

# Learning to validate the quality of detected landmarks

Wolfgang Fuhl<sup>1</sup>, Enkelejda Kasneci<sup>1</sup>,

<sup>1</sup> Eberhard Karls Universität Tübingen, 72076 Tübingen, Germany

## ABSTRACT

We present a new loss function for the validation of image landmarks detected via Convolutional Neural Networks (CNN). The network learns to estimate how accurate its landmark estimation is. This loss function is applicable to all regression-based location estimations and allows the exclusion of unreliable landmarks from further processing. In addition, we formulate a novel batch balancing approach which weights the importance of samples based on their produced loss. This is done by computing a probability distribution mapping on an interval from which samples can be selected using a uniform random selection scheme. We conducted experiments on the 300W, AFLW, and WFLW facial landmark datasets. In the first experiments, the influence of our batch balancing approach is evaluated by comparing it against uniform sampling. In addition, we evaluated the impact of the validation loss on the landmark accuracy based on uniform sampling. The last experiments evaluate the correlation of the validation signal with the landmark accuracy. All experiments were performed for all three datasets.

**Keywords:** landmark detection, landmark validation, deep learning, batch creation

## 1. INTRODUCTION

Landmark localization describes the process of determining the location of characteristic points of an object in an image. These points are usually characterized by a spatial relationship to each other such as the nose and eyes of a face. Under pose changes of the object, landmarks cannot change positions randomly but according to their relative positions towards each other. The problem of landmark localization is common in 2D face alignment, face reconstruction or gesture recognition [1, 2]. These are also preparatory steps for head pose estimation [3, 4], emotion estimation [5, 6] and face recognition [7, 8]. The current state of the art in this area addresses the problems of accurate and robust landmark detection in real and constrained scenarios. Known challenges are the head pose, facial expressions, illumination changes, blurring, and occlusion. Deep neural networks [9, 10] are currently in the focus of attention since they outperformed any other machine learning approach for common computer vision tasks. This includes landmark detection [1, 11] with a multitude of different network architectures, including convolution neuronal networks [12], auto encoders [13], recurrent networks [14], residual networks [15, 16] and hourglass networks [1]. Recent developments emphasize the importance of the loss function as a key element for improving the accuracy and robustness of landmark localization, such as the newly proposed Wing loss [15]. In addition, data balancing plays an important role in training: challenging poses are often massively underrepresented in the training data. Therefore, pose-based data balancing using principal component analysis was proposed by Feng et al. [15] and a cycle generative adversarial network was employed in [16]. Another novel architecture, a local to global context network proposed in [17], outputs heat map candidates and is therefore capable of detecting landmarks for multiple persons in an image. In this work, we address the challenge of facial landmark validation and batch balancing for training. The main contributions of this work are **1:** A novel loss function that trains the network directly to estimate the accuracy alongside with landmark location. The idea behind our validation loss is that we want the network to guess the reliability in terms of distance to the ground truth position. **2:** An online data augmentation strategy that ranks training samples by their produced loss. This addresses the problem of pose-based data balancing [15] by not only normalizing for the occurrences of head poses but also for (a priori unknown) challenges in the dataset like occlusion and illumination.

## 2. RELATED WORK

State of the art methods employing deep learning for facial landmark detection is based on regression, i.e. the network estimates the correct landmark positions directly. We categorize those approaches by their employed architecture, data balancing procedure, cascading, and loss functions. **Architecture:** A straightforward way to use a CNN for direct 2D landmark regression was proposed in [12, 15]. The input to such a network is a face image which is extracted based on preliminary applied face detection. Novel architectures such as residual [15] and recurrent networks [14] have already

improved the state of the art in regression-based landmark detection. Hourglass networks [1] have also led to an increase in both accuracy and robustness of landmark detection. However, instead of a position vector, its output is a heat map, where each pixel represents the probability to be a landmark position. **Data balancing:** Facial landmarks are subject to extreme pose variations. Head poses present in current datasets are however heavily biased towards frontal images; extreme orientations are rarely represented in the data. To overcome this limitation, multiview models were proposed, splitting the problem into frontal and profile faces. Those were already used in traditional approaches, such as ASM [18] and AAM [19], as well as for cascaded regression-based approaches [20, 21]. In [16], a cycle GAN was used to generate images to handle style variations. The third approach to data balancing is multi-task learning. This means that emotion classification, pose estimation, 3D face reconstruction, and landmark detection are trained jointly [22, 23]. This was found to boost individual subtasks. **Cascaded networks:** The aim of these approaches is to use multiple networks stacked upon each other to refine landmark positions. The straightforward approach is to use a coarse positioning network and refine the positions with a second network which receives the coarse positions and a higher resolution image as input [15, 22]. Another approach is to extract image patches surrounding possible facial landmarks and to refine the position within those patches. This was done using a recurrent network trained in an end to end fashion [14]. **Loss functions:** Recently, the wing loss function was proposed [15]. It is a combination of a log loss function if the error is close to zero and otherwise an L1 loss. Other approaches also deal with the validation of landmarks, which is additional challenge for real world applications [24].

