Mask or Non-Mask? Robust Face Mask Detector via Triplet-Consistency Representation Learning

CHUN-WEI YANG, National Yang Ming Chiao Tung University, Taiwan

THANH HAI PHUNG, National Yang Ming Chiao Tung University, Taiwan

HONG-HAN SHUAI, National Yang Ming Chiao Tung University, Taiwan

WEN-HUANG CHENG, National Yang Ming Chiao Tung University and National Chung Hsing University, Taiwan

In the absence of vaccines or medicines to stop COVID-19, one of the effective methods to slow the spread of the coronavirus and reduce the overloading of healthcare is to wear a face mask. Nevertheless, to mandate the use of face masks or coverings in public areas, additional human resources are required, which is tedious and attention-intensive. To automate the monitoring process, one of the promising solutions is to leverage existing object detection models to detect the faces with or without masks. As such, security officers do not have to stare at the monitoring devices or crowds, and only have to deal with the alerts triggered by the detection of faces without masks. Existing object detection models usually focus on designing the CNN-based network architectures for extracting discriminative features. However, the size of training datasets of face mask detection is small, while the difference between faces with and without masks is subtle. Therefore, in this paper, we propose a face mask detection framework that uses the context attention module to enable the effective attention of the feed-forward convolution neural network by adapting their attention maps feature refinement. Moreover, we further propose an anchor-free detector with Triplet-Consistency Representation Learning by integrating the consistency loss and the triplet loss to deal with the small-scale training data and the similarity between masks and occlusions. Extensive experimental results show that our method outperforms the other state-of-the-art methods. The source code is released as a public download to improve public health at https://github.com/wei-1006/MaskFaceDetection.

CCS Concepts: • **Computing methodologies** → **Object detection**.

Additional Key Words and Phrases: Face detection, object detection, deep learning, face occlusion

1 INTRODUCTION

As the saying goes, "An ounce of prevention is worth a pound of cure." To alleviate the burden of healthcare caused by COVID-19, World Health Organization (WHO) suggests that masks should be used as part of a comprehensive

Authors' addresses: Chun-Wei Yang, National Yang Ming Chiao Tung University, 1001 University Road, Hsinchu, Taiwan, wei1997.ee08g@nctu.edu.tw; Thanh Hai Phung, National Yang Ming Chiao Tung University, 1001 University Road, Hsinchu, Taiwan, haipt.eed08g@nctu.edu.tw; Hong-Han Shuai, National Yang Ming Chiao Tung University, 1001 University Road, Hsinchu, Taiwan, hhshuai@nctu.edu.tw; Wen-Huang Cheng, National Yang Ming Chiao Tung University and National Chung Hsing University, 1001 University Road, Hsinchu, Taiwan, whcheng@nctu.edu.tw.

(a) Faces of different orientations   (b) Reflected faces without masks   (c) Crowd scene with different scales

(d) Occlusion and masks in different colors   (e) Faces with different poses   (f) Masks on non-human

(g) Occlusion and blurry faces without masks   (h) Makeup faces without masks   (i) Masks with occlusion
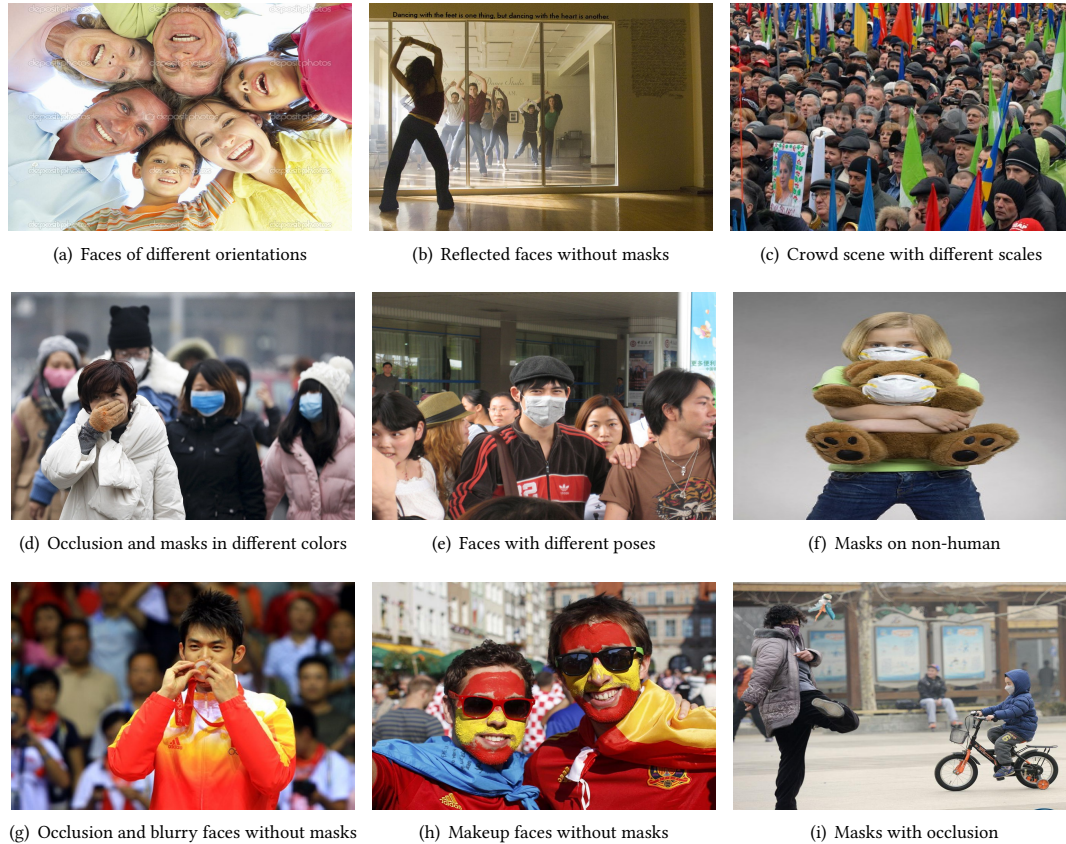
Fig. 1. Illustrative example of difficult cases on the AIZOO dataset.

strategy of measures to suppress transmission and provide protection against COVID-19. Leung et al. [41] also show that surgical masks can protect people from the coronavirus by reducing the probability of airborne transmission. As such, many governments mandate the use of face masks or coverings in public areas, such as retail establishments and public transportation [5, 33]. Nevertheless, to implement the executive order, it requires staff to monitor whether pedestrians wear masks or not at the entrance of restricted areas, which may be a heavy burden for large-scale and long-term surveillance.

With the advance of Artificial Intelligence (AI), it is promising to alleviate the burden by utilizing its machine learning algorithms for implementing an automatic monitoring system. Generally, facial mask detection can be regarded as a special case of object detection [3, 7, 9, 11, 19, 28, 32, 37, 38, 42, 50, 69, 77, 82–85]. That is, the goal is to determine where the faces are (object localization) and whether the faces are with or without masks (object classification) in a given image. Object detection can be categorized into two mainstreams, anchor-based approaches [3, 23, 24, 28, 37, 62, 69] and anchor-free approaches [10, 40, 57, 70, 77, 82, 83]. Anchor-based approaches leverage predefined multiple sizes anchor boxes to detect objects with different scales and aspect ratios, which can also be categorized into one-stage detectors [7, 18, 44, 46, 57, 85] and two-stage detectors [6, 23, 43, 62]. Two-stage detectors adopt region proposals to

extract the region of interest (ROI) and separate the object detection task into object localization and image classification. Due to the high computational complexity of two-stage detectors, one-stage detectors merge the tasks of object localization and image classification into a regression problem by predicting class probabilities and bounding box coordinates simultaneously. On the other hand, anchor-free detectors truncate anchor boxes and directly detect the vital keypoints [14, 40, 57, 70, 77, 82, 83], such as centers and corners of the object.

Generally, existing works of face detection solve a variety of challenging cases as illustrated in Figs. 1(a)-1(c), i.e., faces of different orientations, reflections, crowded scenes of different scales. Nevertheless, it is still challenging to detect faces with or without masks since several issues cannot be well-addressed by most of the previous works of object detection. First, masks are with different colors and styles, e.g., Fig. 1(d) contains dark blue, light blue, white and pink masks, while the masks in Fig. 1(e) and Fig. 1(f) are with different styles. To detect faces with masks of different colors and styles, one naïve solution is to increase the size of training data. However, the number of training data with facial masks is much smaller than of face images in existing datasets. Therefore, it is challenging to devise a machine learning model without a large-scale dataset. Second, images of wearing masks or not are only partially different, while the occlusions are similar to facial masks, e.g., one hand covering mouth in Fig. 1(d), kissing a medal in Fig. 1(g). Moreover, the occlusions of faces are diverse, e.g., makeup in Fig. 1(h), the clothing occluding the mask in Fig. 1(i). One of the possible solutions is to use fine-grained feature extraction models, e.g., [45, 51, 64, 76]. Nevertheless, the embeddings of faces with and without facial masks, as well as occluded faces, are still relatively close to each other, which easily leads to the confusion of the model.

