

# SyNet: An Ensemble Network for Object Detection in UAV Images

Berat Mert Albaba  
Dept. of Mechanical Engineering  
Bilkent University  
Ankara, Turkey  
mert.albaba@bilkent.edu.tr

Sedat Ozer  
Dept. of Computer Engineering  
Bilkent University  
Ankara, Turkey  
sedat@cs.bilkent.edu.tr

**Abstract**—Recent advances in camera equipped drone applications and their widespread use increased the demand on vision based object detection algorithms for aerial images. Object detection process is inherently a challenging task as a generic computer vision problem, however, since the use of object detection algorithms on UAVs (or on drones) is relatively a new area, it remains as a more challenging problem to detect objects in aerial images. There are several reasons for that including: (i) the lack of large drone datasets including large object variance, (ii) the large orientation and scale variance in drone images when compared to the ground images, and (iii) the difference in texture and shape features between the ground and the aerial images. Deep learning based object detection algorithms can be classified under two main categories: (a) single-stage detectors and (b) multi-stage detectors. Both single-stage and multi-stage solutions have their advantages and disadvantages over each other. However, a technique to combine the good sides of each of those solutions could yield even a stronger solution than each of those solutions individually. In this paper, we propose an ensemble network, SyNet, that combines a multi-stage method with a single-stage one with the motivation of decreasing the high false negative rate of multi-stage detectors and increasing the quality of the single-stage detector proposals. As building blocks, CenterNet and Cascade R-CNN with pretrained feature extractors are utilized along with an ensembling strategy. We report the state of the art results obtained by our proposed solution on two different datasets: namely MS-COCO and visDrone with %52.1  $mAP_{IoU=0.75}$  is obtained on MS-COCO *val2017* dataset and %26.2  $mAP_{IoU=0.75}$  is obtained on VisDrone *test-set*. Our code is available at: <https://github.com/mertalbaba/SyNet>

**Index Terms**—Object Detection, Ensemble methods, Deep learning, UAV images

## I. INTRODUCTION

Object detection is an essential yet a challenging task in computer vision field, despite the recent advances using deep learning based techniques. While object detection demonstrated significant success in many applications, those applications are mostly limited in domains using ground taken images. *ImageNet* [1], *COCO* [2] and *PASCAL* [3] datasets are such major datasets driving the deep learning based object detection algorithms and they contain images taken predominantly on the ground. Consequently, they do not include a large variance of the object properties as observed from air. Furthermore, such state of the art object detection algorithms are designed considering the problems associated with the ground images mainly. That is why even the most recent works focusing

on aerial images such as [4], [5], [6] cannot reach to the level where the state of the art object detection algorithms perform today on ground images. With the motivation behind deep neural networks and how object detection is being performed, performance of the state-of-the-art object detectors suffer mainly from two reasons in challenging settings: (1) the images captured has a non-uniform distribution in terms of object types and (2) scale, orientation, shape and the features of the objects can significantly differ from the ones appearing in the ground images.

In recent studies considering class imbalance, most studies addressed the class imbalance problem from considering the imbalance between the foreground and background of the image samples [7], [8]. As another form of class imbalance, the imbalance between the foreground objects also effects detection performance in majority of algorithms, as some of those relevant issues were addressed with different data augmentation techniques as in [8]. Considering the aerial image setting and the existing image samples in *VisDrone* dataset [7] (which we also use in this paper), class imbalance between the foreground objects remains as a common issue that effects the prediction performance negatively. The problem of class imbalance also shows its significance in ground-level view, which can be observed in *MS-COCO* dataset [2]. As also discussed in [9], scale of the objects in datasets is another issue effecting the detection performance.

Typically, each existing object detection algorithm has its advantages and disadvantages over the other existing detection algorithms. Therefore, combining the advantageous properties of multiple detection algorithms can yield a higher performance and that is what we mainly focus on and introduce in this paper for aerial images. We propose an ensemble network that selects the good properties (good decisions) of two recent state of the art detection algorithms namely Cascade R-CNN [10] and CenterNet [11] to achieve even better state of the art results on aerial images. In order to deal with small objects for which assessing detection performance using Intersection over Union (IoU) metric would not be trivial, benefits of the Cascade R-CNN [10] algorithm are utilized in our proposed solution. In addition to that, recently proposed CenterNet [11] is also used in our proposed solution, as CenterNet is another effective detection algorithm with good detection rates

at different scales. CenterNet is oriented at the center of the detected objects and that enables the method to be more robust to scale differences.

Our proposed approach includes two steps for detection: (i) we employ an image augmentation technique inspired from [12] to deal with the imbalance problem making our proposed approach more robust to classes with lower number of samples. (ii) we introduce a synergistic network which combines the predictions of multiple detectors and chooses the most confident score through a weighted bounding box fusion [13] method to detect objects. Combining those two steps offers state of the art prediction performances, when compared to previous recent efforts in the field [14], [15], [11], [10] by considering and dealing with two important problems: the class imbalance problem and the scaling problem.

As verified by our experiments, our proposed method SyNet achieves state of the art performance in object detection, when compared to the performances of the state-of-the-art methods on the VisDrone [7] dataset. In addition to the aerial images, we also observed that proposed approach works well and achieves higher performance for object detection in ground-based images, as our experiments with *MS-COCO* dataset [16] demonstrates. Overall, our main contributions can be listed as:

- introduction of a synergistic approach for better object detection that combines the advantageous properties of two individual and state of the art object detectors,
- demonstrating the performance of our proposed object detection method not only on a drone dataset: *VisDrone* but also on a ground-image based dataset: *MS-COCO*,
- employing an augmentation technique to deal with class imbalance problem in aerial images.

