

Improving Object Detection from Scratch via Gated Feature Reuse

Zhiqiang Shen¹

Honghui Shi^{2,3}

Jiahui Yu²

Hai Phan¹

Rogério Feris^{3,4}

Liangliang Cao⁵

Ding Liu²

Xinchao Wang⁶

Thomas Huang²

Marios Savvides¹

¹ Carnegie Mellon University

² University of Illinois at
Urbana-Champaign

³ IBM Research AI

⁴ MIT-IBM Watson AI Lab

⁵ University of Massachusetts Amherst

⁶ Stevens Institute of Technology

{zhiqians,haithanp,marioss}@andrew.cmu.edu;{hshi10,jyu79,dingliu2,t-huang1}@illinois.edu

Abstract

In this paper, we present a simple and parameter-efficient drop-in module for one-stage object detectors like SSD [15] when learning from scratch (i.e., without pre-trained models). We call our module GFR (Gated Feature Reuse), which exhibits two main advantages. First, we introduce a novel gate-controlled prediction strategy enabled by Squeeze-and-Excitation [14] to adaptively enhance or attenuate supervision at different scales based on the input object size. As a result, our model is more effective in detecting diverse sizes of objects. Second, we propose a feature-pyramids structure to squeeze rich spatial and semantic features into a single prediction layer, which strengthens feature representation and reduces the number of parameters to learn. We apply the proposed structure on DSOD [6, 7] and SSD [15] detection frameworks, and evaluate the performance on PASCAL VOC 2007, 2012, 2012 Comp3 and COCO datasets. With fewer model parameters, GFR-DSOD outperforms the baseline DSOD by 1.4%, 1.1%, 1.7% and 0.6%, respectively. GFR-SSD also outperforms the original SSD and SSD with dense prediction by 3.6% and 2.8% on VOC 2007 dataset. Code is available at: <https://github.com/szq0214/GFR-DSOD>.

1 Introduction

In the recent years, with the wide usage of many electronic devices such as mobile phones, embedded devices, etc, learning compact object detectors with high accuracy has been a urgent-needed breakthrough problem in computer vision and a promising research direction in deep learning. However, most existing state-of-the-art object detection systems are designed based on backbone networks that are pre-trained on classification datasets (e.g., ImageNet [19]), which are fairly heavy on #params and inflexible to adjust appropriate structures for different scenarios and circumstances. In this paper, we tackle this problem by considering the following two aspects: (i) how to improve utilization of features on different scales;

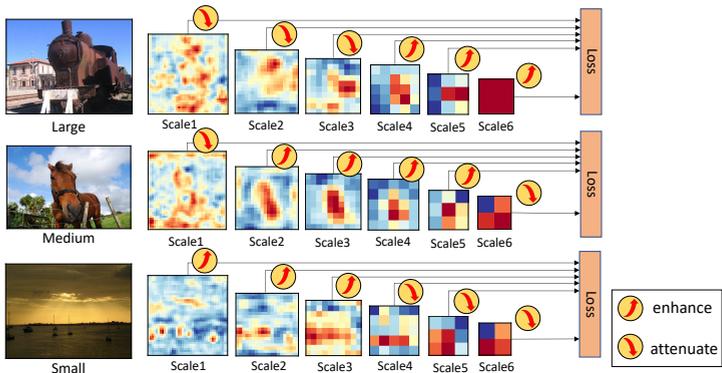


Figure 1: Illustration of the motivation. For a wide range of input object sizes, we adopt gates (the arrows upon the feature maps) on each scale to adaptively recalibrate the supervision intensities based on the input object sizes. Specifically, for large or small objects, our gates will automatically enhance or attenuate feature activations at appropriate scales.

and (ii) how to detect diverse sizes of objects with gate-controlled prediction strategy, in order to build a parameter-efficient object detector. Our design in this paper can further bring incidental advantage that our model is easier to train from scratch as we use fewer parameters and more simple structures in our models.

Some recent works have been explored to learn lite and parameter-efficient network architectures for image classification task, such as MoblileNet [13], MoblileNet V2 [14], ShuffleNet [15], etc. In our work, we adopt the existing backbone networks and focus on designing the detection head, our gated structure can be used in conjunction with any other backbone structure techniques. There are also some recent studies focusing on developing feature-pyramids structures for detectors which have achieved very promising results. One example is the FPN [16] model. Similar to many other detection frameworks such as SSD [17], DSSD [18], DSOD [19, 20], Reconfig [21], MS-CNN [22], Hypernet [23] and ION [24], FPN relies on *feature pyramids* for multiple scale prediction. For different pyramidal layers, a series of prediction operations are conducted to adapt the arbitrary object scales.