To address these issues, in this paper, we present a new framework, namely, "CenterFace", for face mask detection with triplet-consistency representation learning. We first use the pre-trained models, e.g., ResNet-18 [29], MobileNet [30], to extract the basic features. Afterward, to enhance the basic features, the convolution block attention module [78] is leveraged, which is the simple yet effective attention but can blend the cross-channel and spatial information together by operating the interactively informative features along the two significant dimensions. The input feature map can easily concatenate with the attention maps for adaptively refining the features in the feed-forward convolution neural networks module. As the anchor-free models demonstrate promising efficiency and effectiveness, the proposed "CenterFace" uses a keypoint heat map to find the center point of its bounding box. Moreover, we propose a new Triplet-Consistency Representation Learning by integrating the consistency loss [32] and the triplet loss [63] to address the first and second challenges, respectively. Specifically, the consistency loss fully utilizes the labeled data to deal with the various orientations and visual appearance diversity, while the triplet loss emphasizes the difference of face with masks and occluded faces. Experimental results on public datasets show that the proposed "CenterFace" outperforms the state-of-the-art methods by 7.4% and 6.3% for the face with and without masks in terms of the precision, respectively, while the inference time satisfies the real-time constraint.

The main contributions can be summarized as shown bellows:

- Since enforcing a mask policy is important to control the disease propagated through airborne transmission, e.g., COVID-19, and alleviates the burden of healthcare, we propose a new framework, namely, "CenterFace", by utilizing the Triplet-Consistency Representation Learning to mitigate the effect of visual diversity and various orientations of the face masks.
- Experimental results on public datasets manifest that the proposed approach outperforms state-of-the-art methods, while the model size is small and can be adopted on the edge mobile devices. The source codes is released as a public download for public health improvement.

The rest of the paper is organized as follows. We present the related work in Section 2. The methodology is then described in detail in Section 3. Finally, the experimental results are elaborated in Section 4 and we make conclusions and describe future work in Section 5.

2 RELATED WORKS

2.1 Object Detection

Conventional face detection methods are based on the handcrafted feature due to the limitation of computing resources and the lack of large-scale datasets [2, 13, 16, 17, 23–26, 48, 49, 73, 74]. For example, Viola et al. [73] propose the first real-time human face detector without any constraint by combining the Haar features selection with integral images and detection cascades. Dalal et al. [13] further propose an improved model by extracting the histogram of the oriented gradients (HOG)-based scale-invariant feature transformation, which are robust to the translation, scale and illumination features, as well as different sizes of objects [2, 48, 49]. However, to detect objects of different sizes, the HOG requires re-scaling input images several times. Therefore, Felzenszwalb et al. [16, 17] propose the extension of HOG with the advancement of Fourier Transform by integrating the Local Binary Pattern (LBP) to extract the scale-invariant features map. Girsick et al. [23–26] propose the notion that objects can be modeled by parts in a deformable configuration and ensemble the detection of different object parts for the final prediction. Nevertheless, the representative DFM consists of a root-filter and a number of part-filters instead of specifying the size and location of the part filters.

With the advance of deep learning, extracting features from images is now data-driven instead of using the prior knowledge to handcraft the features. Object detection can be grouped into two categories: one-stage detector and two-stage detector. The two-stage detectors [23, 24, 35, 61, 67, 68, 87] first extract a set of object candidate boxes by selective search algorithm [71] as region proposals, and region proposals are warped into a square and fed into a convolutional neural network for extracting discriminative features. However, a large number of overlapped region proposals make redundant feature computation, leading to inefficient detection models. Therefore, several works make progress on improving the computing of the overlapped proposals [23, 43, 61, 61]. Even though, the training process is still multi-stage with computation redundancy at the subsequent detection stage.

On the other hand, one-stage detectors can be further categorized into anchor-based methods and anchor-free methods. Redmon et al. propose you only look once (YOLO) series [4, 58–60], which replace the region proposal network (RPN) with anchor boxes predefined the ratio of width and heights of objects. However, YOLO series still suffers from improving the localization accuracy as compared with two-stage detectors. To solve the problem of one stage detection on localization and small object detection, Liu et al. propose a single shot multibox detector (SSD) [46], which focuses on the multi-scale object detection with a variety of sizes and aspect ratios and computes both the location and class scores using small convolution filters. Despite the improvement of the speed and accuracy, the performance of one-stage detectors is usually inferior to two-stage detectors. A recent line of studies proposes anchor-free detectors, which directly find the objects without using multiple anchors in the input images [37, 39, 44, 70, 86]. There are two popular kinds of anchor-free detection methods. The first kind is keypoint-based methods [39, 86], which bounds the spatial extent of objects by locating several predefined or self-learned keypoints. The second one is center-based methods [37, 70], which defines positive by the center point or region of objects to predict their boundary.

Self-supervised learning has become a popular learning, which creates pretext tasks on unlabeled data and learns in a supervised manner. In other words, the goal of self-supervised learning is to construct the representation of images with semantically meaningful content via pretext task while not requiring the semantic annotation for a large training set of
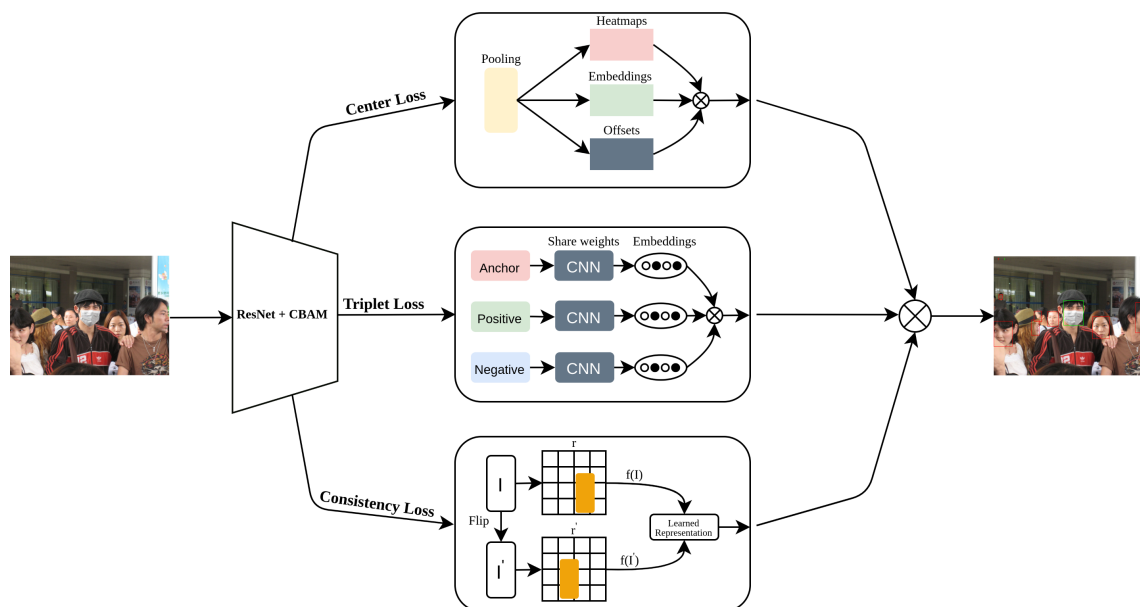
Fig. 2. The overview architecture of CenterFace. We use an output heatmap from a convolution network to detect a face and/or a mask as a pair of the bounding box, the network is trained to predict similar embeddings for faces that belong to the same face.

images. For example, Misra et al. [52] use the pretext tasks with image transformation to encourage the representation of images to be invariant with the image patch perturbation. Xu et al. [80] exploit the similarity from the self-supervised signals as an auxiliary task, which can effectively transfers the hidden information from the teacher to the student network. Gidaris et al. [21] propose a self-supervised learning method based on transforming input images in different rotations, so the framework can learn to estimate the geometric transformations applied to the image, which helps the downstream tasks, such as object detection. Grill et al. [27] leverage two networks, where an online network is trained to predict the representation of another network of another augmented view of the identical image.