### 3. BATCH BALANCING




Figure 1: Exemplary images which represent a balanced set based on the head pose.

Existing pose-based data balancing approaches [15] are a valuable and effective procedure but have one major drawback: the data is balanced only based on the underrepresentation of some head poses with regard to a parameter of histogram bins. Figure 1 shows an example where data is already balanced based on the head pose. But additional challenges, such as reflections (top right) and three different types of occlusions are only represented by one sample each. Our idea is to use the produced loss of a training image to compute the probability of an image being included in the next training batch. Thereby, we can make the network learn especially the challenging cases, even if they are not annotated as such. Using the loss as a criterion will automatically keep a balance between being able to detect the vast majority of data via an appropriate generalization and at the same time dedicate an adequate amount of training towards the especially challenging samples. Figure 1 shows such a training set. Each sample, therefore, has an assigned value range which allows to create a




Figure 2: The proposed ranking example where the produced loss represents the probability of a sample to be included in a training batch for the next epoch.

new batch using a uniform random sampling.

$$B_j = S(k | \sum_{i=1}^k Loss_i \geq P([0; \sum_{i=1}^n Loss_i])) \quad (1)$$

Equation 1 describes our **Loss Probability based Batch Balancing (LPBB)**. For the selection of the  $j$  training sample ( $S$ ) in our batch ( $B$ ) we first select a value in the range  $[0; \sum_{i=1}^n Loss_i]$  using the uniform distribution  $P$ .  $n$  is the amount of all samples in our training set. Afterward, we search for the  $k$ th sample which has a greater or equal sum of the loss values as our randomly selected value. For completeness, it should also be required that  $k$  is the smallest sum that satisfies this condition (Not integrated into the formula for simplification). This sample is then included in the batch. This also means that samples can be assigned twice in a batch but since we use online data augmentation, there is no limitation (Section 5). Figure 2 shows the sample selection. First the sum is computed based on which the range for the uniform distribution is computed. Afterwards, the sample three is selected since it has a greater sum of loss values.

For an exact determination of the probability based on the loss value, all data must be evaluated after each training batch. Since this represents a considerable effort, two different uses were considered in our evaluation (Section 6). An exact description can be found in Section 5.




Figure 3: The used network architecture for the proposed approach.

#### 4. VALIDATION LOSS FORMULATION

The idea behind our validation loss is that we want the network to guess the expected distance between its prediction and the ground truth position as a proxy of the reliability of the prediction. This means that the network not only estimates the landmark position, it also evaluates its accuracy and provides this information in the additional output. Therefore, we compute the distance between the estimated and annotated position (L2, L1, or any task-specific distance metric) and use this as a training label for the next training iteration. For landmarks the conventional loss formulation is  $\frac{1}{2}(gt_i - ev_i)^2$  for the  $i$ th index of the output neuron where  $gt$  is the ground truth and  $ev$  the estimated value. The derivative and therefore the gradient is  $gt_i - ev_i$ . Each landmark position is estimated using two neurons from the last layer as shown in Equation 2.

$$\begin{aligned} LM_{j,x} &= N_i = gt_i - ev_i \\ LM_{j,y} &= N_{i+1} = gt_{i+1} - ev_{i+1} \end{aligned} \quad (2)$$

$LM_{j,x}$  in Equation 2 is the gradient for the  $x$  coordinate of the  $j$ th landmark and  $LM_{j,y}$  the gradient for the  $y$  coordinate which are represented by neuron  $N_i$  and  $N_{i+1}$  respectively. We extended this formulation by a third value which is the validation signal. As metric between the ground truth position and the estimated value we used the Manhattan distance. The computation of the gradients for one landmark position using the validation loss is shown in Equation 3.

