Grand Challenge of 106-Point Facial Landmark Localization

Yinglu Liu[1], Hao Shen[1], Yue Si[1], Xiaobo Wang[1], Xiangyu Zhu[2], Hailin Shi[1]✉ ,
Zhibin Hong[3], Hanqi Guo[3], Ziyuan Guo[3], Yanqin Chen[3], Bi Li[3], Teng Xi[3],
Jun Yu[4], Haonian Xie[4], Guochen Xie[4], Mengyan Li[4], Qing Lu[5], Zengfu Wang[4],
Shenqi Lai[6], Zhenhua Chai[6], Xiaoming Wei[6].

[1]JD AI. [2]Institute of Automation, Chinese Academic of Sciences.
[3]Department of Computer Vision Technology (VIS), Baidu Inc.
[4]University of Science and Technology of China. [5]IFLYTEKCO.,LTD.
[6]Vision and Image Center of Meituan.

Abstract

Facial landmark localization is a very crucial step in numerous face related applications, such as face recognition, facial pose estimation, face image synthesis, etc. However, previous competitions on facial landmark localization (i.e., the 300-W, 300-VW and Menpo challenges) aim to predict 68-point landmarks, which are incompetent to depict the structure of facial components. In order to overcome this problem, we construct a challenging dataset, named JD-landmark. Each image is manually annotated with 106-point landmarks. This dataset covers large variations on pose and expression, which brings a lot of difficulties to predict accurate landmarks. We hold a 106-point facial landmark localization competition[1] on this dataset in conjunction with IEEE International Conference on Multimedia and Expo (ICME) 2019. The purpose of this competition is to discover effective and robust facial landmark localization approaches.

1. Introduction

Facial landmark localization, which is to predict the co-ordinates of a set of pre-defined key points on human face, plays an important role in numerous facial applications. For example, it is commonly used for face geometric normalization which is a crucial step for face recognition. Besides, landmarks are often employed to support more and more interesting applications due to their abundant geometric information, *e.g.*, 3D face reconstruction and face image synthesis. In recent years, the deep learning methods have been largely developed and the performances are continu-

ously improved in facial landmark localization task. However, facial features vary greatly from one individual to another. Even for a single individual, there is a large amount of variations due to the pose, expression, and illumination conditions. There still exist many challenges to be addressed. The iBUG group[2] held several competitions on facial landmark localization. Nevertheless, they all focus on the 68-point landmarks which are incompetent to depict the structure of facial components, *e.g.*, there is no points defined on the lower boundary of eyebrow and the wing of nose. To overcome the above problems, we construct a challenging dataset and hold a competition of 106-point facial landmark localization in conjunction with ICME 2019 on this dataset. The purpose of this competition is to promote the development of research on 106-point facial landmark localization, especially dealing with the complex situations, and discover effective and robust approaches in this field. It has attracted wide attention from both academia and industry. Finally, more than 20 teams participated in this competition. We will introduce the approaches and results of the top three teams in this paper.

2. JD-landmark Dataset

In order to develop advanced approaches for dense landmark localization, we construct a new dataset, named JD-landmark[3]. It consists of about 16,000 images. As Tab. 1 shows, our dataset covers large variations of pose, in particular, the percent of images with pose angle large than 30^o is more than 16%. The training, validation and test sets are described as follows:

- **Training set:** We collect an incremental dataset based

(a) examples of training set   (b) examples of validation set
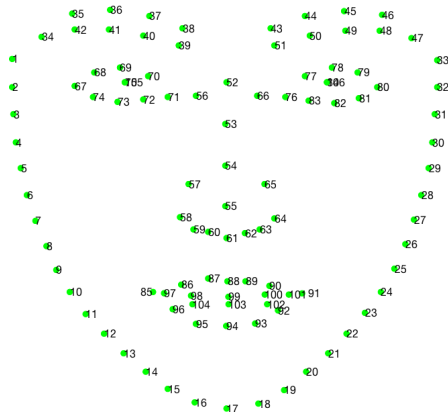
Figure 1. examples of JD-landmark dataset.

Figure 2. The 106-point landmark make-up.

on 300W [11, 10, 16], composed of LFPW [1], AFW [9], HELEN [7] and IBUG [12], and re-annotate them with the 106-point mark-up as Fig. 2 shows. This dataset, containing 11,393 face images, is applied as the training set. It is accessible to the participants (with landmark annotations). Fig.1(a) shows some examples of training set.

- **Validation set:** 2,000 web face images, covering large variations of pose, expression and occlusion, are selected from open source web face database [6]. The participants could optimize the parameters on this set before the final evaluation. Fig. 1(b) shows some examples of validation set.

- **Test set:** It contains 2,000 web face images as well, which is blind to participants throughout the competition. It will be used for the final evaluation.

