GIRAFFEDET: A HEAVY-NECK PARADIGM FOR OBJECT DETECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

In conventional object detection frameworks, a backbone body inherited from image recognition models extracts deep latent features and then a neck module fuses these latent features to capture information at different scales. As the resolution in object detection is much larger than in image recognition, the computational cost of the backbone often dominates the total inference cost. This heavy-backbone design paradigm is mostly due to the historical legacy when transferring image recognition models to object detection rather than an end-to-end optimized design for object detection. In this work, we show that such paradigm indeed leads to sub-optimal object detection models. To this end, we propose a novel heavy-neck paradigm, GiraffeDet, a giraffe-like network for efficient object detection. The GiraffeDet uses an extremely lightweight backbone and a very deep and large neck module which encourages dense information exchange among different spatial scales as well as different levels of latent semantics simultaneously. This design paradigm allows detectors to process the high-level semantic information and low-level spatial information at the same priority even in the early stage of the network, making it more effective in detection tasks. Numerical evaluations on multiple popular object detection benchmarks show that GiraffeDet consistently outperforms previous SOTA models across a wide spectrum of resource constraints.

1 INTRODUCTION

In the past few years, remarkable progress in deep learning based object detection methods has been witnessed. Despite object detection networks getting more powerful by different designing on architecture, training strategy and so on, the meta-goal that detecting all objects with **large-scale variation** has not been changed. For example, the scale of the smallest and largest 10% of object instances in COCO dataset is 0.024 and 0.472 respectively (Singh & Davis, 2018), which scaling in almost 20 times. This presents an extreme challenge to handle such a large-scale variation by using recent approaches. To this end, we aim to tackle this problem by designing a scale robust approach.

To alleviate the problem arising from large-scale variations, an intuitive way is to use multi-scale pyramid strategy for both training and testing. The work of (Singh & Davis, 2018) trains and tests detectors on the same scales of an image pyramid, and selectively back-propagates the gradients of object instances of different sizes as a function of the image scale. Although this approach improves the detection performance of most existing CNN-based methods, it is not very practical, as the image pyramid methods process every scale image, which could be *computationally expensive*. Moreover, the scale of objects between classification and detection datasets remains another challenge in *domain shift* when using pre-trained classification backbones.

Alternatively, the feature pyramid network is proposed to approximate image pyramids with lower computational costs. Recent methods still rely on superior backbone designing, but *insufficient information exchange* between high-level features and low-level features. For example, some work enhances the entire feature hierarchy with accurate localization signals in lower layers by bottom-up path augmentation, however this bottom-up path design might lack exchange between high-level semantic information and low-level spatial information. According to the above challenges, two questions in this task are raised as follows:

- *Is the backbone of the image classification task indispensable in a detection model?*
- *What types of multi-scale representations are effective for detection tasks?*

These two questions motivate us to design a new framework with two sub-tasks two *i.e.,* **efficient feature down-sampling** and **sufficient multi-scale fusion**. First, conventional backbones for scale-sensitive features generation are computationally expensive and exist domain-shift problem. An alternative lightweight backbone can solve these problems. Second, it is crucial for a detector to learn sufficient fused information between high-level semantic and low-level spatial features. According to the above motivations, we design a giraffe-like network, named as **GiraffeDet**, with the following insights: (1) An alternative **lightweight backbone** can extract multi-scale feature transformation without any additional computation costs. (2) A sufficient cross-scale connection, **Queen-Fusion**, like the Queen Piece pathway in chess, to deal with different levels and layers of feature fusion. (3) According to the designed lightweight backbone and flexible FPN, we propose a GiraffeDet family for each FLOPs level. Notably, the experimental results suggest that our GiraffeDet family achieves higher accuracy and better efficiency in each FLOPs level.

In summary, the key **contributions** of our work as follows:
• To the best of our knowledge, we present the first lightweight alternative backbone and flexible FPN combined as a detector. The proposed GiraffeDet family consists of Lightweight S2D-chain and Generalized-FPN, which demonstrates the state-of-the-art performance.
• We design the lightweight space-to-depth chain (S2D-chain) instead of the conventional CNN-based backbone, and controlled experiments demonstrate that the conventional CNN-based backbone is not necessary in object detection task.
• In our proposed Generalized-FPN (GFPN), a novel queen-fusion is proposed as our cross-scale connection style that fuses both level features in previous and current layers, and $\log_2 n$ skip-layer link slightly increase the short distances that can scale to deeper networks.

Based on the light backbone and heavy neck paradigm, the GiraffeDet family models perform well in a wide range of FLOPs-performance trade-offs. In particular, with the multi-scale testing technique, GiraffeDet-D29 achieves 54.1% mAP on the COCO dataset and outperforms other SOTA methods.