$$\begin{aligned} N_i &= gt_i - ev_i \\ N_{i+1} &= gt_{i+1} - ev_{i+1} \\ N_{i+2} &= |gt_i - ev_i| + |gt_{i+1} - ev_{i+1}| - ev_{i+2} \end{aligned} \quad (3)$$

As can be seen each landmark receives an additional neuron ( $N_{i+2}$ ) which learns to evaluate the quality of the  $x, y$  position estimation of the landmark. The higher this signal is, the less accurate is the landmark estimation. The advantages of this formulation are first than the CNN can learn to estimate the landmark quality end to end and that our formulation is not bound to a distance metric. The Manhattan distance ( $|gt_i - ev_i| + |gt_{i+1} - ev_{i+1}|$ ) can also be replaced by the Euclidean distance ( $\sqrt{(gt_i - ev_i)^2 + (gt_{i+1} - ev_{i+1})^2}$ ) or any other distance metric.

#### 5. SYSTEM, TRAINING, AND AUGMENTATION PARAMETERS

For the training of our network, we used DLIB [25]. The training was performed on a POWER8 server from IBM with four Nvidia Tesla P100 GPUs and 1TB of RAM. For the evaluation and measuring the inference runtime, we used a desktop computer with an NVIDIA GeForce GTX 1050ti, 16GB of RAM and an Intel i5-4570 CPU. The network architecture itself is shown in Figure 3 and consists of an initial convolution layer followed by a max pooling. Afterward, three residual blocks are used with 64, 128 and 256 layers, respectively. After each convolution, we used a batch normalization layer followed

by a rectifier linear unit. The end of our architecture consists of two fully connected layers, wherein the last layer has  $3 * l$  outputs with the validation loss and  $2 * l$  without.  $l$  represents the number of landmarks. This means that for the 300W [3] dataset with 68 landmarks our output layer has  $3 * 68 = 204$  neurons using the validation loss. These three values represent the x and y coordinate as well as the validation signal per landmark.

**Naming convention:** For our evaluation we trained four CNNs for each data set (300W [3], AFLW [26], and WFLW [27]). All CNNs follow the same architecture as shown in Figure 3. NON: CNN using no validation signal and no batch balancing. NOBB: CNN using the validation signal and no batch balancing. BBO: CNN using the validation signal and batch balancing (PLBB) with an online update. BB: CNN using the validation signal and batch balancing (PLBB) with an update of the loss value for all training samples every 10% of an epoch.

**Data augmentation:** For data augmentation, we used randomly added noise of up to 20%. The face bounding box was randomly shifted by up to 20% of the image in any direction (corners or between) and rotated between in the range  $[-1, 1]$  radian. Gaussian blur was added to the image in the range  $=[1.0, 1.3]$ . Additionally, we added occlusions which could cover up to 50% of the image. The gray value of the occlusion was also randomly chosen in the range  $[0, 30]$  (dark),  $[200, 255]$  (bright) or randomly between  $[0, 255]$ . The last data augmentation was the contrast which was adjusted with a randomly selected value between  $[-30, 30]$ . All the augmentations are computed online and each augmentation has a chance of 50% to occur in the image.

**Training parameters:** Weight decay was set to  $5 \times 10^4$  and momentum to 0.9 with a batch size of 10. All images during training and evaluation were converted to grayscale. For the training of the NOBB, BBO, and BB networks we used the Manhattan distance for the validation loss as shown in Equation 3. All models were initialized with random values and trained in three steps. The first step used the L2 loss formulation for the x and y coordinate with a learning rate of  $10^{-8}$ . Each model was trained for  $\approx 2000$  epochs with this configuration. This was done since DLIB ends up in a loss of NaN for random initialization networks if the initial learning rate is set too high. Afterward, the learning rate was increased to  $10^{-7}$  for additional  $\approx 2000$  epochs. For the second training step, we changed the loss formulation from L2 to L1 for the coordinate regression. The training was conducted for an additional  $\approx 2000$  epochs. Afterward, the learning rate was increased to  $10^{-6}$  for all models with the exception of those trained on the WFLW [27] dataset. For the WFLW [27] dataset the learning rate was kept at  $10^{-7}$ . The training of all models was continued for an additional  $\approx 2000$  epochs also the models which are trained on WFLW. For the last and third part of our training, we changed the loss function to the wing loss [15] for the coordinate regression. The wing loss function is a combination of a log loss and an L1 loss function. For large errors ( $> w$ ) it behaves like the L1 loss function and for small errors ( $< w$ ) like a log loss function. We used the parameters  $\epsilon = 2, w = 10$  as in the original paper [15] where  $\epsilon$  is used to limit the curvature of the log function. The training was continued for an additional 900 epochs where after each 300 epochs the learning rate was decreased by  $10^{-1}$  to a minimum of  $10^{-8}$ . For the models trained on the WFLW [27] dataset the learning rate was decreased to  $10^{-8}$  after the first 500 epochs and kept at this value for the remaining 400 epochs.