We emphasize that we provide the bounding boxes obtained by our detector for training/validation/test sets. However, participants have the choice of employing other face detectors.

Table 1. Statistics on pose variations.

$0^o \sim 15^o$	$15^o \sim 30^o$	$> 30^o$
34.9%	48.7%	16.4%

3. Evaluation Results

3.1. Evaluation criterion

All submissions are assessed on the total 106-point landmarks as Fig. 2 shows. The average Euclidean point-to-point error normalized by the bounding box size is taken as the metric, which is computed as:

$$NME = \frac{1}{N} \sum_{k=1}^{N} \frac{\|y_k - \hat{y}_k\|_2}{d} \tag{1}$$

where k refers to the index of landmarks. y and $\hat{y}$ denotes the ground truths and the predictions of landmarks for a given face image. In order to alleviate the bias for profile faces caused by the small interocular distance, we employ the square-root of the ground truth bounding box as the normalization factor d, computed as $d = \sqrt{w_{bbox} \times h_{bbox}}$. Here w_{bbox} and h_{bbox} are the width and height of the enclosing rectangle of the ground truth landmarks. If no face is detected, the NME will be set to infinite. The Cumulative Error Distribution (CED) curve corresponding to the percentage of test images of which the error is less than 8% is produced, and the Area-Under-the-Curve (AUC) from the CED curve is calculated as the final evaluation criterion. Besides, further statistics from the CED curves such as the failure rate and average NME are also presented for reference.

3.2. Participation

A total of 23 teams participated in this challenge. Due to the space limitation, we will briefly describe the submitted methods of the top three winners in this subsection.

Hong *et al.* proposed a Multi-Stack Face Alignment method based on autoML [5]. It consists of several stacked hourglass models [8, 15, 3] and performs landmark localization from coarse to fine. The final results are obtained by fusing the outputs based on a voting strategy which could find the most confident cluster and reject outliers. The base models are developed with the help of autoML and trained with a well-designed data augmentation scheme. In addition, one of the base models is jointly trained with segmentation as multi-task learning to take advantage of extra supervision. Equipped with the above designs, the method could perform precise facial landmark localization in various conditions including those with large pose and occlusions.

Yu *et al.* employed a Densely U-Nets Refine Network (DURN) for facial landmark localization. As shown in

2

Table 2. Final results for the 106-point Facial Landmark Localization competition. The top three teams are listed according to their rank in the challenge. Methods are ranked according to the AUC of the CED curve. The Failure rate and NME are also presented for reference.

Rank	Participant	AUC(%)	Failure rate(%)	NME(%)
1	Z. Hong, Z. Guo, Y. Chen, H. Guo, B. Li and T. Xi Department of Computer Vision Technology (VIS), Baidu Inc.	84.01	0.10	1.31
2	J. Yu, H. Xie, G. Xie, M. Li, Q. Lu and Z. Wang University of Science and Technology of China.	82.68	0.05	1.41
3	S. Lai, Z. Chai and X. Wei Vision and Image Center of Meituan	82.22	0.00	1.42

Fig. 3, it involves two sub-networks: DU-Net and Refine-Net. The DU-Net is based on Tang *et al.* [14], where the original intermediate supervision is modified to multi-scale intermediate supervision. It means that each DU-Net employs four intermediate supervision rather than one. The Refine-Net is based on Chen *et al.* [4], and Yu *et al.* add the integral regression [13] after the Refine-Net to obtain the keypoint coordinates instead of the heatmap via argmax function. In addition, the regression loss is computed by the coordinate rather than the heatmap. Finally, Yu *et al.* ensemble 7 models with the similar structure.
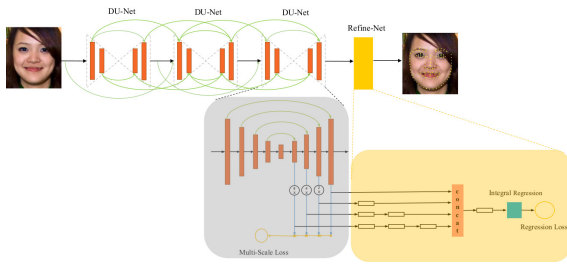

Figure 3. Densely U-Nets Refine Network.