2.2 Attention Mechanism

Attention has become a popular concept and a useful tool in the deep learning community in recent years [12, 15, 31, 47, 72, 75, 78]. The basic idea is that human visual perception only focuses on some specific regions at one time and performs well in object detection. Therefore, to mimic human attention as a sequence of partial glimpses, different kinds of attention mechanisms are proposed to learn "what" and "where" to attend and then focus on the important features and suppressing unnecessary ones. For example, Vaswani et al. [72] use a decoder-encoder architecture with the attention mechanism to refine the feature maps. Hu et al. [31] introduce the inter-channel attention by using global-pooled features to compute their channel-wise attention. Woo et al. [78] demonstrate that channel-wise attention is insufficient and provides the spatial attention to decide "where" to focus, enabling the attention generation process for 3D features map with much less number of parameters.
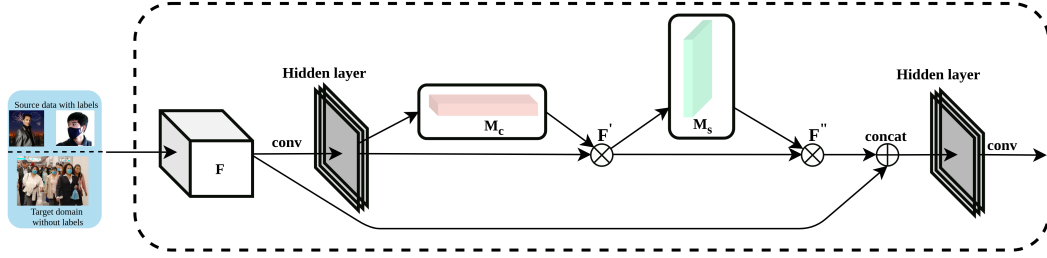
Fig. 3. Convolution Block Attention Module (CBAM) [78] with a block in ResNet [29]. We utilize CBAM to integrate within the convolution outputs in each ResNet block.

3 METHODOLOGY

To effectively detect the faces with and without facial masks, we present a new framework, namely, `"CenterFace"`, for face mask detection with triplet-consistency representation learning. Fig. 2 shows the overview of the proposed `"CenterFace"`. Specifically, we first use the pre-trained models to extract the low-level features since low-level feature extraction is usually shared in different tasks, e.g., texture, edges. Afterward, to enhance the basic features without significantly increasing the number of parameters, the convolution block attention module (CBAM) [78] is leveraged to attend the channel and spatial features. Fig. 3 illustrates the model architecture of CBAM, which contains channel and spatial modules. The channel module utilizes both max-pooling outputs and average-pooling outputs for finding the attention weights on channels, while the spatial module uses similar two outputs and pool along the channel axis to generate the attention weights on spatial dimensions.

Moreover, since the anchor-free models have been proved to be efficient and effective, the proposed `"CenterFace"` uses a key-point heat map to find the center point of its bounding box. Moreover, we propose a new Triplet-Consistency Representation Learning by integrating the consistency loss [32] and the triplet loss [63] to address the first and second challenges, respectively. Specifically, the consistency loss fully utilizes the labeled data to deal with the various orientations and visual appearance diversity, while the triplet loss emphasizes the difference of face with masks and occluded faces. In the following, we first explain the center loss, which leverages the anchor-free models, and then the triplet loss and consistency loss for Triplet-Consistency Representation Learning. Finally, the total loss function is presented, together with the pseudocode. In the following, bold uppercase letters (e.g., $\mathbf{X}$) and lowercase letters (e.g., $\mathbf{x}$) denote matrices and column vectors, respectively. Non-bold letters (e.g., x) and squiggle letters (e.g., $\mathcal{X}$) represent scalars and tensors, respectively.

3.1 Center Loss

Most of the successful object detection models require enumerating candidate object locations and classifying each candidate region, which is computationally expensive. Moreover, since the model may predict several candidate bounding boxes for one object, the predicted bounding boxes require additional post-processing, e.g., non-maximum suppression [53]. To satisfy the requirement of real-time monitoring, we adopt the anchor-free method, which performs object classification and bounding box localization simultaneously. Specifically, let $\mathcal{I} \in \mathbb{R}^{h \times w \times 3}$ denote the input image tensor with the width w and height h. The goal is to generate the key-point heatmap $\hat{\mathcal{Y}} \in [0, 1]^{\frac{h}{s} \times \frac{w}{s} \times c}$ with the stride

s, where $\hat{\mathcal{Y}}_{i,j,k} = 1$ and $\hat{\mathcal{Y}}_{i,j,k} = 0$ respectively represent a predicted keypoint and the background region at position (i, j) belonging to class k, c is the number of keypoint classes,[1]

To train the model for predicting the keypoints, given a ground truth keypoint $\mathbf{p} = (p_x, p_y) \in \mathbb{R}^2$, we first calculate the low resolution equivalence $\tilde{\mathbf{p}} = (\lfloor \frac{p_x}{s} \rfloor, \lfloor \frac{p_y}{s} \rfloor)$. We then smooth the groundtruth keypoint by the 2D Gaussian kernel to derive the heatmap $\mathcal{Y} \in [0, 1]^{\frac{h}{s} \times \frac{w}{s} \times c}$ as follows.

$$y_{i,j,k} = \exp\left(-\frac{(i - \lfloor \frac{p_x}{s} \rfloor)^2 + (j - \lfloor \frac{p_y}{s} \rfloor)^2}{2\delta_p^2} \right), \tag{1}$$

where δ_p is the object-size standard deviation, which is determined by its size to ensure that the pair of radius points can generate the bounding box around the face by the ground truth annotation for each ground truth keypoint $\mathbf{p}$.

The training objective is a pixelwise logistic regression with a variant of focal loss [44], i.e.,

$$L_{pix} = -\frac{1}{N} \sum_{i=1}^{h} \sum_{j=1}^{w} \sum_{k=1}^{c} \begin{cases} (1 - \hat{Y}_{i,j,k})^\alpha \log(\hat{Y}_{i,j,k}) & \text{if } Y_{i,j,k}=1, \\ (1 - Y_{i,j,k})^\beta (\hat{Y}_{i,j,k})^\alpha \log(1 - \hat{Y}_{i,j,k}) & \text{otherwise,} \end{cases} \tag{2}$$

where α and β are the hyperparameters controlling the contribution for each point of the focal loss.[2] N is the number of the faces in input images I, which is applied to normalize all the positive focal instance losses to 1.

As the output stride may cause errors due to the discretization, i.e., $\mathbf{p}$ becomes $\tilde{\mathbf{p}}$, we predict the local offset $\hat{O} \in \mathbb{R}^{\frac{h}{s} \times \frac{w}{s} \times 2}$ for each center point, where the groundtruth offset can be derived by calculating $\frac{\mathbf{p}}{s} - \lfloor \frac{\mathbf{p}}{s} \rfloor$ for each center point $\mathbf{p}$. Let $\mathbf{p}_n$ denote the position of the n-th center. The offset loss can be derived as follows.

$$L_{off} = \frac{1}{N} \sum_{n=1}^{N} \left| \hat{O}_{\mathbf{p}_n} - (\frac{\mathbf{p}_n}{s} - \lfloor \frac{\mathbf{p}_n}{s} \rfloor) \right|. \tag{3}$$

In addition to the centers derived by the keypoint heatmap $\hat{\mathcal{Y}}$ and offset $\hat{O}$, the next goal is to find the bounding boxes. Let $(x_1^n, y_1^n, x_2^n, y_2^n)$ be the coordinates of the bounding box of the object n with class c_n, while the center can be represented by $\mathbf{p}_n = \left(\frac{x_1^n + x_2^n}{2}, \frac{y_1^n + y_2^n}{2} \right)$. We predict the bounding box size $\hat{S}_n \in \mathbb{R}^{\frac{h}{s} \times \frac{w}{s} \times 2}$ for each center point. It is worth noting that $\hat{S}_n$ is shared for the classes of the faces with or without masks to reduce the computational cost. The groundtruth for the n-th center can be calculated by $\mathbf{s}_n = (x_2^n - x_1^n, y_2^n - y_1^n)$. As such, the loss for bounding box size, denoted by L_s, can be derived as follows.

$$L_s = \frac{1}{N} \sum_{n=1}^{N} \left| \hat{S}_{\mathbf{p}_n} - \mathbf{s}_n \right|. \tag{4}$$

The overall center loss for detecting faces with and without masks is:

$$L_{center} = \lambda_{pix} L_{pix} + \lambda_{off} L_{off} + \lambda_s L_s, \tag{5}$$

where λ_{pix}, λ_s, and λ_{off} are the hyperparameters controlling the weights between different prediction targets.[3]

3.2 Triplet Loss

Due to the subtle difference between faces with masks and faces with occlusions, it is difficult to train a model by only using the center loss. Therefore, to further improve the performance on challenging cases, we utilize an online triplet

---

[1]We set s as 4 according to previous work [8, 54, 55], and set c as 2 since the goal is to detect the faces are with or without masks.
[2]Empirically, we set α and β as 2 and 4, respectively.
[3]The hyperparameters are empirically set as $\lambda_{pix} = 1$, $\lambda_{off} = 1$ and $\lambda_s = 0.01$.