Table 1: Results in NME (smaller is better) for the networks NON and NOBB trained from scratch without batch balancing. Bold values are the best results.

Dataset	NON	NOBB
300W Full	<b>4.35</b>	4.43
300W Common	<b>3.67</b>	3.71
300W Challenging	<b>7.14</b>	7.38
AFLW Full	<b>1.74</b>	1.76
WFLW Full	<b>6.62</b>	6.71
WFLW No challenge	<b>4.78</b>	5.17
WFLW Blur	6.99	<b>6.94</b>
WFLW Expression	<b>7.18</b>	7.44
WFLW Illumination	<b>6.26</b>	6.36
WFLW Makeup	6.53	<b>6.39</b>
WFLW Occlusion	7.49	<b>7.20</b>
WFLW Pose	11.76	<b>11.13</b>

Table 2: Average normalized mean error on the 300W dataset in comparison to the state of the art.

Method	Common	Challenging	Full
RCPR [28]	6.18	17.26	8.35
CFAN [13]	5.50	16.78	7.69
ESR [20]	5.28	17.00	7.58
SDM [29]	5.60	15.40	7.52
ERT [30]	-	-	6.40
CFSS [31]	4.73	9.98	5.76
TCDCN [32]	4.80	8.60	5.54
LBF [33]	4.95	11.98	6.32
3DDFA [34]	6.15	10.59	7.01
3DDFA+SDM [34]	5.53	9.56	6.31
DDN [35]	-	-	5.65
RAR [36]	4.12	8.35	4.94
DeFA [37]	5.37	9.38	6.10
TR-DRN [38]	4.36	7.56	4.99
LAB (8-stack) [38]	3.42	6.98	4.12
CNN-6(Wing) [27]	3.35	7.20	4.04
CNN-6/7(Wing) [27]	3.27	7.18	4.10
$SAN_{OD}$ [16]	3.41	7.55	4.24
$SAN_{GT}$ [16]	3.34	6.60	3.98
NON	3.67	7.14	4.35
NOBB	3.71	7.38	4.43
BBO	3.34	6.77	3.93
BB	<b>3.01</b>	<b>6.51</b>	<b>3.69</b>

## 6. EVALUATION

In this section, we evaluate the proposed methodology and compare the resulting networks with the state of the art on the data sets 300W [3], AFLW [26], and WFLW [27]. First, the datasets are described. Afterward, we evaluate the impact of the validation signal to the landmark accuracy. The second experiment evaluates the impact of the different batch balancing approaches (BB and BBO) in comparison to the state of the art. In the last experiment, we evaluate the validation signal quality and its effectiveness to detect inaccurate landmarks.

**Evaluation data:** For the evaluation of our batch balancing approach and the comparison to the state of the art we used the 300W [3], AFLW [26], and WFLW [27] facial landmark dataset [3]. All values are reported with a factor of  $10^2$ . The datasets are described in the following. **300W:** The dataset is an aggregation of multiple face datasets, namely LFPW [41], HELEN [42], AFW [3] and XM2VTS [43]. 68 landmarks are annotated semi automatically [44]. The evaluation procedure is identical to [15, 16, 33]. The training set consists of the annotated images from AFW and the training subset of LFPW and HELEN (3,148 images in total). For evaluation three test sets are considered. The common test set from LFPW and HELEN (554 images). A challenging test set consisting of 135 collected iBUG face images. All images together form the full test set (689 images). We report our results in normalized mean errors (NME). This corresponds to the pixel distance between detected and annotated landmark, normalized by the pixel distance between both eye centers (equal to the inter-pupillary distance if pupils are not annotated [15, 16, 33]). **AFLW:** In our experiments, we used the AFLW-Full protocol [39]. The dataset consists of 20,000 training and 4,386 test images. Each image has 19 manually annotated landmarks and has been widely used for facial landmark localization benchmarking since it contains high pose variations ( $-90^\circ$ - $90^\circ$ ). We report our results in normalized mean errors (NME) which corresponds to the pixel distance between detected and annotated landmark, normalized by the width or height of the squared face bounding box. **WFLW:** The dataset consists of 7,500 images for training and 2,500 images for testing. Each image has 98 manually annotated facial landmarks. In addition, the test set is split up in different challenges e.g. occlusion, pose, make-up, illumination, blur, and expression. We followed the original protocol [27] and evaluated each challenge separately. The results are reported as normalized mean errors (NME). For the original protocol, the pixel distance between detected and annotated landmark is normalized using the inter-ocular distance.

