

# Rethinking the Backbone Architecture for Tiny Object Detection

Jinlai Ning<sup>a</sup>, Haoyan Guan<sup>b</sup> and Michael Spratling<sup>c</sup>

*Department of Informatics, King's College London, London, U.K.*

**Keywords:** Tiny Object Detection, Backbone, Pre-Training.

**Abstract:** Tiny object detection has become an active area of research because images with tiny targets are common in several important real-world scenarios. However, existing tiny object detection methods use standard deep neural networks as their backbone architecture. We argue that such backbones are inappropriate for detecting tiny objects as they are designed for the classification of larger objects, and do not have the spatial resolution to identify small targets. Specifically, such backbones use max-pooling or a large stride at early stages in the architecture. This produces lower resolution feature-maps that can be efficiently processed by subsequent layers. However, such low-resolution feature-maps do not contain information that can reliably discriminate tiny objects. To solve this problem we design “bottom-heavy” versions of backbones that allocate more resources to processing higher-resolution features without introducing any additional computational burden overall. We also investigate if pre-training these backbones on images of appropriate size, using CIFAR100 and ImageNet32, can further improve performance on tiny object detection. Results on TinyPerson and WiderFace show that detectors with our proposed backbones achieve better results than the current state-of-the-art methods.

## 1 INTRODUCTION

Tiny object detection is a sub-field of object detection in the field of computer vision, and has many applications including maritime search-and-rescue, surveillance, and driving assistance. It is an active area of research because classical methods of standard object detection (Ren et al., 2015; Liu et al., 2015; Lin et al., 2017a; Cai and Vasconcelos, 2018; Redmon and Farhadi, 2018; Tian et al., 2019) that work well on datasets such as MS COCO (Lin et al., 2014) and Pascal VOC (Everingham et al., 2015) perform poorly on tiny object datasets such as TinyPerson (Yu et al., 2020) and WiderFace (Yang et al., 2016). Detecting tiny objects is still a challenging task for these standard methods. As a result, many methods have been designed specifically for tiny object detection (Tong and Wu, 2022). As described in (Tong and Wu, 2022), these existing methods adapt standard object detection frameworks to be more suitable for tiny object detection by: using super-resolution techniques (Yang et al., 2019; Tang et al., 2018); exploiting contextual information (Hong et al., 2022); using data augmen-

tation techniques (Yun et al., 2019; Yu et al., 2020; Jiang et al., 2021); employing multi-scale representation learning approaches (Lin et al., 2017b; Gong et al., 2021); using anchor mechanisms more appropriate for small objects (Zhang et al., 2017; Zhang et al., 2018); designing training strategies specific to small objects (Singh et al., 2018; Krishna and Jawahar, 2017); or using loss functions specific for small and tiny objects (Liu et al., 2021a; Lin et al., 2017c). In this article we propose a different approach that does not fit into any of these categories.

Although previous approaches vary from each other in specific details, they all depend on standard deep neural network backbones, such as ResNet (He et al., 2015), to extract features. The features are then used for feature fusion, position regression and category classification. Despite the impressive results previous algorithms have obtained, we believe that by re-deploying the same standard backbones as are used for general object detection (Ren et al., 2015; Lin et al., 2017a; Cai and Vasconcelos, 2018; Tian et al., 2019), these methods suffer from poor feature extraction for tiny object detection. The standard backbone architectures down-sample the feature maps rapidly in the first few layers, and this down-sampling removes much of the information about tiny objects that were present in the original image. As a result, subsequent

<sup>a</sup> <https://orcid.org/0000-0002-0460-7657>

<sup>b</sup> <https://orcid.org/0000-0003-1936-2442>

<sup>c</sup> <https://orcid.org/0000-0001-9531-2813>

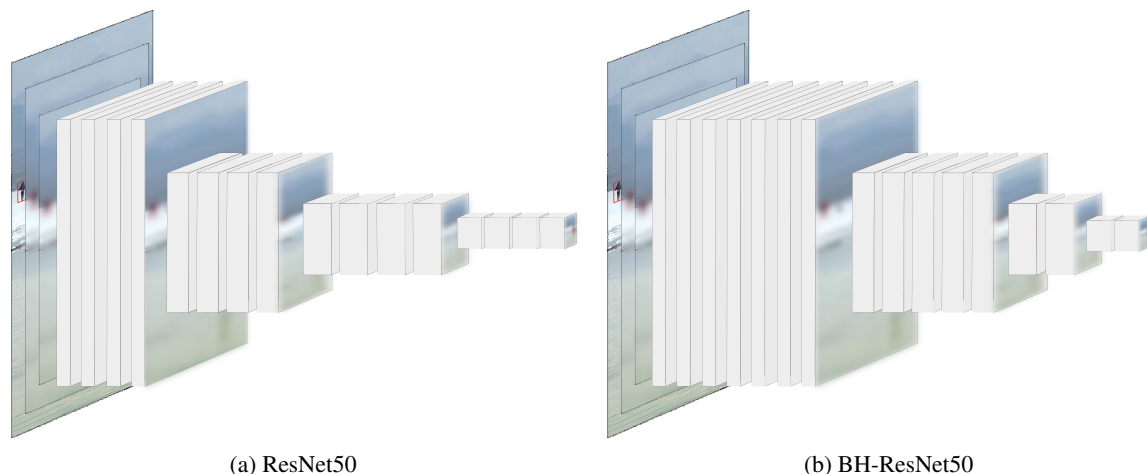


Figure 1: An illustration of (a) an existing backbone, and (b) our modified backbone. Our modified backbone delays down-sampling and changes the number of layers at different depths.

layers in the backbone, and all the methods for improving tiny object detection that fall into the seven categories mentioned in the preceding paragraph, can only work with features that are relatively uninformative and poor at identifying tiny objects.

Down-sampling is a widely applied operation that has proven highly effective in convolutional neural networks (CNNs). It can be achieved using pooling layers that summarize the features in each patch, or by using a convolutional layer with a stride greater than one (Springenberg et al., 2015). Down-sampling can improve translation invariance, avoid over-fitting, and decrease computational costs. For tiny object detection, down-sampling itself is not an issue. Rather it is the improper use of down-sampling that results in poor feature-extraction. Tiny objects occupy very few pixels and down-sampling could potentially remove important features that identify such objects. The only way to preserve information about small features is for convolutional filters in the earliest layers to encode these features and pass this information on to subsequent layers. However, in existing backbones the number of convolutional filters in the early layers is kept to a minimum to reduce the computational burden, and this likely means that not all discriminative features of tiny objects are able to be preserved. For instance, ResNets decrease feature map size by a factor of 4 in less than or equal to 2 layers of convolutions. Using such backbones to handle tiny object detection is likely to result in information about tiny objects disappearing in the feature maps before it is fully extracted.