## II. RELATED WORK

Object detection has received significant amount of attention in the last decade. In this section, most relevant work to ours are summarized under two subcategories: (i) Object Detection and (ii) Image (data) Augmentation.

### A. Object Detection

Following the success of deep convolutional neural networks (CNNs) in image classification [17], CNNs are also used for a more challenging task: object recognition and localization, i.e. object detection. Firstly, single forward-pass networks are proposed for object recognition, which are referred as single-stage detectors. In [18], object detection is considered as a regression problem. After dividing image into grids containing pixels, bounding box and class predictions are made for each cell of the grid and good performance is seen in [19]. In [20], batch normalizations and convolutions with anchor boxes are introduced into [18]. Then, in [21], anchor boxes, fixed size bounding boxes which are used as reference boxes, are utilized along with multi-task loss for object detection and a great performance is achieved in different extensive datasets. In [22], ResNet101 is introduced into SSD and deconvolution layers are used to obtain better detection method, which can perform better for small-scale

objects. In [23], a new loss called *FocalLoss* is utilized to solve the class imbalance problem in the object detection. Then, in [24], a module that refines anchor boxes are proposed. In [25], anchor boxes are fully eliminated, objects are detected as two keypoints defining the bounding box and corner pooling is introduced. In [26], for each class, logistic classifiers are utilized and better feature extractor is used when compared to previous version [20]. Recently, in [11], objects are detected as only a single keypoint, that contains the center location of the object, and other properties such as size are regressed, which results good performance with less computational complexity.

In contrast to single-stage detectors, there are also multi-stage detectors, i.e. double-stage detectors, which divides detection problem into two sub-problems: 1. region proposals and 2. bounding box regression. In [27], nearly 2000 region proposals are generated through selective search and after that, support vector machines and CNN are used to regress bounding boxes and class labels. Then, in [28], region of interest pooling is introduced into [27] and instead of finding regions directly on the image, regions are generated from the feature vector of the image. In [29], for the region proposals, a new approach is presented: Region Proposal Network (RPN), which is a CNN for producing region proposals and separately trainable. Several reference boxes are utilized in [29]. Region proposals generated by RPN are fed into the region of interest pooling layer and a detection network generates bounding box and class label predictions from pooled region proposals. In [30], position sensitive region of interest pooling layers are proposed and region proposals are generated from the full image, instead of from several sub-regions. In [15], region of interest alignment is introduced into [29] and better results are achieved.

In addition to single-stage and multi-stage detectors, there also exists some cascading approaches utilizing sequentially cascaded networks, such as [31] and [10]. In [10], sequential detectors whose quality levels increases with the increase in the depth, i.e. stage number, is proposed and great improvements are seen in state-of-the-art double-stage detectors which are explained previously.

### B. Data Augmentation

To solve the class imbalance problem, data augmentation is a popular technique that is used by applying fundamental operations to images such as rotation, scaling and flipping. Moreover, augmentation methods provide an enrichment on the data that the technique is used. In addition to traditional augmentation techniques like flipping, rotating, translating or coloring, different data augmentation methods are proposed.

In [12], objects which should be detected by object detectors are cutted, i.e. copied, and pasted on several background images, along with altering pixel artifacts to generate new training samples. A technique that randomly removes some parts of the images to create new training samples with occlusions is proposed in [32]. In [33], a overlaying images method named *SamplePairing* is proposed which mixes two different images and takes average values of pixels. In [34],

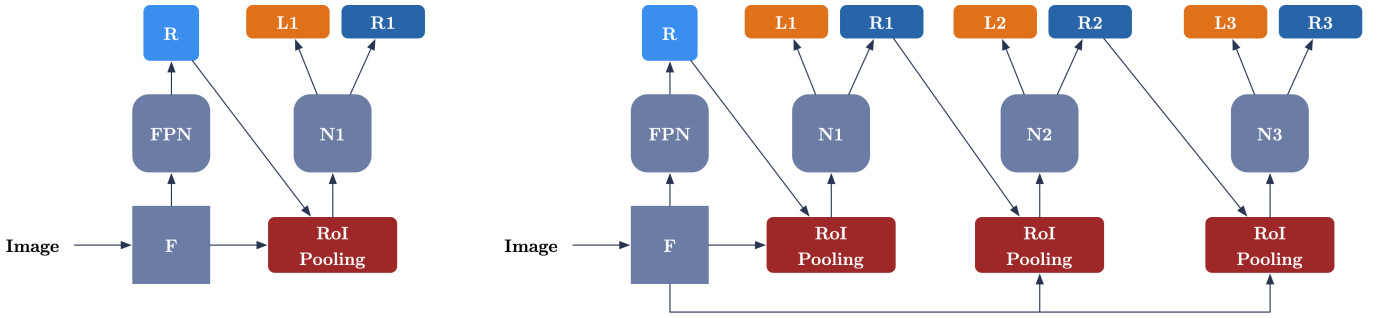


Fig. 1. Architectures of the Faster R-CNN (on the left) and Cascade R-CNN (on the right).  $F$  is the feature extractor,  $N_i$ s are network bodies,  $R$  is the region proposals,  $L_i$ s are label predictions and  $R_i$ s are the bounding box predictions.

generative adversarial neural networks (GANs) are utilized to produce new training samples by mixing different images. In [35], GANs are utilized for data augmentation and it is showed that the generated data may be more beneficial than the original ones. In [36], spatial fusion generative adversarial neural networks (SF-GANs) are used for new image generation from original ones. In order to overcome the problems of distribution in aerial images, a data augmentation approach is proposed by [8], which resulted with significant increase in prediction performance in terms of mAP score. Lastly, in [6] as an augmentation technique, AdaResampling resolves the problems of scale and background mismatch.

