

Face Detection with End-to-End Integration of a ConvNet and a 3D Model

Yunzhu Li¹, Benyuan Sun¹, Tianfu Wu^{2(✉)}, and Yizhou Wang^{1(✉)}

¹ Nat'l Engineering Laboratory for Video Technology,
Key Laboratory of Machine Perception (MoE),
Cooperative Medianet Innovation Center,
Shanghai Sch'l of EECS, Peking University, Beijing 100871, China
{leo.liyunzhu,sunbenyuan,Yizhou.Wang}@pku.edu.cn

² Department of ECE and the Visual Narrative Cluster,
North Carolina State University, Raleigh, USA
tianfu_wu@ncsu.edu

Abstract. This paper presents a method for face detection in the wild, which integrates a ConvNet and a 3D mean face model in an end-to-end multi-task discriminative learning framework. The 3D mean face model is predefined and fixed (e.g., we used the one provided in the AFLW dataset). The ConvNet consists of two components: (i) The face proposal component computes face bounding box proposals via estimating facial key-points and the 3D transformation (rotation and translation) parameters for each predicted key-point w.r.t. the 3D mean face model. (ii) The face verification component computes detection results by pruning and refining proposals based on facial key-points based configuration pooling. The proposed method addresses two issues in adapting state-of-the-art generic object detection ConvNets (e.g., faster R-CNN) for face detection: (i) One is to eliminate the heuristic design of predefined anchor boxes in the region proposals network (RPN) by exploiting a 3D mean face model. (ii) The other is to replace the generic RoI (Region-of-Interest) pooling layer with a configuration pooling layer to respect underlying object structures. The multi-task loss consists of three terms: the classification Softmax loss and the location smooth l_1 -losses of both the facial key-points and the face bounding boxes. In experiments, our ConvNet is trained on the AFLW dataset only and tested on the FDDB benchmark with fine-tuning and on the AFW benchmark without fine-tuning. The proposed method obtains very competitive state-of-the-art performance in the two benchmarks.

Keywords: Face detection · Face 3D model · ConvNet · Deep learning · Multi-task learning

Y. Li and B. Sun contributed equally to this work and are joint first authors.

© Springer International Publishing AG 2016

B. Leibe et al. (Eds.): ECCV 2016, Part III, LNCS 9907, pp. 420–436, 2016.

DOI: 10.1007/978-3-319-46487-9_26

1 Introduction

1.1 Motivation and Objective

Face detection has been used as a core module in a wide spectrum of applications such as surveillance, mobile communication and human-computer interaction. It is arguably one of the most successful applications of computer vision. Face detection in the wild continues to play an important role in the era of visual big data (e.g., images and videos on the web and in social media). However, it remains a challenging problem in computer vision due to the large appearance variations caused by nuisance variabilities including viewpoints, occlusion, facial expression, resolution, illumination and cosmetics, etc.

It has been a long history that computer vision researchers study how to learn a better representation for unconstrained faces [12, 34, 40]. Recently, together with large-scale annotated image datasets such as the ImageNet [8], deep ConvNets [21, 22] have made significant progress in generic object detection [14, 16, 32], as well as in face detection [23, 30]. The success is generally considered to be due to the region proposal methods and region-based ConvNets (R-CNN) [15]. The two factors used to be addressed separately (e.g., the popular combination of the Selective Search [37] and R-CNNs pretrained on the ImageNet), and now they are integrated through introducing the region proposal networks (RPNs) as done in the faster-RCNN [32] or are merged into a single pipeline for speeding up the detection as done in [26, 31]. In R-CNNs, one key layer is the so-called RoI (Region-of-Interest) pooling layer [14], which divides a valid RoI (e.g., an object bounding box proposal) evenly into a grid with a fixed spatial extent (e.g., 7×7) and then uses max-pooling to convert the features inside the RoI into a small feature map. In this paper, we are interested in adapting state-of-the-art ConvNets of generic object detection (e.g., the faster R-CNN [32]) for face detection by overcoming the following two limitations:



Fig. 1. Some example results in the Fddb face benchmark [19] computed by the proposed method. For each testing image, we show the detection results (left) and the corresponding heat map of facial key-points with the legend shown in the right most column. (Best viewed in color) (Color figure online)