Outside the domain of object detection, previous work has addressed information loss caused by down-sampling in a number of different ways. Zhao *et al.* (Zhao et al., 2017) introduced random shifting

in down-sampling during the training process to suppress contextual information loss without extra calculation cost. Ren *et al.* (Jiahuan et al., 2021) proposed a low-rank feature recovery module to try to recover the lost information. In other sub-fields of computer vision, such as hyper-spectral target detection (Bhandari and Tiwari, 2021), and pluralistic image inpainting (Liu et al., 2022), several methods have been designed to handle information loss. Similarly, the chess-playing AI, AlphaGo (Silver et al., 2016), did not use layers with strides bigger than one to avoid losing any spatial information about the chessboard. However, in the domain of tiny object detection, this problem has previously not been considered or has been overlooked due to the convenience of using the same backbone as has been used in general-purpose object detectors (Ren et al., 2015; Liu et al., 2015; Redmon et al., 2015).

In this paper we address this neglected problem and propose alternative backbone architectures that are more appropriate for tiny object detection. We show that simple changes to the architecture (as illustrated in Fig. 1) that delay down-sampling operations, and move more convolutional filters to earlier layers where they can process high-resolution information, can bring clear improvement to object detection performance. We demonstrate the effectiveness of this strategy by making the same modifications to two different backbone architectures: ResNet (He et al., 2015) and HRNet (Wang et al., 2021b). Experiments on TinyPerson (Yu et al., 2020) and WiderFace (Yang et al., 2016) demonstrate that replacing standard backbones with our modified versions always results in an improvement in performance for a number of different object detection frameworks. This is achieved despite the modified backbones having fewer param-

Table 1: The size of objects in several typical general object detection datasets (above the double lines) and tiny object detection datasets (below the double lines). Absolute size is defined as the square root of the object’s absolute bounding box area, measured in pixels. It is reported as a mean±standard deviation.

Dataset	Absolute size	Proportion of small/tiny objects
MS COCO	$99.5 \pm 107.5$	32% (with absolute size $\leq 32$ pixels)
PASCAL VOC	—	10% (with absolute size $\leq 32$ pixels)
CityPersons	$79.8 \pm 67.5$	8% (with absolute size $\leq 20$ pixels)
TinyPerson	$18.0 \pm 17.4$	73% (with absolute size $\leq 20$ pixels)
WiderFace	$32.8 \pm 52.7$	56% (with absolute size $\leq 20$ pixels)

ters than the corresponding original ones. These results support our claim that unsuitable use of down-sampling occurs in the backbones commonly used for tiny object detection and that this problem should be taken seriously to help improve detection accuracy.

## 2 RELATED WORK

Object detection aims to locate, using a bounding box, and predict the category of each object in an image. The majority of state-of-the-art algorithms are based on deep learning techniques (Ren et al., 2015; Tan et al., 2020; Tian et al., 2019; Wang et al., 2021a; Liu et al., 2021c), although classic image processing methods have also made contributions in the early years (Dalal and Triggs, 2005; Felzenszwalb et al., 2008). CNN-based networks have dominated the object detection field for many years. These methods can be classified as one-stage (Redmon et al., 2015; Bochkovskiy et al., 2020; Liu et al., 2015; Tian et al., 2019) or two-stage (Girshick et al., 2014; He et al., 2014; Ren et al., 2015) algorithms, depending on whether or not detection is performed by one end-to-end network. Rather than designing bespoke backbone architectures, methods in both groups use a backbone that is a standard CNN architecture that has been pre-trained on a classification task, and then had the fully-connected layers removed.

Large datasets are the basic resources required for deep learning-based methods. For object detection, MS COCO (Lin et al., 2014) and PASCAL VOC (Everingham et al., 2015) are widely used as benchmarks to evaluate the performance of different algorithms. Both of them contain a variety of natural images, containing a large number of different categories, and can be used to test how object-detection methods perform on general tasks.

Tiny object detection is a sub-field of object detection. It focuses on detecting tiny objects in images because such objects are common in real-world scenarios, but are hard to detect with standard object detection methods. The main difference between

tiny object detection datasets such as TinyPerson (Yu et al., 2020) and WiderFace (Yang et al., 2016) and general object detection datasets is the scale of the target objects. As shown in Table 1, tiny object datasets contain more small and tiny objects while general datasets contain objects at a wider range of scales. As mentioned in the Introduction, previous methods for tiny object detection fall into seven categories (Tong and Wu, 2022). Representative examples from each category will be reviewed in the following paragraphs, with results summarized in Tables 2 and 3.

Super-resolution techniques increase the resolution of the image, to enable standard techniques for larger object detection to be applied successfully to images containing tiny targets. Increasing the resolution can be done for the whole image, for example using a generative neural network (Bai et al., 2018), or can be used on a small region of interests (Yang et al., 2019).

Contextual information (*ie.* features from the surrounding region of the image) can be used to help detect small objects. Hence, many techniques include context in the computation. One example is the scale selection pyramid network (SSPNet) that consists of a context attention module (CAM), scale enhancement module (SEM), and scale selection module (SSM) (Hong et al., 2022). CAM includes context information by generating hierarchical attention heat-maps. SEM ensures feature maps of different scales only focus on objects of suitable scale rather than other objects and the background. SSM exploits feature fusion between shallow and deep features to keep the information shared by different layers.

Data augmentation aims to improve performance by extending the size of training dataset through identity-preserving image transformations. Yu *et al.* (Yu et al., 2020) proposed a simple yet efficient augmentation approach named scale match. This method can transform the distribution of object sizes from a general dataset to be similar to a task-specific dataset. In that paper, MS COCO was transformed with scale match and then used as additional data for training a detector for the TinyPerson detection task. One step in scale match is to sample a size-bin from the his-

Table 2: Detection performance of previous methods, including the current state-of-the-art methods, on TinyPerson, sorted by performance on the  $mAP_{tiny}^{50}$  metric.