### III. METHODS

For object detection, different detectors with different architectures, some of which have multiple steps, are utilized in the literature (as discussed in the previous section). In this work, a method that combines single-stage and multi-stage object detection architectures synergistically is proposed. Furthermore, image augmentation is utilized in order to enhance the results. Therefore, next, we describe the details of the used methods in our proposed solution.

#### A. Object Detection

In this paper, a neural network based object detector is proposed that aims to localize and recognize several objects on given images. A detector computes a bounding box, which, typically, is a quadruple  $(x_{tl}^p, y_{tl}^p, x_{br}^p, y_{br}^p)$  where  $x_{tl}^p, y_{tl}^p$  are the coordinates of the top-left corner of the prediction,  $p$ , and  $x_{br}^p, y_{br}^p$  are the coordinates of the bottom right corner. For evaluating the quality of the prediction, as a common metric, intersection of union ( $IoU$ ) is utilized, which can be defined as  $IoU(b^p, b^{gt}) = \frac{a(b^p \cap b^{gt})}{a(b^p \cup b^{gt})}$ , where  $a$  is the area function that calculates the area of the given bounding box,  $(x_{tl}, y_{tl}, x_{br}, y_{br})$ ,  $b^p$  is the predicted bounding box and  $b^{gt}$  is the true bounding box. Fig. 2 shows the general process.

For the evaluation of the prediction based on the  $IoU$  metric, a threshold  $t$  is selected and if the  $IoU(b^p, b^{gt})$  score is larger than the threshold, prediction is considered as the successful. Usually, prediction threshold  $t$  is referred as the quality level of the detection network.

In our proposed solution, two individual detectors are utilized: Cascade R-CNN and CenterNet. Next, we describe each of those detectors.

1) *Cascade R-CNN*: Cascade R-CNN is a cascaded network, that contains multiple repeated networks, which are connected sequentially, and based on the Faster R-CNN [14] method. In other words, Cascade R-CNN is the sequentially repeated version of the Faster R-CNN. In this work, Cascade R-CNN is utilized as the first object detection building block. Detailed explanation of Cascade R-CNN can be found in [10] and this method is briefly explained here.

For object detection, there are two main problems: (1) a detection network optimized for a certain intersection of union (IoU) score  $q$ , may not perform well in the test dataset, since the network is optimized for  $q$  and IoU scores of the test samples may significantly differ from  $q$ ; and (2) number of samples decreases seriously with the increase in the overall IoU scores of the samples and this makes a high quality (IoU score) detector prone to over-fitting problem since the number of samples are not enough. To overcome these problems, a cascaded network is proposed to increase the number of high IoU score samples and allow the detector to perform well for different IoU scores, whose architecture is presented in Fig. 1.

For the bounding-box predictions, a single detector,  $d$ , tries to minimize bounding-box regression loss [10], which is defined as

$$L(d) = \sum_i G(b_i^p, b_i^{gt}) \quad (1)$$

where

$$G(b^1, b^2) = H(x_{tl}^1 - x_{tl}^2) + H(y_{tl}^1 - y_{tl}^2) + H(x_{br}^1 - x_{br}^2) + H(y_{br}^1 - y_{br}^2) \quad (2)$$

and

$$H = \begin{cases} 0.5x^2, & \text{if } |x| < 1 \\ |x| - 0.5, & \text{otherwise} \end{cases} \quad (3)$$

in which  $b_i^p$  is the bounding box prediction of  $i^{th}$  object,  $b_i^{gt}$  is the true bounding box of  $i^{th}$  object, and  $(x_{tl}, y_{tl})$  and  $(x_{br}, y_{br})$  are the top-left corner coordinates and the bottom-right coordinates of a bounding box, respectively. In cascaded structure, as shown in Fig. 1,  $N$  bounding box regressors are

utilized instead of a single regressor and the final bounding box predictions are generated as

$$b^p = d_N \circ d_{N-1} \circ \dots \circ d_1(x, b^{gt}) \quad (4)$$

where  $x$  is the input image and  $\circ$  is an iterative operator such that each regressor uses the outputs of the previous regressor. As showed in [10], this allows the network to perform well in different IoU scores of inputs since each regressor are trained with different sample qualities and optimized for different IoU scores. Furthermore, in Cascade R-CNN, a detector  $d_n$  at stage  $n$  minimizes:

$$L(x_n, b^{gt}) = L_C(c_n(x_n), y_n) + I[y_n \geq 1]L_L(d_n(x_n, b_n^p), b^{gt}) \quad (5)$$

where  $c_n$  is the classifier at stage  $n$ ,  $b_n^p = d_{n-1}(x_{n-1}, b_{n-1}^p)$ ,  $x_n$  is the input to stage  $n$ ,  $y_n$  is the label of  $x_n$ ,  $L_C(c_n(x_n), y_n)$  is classification loss defined as

$$L_C(c_n(x_n), y_n) = -\log(p(y = y_n|x)) \quad (6)$$

and  $L_L(d_n(x_n, b_n^p), b^{gt})$  is defined as

$$L_L(d_n(x_n, b_n^p), b^{gt}) = G(d_n(x_n, b_n^p), b^{gt}). \quad (7)$$

In this structure, since the outputs of the previous detector  $d_{n-1}$ , which is optimized for IoU score  $1_{n-1}$  are fed to stage  $n$ , which is optimized for  $q_n$  such that  $q_n > q_{n-1}$ , IoU score of the generated samples are gradually increased. In other words, as going deeper in the cascaded structures, number of high IoU score samples increases, which provides enough samples to higher level detectors and prevent over-fitting issue. To clarify, IoU score level increases as going deeper in the cascaded structure since the IoU score of inputs increases.

