

Two-Stage Convolutional Part Heatmap Regression for the 1st 3D Face Alignment in the Wild (3DFAW) Challenge

Adrian Bulat^(✉) and Georgios Tzimiropoulos

Computer Vision Laboratory, University of Nottingham, Nottingham, UK
{adrian.bulat,yorgos.tzimiropoulos}@nottingham.ac.uk

Abstract. This paper describes our submission to the 1st 3D Face Alignment in the Wild (3DFAW) Challenge. Our method builds upon the idea of convolutional part heatmap regression (Bulat and Tzimiropoulos, 2016), extending it for 3D face alignment. Our method decomposes the problem into two parts: (a) X,Y (2D) estimation and (b) Z (depth) estimation. At the first stage, our method estimates the X,Y coordinates of the facial landmarks by producing a set of 2D heatmaps, one for each landmark, using convolutional part heatmap regression. Then, these heatmaps, alongside the input RGB image, are used as input to a very deep subnetwork trained via residual learning for regressing the Z coordinate. Our method ranked 1st in the 3DFAW Challenge, surpassing the second best result by more than 22%. Code can be found at <http://www.cs.nott.ac.uk/~psxab5/>.

Keywords: 3D face alignment · Convolutional Neural Networks · Convolutional part heatmap regression

1 Introduction

Face alignment is the problem of localizing a set of facial landmarks in 2D images. It is a well-studied problem in Computer Vision research, yet most of prior work [1, 2], datasets [3, 4] and challenges [4, 5] have focused on frontal images. However, under a totally unconstrained scenario, faces might be in arbitrary poses. To address this limitation of prior work, recently, a few methods have been proposed for large pose face alignment [6–9]. 3D face alignment goes one step further by treating the face as a full 3D object and attempting to localize the facial landmarks in 3D space. To boost research in 3D face alignment, the 1st Workshop on 3D Face Alignment in the Wild (3DFAW) & Challenge is organized in conjunction with ECCV 2016 [10]. In this paper, we describe a Convolutional Neural Network (CNN) architecture for 3D face alignment, that ranked 1st in the 3DFAW Challenge, surpassing the second best result by more than 22%.

Our method is a CNN cascade consisting of three connected subnetworks, all learned via residual learning [11, 12]. See Fig. 1. The first two subnetworks perform residual part heatmap regression [13] for estimating the X,Y coordinates of

the facial landmarks. As in [13], the first subnetwork is a part detection network trained to detect the individual facial landmarks using a per-pixel softmax loss. The output of this subnetwork is a set of N landmark detection heatmaps. The second subnetwork is a regression subnetwork that jointly regresses the landmark detection heatmaps stacked with image features to confidence maps representing the location of the landmarks. Then, on top of the first two subnetworks, we added a third very deep subnetwork that estimates the Z coordinate of each fiducial point. The newly introduced network is guided by the heatmaps produced by the 2D regression subnetwork and subsequently learns where to “look” by explicitly exploiting information about the 2D location of the landmarks. We show that the proposed method produces remarkable fitting results for both X, Y and Z coordinates, securing the first place on the 3DFAW Challenge.

This paper is organized as follows: Sect. 2 describes our system in detail. Section 3 describes the experiments performed and our results on the 3DFAW dataset. Finally, Sect. 4 summarizes our contributions and concludes the paper.

2 Method

The proposed method adopts a *divide et impera* technique, splitting the 3D facial landmark estimation problem into two tasks as follows: The first task estimates the X, Y coordinates of the facial landmarks and produces a series of N regression heatmaps, one for each landmark, using the network described in Sect. 2.1. The second task, described in Sect. 2.2, predicts the Z coordinate (i.e. the depth of each landmark), using as input the stacked heatmaps produced by the first task, alongside the input RGB image. The overall architecture is illustrated in Fig. 1.

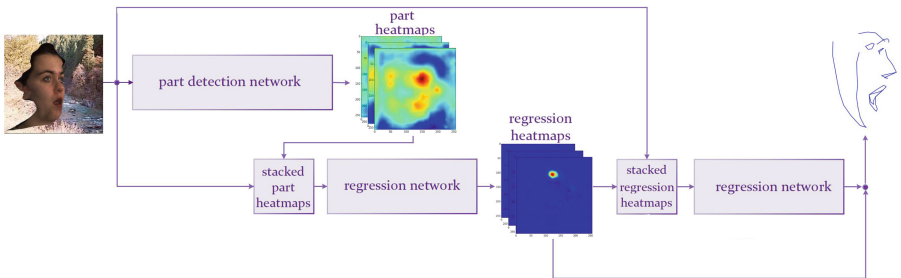


Fig. 1. The system submitted to the 3DFAW Challenge. The part detection and regression subnetworks implement convolutional part heatmap regression, as described in [13], and produce a series of N heatmaps, one for the X, Y location of each landmark. They are both very deep networks trained via residual learning [12]. The produced heatmaps are then stacked alongside the input RGB image, and used as input to the Z regressor which regresses the depth of each point. The architecture for the Z regressor is based on ResNet [12], as described in Sect. 2.2.