Method	$mAP_{tiny}^{50}$	$mAP_{tiny1}^{50}$	$mAP_{tiny2}^{50}$	$mAP_{tiny3}^{50}$	$mAP_{small}^{50}$
Faster R-CNN-FPN (Ren et al., 2015)	47.35	30.25	51.58	58.95	63.18
Faster R-CNN-FPN with S- $\alpha$ (Gong et al., 2021)	48.39	31.68	52.20	60.01	65.15
Faster R-CNN-RFLA (Xu et al., 2022)	48.86	30.35	54.15	61.28	66.69
Faster R-CNN-FPN-SM (Yu et al., 2020)	51.33	33.91	55.16	62.58	66.96
Faster R-CNN-FPN-SM+ (Jiang et al., 2021)	51.46	33.74	55.32	62.95	67.37
RetinaNet-SSPNet (Hong et al., 2022)	54.66	42.72	60.16	61.52	65.24
Faster R-CNN with SFRF (Liu et al., 2021b)	57.24	51.59	64.51	67.78	65.33
Cascade R-CNN-SSPNet (Hong et al., 2022)	58.59	45.75	62.03	65.83	71.80
Faster R-CNN-SSPNet (Hong et al., 2022)	59.13	47.56	62.36	66.15	71.17

Table 3: Detection performance of previous methods, including the current state-of-the-art methods, on WiderFace, sorted by performance on the  $mAP_{medium}$  metric.

Method	$mAP_{easy}$	$mAP_{medium}$	$mAP_{hard}$
HR (Hu and Ramanan, 2017)	92.5	91.0	80.6
S3FD (Zhang et al., 2017)	93.7	92.4	85.2
FaceGAN (Bai et al., 2018)	94.4	93.3	87.3
SFA (Luo et al., 2019)	94.9	93.6	86.6
LSFHI (Zhang et al., 2018)	95.7	94.9	89.7
Pyramid-Box (Tang et al., 2018)	96.1	95.0	88.9
RetinaFace (Deng et al., 2020)	96.5	95.6	90.4
TinaFace (Zhu et al., 2020)	96.3	95.7	93.1

togram of object sizes, this sampling is done at random with the probability based on a given image in the external dataset. To avoid the range of sampled size-bins having a big difference with the size of objects in the external image, monotone scale match (MSM) is utilized to sample the range of sizes from small to big monotonically. Jiang *et al.* (Jiang et al., 2021) proposed an enhanced version of scale match, SM+, that improves the scale match from image level to instance level and uses probabilistic structure inpainting (PSI) to handle the background.

Multi-scale representation learning aims to enhance detection performance by making use of features at different scales, represented in different layers of the CNN. This technique is also commonly applied in general object detection, for example through the feature pyramid network (FPN) (Lin et al., 2017b). Gong *et al.* (Gong et al., 2021) argue that the FPN brings not only a positive impact but also a negative one caused by the top-down connections. They design a statistic-based fusion factor to adaptively adjust the weight of different layers during the feature fusion. Tang *et al.* (Tang et al., 2018) argue that not all high-level semantic features are really helpful to smaller targets. They therefore modify FPN to a low-level feature pyramid network (LFPN) that starts the top-down structure from a middle layer rather than the

high layer. Liu *et al.* (Liu et al., 2021b) devise a feature rescaling and fusion (SFRF) network that selects and generates a new resized feature map with a high-density distribution of tiny objects through the use of a Nonparametric Adaptive Dense Perceiving Algorithm (NADPA) module.

Anchors are predefined locations where objects might be detected and are widely applied in two-step detectors such as Faster R-CNN (Ren et al., 2015) and Cascade R-CNN (Cai and Vasconcelos, 2018). Zhang *et al.* (Zhang et al., 2017) guarantees that different scales of anchor have the same density on the image, so that various scales can approximately match the same number of anchors. They also propose a scale compensation anchor matching strategy to ensure that all scales have enough anchors matched. Deng *et al.* (Deng et al., 2020) specify different sizes of anchors in different layers of features to ensure all sizes of objects could be matched properly. Zhang *et al.* (Zhang et al., 2018) develop a single-level face detection framework to specifically detect small faces, which uses different dilation rates for anchors with different sizes and performed anchor-level hard example mining. Xu *et al.* (Xu et al., 2022) design a novel label assignment method to mitigate the problems of lack of positive samples and gap between the uniformly distributed prior anchors and the Gaussian

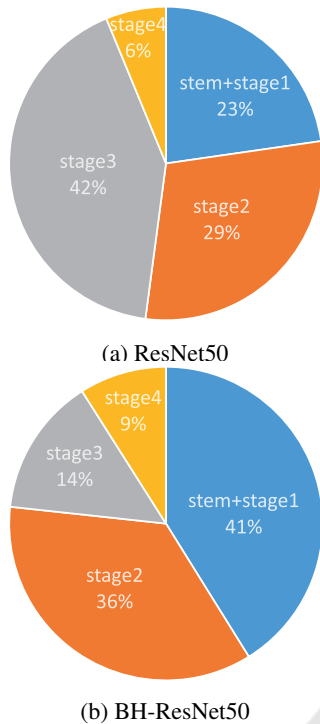


Figure 2: The proportion of FLOPs per stage for (a) standard ResNet50, and (b) our modified BH-ResNet50.

distributed receptive field.

Adapting the training strategy to be more appropriate for tiny object detection can be beneficial. For example, Luo *et al.* (Luo et al., 2019) propose a multi-branch small face attention (SFA) detector. It trains the model with multi-scale training images to improve robustness and performance. Hu *et al.* (Hu and Ramanan, 2017) train separate face detectors efficiently by utilizing features extracted from multiple layers of a single feature hierarchy for different scales.

Optimizing the loss function is an effective strategy to improve the detection performance of tiny objects. Zhu *et al.* (Zhu et al., 2020) modify the loss function of bounding box regression from the widely applied smooth L1 loss (Girshick, 2015) to DIoU loss (Zheng et al., 2020). It ensures the loss function is consistent with the objective of bounding box regression by including the IoU metrics.

Our proposed approach, of modifying the backbone, is orthogonal to all this previous work. Hence, potentially all these diverse techniques might be further improved by combining them with our proposed approach.