**Impact of validation loss on the accuracy:** Table 1 shows the comparison of the models NON and NOBB. Both do not use the batch balancing but the model NOBB uses the proposed validation signal regression. As can be seen for the datasets 300W and AFLW the validation signal reduces the accuracy of the landmark regression. This due to the competition between the landmark position and its inaccuracy. For the dataset WFLW this is slightly different. Overall the model NON is more accurate but not for all challenges. For blur, makeup, occlusion, and pose the model NOBB outperforms NON.

**Impact of PLBB:** Table 2 shows the results on the 300W dataset. Comparing the results of NOBB to BBO and BB it can be seen that our batch balancing approach improves the results further. For BBO the results are slightly worse compared to BB because the loss distribution is less accurate. However, the computational effort for BBO is only one fifth. Comparing both (BB and BBO) to the state-of the art it can be seen that they are equally or more accurate.

The models CNN-6 [27] and CNN-6/7 [27] where proposed together with the wing loss which we are using too. The authors additionally proposed a dataset balancing strategy (PDB) which shows to be effective by comparing these two networks to our NON and NOBB models which do not use any batch balancing strategy. The comparison of the CNN-6/7 [27] to our BBO model it can be seen that the main improvement of our batch balancing strategy is in the challenging part of

Table 3: Average normalized mean error on the AFLW dataset in comparison to the state of the art.

Method	Full
CCL [39]	2.72
DAC-CSR [40]	2.27
TR-DRN [38]	2.17
LAB w/o boundary [27]	1.85
CNN-6 (Wing) [15]	1.83
CNN-6/7 (Wing) [15]	1.65
$SAN_{GT}$ [16]	1.91
NON	1.74
NOBB	1.76
BBO	1.61
BB	<b>1.56</b>

Table 4: Average normalized mean error on the WFLW dataset in comparison to the state of the art.

	Full	Blur	Expression	Illumination	Makeup	Occlusion	Pose
ESR [20]	11.13	12.20	11.47	10.49	11.05	13.75	25.88
SDM [29]	10.29	11.28	11.45	9.32	9.38	13.03	24.10
CFSS [31]	9.07	9.96	10.09	8.30	8.74	11.76	21.36
DVLN [45]	6.08	6.88	6.78	5.73	5.98	7.33	11.54
LAB [27]	<b>5.27</b>	6.32	<b>5.51</b>	<b>5.23</b>	<b>5.15</b>	6.79	10.24
NON	6.62	6.99	7.18	6.26	6.53	7.49	11.76
NO	6.71	6.94	7.44	6.36	6.39	7.20	11.13
BBO	5.86	6.01	6.28	5.53	5.68	6.28	9.67
BB	5.54	<b>5.71</b>	5.94	5.26	5.37	<b>5.94</b>	<b>9.24</b>

the 300W dataset. This is due to similar images in the training set produce higher loss values and are more likely to be in a batch.

For the LAB [38] landmark detection, there are also results using the ground truth boundary information on the 300W dataset. This combination together with the pretrained ResNet-50 [27] is compared to our approach in Table 6 where our model uses the validity signal to rule out inaccurate landmarks. Table 3 shows the comparison of our approach on the AFLW dataset. As in Table 2 it can be seen that the batch balancing approach used for CNN-6/7 [15] is effective in comparison to our models NON and NOBB. The proposed batch balancing approach used in our models BBO and BB further improves the results on the AFLW dataset wherein compared to Table 2, BBO is slightly more accurate than the network CNN-6/7 [15]. The results for LAB [27] with boundary information on the AFLW dataset will be shown in Table 6 as well as the results for the pretrained ResNet-50 [15]. Table 4 shows the results of our models in comparison to the state of the art. It can be seen that overall the LAB [27] model performs best on the WFLW dataset with which it was published together. It has to be mentioned that the landmarks on this dataset were annotated so that they always lie on the boundary of the face. This makes the LAB [27] model especially effective on this dataset. In comparison to our models it can be seen that our batch balancing approach proves to be effective especially for the challenging pose dataset. Both models BB and BBO outperform the LAB [27] approach on this challenge. In addition, blur as well as occlusions show to be more accurate using our batch balancing. This has the same reason as for the challenging iBUG images from the 300W dataset. For all approaches in the Table 4 it can be seen that blur, occlusion and pose are the most inaccurate which means that similar images produce the highest loss values.