Fig. 2. Illustration of the proposed method of an end-to-end integration of a ConvNet and a 3D model for face detection (Top), and some intermediate and the final detection results for an input testing image (Bottom). See the legend for the classification score heat map in Fig. 1. The 3D mean face model is predefined and fixed in both training and testing. The key idea of the proposed method is to learn a ConvNet to estimate the 3D transformation parameters (rotation and translation) w.r.t. the 3D mean face model to generate accurate face proposals and predict the face key points. The proposed ConvNet is trained in a multi-task discriminative training framework consisting of the classification Softmax loss and the location smooth l_1 -losses [14] of both the facial key-points and the face bounding boxes. It is surprisingly simple w.r.t. its competitive state-of-the-art performance compared to the other methods in the popular FDDB benchmark [19] and the AFLW benchmark [41]. See text for details. (Color figure online)

- (i) RPNs need to predefine a number of anchor boxes (with different aspect ratios and sizes), which requires potentially tedious parameter tuning in training and is sensitive to the (unknown) distribution of the aspect ratios and sizes of the object instances in a random testing image.
- (ii) The RoI pooling layer in R-CNNs is predefined and generic to all object categories without exploiting the underlying object structural configurations, which either are available from the annotations in the training dataset (e.g., the facial landmark annotations in the AFLW dataset [20]) as done in [41] or can be pursued during learning (such as the deformable part-based models [10, 30]).

To address the two above issues in learning ConvNets for face detection, we propose to integrate a ConvNet and a 3D mean face model in an end-to-end multi-task discriminative learning framework. Figure 1 shows some results of the proposed method.

1.2 Method Overview

Figure 2 illustrates the proposed method. We use 10 facial key-points in this paper, including “LeftEyeLeftCorner”, “RightEyeRightCorner”, “LeftEar”, “NoseLeft”, “NoseRight”, “RightEar”, “MouthLeftCorner”, “MouthRightCorner”, “ChinCenter”, “CenterBetweenEyes” (see an example image in the left-top of Fig. 2). The 3D mean face model is then represented by the corresponding ten 3D facial key-points. The architecture of our ConvNet is straight-forward when taking into account a 3D model (see Sect. 3.2 for details).

The key idea is to learn a ConvNet to (i) estimate the 3D transformation parameters (rotation and translation) w.r.t. the 3D mean face model for each detected facial key-point so that we can generate face bounding box proposals and (ii) predict facial key-points for each face instance more accurately. Leveraging the 3D mean face model is able to “kill two birds with one stone”: Firstly, we can eliminate the manually heuristic design of anchor boxes in RPNs. Secondly, instead of using the generic RoI pooling, we devise a “configuration pooling” layer so as to respect the object structural configurations in a meaningful and principled way. In other words, we propose to learn to compute the proposals in a straight-forward top-down manner, instead of to design the bottom-up heuristic and then learn related regression parameters. To do so, we assume a 3D mean face model is available and facial key-points are annotated in the training dataset. Thanks to many excellent existing work in collecting and annotating face datasets, we can easily obtain both for faces nowadays. In learning, we have multiple types of losses involved in the objective loss function, including classification Softmax loss and location smooth l_1 -loss [14] of facial key-points, and location smooth l_1 -loss of face bounding boxes respectively, so we formulate the learning of the proposed ConvNet under the multi-task discriminative deep learning framework (see Sect. 3.3).

In summary, we provide a clean and straight-forward solution for end-to-end integration of a ConvNet and a 3D model for face detection¹. In addition to the competitive performance w.r.t the state-of-the-art face detection methods on the FDDB and AFW benchmarks, the proposed method is surprisingly simple and it is able to detect challenging faces (e.g., small, blurry, heavily occluded and extreme poses).

Potentially, the proposed method can be utilized to learn to detect other rigid or semi-rigid object categories (such as cars) if the required information (such as the 3D model and key-point/part annotation) are provided in training.

¹ We use the open source deep learning package, MXNet [5], in our implementation. The full source code is released at <https://github.com/tfwu/FaceDetection-ConvNet-3D>.

2 Related Work

There are a tremendous amount of existing works on face detection or generic object detection. We refer to [40] for a more thorough survey on face detection. We discuss some of the most relevant ones in this section.