### 3 METHODOLOGY

A natural idea to preserve the high-resolution features required for tiny object detection is to dispense with several down-samplings. However, such a simplistic approach would significantly increase the computational complexity of the backbone as the number of calculations increases quadratically with the size of a feature-map. To avoid this issue, our proposed bottom-heavy (BH) architectures perform as much down-sampling as the corresponding original architecture, but the number of convolutional filters is increased in the earlier stages of the network, and decreased in the later stages. This decrease in the number of convolutional filters in the later layers reduces the amount of computation performed in the deeper layers to compensate for the increased computation performed in the earlier layers. Specifically, in our BH architecture, the number of convolutional layers is increased before the feature-map is down-sampled for the third time, and decreased after this point. The proportion of FLOPs performed in different stages of a ResNet50 and a BH-ResNet50 is shown in Fig 2. The change in the distribution of the convolutional layers is designed to ensure that the original backbone and the modified version perform the same number of floating-point operations, to allow a fair comparison of the proposed backbones with existing ones. Another result of these modifications is that the BH networks have fewer parameters than their original counter-parts. Full details are provided in Table 4. These modifications ensure that high-resolution information that can distinguish tiny objects is processed more thoroughly, using more convolutional layers.

## 4 EXPERIMENTS

### 4.1 Pre-Training

Pre-training is a process that initializes model parameters with previously learned ones. For object detection tasks, it is common to start training the whole model with a backbone that has been pre-trained on an image classification task, typically ImageNet. Compared with training the whole detection network from scratch, using a pre-trained backbone is better to maintain a stable and fast training process (Pan and Yang, 2010; Han et al., 2021). Additionally, the learned knowledge from classification also provides a solid foundation for detection. Our backbones are pre-trained on CIFAR-100 (Krizhevsky and Hinton, 2009), ImageNet (Deng et al., 2009) and a down-sampled variant of ImageNet, ImageNet32

Table 4: The architectures of three standard deep neural network backbones (ResNet50, HRNet32, and HRNet18) and the proposed bottom-heavy (BH) versions. Convolutional layers are specified using the notation width  $\times$  height, number of channels, and stride. Pooling layers are specified using the notation width  $\times$  height, type, and stride, where type is M for max pooling. Repeated blocks are shown using squared brackets with the number of repeats indicated by the number following the multiply sign after the brackets. Where stride is unspecified it is equal to one. ‘\*’ indicates the stride of that layer equals two but only for the first block in a set of repeated layers. Rather than illustrating all parallel branches of HRNet, only the deepest branch is shown for clarity. The number of giga-floating-point operations (GFLOPs) is calculated for an input size of 640 pixels  $\times$  512 pixels.

	Stem	Stage 1	Stage 2	Stage 3	Stage 4	Params & GFLOPs
ResNet50	$7 \times 7, 64, 2$ $3 \times 3, M, 2$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128, 2^* \\ 1 \times 1, 512 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1, 256 \\ 3 \times 3, 256, 2^* \\ 1 \times 1, 1024 \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1, 512 \\ 3 \times 3, 512^* \\ 1 \times 1, 2048 \end{bmatrix} \times 3$	23.23M 26.91
BH-ResNet50	$3 \times 3, 64, 1$ $3 \times 3, M, 2$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64, 2^* \\ 1 \times 1, 256 \end{bmatrix} \times 7$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128, 2^* \\ 1 \times 1, 512 \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1, 256 \\ 3 \times 3, 256, 2^* \\ 1 \times 1, 1024 \end{bmatrix} \times 2$	$\begin{bmatrix} 1 \times 1, 512 \\ 3 \times 3, 512, 2^* \\ 1 \times 1, 2048 \end{bmatrix} \times 1$	10.91M 27.12
HRNet32	$3 \times 3, 64, 2$ $3 \times 3, 64, 2$	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 4$	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 4$	$\begin{bmatrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{bmatrix} \times 16$	$\begin{bmatrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{bmatrix} \times 12$	29.31M 51.93
BH-HRNet32	$3 \times 3, 64, 1$ $3 \times 3, 64, 2$ $3 \times 3, M, 2$	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 4$	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 4$	$\begin{bmatrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{bmatrix} \times 8$	$\begin{bmatrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{bmatrix} \times 3$	12.36M 51.89
HRNet18	$3 \times 3, 64, 2$ $3 \times 3, 64, 2$	$\begin{bmatrix} 3 \times 3, 36 \\ 3 \times 3, 36 \end{bmatrix} \times 4$	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 4$	$\begin{bmatrix} 3 \times 3, 72 \\ 3 \times 3, 72 \end{bmatrix} \times 16$	$\begin{bmatrix} 3 \times 3, 144 \\ 3 \times 3, 144 \end{bmatrix} \times 12$	9.56M 21.71
BH-HRNet18	$3 \times 3, 64, 1$ $3 \times 3, 64, 2$ $3 \times 3, M, 2$	$\begin{bmatrix} 3 \times 3, 36 \\ 3 \times 3, 36 \end{bmatrix} \times 4$	$\begin{bmatrix} 3 \times 3, 36 \\ 3 \times 3, 36 \end{bmatrix} \times 4$	$\begin{bmatrix} 3 \times 3, 72 \\ 3 \times 3, 72 \end{bmatrix} \times 12$	$\begin{bmatrix} 3 \times 3, 144 \\ 3 \times 3, 144 \end{bmatrix} \times 6$	5.85M 21.66

Table 5: Classification top-1 accuracy of backbone architectures after pre-training on CIFAR-100, ImageNet 32 and ImageNet.

Method	CIFAR-100	ImageNet32	ImageNet
ResNet50	65.82%	43.12%	75.92%
BH-ResNet50	68.57%	48.75%	74.68%
HRNet18	55.75%	—	76.80%
BH-HRNet18	59.91%	50.68%	76.62%
HRNet32	55.08%	—	78.50%
BH-HRNet32	61.71%	56.00%	78.23%

(Chrabaszcz et al., 2017).

The CIFAR-100 dataset (Krizhevsky and Hinton, 2009) consists of 60000  $32 \times 32$  colour images in 100 classes, separated into 50000 training and 10000 test images. It is chosen for pre-training because: (1) the image size of CIFAR-100 is consistent with the size of target objects in the tiny object detection datasets (2) it has a reasonable training time that is appropriate for prototyping a new backbone; (3) some of the categories in CIFAR-100 such as man, woman, baby, boy, girl, sea etc. are similar to those in the TinyPerson and WiderFace datasets, and hence, it is expected that the parameters will transfer from image recognition to object detection.

