alex Sergeev, and Jason Yosinski. An intriguing failing of convolutional neural networks and the coordconv solution, 2018. 6
- [15] Jiangjing Lv, Xiaohu Shao, Junliang Xing, Cheng Cheng, and Xi Zhou. A deep regression architecture with two-stage re-initialization for high performance facial landmark detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3317–3326, 2017. 5
- [16] Xin Miao, Xiantong Zhen, Xianglong Liu, Cheng Deng, Vassilis Athitsos, and Heng Huang. Direct shape regression networks for end-to-end face alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5040–5049, 2018. 5
- [17] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 3, 7
- [18] Shaoqing Ren, Xudong Cao, Yichen Wei, and Jian Sun. Face alignment at 3000 fps via regressing local binary features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1685–1692, 2014. 5
- [19] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 397–403, 2013. 7
- [20] Zhiqiang Tang, Xi Peng, Shijie Geng, Lingfei Wu, Shaoting Zhang, and Dimitris Metaxas. Quantized densely connected u-nets for efficient landmark localization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 339–354, 2018. 5
- [21] George Trigeorgis, Patrick Snape, Mihalios A Nicolaou, Epameinondas Antonakos, and Stefanos Zafeiriou. 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*, pages 4177–4187, 2016. 5
- [22] Roberto Valle, Jose M Buenaposada, Antonio Valdés, and Luis Baumela. A deeply-initialized coarse-to-fine ensemble of regression trees for face alignment. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 585–601, 2018. 5
- [23] Mei Wang and Weihong Deng. Deep face recognition: A survey. *CoRR*, abs/1804.06655, 2018. 1
- [24] Xinyao Wang, Liefeng Bo, and Li Fuxin. Adaptive wing loss for robust face alignment via heatmap regression. *arXiv preprint arXiv:1904.07399*, 2019. 2, 3, 4, 5, 7, 8
- [25] Wayne Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, and Qiang Zhou. Look at boundary: A boundary-aware face alignment algorithm. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2129–2138, 2018. 1, 3, 4, 5, 6, 7, 8
- [26] Wenyan Wu and Shuo Yang. Leveraging intra and inter-dataset variations for robust face alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 150–159, 2017. 6
- [27] Yue Wu and Qiang Ji. Robust facial landmark detection under significant head poses and occlusion. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3658–3666, 2015. 5
- [28] Shengtao Xiao, Jiashi Feng, Junliang Xing, Hanjiang Lai, Shuicheng Yan, and Ashraf Kassim. Robust facial landmark detection via recurrent attentive-refinement networks.

- In *European conference on computer vision*, pages 57–72. Springer, 2016. 5
- [29] Xuehan Xiong and Fernando De la Torre. Supervised descent method and its applications to face alignment. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 532–539, 2013. 5, 6
 - [30] Jie Zhang, Shiguang Shan, Meina Kan, and Xilin Chen. Coarse-to-fine auto-encoder networks (cfan) for real-time face alignment. In *European conference on computer vision*, pages 1–16. Springer, 2014. 5
 - [31] Richard Zhang. Making convolutional networks shift-invariant again, 2019. 2, 3, 6
 - [32] Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Facial landmark detection by deep multi-task learning. In *European conference on computer vision*, pages 94–108. Springer, 2014. 4
 - [33] Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Learning deep representation for face alignment with auxiliary attributes. *IEEE transactions on pattern analysis and machine intelligence*, 38(5):918–930, 2015. 5
 - [34] Shizhan Zhu, Cheng Li, Chen Change Loy, and Xiaoou Tang. Face alignment by coarse-to-fine shape searching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4998–5006, 2015. 5, 6