

# FemtoDet: An Object Detection Baseline for Energy Versus Performance Tradeoffs

Peng Tu<sup>1</sup>, Xu Xie<sup>1</sup>, GUO AI<sup>1</sup>, Yuexiang Li<sup>2</sup>, Yawen Huang<sup>2\*</sup>, Yefeng Zheng<sup>2</sup>

<sup>1</sup> MicroBT Inc. <sup>2</sup> Jarvis Lab, Tencent

## Abstract

*Efficient detectors for edge devices are often optimized for parameters or speed count metrics, which remain in weak correlation with the energy of detectors. However, some vision applications of convolutional neural networks, such as always-on surveillance cameras, are critical for energy constraints. This paper aims to serve as a baseline by designing detectors to reach tradeoffs between energy and performance from two perspectives: 1) We extensively analyze various CNNs to identify low-energy architectures, including selecting activation functions, convolutions operators, and feature fusion structures on necks. These underappreciated details in past work seriously affect the energy consumption of detectors; 2) To break through the dilemmatic energy-performance problem, we propose a balanced detector driven by energy using discovered low-energy components named FemtoDet. In addition to the novel construction, we improve FemtoDet by considering convolutions and training strategy optimizations. Specifically, we develop a new instance boundary enhancement (IBE) module for convolution optimization to overcome the contradiction between the limited capacity of CNNs and detection tasks in diverse spatial representations, and propose a recursive warm-restart (RecWR) for optimizing training strategy by considering the data shift produced in popular augmentations. As a result, FemtoDet with only 68.77k parameters achieves a competitive score of 46.3 AP50 on PASCAL VOC and 1.11 W & 64.47 FPS on Qualcomm Snapdragon 865 CPU platforms. Extensive experiments on COCO and TJU-DHD datasets indicate that the proposed method achieves competitive results in diverse scenes.*

## 1. Introduction

The deployment of efficient convolutional neural networks (CNNs) enabled immense progress [23, 29, 30, 31, 4, 9, 41] in vision detectors for edge devices, in which they

consistently reduce parameters and speed counts for improving accuracy. However, these metrics are not correlated well with the efficiency of the models in terms of energy. Evaluation metrics, such as parameters, do not take into account the energy cost of models, resulting in a nontrivial effect on the energy cost of detectors. Compared with the same architecture, the parameters of models are positively correlated with their energy cost (shown in Table 2). However, in the case of equal model parameters, their energy consumption may be negatively correlated or even irrelevant to the model parameters (shown in Table 1). Considering that various activation functions, convolution operators, and feature fuse structures may not increase model parameters, but generate more energy costs. Similarly, the speed count is also not well correlated with energy, as it can be optimized by the degree of parallelism. These disconnections will leave customized efficient detectors unavailable, when they are deployed under severe energy constraints like the always-on surveillance cameras.

This paper aims to reduce the energy costs of efficient object detectors, while improving their performances by achieving energy versus performance tradeoffs. Specifically, these bottlenecks fall into several categories:

1) *Detector components with unknown energy.* Most current object detection methods focus either on latency-oriented [5, 34, 44] or accuracy-oriented [9, 38, 32, 36] works. There are very limited works toward the energy cost of the detector components, which is the first obstacle to designing energy-performance-balanced detectors. To identify energy-efficient components of detectors, we followed [42] to benchmark their energy metrics from three types of structures, *i.e.*, activation functions, convolution operators, and necks of detectors. The following finding is acquired: Firstly, some activation functions have been widely used, because of their ability to improve models without adding more parameters. They have increased costs, yet received very limited attention. As shown in Table 1, we set ReLU [1] based models as a baseline. When replacing the ReLU with the GELU [16], 4.40% performance is gained, while 12.50% energy cost is increased, and the corresponding factor of the mean energy versus the perfor-