We also perform experiments using backbones pre-trained on ImageNet. This allows us to compare the performance of the new architecture with published results for existing object detectors that are pre-trained on ImageNet. ImageNet is a large visual

dataset containing more than 20000 categories and 14 million images. It is considered a standard choice to pre-train a backbone for object detection tasks. In addition, we also repeat some experiments on a down-sampled variant of ImageNet, named ImageNet32. ImageNet32 contains exactly the same number of images as the original ImageNet, but all images are resized to  $32 \times 32$  pixels. This dataset has the same advantages as CIFAR100, but provides more training images.

For data augmentation, we use random horizontal flipping, random cropping and resizing. The classification results at the end of pre-training are shown in Table 5. It can be seen that the proposed BH networks produce more accurate classifications for the datasets with small image sizes, compared to the corresponding original networks. This supports our claim that these bottom-heavy networks are more appropriate for the discrimination of small objects. This improvement in performance on small objects comes at the cost of a small decrease in classification accuracy for large images. There may be scope for improving these pre-training results (and as a consequence improving the image detection results) as no great effort has been taken to optimise the hyper-settings used during pre-training: we had insufficient computational resources to do so, and searching for the best settings is not the main point of this paper.

Table 6: Detection performance of previous methods and our modified (BH) architecture on TinyPerson. Pre-training is done on CIFAR-100.

Method	Backbone	$mAP_{tiny}^{50}$	$mAP_{tiny1}^{50}$	$mAP_{tiny2}^{50}$	$mAP_{tiny3}^{50}$	$mAP_{small}^{50}$	FLOPs
Faster R-CNN-FPN (Ren et al., 2015)	ResNet50	43.02	26.09	47.43	54.24	<b>60.13</b>	75.58
	BH-ResNet50	<b>46.53</b>	<b>31.65</b>	<b>50.51</b>	<b>56.60</b>	59.45	75.79
	HRNet32	37.67	24.62	41.84	47.29	51.39	100.66
	BH-HRNet32	<b>41.82</b>	<b>27.19</b>	<b>46.03</b>	<b>51.85</b>	<b>55.39</b>	100.62
	HRNet18	36.94	24.05	41.46	45.94	50.57	69.33
	BH-HRNet18	<b>40.27</b>	<b>25.56</b>	<b>44.50</b>	<b>50.18</b>	<b>53.65</b>	69.28
Faster R-CNN-SM (Yu et al., 2020)	ResNet50	43.21	27.12	47.44	54.52	59.78	75.58
	BH-R50	<b>46.69</b>	<b>31.19</b>	<b>50.88</b>	<b>57.13</b>	<b>61.09</b>	75.59
Adap-FCOS (ucas-vg, 2020)	ResNet50	36.29	21.70	40.60	46.10	49.87	174.74
	BH-ResNet50	<b>41.42</b>	<b>25.92</b>	<b>45.49</b>	<b>51.33</b>	<b>57.08</b>	174.95
Faster R-CNN-RFLA (Xu et al., 2022)	ResNet50	42.41	24.73	45.39	55.11	61.50	75.59
	BH-ResNet50	<b>45.34</b>	<b>27.09</b>	<b>49.27</b>	<b>58.41</b>	<b>62.63</b>	75.58
RetinaNet-SPPNet (Hong et al., 2022)	ResNet50	48.37	<b>36.78</b>	54.47	55.05	60.74	260.06
	BH-ResNet50	<b>49.23</b>	36.01	<b>54.47</b>	<b>56.17</b>	<b>61.60</b>	260.27
Cascade R-CNN-SSPNet (Hong et al., 2022)	ResNet50	48.79	35.59	50.75	<b>57.59</b>	<b>65.63</b>	186.62
	BH-ResNet50	<b>48.96</b>	<b>37.69</b>	<b>52.15</b>	55.69	62.62	186.83
Faster R-CNN-SSPNet (Hong et al., 2022)	ResNet50	48.05	36.40	50.38	56.51	62.94	158.82
	BH-ResNet50	<b>49.78</b>	<b>39.99</b>	<b>52.26</b>	<b>57.14</b>	<b>64.44</b>	159.03

## 4.2 Fine-Tuning

We trained object detectors using the TinyPerson and WiderFace datasets. TinyPerson has 1610 large-scale seaside images taken from a long distance. It contains 794 and 816 images for training and testing, respectively. The average absolute size of target objects is about 18 pixels and non-small objects (absolute size  $> 32$  pixels) are rare. TinyPerson divides objects into different intervals based on the absolute size (measured in pixels): tiny [2, 20], small [20, 32] and all [2,  $\infty$ ]. The tiny set is partitioned into three (overlapping) subsets: tiny1 [2, 8], tiny2 [8, 12] and tiny3 [12, 20]. The WiderFace dataset is a face detection benchmark containing 32,203 images of different human faces with a high degree of variability in scale, pose and occlusion. The average size of objects in WiderFace is about 32 pixels. It separates objects into three evaluation sets (easy, medium and hard) based on the difficulty of detecting them.

The resolution of the images in TinyPerson ranges from  $497 \times 700$  pixels to  $4064 \times 6354$  pixels. Directly feeding large images into the backbone would produce out-of-memory issues. Meanwhile, simply resizing such images to make them more manageable would risk losing information about tiny objects. Therefore we follow the procedure described in (Yu et al., 2020) for training with this dataset. Specifically, images are cut into overlapping patches of size  $640 \text{ pixels} \times 512 \text{ pixels}$  with 30 pixels overlap. We use the default training schedule for Faster-RCNN and Adap-FCOS(ucas-vg, 2020) on TinyPerson: 12

epochs with stochastic gradient descent (SGD) as the optimizer. The learning rate is initialized to 0.01 and decreased by a factor of 0.1 after 8 and 11 epochs for training on 2 GPUs. The only difference for training different architectures of SSPNet is that the default schedule is 10 epochs. For WiderFace, the input size is  $640 \text{ pixels} \times 640 \text{ pixels}$ . The default training schedule of 630 epochs with SGD (Zhu et al., 2020) is used. The learning rate is initialized to 0.00375 and follows a cosine restart scheduler that restarts the learning rate every 30 epochs. We do not apply any complicated image augmentation methods but follow the same settings used by the corresponding state-of-the-art methods.

## 5 RESULTS

### 5.1 Experimental Setup and Evaluation Metrics

Mean average precision (mAP) is a standard criterion for evaluation in object detection. AP calculates the average precision among a set of different recalls. mAP is the mean AP among all classes. Whether a bounding box is marked as a positive or negative prediction is determined by whether the intersection-over-union (IoU) of the predicted and ground-truth bounding boxes is greater than a set threshold. The mAP on a set of objects  $i$  with the percentage IoU threshold  $j$  is noted as  $mAP_i^j$ . For example,  $mAP_{tiny}^{50}$

Table 7: Detection performance of previous methods and our modified (BH) architecture on TinyPerson. Pre-training is done on ImageNet and downsampled ImageNet (indicated with †). Only results with our proposed backbone are trained by ourselves while all others are previously reported results.