In human/animal vision, how the brain distills a representation of objects from retinal input is one of the central challenges for systems neuroscience, and many works have been focused on the ecologically important class of objects—faces. Studies using fMRI experiments in the macaque reveal that faces are represented by a system of six discrete, strongly interconnected regions which illustrates hierarchical information processing in the brain [12], as well as some other results [34]. These findings provide some biologically-plausible evidences for supporting the usage of deep learning based approaches in face detection and analysis.

The seminal work of Viola and Jones [38] made face detection by a computer vision system feasible in real world applications, which trained a cascade of AdaBoost classifiers using Haar wavelet features. Many works followed this direction with different extensions proposed in four aspects: appearance features (beside Haar) including Histogram of Oriented Gradients (HOG) [7], Aggregate Channel Features (ACF) [9], Local Binary Pattern (LBP) features [1] and SURF [3], etc.; detector structures (beside cascade) including the the scalar tree [11] and the width-first-search tree [18], etc.; strong classifier learning (beside AdaBoost) including RealBoost [33] and GentleBoost [13], ect; weak classifier learning (beside stump function) including the histogram method [25] and the joint binarizations of Haar-like feature [28], etc.

Most of the recent face detectors are based on the deformable part-based model (DPM) [10,27,41] with HOG features used, where a face is represented by a collection of parts defined based on either facial landmarks or heuristic pursuit as done in the original DPM. [27] showed that a properly trained vanilla DPM can yield significant improvement for face detection.

More recent advances in deep learning [21,22] further boosted the face detection performance by learning more discriminative features from large-scale raw data, going beyond those handcrafted ones. In the Fddb benchmark, most of the face detectors with top performance are based on ConvNets [23,30], combining with cascade [23] and more explicit structure [39].

3D information has been exploited in learning object models in different ways. Some works [29,35] used a mixture of 3D view based templates by dividing the view sphere into a number of sectors. [17,24] utilized 3D models in extracting features and inferring the object pose hypothesis based on EM or DP. [36] used a 3D face model for aligning faces in learning ConvNets for face recognition. Our work resembles [2] in exploiting 3D model in face detection, which obtained very good performance in the Fddb benchmark. [2] computes meaningful 3D pose candidates by image-based regression from detected face key-points with traditional handcrafted features, and verifies the 3D pose candidates by a parameter sensitive classifier based on difference features relative to the 3D pose. Our work

integrates a ConvNet and a 3D model in an end-to-end multi-task discriminative learning fashion, which is more straightforward and simpler compared to [2].

Our Contributions. The proposed method contributes to face detection in three aspects.

- (i) It presents a simple yet effective method to integrate a ConvNet and a 3D model in an end-to-end learning with multi-task loss used for face detection in the wild.
- (ii) It addresses two limitations in adapting the state-of-the-art faster RCNN [32] for face detection: eliminating the heuristic design of anchor boxes by leveraging a 3D model, and replacing the generic and predefined *RoI pooling* with a *configuration pooling* which exploits the underlying object structural configurations.
- (iii) It obtains very competitive state-of-the-art performance in the FDDB [19] and AFW [41] benchmarks.

Paper Organization. The remainder of this paper is organized as follows. Section 3 presents the method of face detection using a 3D model and details of our ConvNet including its architecture and training procedure. Section 4 presents details of experimental settings and shows the experimental results in the FDDB and AFW benchmarks. Section 5 first concludes this paper and then discuss some on-going and future work to extend the proposed work.

3 The Proposed Method

In this section, we introduce the notations and present details of the proposed method.

3.1 3D Mean Face Model and Face Representation

In this paper, a 3D mean face model is represented by a collection of n 3D key-points in the form of (x, y, z) and then is denoted by a $n \times 3$ matrix, $F^{(3)}$. Usually, each key-point has its own semantic name. We use the 3D mean face model in the AFLW dataset [20] which consists of 21 key-points. We select 10 key-points as stated above.