2 Related work

It is crucial for the object detector to recognize and localize objects by learning scale-sensitive features. Traditional solutions for the large-scale variation problem mainly based on improved convolutional neural networks. CNN-based object detectors are mainly categorized by two-stage detectors and one-stage detectors. Two-stage detectors (Ren et al., 2015; Dai et al., 2016; He et al., 2017; Cai & Vasconcelos, 2018; Pang et al., 2019) predict region proposals and then refine them by a sub-network, and one-stage detectors (Liu et al., 2016; Lin et al., 2017b; Redmon et al., 2016; Redmon & Farhadi, 2017; Tian et al., 2019; Zhu et al., 2019; Zhang et al., 2020; 2019; Ge et al., 2021) directly detecting bounding-boxes without the proposal generation step. In this work, we mainly conduct experiments based on one-stage detector methods.

Recently, the main research line is utilizing pyramid strategy, including image pyramid and feature pyramid. The image pyramid strategy is used for detecting instances by scaling images. For example, SNIPER (Singh et al., 2018) propose a fast multi-scale training method, which samples the foreground regions around ground-truth object and background regions for different scale training. Unlike image pyramid methods, feature pyramid methods fuse pyramidal representations that cross different scales and different semantic information layers. For instance, PANet (Liu et al., 2018) enhances the feature hierarchies on top of feature pyramid network by additional bottom-up path augmentation. In addition, NAS-FPN (Ghiasi et al., 2019) utilize neural architecture search to automatically search feature pyramid network topology Our work focuses on feature pyramid strategy and proposes a sufficient high-level semantic and low-level spatial information fusion method.

Some researchers start working on designing new architectures to solve the large-scale variation problem instead of "backbone-neck-head" architecture in detection tasks. The work of Sun et al. (2019b) proposed the FishNet as an encoder-decoder architecture with skip connections to fuse multi-scale features. SpineNet (Du et al., 2020) designed as a backbone with scale-permuted intermediate features and cross-scale connections that is learned on an object detection task by Neural Architecture Search. Our work is inspired by these methods and proposes a lightweight space-to-depth backbone instead of a CNN-based backbone. However, our GiraffeDet still designed as the "backbone-neck-head" architecture. Because this typical architecture is widely used and proved effective in detection tasks.

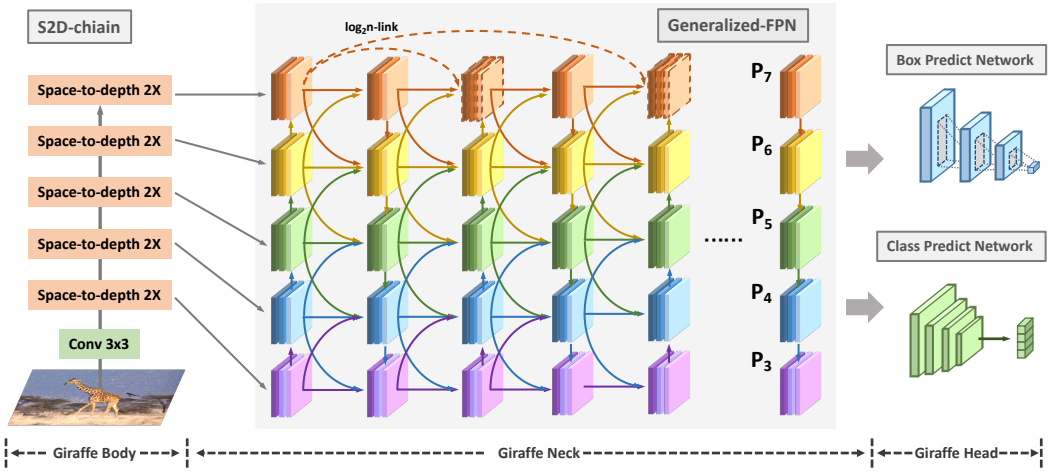

Figure 1: Overview of the GiraffeDet which has three parts: 1) *Body* contains lightweight S2D-chain and image preprocessing 2) *Heavy neck* refines and fuses high-level semantic and low-level spatial features 3) *Head* predicts the bounding box and class label of exist objects.

3 THE GIRAFFEDET

Although extensive research has been carried out to investigate efficient object detection, large-scale variation still remains a challenge. To achieve the goal of sufficient multi-scale information exchange efficiently, we proposed the GiraffeDet for efficient object detection, and the "giraffe" consists of lightweight space-to-depth chain, generalized-FPN and prediction networks. The overall framework is shown in Figure 1, and details of architecture are introduced below.

3.1 LIGHTWEIGHT SPACE-TO-DEPTH CHAIN

Most feature pyramid networks apply conventional CNN-based networks as the backbone to extract multi-scale feature maps and even learn information exchange. However, recent backbones became much heavier with the development of CNN, it is computationally expensive to utilize them. Moreover, most applied backbones are mainly pre-trained on classification dataset, *e.g.,* ResNet50 pre-trained on ImageNet, and we argue these pre-trained backbones are inappropriate in detection task and remains the domain-shift issue. Alternatively, FPN more emphasis on high-level semantic and low-level spatial information exchange. Therefore, we assume that FPN is more critical and that the conventional backbone is not necessary in the object detection model.