The current *feature pyramids* design, however, has two major limitations. First is that each pyramid has a fixed and thus non-adaptive contribution to the final supervision signals. Intuitively, objects at small scales may be easier to detect with fine-resolution (lower-level) features, and thus signals from those lower-level features should be enhanced; Similarly, large scale objects could be easier to detect at coarse-resolution (higher-level) feature maps, and thus signals from those higher-level feature maps should be enhanced. Nevertheless, these are unfortunately not taken into account by the state-of-the-art detector pyramids. The second limitation is the naive single-scaled feature representation in each pyramidal layer, as done in SSD [17] and FPN [16], where the pyramidal layers are independent without any interactions.

In this paper, we introduce a novel gated feature reuse (GFR) mechanism that explicitly tackles the above two challenges. To address the first challenge, we design a gate structure enabled by Squeeze-and-Excitation [25] that dynamically adjusts supervision at different scales. As illustrated in Fig. 1, for large or small objects, our gates automatically intensify or diminish the feature activations at proper scales. For the second challenge, we propose *feature reuse*, a network that concatenates high-level semantic features and low-level spatial

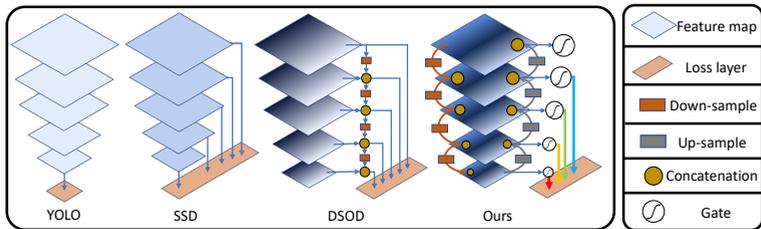


Figure 2: An overview of our proposed GFR-DSOD together with three one-stage detector methods (YOLO [27], SSD [25] and DSOD [61, 62]). YOLO is a single-scale detector, SSD and DSOD are two multi-scale detectors that localize objects based on various resolution features. DSOD further adopts dense connections on prediction layers to combine different resolution features into one scale. Our GFR-DSOD consists of two modules: *Iterative Feature Pyramids* and *Gating Mechanism*. More details are given in §3.

features in a single pyramid, so that the features may be complementary with each other and jointly lead to better results. Our ablation experiments show that this simple design vastly boost the performance of object detection and reduce the model parameters. We incorporate our proposed GFR mechanism into DSOD and SSD, resulting in higher detection accuracy, fewer parameters, and faster convergence, as shown in Fig. 6. Such incorporation is done without bells and whistles, and could be readily extended to other detection frameworks.

Our main contributions are summarized as follows:

- (1) We propose the *Feature Reuse*, a novel structure for learning detectors from scratch, that allows the features of different scales to interact, further leading to less model parameters and faster convergence speed.
- (2) We propose *Gating Mechanism* for generic object detection, which is, to our best knowledge, the first successful attempt on recalibrating supervision signals for detection.
- (3) We apply *GFR* structure on two detection frameworks DSOD and SSD, the resulting GFR-DSOD method achieves state-of-the-art (under similar #params) for learning detection from scratch with real-time processing speed and more compact models.

2 Related Work

Object Detection. Generally, modern object detection frameworks fall into two groups. One is the two-stage detectors like R-CNN [9], Fast RCNN [8], R-FCN [9], Faster RCNN [28], Deformable CNN [6], Mask RCNN [12], etc. In this paper we focus on another group: one-stage detectors (also known as proposal-free detectors). OverFeat [60] is regarded as one of the first CNN-based one-stage detectors. After that, a number of recent detectors have been proposed, such as YOLO [26, 27], SSD [25], RetinaNet [23], DSOD [61, 62], CornerNet [20], ExtremeNet [62], Refinement [69], UnitBox [67], STDN [41], RFB [24], etc. The key advantages of one-stage detectors is the straight-forward structures and high speed but their accuracy is moderate. Thus, the aim of this work is to further boost the performance of one-stage detectors under the setting of training from scratch.