Table 5: The results of our approach using different amounts of landmark positions based on the validation signal on the 300W dataset. For the model NON, we randomly selected 10%, 20% and 30% as a baseline since it has no validation signal.

Discard Method		Common	Challenging	Full
10%	NON	3.72	7.23	4.41
	NOBB	3.32	6.49	3.94
	BBO	2.91	6.02	3.52
	BB	2.72	5.79	3.32
20%	NON	3.72	7.23	4.40
	NOBB	2.99	5.76	3.53
	BBO	2.64	5.40	3.18
	BB	2.48	5.18	3.01
30%	NON	3.71	7.18	4.39
	NOBB	2.72	5.17	3.20
	BBO	2.42	4.87	2.90
	BB	2.27	4.66	2.74




Figure 4: Plots for the validation signal (green) to each landmark point error (red) in ascending order of the error (red). The first, second, and third row are the plots for the 300W, AFLW, and WFLW data set respectively. Each row shows the results for the models BBO and BB from left to right.

**Validation signal accuracy:** Figure 4 shows the normalized error (red) against the normalized validation signal (green) for all landmarks on the datasets 300W, AFLW, and WFLW. We sorted the values ascending based on the normalized error (red) for a better visualization. For the models BBO and BB it can be seen that the validation signal correlates with the error

on all datasets. To show the effectiveness of the validity signal to remove inaccurate landmarks, we made an additional experiment on the 300W and AFLW datasets. Table 5 shows the results for our networks by discarding different amounts of landmarks based on the validation signal. This was done by ignoring the landmarks per face with the highest validation signal values. As a baseline we also included the NON model which has no validation signal. Therefore, we selected the 10%, 20%, and 30% of the landmarks per face randomly using a uniform distribution. As can be seen in Table 5 the model NON has only a slight improvement by discarding 30% of the landmarks. The comparison to the models NOBB, BBO, and BB, which improve significantly, shows the efficiency of our validation loss which can improve all of the other state of the art approaches and is applicable to every location based regression using convolution neuronal networks. In addition, it can be seen that the highest impact is for the most challenging part of the evaluation the iBUG dataset.

**Runtime comparison:** Table 6 shows the results on the 300W and AFLW dataset together with runtime. The models ResNet-50 [15] and LAB [27] were pretrained on other data. In addition, the results for the model LAB [27] on the 300W dataset were obtained using the ground truth boundary which are the reasons why we did not include those results into the other tables. For our models the runtime was obtained on a NVIDIA GeForce GTX 1050ti, for LAB [27], ResNet-50 [15], CNN-6 [15] and CNN-6/7 [15] on a NVIDIA GeForce GTX Titan X. As can be seen in Table 6 the fastest model is the CNN-6 [15] with 400fps directly followed by our architecture (333fps). The best results are obtained by the model BB discarding 30% of the landmarks per face by using the validation signal. For all landmarks the best result is obtained by LAB [27] with the usage of the ground truth boundary.

## 7. CONCLUSION

We proposed a novel loss formulation that allows regression based point estimation alongside with a reliability estimate of the result. This improves the usability of the signal as well as the accuracy of subsequent steps that rely on it. Additionally, the signal can be used for a secondary validation of face detection. To tackle the problem of unbalanced training data in available datasets, we proposed an approach to automatically balance the training data based on the produced loss of samples. The advantage of this approach is that it does not account only for the head pose but also for other challenges and imbalances in the data.

**Acknowledgments:** Work of the authors is supported by the Institutional Strategy of the University of Tübingen (Deutsche Forschungsgemeinschaft, ZUK 63). This research was supported by an IBM Shared University Research Grant including an IBM PowerAI environment. We especially thank our partners Benedikt Rombach, Martin Mähler and Hildegard Gerhardy from IBM for their expertise and support.