\*Corresponding author; First author's Email: yh.peng.tu@gmail.com

mance tradeoffs (Eq. 1) drops 7.16%. In addition, although large kernel convolutions [25] can improve models (Table 2), the increased energy cost is unacceptable. As seen in Table 2, the large kernel convolutions are equipped to enable the model to rise by around 1.87% in performance, but with more energy costs (16.20%) than in the small kernel size of convolutions (shown in the second row of Table 2). Finally, the standard FPN [21, 22, 39, 10, 37] in the detectors also causes significant energy costs. We observe that this situation is happened due to multi-feature fusion paths following either bottom-to-up or top-to-down way, resulting in frequent data reading in memory for data coverage. However, it is shown that the up-down fusion between multi-features may not necessary. As shown in Table 3, we propose a simple SharedNeck for replacing FPN to reduce around 5.77% energy costs, while obtaining 6.25% performance gains. SharedNeck uses an adaptive fusion between multi-features by learning convolution, instead of feature-by-feature bottom-to-up or top-to-down fusions, as operated in FPN. Based on these analyses, we further build upon a low-energy detector, named FemtoDet. Surprisingly, FemtoDet just has 68.77k parameters and 1.1W power on the platform. 2) *Optimization for CNN* is another bottleneck, since obtaining a favorable detector is very challenging, especially for a small number of parameters. Light-weight detectors are restricted by their limited capacities [6], which will cause obfuscation on feature maps of the instances' boundaries of interesting objects, as shown in Figure 2 (b). Ambiguous instance boundaries on the features may raise the risk of false detection in the models. For this problem, we propose a novel instance boundary enhancement module (IBE), which emphasizes the potential of the object boundary information: The parameter reuse mechanism is first applied to integrate local descriptors with convolution operation to make robust and diverse instance boundary feature representation. But this operation would cause feature un-alignment between normal features and integrated features. Then, in order to sufficiently exploit these two types of features, we design dual-normalization in IBE, as to realign the features (Fig. 4 of the Appendix). IBE employs a shared convolution yet an independent batch normalization layer for separate normalization is also adopted. Finally, we add normal and integrated features together behind the dual-normalization layer. IBE provides a novel form of parameter reuse to generate new local descriptors from shared convolution operators. This method can capture object boundary information by integrating gradient cues around the learned convolution. After trained, IBE modules can be folded into simple convolution operators, thus requiring no extra computations anymore; In addition, data augmentation is a common approach for efficient detector training. Using well-designed strong augmentations is conducive to improving the generalization of mod-

els [31, 9]. However, how to prevent data shifts between training and validation images has not been explored, while training data suffers strong augmentations. It is found that the produced data shifts by strong augmentations would prevent light-weight detectors from moving toward global optimization. A common view is that strong image augmentations can effectively encourage networks to learn diverse features [11]. But for light-weight detectors, these diverse features cannot assist the models in making better generalizations in the validation set. In other words, light-weight detectors are more susceptible to these ignored data shifts because of their limited capacities. Furthermore, we propose an effective training strategy, namely recursive warm-restart (RecWR), to adapt these diverse features to improve the model's generalization. RecWR works based on multi-stage training while gradually weakening data augmentation strength. This method can help the limited-capacity detectors jump out of local optima under the assistance of high-dimension & diverse features. The effectiveness of the IBE and RecWR has been evaluated on PASCAL VOC dataset. And the experimental results show that IBE can improve FemtoDet  $\sim 7.72\%$  performance, while incurring no extra parameter burdens, as compared to the original architecture; RecWR can improve FemtoDet  $\sim 6.19\%$  performance by gradually weakening the data augmentation strength in multi-stage learning. By jointly training FemtoDet with IBE and RecWR, our proposed methodology can go over the YOLOX [9] by 51.34% in performance when using the same-level parameters.

Something worth mentioning is that FemtoDet is specially designed for hierarchical intelligent chips to enable always-on alerts: the always-on low power, high-recall, and decent accuracy - it achieved 85.8 AR20 & 76.3 AP20 (Table 7) while performing pedestrian detection on TJU-DHD. In addition, FemtoDet is poor in detecting small objects but is excellent in detecting medium and large objects, which they are more interested in, with 88.8 AR20-m & 94.1 AP20-m / 95.3 AR20-l & 98.6 AP20-l (Table 7) on TJU-DHD. Information can be passed to other models after identifying possible objects of interest, and high-precision robust models are then initiated for accurate recognition. The always-on smart products have a wide range of applications, e.g., in-home monitoring or robots. Hence, the loose metric (e.g., AP50 or AP20) and data scenario with moderate difficulty (e.g., VOC) can well reflect the application ability of FemtoDet. Further, the experiments on COCO datasets have verified that the proposed method is suitable for diverse scenarios resulting in competitive results.