Feature Pyramids. Adopting multiple layers for generic object detection is a common practice in a range of recently proposed approaches [7, 22, 25, 62]. For instance, DSOD [61, 62] applies dense connections in prediction layers to combine different resolution features for detecting objects. FPN [22] develops a top-down architecture with lateral connections to

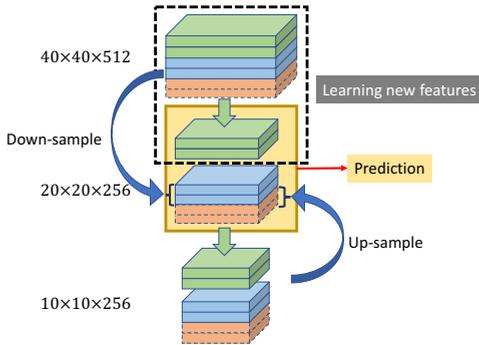


Figure 3: A building block illustrating the iterative feature pyramids, including down-sampling pathway, up-sampling pathway and the concatenation operation. Green color maps are the learned new features, blue and red ones are the re-used old features from previous and following layers, respectively.

build pyramidal hierarchy semantic features at all scales. DSSD [4] involves extra deconvolutional layers to capture additional large-scale context. In this paper, we propose a novel iterative feature pyramids structure to not only improve accuracy but also reduce parameters.

Gating Mechanism. Gating (or attention) can be viewed as a process to adaptively adjust or allocate resource intensity towards the most informative or useful components of inputs. There are several methods for exploiting gating mechanism to improve image classification [24, 35, 36] and detection [38]. GBD-Net [38] proposes a gated bi-directional CNN for object detection that passes messages between features from different regions and uses gated functions to control message transmission. SENet [24] uses gating mechanism to model channel-wise relationships and enhances the representation power of modules throughout the networks. In this paper, we introduce an effective gating mechanism that combines attention both locally and globally.

3 Method

We begin with presenting Iterative Feature Reuse, a structure that combines adjacent layer features iteratively for object detection. Following that, we introduce how an elaborately designed gating mechanism is used to adaptively control supervision at multiple scales in a deep network. Finally, we show how the above two structures can be applied in DSOD [6], [32] and SSD [25] seamlessly to obtain GFP-DSOD [32]/SSD [25].

3.1 Iterative Feature Re-Utilization

Our goal is to utilize multi-scale CNN features in a single prediction pass that combines detailed shape and appearance cues from lower level layers and semantics from higher levels. Our proposed feature reuse structure is split into three parts, shown in Fig. 3. First, a down-sampling pathway takes the low level feature maps, and through a down-sampling block outputs the low-resolution feature maps that should be concatenated with the current features. Then, to incorporate higher-level semantics, an up-sampling pathway is adopted to concatenate high level features to the current layer. Finally, we repeat the concatenation operation in each scale of prediction layers in an iteration fashion.

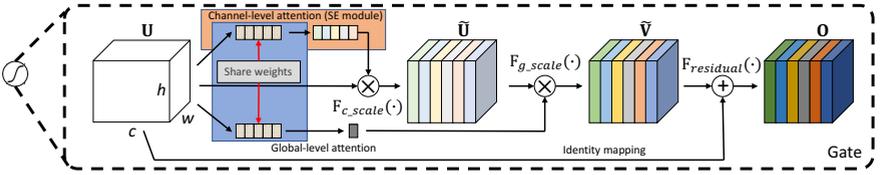


Figure 4: Illustration of the structure of a gate, including: (i) channel-level attention; (ii) global-level attention; and (iii) identity mapping. More details are given in §3.2.

Down/Up-Sampling Pathways. The down-sampling pathway consists mainly of a max-pooling layer ($kernel\ size = 2 \times 2, stride = 2$), followed by a conv-layer ($kernel\ size = 1 \times 1, stride = 1$) to reduce channel dimensions, which is similar to the DSOD down-sampling block. The up-sampling pathway generates higher resolution features by upsampling spatially coarser, but semantically stronger features from the adjacent scale (we use nearest resolution in the upper layer for simplicity). We conduct a deconvolutional operation via bilinear upsampling followed by a conv-layer ($kernel\ size = 1 \times 1, stride = 1$) on the spatial resolution features maps. The upsampled maps are then concatenated with features from the down-sampling pathway and the current layer. Hence, each block has very rich multi-resolution features.

