

Sub-pixel face landmarks using heatmaps and a bag of tricks

Samuel W. F. Earp* Aubin Samacoits Sanjana Jain Pavit Noinongyao
Siwa Boonpunmongkol

Sertis Vision Lab[†]

Abstract

Accurate face landmark localization is an essential part of face recognition, reconstruction and morphing. To accurately localize face landmarks, we present our heatmap regression approach. Each model consists of a MobileNetV2 backbone followed by several upscaling layers, with different tricks to optimize both performance and inference cost. We use five naïve face landmarks from a publicly available face detector to position and align the face instead of using the bounding box like traditional methods. Moreover, we show by adding random rotation, displacement and scaling—after alignment—that the model is more sensitive to the face position than orientation. We also show that it is possible to reduce the upscaling complexity by using a mixture of deconvolution and pixel-shuffle layers without impeding localization performance. We present our state-of-the-art face landmark localization model (ranking second on The 2nd Grand Challenge of 106-Point Facial Landmark Localization validation set). Finally, we test the effect on face recognition using these landmarks, using a publicly available model and benchmarks.

1 Introduction

Landmark localization is the task of finding the positions of predefined key points in an image. One of its main applications is Human Pose Estimation (HPE), a major research topic for computer vision with applications such as action recognition, people tracking and sports analytics. Landmark localization is a challenging problem due to variations in pose, illumination and occlusion (Dibeklioglu et al., 2008). Early solutions based on the pictorial structure model (Fischler and Elschlager, 1973) define objects as a set of landmarks connected in a deformable structure. These methods combine an appearance model to localize the landmarks with a predefined structural model of the object. This approach produces landmarks that satisfy both appearance features and spatial configuration of the object (Eichner and Ferrari, 2009; Felzenszwalb et al., 2010; Yang and Ramanan, 2013).

Deep learning algorithms have dramatically improved the performance of HPE. Toshev and Szegedy (2014) proposed one of the first methods, where the HPE is formulated as a regression problem over the normalized joint coordinates. Although predicting joint coordinates directly from images is very difficult, numerous improvements have been discovered. For example, Sun et al. (2017) combined a regression approach with prior information

about the human body structure to achieve more robust results. Tompson et al. (2014b) approached HPE as a heatmap regression problem, where feature maps are up-scaled to generate spatial probability maps, with pixel values corresponding to the probability of the landmark residing at a particular position. These heatmaps are created by convolving the location of each joint with a 2D Gaussian (Tompson et al., 2014a,b). This novel approach outperformed the existing state of the art methods (Dantone et al., 2013; Toshev and Szegedy, 2014) on both FLIC and LSP datasets. Wei et al. (2016) also presented a heatmap regression network consisting of different stages. Each stage receives the original image and the heatmap from the previous stage. This approach outperforms Tompson et al. (2014b) on the FLIC, LSP and MPII (Andriluka et al., 2014) datasets. A heatmap based stacked hourglass network comprised of stacked modules was proposed by Newell et al. (2016), each with pooling layers followed by upsampling layers. They achieve between 0.4 – 4.6 per cent improvement on different joints in comparison to Wei et al. (2016). Xiao et al. (2018) use a ResNet (He et al., 2016) backbone followed by deconvolutional layers to generate a set of heatmaps representing the probability map of each joint. Whereas, Sun et al. (2019) combine multiple sub-nets and capture information from different scales, im-

*Email: searp@sertiscorp.com

[†]597/5 Sukhumvit Road, Watthana, Bangkok, 10110, Thailand

proving the mAP on COCO by 3.3 mAP compared to [Xiao et al. \(2018\)](#). Recent studies focus on the target heatmaps and the accuracy of the coordinates retrieved from them. [Zhang et al. \(2020\)](#) extract joint coordinates from their respective heatmap using sub-pixel Gaussian fitting, resulting in a performance increase for different state of the art models on both COCO ([Lin et al., 2014](#)) and MPII datasets. Another direction of amelioration, presented in [Huang et al. \(2020\)](#), consists of exploring the effect of different training schedules on various HPE methods. The authors found schedules that use information dropping augmentations (e.g. random erasing) after a certain number of epochs yield the best performance. They report a 0.6 AP increase on the COCO dataset when they apply their suggested training schedule to the state of the art HRNet-W48.

Facial landmark localization focuses on predicting pre-defined face key points (e.g. contours of the face, eyes, nose, mouth, eyebrows, etc.) and is used for a wide variety of face-related tasks. Face recognition uses these landmarks for alignment, this has been an intrinsic part of these systems for over a decade ([An et al., 2020](#); [Deng et al., 2018](#); [Kumar et al., 2009](#); [Liu et al., 2017](#); [Parkhi et al., 2015](#); [Schroff et al., 2015](#); [Wang et al., 2018](#); [Wolf et al., 2010](#)). They are also used in face reconstruction where 3D annotations are not available (e.g. [Dou et al., 2017](#); [Feng et al., 2018](#); [Roth et al., 2015](#)), as well as, face animation (e.g. [Cao et al., 2013](#)), emotion classification ([Day, 2016](#); [Munasinghe, 2018](#)), detecting synthesized faces ([Yang et al., 2019](#)) and facial action unit detection ([Hinduja and Canavan, 2020](#)).