Conventional down-sampling method utilizes interpolation or max-pooling to down-sample feature maps. Also, CNN-based backbone down-samples feature maps with non-linear transformation and information. Inspired by (Shi et al., 2016; Sajjadi et al., 2018), we propose a Space-to-Depth chain (S2D-chain) as our lightweight backbone by sampling and re-organizing features with a fixed gap uniformly. The S2D-chain moves spatial-dimension information to depth-dimension so as to down-sample features, as shown in Figure 2. Compared with conventional down-sampling method and CNN-based method, S2D-chain is computationally cheap without any additional parameters.

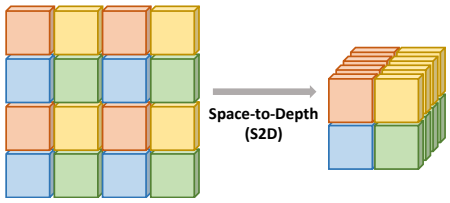

Figure 2: Illustration of the space-to-depth transformation. The S2D operation moves the activation from the spatial dimension to the channel dimension

To verify our assumption, we conduct controlled experiments on different backbone and FPN neck ratios in multiple object detection of the same FLOPs in Section 4. The results show that increasing backbone computations is not effective for object detection tasks.

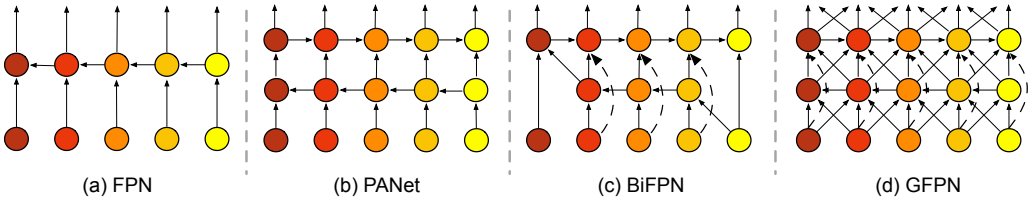

Figure 3: Feature pyramid network evolution design from level 3 to level 7 (P3 - P7). (a) FPN (Lin et al., 2017a) introduces a top-down pathway to fuse multi-scale features; (b) PANet (Liu et al., 2018) adds an additional bottom-up pathway on top of FPN; (c) BiFPN (Tan et al., 2020) introduces a bidirectional cross-scale pathway; (d) our GFPN contains both queen-fusion style pathway and skip-layer connection.

3.2 GENERALIZED-FPN

In feature pyramid network, multi-scale feature fusion aims to aggregate different resolution features that are extracted from the backbone. Figure 3 shows the evolution of feature pyramid network design. Conventional FPN (Lin et al., 2017a) introduces a top-down pathway to fuse multi-scale features from level 3 to 7. Considering the limitation of one-way information flow, PANet(Liu et al., 2018) adds an extra bottom-up path aggregation network, but with more computational cost. Besides, BiFPN (Tan et al., 2020) removes nodes that only have one input edge, and add extra edge from the original input on the same level. However, we observe that previous methods focus only on feature fusion, but lack the inner block connection. Therefore, we design a novel pathway fusion including skip-layer and cross-scale connections, as shown in Figure 3(d).

Skip-layer Connection. Compared to other connection methods, skip connections have short distances among feature layers during back-propagation. In order to reduce gradient vanish in such a heavy "giraffe" neck, we propose two feature link methods: dense-link and $\log_2 n$-link in our proposed GFPN, as shown in Figure 4.

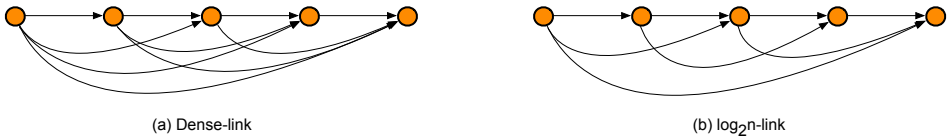

Figure 4: Two link mode of skip-layer connection: (a) dense-link: the concatenation of all preceding layers (b) $\log_2$n-link: the concatenation of at most $log_2 l + 1$ layers.

• ***dense-link***: Inspired by DenseNet (Huang et al., 2017), for each scale feature P_k^l in level k, Consequently, the lth layer receives the feature-maps of all preceding layers:

$$P_k^l = Conv(Concat(P_k^0, \ldots, P_k^{l-1})), \tag{1}$$

where $Concat()$ refers to the concatenation of the feature-maps produced in all preceding layers, and $Conv()$ represents a 3x3 convolution.

• ***log$_2$n-link***: Specifically, in each level k, the lth layer receives the feature-maps from at most $log_2 l + 1$ number of preceding layers, and these input layers are exponentially apart from depth i with base 2, as denoted:

$$P_k^l = Conv(Concat(P_k^{l-2^n}, \ldots, P_k^{l-2^1}, P_k^{l-2^0})), \tag{2}$$

where $l - 2^n \geq 0$, $Concat()$ and $Conv()$ also represent concatenation and 3x3 convolution respectively. Compare to dense-link at depth l, the time complexity of $\log_2 n$-link only cost $O(l \cdot log_2 l)$, instead of $O(l^2)$. Moreover, $log_2 n$-link only increase the short distances among layers during back-propagation from 1 to 1+$\log_2 l$. Hence, $\log_2 n$-link can scale to deeper networks.