A face, denoted by f , is presented by its 3D transformation parameters, Θ , for rotation and translation, and a collection of 2D key-points, $F^{(2)}$, in the form of (x, y) (with the number being less than or equal to n). Hence, $f = (\Theta, F^{(2)})$. The 3D transformation parameters Θ are defined by,

$$\Theta = (\mu, s, A^{(3)}), \quad (1)$$

where μ represents a 2D translation (dx, dy) , s a scaling factor, and $A^{(3)}$ a 3×3 rotation matrix. We can compute the predicted 2D key-points by,

$$\hat{F}^{(2)} = \mu + s \cdot \pi(A^{(3)} \cdot F^{(3)}), \quad (2)$$

where $\pi()$ projects a 3D key-point to a 2D one, that is, $\pi : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ and $\pi(x, y, z) = (x, y)$. Due to the projection $\pi()$, we only need 8 parameters out of the original 12 parameters. Let $A^{(2)}$ denote a 2×3 matrix, which is composed by the top two rows of $A^{(3)}$. We can re-produce the predicted 2D key-points by,

$$\hat{F}^{(2)} = \mu + A^{(2)} \cdot F^{(3)} \quad (3)$$

which makes it easy to implement the computation of back-propagation in training our ConvNet.

Note that we use the first sector in a 4-sector X - Y coordinate system to define all the positions, that is, the origin point $(0, 0)$ is defined by the left-bottom corner in an image lattice.

In face datasets, faces are usually annotated with bounding boxes. In the FDDB benchmark [19], however, faces are annotated with ellipses and detection performance are evaluated based on ellipses. Given a set of predicted 2D key-points $\hat{F}^{(2)}$, we can compute proposals in both ellipse form and bounding box form.

Computing a Face Ellipse and a Face Bounding Box based on a set of Predicted 2D Key-Points. For a given $\hat{F}^{(2)}$, we first predict the position of the top of head by,

$$\begin{pmatrix} x \\ y \end{pmatrix}_{\text{TopOfHead}} = 2 \times \begin{pmatrix} x \\ y \end{pmatrix}_{\text{CenterBetweenEyes}} - \begin{pmatrix} x \\ y \end{pmatrix}_{\text{ChinCenter}}.$$

Based on the keypoints of a face proposal, we can compute its ellipse and bounding box.

Face Ellipse. We first compute the outer rectangle. We use as one axis the line segment between the top-of-the-head key-point and the chin key-point, and then compute the minimum rectangle, usually a rotated rectangle, which covers all the key-points. Then, we can compute the ellipse using the two edges of the (rotated) rectangle as the major and minor axes respectively.

Face Bounding Box. We compute a face bounding box by the minimum upright rectangle which covers all the key-points, which is also adopted in the FDDB benchmark [19].

3.2 The Architecture of Our ConvNet

As illustrated in Fig. 2, the architecture of our ConvNet consists of:

- (i) *Convolution, ReLu and MaxPooling Layers.* We adopt the VGG [4] design in our experiments which has shown superior performance in a series of tasks. There are 5 groups and each group has 3 convolution and ReLu consecutive layers followed by a MaxPooling layer except for the 5th group. The spatial extent of the final feature map is of 16 times smaller than that of an input image due to the sub-sampling.
- (ii) *An Upsampling Layer.* Since we will measure the location difference between the input facial key-points and the predicted ones, we add an

upsampling layer to compensate the sub-sampling effects in previous layers. It is implemented by deconvolution. We upsample the feature maps to 8 times bigger in size (i.e., the upsampled feature maps are still quarter size of an input image) considering the trade-off between key-point location accuracy, memory consumption and computation efficiency.

- (iii) *A Facial Key-point Label Prediction Layer.* There are 11 labels (10 facial key-points and 1 background class). It is used to compute the classification Softmax loss based on the input in training.
- (iv) *A 3D Transformation Parameter Estimation Layer.* This is the key observation in this paper. Originally, there are 12 parameters in total consisting of 2D translation, scaling and 3×3 rotation matrix. Since we focus on the 2D projected key-points, we only need to account for 8 parameters (see the derivation above).
- (v) *A Face Proposal Layer.* At each position, based on the 3D mean face model and the estimated 3D transformation parameters, we can compute a face proposal consisting of 10 predicted facial key-points and the corresponding face bounding box. *The score of a face proposal* is the sum of *log* probabilities of the 10 predicted facial key-points. The predicted key-points will be used to compute the smooth l_1 loss [14] w.r.t. the ground-truth key-points. We apply the non-maximum suppression (NMS) to the face proposals in which the overlap between two bounding boxes a and b is computed by $\frac{|a \cap b|}{|b|}$ (where $|\cdot|$ represents the area of a bounding box), instead of the traditional intersection-over-union, accounting for the fact that it is rarely observed that one face is inside another one.
- (vi) *A Configuration Pooling Layer.* After NMS, for each face proposal, we pool the features based on the predicted 10 facial key-points. Here, for simplicity, we use all the 10 key-points without considering the invisibilities of certain key-points in different face examples.
- (vii) *A Face Bounding Box Regression Layer.* It is used to further refine face bounding boxes in the spirit similar to the method [14]. Based on the configuration pooling, we add two fully-connected layers to implement the regression. It is used to compute the smooth l_1 loss of face bounding boxes.