## 2. Related Works

### 2.1. Object Detection

Object detection is a classic computer vision task for identifying the category and location of object in pictures

or videos. The existing object detectors can be divided into two categories: two-stage detectors and one-stage detectors. Two-stage detectors [32, 36, 14, 21] are anchor-based devices that generate region proposals from the image and then produce the final prediction boxes from these proposals. Further, FPN [21] improves the two-stage detectors by fusing multi-level features. Despite their higher accuracy compared to the one-stage detectors, two-stage detectors are still difficult to achieve low latency when deployed on edge devices.

Specifically, there are two types of one-stage detectors: anchor-based and anchor-free, which depends on anchor priors injecting into whole images to enable box regression. SSD [23], as a classical anchor-based one-stage detector, discretizes the output space of bounding boxes into a set of default anchors with different aspect ratios and scales per feature map location. This operation helps detecting tiny objects. Another typical anchor-based detector is the YOLO series [30, 31, 4]. YOLOv2 [30] explores bag of freebies to improve the performances of one-stage detectors. YOLOv3 [31] proposes cross-scale features and novel nms (Non-Maximum Suppression) to obtain more confident predictions. YOLOv4 [4] finds that the repeated gradient of network optimization aggravates detector latency, and thus designs Cross-Stage-Partial-connection modules to reduce the latency of detectors while maintaining detection performance. Anchor-free detectors [29, 9, 19, 38] aim to eliminate the pre-defined set of anchor boxes. YOLO [29], as an anchor-free one-stage detector, divides an image into multiple grids and predict boxes at grids near the object's center. CornerNet [19] detects an object as a pair of key points (through the top-left corner and bottom-right corner of the bounding box). A single convolutional network was leveraged in CornerNet to predict a heatmap for all top-left corners of instances of the same object category, a heatmap for all bottom-right corners, and an embedding vector for each detected corner. Indeed, CornerNet is a new pipeline for object detection. FCOS [38] eliminates the anchor setting by proposing a fully convolutional one-stage object detector, which can solve object detection in a per-pixel prediction manner. The anchor-free detectors solve the problem happened in anchor-based detectors, which reduce the memory cost and increase accuracy of the bounding box.

Both one-stage and two-stage object detection methods have obtained high performance on many challenging public datasets, *i.e.*, COCO and TJU-DHD [27]. These methods provide the accuracy-oriented detectors. However, the critical problem for detectors serving edge devices is detection latency and its power. In other words, detectors obtain higher performance in challenging scenes, which is not a necessary option for detectors deployed on edge devices. For the detection latency, a lot of effort [5, 34, 44] have been devoted to achieving the balance between accuracy

and efficiency. FastYOLO [34] is an optimized architecture extended from YOLOv2 with fewer parameters and 2% performance drop, which makes FastYOLO run at an average of  $\sim 18$ FPS on an Nvidia Jetson TX1 embedded system. YOLObite [5] focuses on pruning redundancy object detectors to be real-time by designing in a compression-compilation way. NanoDet uses ShuffleNetV2 [26] as its backbone to make the model lighter, and further uses ATSS [47] and GFL [20] to improve its accuracy. It is worth noting that NanoDet also reached  $\sim 60$ FPS on an ARM CPU. These light-weight detectors are latency-oriented. Actually, neither accuracy-oriented nor latency-oriented detectors take their energy cost into account. High energy cost detectors are unfriendly to devices deployed at edges.

This paper aims to develop a light-weight detector that can achieve tradeoffs between energy and performance.

## 2.2. Energy-Oriented Convolutional Neural Networks

Besides the manual design of deep neural networks, a great deal of work has been done to improve the effectiveness and efficiency of deep neural networks via network pruning [45], quantization [13, 2], architecture search [12, 40]. ECC [42] proposes the end-to-end classifier training framework, which provides quantitative energy consumption guarantees by using weighted sparse projection and input masking. Zhang et al. [46] specialized in video broadcast and leveraged distillation techniques to reduce memory consumption by approximating the video frame data. MIME [3] is an algorithm-hardware co-design approach that reuses the weight parameters of a trained parent task and learns task-specific threshold parameters for inferring on multiple child tasks. However, the above works are all in service for image classification, and we can observe that fewer parameters represent lower energy costs in the same architecture. Object detection on edge devices is still a challenging topic, considering that the conflict between the limited capacity of CNNs and diverse spatial representations. This paper is the first work to provide a systematized solution for energy-oriented light-weight detectors.