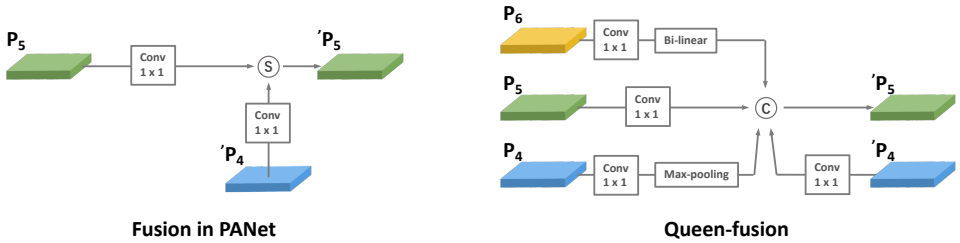

Figure 5: Illustration of cross-scale connection between PANet and our Queen-fusion in GFPN. S and C represent summation and concatenation fusion style, and $'P_k$ denotes node in next layer.

Cross-scale Connection. Based on our assumption, our designed sufficient information exchange should contains not only skip-layer connection, but also cross-scale connection, to overcome large-scale variation. Previous works in connecting features between adjacent layers only consider same level feature (Liu et al., 2018) or previous level feature (Tan et al., 2020). Therefore, we propose a new cross-scale fusion named as **Queen-fusion**, that considering both same level and neighbor level features as shown in Figure 3(d), like playing the queen piece in chess. As an example shown in Figure 5(b), the concatenation of Queen-fusion in P_5 consists previous layer P_4 down-sampling, previous layer P_6 up-sampling, previous layer P_5 and current layer P_4. In this work, we apply *bilinear interpolation* and *max-pooling* as our up-sampling and down-sampling functions respectively.

Therefore, in the extreme large-scale variation scenario, it requires that the model has sufficient high-level and low-level information exchange. Based on the mechanism of our skip-layer and cross-scale connections, the proposed generalized-FPN can be expanded as long as possible, just like the "giraffe neck". With such a "heavy neck" and a lightweight backbone, our GiraffeDet can balance higher accuracy and better efficiency trade-off.

3.3 GIRAFFEDET FAMILY

According to our proposed S2D-chain and Generlized-FPN, we can develop a family of different GiraffeDet scaling models that can overcome a wide range of resource constraints. Previous work scale up its detector in the inefficient way, as changing bigger backbone networks like ResNeXt (Xie et al., 2017), or stacking FPN blocks *e.g.,* NAS-FPN (Ghiasi et al., 2019). Specially, EfficientDet (Tan et al., 2020) start using compound coefficient ϕ to jointly scale up all dimensions of backbone. Inspired by EfficientDet, we also apply compound coefficient ϕ for scaling, but we only focus on the scaling of GFPN layers instead of the whole framework including lightweight backbone.

Based on our GFPN and eS2D chain, we have developed a GiraffeDet family. Previous work mostly scale up a baseline detector by changing bigger backbone networks (*e.g.,* ResNeXtXie et al. (2017) or AmoebaNetReal et al. (2019), training with larger input images, or stacking more FPN layers Ghiasi et al. (2019). These methods are usually ineffective since they only focus on a single or limited scaling dimensions. As we assume backbone is not critical for object detection task, the GiffeDet family only focus on the scaling of generalized-FPN. Two multipliers are proposed to control the depth (# of layers) and width (# of channels) for GFPN:

$$D_{gfpn} = \phi, \;\; W_{gfpn} = 64 * (log_2 \phi), \tag{3}$$

Following above setting and equation. 3, we have developed six architectures of GiraffeDet, as shown in Table 1. GiraffeDet-D7,11,14,16 have the same level FLOPs with ResNet-series based model and we compare performance of GiraffeDet family with SOTA models in the next section.

Table 1: Scaling configs for GiraffeDet family $-$ ϕ is the hyper$-$parameter that denotes the depth (# of layers) of GFPN. The width (# of channels) of GFPN can be logarithmically calculated by by equation. 3.

	Generalized-FPN	
	D_{gfpn}	W_{gfpn}
D7 (ϕ=7)	7	179
D11 (ϕ=11)	11	221
D14 (ϕ=14)	14	243
D16 (ϕ=16)	16	256
D25 (ϕ=25)	25	297
D29 (ϕ=29)	29	310

4 EXPERIMENTS

In this section, we first introduce the implementation details and present our experimental result on the COCO dataset (Lin et al., 2014). Then compare our proposed GiraffeDet family with other state-of-the-art methods, and an in-depth analysis is provided to better understand our framework. Previous work mostly scales up a baseline detector by changing bigger backbone networks

4.1 IMPLEMENTATION DETAILS

COCO dataset. We evaluate GiraffeDet on COCO 2017 detection dataset with 80 object categories. It includes 115k images for training ($train$), 5k images for validation (val) and 41k images with no public ground-truth for testing ($test - dev$). The training of all methods is conducted on the 115k training images. We report results on the validation dataset for ablation study and results of the test-dev dataset from the evaluation server for state-of-the-art comparison.