Denote by ω all the parameters in our ConvNet, which will be estimated through multi-task discriminative end-to-end learning.

3.3 The End-to-End Training

Input Data. Denote by $C = \{0, 1, \dots, 10\}$ as the key-point labels where $\ell = 0$ represents the background class. We use the image-centric sampling trick as done in [14, 32]. Without loss of generality, considering a training image with only one face appeared, we have its bounding box, $B = (x, y, w, h)$ and m 2D key-points ($m \leq 10$), $\{(\mathbf{x}_i, \ell_i)_{i=1}^m\}$ where $\mathbf{x}_i = (x_i, y_i)$ is the 2D position of the i th key-point and $\ell_i \geq 1 \in C$. We randomly sample m locations outside the face bounding box B as the background class, $\{(\mathbf{x}_i, \ell_i)_{i=m+1}^{2m}\}$ (where $\ell_i = 0, \forall i > m$). Note that in our ConvNet, we use the coordinate of the upsampled feature map which is

half size along both axes of the original input. All the key-points and bounding boxes are defined accordingly based on ground-truth annotation.

The Classification Softmax Loss of Key-point Labels. At each position \mathbf{x}_i , our ConvNet outputs a discrete probability distribution, $p^{\mathbf{x}_i} = (p_0^{\mathbf{x}_i}, p_1^{\mathbf{x}_i}, \dots, p_{10}^{\mathbf{x}_i})$, over the 11 classes, which is computed by the Softmax over the 11 scores as usual [21]. Then, we have the loss,

$$\mathcal{L}_{cls}(\omega) = -\frac{1}{2m} \sum_{i=1}^{2m} \log(p_{\ell_i}^{\mathbf{x}_i}) \quad (4)$$

The Smooth l_1 Loss of Key-point Locations. At each key-point location \mathbf{x}_i ($\ell_i \geq 1$), we compute a face proposal based on the estimated 3D parameters and the 3D mean face, denoted by $\{(\hat{\mathbf{x}}_j^{(i)}, \hat{\ell}_j^{(i)})_{j=1}^{10}\}$ the predicted 10 keypoints. So, for each key-point location \mathbf{x}_i , we will have m predicted locations, denoted by $\hat{\mathbf{x}}_{i,j}$ ($j = 1, \dots, m$). We follow the definition in [14] to compute the smooth l_1 loss for each axis individually.

$$\mathcal{L}_{loc}^{pt}(\omega) = \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m \sum_{t \in \{x,y\}} \text{Smooth}_{l_1}(t_i - \hat{t}_{i,j}) \quad (5)$$

where the smooth term is defined by,

$$\text{Smooth}_{l_1}(a) = \begin{cases} 0.5a^2 & \text{if } |a| < 1 \\ |a| - 0.5 & \text{otherwise.} \end{cases} \quad (6)$$

Faceness Score. The faceness score of a face proposal in our ConvNet is computed by the sum of log probabilities of the predicted key-points,

$$\text{Score}(\hat{\mathbf{x}}_i, \hat{\ell}_i) = \sum_{i=1}^{10} \log(p_{\ell_i}^{\hat{\mathbf{x}}_i}) \quad (7)$$

where for simplicity we do not account for the invisibilities of certain key-points. So the current faceness score has the issue of potential double-counting, especially for low-resolution faces. We observed that it hurts the quantitative performance in our experiments. We will address this issue in future work. See some heat maps of key-points in Fig. 1.