Early face landmark localization focused on using statistical models or component detectors. One statistical method is the Active Shape Model (ASM), proposed by [Cootes et al. \(1995\)](#). This method focuses on generating a mean face shape model to determine the best position for the facial landmarks. ASM computes the initial face shape using the shape model, then fine-tunes the shape parameters for the image by examining the region around each landmark. Similarly to ASM, [Cootes et al. \(1998\)](#) focus on generating a global face shape model; while also attempting to incorporate texture information using an appearance model. The appearance and shape parameters get fitted by reducing the difference between the image and its synthetically generated counterpart. Some researches also explored training patch-based detectors or component detectors to predict each landmark on local patches or anatomical components on the face image, respectively ([Amberg and Vetter, 2011](#); [Belhumeur et al., 2013](#); [Efraty et al., 2011](#); [Liang et al., 2008](#); [Zhu and Ramanan, 2012](#)). These approaches require constraints on the face shape to obtain the best landmark configuration due to the lack of global contextual information of the face.

Recently, Convolutional Neural Networks (CNNs) have supplanted classical approaches due to their ability to extract contextual information. Generally speaking, there are two widely adopted approaches; coordinate regression and heatmap fitting—just like HPE (e.g. [Tompson et al., 2014a,b](#); [Xiao et al., 2018](#); [Zhang et al., 2020](#)).

For coordinate regression, a dense layer is added at the end of the CNN to predict the coordinates of each landmark. Among the first of these approaches, [Sun et al. \(2013\)](#) fit a three-level cascaded CNN to localize five facial landmarks: the corners of the eyes, the tip of the nose and corners of the mouth. The first level estimates the coordinates then the subsequent levels fine-tune them. While outperforming classical methods, the cascaded CNN architecture is complicated. The first level contains three CNNs, each predicting a different subset of landmark, and the following two levels each have ten CNNs. [Zhou et al. \(2013\)](#) proposed a four-level cascaded coarse-to-fine CNN, where the first level estimates two bounding boxes for each face, while the following layers predict the landmarks for the bounding boxes with refinement on the inner parts of the face (e.g., eyes, nose, and mouth). Both sets of landmarks are then combined to obtain the final positions. This method works well in challenging conditions, such as high pose variance, poor illumination, and occlusion. The authors achieve a significant performance improvement on the 300W challenge ([Sagonas et al., 2013](#); [Sagonas et al., 2016](#)) compared to the baseline. Similarly, [Zhang et al. \(2014\)](#) proposed a coarse-to-fine cascade of stacked auto-encoders, where the first level performs initial landmarks estimation on lower resolution face image and subsequent layers refine these using higher-resolution local patches for each landmark. Their approach outperforms both [Sun et al. \(2013\)](#) and [Xiong and De la Torre \(2013\)](#), while being significantly faster. More recently, multi-task learning has been utilized to reduce architectural complexity. By combining: the landmarks with the pose, facial expression, gender, and other attributes, [Zhang et al. \(2016\)](#) built a lighter CNN, which is robust to pose and occlusion. Similarly, [Ranjan et al. \(2016\)](#) proposed using multi-task learning on related tasks such as face detection, landmarks localization, pose estimation, and gender classification to achieve state-of-the-art performance on individual tasks, including a 0.42 per cent reduction in error for facial landmarks localization.

Most regression-based approaches suffer from spatial information loss due to the compression of feature maps before the fully-connected layers. This shortcoming has inspired researchers to leverage the encoder-decoder architecture and propose heatmap based approaches. [Kowalski et al. \(2017\)](#) proposed using landmark heatmaps for their alignment network to transfer information between different network stages. Instead of using local patches leading to local minima, their system works on entire images, handling large pose variations and achieving a 72 per cent reduction in the failure rate on the 300W dataset. Furthermore, [Mahpod et al. \(2018\)](#) proposed a cascaded CNN architecture comprising two CNNs followed by two cascaded subnetworks: a heatmap based network and a regression-based network to refine the heatmap localizations. Although these methods have led to performance increases, they rely on generating large heatmaps; which can suffer from increased post-processing complexity. The ground-truth heatmaps are often dominated by background pixels with tiny positive-valued regions at the landmarks, leading

to slower convergence during training. Xiong et al. (2020) proposed using a quasi-Gaussian distribution to represent ground-truth landmark positions as vectors, addressing the foreground-background imbalance. They also convert the output heatmaps into vectors, which incorporates spatial information and reduces sensitivity. The authors demonstrate that their approach leads to better convergence, reduced post-processing complexity, and achieves state-of-the-art performance on multiple evaluation datasets, including ranking second on the JD-landmark challenge (Liu et al., 2019).

We organize the paper as follows. Section 2 presents some of the most recent work that we will draw upon. Section 3 outlines our training and evaluation methods and our baseline parameters. Section 4 compares the results using different techniques to increase and reduce both landmark accuracy and computational complexity, respectively. Section 5 presents our final models, the additional tricks we employed and investigates the potential impact on facial recognition performance. Section 5 presents our conclusions.

2 Related Work

2.1 Heatmap fitting

As discussed in the previous section, coordinate regression and heatmap based approaches have been used widely for landmark localization. However, regression-based approaches (Sun et al., 2013; Zhang et al., 2014) tend to lead to spatial information loss. This reason, combined with the recent success of Kowalski et al. (2017) and Xiong et al. (2020), leads us to employ the heatmap fitting approach.

2.2 Sub-pixel inference

The last step of landmark localization consists of extracting a set of coordinates from the estimated heatmaps. The simplest approach is to use the argmax of the heatmaps:

$$(x, y)_i = \underset{m < w, n < h}{\operatorname{argmax}} (H_i(m, n)), \quad (1)$$

where $H_i(m, n)$ is the estimated value of the i^{th} landmark heatmap in (m, n) and w, h are the width and height of H_i , respectively. This approach does not allow for sub-pixel localization, as the accuracy is limited by the resolution of the heatmap. A simple way to reduce this effect is to modify Equation 1 using the gradient of the heatmap (Xiao et al., 2018),

$$(x, y)_i = \underset{m < w, n < h}{\operatorname{argmax}} (H_i(m, n)) + c * \frac{\partial H}{\partial x \partial y}, \quad (2)$$

where c is a correction coefficient applied to the gradient of the heatmap around (x, y) from Equation 1. This approach incorporates the region around the heatmap peak, shifting the position by a sub-pixel distance using the gradient.