## 3. FemtoDet

This section will be divided into two subsections to describe how FemtoDet achieves energy versus performance tradeoffs. 1) *Benchmarks for Low-energy Detectors*: We provide benchmarks for designing low-energy detectors, including the exploring of activation functions, convolution operators, and neck of detectors. 2) *Energy-Oriented FemtoDet*: Based on the benchmarks for designing low-energy detectors, an energy-oriented light-weight detector named FemtoDet is provided, which stacks layers from depthwise separable convolutions (DSC), BN and ReLU. FemtoDet involves only 68.77k parameters and power of 1.11W on

the Qualcomm Snapdragon 865 CPU platform. Moreover, FemtoDet can be optimized with two designs: Firstly, instance boundary enhancement (IBE) modules are used to improve the DSC in FemtoDet, overcoming the bottleneck of light-weight models’ representation optimization (*i.e.*, models will learn obfuscation features because of their limited capacity, as shown in Fig. 2 (b)). Second, the recursive warm-restart (RecWR) training strategy is a multi-stage recursive warm-restart learning procedure, which can surmount data shifts produced from strong data augmentations.

### 3.1. Benchmarks for Low-energy Detectors’ Designing

#### 3.1.1 Evaluation Metrics

Top1-Acc (top1 accuracy, for image classification), and AP (average precision, for object detection) are widely used measures to evaluate the performance of CNN. In addition to the commonly used metrics top1 – Acc and AP, we propose the Power (energy costs) and mEPT (mean energy versus performance tradeoffs) to comprehensively evaluate the energy costs of the models and their ability to achieve energy versus performance tradeoffs. Specifically, Power and mEPT are represented by

$$\text{mEPT} = \left( \frac{1}{N} \sum_i^N P \right) / (\text{Power}), \quad (1)$$

$$\text{s.t. Power} = \frac{1}{T_{\text{time}}} \sum_i^N (\varepsilon(\mathbf{W}, x_i) - \varepsilon(\hat{\mathbf{W}}, x_i)), \quad (2)$$

where  $T_{\text{time}}$  indicates the total times required to evaluate  $N$  images,  $\varepsilon(\mathbf{W}, x_i)$  represents the energy costs while models evaluate  $i$ -th image  $x_i$  with trained parameters  $\mathbf{W}$ .  $\varepsilon(\hat{\mathbf{W}}, x_i)$  is similar to  $\varepsilon(\mathbf{W}, x_i)$ . The difference is that we set a value of 1 for each layer of channels, where the model parameter is  $\hat{\mathbf{W}}$ .  $\varepsilon(\hat{\mathbf{W}}, x_i)$  indicates that the model is an empty state.  $P$  denotes the model performance (top1-Acc for image classification while evaluating activation functions and convolution operators, AP50 for object detection while evaluating necks of detectors in Sec. 3.1.2) in image  $x_i$ . For the power metric, the value is expected to be as small as possible. For the mEPT metric, the result is expected as greater as possible.

#### 3.1.2 Components to Be Evaluated

To identify low-energy components of detectors, we followed [42] to benchmark their energy costs from three types of structures: activation functions, convolution operators and necks of detectors:

*Activation Functions.* Rectified Linear Unit (ReLU) [29, 23], GELU [16], Swish [9], SiLU [8], and *etc.* are widely used for object detection because they have fewer

parameters and flops. We calculate the energy costs of different activation functions in the same architecture to explore which activation function is more friendly for designing an energy-oriented detector.

*Convolution Operators* can automatically learn filter weights while CNNs are trained to summarize helpful information, including horizontal, vertical, edged, diagonal, and *etc.* features from images. Vanilla convolutions (vanConv) [18], depthwise separable convolutions (DSC) [17] and their large kernel size versions [25] are widely used. Ding *et al.* [7] pointed out that when CNNs are constructed based on enough large convolution kernels, results can be comparable to the transformer [24]. Here, we set up energy consumption comparison experiments with different convolutions (vanCon and DSC) and multi-scale convolution kernel sizes ( $3 \times 3$ , and  $5 \times 5$ ) on different network structures (ResNet and MobileNetV2). It is important that for each experiment case, the convolution operator type or kernel size would be the only variable.