The Smooth l_1 Loss of Bounding Boxes. For each face bounding box proposal \hat{B} (after NMS), our ConvNet computes its bounding box regression offsets, $t = (t_x, t_y, t_w, t_h)$, where t specifies a scale-invariant translation and log-space height/width shift relative to a proposal, as done in [14, 32]. For the ground-truth bounding box B , we do the same parameterization and have $v = (v_x, v_y, v_w, v_h)$. Assuming that there are K bounding box proposals, we have,

$$\mathcal{L}_{loc}^{box}(\omega) = \frac{1}{K} \sum_{k=1}^K \sum_{i \in \{x,y,w,h\}} \text{Smooth}_{l_1}(t_i - v_i) \quad (8)$$

So, the overall loss function is defined by,

$$\mathcal{L}(\omega) = \mathcal{L}_{cls}(\omega) + \mathcal{L}_{loc}^{pt}(\omega) + \mathcal{L}_{loc}^{box}(\omega), \quad (9)$$

where the third term depends on the output of the first two terms, which makes the loss minimization more challenging. We adopt a method to implement the differentiable bounding box warping layer, similar to [6].

4 Experiments

In this section, we present the training procedure and implementation details and then show evaluation results on the FDDB [19] and AFW [41] benchmarks.

4.1 Experimental Settings

The Training Dataset. The only dataset we used for training our model is the AFLW dataset [20], which contains 25,993 annotated faces in real-world images. The facial key-points are annotated upon visibility w.r.t. a 3D mean face model with 21 landmarks. Of the images 70% are used for training while the remaining is reserved as a validation set.

Training process. For convenience, the short edge of every image is resized to 600 pixels while preserving the aspect ratio (as done in the faster RCNN [32]), thus our model learns how to handle faces under various scale. To handle faces of different resolution, we randomly blur images using Gaussian filters in pre-processing. Apart from the rescaling and blurring, no other preprocessing mechanisms (e.g., random crop or left-right flipping) are used.

We adopt the method of image-centric sampling [14, 32] which uses one image at a time in training. Under the consideration that grids around the labeled position share almost the same context information, thus the 3×3 grids around every labeled key-point's position are also regarded as the same positive examples, and we randomly choose the same amount of background examples outside the bounding boxes. The convolution filters are initialized by the VGG-16 [4] pretrained on the ImageNet [8]. We train the network for 13 epoch, and during the process, the learning rate is modified from 0.01 to 0.0001.


4.2 Evaluation of the Intermediate Results

Key-points classification in the validation dataset. As are shown by the heat maps in Fig. 1, our model is capable of detecting facial key-points with rough face configurations preserved, which shows the effectiveness of exploiting the 3D mean face model. Table 1 shows the key-point classification accuracy on the validation set in the last epoch in training.

Face proposals. To evaluate the quality of our face proposals, we first show some qualitative results on the FDDB dataset in Fig. 3. These ellipses are directly

Table 1. Classification accuracy of the key-points in the AFLW validation set at the end training.

Category	Accuracy	Category	Accuracy
Background	97.94 %	LeftEyeLeftCorner	99.12 %
RightEyeRightCorner	94.57 %	LeftEar	95.50 %
NoseLeft	98.48 %	NoseRight	97.78 %
RightEar	91.44 %	MouthLeftCorner	97.97 %
MouthRightCorner	98.64 %	ChinCenter	98.65 %
CenterBetweenEyes	96.04 %	AverageDetectionRate	97.50 %

**Fig. 3.** Examples of face proposals computed using predicted 3D transformation parameters without non-maximum suppression. For clarity, we randomly sample 1/30 of the original number of proposals.

calculated from the predicted 3D transformation parameters, forming several clusters around face instances. We also evaluate the quantitative results of face proposals. After a non-maximum suppression of IoU 0.7, the recall rate of 93.67 % is obtained with average 34.4 proposals per image.

4.3 Face Detection Results

To show the effectiveness of our method, we test our model on two popular face detection benchmarks: FDDDB [19] and AFW [41].

Results on FDDDB. FDDDB is a challenge benchmark for face detection in unconstrained environment, which contains the annotations for 5171 faces in a set of 2845 images. We evaluate our results by using the evaluation code provided by the FDDDB authors. The results on the FDDDB dataset are shown in Fig. 4. Our result is represented by “Ours-Conv3D”, which surpasses the recall rate of 90 % when encountering 2000 false positives and is competitive to the state-of-the-art methods. We compare with published methods only. Only DP2MFD [30] is slightly better than our model on discrete scores. It’s worth noting that we beat all other methods on continuous scores. This is partly caused by the predefined 3D face model helps us better describe the pose and part locations of faces.




Fig. 4. FDDDB results based on discrete scores using face bounding boxes in evaluation. The recall rates are computed against 2000 false positives.