Learning one-third and reusing two-thirds. Fig. 3 shows the building block that constructs our feature-pyramids. With coarser-resolution and fine-resolution features, we introduce a bottleneck block with a 1×1 conv-layer plus a 3×3 conv-layer to learn new features. The number of parameters is one-third compared with DSOD.

Concatenation with A Reduplicative Scheme. Each concatenation operation merges feature maps of the same spatial resolution from the down-sampling pathway and the up-sampling pathway. The process is iterated until the coarsest resolution block is generated. For 320×320 input images, we use six resolutions of features for predicting objects. The finest resolution is 40×40 and the coarsest resolution is 2×2 . To start the iteration, we simply choose two adjacent scales of the current resolution as the inputs of down-sampling and up-sampling pathways. We also apply an extra 160×160 resolution as the input of the down-sampling pathway for the finest resolution (40×40) to improve the ability of detecting small objects.

3.2 Gate-Controlled Adaptive Recalibration

Motivation. Our goal is to ensure that the object detection network is able to adaptively select the meaningful scales for objects with different sizes, so that it can enhance the useful features in an appropriate resolution and suppress less useful ones. The gate function \mathbf{F}_{gate} can take any form, such as a fully-connected network or a convolutional network, but should be differentiable. In this paper, we propose to achieve this by using a two-level attention mechanism and an identity mapping before each prediction layer, partially inspired by Squeeze-and-Excitation [14]. A diagram of our gate structure is shown in Fig. 4 and will be described in more detail in the following sections. We will verify the benefit of each component of our gate structure design in §4.1.

Gate Definition. A gate is a series of transformation \mathbf{F}_{gate} that transforms the input feature maps \mathbf{U} to outputs \mathbf{O} ($\mathbf{U} \rightarrow \mathbf{O}$). Let $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_c]$ denote the set of filter maps. Suppose $\mathbf{U}, \mathbf{O}, \tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}} \in \mathbb{R}^{w \times h \times c}$ are all with width w , height h and c channels, where $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ are intermediate states. Denote \mathbf{F}_{c_scale} (\mathbf{F}_c as abbreviation, similarly hereinafter): $\mathbf{U} \rightarrow \tilde{\mathbf{U}}$;

$\mathbf{F}_{g_scale}(\mathbf{F}_g): \tilde{\mathbf{U}} \rightarrow \tilde{\mathbf{V}}$ and $\mathbf{F}_{residual}(\mathbf{F}_r): \tilde{\mathbf{V}} \rightarrow \mathbf{O}$ and therefore a gate can be formulated as:

$$\mathbf{O} = \mathbf{F}_{gate}(\mathbf{U}) = \mathbf{F}_r(\mathbf{F}_g(\mathbf{F}_c(\mathbf{U}))) \quad (1)$$

Channel-level and Global-level Attention. The aim of channel-level attention is to model relationships between channels and the global-level attention is to adaptively enhance or attenuate different scale supervision. We apply Squeeze-and-Excitation block [14] as our channel-level attention which consists of: (i) a *squeeze* stage \mathbf{F}_{sq} for global information embedding; and (ii) an *excitation* stage \mathbf{F}_{ex} for channel-level recalibration. Therefore we can formulate the channel-level outputs as:

$$\tilde{\mathbf{U}} = \mathbf{F}_{ex}(\mathbf{F}_{sq}(\mathbf{U})) \quad (2)$$

The *squeeze* stage can be formulated as a global pooling operation on each channel:

$$s_c = \mathbf{F}_{sq}(\mathbf{u}_c) = \frac{1}{w \times h} \sum_{i=1}^w \sum_{j=1}^h u_c(i, j) \quad (3)$$

where s_c is the c -th element of \mathbf{s} . $\mathbf{s} \in \mathbb{R}^c$ is a vector calculated by global-pooling filter \mathbf{u} . The *excitation* stage is two fully-connected layers plus a sigmoid activation:

$$\mathbf{e} = \mathbf{F}_{ex}(\mathbf{s}) = \sigma(f_c(f_{\frac{c}{16}}(\mathbf{s}))) \quad (4)$$