*Necks of Detectors:* Feature pyramid networks (FPN) are essential components for two-stage or one-stage detectors. FPN-based necks of detectors fuse multiple low-resolution and high-resolution feature inputs for better representations, leading to a series of studies for designing manually complex fusion methods [21, 22, 39]. However, they only presented favorable results brought by FPN but ignored its energy costs. We form a series of experiments on FemtoDet by using different neck types, *e.g.* FPN [21], PAN [22], and our proposed SharedNeck (details can be found in Sec. 3.2.1) to study corresponding metrics (including energy costs, parameters, and detection precision on PASCAL VOC) changes.

All related experiments can be found in Sec. 4.1. We can observe that DSC and ReLU are more energy friendly than vanCon and other activation functions. In addition, by comparing with the necks of detectors such as FPN and PAN, the proposed SharedNeck is more suitable for light-weight detectors.

### 3.2. Energy-Oriented FemtoDet

#### 3.2.1 Building FemtoDet

*Backbone:* The backbone of FemtoDet contains an initial full convolution layer with 8 filters. We use ReLU as the non-linearity and BN as the batch normalization. From the second layer beginning, all the convolution operators use DSC. The reason is that we follow the results of our benchmarks (the results can be seen in Sec. 4.1) to select the energy friendly components. The whole backbone is described in Appendix Table 9; *Neck:* We customize a neck for FemtoDet to achieve better tradeoffs between energy and performance. In detail, SharedNeck first aligns the scale information of inputs from backbones, and then merges these alignment features with elements addition. Fi-

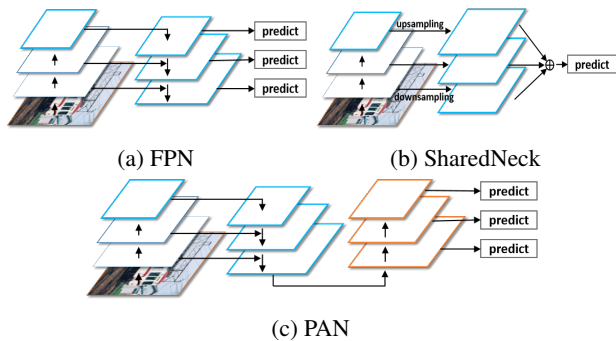


Figure 1: (a) FPN: A top-down multi-scale feature fusion architecture that outputs multiple predictions; (b) Shared-Neck: Simple adaptive multi-scale feature fusion architecture with single prediction, enabling tradeoffs between energy and performance; (c) PAN [22]: Exploring extra bottom-to-up multi-scale feature fusions with many predictions based on FPN.

nally, a DSC performs adaptive multi-scale information fusion between merged features. The implementation details of SharedNeck, and how it differs from other necks are shown in Fig. 1; *Head and Training Loss*: Here, we use the decouple head of YOLOX as the head of our proposed detector with the same training loss as YOLOX.

### 3.2.2 Instance Boundary Enhancement Module

Optimizing light-weight detectors is a notoriously challenging problem. The reason for this is that limited by their representation, features learned by detectors are diffuse, as shown in Fig. 2 (b). Instance boundary enhancement (IBE) modules are designed to improve the DSC in FemtoDet, thus overcoming the bottleneck of light-weight models’ representation optimization. IBE is similar to modules introduced in [45, 35, 43], except that our modules are designed for convolutional layers that are factorized into depthwise and pointwise layers. We further introduce a dual-normalization mechanism, such that the IBE modules can be used for object detection, while the modules in [45, 35, 43] just working for low-level tasks. Based on the DSC block, the basic block of IBE includes  $3 \times 3$  depthwise convolution followed by  $1 \times 1$  convolution. IBE enhances DSC by designing a new local descriptor, semantic projector, and dual-normalization layer. Specifically, the  $1 \times 1$  local descriptor is a parameter reusing mechanism generated by the linear transformation of the integrating gradient cues around shared depthwise convolution. Therefore, the object boundary information can be found in the local descriptors. Previously, the local descriptor is known as the difference convolution [45, 35, 43] used for low-level tasks. Although the idea is easy to get that using these object boundary information to enhance the noisy feature representation (shown in Fig. 2 (b)) of the above standard operators, such