Zhang et al. (2019) use the prior Gaussian information of the heatmaps to increase the accuracy of the predicted joint coordinates. As the target heatmaps are generated

by a Gaussian convolution, the model is incentivized to also produce Gaussian probability maps. These estimated heatmaps can be approximated by,

$$G(X, \mu, \Sigma) = \frac{\exp(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu))}{(2\pi)^{\frac{1}{2}} |\Sigma|^{\frac{1}{2}}}, \quad (3)$$

where X is the location and μ is the ground truth coordinate of the landmark. The $(x, y)_i$ is computed for each landmark to minimize the error between the Equation 3 applied to $(x, y)_i$ and the predicted heatmap H_i . This approach results in sub-pixel localization and hence increases the accuracy of the landmarks. Due to this performance increase, we will employ this method for our post-processing.

2.3 Pixel shuffle

Single image super-resolution (SISR) focuses on generating an image with r times higher resolution. Previous approaches use interpolation (e.g., bicubic interpolation) to perform upsampling either at the first layer of the network (Chen and Pock, 2017; Dong et al., 2016; Wang et al., 2015) or gradually through the network (Osendorfer et al., 2014); increasing the computational complexity. Shi et al. (2016) propose efficient sub-pixel convolutional layers that learn an effective upsampling filter instead of using a pre-defined interpolation. These layers contain convolutions and pixel rearrangement and are known as pixel shuffle layers. SISR can be performed by passing the original image through a three-layer convolutional network, where the output feature maps have the same width and height as the input image. Afterwards, the r^2 feature maps are rearranged to generate a feature map of the desired size. Their proposed approach outperforms Dong et al. (2016) on all and Chen and Pock (2017) on most benchmark datasets while running at least ten times faster on these datasets for SISR. We leverage the pixel-shuffle layers to build an efficient model architecture; due to the reduction in inference time and computational complexity.

3 Methodology

3.1 Training dataset

Earlier face landmark datasets, including MULTI-PIE (Gross et al., 2008), 300W (Sagonas et al., 2013; Sagonas et al., 2016), LFW (Huang et al., 2008a) and the Menpo benchmark (Deng et al., 2019b), mainly focus on 68 landmarks. These landmarks lack some of the core features of the face. For example, the lower eyebrow and border/wings of the nose are all ignored. To correct this Liu et al. (2019) presented the 106-point JD-landmark dataset. The dataset consists of 11,393, 2,000, and 2,000 images for the training, validation and test sets, with large variations in the pose. The updated training dataset was released as part of the 2nd Grand Challenge of 106-Point Facial Landmark Localization (JD-landmark-2)³.

We preprocess the training and validation images using the publicly available ResNet50 face detector from

³<https://flc-icpr2020.github.io>

Deng et al. (2019a). From this model, we get the bounding box and the five naïve landmarks for the eyes, nose and mouth of the centre most face. We calculate the affine transformation between these five landmarks and a set of reference landmarks and apply this transform to the image and the ground truth landmarks. We choose the reference landmarks such that the face is in the centre of the image with the eyes and mouth horizontal. Figure 1 compares the example images from the validation set which are presented in Liu et al. (2019) and our aligned versions. We train our model to predict using the aligned image then recover the original landmark positions by applying the inverse affine transform to the predicted landmarks.



Figure 1: Left: the original images, right: images aligned using the landmarks from the face detector.

3.2 Evaluation

Following Liu et al. (2019) we calculate the normalized mean Euclidean error (NME) across all landmarks,

$$\text{NME} = \frac{1}{N} \sum_{k=1}^n \frac{\|y_k - \hat{y}_k\|_2}{d}, \quad (4)$$

where n , N , y and \hat{y} are the number of landmarks, number of images, ground truth landmarks and predicted landmarks, respectively. The normalization factor d is given by $\sqrt{\text{bbox}_w \text{bbox}_h}$ where bbox_w and bbox_h are the width and height of the ground truth bounding box, respectively. We then compute the cumulative error distribution (CED) curve for images with NME less than 0.08. We define our final evaluation metrics to be the area-under-the-curve (AUC) for the CED curve, the failure rate where the NME is above 0.08, and the average NME; please refer to Liu et al. (2019) for a more detailed description.

For the evaluation dataset, we use the publicly available⁴ JD-landmark test set. These 2000 test images are in the JD-landmark-2 training dataset. Therefore, we remove these images from JD-landmark-2 to create our training dataset.

3.3 Baseline settings

To train our models we use MXNET⁵ (Chen et al., 2015) and GLUONCV⁶ (Guo et al., 2020). We train our models using ADAM (Kingma and Ba, 2014) for 40 epochs with a

learning rate of 0.001 decreasing by an order magnitude at epochs 20 and 30. All training is done on a single NVIDIA Titan X, using a batch size of 16. During training, the training dataset is randomly augmented using both the AlexNet style PCA colour augmentation (Krizhevsky et al., 2012) with $\sigma = 0.05$ and using the GLUONCV random colour jitter augmentation with a value of 0.4 for brightness, contrast and saturation.