where $\mathbf{e} \in \mathbb{R}^c$ is the output, σ is the sigmoid function. f_c and $f_{\frac{c}{16}}$ are the two fully-connected layers with output dimensions of c and $\frac{c}{16}$, respectively. Then, we can calculate $\tilde{\mathbf{U}}$ by:

$$\tilde{\mathbf{U}} = \mathbf{F}_c(\mathbf{U}) = \mathbf{e} \otimes \mathbf{U} \quad (5)$$

where \otimes denotes channel-wise multiplication. More details can be referred to the SENets [14] paper. Our global attention takes \mathbf{s} (the output of *squeeze* stage) as input, and we modify the *excitation* stage by generating only one element. The new *excitation* stage $\bar{\mathbf{F}}_{ex}$ (for global attention) can be formulated as:

$$\bar{\mathbf{e}} = \bar{\mathbf{F}}_{ex}(\mathbf{s}) = \sigma(f_1(f_{\frac{c}{16}}(\mathbf{s}))) \quad (6)$$

where $\bar{\mathbf{e}} \in \mathbb{R}^1$ is the global attention. The weight of $f_{\frac{c}{16}}$ is shared between \mathbf{F}_{ex} and $\bar{\mathbf{F}}_{ex}$. Finally, $\tilde{\mathbf{V}}$ is calculated by:

$$\tilde{\mathbf{V}} = \mathbf{F}_g(\tilde{\mathbf{U}}) = \bar{\mathbf{e}} \otimes \tilde{\mathbf{U}} \quad (7)$$

Identity Mapping. We use an element-wise addition operation [10] to obtain the final outputs:

$$\mathbf{O} = \mathbf{U} \oplus \tilde{\mathbf{V}} \quad (8)$$

where \oplus denotes element-wise addition. Fig. 5 shows several examples of feature map visualization before and after the gating operation.

3.3 Feature Reuse for DSOD and SSD

Our proposed method is a generic solution for building iterative feature pyramids and gates inside deep convolutional neural networks based detectors, thus it's very easy to apply to existing frameworks, such as SSD [25], DSOD [61, 62], FPN [22], etc. In the following, we

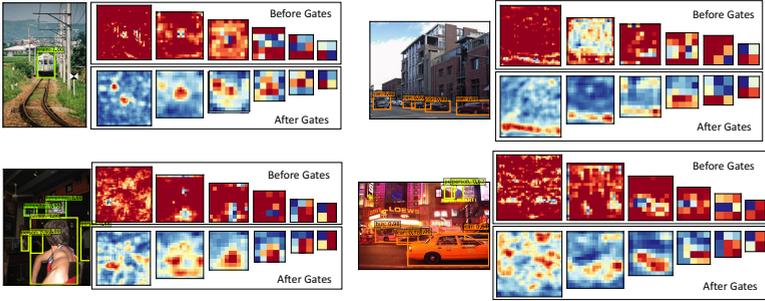


Figure 5: Visualization of feature maps **before** and **after** gates. In each block, left is the input image with detection results. The right-top are the feature maps before gates and the right-bottom are the maps after gates.

Table 1: Ablation Experiments of gate structure design on PASCAL VOC 2007.

Method	mAP (%)
+ Channel-level attention	+0.4 (78.2)
+ Global-level attention	+0.2 (78.4)
+ Identity mapping	+0.2 (78.6)

adopt our method on DSOD [61, 62] and SSD [25] for general object detection, in order to demonstrate the effectiveness and advantages of our method.

There are two steps to adapt *Gated Feature Reuse* for DSOD. First, we apply *iterative feature reuse* to replace the dense connection in DSOD prediction layers. Following that, we add gates in each prediction layer to obtain GFR-DSOD. Other principles in DSOD are inherited in GFR-DSOD like *Stem*, *Dense Block*, etc. For SSD, similar operations are conducted to obtain GFR-SSD. Specifically, we replace the extra layers in SSD with GFR structure and cascade gates in prediction layers. Implementation details and empirical results are given in the next section.

4 Experiments

We conduct experiments on three widely used benchmarks: 20-category PASCAL VOC 2007, 2012 [6] and 80-category MS COCO detection datasets [21]. Following the previous practice of learning detection from scratch [25, 51, 62], we train using the union of VOC 2007 `trainval` and VOC 2012 `trainval` (“07+12”) and test on VOC 2007 test set. For VOC 2012, we use VOC 2012 `trainval` and VOC 2007 `trainval` + `test` for training, and test on VOC 2012 test set. For COCO with 80k images in training set, 40k in validation set and 20k in testing set (`test-dev`). In our study, all experiments are trained from scratch without ImageNet [9] pre-trained models. We adopt the backbone network proposed by DSOD [62] (GFR-DSOD) or VGGNet [64] (GFR-SSD) to ensure fair comparisons.