Method	Backbone	$mAP_{tiny}^{50}$	$mAP_{tiny1}^{50}$	$mAP_{tiny2}^{50}$	$mAP_{tiny3}^{50}$	$mAP_{small}^{50}$	FLOPs
Faster R-CNN-FPN (Ren et al., 2015)	ResNet50†	47.81	31.78	53.54	58.34	64.56	75.58
	ResNet50	47.35	30.25	51.58	58.95	63.18	75.58
	BH-ResNet50†	<b>52.60</b>	<b>38.18</b>	<b>58.18</b>	<b>61.66</b>	<b>67.42</b>	75.59
	BH-ResNet50	52.03	37.95	57.58	60.99	65.71	75.59
	HRNet32	53.05	35.61	58.46	63.87	68.02	100.66
	BH-HRNet32†	52.62	<b>38.76</b>	57.64	62.28	67.42	100.62
	BH-HRNet32	<b>53.29</b>	38.08	<b>58.92</b>	<b>63.08</b>	<b>67.72</b>	100.62
	HRNet18	52.28	34.92	58.39	63.36	<b>67.86</b>	69.33
Faster R-CNN-SM (Yu et al., 2020)	BH-HRNet18†	50.04	35.43	55.10	60.04	65.03	69.28
	BH-HRNet18	<b>52.69</b>	<b>36.12</b>	<b>59.13</b>	<b>63.48</b>	67.45	69.28
	ResNet50†	50.65	33.68	55.69	61.40	<b>67.39</b>	75.58
	ResNet50	51.33	33.91	55.16	<b>62.58</b>	66.96	75.58
Adap FCOS (ucas-vg, 2020)	BH-ResNet50†	<b>52.04</b>	<b>37.87</b>	57.01	61.32	65.91	75.59
	BH-ResNet50	51.48	37.35	<b>57.35</b>	60.49	65.47	75.59
	ResNet50†	42.13	22.62	46.50	54.32	57.61	174.74
	ResNet50	47.42	29.05	52.06	59.15	65.06	174.74
Faster R-CNN-RFLA (Xu et al., 2022)	BH-ResNet†	48.20	28.56	53.00	59.87	66.51	174.95
	BH-ResNet	<b>49.16</b>	<b>31.13</b>	<b>53.89</b>	<b>60.09</b>	<b>67.60</b>	174.95
	ResNet50 †	45.78	27.72	49.09	58.33	64.38	75.58
	ResNet50	48.86	30.35	54.15	61.28	66.69	75.58
RetinaNet-SPPNet (Hong et al., 2022)	BH-ResNet50 †	50.07	32.29	54.83	62.60	<b>68.21</b>	75.59
	BH-ResNet50	<b>52.27</b>	<b>34.36</b>	<b>56.62</b>	<b>64.48</b>	67.77	75.59
	ResNet50†	51.28	39.38	56.60	58.55	65.01	260.06
	ResNet50	50.98	38.97	57.26	57.68	63.38	260.06
Cascade R-CNN-SSPNet (Hong et al., 2022)	BH-ResNet50†	<b>55.34</b>	<b>42.45</b>	<b>61.04</b>	<b>62.36</b>	67.63	260.27
	BH-ResNet50	52.70	39.61	58.82	60.35	<b>69.00</b>	260.27
	ResNet50†	54.31	39.25	56.81	63.54	69.40	186.62
	ResNet50	56.60	42.41	59.56	65.36	<b>71.20</b>	186.62
Faster R-CNN-SSPNet (Hong et al., 2022)	BH-ResNet50†	56.54	43.65	59.54	64.45	69.59	186.83
	BH-ResNet50	<b>57.89</b>	<b>44.39</b>	<b>60.86</b>	<b>66.25</b>	71.03	186.83
	ResNet50†	50.48	36.81	51.08	60.87	67.38	158.82
	ResNet50	58.58	44.92	<b>62.22</b>	67.19	71.88	158.82
Faster R-CNN-SSPNet (Hong et al., 2022)	BH-ResNet50†	56.22	46.07	58.03	63.36	69.54	159.03
	BH-ResNet50	<b>58.97</b>	<b>47.22</b>	61.61	<b>67.45</b>	<b>72.37</b>	159.03

represents the mAP among objects in the *tiny* interval (absolute size  $\in [2, 20]$ ) with IoU threshold 0.5. Floating-point operations (FLOPs) are used to measure the computational complexity of a model. The higher the number of FLOPs, the more calculations are performed by the method, for a given size of the image. In Tables 6 and 7, FLOPs are measured for an image size of  $640 \times 512$  pixels.

## 5.2 Comparison to State-of-the-Arts

To evaluate our proposed backbone, we tested several architectures that have previously been shown to produce state-of-the-art performance on TinyPerson and WiderFace, replacing the standard backbone with our proposed bottom-heavy (BH) equivalent. The results for TinyPerson are shown in Tables 6 and 7.

In each case better performance was achieved using the proposed backbones than with the original ones, as summarized in Table 8. This was achieved without increasing the computational complexity and using backbones containing fewer parameters than the originals. It can also be seen from Table 8 that the proposed BH modifications produce a significant improvement in detection performance compared to standard backbones, and that this improvement is highest when using CIFAR-100 and ImageNet32 for pre-training.

A pre-training dataset should be consistent with the target task (Pan and Yang, 2010). Down-sampled ImageNet is a simple alternative to ImageNet that has higher consistency with tiny object detection from the perspective of object size. However, from our experiments, only for some methods did pre-training on Im-



Table 8: The difference of  $mAP_{tiny}$  between the standard backbones and our ‘Bottom Heavy’ versions backbones, calculated by calculating  $mAP_{tiny}$  (for the BH architecture) minus  $mAP_{tiny}$  (for the corresponding standard architecture).