2.1 2D (X,Y) Landmark Heatmap Regression

The first part of our system estimates the X,Y coordinates of each facial landmark using part heatmap regression [13]. The network consists of two connected subnetworks, the first of which performs landmark detection while the second one refines the initial estimation of the landmarks’ location via regression. Both networks are very deep been trained via residual learning [11,12]. In the following, we briefly describe the two subnetworks; the exact network architecture and layer specification for each of them are described in detail in [13].

Landmark detection subnetwork. The architecture of the landmark detection network is based on a ResNet-152 model [11]. The network was adapted for landmark localization by: (1) removing the fully connected layer alongside the preceding average pooling layer, (2) changing the stride of the 5th block from 2 to 1 pixels, and (3) adding at the end a deconvolution [14] followed by a fully convolutional layer. These changes convert the model to a fully convolutional network, recovering to some extent the lost spatial resolution. The ground truth was encoded as a set of N binary maps (one for each landmark), where the values located within the radius of the provided ground truth landmark are set to 1, while the rest to 0. The radius defining the “correct location” was empirically set to 7 pixels, for a face with a bounding box height equal to approximately 220 pixels. The network was trained using the pixel wise sigmoid cross entropy loss function.

Landmark regression subnetwork. As in [13], the regression subnetwork plays the role of a graphical model aiming to refine the initial prediction of the landmark detection network. It is based on a modified version of the “hourglass network” [15]. The hourglass network starts from the idea presented in [16], improving a few important concepts: (1) it updates the model using residual learning [11,12], and (2) introduces an efficient way to analyze and recombine features at different resolutions. Finally, as in [13], we replaced the original nearest neighbour upsampling of [15] by learnable deconvolutional layers, and added another deconvolutional layer in the end that brings the output to the input resolution. The network was then trained using a pixel wise L2 loss function.



Fig. 2. The architecture of the Z regression subnetwork. The network is based on ResNet-200 (with preactivation) and its composing blocks. The blocks B1–B6 are defined in Table 1. See also text.

Table 1. Block specification for the Z regression network. Torch notations (channels, kernel, stride) and (kernel, stride) are used to define the conv and pooling layers. The bottleneck modules are defined as in [12].

B1	B2	B3	B4	B5	B6
1x conv layer (64,7 × 7,2 × 2)	3x bottleneck modules	24x bottleneck modules	38x bottleneck modules	3x bottleneck modules	1x fully connected layer (66)
1x pooling (3 × 3, 2 × 2)	[(64,1 × 1), (64,3 × 3), (256,1 × 1)]	[(128,1 × 1), (128,3 × 3), (512,1 × 1)]	[(256,1 × 1), (256,3 × 3), (1024,1 × 1)]	[(512,1 × 1), (512,3 × 3), (2048,1 × 1)]	

2.2 Z Regression

In this section, we introduce a third subnetwork for estimating the Z coordinate i.e. the depth of each landmark. As with the X,Y coordinates, the estimation of the Z coordinate is performed jointly for all landmarks. The input to the Z regressor subnetwork is the stacked regression heatmaps produced by the regression subnetwork alongside the input RGB image. The use of the stacked heatmaps is a key feature of the subnetwork as they provide pose-related information (encoded by the X,Y location of all the landmarks) and guide the network where to “look”, explicitly showing where the depth should be estimated.

We encode each landmark as a heatmap using a 2D Gaussian with std = 6 pixels centered at the X,Y coordinates of that landmark. The proposed Z regression network is based on the latest ResNet-200 network with preactivation modules [12] modified as follows: in order to adapt the model for 1D regression, we replaced the last fully connected layer (used for classification) with one that has N output channels, one for each landmark. Additionally, the first convolutional layer of the network was modified to accommodate 3+N input channels. The network is described in detail in Fig. 2 and Table 1. All newly introduced filters were initialized randomly using a Gaussian distribution. Finally, the network was trained using the L2 loss:

$$l_2 = \frac{1}{N} \sum_{n=1}^N (\tilde{z}_n - z_n)^2, \quad (1)$$

where \tilde{z}_n and z_n are the predicted and ground truth Z values (in pixels) for the n th landmark.