4 Results

4.1 Baseline results

For our baseline, we train MobileNetV2 (Sandler et al., 2018) with a series of four deconvolutions—all with 256 filters and a stride of two—each followed by a batch normalization layer and a ReLU activation, and finally a 2D convolution with 106 filters to generate the heatmaps. We set the input size to $192 \times 192 \times 3$, and the generated heatmaps have a size of 96×96 . Table 1 presents our baseline model results compared to the top three submissions from Liu et al. (2019)⁴. We also experimented with both MobileNet (Howard et al., 2017) and MobileNetV3 (Howard et al., 2019) but found they performed significantly worse than MobileNetV2.

Model	AUC (%)	NME (%)
Baidu-VIS [†]	84.01	1.31
Xiong et al. (2020)	83.34	1.35
USTC [†]	82.68	1.41
VIC [†]	82.22	1.42
MNetV2 _{0.25}	82.20	1.42

Table 1: This table presents the performance of our baseline models compared to the top three submissions from Liu et al. (2019)[†] and Xiong et al. (2020).

4.2 Downsizing

Pixelshuffle

To test the performance using pixel-shuffle layers (Shi et al., 2016), we replace our baseline deconvolution layers with four upsampling pixel shuffle blocks. Each block consists of a 2D convolution, a batch normalization, a ReLU activation, and a final pixel shuffle layer. We generate the heatmaps by adding a 2D convolutional layer before the final block. Our results are presented in Table 2.

Strategy	AUC (%)	GFLOPS	Size (MB)
SSSS	79.18	0.56	6.12
DDDD	82.20	3.50	18.26

Table 2: This table presents a comparison between pixel-shuffle and deconvolution upsampling strategies. All models are based on a MobileNetV2_{0.25} backbone and are trained with the baseline settings described in Section 4.1.

⁴<https://facial-landmarks-localization-challenge.github.io>

⁵<https://github.com/apache/incubator-mxnet>

⁶<https://github.com/dmlc/gluon-cv>

We find that using pixel-shuffle layers give worse performance than deconvolutions, however, it does significantly reduce the model size and number of FLOPS. We also experimented using different numbers of upsampling layers but found four to give us the best results for this backbone.

Intermittent shuffling

In the previous section, we compared deconvolution layers and pixel shuffle layers, finding that deconvolutions give better results but are less efficient. Therefore, we propose a new approach; we combine deconvolution and pixel shuffle layers. Table 3 shows the results from different layer arrangements, where S and D in the strategy column denote pixel-shuffle and 2D deconvolution layers, respectively. We notice two things:

1. having the deconvolution layers near the end improves performance,
2. having pixel shuffle layers after deconvolution layers degrades performance.

For example, DDSS and DSDS are comparable in FLOPS to SSSD and SDSD, however, the latter have far better performance. Comparing DDDD to SDSD we can see that incorporating pixel-shuffle layers into the upsampling strategy can almost half the number of FLOPS, with only a slight reduction in performance.

Strategy	AUC (%)	GFLOPS	Size (MB)
SSSS	79.18	0.56	6.12
DSSS	80.72	0.64	9.56
DDSS	80.66	0.96	12.00
DSSD	81.55	1.10	10.08
DSDS	80.92	1.73	11.69
DDDS	80.81	2.25	16.44
SSSD	81.71	1.02	6.64
SDSD	81.90	1.29	8.76
SSDD	81.10	2.90	10.08
SDDD	81.89	3.36	13.51
DDDD	82.20	3.50	18.26

Table 3: This table presents a comparison between our stacked pixel-shuffle and deconvolution upsampling strategies. All models are based on a MobileNetV2_{0.25} backbone and are trained with the baseline settings described in Section 4.1.

Upsampling filters

An easy way to reduce the FLOPS is to reduce the number of filters in the deconvolution layers. We found that reducing the number of filters to 128, using the DDDD upsampling strategy, surprisingly gives us a small boost in AUC of ~ 0.05 with no increase in failure rate or NME.

⁷<https://github.com/apache/tvm>

4.3 Upsizing

By reducing the number of filters and using intermittent shuffling we have drastically reduced the FLOPS of our model. Therefore, we can increase the size of our backbone; hopefully resulting in better feature extraction. Table 4 presents our results from the final MobileNet models. We find that using just three upsampling layers consisting of one pixel-shuffle layer and two deconvolutions gives the best performance.

Backbone	Strategy	AUC (%)	GFLOPS
MNetV2 _{1.0}	SSD	84.56	0.32
MNetV2 _{1.0}	SDD	84.67	0.43
MNetV2 _{1.0}	SSSD	84.43	0.55
MNetV2 _{1.0}	SDSD	84.46	0.61

Table 4: This table shows the results for our stacked pixel shuffle and deconvolution models. All these models are trained using the same baseline settings described in Section 4.1.

5 Final models

5.1 Opening our bag of tricks

To get our final models we use four more tricks. First, we extend the training dataset by including the horizontal flip of each image and their corresponding landmarks; this resulted in a performance increase of ~ 0.7 in AUC(%). Second, we run inference on a batch containing both the original image and a horizontally flipped copy. We then average these two predictions, resulting in an AUC(%) gain of ~ 0.6 . We also tried to stack the two heatmaps generated for each landmark, then predict using the combined heatmap. However, we found that this performed almost the same or worse than not using the flipped image. Third, we employed the random erasing strategy described by Huang et al. (2020), resulting in a further AUC(%) increase of ~ 0.2 . Last, we incorporate random rotation ($\pm 1^\circ - \pm 15^\circ$), scaling ($1 \pm 0.05 - \pm 0.2$) and repositioning ($\pm 5 - \pm 20$ pixels) to each image and it’s corresponding landmarks during training. We found that both random scaling and repositioning result in ~ 1.0 AUC(%) performance loss. Yet, random rotation resulted in a further AUC(%) increase of ~ 0.2 . This difference in performance gain/loss indicates that the model is more receptive to consistent face position and scale than rotation.