Fig. 5. FDDDB results based on continuous scores using face ellipses in evaluation. The recall rates are computed against 2000 false positives.

We refer to the FDDDB result webpage² for details of the published methods evaluated on it (Figs. 4 and 5).

When comparing with recent work Faceness [39], we both recognize that one of the central issues to alleviate the problems of the occlusion and pose variation is to introduce facial part detector. However, our mechanism of computing face bounding box candidates is more straight forward since we explicitly integrate the structural information of a 3D mean face model instead of using a heuristic way of assuming the facial part distribution over a bounding box.

² <http://vis-www.cs.umass.edu/fddb/results.html>.




Fig. 6. Precision-recall curves on the AFW dataset (AP = average precision) without configuration pool and face bounding box regression used.




Fig. 7. Some qualitative results on the FDDB dataset

Results on AFW. AFW dataset contains 205 images with faces in various poses and view points. We use the evaluation toolbox provided by [27], which contains updated annotations for the AFW dataset where the original annotations are not comprehensive enough. Since the method of labeling face bounding boxes in AFW is different from that of in FDDB, we only use face proposals without configuration pooling and bounding box regression. The results on AFW are shown in Fig. 6.



Fig. 8. Some qualitative results on the AFW dataset

In our current implementation, there is one major limitation that prevents us from achieving better results. We do not explicitly handle invisible facial parts, which would be harmful when calculating the faceness score according to Eq. 7, we will refine the method and introduce mechanisms of handling the invisible problem in future work. More detection results on both datasets are shown in Figs. 7 and 8.

5 Conclusion and Discussion

We have presented a method of end-to-end integration of a ConvNet and a 3D model for face detection in the wild. Our method is a clean and straightforward solution when taking into account a 3D model in face detection. It also addresses two issues in state-of-the-art generic object detection ConvNets: eliminating heuristic design of anchor boxes by leveraging a 3D model, and overcoming generic and predefined RoI pooling by configuration pooling which exploits underlying object configurations. In experiments, we tested our method on two benchmarks, the FDDB dataset and the AFW dataset, with very compatible state-of-the-art performance obtained. We analyzed the experimental results and pointed out some current limitations.

In our on-going work, we are working on addressing the doubling-counting issue of the faceness score in the current implementation. We are also working on extending the proposed method for other types of rigid/semi-rigid object classes (e.g., cars). We expect that we will have a unified model for cars and faces which can achieve state-of-the-art performance, which will be very useful in a lot of practical applications such as surveillance and driveless cars.

Acknowledgement. Y. Li, B. Sun and Y. Wang were supported in part by China 973 Program under Grant no. 2015CB351800, and NSFC-61231010, 61527804, 61421062, 61210005. T. Wu was supported by the ECE startup fund 201473-02119 at NCSU.

References

1. Ahonen, T., Hadid, A., Pietikäinen, M.: Face description with local binary patterns: application to face recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **28**(12), 2037–2041 (2006)
2. Barbu, A., Gramajo, G.: Face detection using a 3D model on face keypoints. *CoRR abs/1404.3596* (2014)
3. Bay, H., Ess, A., Tuytelaars, T., Gool, L.J.V.: Speeded-up robust features (SURF). *Comput. Vis. Image Underst.* **110**(3), 346–359 (2008)
4. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: delving deep into convolutional nets. In: *British Machine Vision Conference* (2014)
5. Chen, T., Li, M., Li, Y., Lin, M., Wang, N., Wang, M., Xiao, T., Xu, B., Zhang, C., Zhang, Z.: MXNet: a flexible and efficient machine learning library for heterogeneous distributed systems. *CoRR abs/1512.01274* (2015)
6. Dai, J., He, K., Sun, J.: Instance-aware semantic segmentation via multi-task network cascades. In: *CVPR* (2016)
7. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *CVPR* (2005)
8. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Li, F.: Imagenet: a large-scale hierarchical image database. In: *CVPR*, pp. 248–255 (2009)
9. Dollár, P., Appel, R., Belongie, S.J., Perona, P.: Fast feature pyramids for object detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(8), 1532–1545 (2014)
10. Felzenszwalb, P.F., Girshick, R.B., McAllester, D.A., Ramanan, D.: Object detection with discriminatively trained part-based models. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(9), 1627–1645 (2010)
11. Fleuret, F., Geman, D.: Coarse-to-fine face detection. *Int. J. Comput. Vis.* **41**(1/2), 85–107 (2001)
12. Freiwald, W.A., Tsao, D.Y.: Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science* **330**(6005), 845–851 (2010)
13. Friedman, J., Hastie, T., Tibshirani, R.: Additive logistic regression: a statistical view of boosting. *Ann. Stat.* **28**, 2000 (1998)
14. Girshick, R.: Fast R-CNN. In: *ICCV* (2015)
15. Girshick, R.B., Donahue, J., Darrell, T., Malik, J.: Region-based convolutional networks for accurate object detection and segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(1), 142–158 (2016)