For fair comparison, all results are produced under mmdetection (Chen et al., 2019) and the standard COCO-style evaluation protocol is adopted. Following the the work of (He et al., 2019), all models are trained from scratch to reduce the influence of pre-train backbones on ImageNet. The shorter side of input images is resized to 800 and the maximum size is restricted within 1333. SGD optimizer with momentum 0.9 and weight decay 1e-4 is utilized, and initial learning rate is set to 0.02 with 0.1 step decreasing ratio. To enhance the stability of scratch training, we adopt multi-scale training for all models, including: 3x scratch learning schedule (36 epochs, decays at 28 and 33 epochs) is applied in ablation study with a multi-scale range between [0.8, 1.0] on the short side, and 6x scratch learning schedule (72 epochs, decays at 65 and 71 epochs) is applied in state-of-the-art comparison. Each model is trained on 8 Tesla V100 GPUs with the batch size of 16 (2 per GPU).

4.2 QUANTITATIVE EVALUATION ON COCO DATASET

We compare GiraffeDet with state-of-the-art approaches under the single-model single-scale setting with no test-time augmentation in Table. 2. We report accuracy for both test-dev (41k images with no public ground-truth) and val with 5k validation images. FLOPs denote the number of multiply-adds. We group models together if they have similar FLOPs and compare their accuracy in each group. Notably, model performance depends on both network architecture and training settings. We refer most models from their paper. But for a fair comparison, we also reproduce some of RetinaNet (Lin et al., 2017b), FCOS (Tian et al., 2019), HRNet (Sun et al., 2019a), with 6x training from scratch, which denoted as †.

Large-scale Variance. According to the performance of Figure 6, we can observe that our proposed GiraffeDet achieves the best performance in each pixel scale range, which indicates that the light backbone and heavy-neck paradigm, as well as our proposed GFPN, can effectively solve a large-scale variance problem. Also, under the skip-layer and cross-scale connections, high-level semantic information and low-level spatial information can be sufficiently exchanged. Many object instances are smaller than 1% of the image area in the COCO dataset, making detectors difficult to detect. Even though extremely small instance are difficult to detect, our method still outperforms 5.7% mAP than RetinaNet in the pixel range 0-32, which outperforms the same mAP in the middle pixel range 80-144. Notably, in the scale of pixel range 192-256, the proposed GiraffeDet outperforms the most than other methods, which proves that our design can learn scale-sensitive features effectively.

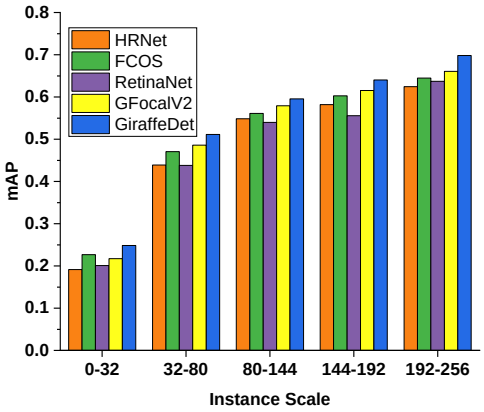

Figure 6: mAP on all scale of object instances (pixels) in five different models under R50 FLOPs level and 6x scratch training, including HRNet (Sun et al., 2019a), GFocalV2 (Li et al., 2020), RetinaNet (Lin et al., 2017b), FCOS (Tian et al., 2019) and our proposed GiraffeDet.

Table 2: **GiraffeDet performance on COCO** - Results are for single-model single-scale. test-dev is the COCO test set and val is the validation set. †means that the results are reproduced by 6x scratch training, others are referred from their paper. We group models together if they have similar FLOPs, and compare their accuracy in each group. $\mathbf{MS}_{test}$: multiscale testing, **R:** ResNet, **X:** ResNext, and **W:** low-level features map width in HRNet (# of channels).