As a summary, by resampling the initial samples, the IoU scores of the samples are increased sequentially and a network robust to over-fitting is obtained. Moreover, since detectors at different stages are optimized for different IoU scores, mismatches between training and test performances is eliminated at a higher level. Hence, Cascade R-CNN offers high quality detection results.

2) *CenterNet*: In most of the object detection architectures, objects are detected by utilizing some anchors, which are fixed size bounding boxes. In contrast, as shown in Fig. 2, CenterNet [37] approaches object detection as a keypoint estimation problem, which allows more flexible, faster object detection.

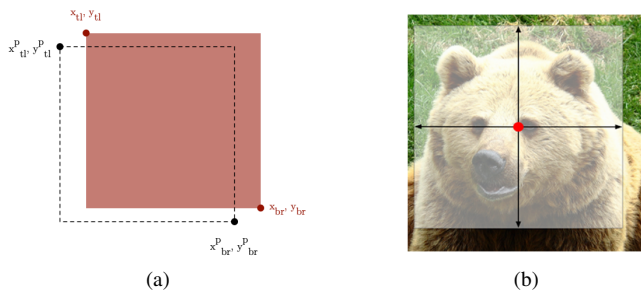


Fig. 2. (a) The object (red box) and the predicted bounding box (dashed box). (b) Estimated center point and sizes of the object by CenterNet.

In CenterNet approach, first of all, a heatmap of keypoints is predicted as  $Y^p$ , in which 1 represents a keypoint detection and 0 is background label. For a true keypoint tuple  $gt_i = (x_i^{gt}, y_i^{gt}) \in R^2$ , a heatmap  $Y_i^{gt}$  is created as

$$Y_i^{gt} = e^{-\frac{(x_i^{gt} - k_x^{quant_i})^2 - (y_i^{gt} - k_y^{quant_i})^2}{2\sigma_{k^{quant_i}}^2}} \quad (8)$$

where  $k^{quant_i} = \lfloor x_i^{gt}/R, y_i^{gt}/R \rfloor$ , where  $R$  is the output stride, and  $\sigma$  is standard deviation which is size adaptive. Then, as explained in [37], a network that predicts keypoints is trained by minimizing the loss defined as

$$L = \frac{-1}{K} \sum_{gt_i} g[Y_i^{gt}] (1 - Y_{gt_i}^p)^2 \log(Y_{gt_i}^p) + g[Y_i^{gt}] (1 - Y_i^{gt})^4 (Y_{gt_i}^p)^2 \log(1 - Y_{gt_i}^p) \quad (9)$$

where  $g(x) = x$  if  $x = 1$  and  $g(x) = 1 - x$  otherwise and  $K$  is number of keypoints. In addition to keypoints, an offset,  $O_{gt_i}$  is predicted for each point in order to recover error from discretization by minimizing  $L_o = \frac{1}{K} \sum_{gt_i} |O_{k^{quant_i}} - k^{quant_i} - gt_i/R|$ .

After keypoint predicting network is trained, for each object, keypoints are utilized as center points of the objects, which have a bounding box represented as quadruple  $(x_{tl}, y_{tl}, x_{br}, y_{br})$ . For each object classes,  $c$ , width and height predictions are made as  $v_c^p = (w_c^p, h_c^p)$  to predict true width and height of  $i^{th}$  object from class  $c$ ,  $(v^{gt_c})_i = (x_{tl}^i - x_{br}^i, y_{tl}^i - y_{br}^i)$  by minimizing

$$L_v = \frac{1}{n} \sum_i |v_c^p - (v^{gt_c})_i| \quad (10)$$

for each class. Then, a single loss is defined as weighted sum of  $L$ ,  $L_v$  and  $L_o$ , which is used in a single network for all predictions [37]. For bounding box predictions, detected local maximum keypoints are used as center points for each class and a bounding box quadruple is predicted as  $(x_{tl}^i + o^i - w^i/2, y_{tl}^i + q^i - h^i/2, x_{br}^i + o^i + w^i/2, y_{br}^i + q^i + h^i/2)$ , where  $(o, q)$  is the offset predictions and  $(w, h)$  is the size predictions. As result, a powerful and fast object detector is obtained which does not require any post-processing like bounding box suppressions and any fixed anchor boxes.

3) *Weighted Bounding Box Fusion*: For the combination of different methods, an ensemble method named weighted boxes fusion approach [13] is used. As explained in [13], instead of removing some bounding box predictions like non-maximum suppression, this method utilizes all predictions in order to find bounding box clusters.