5.2 Results on JD-landmark datasets

Table 5 reports our final results on the JD-landmark test set and the JD-landmark-2 validation set. All of our models perform exceptionally well on the JD-landmark test set, surpassing all the models on the leaderboard⁴. Our MNetV2_{1.0} is able to rank second on the JD-landmark-2 validation set³, while our ResNet18 and ResNet50 are reported to show the upper limit of our approach.

	Backbone	AUC (%)	Failure Rate (%)	NME (%)
Challenge 1 ⁴	Our-ResNet50	87.06	0.00	1.03
	Our-ResNet18	86.87	0.00	1.07
	Our-MNetV2 _{1.0}	86.49	0.00	1.08
	†Baidu-VIS	84.01	0.10	1.31
	Xiong et al. (2020)	83.34	0.10	1.35
	‡USTC	82.68	0.05	1.41
Challenge 2 ³	Our-ResNet50	81.64	0.15	1.47
	Our-ResNet18	81.42	0.05	1.49
	†Sogou AI	80.96	0.10	1.52
	Our-MNetV2 _{1.0}	80.81	0.05	1.54
	‡OPPO Research Institute	80.46	0.00	1.56

Table 5: Top: the evaluation results for different backbones on the JD-landmark test set, bottom: the results for each approach on the JD-landmark-2 validation set; † and ‡ denote the first and second place entries to each challenge.

5.3 Inference time

We report the inference times of our models in 6. To optimize our models we use the Apache TVM⁷ compiler framework to optimize our models (Chen et al., 2018). Our MNetV2_{1.0} can achieve an inference time of 65.70 ms on one core of a desktop CPU (intel i5-9300H) and 4.2 ms on a desktop GPU (NVIDIA GTX 1650). We also test the inference time of our MNetV2_{1.0} on a mobile device, a Pocophone F1 with a Snapdragon 845 chipset (Kryo 385 CPU and Adreno 630 GPU), achieving an inference time of 67.3 ms.

	MNetV2 _{1.0}	ResNet18	ResNet50
GPU (ms)	4.20	4.83	8.71
CPU (ms)	65.70	131.5	241.5
GFLOPS	0.43	1.13	3.23
Size (MB)	12.16	45.56	94.10

Table 6: Inference times for our models measured on an NVIDIA GTX 1650 and an intel i5-9300H, for GPU and CPU, respectively.

5.4 Effect on face recognition

Face recognition systems usually contain four stages: face detection, alignment, embedding and then distance calculation. In this section, we explore using our landmark localization model to aid the alignment process. After we detected the face, we then pass the image to our landmark localization model to obtain the 106 facial landmarks. Following the conventional five landmarks alignment (An et al., 2020; Deng et al., 2018; Liu et al., 2017; Schroff et al., 2015; Wang et al., 2018; Wolf et al., 2010), we take a subset of five landmarks containing the centres of each eye, the tip of the nose and corners of the mouth.

We evaluate the performance using two different benchmarks for face recognition. For the first benchmark, we report the face verification accuracy on four public datasets, LFW, CFP-FP, CALFW, CPLFW (Huang et al.,

2008b; Sengupta et al., 2016; Zheng and Deng, 2018; Zheng et al., 2017). For the second benchmark we follow the IJB-B and IJB-C protocol from ArcFace (see An et al., 2020; Deng et al., 2018, for more details)⁸, incorporating both detector score and feature normalization. To detect the faces we use our MNetV2 model from Earp et al. (2019), a publicly available ResNet50 model from (Deng et al., 2019a)⁸. For the face embedding network, we use the pretrained LResNet100E-IR from Insightface⁸ and for the landmark model (referred to as 106p) we use our MNetV2_{1.0} presented in Section 5.2. Our results are shown in Table 7. We give the face detector backbones in the first column, +106p indicates that we replace the detector landmarks using the landmark model.

For LFW and CALFW we see little performance change. However, for CFP-FP and CPLFW we see accuracy gains of 0.23 and 0.44 per cent for MNetV2_{1.0} and 0.07 and 0.2 per cent for ResNet50, respectively. For IJB-B we report the True Acceptance Rate (TAR) at a False Acceptance Rate (FAR) of $1e-4$, finding an improvement of 0.09 per cent and 0.08 per cent for MNetV2_{1.0} and ResNet50, respectively. Similarly for IJB-C we find a TAR (FAR = $1e-4$) improvement of 0.13 per cent and 0.06 per cent for MNetV2_{1.0} and ResNet50, respectively. We also report the combined inference cost of the face detector and the localization model. The face detection model is optimized using the same procedure as the landmark model (see Section 5.3).