Model	#FLOPs Backbone	#FLOPs Neck	#FLOPs total	test-dev $\mathbf{AP_{test}}$	$\mathbf{AP}_S$	$\mathbf{AP}_M$	$\mathbf{AP}_L$	val2017 $\mathbf{AP_{val}}$
Yolov3-darknet53	139.82	32.54	187.73	33.0	18.3	35.4	41.9	33.8
RetinaNet-R50†	76.15	16.6	229.59	40.4	23.1	43.3	52.2	40.2
FCOS-R50†	76.15	16.6	192.39	42.9	26.6	46.5	53.8	42.7
GFLV2-R50	76.15	16.6	199.96	44.3	26.8	47.7	54.1	43.9
HRNetV2p-W18	65.13	16.34	182.06	38.3	23.1	41.5	49.2	38.1
HRNetV2p-W18†	65.13	16.34	182.06	40.6	25.7	43.3	50.6	40.2
GiraffeDet-D7	3.89	76.13	183.67	45.6	28.8	48.7	55.6	44.9
RetinaNet-R101	153.74	16.6	302.62	39.1	21.8	42.7	50.2	38.9
FCOS-R101	153.74	16.6	265.38	41.5	24.4	44.8	51.6	40.8
ATSS-R101	153.74	16.6	269.94	43.6	26.1	47.0	53.6	41.5
PAA-R101	153.74	16.6	269.94	44.8	26.5	48.8	56.3	43.5
GFLV2-R101	153.74	16.6	272.99	46.2	27.8	49.9	57.0	45.9
HRNetV2p-W32	155.8	19.64	276.03	40.5	23.4	42.6	51.0	40.3
HRNetV2p-W32†	155.8	19.64	276.03	44.6	27.9	48.0	56.8	44.1
GiraffeDet-D11	3.89	166.73	274.26	46.9	29.9	51.1	58.4	46.6
RetinaNet-R152	226.84	16.6	375.72	45.1	28.4	48.8	58.2	-
HRNetV2p-W40	229.57	21.52	351.69	42.8	27.0	46.4	54.5	42.7
GiraffeDet-D14	3.89	251.9	359.43	47.7	30.9	51.6	60.3	47.3
FSAF-X101-64x4d	304.68	16.6	421.86	42.9	26.6	46.2	52.7	42.4
libraRCNN-X101-64x4d			424.32	43.0	25.3	45.6	54.6	42.7
FreeAnchor-X101-64x4d	304.68	16.6	458.07	44.9	26.5	48.0	56.5	-
FCOS-X101-64x4d	304.68	16.6	420.87	43.2	26.5	46.2	53.3	42.6
ATSS-X101-64x4d	304.68	16.6	425.43	45.6	28.5	48.9	55.6	-
OTA-X101-64x4d	304.68	16.6	453.55	47.0	29.2	50.4	57.9	-
GiraffeDet-D16	3.89	315.69	423.22	48.7	31.7	52.4	61.3	48.3
GiraffeDet-D25	3.89	681.02	788.55	50.5	32.2	54.2	63.5	49.9
GiraffeDet-D29	3.89	865.44	972.97	51.3	33.1	54.9	64.9	51.0
GiraffeDet-D29+MS$_{test}$	3.89	865.44	972.97	54.1	35.9	56.8	67.2	53.9

Comparison with State-of-the-art Methods. Table 2 shows that our GiraffeDet family achieves better performance than previous detectors in each same level of FLOPs, which indicates that our method can detect objects effectively and efficiently. 1) Compared to ResNet-based methods on the low-level FLOPs scale, we found that, even if the overall performance is obviously not increased too much, our method has a significant performance in detecting small and large object cases. It indicates our method performs better in large-scale variation dataset. 2) Compared to ResNext-based methods in high-level FLOPs scale, we find that GiraffeDet has a higher performance than in low-level FLOPs slot, which indicates that a good design of FPN can be more crucial than a heavy backbone. 3) Compared to other methods, the proposed GiraffeDet family also has the SOTA performance that proves our design achieves higher accuracy and better efficiency in each FLOPs level. Besides, the NAS-based method consumes a ton of computational resources to cover the search space in the training process, and therefore we do not consider comparing our method with them. Finally, with the multi-scale test protocol, our GiraffeDet achieve 54.1% mAP, especially AP_S increases 2.8% and AP_L increases 2.3% much more than 1.9% in AP_M.

4.3 Ablation Study

The success of our GiraffeDet can be attributed to both framework design and technical improvements in each component. To analyze the effect of each component in GiraffeDet, we construct ablation studies including: 1) Connection analysis in generalized-FPN; 2) Fusion-style analysis in generalized-FPN; 3) Backbone discussion.

Table 3: Ablation study on the Connection analysis. The model of "GFPN w/o skip" neck without any skip-layer connection, "GFPN-dense" neck model utilizes dense-link and "GFPN-$\log_2$n" neck model utilizes dense-$\log_2$n. "s-3x" training represents 3x scratch learning schedule.

Backbone	Neck	training	FLOPs(G)	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
S2D chain	stacked FPN	s-3x	276.11	38.8	54.3	42.2	23.1	41.9	50.6
S2D chain	PANet	s-3x	275.32	40.5	56.3	43.9	24.3	43.8	52.1
S2D chain	stacked BiFPN	s-3x	273.1	41.0	57.1	44.3	24.0	43.6	51.9
S2D chain	GFPN w/o skip	s-3x	273.51	41.2	57.0	44.3	25.7	43.5	51.9
S2D chain	GFPN-dense	s-3x	273.43	41.3	57.1	44.4	26.0	43.8	52.2
S2D chain	**GFPN-$\log_2$n**	s-3x	274.26	41.8	58.1	45.7	26.4	44.9	52.7

Connection Analysis There are multiple options for constructing pathways between nodes, which are mainly based on graph theory design and human empirical design. Different connections represent different exchanges of information on feature maps. We construct ablation study models and conduct experiments to investigate the effects of our proposed connections. In addition, we stacked basic FPN and BiFPN several times for fair comparison on the same FLOPs level and used the same backbone and prediction head. Results are given in Table 3.