16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
17. Hu, W., Zhu, S.: Learning 3D object templates by quantizing geometry and appearance spaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **37**(6), 1190–1205 (2015)
18. Huang, C., Ai, H., Li, Y., Lao, S.: High-performance rotation invariant multiview face detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(4), 671–686 (2007)
19. Jain, V., Learned-Miller, E.: FDDB: a benchmark for face detection in unconstrained settings. Technical report UM-CS-2010-009, University of Massachusetts, Amherst (2010)
20. Koestinger, M., Wohlhart, P., Roth, P.M., Bischof, H.: Annotated facial landmarks in the wild: a large-scale, real-world database for facial landmark localization. In: First IEEE International Workshop on Benchmarking Facial Image Analysis Technologies (2011)
21. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
22. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–2324 (1998)
23. Li, H., Lin, Z., Shen, X., Brandt, J., Hua, G.: A convolutional neural network cascade for face detection. In: CVPR (2015)
24. Liebelt, J., Schmid, C.: Multi-view object class detection with a 3D geometric model. In: CVPR (2010)
25. Liu, C., Shum, H.: Kullback-Leibler boosting. In: CVPR (2003)
26. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: single shot multibox detector. arXiv preprint (2015). [arXiv:1512.02325](https://arxiv.org/abs/1512.02325)
27. Mathias, M., Benenson, R., Pedersoli, M., Gool, L.: Face detection without bells and whistles. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8692, pp. 720–735. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10593-2_47](https://doi.org/10.1007/978-3-319-10593-2_47)
28. Mita, T., Kaneko, T., Hori, O.: Joint Haar-like features for face detection. In: ICCV (2005)
29. Payet, N., Todorovic, S.: From contours to 3D object detection and pose estimation. In: ICCV (2011)
30. Ranjan, R., Patel, V.M., Chellappa, R.: A deep pyramid deformable part model for face detection. In: IEEE 7th International Conference on Biometrics Theory, Applications and Systems (2015)
31. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: unified, real-time object detection. In: CVPR (2016)
32. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: NIPS (2015)
33. Schapire, R.E., Singer, Y.: Improved boosting algorithms using confidence-rated predictions. *Mach. Learn.* **37**(3), 297–336 (1999)
34. Sinha, P., Balas, B., Ostrovsky, Y., Russell, R.: Face recognition by humans: 19 results all computer vision researchers should know about. *Proc. IEEE* **94**(11), 1948–1962 (2006)
35. Su, H., Sun, M., Li, F., Savarese, S.: Learning a dense multi-view representation for detection, viewpoint classification and synthesis of object categories. In: ICCV (2009)
36. Taigman, Y., Yang, M., Ranzato, M., Wolf, L.: Deepface: closing the gap to human-level performance in face verification. In: CVPR (2014)
37. Uijlings, J.R.R., van de Sande, K.E.A., Gevers, T., Smeulders, A.W.M.: Selective search for object recognition. *Int. J. Comput. Vis.* **104**(2), 154–171 (2013)

38. Viola, P.A., Jones, M.J.: Robust real-time face detection. *Int. J. Comput. Vis.* **57**(2), 137–154 (2004)
39. Yang, S., Luo, P., Loy, C.C., Tang, X.: From facial parts responses to face detection: a deep learning approach. In: *ICCV* (2015)
40. Zafeiriou, S., Zhang, C., Zhang, Z.: A survey on face detection in the wild. *Comput. Vis. Image Underst.* **138**(C), 1–24 (2015)
41. Zhu, X., Ramanan, D.: Face detection, pose estimation, and landmark localization in the wild. In: *CVPR* (2012)