2.3 Training

For training, all images were cropped around an extended (by 20–25%) bounding box, and then resized so that the final cropped image had a resolution of 384×384 pixels. While batch normalization is known to prevent overfitting to some extent, we additionally augmented the data with a set of image transformations applied randomly: flipping, in-plane rotation (between -35° and 35°), scaling (between 0.85 and 1.15) and colour jittering. While the entire system,

shown in Fig. 1, can be trained jointly from the beginning, in order to accelerate convergence we trained each task independently.

For 2D (X,Y) landmark heatmap regression, the network was fine-tuned from a pretrained model on the large ImageNet [17] dataset, with the newly introduced layers initialized with zeros. The detection component was then trained for 30 epochs with a learning rate progressively decreasing from $1e - 3$ to $2.5e - 5$. The regression subnetwork, based on the ‘‘hourglass’’ architecture was trained for 30 epochs using a learning rate that varied from $1e - 4$ to $2.5e - 5$. During this, the learning rate for the detection subnetwork was frozen. All the newly introduced deconvolutional layers were initialised using the bilinear upsampling filters. Finally, the subnetworks were trained jointly for 30 more epochs.

For Z regression, we again fine-tuned from a model previously trained on ImageNet [17]; this time we used a ResNet-200 network [12]. The newly introduced filters in the first convolutional layer were initialized from a Gaussian distribution with $\text{std} = 0.01$. The same applied for the fully connected layer added at the end of the network. During training, we used as input to the Z regression network both the heatmaps generated from the ground truth landmark locations and the ones estimated by the first task. We trained the subnetwork for about 100 epochs with a learning rate varying from $1e - 2$ to $2.5e - 4$.

The network was implemented and trained using Torch7 [18] on two Titan X 12Gb GPUs using a batch of 8 and 16 images for X,Y landmark heatmap regression and Z regression, respectively.

3 Experimental Results

In this section, we present the performance of our system on the 3DFAW dataset.

Dataset. We trained and tested our model on the 3D Face Alignment in the Wild (3DFAW) dataset. [19–22] The dataset contains images from a wide range of conditions, captured in both controlled and in-the-wild settings. The dataset includes images from MultiPIE [19] and BP4D [21] as well as images collected from the Internet. All images were annotated in a consistent way with 66 3D fiducial points. The final model was trained on the training set (13672 images) and the validation set (4725 images), and tested on the test set containing 4912 images. We also report results on the validation set with a model trained on the training set, only.

Metrics. Evaluation was performed using two different metrics: Ground Truth Error(GTE) and Cross View Ground Truth Consistency Error(CVGTCE).

GTE measures the average point-to-point Euclidean error normalized by the inter-ocular distance, as in [4]. GTE is calculated as follows:

$$E(\mathbf{X}, \mathbf{Y}) = \frac{1}{N} \sum_{n=1}^N \frac{\|\mathbf{x}_n - \mathbf{y}_n\|_2}{d_i}, \quad (2)$$

where \mathbf{X} is the predicted set of points, \mathbf{Y} is their corresponding ground truth and d_i denotes the interocular distance for the i th image.

Table 2. Performance on the 3DFAW test set.

Method	GTE (%)	CVGTCE (%)
Ours	4.5623	3.4767
Second best	5.8835	3.9700

CVGTCE evaluates the cross-view consistency of the predicted landmarks of the same subject and is defined as follows:

$$E_{vc}(\mathbf{X}, \mathbf{Y}, T) = \frac{1}{N} \sum_{n=1}^N \frac{\|s\mathbf{R}\mathbf{x}_n - \mathbf{y}_n\|_2}{d_i}, \quad (3)$$

where $P = \{s, \mathbf{R}, \mathbf{t}\}$ denotes scale, rotation and translation, respectively. CVGTCE is computed as follows:

$$\{s, \mathbf{R}, \mathbf{t}\} = \operatorname{argmin}_{s, \mathbf{R}, \mathbf{t}} \sum_{n=1}^N \|\mathbf{y}_k - (s\mathbf{R}\mathbf{x}_k + \mathbf{t})\|_2^2. \quad (4)$$

Results. Table 2 shows the performance of our system on the 3DFAW test set, as provided by the 3DFAW Challenge team. As it can be observed, our system outperforms the second best method (the result is taken from the 3DFAW Challenge website) in terms of GTE by more than 22%.




Fig. 3. GTE vs fraction of test images on the 3DFAW validation set, on (X,Y), (X,Y,Z), X alone, Y alone and Z alone.

Table 3. Performance on the 3DFAW validation set, on (X,Y), (X,Y,Z), X alone, Y alone and Z alone.