6 Conclusions

We have shown that replacing all the deconvolution layers with pixel-shuffle layers reduces the total number of FLOPS, but this approach significantly impacts the model’s performance. Therefore, we propose the stacked pixel-shuffle and deconvolution upsampling strategy reducing the total number of FLOPS with only a small impact on localization performance. We applied the random erasing strategy proposed by Huang et al. (2020) to help improve landmark localization for human joint localization, find-

⁸<https://github.com/deepinsight/insightface>

	LFW	CFP-FP	CALFW	CPLFW	IJB-B	IJB-C	CPU _{total} (ms)
MNetV2 _{1.0}	99.87	98.36	95.82	93.58	94.81	96.14	62.66
+ 106p	99.87	98.59	95.80	94.02	94.90	96.27	128.4
ResNet50	99.85	98.63	95.78	93.70	94.85	96.23	261.2
+ 106p	99.87	98.70	95.77	93.90	94.93	96.29	326.9

Table 7: Face recognition results on benchmark datasets with two different detector backbones: MNetV2_{1.0} and ResNet50, where +106p indicates the landmarks have been adjusted by our MNetV2_{1.0} landmark model.

ing that this approach also affective for facial landmark localization. To test the model’s dependence on the initial alignment, we randomly rotated, scaled and repositioned each image during training. We found that both random scaling and repositioning were detrimental to model performance. On the other hand, random rotation increased the model’s performance. Indicating that the network benefits more from the consistent positioning of our preprocessing method; and is not relying on the alignment.

We present three final models, each with a different backbone, MobileNetV2, ResNet18 and ResNet50. Our smallest model performs exceptionally well on the JD-landmark-2 validation set, ranking second place only 0.15 per cent AUC behind first place. This MobileNetV2 model uses just 43 per cent of the allowed GFLOPS and runs in 65.7 ms on an intel i5-9300 CPU. Our ResNet18 surpasses the first place entry by 0.46 per cent in AUC and 0.03 per cent in NME, while only falling 13 per cent above the GFLOPS limit and only 13 per cent slower than the MobileNetV2 on GPU. We also present our ResNet50 version, which only has a slight lead in performance over ResNet18; but is significantly slower.

We have shown adjusting the face landmarks using our landmark localization model leads to improvements in face recognition performance. We see accuracy gains on CFP-FP (0.07 – 0.23%) and CPLFW (0.2 – 0.44%), and also TAR (FAR= $1e-4$) for IJB-B (0.08 – 0.09%) and IJB-C (0.06 – 0.13%). However, we see little change in recognition performance on LFW and CALFW. One possible cause for this discrepancy is the emphasis both CFP-FP and CPLFW place on using pairs of different poses. Therefore, we conclude that either the recognition model is more particular about the alignment of side faces; or that the detector is insufficient to align some of these faces.

Although our models can achieve high accuracy and fast inference speeds, post-processing the landmarks is still expensive. Future works should direct their attention towards trying to reduce this computational overhead. For example, Xiong et al. (2020) presented a method where the heatmaps are replaced by vectors, drastically reducing post-processing and still achieving high landmark accuracy. Further research should be continued, in this vein, to break away from the expensive post-processing while maintaining high accuracy.

Acknowledgements

We would like to thank Yinglu Liu from JD AI Research for evaluating our models after the deadline had passed. We would also like to thank Jeff Hnybida & Moses Kim for useful conversations during the writing of this paper.

References

- B. Amberg and T. Vetter. Optimal landmark detection using shape models and branch and bound. pages 455–462, 11 2011. doi: 10.1109/ICCV.2011.6126275.
- X. An, X. Zhu, Y. Xiao, L. Wu, M. Zhang, Y. Gao, B. Qin, D. Zhang, and Y. Fu. Partial FC: Training 10 Million Identities on a Single Machine. *arXiv e-prints*, art. arXiv:2010.05222, Oct. 2020.
- M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3686–3693, 2014. doi: 10.1109/CVPR.2014.471.
- P. Belhumeur, D. Jacobs, D. Kriegman, and N. Kumar. Localizing parts of faces using a consensus of exemplars. *IEEE transactions on pattern analysis and machine intelligence*, 35:2930–40, 12 2013. doi: 10.1109/TPAMI.2013.23.
- C. Cao, Y. Weng, S. Lin, and K. Zhou. 3d shape regression for real-time facial animation. *ACM Trans. Graph.*, 32(4), July 2013. ISSN 0730-0301. doi: 10.1145/2461912.2462012. URL <https://doi.org/10.1145/2461912.2462012>.
- T. Chen, M. Li, Y. Li, M. Lin, N. Wang, M. Wang, T. Xiao, B. Xu, C. Zhang, and Z. Zhang. MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems. *arXiv e-prints*, art. arXiv:1512.01274, Dec. 2015.
- T. Chen, T. Moreau, Z. Jiang, L. Zheng, E. Yan, M. Cowan, H. Shen, L. Wang, Y. Hu, L. Ceze, C. Guestrin, and A. Krishnamurthy. TVM: An Automated End-to-End Optimizing Compiler for Deep Learning. *arXiv e-prints*, art. arXiv:1802.04799, Feb. 2018.
- Y. Chen and T. Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1256–1272, 2017. doi: 10.1109/TPAMI.2016.2596743.