mining method [63], which enforces that the distance between a pair of samples with the same label is smaller than that between a pair of samples with different labels. Indeed, the triplet loss works directly on features. A triplet is formed by 1) anchor input, 2) positive input (samples with the same label) and 3) negative input (samples with the different labels or random cropped samples). The distance between the anchor input and the positive input is enforced to be smaller than the distance between the anchor input and the negative input. Specifically, let $\mathcal{R}_i^a$, $\mathcal{R}_i^+$ and $\mathcal{R}_i^-$ respectively denote anchor image region, positive image region and negative image region in the i-th triplet. The triplet loss can be derived as follows.

$$L_{tri} = \sum_i \left[\left\| f(\mathcal{R}_i^a) - f(\mathcal{R}_i^+) \right\|_2^2 - \left\| f(\mathcal{R}_i^a) - f(\mathcal{R}_i^-) \right\|_2^2 + \epsilon \right]. \tag{6}$$

The input image regions $\mathcal{R}_i^a$, $\mathcal{R}_i^+$ and $\mathcal{R}_i^-$ are embedded by the same mapping function f. We ensure that the features of the anchor region $\mathcal{R}_i^a$ is particularly close to other positive samples $\mathcal{R}_i^+$ and far away from any of the negative images $\mathcal{R}_i^-$ by using the margin ϵ to separate positive and negative pairs. In other words, when faces are regarded as anchors and positive examples, the masks and some regions near the face and mask areas become negative examples. On the other hand, when the face with masks are regarded as anchors and positive examples, the face area and some regions near face areas become negative examples. Eventually, the distance from the anchor input to the positive samples is minimized, while the distance from the anchor input to the negative samples is maximized.

3.3 Consistency Loss

In addition to the triplet loss, we further impose the consistency constraint on the predicted heatmap of the original image and the predicted heatmap of the horizontally-flipped images to simultaneously stabilize the prediction results and enlarges the training data. The consistency constraint is especially effective when incorporating the triplet loss since multiple positive samples can be created at the same time. Specifically, let $\hat{\mathcal{Y}}$ and $\hat{\mathcal{Y}}'$ denote the predicted heatmap of the original image and the predicted heatmap of the horizontally-flipped images, respectively. The prediction probability at position (i, j), i.e., $\hat{\mathcal{Y}}_{i,j,:}$, should be close to the prediction probability at the horizontally opposite position (i', j'), i.e., $\hat{\mathcal{Y}}'_{i',j',:}$. One way to measure the closeness is to use L2 distance. However, L2 loss regards all the classes equally so that the irrelevant classes with a low probability may also highly affects the results. Therefore, inspired by [32], we also use Jensen-Shannon divergence (JSD) to measure the difference between $\hat{\mathcal{Y}}_{i,j,:}$ and $\hat{\mathcal{Y}}'_{i',j',:}$. The consistency loss for classification, denoted by L_{con-c}, is obtained by calculating the expectation of all bounding box pairs can thus be derived as follows.

$$L_{con-c} = \mathbb{E}_n[JSD(\hat{\mathcal{Y}}_{\mathbf{p}_n}, \hat{\mathcal{Y}}'_{\mathbf{p}'_n}))]. \tag{7}$$

On the other hand, the localization result of the candidate box is based on the offset $\hat{O}$ and size $\hat{S}$. Let $\hat{O}'$ and $\hat{S}'$ denote the predicted offset and size of the horizontally-flipped images, respectively. Unlike the classification, a simple modification is required to make the prediction equivalent to each other. As the flipping transformation makes the offset change in the opposite direction, a negation is applied to correct the groundtruth, i.e.,

$$\hat{O}_{\mathbf{p}_n,1} \Leftrightarrow -\hat{O}'_{\mathbf{p}'_n,1}$$

$$\hat{O}_{\mathbf{p}_n,2}, \hat{S}_{\mathbf{p}_n} \Leftrightarrow \hat{O}'_{\mathbf{p}'_n,2}, \hat{S}'_{\mathbf{p}'_n}$$

Afterward, we derive the localization consistency loss by calculating the expectation of the offset and size difference for each single pair of the candidate box center at position $\mathbf{p}_n$ and $\mathbf{p}'_n$ as below.

$$L_{con-l} = \mathbb{E}_n[\left\| \hat{O}_{\mathbf{p}_n,1} - (-\hat{O}'_{\mathbf{p}'_n,1}) \right\|^2 + \left\| \hat{O}_{\mathbf{p}_n,2} - (\hat{O}'_{\mathbf{p}'_n,2}) \right\|^2 + \left\| \hat{S}_{\mathbf{p}_n} - \hat{S}'_{\mathbf{p}'_n} \right\|^2]. \tag{8}$$

Finally, when the consistency loss is computed with all candidates, the results may easily be dominated by the backgrounds, which deteriorates the performance of the foreground classification candidates. As such, we have to build a mask with the same size as the heatmap for every groundtruth bounding box to exclude the boxes with a high probability of background class. The total consistency loss, denoted by L_{con}, simply summarizes L_{con-c} and L_{con-l}, i.e., $L_{con} = L_{con-c} + L_{con-l}$.

3.4 Overall Objective Function

The total loss function is summarized over the center loss, triplet loss and consistency loss, i.e.,

$$L_{total} = L_{center} + \lambda_{tri}L_{tri} + \lambda_{con}L_{con}, \tag{9}$$

where λ_{tri} and λ_{con} are hyper-parameters controlling the contribution of different loss terms.[4] In summary, the keypoint heatmap prediction works as a general-purpose object detector which extends the keypoint estimator to generate the face's bounding box. On the other hand, the triplet loss, which enforces the face margin between each pair of face to roughly align matching/non-matching face, along with the consistency classification and localization loss are proposed to localize not only the face classification but also the position. The overall objective loss facilitates the detection model to be robust enough for any complicated situations among various representation of the face detection, especially from the masked face detection. The pseudocode of "CenterFace" is presented in Algorithm 1.

Algorithm 1 CenterFace algorithm

Input:
  Input image set with groundtruth
  Hyperparameters α and β (focal loss), stride s, margin ϵ, λ_{pix}, λ_s, λ_{off}, λ_{tri}, λ_{con}, number of iterations n_{iter}
Output:
  Bounding box with class represented by $\hat{\mathcal{Y}}^*$, $\hat{S}^*$ and $\hat{O}^*$
1: **for** $iter = 1...n_{iter}$ **do**
2:     Generate heatmaps $\hat{\mathcal{Y}}$ with size $\hat{S}$ and offset $\hat{O}$ from L_{center} according to Section 3.1
3:     **for each** $\mathcal{I}$ **do**
4:         Compute L_{tri} according to Eq. 6
5:         Compute L_{con-c} according to Eq. 7
6:         Compute L_{con-l} according to Eq. 8
7:         Optimize $L_{total} = L_{center} + \lambda_{tri}L_{tri} + \lambda_{con}L_{con}$
8:     **end for**
9:     **if** L_{total} is the smallest so far **then**
10:         $\hat{\mathcal{Y}}^*, \hat{S}^*, \hat{O}^* \leftarrow \hat{\mathcal{Y}}, \hat{S}, \hat{O}$
11:     **end if**
12: **end for**
13: **return** $\hat{\mathcal{Y}}^*$, $\hat{S}^*$ and $\hat{O}^*$

---

4 EXPERIMENTS

In this section, we first present the experimental setup in detail. Next, we evaluate the proposed framework on public datasets with diverse face mask representations including synthetic data, visual diversity, and various orientations. Finally, the ablation studies are presented to show the effectiveness of each component in the proposed framework.

---

[4]Empirically, we set $\lambda_{tri} = 1$ and $\lambda_{con} = 100$.

Table 1. Detection Accuracy of CenterFace

Model	Face		Mask	
	Precision	Recall	Precision	Recall
baseline [1]	89.6%	85.3%	91.9%	88.6%
RetinaMask + MobileNet [34]	83.0%	95.6%	82.3%	89.1%
RetinaMask + ResNet50 [34]	91.9%	**96.3%**	93.4%	94.5%
CenterFace (MobileNet)	96.8%	88.0%	88.0%	95.2%
CenterFace (ResNet18)	**99.3%**	96.2%	**99.7%**	**96.7%**

4.1 Experimental Setup

Datasets. In this section, we evaluate our detection approach on the AIZOO dataset [1], WiderFace [81] and MAFA dataset [20] to verify the effectiveness of our "CenterFace" module with baselines. AIZOO dataset is introduced when the pandemic situation just ransomed for the hope that people in the world can defeat the pandemic as soon as possible. The data set includes 7,959 images with various illuminations, occlusion, and different poses, and splits into the training, testing, and validation sets with 4096, 1,226, and 1,839, respectively. Moreover, WiderFace dataset contains 393,703 annotated faces with a variety in poses, crowd, face expression, occlusion, and scenarios from 32,203 images. These images are split into three subsets: training (40%), validation (10%), and testing (50%) set along with three difficulty levels: easy, medium, and hard based on the detection rate of EdgeBox benchmark [88]. On the other hand, MAFA Dataset is a popular facial mask dataset that contains 35,806 masked faces in 30,811 images. This image dataset has been covered almost incident cases of occluded faces in real-life. In detail, it contains 60 cases of occluded faces with three different levels of occlusions in daily scenarios. Furthermore, four types of masks and five face orientations are also provided to diversify the dataset, and if the faces are considerably blurred or smaller than 32 pixels, it is ignored. All of those properties are divided into three subsets: the whole subset contains 6354 masked faces, 996 unmasked faces, and 2683 ignored faces. The masked subset and the remaining subset consist of both masked faces and unmasked faces.