Axes	GTE (%)
xy	3.6263
xyz	4.9408
x	2.12
y	2.48
z	2.77



Fig. 4. Fitting results produced by our system on the 3DFAW test set. Observe that our method copes well with a large variety of poses and facial expressions on both controlled and in-the-wild images. Best viewed in colour. (Color figure online)

In order to better understand the performance of our system, we also report results on the validation set using a model trained on the training set, only. To measure performance, we used the Ground Truth Error measured on (X,Y) , (X,Y,Z) , X alone, Y alone and Z alone, and report the cumulative curve corresponding to the fraction of images for which the error was less than a specific value. Results are reported in Fig. 3 and Table 3. It can be observed that our system performs better on predicting the X and Y coordinates (compared to Z), but this difference is quite small and, to some extent, expected as Z is estimated at a later stage of the cascade. Finally, Fig. 4 shows a few fitting results produced by our system.

4 Conclusions

In this paper, we proposed a two-stage CNN cascade for 3D face alignment. The method is based on the idea of splitting the 3D alignment task into two separate subtasks: 2D landmark estimation and 1D (depth) estimation, where the first one guides the second. Our system secured the first place in the 1st 3D Face Alignment in the Wild (3DFAW) Challenge.

Acknowledgments. This work was supported in part by the EPSRC project EP/M02153X/1 Facial Deformable Models of Animals.

References

1. Cao, X., Wei, Y., Wen, F., Sun, J.: Face alignment by explicit shape regression. In: CVPR (2012)
2. Xiong, X., De la Torre, F.: Supervised descent method and its applications to face alignment. In: CVPR, pp. 532–539(2013)
3. Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., Pantic, M.: A semi-automatic methodology for facial landmark annotation. In: CVPR-W (2013)
4. Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., Pantic, M.: 300 faces in-the-wild challenge: the first facial landmark localization challenge. In: CVPR, pp. 397–403 (2013)
5. Shen, J., Zafeiriou, S., Chrysos, G., Kossaifi, J., Tzimiropoulos, G., Pantic, M.: The first facial landmark tracking in-the-wild challenge: benchmark and results. In: ICCV-W (2015)
6. Jourabloo, A., Liu, X.: Pose-invariant 3D face alignment. In: CVPR, pp. 3694–3702 (2015)
7. Jourabloo, A., Liu, X.: Large-pose face alignment via CNN-based dense 3D model fitting. In: CVPR (2016)
8. Zhu, S., Li, C., Change Loy, C., Tang, X.: Face alignment by coarse-to-fine shape searching. In: CVPR, pp. 4998–5006 (2015)
9. Bulat, A., Tzimiropoulos, G.: Convolutional aggregation of local evidence for large pose face alignment. In: BMVC (2016)
10. 1st Workshop on 3D Face Alignment in the Wild (3DFAW) & Challenge. <http://mhug.disi.unitn.it/workshop/3dfaw/>. Accessed 30 Aug 2016

11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. arXiv preprint [arXiv:1512.03385](https://arxiv.org/abs/1512.03385) (2015)
12. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. arXiv preprint [arXiv:1603.05027](https://arxiv.org/abs/1603.05027) (2016)
13. Bulat, A., Tzimiropoulos, G.: Human pose estimation via convolutional part heatmap regression. In: ECCV (2016)
14. Zeiler, M.D., Taylor, G.W., Fergus, R.: Adaptive deconvolutional networks for mid and high level feature learning. In: ICCV, pp. 2018–2025. IEEE (2011)
15. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. arXiv preprint [arXiv:1603.06937](https://arxiv.org/abs/1603.06937) (2016)
16. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR, pp. 3431–3440 (2015)
17. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR, pp. 248–255. IEEE (2009)
18. Collobert, R., Kavukcuoglu, K., Farabet, C.: Torch7: A matlab-like environment for machine learning. In: NIPS-W. Number EPFL-CONF-192376 (2011)
19. Gross, R., Matthews, I., Cohn, J., Kanade, T., Baker, S.: Multi-pie. *Image Vis. Comput.* **28**(5), 807–813 (2010)
20. Yin, L., Chen, X., Sun, Y., Worm, T., Reale, M.: A high-resolution 3D dynamic facial expression database. In: 8th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2008, pp. 1–6. IEEE (2008)
21. Zhang, X., Yin, L., Cohn, J.F., Canavan, S., Reale, M., Horowitz, A., Liu, P., Girard, J.M.: BP4D-spontaneous: a high-resolution spontaneous 3D dynamic facial expression database. *Image Vis. Comput.* **32**(10), 692–706 (2014)
22. Jeni, L.A., Cohn, J.F., Kanade, T.: Dense 3D face alignment from 2D video for real-time use. *Image and Vision Computing* (2016)