Method	CIFAR-100	ImageNet32	ImageNet
Faster R-CNN (Ren et al., 2015)	+3.51	+4.79	+4.68
Faster R-CNN-SM (Yu et al., 2020)	+3.48	+1.39	+0.15
Adap FCOS (ucas-vg, 2020)	+5.13	+6.07	+1.74
Faster R-CNN-RFLA (Xu et al., 2022)	+2.93	+4.29	+3.41
RetinaNet-SPPNet (Hong et al., 2022)	+0.86	+4.06	+1.72
Cascade R-CNN-SSPNet (Hong et al., 2022)	+0.17	+2.23	+1.29
Faster R-CNN-SSPNet (Ren et al., 2015)	+1.73	+5.74	+0.39
Average change in $mAP_{tiny}$	+2.54	+3.96	+2.04

Table 9: Detection performance of previous methods and our modified architecture on WiderFace. Pre-training is done on ImageNet. BN, GN &amp; DCN indicates the batch normalization, group normalization and deformable convolution networks respectively.

Method	Backbone	$mAP_{easy}$	$mAP_{medium}$	$mAP_{hard}$	FLOPs
TinaFace w BN(Zhu et al., 2020)	ResNet50	<b>95.77</b>	95.43	92.23	191.30
	BH-ResNet50	95.56	<b>95.71</b>	<b>92.28</b>	191.54
TinaFace w GN & DCN (Zhu et al., 2020)	ResNet50	96.27	95.67	93.06	184.29
	BH-ResNet50	<b>96.61</b>	<b>95.98</b>	<b>93.26</b>	188.60

ageNet32 outperform pre-training on ImageNet. It is likely that ImageNet32 is more relevant to tiny object detection but the blurring of images caused by down-sampling also makes it harder to learn useful representations from this dataset. The low top-1 classification accuracy, as shown in Table 5, provides evidence for this. CIFAR-100 also uses images that are consistent in size to the target objects in tiny object detection, but detection is much worse overall when using CIFAR-100 for pre-training, presumably because the small size of this dataset means that the representations that are learnt generalise more poorly.

When pre-training on ImageNet, although our method still outperforms the state-of-the-art methods, the improvement is not as significant as when pre-training with CIFAR-100 and ImageNet32. This difference is likely due to two reasons. One is that our backbones work better when they are pre-trained on a dataset that is more consistent with the tiny objects because our backbones are designed specifically for such objects. Another reason is we have simply used the same hyper-parameters to pre-train our BH networks on ImageNet as are used for the corresponding standard architectures, as we did not have the computational resources to search for more optimum hyper-parameters appropriate for pre-training BH networks on such a large dataset.

The results for WiderFace are shown in Table 9. We evaluate performance using the current best method, TinaFace (Zhu et al., 2020), with two different settings. Our backbones outperform the standard backbones in both cases in medium and hard tasks and

have achieved the state-of-the-art result in the easy task at 96.61%.

Overall, improved performance was generated using different backbone families (ResNet and HRNet), different network depths, pre-training with different datasets (CIFAR-100, ImageNet and down-sampled ImageNet), and when integrating the proposed backbones into several state-of-the-art tiny object detection frameworks. These results demonstrate that the proposed method of network modification generalises, and could potentially be used with other architectures to make them more suitable for small and tiny object classification and detection. Furthermore, our results demonstrate that the use of standard backbones is harmful for tiny object detection. The results also suggest that current architectures are potentially wasting computational resources by performing unnecessary computations in deep layers.

## 6 CONCLUSION

This paper proposes that standard deep neural network architectures used as feature-extraction front-ends in object-detection algorithms produce poor features for tiny object detection tasks. Our claim is that early down-sampling in such architectures results in information loss and features that are poor at representing small objects. To test this claim we designed bottom-heavy versions of popular backbone architectures, ResNet and HRNet, that increase the

number of convolutional layers in the shallow, high-resolution, stages of the network, and have fewer convolutional layers at later stages in the network. These changes in the distribution of the convolutional layers are made to ensure that the computational complexity of the modified backbones matches that of the original, standard, network. Experimental results show that these changes, despite reducing the number of parameters in the networks, result in more accurate object detection across a number of object-detectors, using different backbones, and pre-training schemes (CIFAR100, ImageNet and ImageNet32), for two standard benchmark datasets (TinyPerson and WiderFace). The architectures proposed in this paper are not a final answer to the information loss problem. Rather they serve as a motivation for developing improved feature-extraction backbones appropriate to the task. Hopefully, our current results and insights will inspire the development of even better backbones for tiny object detectors in future research.

## ACKNOWLEDGEMENTS

The authors acknowledge the use of the research computing facilities at King's College London, Rosalind, the King's Computational Research, Engineering and Technology Environment (CREATE) (London, 2022), and the Joint Academic Data science Endeavour (JADE) facility. This research was funded by the King's - China Scholarship Council (K-CSC).

## REFERENCES

- Bai, Y., Zhang, Y., Ding, M., and Ghanem, B. (2018). Finding tiny faces in the wild with generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Bhandari, A. and Tiwari, K. (2021). Loss of target information in full pixel and subpixel target detection in hyperspectral data with and without dimensionality reduction. *Evolving Systems*, 12.
- Bochkovskiy, A., Wang, C.-Y., and Liao, H.-Y. M. (2020). Yolov4: Optimal speed and accuracy of object detection.
- Cai, Z. and Vasconcelos, N. (2018). Cascade r-cnn: Delving into high quality object detection. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6154–6162.
- Chrabaszcz, P., Loshchilov, I., and Hutter, F. (2017). A downsampled variant of imagenet as an alternative to the cifar datasets. *ArXiv*, abs/1707.08819.
- Dalal, N. and Triggs, B. (2005). Histograms of oriented gradients for human detection. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 886–893 vol. 1.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- Deng, J., Guo, J., Ververas, E., Kotsia, I., and Zafeiriou, S. (2020). Retinaface: Single-shot multi-level face localisation in the wild. In *CVPR*.
- Everingham, M., Eslami, S. M. A., Van Gool, L., Williams, C. K. I., Winn, J., and Zisserman, A. (2015). The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136.
- Felzenszwalb, P., McAllester, D., and Ramanan, D. (2008). A discriminatively trained, multiscale, deformable part model. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8.
- Girshick, R. (2015). Fast R-CNN. In *International Conference on Computer Vision (ICCV)*.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 580–587.
- Gong, Y., Yu, X., Ding, Y., Peng, X., Zhao, J., and Han, Z. (2021). Effective fusion factor in fpn for tiny object detection. *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1159–1167.
- Han, X., Zhang, Z., Ding, N., Gu, Y., Liu, X., Huo, Y., Qiu, J., Yao, Y., Zhang, A., Zhang, L., Han, W., Huang, M., Jin, Q., Lan, Y., Liu, Y., Liu, Z., Lu, Z., Qiu, X., Song, R., Tang, J., Wen, J.-R., Yuan, J., Zhao, W. X., and Zhu, J. (2021). Pre-trained models: Past, present and future. *AI Open*, 2:225–250.
- He, K., Zhang, X., Ren, S., and Sun, J. (2014). Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition. *arXiv e-prints*, page arXiv:1406.4729.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition. *CoRR*, abs/1512.03385.
- Hong, M., Li, S., Yang, Y., Zhu, F., Zhao, Q., and Lu, L. (2022). Sspnet: Scale selection pyramid network for tiny person detection from uav images. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5.
- Hu, P. and Ramanan, D. (2017). Finding tiny faces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jiahuan, R., Zhao, Z., Jicong, F., Haijun, Z., Mingliang, X., and Meng, W. (2021). Robust low-rank deep feature recovery in CNNs: Toward low information loss and fast convergence. In *2021 IEEE International Conference on Data Mining (ICDM)*, pages 529–538.
- Jiang, N., Yu, X., Peng, X., Gong, Y., and Han, Z. (2021). Sm+: Refined scale match for tiny person detection. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1815–1819.
- Krishna, H. and Jawahar, C. (2017). Improving small object detection. In *2017 4th IAPR Asian Conference on Pattern Recognition (ACPR)*, pages 340–345.