Baselines. To the best of our knowledge, only one published paper that focuses on face mask detection with available source codes. Specifically, in AIZOO dataset, the authors verify their dataset by deploying the structure of SSD [46] with a light-weight backbone network, which contains 24 layers, of which 8 layers are convolution layers. The input size is set as 260×260, while the total number of parameters is only 1.01M.

Evaluation Metrics We evaluate the above approach by calculating the precision and recall metrics for the facial detection model and masked facial classifier, which is defined as follows.

$$Precision = \frac{TP}{TP+FP} \times 100\%$$

$$Recall = \frac{TP}{TP+FN} \times 100\% \tag{10}$$

4.2 Implementation Details

In the experiment, we follow the evaluation protocol in [79], and use ResNet-18 [29] as the backbone network. The model is trained on an input image resolution 520×520 with a batch size of 10 with Adam optimizer [36]. The learning rate is $1.25e^{-4}$ for 140 epochs, with the learning rate dropped 10× at 30 and 80 epochs, respectively. We implement the proposed CenterFace with Pytorch [56]. Then, to strengthen the effective attention for feed-forward convolutional

Table 2. Flops and Number of Parameters of Models (AIZOO dataset).

Model	FLOPs	MACs
baseline [1]	1.01 M	–
RetinaMask + MobileNet [34]	3.5 M	1.67 G
RetinaMask + ResNet50 [34]	25.56 M	21.53 G
CenterFace (MobileNet)	3.5 M	1.67 G
CenterFace (ResNet18)	11.69 M	9.52 G

neural network, we utilize CBAM [78] to every single block of the CNNs. The attention module not only improves the effectiveness of the CNNs but also gives a solution to reduce the computation and time consuming of the framework.

4.3 Experiment Results

Performance Comparison on AIZOO Dataset. We evaluate the proposed model with two light-weight backbone networks, ResNet-18 [29] and MobileNet [30], which are capable to deploy on the edge mobile devices. Table 1 compares the proposed method with different approaches in terms of precision and recall. The results manifest that our model significantly outperforms other candidate methods. Specifically, equipped with the proposed Triplet-Consistency Representation Learning, MobileNet and ResNet improve the RetinaMask [34] by 7.4% and 6.3% in both face and mask classes in terms of precision. On the other hand, the recall of the proposed approach is almost equaled to the face class and 2.2% higher in the mask class.

Efficiency Analysis. To evaluate the efficiency of different approaches, we use FLOPs and MACs to measure the complexity of different models. As shown in Table 2, the results manifest that the number of parameters for ours with RestNet-18 network is much less than the state-of-the-art results using ResNet-50 since the model size of ResNet-50 is greater than the of ResNet-18. Nevertheless, even with a deep network, Table 3 still shows that a shallower network with the proposed triplet-consistency representation learning performs better than that with a deeper network. Even more, when using the same MobileNet network, our result still outperforms the current state-of-the-artwork. This is because of the efficiency of the anchor-free techniques and the attention module we utilize in this work, which saves the computation cost by enhancing the features with fewer parameters, instead of directly increasing the number of layers.

Analysis on Qualitative Results. Fig. 4 demonstrates the qualitative detection results of our methods on the face class, as well as the results of the baseline. The demonstration cases show the difficulties and challenges of face detection such as intra-class variation, occlusion, and multi-scale detection [89]. For example, the intra-class variation of human face may present as a variety of expressions, movements, poses, and skin colors as shown in Figs. 4(c), 4(e), and 4(g). Moreover, the faces may be occluded by other things as shown in Figs. 4(d) and 4(h) or in a large variety of scales as shown in Figs. 4(a), 4(b), and 4(f). Although these references are from various perspectives, our model is still able to extract the salient parts, while the baselines may miss one or two faces. Fig. 5 further demonstrates the qualitative detection results of our methods on the mask class, as well as the results of the baseline. In details, the qualitative results of the baseline are worse than that of the proposed method. For example, in various cases such as occluded faces in Figs. 5(d), 5(f) 5(g), the intra-class issues such as various of expression, blur face, small face as shown in the rest of Fig. 5.

Due to the lack of public video benchmarks, we collect the online video clips with face masks and also conduct a laboratory experiment. For the qualitative video results, please refer to the following URLs in the footnotes.[5][6] The

---

[5]Example 1: https://drive.google.com/file/d/1Tl-8iRXgxMqeM8sbSll0tU9RubccT-lc/view.
[6]Example 2: https://drive.google.com/file/d/1ld1Nw3-cr1ODvzEBGkKL6hUH2dVjmojQ/view.

(a) Crowd scene with different scales

(b) Multi-pose facial detection

(c) Occlusion faces

(d) Reflected faces without masks

(e) Faces of different orientations

(f) Small faces in the meeting scene
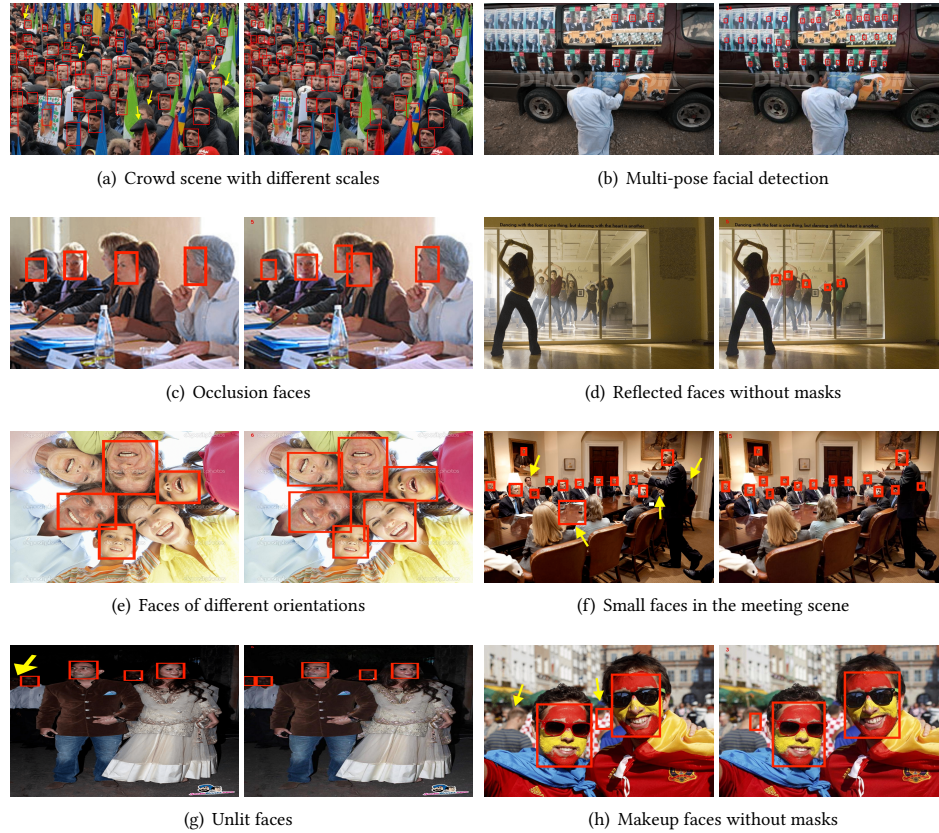
(g) Unlit faces

(h) Makeup faces without masks

Fig. 4. Demonstrations of the detection results on the target image on the Face class, where the results of the baseline and ours are shown on the left-hand side and right-hand side, respectively. Best views in color.

quantitative results show that 5/40 people in the video clips are missed in some frames. However, most of the missing ones are in a far distance. In contrast, 10/10 people in the laboratory experiment are correctly detected. The results manifest that the proposed model is suitable for monitoring the entrance of public transports.