• *Skip-layer Connection*. According to the results of GFPN-dense and GFPN-$\log_2$n neck of GiraffeDet, we observe that $\log_2$n connection has achieved the best performance, and dense connection only performs slightly better than without any skip-layer connection. It indicates that the $\log_2$n connection preserves the information exchange from early nodes to later, and dense connection provides the redundant information exchange and random graph connection, which is unpredictable. Notably, both generalized-FPN connections obtain higher performance than stacked BiFPN on the same level of FLOPs, which proves that our proposed GiraffeDet can be more efficient.

• *Cross-scale Connection*. From the table, we can see that our GiraffeDet model can achieve better performance, which proves that our Queen-fusion can effectively obtain information from previous nodes. Especially, even without skip-layer connection, our generalized-FPN can outperform other methods, which indicates the effectiveness of Queen-fusion. In addition, stacked PANet and stacked BiFPN can achieve higher accuracy than their basic structure with bidirectional information flow.

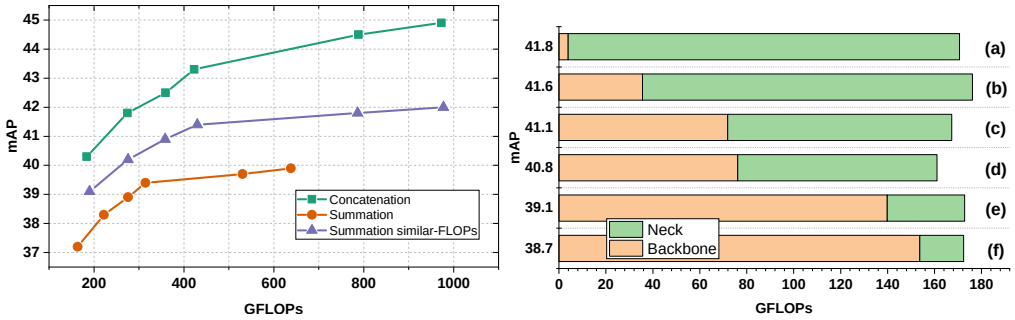

Figure 7: Ablation study on the fusion-style analysis and different backbone GiraffeDet. Fusion-style analysis consists three models: 1) "Concatenation" model: GiraffeDet utilizes concatenation fusion style; 2) "Summation" model: GiraffeDet utilizes summation fusion style; 3) "Summation smilar-FLOPs" model: same FLOPs level with "Concatenation" model. Six Different backbone GiraffeDet-like models conducted on same FLOPs level, they consists: a) S2D-chain and GFPN-D11; b) ResNet-18 and GFPN-D10; c) ResNet-34 and GFPN-D8; d) ResNet-50 and GFPN-D7; e) DarkNet and GFPN-D4; f) ResNet-101 and GFPN-D2.

Fusion-style Analysis. Figure 7 left shows the performance of using summation-based feature fusion and concatenation-based feature fusion style. We can observe that the concatenation-based fusion style of features can achieve better performance in the same FLOPs level. Although summation-based feature fusion has fewer FLOPs than concatenation-based style, performance is significantly lower. We think it is not worth sacrificing mAP to have fewer FLOPs. Notably, the performance of the "Summation" model is growing slightly after GFLOPs over 300, which indicates the concatenation-based feature fusion style can be more accurate and efficient again.

Backbone Discussion. Figure 7 right shows the performance of different GFPN depth and differnet backbone in the same FLOPs level. The results show that the combination of S2D-chain and GFPN outperforms other backbone models, which verify our assumption that FPN is more crucial and that conventional backbone is not necessary in object detection. In particular, we can observe that performance even decreases with the growth of the backbone model. We consider this might because the *domain-shift* problem remains higher in a large backbone, and it also proves our assumption.