- T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models-their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995. ISSN 1077-3142. doi: <https://doi.org/10.1006/cvui.1995.1004>. URL <http://www.sciencedirect.com/science/article/pii/S1077314285710041>.
- T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 484–498. Springer, 1998.
- M. Dantone, J. Gall, C. Leistner, and L. van Gool. Human pose estimation using body parts dependent joint regressors. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3041–3048, Portland, OR, USA, June 2013. IEEE.
- M. Day. Exploiting facial landmarks for emotion recognition in the wild. *CoRR*, abs/1603.09129, 2016. URL <http://arxiv.org/abs/1603.09129>.
- J. Deng, J. Guo, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *CoRR*, abs/1801.07698, 2018. URL <http://arxiv.org/abs/1801.07698>.
- J. Deng, J. Guo, Z. Yuxiang, J. Yu, I. Kotsia, and S. Zafeiriou. Retinaface: Single-stage dense face localisation in the wild. In *arxiv*, 2019a.
- J. Deng, A. T. Roussos, G. Chrysos, E. Ververas, I. Kotsia, J. Shen, and S. Zafeiriou. The menpo benchmark for multi-pose 2d and 3d facial landmark localisation and tracking. *International Journal of Computer Vision*, 127, 06 2019b. doi: 10.1007/s11263-018-1134-y.
- H. Dibeklioglu, A. A. Salah, and L. Akarun. 3d facial landmarking under expression, pose, and occlusion variations. In *2008 IEEE Second International Conference on Biometrics: Theory, Applications and Systems*, pages 1–6, 2008. doi: 10.1109/BTAS.2008.4699324.
- C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2016. doi: 10.1109/TPAMI.2015.2439281.
- P. Dou, S. K. Shah, and I. A. Kakadiaris. End-to-end 3D face reconstruction with deep neural networks. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-January:1503–1512, apr 2017. URL <http://arxiv.org/abs/1704.05020>.
- S. W. F. Earp, P. Noinongyao, J. A. Cairns, and A. Ganguly. Face Detection with Feature Pyramids and Landmarks. *arXiv e-prints*, art. arXiv:1912.00596, Dec. 2019.
- B. A. Efraty, M. Papadakis, A. Proffitt, S. Shah, and I. A. Kakadiaris. Facial component-landmark detection. In *2011 IEEE International Conference on Automatic Face Gesture Recognition (FG)*, pages 278–285, 2011. doi: 10.1109/FG.2011.5771411.
- M. Eichner and V. Ferrari. Better appearance models for pictorial structures. 23, 01 2009. doi: 10.5244/C.23.3.
- P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2010. doi: 10.1109/TPAMI.2009.167.
- Y. Feng, F. Wu, X. Shao, Y. Wang, and X. Zhou. Joint 3D Face Reconstruction and Dense Alignment with Position Map Regression Network. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11218 LNCS:557–574, mar 2018. URL <http://arxiv.org/abs/1803.07835>.
- M. A. Fischler and R. A. Elschlager. The representation and matching of pictorial structures. *IEEE Transactions on Computers*, C-22(1):67–92, 1973. doi: 10.1109/T-C.1973.223602.
- R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. In *2008 8th IEEE International Conference on Automatic Face Gesture Recognition*, pages 1–8, 2008. doi: 10.1109/AFGR.2008.4813399.
- J. Guo, H. He, T. He, L. Lausen, M. Li, H. Lin, X. Shi, C. Wang, J. Xie, S. Zha, A. Zhang, H. Zhang, Z. Zhang, Z. Zhang, S. Zheng, and Y. Zhu. Gluoncv and gluonnlp: Deep learning in computer vision and natural language processing. *Journal of Machine Learning Research*, 21(23):1–7, 2020. URL <http://jmlr.org/papers/v21/19-429.html>.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90.
- S. Hinduja and S. Canavan. Facial Action Unit Detection using 3D Facial Landmarks. *arXiv e-prints*, art. arXiv:2005.08343, May 2020.
- A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam. Searching for MobileNetV3. may 2019. URL <http://arxiv.org/abs/1905.02244>.
- A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. apr 2017. URL <http://arxiv.org/abs/1704.04861>.
- G. Huang, M. Mattar, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. *Tech. rep.*, 10 2008a.
- G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition*, Marseille, France, Oct. 2008b. Erik Learned-Miller and Andras Ferencz and Frédéric Jurie. URL <https://hal.inria.fr/inria-00321923>.

- J. Huang, Z. Zhu, G. Huang, and D. Du. Aid: Pushing the performance boundary of human pose estimation with information dropping augmentation. *arXiv preprint arXiv:2008.07139*, 2020.
- D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, art. arXiv:1412.6980, Dec. 2014.
- M. Kowalski, J. Naruniec, and T. Trzcinski. Deep alignment network: A convolutional neural network for robust face alignment. *CoRR*, abs/1706.01789, 2017. URL <http://arxiv.org/abs/1706.01789>.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 1097–1105. Curran Associates, Inc., 2012. URL <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>.
- N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar. Attribute and simile classifiers for face verification. In *2009 IEEE 12th International Conference on Computer Vision*, pages 365–372, 2009. doi: 10.1109/ICCV.2009.5459250.
- L. Liang, R. Xiao, F. Wen, and J. Sun. Face alignment via component-based discriminative search. pages 72–85, 10 2008. ISBN 978-3-540-88685-3. doi: 10.1007/978-3-540-88688-4_6.
- T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. URL <http://arxiv.org/abs/1405.0312>.
- W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. SpheroFace: Deep hypersphere embedding for face recognition. *CoRR*, abs/1704.08063, 2017. URL <http://arxiv.org/abs/1704.08063>.
- Y. Liu, H. Shen, Y. Si, X. Wang, X. Zhu, H. Shi, Z. Hong, H. Guo, Z. Guo, Y. Chen, B. Li, T. Xi, J. Yu, H. Xie, G. Xie, M. Li, Q. Lu, Z. Wang, S. Lai, Z. Chai, and X. Wei. Grand Challenge of 106-Point Facial Landmark Localization. *arXiv e-prints*, art. arXiv:1905.03469, May 2019.
- S. Mahpod, R. Das, E. Maiorana, Y. Keller, and P. Campisi. Facial landmark point localization using coarse-to-fine deep recurrent neural network. *ArXiv*, abs/1805.01760, 2018.
- M. I. N. P. Munasinghe. Facial expression recognition using facial landmarks and random forest classifier. In *2018 IEEE/ACIS 17th International Conference on Computer and Information Science (ICIS)*, pages 423–427, 2018. doi: 10.1109/ICIS.2018.8466510.
- A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9912 LNCS:483–499, mar 2016. ISSN 16113349. doi: 10.1007/978-3-319-46484-8_29. URL <http://arxiv.org/abs/1603.06937>.
- C. Osendorfer, H. Soyer, and P. van der Smagt. Image super-resolution with fast approximate convolutional sparse coding. 11 2014. ISBN 978-3-319-12642-5. doi: 10.1007/978-3-319-12643-2_31.
- O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 41.1–41.12. BMVA Press, September 2015. ISBN 1-901725-53-7. doi: 10.5244/C.29.41. URL <https://dx.doi.org/10.5244/C.29.41>.
- 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. *CoRR*, abs/1603.01249, 2016. URL <http://arxiv.org/abs/1603.01249>.
- J. Roth, Y. Tong, and X. Liu. Unconstrained 3D face reconstruction. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 07-12-June-2015, pages 2606–2615. IEEE Computer Society, oct 2015. ISBN 9781467369640. doi: 10.1109/CVPR.2015.7298876.
- C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *2013 IEEE International Conference on Computer Vision Workshops*, pages 397–403, 2013. doi: 10.1109/ICCVW.2013.59.
- C. Sagonas, E. Antonakos, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 faces in-the-wild challenge: database and results. *Image and Vision Computing*, 47, 01 2016. doi: 10.1016/j.imavis.2016.01.002.
- M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. MobileNetV2: Inverted Residuals and Linear Bottlenecks. jan 2018. URL <http://arxiv.org/abs/1801.04381>.
- F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. *CoRR*, abs/1503.03832, 2015. URL <http://arxiv.org/abs/1503.03832>.
- S. Sengupta, J. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs. Frontal to profile face verification in the wild. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9, March 2016. doi: 10.1109/WACV.2016.7477558.
- W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1874–1883, 2016. doi: 10.1109/CVPR.2016.207.