Failure Case Analysis. The objective loss function of "CenterFace" module is comprised of three different losses, which have different effects on the performance of the detection. For example, a face may be missed if 1) the anchor is missed, 2) imprecise embeddings relate to multiple false bounding boxes, or 3) the offset threshold is not correct. To understand how each part contributes to errors, we take a look at Fig. 6. The major causes of qualitative error results happen when faces are hidden by humans or things as shown in Figs. 6(c), 6(d), and 6(f). In these cases, the detectors cannot recognize a face with or without facial masks. On the other hand, in Figs. 6(a), 6(b), and 6(e), faces are misrecognized when the needed detection faces are too small or in cozy environment.

4.4 Ablation Studies

Here, the ablation studies are presented to show the effectiveness of each component in the proposed framework. We perform the ablation studies on the effectiveness of objective function and training setting. We use WiderFace and

(a) Faces with different poses

(b) Masks with occlusion

(c) The street view scene with and without masks

(d) Occlusion and masks of different colors

(e) Selfie faces

(f) An obscure faces by thing
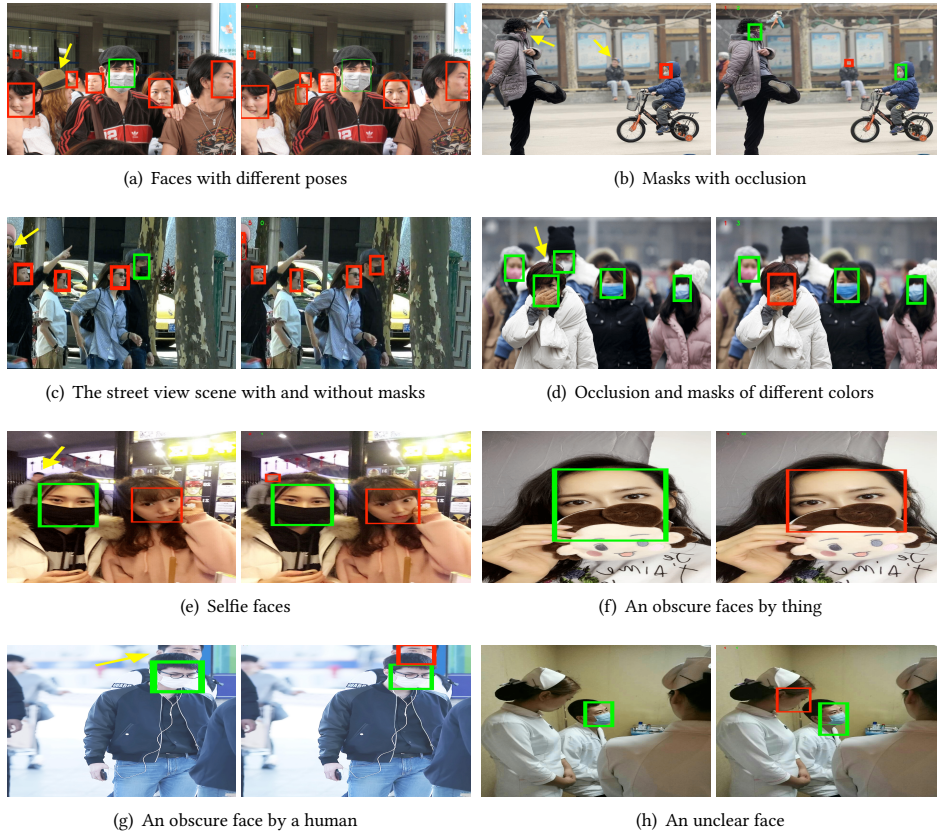
(g) An obscure face by a human

(h) An unclear face

Fig. 5. Demonstrations of the detection results on the target image on the Mask class, where the results of the baseline and ours are shown on the left-hand side and right-hand side, respectively. Best views in color.

Table 3. State-of-the-art comparison

Model	Wider Face			MAFA
	Easy Set	Medium Set	Hard Set	
ResNet50 [34]	59.4%	48.9%	26.5%	94.5%
ResNet18	72.6%	71.4%	46.0%	91.5%
CBAM sequence spatial-channel	19.4%	19.6%	10.4%	22.4%
ResNet18 + triplet loss	76.4%	74.0%	47.1%	91.0%
ResNet18 + consistency loss	74.0%	72.2%	45.6%	91.8%
ResNet18 + triplet loss + consistency loss	75.8%	73.0%	46.3%	91.4%
CenterFace (all modules)	90.2%	85.4%	54.1%	91.5%
ResNet18 + self-supervised rotation [22]	65.5%	63.0%	38.3%	91.4%
ResNet18 + self-supervised grayscale [80]	60.2%	61.4%	39.0%	91.5%
ResNet18 + self-supervised random crop [80]	60.3%	58.9%	37.8%	91.9%

(a) Flipped faces

(b) Small faces

(c) Hidden faces by things

(d) Occlusion faces by humans
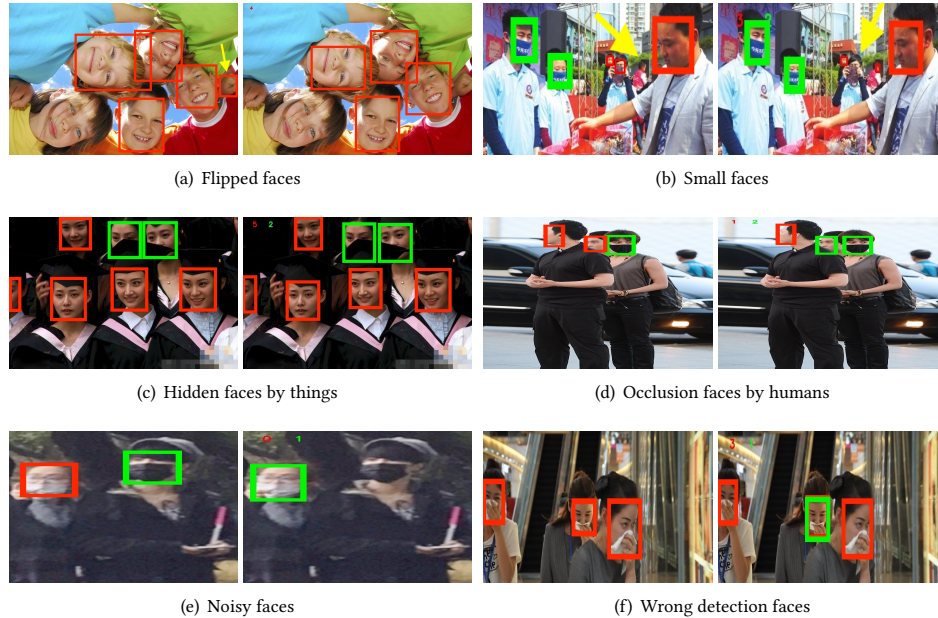
(e) Noisy faces

(f) Wrong detection faces

Fig. 6. Failure cases of both classes, where the results of the baseline and ours are shown on the left-hand side and right-hand side, respectively.

MAFA datasets to show the results. We compare with the recently state-of-the-art detector in WiderFace and MAFA test-dev in Table 3. We also adopt three objective loss functions in Eq. 9 to evaluate their individual performance.

Please note that the ResNet backbone here is utilized alone without the attention mechanism to certify the performance of objective loss functions. The results show that our module outperforms the baseline without adding any objective losses. Afterward, we add the proposed losses one-by-one to observe their effects on the accuracy. This study shows that the proposed objective loss functions significantly boost the performance on the WiderFace dataset.

Attention Mechanism. In this study, We compare the two different ways of arranging the channel and spatial attention sub-module in the convolution neural network architecture. As each module has a different setting function, the order may affect the overall performance. For example, the spatial attention works locally while the channel attention is applied globally. Indeed, in Table 3 our experiment shows that it yields a lower accuracy if we arrange the sub-modules channel and spatial in different ways. The reason is that in traditional architecture the channel attention focus on "what" is an informative part, e.g., the coordinate part in the feature space, which is complementary to the spatial attention. The informative parts then concatenate with the computed pooling from the spatial attention module to generate the efficient feature descriptor. the vise-versa architecture of the attention module may make the descriptor miss to learn some crucial features such as corner area or flipped angle of objects.