Implementation details. We adopt SGD for training our models on 8 GPUs. Following [25, 62], we use a weight decay of 0.0005 and a momentum of 0.9. All conv-layers are initialized with the “xavier” method [10]. For other settings, we followed the same implementation as in the original DSOD [62] and SSD [25] papers.

Table 2: Ablation Experiments on PASCAL VOC 2007. “IFP” denotes our iterative feature pyramids. We add additional aspect ratio 1.6 for default boxes at every prediction layer. While for DSOD, we found there was no improvement when adding more aspect ratios.

Method	mAP (%) / #params
DSOD300 [52]	77.7 (14.8M)
GFR-DSOD300	78.9 (14.1M)
DSOD320 [52]	77.8
DSOD320 (using FR only)	78.6
DSOD320 (using gates only)	78.6
GFR-DSOD320	79.2

Table 3: Ablation Experiments of SSD300 from scratch on PASCAL VOC 2007.

Method	# params	mAP (%)
SSD300 [52]	26.3M	69.6
SSD300 (dense pred.) [52]	26.0M	70.4
GFR-SSD300	24.9M	73.2

Table 4: Comparisons of two-stage detectors on MS COCO 2015 test-dev set.

Method	network	pre-train	# param	COCO (Avg. Precision, IoU:)		
				0.5:0.95	0.5	0.75
One-Stage Detectors:						
SSD300 [52]	VGGNet	✓	34.3M	23.2	41.2	23.4
SSD320* [52]	VGGNet	✓	34.3M	25.1	43.1	25.8
DSSD321 [6]	ResNet-101	✓	256M	28.0	46.1	29.2
DSOD320 [52]	DSOD	✗	21.9M	29.4	47.5	30.7
Ours (GFR-DSOD320)	DSOD	✗	21.2M	30.0	47.9	31.8
Two-Stage Detectors:						
FRCNN320/540 [23]	ResNet-101	✓	62.6M	23.3	42.8	23.1
R-FCN320/540+OHEM [8]	ResNet-101	✓	54.4M	25.8	45.8	26.4
Deformable FRCNN320/540 [8, 23]	ResNet-101	✓	64.8M	28.3	46.5	30.2
Deformable R-FCN320/540 [8]	ResNet-101	✓	63.6M	29.5	47.6	30.7
FPN320/540 [24]	ResNet-101	✓	121.2M	29.7	48.1	31.0

4.1 Ablation Experiments on PASCAL VOC 2007

In this section, we investigate the effectiveness of each component of our GFR-DSOD framework. We design several controlled experiments on PASCAL VOC 2007 for the ablation study, including: (i) iterative feature pyramids; (ii) gates; and (iii) two level attention and identity mapping. In these experiments, we train on the union of VOC 2007 `trainval` and 2012 `trainval` (“07+12”), test on the VOC 2007 `test` set.

Effectiveness of Channel attention, Global attention and Identity mapping. Tab. 1 shows the ablation results of gate structure design. After adopting channel attention, global attention and identity mapping, we obtain gains of 0.4%, 0.2% and 0.2%, respectively.

Effectiveness of Iterative Feature Reuse. Tab. 2 (row 4) shows the results of our feature pyramids without the gates. The result (78.6%) is on par with GFR-DSOD320 (row 6) and achieves 0.8% improvement comparing with baseline (77.8%). It indicates that our feature reuse structure contribute a lot on boosting the final detection performance.

Effectiveness of Gates. Tab. 2 (row 5) shows the results of adding gates without the iterative feature pyramids. The result (78.6%) also outperforms the baseline result by 0.8% mAP.

We also show GFR-SSD results in Tab. 3. We can see that our GFR structure helps the original SSD to improve the performance by a large margin. We conjecture the reasons are

two-fold: First, the baseline (69.6%) is still at a very low level. So there is a very large room for performance improvement. Second, VGGNet backbone is a plain network, our GFR structure (skip connection from very low-level to high-level layers) helps this kind of structure greatly. After using more channels and batch norm [16] in each GFR pyramid, our result is further improved to **75.8%** ($\uparrow 6.2\%$ **mAP**) with only **25.4M** model parameters.