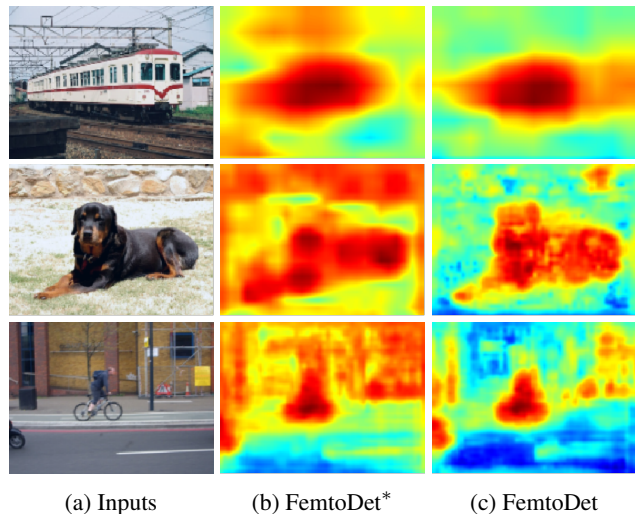


Figure 2: Feature visualization of the trained light-weight detector. (a) Inputs: RGB images from PASCAL VOC; (b) FemtoDet\*: Trained FemtoDet detector consisting of pure depthwise separable convolution; (c) FemtoDet: Trained FemtoDet detector, which was trained with the IBE modules.

as depthwise convolution, in practice, the original difference convolution architecture cannot serve for high-level semantic tasks. We observe that features obtained by the difference convolution are unaligned with the features obtained by standard convolution (the figure shown in Fig. 4 of the Appendix). On the one hand, the difference convolution captures diverse information about object boundaries by integrating gradient cues around  $3 \times 3$  or with larger convolution. On the other hand, high-level tasks encourage standard convolution to summarize the abstract semantic information of images. That is the reason that the classical difference convolution cannot serve high-level tasks. To simultaneously solve this problem, we propose a semantic projector and a dual-normalization layer. The semantic projector is the linear transformation of depthwise convolution to transfer operators related to semantic information extraction; The dual-normalization layer consists of two independent batch normalization modules specifically designed for aligning the unaligned features. After that, we incorporate the object boundary cues from the feature addition between the local descriptor and depthwise convolution to guide the model in learning an effective representation of the instance (refined results can be seen in Fig. 2 (c)).

The description of IBE modules is given in Fig. 3 of the Appendix. For the depthwise convolution of kernel size 3,  $C_{in}$  denotes the input channel dimension,  $C_{out}$  represents the output channel dimension, the weight matrix can be denoted as  $\mathbf{W}_{conv} \in \mathbb{R}^{1 \times C_{in} \times 3 \times 3}$ , and the bias is represented as  $\mathbf{b} \in \mathbb{R}^D$ . The local descriptor of kernel size 1 can be

generated from the depthwise convolution by reusing parameters. Its weight matrix is denoted as  $\mathbf{W}_{\text{des}}$ , which is integrating the gradient cues around  $\mathbf{W}_{\text{conv}}$ :

$$\mathbf{W}_{\text{des}} = -\theta_1 \cdot \sum_{p_n \in \mathbf{R}} \mathbf{W}_{\text{conv}}(p_n), \quad (3)$$

where  $\mathbf{W}_{\text{des}} \in \mathbb{R}^{C_{\text{out}} \times C_{\text{in}}}$ ,  $\theta_1 \in [0, 1]$ ,  $\in \mathbb{R}^{1 \times 1}$  is a learnable parameter projection factor.  $\mathbf{R}$  represents the  $3 \times 3$  convolution in  $\mathbf{W}_{\text{conv}}$ , and  $p_n$  is the  $n$ -th weight value. Integrating the gradient cues around  $\mathbf{W}_{\text{conv}}$  can help the local descriptors capture the object boundary information effectively. The semantic projector is generated from the depthwise convolution, while diverse semantic presentations can be obtained through a learnable linear transformation, shown as follows:

$$\mathbf{W}_{\text{pro}} = \theta_2 \cdot \mathbf{W}_{\text{conv}}, \mathbf{W}_{\text{pro}} \in \mathbb{R}^{C_{\text{out}} \times C_{\text{in}} \times 3 \times 3}, \quad (4)$$

where  $\theta_2 \in [0, 1]$ ,  $\mathbb{R}^{1 \times 1}$  is another learnable parameter as the projection factor. Subsequently, the IBE modules perform four steps to refine the feature representation under the guidance of object boundary information. 1) inputs  $x_{\text{in}}$  will be convolved with the obtained three convolution operators, and the corresponding results are denoted as  $x_{22}$ ,  $x_{21}$ , and  $x_{23}$ , respectively; 2) elements addition are conducted between  $x_{22}$  and  $x_{21}$ ; 3) the dual-normalization layer enables feature distribution normalization on  $x_{21}$  and  $x_{23}$ ; 4) the two outputs from the dual-normalization layer can be added together; As a result, the aforementioned added features are convolved through pointwise convolution as the final output of the IBE module.