Comparisons with Self-Supervised Learning. In addition, we conduct another experiment for demonstrating the performance of consistency loss. Specifically, we use the other widely-used self-supervised pretext task for face mask detection, i.e., rotation [22]. Furthermore, following the previous work [80], we also apply other augmentation methods like grayscale and random crop as the self-supervision tasks. Table 3 shows that all other self-supervision tasks decrease the accuracy of the framework on the WiderFace dataset while keeping a similar accuracy on the MAFA dataset. This is

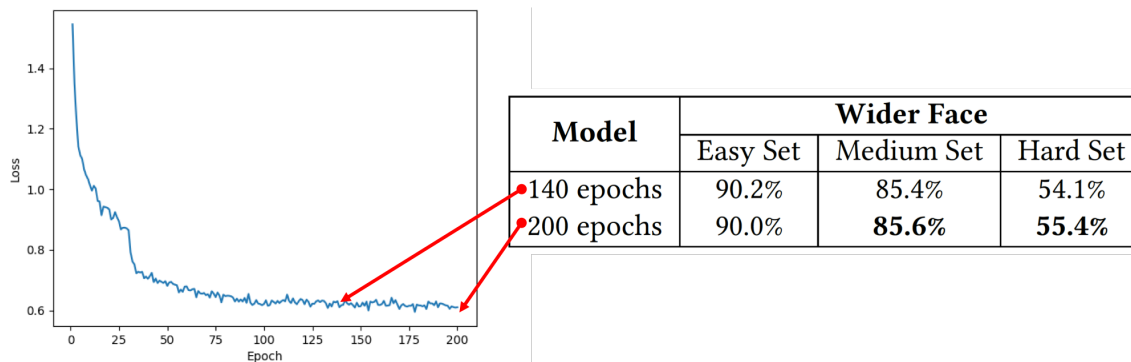| Model	Wider Face		
	Easy Set	Medium Set	Hard Set
140 epochs	90.2%	85.4%	54.1%
200 epochs	90.0%	**85.6%**	**55.4%**

Fig. 7. Training curve (loss with different numbers of epochs).

Table 4. Training performance comparison, where L1 loss works a bit better than Smooth L1 loss.

Model	Wider Face			MAFA
	Easy Set	Medium Set	Hard Set	
Smooth L1	90.0%	85.3%	**55.2%**	91.2%
L1	**90.2%**	**85.4%**	54.1%	**91.5%**

because of the characteristics of the dataset, i.e., the WiderFace dataset includes people in different angles and various appearance diversity while most images in the MAFA dataset are the small group or single people with or without masks. As such, knowing the knowledge from the self-supervised tasks cannot further improve or even deteriorate the performance. It is also worth noting that the objects are different when applying the self-supervised learning in the original paper. Therefore, knowing the rotation angles means that the model understands the semantics of the objects. However, the subjects in the face mask detection are almost about people. As such, knowing the rotation does not mean that the model understand the semantic information. In contrast, imposing the consistency constraint on the prediction results of original images and their horizontally-flipped images simultaneously stabilizes the prediction results and enlarges the training data. Therefore, the consistency loss facilitates the learning process.

Training Curve. The default number of training epochs is 140 epochs with a learning drop at 30 and 80 epochs. We further train the model to 200 epochs to observe if there is any improvement in performance. As shown in Figure 7, the results manifest that the medium and hard sets of the WiderFace dataset further increases by 0.2% and 0.3% while the result of MAFA dataset increases by 0.2%. The improvement is relatively minor while the training requires much longer training time. Therefore, to keep the performance of the algorithm without wasting the computational resource, we suggest to train the the proposed model with 140 epochs.

Regression Loss. The regression loss is the distance loss, which enforces the distance between the anchor to the positive and negative pair. Certainly, the smooth L1 loss is usually a common choice for classification, regression, and distance loss in object detection problems. Similar with the observation from previous works [65, 66], L1 loss yields better accuracy at a fine scale as compared to Smooth L1 loss [39]. We compare an L1 loss to a Smooth L1 Loss in Table 4. The results show that the L1 loss is a bit better than Smooth L1 loss. It is because the L1 loss has a higher tolerance and thus is more robust to the noise. In contrast, the Smooth L1 loss is smoother but sensitive to outliers.

Table 5. Sensitivity test of different hyperparameters λ_s and λ_{con}.

Weights	Wider Face			MAFA
	Easy Set	Medium Set	Hard Set	
$\lambda_s = 0.01$	90.2%	85.4%	54.1%	91.5%
$\lambda_s = 0.02$	89.9%	84.9%	53.4%	91.7%
$\lambda_s = 0.2$	90.0%	85.5%	54.7%	90.6%
$\lambda_{con} = 1$	89.9%	85.3%	54.4%	91.5%
$\lambda_{con} = 10$	89.9%	84.9%	53.4%	91.2%
$\lambda_{con} = 100$	90.2%	85.4%	54.1%	91.5%

Sensitivity Test. The default values of hyperparameters are set to $\lambda_{pix}, \lambda_{off} = 1$, $\lambda_s = 0.01$, $\lambda_{tri} = 1$, and $\lambda_{con} = 100$. Table 5 shows the sensitivity test on different λ_s and λ_{con}, which manifests that the accuracy slightly changes with different hyperparameter weights. Therefore, λ_s and λ_{con} are respectively set to 0.01 and 100 according to the results.

5 CONCLUSIONS

To prevent the massive infections of the coronavirus and reduce the overloading of healthcare, we propose a new framework named "CenterFace" to automate the monitoring of wearing face masks. "CenterFace" uses the context attention module to enable the effective attention of the feed-forward convolution neural network by adapting their attention maps feature refinement. Moreover, we further propose an anchor-free detector with Triplet-Consistency Representation Learning by integrating the consistency loss and the triplet loss to deal with the small-scale training data and the similarity between masks and occlusions. Experimental results show that "CenterFace" outperforms the other state-of-the-art methods, which is released as a public download to improve public health. In the future, we plan to study the problem of face recognition with masks and jointly consider the privacy issues.

ACKNOWLEDGMENTS

REFERENCES

[1] Daniell Chiang AIZOOTech. [n.d.]. AIZOOTech/FaceMaskDetection. https://github.com/AIZOOTech/FaceMaskDetection
[2] Serge Belongie, Jitendra Malik, and Jan Puzicha. 2002. Shape matching and object recognition using shape contexts. *IEEE transactions on pattern analysis and machine intelligence* 24, 4 (2002), 509–522.
[3] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. 2020. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934* (2020).
[4] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. 2020. YOLOv4: YOLOv4: Optimal Speed and Accuracy of Object Detection. *arXiv* (2020).
[5] Neel Ramakant Borkar and Sonia Kuwelkar. 2017. Real-time implementation of face recognition system. In *2017 International Conference on Computing Methodologies and Communication (ICCMC)*. IEEE, 249–255.
[6] Zhaowei Cai and Nuno Vasconcelos. 2018. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 6154–6162.
[7] Jiale Cao, Hisham Cholakkal, Rao Muhammad Anwer, Fahad Shahbaz Khan, Yanwei Pang, and Ling Shao. 2020. D2det: Towards high quality object detection and instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11485–11494.

[8] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2017. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7291–7299.

[9] Qiang Chen, Yingming Wang, Tong Yang, Xiangyu Zhang, Jian Cheng, and Jian Sun. 2021. You Only Look One-level Feature. *arXiv preprint arXiv:2103.09460* (2021).

[10] Yihong Chen, Zheng Zhang, Yue Cao, Liwei Wang, Stephen Lin, and Han Hu. 2020. RepPoints V2: Verification Meets Regression for Object Detection. In *Neural Information Processing Systems (NeurIPS)*.

[11] G Jignesh Chowdary, Narinder Singh Punn, Sanjay Kumar Sonbhadra, and Sonali Agarwal. 2020. Face mask detection using transfer learning of inceptionv3. In *International Conference on Big Data Analytics*. Springer, 81–90.

[12] Tao Dai, Jianrui Cai, Yongbing Zhang, Shu-Tao Xia, and Lei Zhang. 2019. Second-order attention network for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 11065–11074.

[13] Navneet Dalal and Bill Triggs. 2005. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, Vol. 1. IEEE, 886–893.

[14] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. 2019. CenterNet: Keypoint Triplets for Object Detection. In *IEEE International Conference on Computer Vision (ICCV)*.

[15] Gamaleldin Elsayed, Simon Kornblith, and Quoc V Le. 2019. Saccader: improving accuracy of hard attention models for vision. In *Advances in Neural Information Processing Systems*. 702–714.

[16] Pedro Felzenszwalb, David McAllester, and Deva Ramanan. 2008. A discriminatively trained, multiscale, deformable part model. In *2008 IEEE conference on computer vision and pattern recognition*. IEEE, 1–8.

[17] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. 2009. Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence* 32, 9 (2009), 1627–1645.

[18] Cheng-Yang Fu, Mykhailo Shvets, and Alexander C Berg. 2019. RetinaMask: Learning to predict masks improves state-of-the-art single-shot detection for free. *arXiv preprint arXiv:1901.03353* (2019).

[19] Jenil Gathani and Krish Shah. 2020. Detecting Masked Faces using Region-based Convolutional Neural Network. In *2020 IEEE 15th International Conference on Industrial and Information Systems (ICIIS)*. IEEE, 156–161.

[20] Shiming Ge, Jia Li, Qiting Ye, and Zhao Luo. 2017. Detecting masked faces in the wild with lle-cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2682–2690.