Initially, clustered boxes are stored in  $C$ .  $C_i$  contains the matched boxes whose labels are  $i$ . Iteratively, each bounding box prediction with  $j^{th}$  label is compared with  $F_j$ , which is the fusion of boxes in  $C_j$  and represented by quintuple  $(s, x_{tl}, x_{br}, y_{tl}, y_{br})$ , where  $s$  is the confidence score. When a

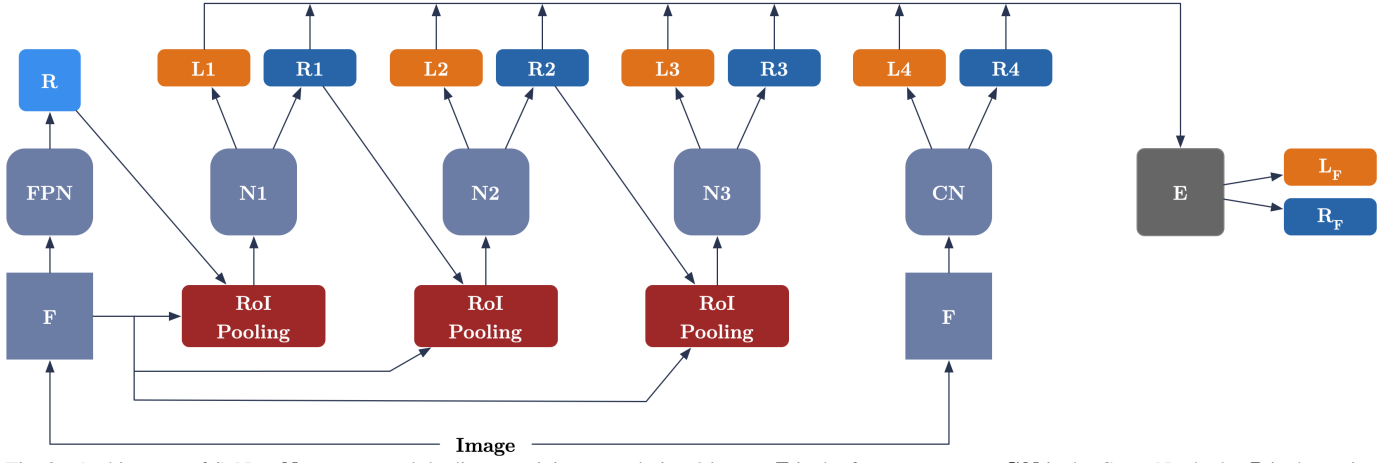


Fig. 3. Architecture of SyNet.  $N_i$ s are network bodies containing convolutional layers,  $F$  is the feature extractor,  $CN$  is the CenterNet body,  $R$  is the region proposals,  $L_i$ s are label predictions,  $R_i$ s are bounding box predictions,  $E$  is the ensemble function and  $L_F, R_F$  are final label and bounding box predictions, respectively.

box  $b$  is matched with  $F_j$  and added to  $C_j$ , fusion box  $F_j$  confidence is recalculated as

$$F_j(s) = \frac{\sum_{A_k \in C_j} A_k(s)}{N} \quad (11)$$

where  $N$  represents the total number of boxes in  $C_j$  and  $A_k$  is bounding box sample, i.e. a quintuple  $(s, x_{tl}, x_{br}, y_{tl}, y_{br})$ . Coordinates of the fusion box are updated as follows:

$$F_j(x_1) = \frac{\sum_{A_k \in C_j} A_k(s) A_k(x_{tl})}{s_1 + s_2 + \dots + s_N} \quad (12)$$

$$F_j(x_2) = \frac{\sum_{A_k \in C_j} A_k(s) A_k(x_{br})}{s_1 + s_2 + \dots + s_N} \quad (13)$$

$$F_j(y_1) = \frac{\sum_{A_k \in C_j} A_k(s) A_k(y_{tl})}{s_1 + s_2 + \dots + s_N} \quad (14)$$

$$F_j(y_2) = \frac{\sum_{A_k \in C_j} A_k(s) A_k(y_{br})}{s_1 + s_2 + \dots + s_N} \quad (15)$$

where  $s_k = A_k(s)$  is the confidence score of  $A_k$ . As these equations show, confidence scores of each bounding box is utilized as weights in order to obtain meaningful final bounding box for each class. After all the predicted boxes are compared with clusters sequentially, box confidence scores are rescaled as explained in [13], where it is presented that when compared to common methods such as [38] and [39].

## B. SyNet

In this work, a synergistic approach is proposed that combines a single stage detector with a multi-stage detector for better object detection. The main motivation is that multi-stage detectors tend to produce more false negatives, which means that multi-stage detectors fails to detect some objects. On the other hand, single-stage detectors generally propose more bounding boxes with less quality. Hence, the combination of these two different architectures may predict more bounding

boxes than multi-stage detectors and quality of the single-stage detector predictions may be increased by the multi-stage one. Thus, an approach that combines Cascade R-CNN [10] and CenterNet [37] through weighted box ensemble [13] is proposed.

General architecture of the proposed method is presented in Fig. 3, where *Image* represents the images in training set, *FPN* is the feature pyramid network [40] for initial features, *R* is the proposal generated by *FPN*, *F* is the feature extractor network, *RoIPooling* is the region of interest pooling layer as described in [29], *N1, N2, N3* are the network bodies consisting several convolutional layers, *CN* is the CenterNet network body, *L1, L2, L3, L4* are the label predictions, *R1, R2, R3, R4* defines the bounding box predictions, *E* is the weighted box ensemble function that ensembles inputs based on weighted bounding box fusion as explained in the previous subsection and  $L_F, R_F$  are the final class and bounding box predictions, respectively. As feature extractors, ResNet101 [41] pretrained on ImageNet [42] is used in Cascade R-CNN and DLA [43] is utilized in CenterNet.