- Krizhevsky, A. and Hinton, G. (2009). Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario.
- Lin, T., Goyal, P., Girshick, R. B., He, K., and Dollár, P. (2017a). Focal loss for dense object detection. *CoRR*, abs/1708.02002.
- Lin, T., Maire, M., Belongie, S. J., Bourdev, L. D., Girshick, R. B., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312.
- Lin, T.-Y., Dollár, P., Girshick, R. B., He, K., Hariharan, B., and Belongie, S. J. (2017b). Feature pyramid networks for object detection. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 936–944.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollar, P. (2017c). Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Liu, G., Han, J., and Rong, W. (2021a). Feedback-driven loss function for small object detection. *Image and Vision Computing*, 111:104197.
- Liu, J., Gu, Y., Han, S., Zhang, Z., Guo, J., and Cheng, X. (2021b). Feature rescaling and fusion for tiny object detection. *IEEE Access*, 9:62946–62955.
- Liu, Q., Tan, Z., Chen, D., Chu, Q., Dai, X., Chen, Y., Liu, M., Yuan, L., and Yu, N. (2022). Reduce information loss in transformers for pluralistic image inpainting. *ArXiv*, abs/2205.05076.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S. E., Fu, C., and Berg, A. C. (2015). SSD: single shot multi-box detector. *CoRR*, abs/1512.02325.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. (2021c). Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- London, K. C. (2022). King’s computational research, engineering and technology environment (CREATE). <https://doi.org/10.18742/rnvf-m076/>.
- Luo, S., Li, X., Zhu, R., and Zhang, X. (2019). Sfa: Small faces attention face detector. *IEEE Access*, 7:171609–171620.
- Pan, S. J. and Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.
- Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. (2015). You Only Look Once: Unified, Real-Time Object Detection. *arXiv e-prints*, page arXiv:1506.02640.
- Redmon, J., Divvala, S. K., Girshick, R. B., and Farhadi, A. (2015). You only look once: Unified, real-time object detection. *CoRR*, abs/1506.02640.
- Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *CoRR*, abs/1804.02767.
- Ren, S., He, K., Girshick, R. B., and Sun, J. (2015). Faster R-CNN: towards real-time object detection with region proposal networks. *CoRR*, abs/1506.01497.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489.
- Singh, B., Najibi, M., and Davis, L. S. (2018). SNIPER: Efficient multi-scale training. *NeurIPS*.
- Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. A. (2015). Striving for simplicity: The all convolutional net. *CoRR*, abs/1412.6806.
- Tan, M., Pang, R., and Le, Q. V. (2020). Efficientdet: Scalable and efficient object detection. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10778–10787.
- Tang, X., Du, D. K., He, Z., and Liu, J. (2018). Pyramidbox: A context-assisted single shot face detector. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Tian, Z., Shen, C., Chen, H., and He, T. (2019). Fcos: Fully convolutional one-stage object detection. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9626–9635.
- Tong, K. and Wu, Y. (2022). Deep learning-based detection from the perspective of small or tiny objects: A survey. *Image and Vision Computing*, 123:104471.
- ucas-vg (2020). TOV\_mmdetection.
- Wang, C.-Y., Bochkovskiy, A., and Liao, H.-Y. M. (2021a). Scaled-YOLOv4: Scaling cross stage partial network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13029–13038.
- Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., Liu, W., and Xiao, B. (2021b). Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43:3349–3364.
- Xu, C., Wang, J., Yang, W., Yu, H., Yu, L., and Xia, G.-S. (2022). Rfla: Gaussian receptive based label assignment for tiny object detection. In *European Conference on Computer Vision (ECCV)*.
- Yang, S., Luo, P., Loy, C. C., and Tang, X. (2016). Wider face: A face detection benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yang, Z., Chai, X., Wang, R., Guo, W., Wang, W., Pu, L., and Chen, X. (2019). Prior knowledge guided small object detection on high-resolution images. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 86–90.
- Yu, X., Gong, Y., Jiang, N., Ye, Q., and Han, Z. (2020). Scale match for tiny person detection. *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1246–1254.
- Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and Yoo, Y. (2019). Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.

- Zhang, S., Zhu, X., Lei, Z., Shi, H., Wang, X., and Li, S. Z. (2017). S3fd: Single shot scale-invariant face detector. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Zhang, Z., Shen, W., Qiao, S., Wang, Y., Wang, B., and Yuille, A. L. (2018). Robust face detection via learning small faces on hard images. *CoRR*, abs/1811.11662.
- Zhao, G., Wang, J., and Zhang, Z. (2017). Random shifting for CNN: a solution to reduce information loss in down-sampling layers. In Sierra, C., editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 3476–3482. ijcai.org.
- Zheng, Z., Wang, P., Liu, W., Li, J., Ye, R., and Ren, D. (2020). Distance-IoU loss: Faster and better learning for bounding box regression. In *The AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhu, Y., Cai, H., Zhang, S., Wang, C., and Xiong, Y. (2020). Tinaface: Strong but simple baseline for face detection. *arXiv preprint arXiv:2011.13183*.