[21] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. 2018. Unsupervised representation learning by predicting image rotations. In *The International Conference on Learning Representations (ICLR)*.

[22] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. 2018. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728* (2018).

[23] Ross Girshick. 2015. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*. 1440–1448.

[24] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 580–587.

[25] Ross Girshick, Pedro Felzenszwalb, and David McAllester. 2011. Object detection with grammar models. *Advances in neural information processing systems* 24 (2011), 442–450.

[26] Ross B Girshick, Pedro F Felzenszwalb, and David McAllester. 2012. Discriminatively trained deformable part models, release 5. (2012).

[27] et al. Grill, J. B. 2020. Bootstrap your own latent: A new approach to self-supervised learning. In *Advances in Neural Information Processing Systems (NeurIPS)*. Ieee.

[28] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9729–9738.

[29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[30] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).

[31] Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7132–7141.

[32] Jisoo Jeong, Seungeui Lee, Jeesoo Kim, and Nojun Kwak. 2019. Consistency-based semi-supervised learning for object detection. In *Advances in neural information processing systems*. 10759–10768.

[33] Zhang Jian and Song Wan-Juan. 2010. Face detection for security surveillance system. In *2010 5th International Conference on Computer Science & Education*. IEEE, 1735–1738.

[34] Mingjie Jiang, Xinqi Fan, and Hong Yan. 2020. RetinaMask: A Face Mask detector. arXiv:2005.03950 [cs.CV]

[35] Aniruddha Srinivas Joshi, Shreyas Srinivas Joshi, Goutham Kanahasabai, Rudraksh Kapil, and Savyasachi Gupta. 2020. Deep Learning Framework to Detect Face Masks from Video Footage. In *2020 12th International Conference on Computational Intelligence and Communication Networks (CICN)*. IEEE, 435–440.

[36] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[37] Tao Kong, Fuchun Sun, Huaping Liu, Yuning Jiang, Lei Li, and Jianbo Shi. 2020. Foveabox: Beyound anchor-based object detection. *IEEE Transactions on Image Processing* 29 (2020), 7389–7398.

[38] Ashu Kumar, Amandeep Kaur, and Munish Kumar. 2019. Face detection techniques: a review. *Artificial Intelligence Review* 52, 2 (2019), 927–948.

[39] Hei Law and Jia Deng. 2018. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 734–750.

[40] Youngwan Lee and Jongyoul Park. 2020. Centermask: Real-time anchor-free instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 13906–13915.

[41] Nancy HL Leung, Daniel KW Chu, Eunice YC Shiu, Kwok-Hung Chan, James J McDevitt, Benien JP Hau, Hui-Ling Yen, Yuguo Li, Dennis KM Ip, JS Malik Peiris, et al. 2020. Respiratory virus shedding in exhaled breath and efficacy of face masks. *Nature medicine* 26, 5 (2020), 676–680.

[42] Shaohui Lin, Ling Cai, Xianming Lin, and Rongrong Ji. 2016. Masked face detection via a modified LeNet. *Neurocomputing* 218 (2016), 197–202.

[43] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2117–2125.

[44] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*. 2980–2988.

[45] Chuanbin Liu, Hongtao Xie, Zhengjun Zha, Lingyun Yu, Zhineng Chen, and Yongdong Zhang. 2020. Bidirectional Attention-Recognition Model for Fine-Grained Object Classification. *IEEE transactions on multimedia* 22, 7 (2020), 1785–1795.

[46] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*. Springer, 21–37.

[47] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. 2020. Object-centric learning with slot attention. *Advances in Neural Information Processing Systems* 33 (2020).

[48] David G Lowe. 1999. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision*, Vol. 2. Ieee, 1150–1157.

[49] David G Lowe. 2004. Distinctive image features from scale-invariant keypoints. *International journal of computer vision* 60, 2 (2004), 91–110.

[50] Hanan A Hosni Mahmoud and Hanan Abdullah Mengash. 2020. A novel technique for automated concealed face detection in surveillance videos. *Personal and Ubiquitous Computing* (2020), 1–12.

[51] Shaobo Min, Hantao Yao, Hongtao Xie, Zheng-Jun Zha, and Yongdong Zhang. 2020. Multi-Objective Matrix Normalization for Fine-Grained Visual Recognition. *IEEE transactions on image processing* 29 (2020), 4996–5009.

[52] I. Misra and L. V. D. Maaten. 2020. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

[53] Alexander Neubeck and Luc Van Gool. 2006. Deep residual learning for image recognition. In *Proceedings of IEEE International Conference on Pattern Recognition (ICPR)*.

[54] Alejandro Newell, Kaiyu Yang, and Jia Deng. 2016. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*. Springer, 483–499.

[55] George Papandreou, Tyler Zhu, Nori Kanazawa, Alexander Toshev, Jonathan Tompson, Chris Bregler, and Kevin Murphy. 2017. Towards accurate multi-person pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4903–4911.

[56] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*. 8026–8037.

[57] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You Only Look Once: Unified, Real-Time Object Detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[58] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 779–788.

[59] Joseph Redmon and Ali Farhadi. 2017. YOLO9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7263–7271.

[60] Joseph Redmon and Ali Farhadi. 2018. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767* (2018).

[61] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.

[62] S. Ren, K. He, R. Girshick, and J. Sun. 2017. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 6 (2017), 1137–1149.

[63] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 815–823.

[64] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. 2019. Deep High-Resolution Representation Learning for Human Pose Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[65] Xiao Sun, Jiaxiang Shang, Shuang Liang, and Yichen Wei. 2017. Compositional human pose regression. In *Proceedings of the IEEE International Conference on Computer Vision*. 2602–2611.

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

[67] Yi Sun, Xiaogang Wang, and Xiaoou Tang. 2015. Deeply learned face representations are sparse, selective, and robust. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2892–2900.

[68] Yaniv Taigman, Ming Yang, Marc'Aurelio Ranzato, and Lior Wolf. 2014. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1701–1708.

[69] Mingxing Tan, Ruoming Pang, and Quoc V Le. 2020. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10781–10790.

[70] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. 2019. FCOS: Fully Convolutional One-Stage Object Detection. In *IEEE International Conference on Computer Vision (ICCV)*.

[71] Koen EA Van de Sande, Jasper RR Uijlings, Theo Gevers, and Arnold WM Smeulders. 2011. Segmentation as selective search for object recognition. In *2011 International Conference on Computer Vision*. IEEE, 1879–1886.

[72] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.

[73] Paul Viola and Michael Jones. 2001. Rapid object detection using a boosted cascade of simple features. In *Proceedings of the 2001 IEEE computer society conference on computer vision and pattern recognition. CVPR 2001*, Vol. 1. IEEE, I–I.

[74] Paul Viola and Michael J Jones. 2004. Robust real-time face detection. *International journal of computer vision* 57, 2 (2004), 137–154.

[75] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. 2017. Residual attention network for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3156–3164.

[76] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. 2019. Deep High-Resolution Representation Learning for Visual Recognition. *TPAMI* (2019).

[77] Tiancai Wang, Tong Yang, Martin Danelljan, Fahad Shahbaz Khan, Xiangyu Zhang, and Jian Sun. 2020. Learning human-object interaction detection using interaction points. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4116–4125.

[78] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. 2018. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*. 3–19.

[79] Bin Xiao, Haiping Wu, and Yichen Wei. 2018. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*. 466–481.

[80] Guodong Xu, Ziwei Liu, Xiaoxiao Li, and Chen Change Loy. 2020. Knowledge distillation meets self-supervision.. In *Proceedings of the European conference on computer vision (ECCV)*.

[81] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. 2016. WIDER FACE: A Face Detection Benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[82] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. 2019. RepPoints: Point Set Representation for Object Detection. In *IEEE International Conference on Computer Vision (ICCV)*.

[83] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. 2020. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9759–9768.

[84] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. 2019. Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems* 30, 11 (2019), 3212–3232.

[85] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. 2020. Distance-IoU loss: Faster and better learning for bounding box regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 12993–13000.

[86] Xingyi Zhou, Jiacheng Zhuo, and Philipp Krahenbuhl. 2019. Bottom-up object detection by grouping extreme and center points. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 850–859.

[87] Zhenyao Zhu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2014. Recover canonical-view faces in the wild with deep neural networks. *arXiv preprint arXiv:1404.3543* (2014).

[88] C Lawrence Zitnick and Piotr Dollár. 2014. Edge boxes: Locating object proposals from edges. In *European conference on computer vision*. Springer, 391–405.

[89] Zhengxia Zou, Zhenwei Shi, Yuhong Guo, and Jieping Ye. 2019. Object detection in 20 years: A survey. *arXiv preprint arXiv:1905.05055* (2019).