First of all, images in the train dataset are augmented by including additional images, which are generated with the procedure explained in Section 3.1. Then, these are fed into both CenterNet and Cascade R-CNN and these networks are trained. For the Cascade R-CNN, as the feature extractor, ResNet 101 is used and for CenterNet, DLA is utilized. After the trainings, bounding box predictions of these two networks are obtained for each element in the inference set and these are combined with weighted box fusion method. As result, final bounding box predictions are obtained. Since the proposed approach requires the training of CenterNet and Cascade R-CNN separately, memory consumption is significantly higher than both Cascade R-CNN and CenterNet. However, time consumption is similar to that of Cascade R-CNN.

Although CenterNet and Cascade R-CNN are utilized as single-stage and multi-stage detectors, other detectors may also



Fig. 4. (a) The original image. (b) The augmented version.

be introduced into the network.

### C. Image Augmentation

With the rapid development of object detectors, the need for extensive data with several local and global features is increased extremely. For this purpose, in this work, data augmentation approach presented in [12] utilized.

In this approach, the main goal is that generating augmented *real-like* images to allow object detectors work better. As explained in [12], this method consists four main steps. At first, images of objects, which should be predicted by the object detector, are collected. In this work, object images in datasets are utilized. To exemplify, for VisDrone dataset [44], for *car* class, image parts limited by bounding box coordinates whose labels are *car* are collected as object instances. After this, scene images should be collected for this approach and in this work, as scene images, raw training dataset images are utilized. Then, masks should be predicted for object images collected at first step. For this purpose, convolutional neural networks are utilized in [12], however, in this work, object segmentations in datasets are used as masks after converting them to binary images if available. Lastly, convolution of masks and objects are placed on randomly chosen scene images.

In this work, for each image in a dataset, two augmented images, that contains more than 10 and less than 30 pasted additional objects are created to improve training. Fig. 4 shows an example of the augmentation.

## IV. DATASETS

In this work, for the evaluation of the proposed approach, two different datasets are used, one of which is an aerial dataset and the other is a ground dataset.

### A. VisDrone Dataset

VisDrone dataset [44] for aerial object detection consists more than 6000 aerial images taken by camera equipped unmanned air vehicles. In total, there are ten classes which are used in the evaluation in both [44] and in this work, which may be listed as: 1. pedestrian, 2. people, 3. bicycle, 4. car, 5. van, 6. truck, 7. tricycle, 8. awning-tricycle, 9. bus and 10. motor. For each sample, bounding box coordinates are given along with truncation and occlusion information. In this work, truncation and occlusion information are not utilized since

TABLE I  
RESULTS ON *test-set* OF VISDRONE [44] DATASET.

	$mAP_C$	$mAP_{0.50}$	$mAP_{0.75}$
SyNet (ours)	<b>25.1</b>	<b>48.4</b>	<b>26.2</b>
Cascade R-CNN	24.7	43.7	24.3
CenterNet	14.3	26.6	13.1

the goal is building a robust detector which may also detect partially truncated or occluded objects. Since the aerial image detection task is still challenging because of object-image size mismatch and class imbalance, this dataset is utilized in this work for the validation of the proposed approach.

For the training, training set is utilized that consists 6471 aerial images. Proposed image augmentation method is applied only to the training set. For the validation, for which the results generated, test set is used that consists 1580 images.

### B. MS-COCO Dataset

For the evaluation of the proposed approach in ground images, Microsoft Common Objects in Context dataset [16] (MS-COCO) is utilized. In this work, 2017 version is used, which contains more than 100000 images from 79 classes. Since the most of the recent state-of-the-art object detection algorithms are evaluated by using MS-COCO, this dataset is selected for evaluation of the proposed method.

For the training, training set consisting more than 118000 images is utilized after applying the proposed image augmentation techniques. For the validation, validation set that contains more than 5000 images is utilized.

## V. EXPERIMENTAL RESULTS

We evaluated the prediction performance of our proposed method on both *test-dev-set* of the VisDrone dataset [44] and *val-2017* set of the MS-COCO dataset [16]. For the evaluation metric, a common metric: mean average precision (mAP) is utilized. Average precision is defined as the area under the precision-recall curve. We computed three different average precision metrics:  $mAP_{0.50}$ ,  $mAP_{0.75}$  and  $mAP_C$ . For  $mAP_{0.50}$ , in order to consider a bounding box prediction as true, intersection of union score (IoU) between the predicted and the ground truth bounding box must be higher than 0.50. Similarly, for  $mAP_{0.75}$ , bounding box predictions with IoU scores higher than 0.75 are considered as true. Lastly, for  $mAP_C$ , average precisions are calculated for different IoU scores varying in the range of: 0.50 : 0.05 : 0.95, and then, the mean value of those computed average precisions is used.

Results for VisDrone dataset are presented in Table 1. To obtain Cascade R-CNN and CenterNet results on VisDrone dataset, those methods are trained on augmented VisDrone dataset. As this table shows, SyNet may be considered as state-of-the-art object detection method for aerial object detection since the proposed approach demonstrates a better performance for  $mAP_C$ ,  $mAP_{0.50}$  and  $mAP_{0.75}$  than the recent state-of-the-art methods, Cascade R-CNN and CenterNet. Furthermore, in order to present the performance of the proposed

TABLE II  
RESULTS ON INDIVIDUAL CLASSES USING VISDRONE [44]  
*test - dev - set (FOR  $mAP_{0.50}$ ).*

	SyNet (ours)	Cascade R-CNN	CenterNet
Pedestrian	<b>48.1</b>	42.6	22.6
People	<b>37.8</b>	33.1	20.6
Bicycle	<b>23.8</b>	21.2	14.6
Car	<b>83.2</b>	79.8	59.7
Van	<b>55.4</b>	49.3	24.0
Truck	<b>49.3</b>	43.5	21.3
Tricycle	<b>34.2</b>	31.6	20.1
Awning	<b>24.2</b>	21.5	17.4
Bus	<b>66.0</b>	61.9	37.9
Motor	<b>44.8</b>	43.1	23.7