- K. Sun, B. Xiao, D. Liu, and J. Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019.
- X. Sun, S. Jiaxiang, S. Liang, and Y. Wei. Compositional human pose regression. *Computer Vision and Image Understanding*, 176-177, 04 2017. doi: 10.1016/j.cviu.2018.10.006.
- Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3476–3483, 2013. doi: 10.1109/CVPR.2013.446.
- J. Tompson, A. Jain, Y. Lecun, and C. Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. 06 2014a.
- J. Tompson, A. Jain, Y. LeCun, and C. Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. *CoRR*, abs/1406.2984, 2014b. URL <http://arxiv.org/abs/1406.2984>.
- A. Toshev and C. Szegedy. DeepPose: Human pose estimation via deep neural networks. pages 1653–1660, 2014. doi: 10.1109/CVPR.2014.214.
- H. Wang, Y. Wang, Z. Zhou, X. Ji, Z. Li, D. Gong, J. Zhou, and W. Liu. Cosface: Large margin cosine loss for deep face recognition. *CoRR*, abs/1801.09414, 2018. URL <http://arxiv.org/abs/1801.09414>.
- Z. Wang, D. Liu, J. Yang, W. Han, and T. Huang. Deeply improved sparse coding for image super-resolution. 07 2015.
- S. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4732, 2016. doi: 10.1109/CVPR.2016.511.
- L. Wolf, T. Hassner, and Y. Taigman. Similarity Scores Based on Background Samples. In H. Zha, R.-i. Taniguchi, and S. Maybank, editors, *Computer Vision – ACCV 2009*, pages 88–97, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-12304-7.
- B. Xiao, H. Wu, and Y. Wei. Simple baselines for human pose estimation and tracking. *ArXiv*, abs/1804.06208, 2018.
- X. Xiong and F. De la Torre. Supervised descent method and its applications to face alignment. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 532–539, 06 2013. doi: 10.1109/CVPR.2013.75.
- Y. Xiong, Z. Zhou, Y. Dou, and Z. Su. Gaussian vector: An efficient solution for facial landmark detection, 2020.
- X. Yang, Y. Li, H. Qi, and S. Lyu. Exposing gan-synthesized faces using landmark locations. *CoRR*, abs/1904.00167, 2019. URL <http://arxiv.org/abs/1904.00167>.
- Y. Yang and D. Ramanan. Articulated human detection with flexible mixtures of parts. *IEEE transactions on pattern analysis and machine intelligence*, 35:2878–90, 12 2013. doi: 10.1109/TPAMI.2012.261.
- F. Zhang, X. Zhu, H. Dai, M. Ye, and C. Zhu. Distribution-aware coordinate representation for human pose estimation, 2019.
- F. Zhang, X. Zhu, H. Dai, M. Ye, and C. Zhu. Distribution-aware coordinate representation for human pose estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- J. Zhang, S. Shan, M. Kan, and X. Chen. Coarse-to-fine auto-encoder networks (cfan) for real-time face alignment. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 1–16, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10605-2.
- Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Learning deep representation for face alignment with auxiliary attributes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(5):918–930, 2016. doi: 10.1109/TPAMI.2015.2469286.
- T. Zheng and W. Deng. Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments. *Beijing University of Posts and Telecommunications, Tech. Rep*, 5, 2018.
- T. Zheng, W. Deng, and J. Hu. Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments. *arXiv preprint arXiv:1708.08197*, 2017.
- E. Zhou, H. Fan, Z. Cao, Y. Jiang, and Q. Yin. Extensive facial landmark localization with coarse-to-fine convolutional network cascade. In *2013 IEEE International Conference on Computer Vision Workshops*, pages 386–391, 2013. doi: 10.1109/ICCVW.2013.58.
- X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2879–2886, 2012. doi: 10.1109/CVPR.2012.6248014.