Lai *et al.* proposed an end-to-end trainable facial landmark localization framework, which has achieved promising localization accuracy even under challenging wild environments (*e.g.* unconstrained pose, expression, lighting and occlusion). Different from the classical four stage stacked HGs [8], they propose to use the hierarchical module [2] rather than the standard residual block, which will generate the probability heatmap for each landmark and can make the non-linearity stronger. Besides, in previous work [2] researchers use argmax and post-process operations (e.g. rescale) to get final results, which may decrease the performance by the coordinate quantization. In order to overcome this problem, dual soft argmax function is proposed to map probability of heatmap to numerical coordinates, which is shown in Fig. 4. For a gaussian response in an image, with matrix X and matrix Y, the coordinates x and y can be computed directly. Finally, three models (*i.e.* 64×64 SA, 64×64 DSA and 128×128 DSA) are trained in total, where the number means the size of output heatmaps, SA means soft argmax and DSA means dual soft argmax. The weighted predictions of three models will be used as the final results.
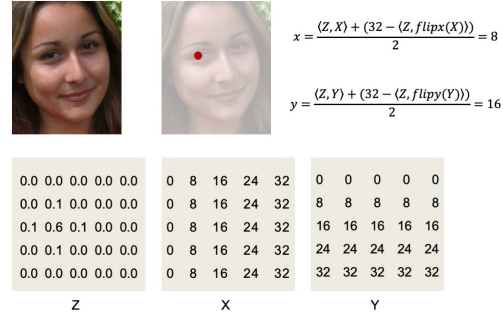

Figure 4. Flipx means flipping matrix horizontally and flipy means flip matric vertically.

3.3. Results

As is mentioned in Sec. 3.1, the submissions are ranked according to the AUC of the CED curve with the threshold of 8%. The winner is Hong *et al.* from Baidu Inc. Second place goes to Yu *et al.* from University of Science and Technology of China. Lai *et al.* from Meituan achieves the 3rd place. Fig. 5 draws the CED curves of the top three teams on the JD-landmark test set. In order to comprehensively evaluate the submissions, we also report the average NME defined as Eq. (1) and the failure rate (if the average NME is larger than 8%, the picture will be taken as a failure prediction.) in Tab. 2. We can see that Hong *et al.* achieved the highest AUC of 84.01%, higher than Yu *et al.* and Lai *et al.* by 1.33% and 1.79%, respectively. Hong *et al.* also performed the best on NME (1.31%), lower than Yu *et al.* and Lai *et al.* by 0.1% and 0.11%, respectively.

4. Conclusion

In this paper, we summarize the grand challenge of 106-point facial landmark localization in conjunction with ICME 2019. We construct and release a new facial landmark dataset, named JD-landmark. Compared with previous challenges on facial landmark localization, our work pays attention on 106-point landmarks which contain more structure information than the 68-point landmarks. Meanwhile, our dataset covers large variations of poses and expressions, which bring a lot of difficulties for participants. Finally, more than 20 teams submitted their binaries or
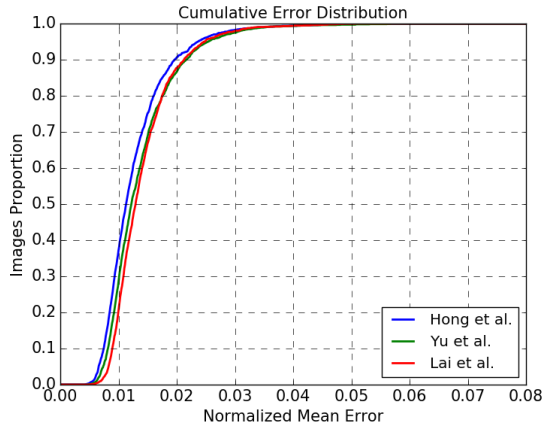
Figure 5. The Cumulative Errors Distribution curves of top three teams on the JD-landmark test set.

models. We introduced the methods together with the performance of top three teams in this paper. We hope this work could make contributions on the development of facial landmark localization.

References

[1] Peter N Belhumeur, David W Jacobs, David J Kriegman, and Neeraj Kumar. Localizing parts of faces using a consensus of exemplars. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2930–2940, 2013.

[2] 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] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks). In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1021–1030, 2017.

[4] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7103–7112, 2018.

[5] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision*, pages 784–800, 2018.

[6] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4873–4882, 2016.

[7] Vuong Le, Jonathan Brandt, Zhe Lin, Lubomir Bourdev, and Thomas S Huang. Interactive facial feature localization. In *Proceedings of the European conference on computer vision*, pages 679–692. Springer, 2012.

[8] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Proceedings of the European Conference on Computer Vision*, pages 483–499. Springer, 2016.

[9] Deva Ramanan and Xiangxin Zhu. Face detection, pose estimation, and landmark localization in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2879–2886. Citeseer, 2012.

[10] Christos Sagonas, Epameinondas Antonakos, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: Database and results. *Image and vision computing*, 47:3–18, 2016.

[11] 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.

[12] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. A semi-automatic methodology for facial landmark annotation. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 896–903, 2013.

[13] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings of the European Conference on Computer Vision*, pages 529–545, 2018.

[14] 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*, pages 339–354, 2018.

[15] Jing Yang, Qingshan Liu, and Kaihua Zhang. Stacked hourglass network for robust facial landmark localisation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 79–87, 2017.

[16] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z Li. Face alignment across large poses: A 3d solution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 146–155, 2016.