## REFERENCES

- [1] A. Bulat and G. Tzimiropoulos, “How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks),” in ICCV, **1**, 4, IEEE (2017).
- [2] P. Koppen, Z.-H. Feng, J. Kittler, M. Awais, W. Christmas, X.-J. Wu, and H.-F. Yin, “Gaussian mixture 3d morphable face model,” Pattern Recognition **74**, 617–628 (2018).
- [3] X. Zhu and D. Ramanan, “Face detection, pose estimation, and landmark localization in the wild,” in CVPR, IEEE (2012).
- [4] J. Wu and M. M. Trivedi, “A two-stage head pose estimation framework and evaluation,” Pattern Recognition **41**(3) (2008).
- [5] R. Walecki, O. Rudovic, V. Pavlovic, and M. Pantic, “Copula ordinal regression for joint estimation of facial action unit intensity,” in CVPR, 4902–4910, IEEE (2016).
- [6] S. Li, W. Deng, and J. Du, “Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild,” in CVPR, 2584–2593, IEEE (2017).
- [7] J. Yang, P. Ren, D. Zhang, D. Chen, F. Wen, H. Li, and G. Hua, “Neural aggregation network for video face recognition,” in CVPR, **4**, 7, IEEE (2017).
- [8] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, “Sphereface: Deep hypersphere embedding for face recognition,” in CVPR, **1**, 1, IEEE (2017).
- [9] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in CVPR, **1**, 3, IEEE (2017).

Table 6: Results in NME (smaller is better) in comparison to the state of the art on 300W and AFLW datasets together with the runtime for each model. \* marks models pretrained on other images or using the ground truth boundary. The best results per row are bold.

Dataset	*ResNet-50 [15]	*LAB [27]	CNN-6 [15]	CNN-6/7 [15]	BB	BB (Discard 30%)
300W Full	3.60	2.99	4.10	4.04	3.69	<b>2.74</b>
300W Common	3.01	2.57	3.35	3.27	3.01	<b>2.27</b>
300W Challenging	6.01	4.72	7.20	7.18	6.51	<b>4.66</b>
AFLW Full	1.47	1.25	1.83	1.65	1.56	<b>1.24</b>
Runtime(fps)	30	16,66	<b>400</b>	170	333	333