4.2 Results on PASCAL VOC 2007&2012

Tab. 2 shows our results on VOC2007 test set. Our GFR-DSOD achieves 79.2%, which is better than baseline method DSOD (77.8%). Fig. 7 shows some qualitative detection examples on VOC 2007 test set with DSOD and our GFR-DSOD models. Our method achieves better results on both small objects and dense scenes

Convergence Speed Comparison. In GFR-DSOD, we observe that models always obtain the best accuracy with 62k iterations. In DSOD, however, the models need around 100k iterations to achieve final convergence, with the same batch size. Thus, GFR-DSOD has relative 38% faster convergence speed than DSOD. Fig. 6 shows the comparison of training and testing accuracy with DSOD method. For the inference time, With 300×300 input, our full GFR-DSOD can run an image at 17.5 *fps* on a single Titan X GPU with batch size 1. The speed is similar to DSOD300 with the dense prediction structure. When enlarging the input size to 320×320 , the speed decrease to 16.7 *fps* and 16.3 *fps* (with more default boxes). As comparisons, SSD321 runs at 11.2 *fps* and DSSD321 runs at 9.5 *fps* with ResNet-101 [14] backbone network. Our method is much faster than these two competitors.

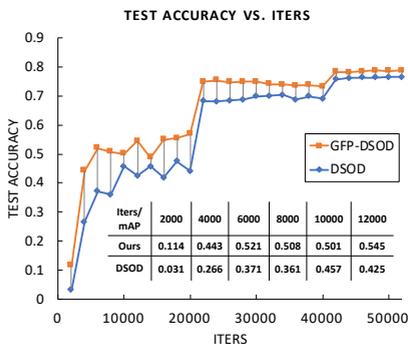


Figure 6: Comparison of DSOD [5] (blue diamonds) and GFR-DSOD (orange squares) on VOC 2007 test set. We show results (mAP) of GFR-DSOD and DSOD at six different iterations. GFR-DSOD obtains much higher accuracy with same training iterations and also achieves faster convergence than the baseline DSOD. Details are given in §4.2.

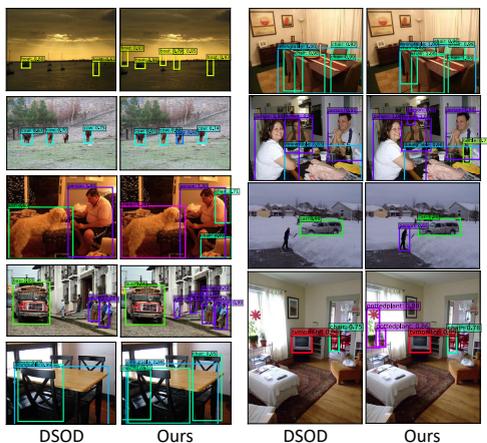


Figure 7: Detection examples on VOC 2007 test set with DSOD / GFR-DSOD models. For each pair, the left is the result of DSOD and right is the result of GFR-DSOD. We show detections with scores higher than 0.6. Each color corresponds to an object category in that image. Our method achieves better performance on both small objects and dense scenes.

We also perform our experiments on VOC 2012 `test` set with two different training subsets of related experimental results: one with models only trained from scratch with VOC 2012 provided data (VOC 2012 `Comp3` Challenge), and another one with additional VOC 2007 data for training. On PASCAL VOC 2012 `Comp3` Challenge, our result (72.5%) outperforms the previous state-of-the-art DSOD (70.8%) [52] by 1.7% mAP. After adding VOC 2007 as training data, our GFR-DSOD320 (77.5% mAP) is consistently better than baseline DSOD320 (76.4%) and some other state-of-the-art methods using VGGNet or ResNet-101 pre-trained models like SSD321 (75.4%) and DSSD321 (76.3%).

4.3 Results on MS COCO

Finally we evaluate our GFR-DSOD on the MS COCO dataset [21]. The batch size is set to 128. Results are summarized in Tab. 4. We first compare our method with state-of-the-art one-stage detectors. We observe that GFR-DSOD can achieve higher performance than the baseline method DSOD (30.0% vs. 29.4%) with fewer parameters (21.2M vs. 21.9M). We then compare our method with state-of-the-art two-stage detectors. It can be observed that our model is much more compact than that of the two-stage detectors using ResNet-101 as backbone. For instance, our results is comparable with FPN320/540 [27] (30.0% vs. 29.7%), but the parameters of our model is only 1/6 of FPN.