TABLE III  
RESULTS ON *val - 2017 SET OF MS-COCO [16] DATASET.*

	Backbone	$mAP_C$	$mAP_{50}$	$mAP_{75}$
SyNet (ours)	ResNet101 + DLA-34	<b>47.2</b>	<b>66.4</b>	<b>52.1</b>
Cascade R-CNN	ResNet101	42.7	61.6	46.6
	ResNet50	40.3	59.4	43.7
CenterNet	Hourglass-104	40.3	59.1	44.0
	DLA-34	37.4	55.1	40.8
Faster R-CNN	ResNet101	38.5	60.3	41.6
	ResNet50	36.4	58.4	39.1
Mask R-CNN	ResNet101	39.4	60.9	43.3
	ResNet50	37.3	59.0	40.2
Retina Net	ResNet101	37.7	57.5	40.4
	ResNet50	35.6	55.5	38.3
Cascade Mask R-CNN	ResNet101	42.6	60.7	46.7
	ResNet50	41.2	59.1	45.1
Hybrid Task	ResNet101	44.9	63.8	48.7
Cascade	ResNet50	43.2	62.1	46.8

network, average precision scores are presented for each class in VisDrone dataset in Table 2, which also shows that SyNet performs better than Cascade R-CNN and CenterNet.

Results for MS-COCO dataset are presented in Table 2. Results for other methods are taken from [10], [45] and [37]. SyNet achieved the best performance in the literature and outperformed many state-of-the-art methods for  $mAP_C$ ,  $mAP_{0.50}$  and  $mAP_{0.75}$  metrics. However, a version of the EfficientDet [46], referred as EfficientDet-D7 (1536) achieved better result on  $mAP_C$  metric, whose  $mAP_{0.50}$  and  $mAP_{0.75}$  results on validation set are not calculated thus omitted in the table. As result, SyNet not only performs well in aerial images, but performs well in both aerial images and ground images.

Besides showing some quantitative results, we also demonstrate some qualitative results as shown in Figure 5. For instance, second example shows that there are four vans in the image, Cascade R-CNN and CenterNet predicted only one of them correctly but SyNet predicted two of them correctly.

## VI. CONCLUSION

In this paper, we introduced a novel ensemble network for object detection yielding state of the art results when compared to the existing techniques on drone images. Our proposed method, SyNet, combines a multi-stage detector and a single-stage detector and makes a prediction by combining

the individual predictions of each algorithm through a fusion stage. By carefully choosing the the individual algorithms, we made use of our their advantageous properties in detection and combined them in our ensemble network. As demonstrated in our experiments, our proposed approach yields the highest  $mAP_C$  and  $mAP_{0.75}$  scores in both datasets: VisDrone and in MS-COCO, when compared to most recently proposed state of the art detection algorithms.

While our SyNet performs better in Table II and yields higher results than Cascade R-CNN and CenterNet solutions in all classes on VisDrone dataset, as the results of both Table III and Table II suggests, the average  $mAP_{0.50}$  value in Table III remain significantly lower than the average  $mAP_{0.50}$  value of Table II. That analysis suggests that the object detection task with the existing state of the art object detection algorithms still remains as a more challenging problem in drone images, when compared to detecting objects in ground taken images.

## ACKNOWLEDGMENT

This paper has been produced benefiting from the 2232 International Fellowship for Outstanding Researchers Program of TÜBİTAK (Project No:118C356). However, the entire responsibility of the paper belongs to the owner of the paper. The financial support received from TÜBİTAK does not mean that the content of the publication is approved in a scientific sense by TÜBİTAK.

## REFERENCES

- [1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [2] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollar, "Microsoft coco: Common objects in context," *arXiv preprint arXiv:1405.0312v3*, 2014.
- [3] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [4] Z. Wu, K. Suresh, P. Narayanan, H. Xu, H. Kwon, and Z. Wang, "Delving into robust object detection from unmanned aerial vehicles: A deep nuisance disentanglement approach," *arXiv preprint arXiv:1908.03856*, 2019.
- [5] S. Hong, S. Kang, and D. Cho, "Patch-level augmentation for object detection in aerial images," *The IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [6] C. Chen, Y. Zhang, Q. Lv, S. Wei, X. Wang, X. Sun, and J. Dong, "Rrnet: A hybrid detector for object detection in drone-captured images," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2019, pp. 0–0.
- [7] P. Zhu, L. Wen, D. Du, X. Bian, Q. Hu, and H. Ling, "Vision meets drones: Past, present and future," *arXiv preprint arXiv:2001.06303*, 2020.
- [8] S. Hong, S. Kang, and D. Cho, "Patch-level augmentation for object detection in aerial images," *The IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [9] F. Yang, H. Fan, P. Chu, E. Blasch, and H. Ling, "Clustered object detection in aerial images," in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [10] Z. Cai and N. Vasconcelos, "Cascade r-cnn: High quality object detection and instance segmentation," *arXiv preprint arXiv:1906.09756*, 2019.
- [11] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, "Centernet: Keypoint triplets for object detection," pp. 6569–6578, 2019.
- [12] D. Dwibedi, I. Misra, and M. Hebert, "Cut, paste and learn: Surprisingly easy synthesis for instance detection," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1301–1310.
- [13] R. Soloviyev and W. Wang, "Weighted boxes fusion: ensembling boxes for object detection models," *arXiv preprint arXiv:1910.13302*, 2019.