- [10] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in ICCV, 2223–2232, IEEE (2017).
- [11] J. Yang, Q. Liu, and K. Zhang, “Stacked hourglass network for robust facial landmark localisation,” in CVPR Workshops, 2025–2033, IEEE (2017).
- [12] Y. Sun, X. Wang, and X. Tang, “Deep convolutional network cascade for facial point detection,” in CVPR, 3476–3483, IEEE (2013).
- [13] J. Zhang, S. Shan, M. Kan, and X. Chen, “Coarse-to-fine auto-encoder networks (cfan) for real-time face alignment,” in ECCV, 1–16, Springer (2014).
- [14] G. Trigeorgis, P. Snape, M. A. Nicolaou, E. Antonakos, and S. Zafeiriou, “Mnemonic descent method: A recurrent process applied for end-to-end face alignment,” in CVPR, 4177–4187, IEEE (2016).
- [15] Z. Feng, J. Kittler, M. Awais, P. Huber, and X. Wu, “Wing loss for robust facial landmark localisation with convolutional neural networks,” in CVPR, 1, 3, IEEE (2018).
- [16] X. Dong, Y. Yan, W. Ouyang, and Y. Yang, “Style aggregated network for facial landmark detection,” in CVPR, 2, 6, IEEE (2018).
- [17] D. Merget, M. Rock, and G. Rigoll, “Robust facial landmark detection via a fully-convolutional local-global context network,” in CVPR, 781–790, IEEE (2018).
- [18] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, “Active shape models-their training and application,” *Computer Vision and Image Understanding* **61**(1), 38–59 (1995).
- [19] T. F. Cootes, G. J. Edwards, and C. J. Taylor, “Active appearance models,” *PAMI* **6** (2001).
- [20] X. Cao, Y. Wei, F. Wen, and J. Sun, “Face alignment by explicit shape regression,” *Journal of Computer Vision* (2014).
- [21] Z.-H. Feng, P. Huber, J. Kittler, W. Christmas, and X.-J. Wu, “Random cascaded-regression cope for robust facial landmark detection,” *Signal Processing Letters* **22**(1), 76–80 (2015).
- [22] X. Xu and I. A. Kakadiaris, “Joint head pose estimation and face alignment framework using global and local cnn features,” in *Automatic Face and Gesture Recognition*, 642–649, IEEE (2017).
- [23] R. Ranjan, V. M. Patel, and R. Chellappa, “Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition,” *PAMI* (2017).
- [24] A. Steger and R. Timofte, “Failure detection for facial landmark detectors,” in *Asian Conference on Computer Vision*, 361–376, Springer (2016).
- [25] D. E. King, “Dlib-ml: A machine learning toolkit,” *Journal of Machine Learning Research* **10**, 1755–1758 (Dec. 2009).
- [26] M. Koestinger, P. Wohlhart, P. M. Roth, and H. Bischof, “Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization,” in ICCV Workshops, 2144–2151, IEEE (2011).
- [27] W. Wu, C. Qian, S. Yang, Q. Wang, Y. Cai, and Q. Zhou, “Look at boundary: A boundary-aware face alignment algorithm,” in CVPR, IEEE (2018).
- [28] X. P. Burgos-Artizzu, P. Perona, and P. Dollár, “Robust face landmark estimation under occlusion,” in ICCV, IEEE (2013).
- [29] X. Xiong and F. De la Torre, “Supervised descent method and its applications to face alignment,” in CVPR, 532–539, IEEE (2013).
- [30] V. Kazemi and J. Sullivan, “One millisecond face alignment with an ensemble of regression trees,” in CVPR, IEEE (2014).
- [31] S. Zhu, C. Li, C. Change Loy, and X. Tang, “Face alignment by coarse-to-fine shape searching,” in CVPR, IEEE (2015).
- [32] Z. Zhang, P. Luo, C. C. Loy, and X. Tang, “Facial landmark detection by deep multi-task learning,” in ECCV, Springer (2014).
- [33] S. Ren, X. Cao, Y. Wei, and J. Sun, “Face alignment via regressing local binary features,” *Image Processing* **25**(3) (2016).
- [34] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li, “Face alignment across large poses: A 3d solution,” in CVPR, 146–155, IEEE (2016).
- [35] X. Yu, F. Zhou, and M. Chandraker, “Deep deformation network for object landmark localization,” in ECCV, Springer (2016).
- [36] S. Xiao, J. Feng, J. Xing, H. Lai, S. Yan, and A. Kassim, “Robust facial landmark detection via recurrent attentive-refinement networks,” in ECCV, 57–72, Springer (2016).
- [37] Y. Liu, A. Jourabloo, W. Ren, and X. Liu, “Dense face alignment,” in ICCV Workshops, 1619–1628, IEEE (2017).
- [38] J. Lv, X. Shao, J. Xing, C. Cheng, and X. Zhou, “A deep regression architecture with two-stage re-initialization for high performance facial landmark detection,” in CVPR, IEEE (2017).
- [39] S. Zhu, C. Li, C.-C. Loy, and X. Tang, “Unconstrained face alignment via cascaded compositional learning,” in CVPR, IEEE (2016).
- [40] Z.-H. Feng, J. Kittler, W. Christmas, P. Huber, and X.-J. Wu, “Dynamic attention-controlled cascaded shape regression exploiting training data augmentation and fuzzy-set sample weighting,” in CVPR, IEEE (2017).
- [41] P. N. Belhumeur, D. W. Jacobs, D. J. Kriegman, and N. Kumar, “Localizing parts of faces using a consensus of exemplars,” *PAMI* **35**(12), 2930–2940 (2013).
- [42] V. Le, J. Brandt, Z. Lin, L. Bourdev, and T. S. Huang, “Interactive facial feature localization,” in ECCV, 679–692, Springer (2012).
- [43] K. Messer, J. Matas, J. Kittler, J. Luettin, and G. Maitre, “Xm2vtsdb: The extended m2vts database,” in *Audio and Video-based Biometric Person Authentication*, **964**, 965–966 (1999).
- [44] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, “A semi-automatic methodology for facial landmark annotation,” in CVPR Workshops, 896–903, IEEE (2013).
- [45] W. Wu and S. Yang, “Leveraging intra and inter-dataset variations for robust face alignment,” in CVPR Workshops, 150–159, IEEE (2017).



### **AUTHORS' BACKGROUND**

Name	Title	Research Field	Personal website
Wolfgang Fuhl	Dr.	Computer science	<a href="http://www.ti.uni-tuebingen.de/Wolfgang-Fuhl.1651.0.html">http://www.ti.uni-tuebingen.de/Wolfgang-Fuhl.1651.0.html</a>
Enkelejda Kasneci	Jun. Prof.	Computer science	<a href="http://www.ti.uni-tuebingen.de/Prof-Dr-Enkelejda-Kasneci.956.0.html">http://www.ti.uni-tuebingen.de/Prof-Dr-Enkelejda-Kasneci.956.0.html</a>