5 Conclusions

We have presented *gated feature reuse*, a novel structure design that is more fit for learning detection from scratch, and improves performance in model parameters, convergence speed and accuracy. Extensive experiments on PASCAL VOC 2007, 2012 and MS COCO demonstrate the effectiveness of our proposed method. Since learning from scratch uses limited training data, our future work will focus on adopting GAN-based image generation and data augmentation method like MUNIT [15], INIT [33], etc. to enlarge the diversity of training objects, in order to obtain better detection performance when learning from scratch.

References

- [1] Sean Bell, C Lawrence Zitnick, et al. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *CVPR*, 2016.
- [2] Zhaowei Cai, Quanfu Fan, Rogerio S Feris, and Nuno Vasconcelos. A unified multi-scale deep convolutional neural network for fast object detection. In *ECCV*, 2016.
- [3] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, and Guodong Zhang. Deformable convolutional networks.
- [4] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *NIPS*, 2016.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, et al. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [6] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010.

-
- [7] Cheng-Yang Fu, Wei Liu, Ananth Ranga, Ambrish Tyagi, and Alexander C Berg. Dssd: Deconvolutional single shot detector. *arXiv preprint arXiv:1701.06659*, 2017.
- [8] Ross Girshick. Fast r-cnn. In *ICCV*, 2015.
- [9] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [10] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feed-forward neural networks. In *AISTATS*, 2010.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017.
- [13] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [14] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. *arXiv preprint arXiv:1709.01507*, 2017.
- [15] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018.
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [17] Tao Kong, Anbang Yao, Yurong Chen, and Fuchun Sun. Hypernet: Towards accurate region proposal generation and joint object detection. In *CVPR*, 2016.
- [18] Tao Kong, Fuchun Sun, Chuanqi Tan, Huaping Liu, and Wenbing Huang. Deep feature pyramid reconfiguration for object detection. In *ECCV*, 2018.
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [20] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, 2018.
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, et al. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [22] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [23] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [24] Songtao Liu, Di Huang, et al. Receptive field block net for accurate and fast object detection. In *ECCV*, 2018.

- [25] Wei Liu, Dragomir Anguelov, Dumitru Erhan, et al. Ssd: Single shot multibox detector. In *ECCV*, 2016.
- [26] Joseph Redmon and Ali Farhadi. Yolo9000: Better, faster, stronger. In *CVPR*, 2017.
- [27] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [29] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018.
- [30] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2013.
- [31] Z. Shen, Z. Liu, J. Li, Y. Jiang, Y. Chen, and X. Xue. Object detection from scratch with deep supervision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019.
- [32] Zhiqiang Shen, Zhuang Liu, Jianguo Li, Yu-Gang Jiang, Yurong Chen, and Xiangyang Xue. Dsod: Learning deeply supervised object detectors from scratch. In *ICCV*, 2017.
- [33] Zhiqiang Shen, Mingyang Huang, Jianping Shi, Xiangyang Xue, and Thomas Huang. Towards instance-level image-to-image translation. In *CVPR*, 2019.
- [34] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [35] Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.
- [36] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *CVPR*, 2017.
- [37] Jiahui Yu, Yuning Jiang, Zhangyang Wang, Zhimin Cao, and Thomas Huang. Unitbox: An advanced object detection network. In *Proceedings of the 24th ACM international conference on Multimedia*. ACM, 2016.
- [38] Xingyu Zeng, Wanli Ouyang, Bin Yang, Junjie Yan, and Xiaogang Wang. Gated bi-directional cnn for object detection. In *ECCV*, 2016.
- [39] Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, and Stan Z Li. Single-shot refinement neural network for object detection. In *CVPR*, 2018.
- [40] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018.
- [41] Peng Zhou, Bingbing Ni, Cong Geng, Jianguo Hu, and Yi Xu. Scale-transferrable object detection. In *CVPR*, 2018.
- [42] Xingyi Zhou, Jiacheng Zhuo, and Philipp Krähenbühl. Bottom-up object detection by grouping extreme and center points. *arXiv preprint arXiv:1901.08043*, 2019.