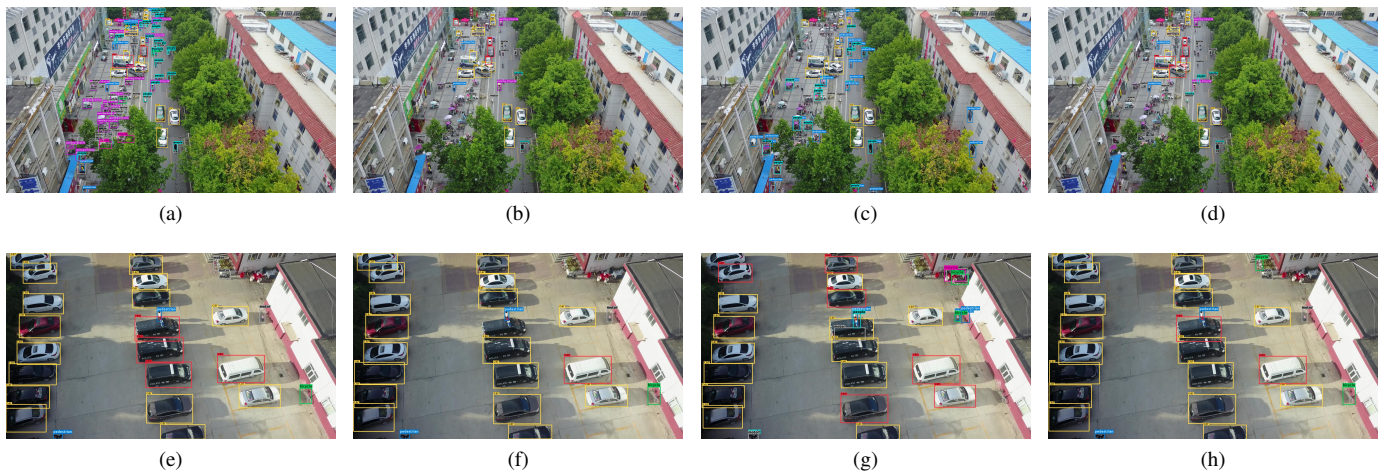


Fig. 5. Results on two sample images, in which only bounding boxes with more than 0.5 confidence score are shown: (a) Ground truth 1, (b) Cascade R-CNN result 1, (c) CenterNet result 1, (d) SyNet result 1, (e) Ground truth 2, (f) Cascade R-CNN result 2, (g) CenterNet result 2 and (h) SyNet result 2.

- [14] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [15] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [16] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [18] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [19] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results," <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>.
- [20] J. Redmon and A. Farhadi, "Yolo9000: better, faster, stronger," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7263–7271.
- [21] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision*. Springer, 2016, pp. 21–37.
- [22] C.-Y. Fu, W. Liu, A. Ranga, A. Tyagi, and A. C. Berg, "Dssd: Deconvolutional single shot detector," *arXiv preprint arXiv:1701.06659*, 2017.
- [23] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [24] S. Zhang, L. Wen, X. Bian, Z. Lei, and S. Z. Li, "Single-shot refinement neural network for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4203–4212.
- [25] H. Law and J. Deng, "Cornersnet: Detecting objects as paired keypoints," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 734–750.
- [26] J. R. A. Farhadi and J. Redmon, "Yolov3: An incremental improvement," *Retrieved September*, vol. 17, p. 2018, 2018.
- [27] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 580–587.
- [28] R. Girshick, "Fast r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1440–1448.
- [29] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [30] J. Dai, Y. Li, K. He, and J. Sun, "R-fcn: Object detection via region-based fully convolutional networks," in *Advances in neural information processing systems*, 2016, pp. 379–387.
- [31] S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia, "Path aggregation network for instance segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8759–8768.
- [32] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, "Random erasing data augmentation," *arXiv preprint arXiv:1708.04896*, 2017.
- [33] H. Inoue, "Data augmentation by pairing samples for images classification," *arXiv preprint arXiv:1801.02929*, 2018.
- [34] D. Liang, F. Yang, T. Zhang, and P. Yang, "Understanding mixup training methods," *IEEE Access*, vol. 6, pp. 58 774–58 783, 2018.
- [35] F. H. K. d. S. Tanaka and C. Aranha, "Data augmentation using gans," *arXiv preprint arXiv:1904.09135*, 2019.
- [36] F. Zhan, H. Zhu, and S. Lu, "Spatial fusion gan for image synthesis," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3653–3662.
- [37] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as points," *arXiv preprint arXiv:1904.07850*, 2019.
- [38] H. Zhou, Z. Li, C. Ning, and J. Tang, "Cad: Scale invariant framework for real-time object detection," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2017, pp. 760–768.
- [39] N. Bodla, B. Singh, R. Chellappa, and L. S. Davis, "Soft-nms—improving object detection with one line of code," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5561–5569.
- [40] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [41] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [42] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*, 2009.
- [43] F. Yu, D. Wang, E. Shelhamer, and T. Darrell, "Deep layer aggregation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2403–2412.
- [44] D. R. Pailla, "Visdrone-det2019: The vision meets drone object detection in image challenge results," 2019.
- [45] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu *et al.*, "Mmdetection: Open mmlab detection toolbox and benchmark," *arXiv preprint arXiv:1906.07155*, 2019.
- [46] M. Tan, R. Pang, and Q. V. Le, "Efficientdet: Scalable and efficient object detection," *arXiv preprint arXiv:1911.09070*, 2019.