4.4 Qualitative Evaluation

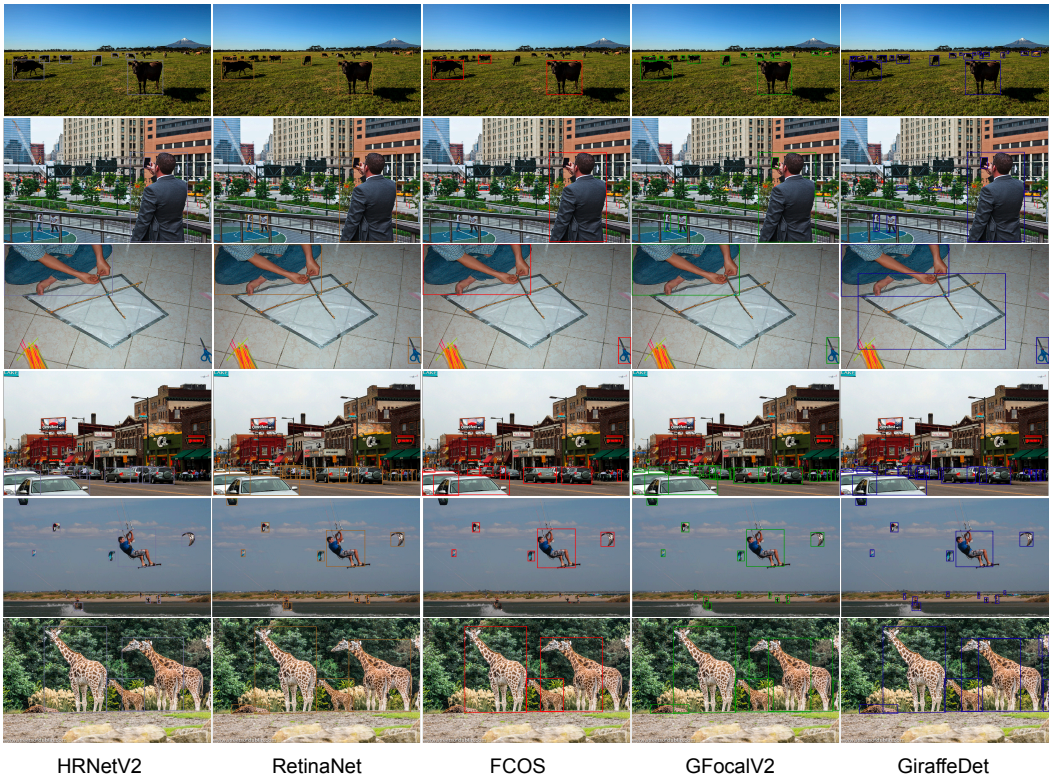

Figure 8: Qualitative Evaluation of different approaches for object detection on COCO dataset.

To better illustrate the performance of different approaches, we provide qualitative result in Figure 8. Overall, we can observe that all methods can detect object instances from each image. Furthermore, GiraffeDet can detect more instances than other SOTA methods, especially small object instances, which proves that our designed FPN can be effective in the large-scale variation dataset.

5 Conclusion

In this paper, we propose a novel heavy-neck paradigm framework, GiraffeDet, a giraffe-like network, to address the problem of large-scale variation. In particular, GiraffeDet uses a lightweight spatial-to-depth chain as a backbone, and the proposed generalized-FPN as a heavy neck. The spatial-to-depth chain is applied to extract multi-scale image features in a lightweight way, and the generalized-FPN is proposed to learn sufficient high-level semantic information and low-level spatial information exchange. Extensive results manifested that the proposed GiraffeDet family achieves higher accuracy and better efficiency, especially detecting small and large object instances.

REFERENCES

Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6154–6162, 2018.

Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.

Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *Advances in neural information processing systems*, pp. 379–387, 2016.

Xianzhi Du, Tsung-Yi Lin, Pengchong Jin, Golnaz Ghiasi, Mingxing Tan, Yin Cui, Quoc V Le, and Xiaodan Song. Spinenet: Learning scale-permuted backbone for recognition and localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11592–11601, 2020.

Zheng Ge, Songtao Liu, Zeming Li, Osamu Yoshie, and Jian Sun. Ota: Optimal transport assignment for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 303–312, 2021.

Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Nas-fpn: Learning scalable feature pyramid architecture for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7036–7045, 2019.

Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, 2017.

Kaiming He, Ross Girshick, and Piotr Dollár. Rethinking imagenet pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4918–4927, 2019.

Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.

Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *arXiv preprint arXiv:2006.04388*, 2020.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.

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

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

Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8759–8768, 2018.

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

Jiangmiao Pang, Kai Chen, Jianping Shi, Huajun Feng, Wanli Ouyang, and Dahua Lin. Libra r-cnn: Towards balanced learning for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 821–830, 2019.

Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pp. 4780–4789, 2019.

Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7263–7271, 2017.

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

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28: 91–99, 2015.

Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6626–6634, 2018.

Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1874–1883, 2016.

Bharat Singh and Larry S Davis. An analysis of scale invariance in object detection snip. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3578–3587, 2018.

Bharat Singh, Mahyar Najibi, and Larry S Davis. Sniper: Efficient multi-scale training. *arXiv preprint arXiv:1805.09300*, 2018.

Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5693–5703, 2019a.

Shuyang Sun, Jiangmiao Pang, Jianping Shi, Shuai Yi, and Wanli Ouyang. Fishnet: A versatile backbone for image, region, and pixel level prediction. *arXiv preprint arXiv:1901.03495*, 2019b.

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

Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9627–9636, 2019.

Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.

Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9759–9768, 2020.

Xiaosong Zhang, Fang Wan, Chang Liu, Rongrong Ji, and Qixiang Ye. Freeanchor: Learning to match anchors for visual object detection. *arXiv preprint arXiv:1909.02466*, 2019.

Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 840–849, 2019.