In addition, we leverage the homogeneity and additivity of the convolutions to fold IBM modules into the simpler depthwise separable convolution in the inference stage without performance degradation. The details can be seen in Appendix C.

### 3.3. Recursive Warm-restart Training Strategy

Strong augmentations (SA) are widely used for object detection, and [9, 31] shows well-designed SA effectively improves detectors. However, we find that well-designed SA does not always benefit light-weight detectors, because the current training strategies cannot sufficiently exploit diverse training representations to boost the generalization ability on real validation data. For example, YOLOX [9] points out MixUp reduces YOLOX-nano’s performance by about 5%. We argue that the limited capacity detectors fit the diverse data generated by SA as much as possible during training, such that the modules have no surplus capacity to adjust learned features to boost the generalization ability of modules on real validation data. In other words, SA produces diverse representations that are equal to the data

shifts, thus breaking the generalization ability of the module. We also propose an effective training strategy to adjust these diverse features to improve the generalization ability, namely recursive warm-restart (RecWR).

As shown in Fig. 5 of the Appendix, the whole training procedure can be divided into four stages. Starting from the 1-st stage to 4-th stage training, the intensity of image augmentations gradually decreases. Specifically, in the 1-st training stage, some SA types will be combined, such as MixUp, Mosaic, and RandomAffine. Starting from the 2-nd training stage, the above SA types are gradually unloaded in each training stage, until the 4-th training stage. For the last training stage, only the random flip and random scales are performed on the training data. In addition, before starting each training stage, the waiting training detectors load the trained weights of the previous training stage as the initialization.

We can see that in Appendix D, after training FemtoDet with RecWR, MixUp also help such extremely small detectors obtain better performance. In other words, RecWR takes advantage of diversity features learned from SA to get FemtoDet out of sub-optimization.

## 4. Experiments

We used VOC, COCO, TJU, and ImageNet datasets to conduct experiments and validate the proposed method. Specifically, whole experiments can be divided into two sections, i.e., identifying low-energy components of detectors mentioned in Sec. 3.1.2, and validating the effectiveness of our designed FemtoDet, where the Power metric was measured in GTX 3090.

Table 1: Comparisons of energy-related metrics in MobileNetV2 0.25 (MBV2) [33] used different activation functions. We trained a classifier on ImageNet to evaluate relevant metrics.

Activation Functions	Param (M)	Power (W)	Top1 (Acc)	mEPT
ReLU [1]		<b>5.04</b>	45.22	<b>8.97</b>
GELU [16]		5.67	47.21	8.33
Swish [9]	1.48	5.47	47.45	8.67
HSwish [44]		5.33	<b>47.60</b>	8.93
SiLU [8]		5.73	47.25	8.25

### 4.1. Benchmark to Find Low-energy Components

#### 4.1.1 Activation Functions

Table 1 shows the energy-related metric results using different activation functions under the same architecture (MobileNetV2 0.25). We observe that the activation functions can significantly improve models without extra parameters (Param). On the other hand, such activation functions incur

unacceptable overhead on some metrics which are inconvenient to measure. For example, Swish improves ReLU-based models performance by 4.93%, but produces more than 8.53% more energy cost. The mEPT metric also demonstrate that ReLU results in the best energy versus performance tradeoffs.

Table 2: Comparisons of energy-related metrics in MobileNetV2 0.75 and ResNet18 [15]. The two CNNs can present the two typical convolution operators: vanilla convolutions and depthwise separable convolutions. We compare different convolution operators with different kernel sizes on the same convolution operators. Here, ♣, and ◇ indicate two kernel sizes of 3×3, and 5×5. In this experiment, we also evaluate models by training them on ImageNet.

CNNs	Convolution Operators	Param (M)	Power (W)	Top1 (Acc)	mEPT
Res18	vanConv♣	12.46	170.64	69.93	0.41
	vanConv◇	31.99	198.28	71.24	0.36
MBV2	DSC♣	2.29	5.56	69.51	12.50
	DSC◇	2.35	6.07	70.96	11.69

#### 4.1.2 Convolution Operators

This sub-section compares energy-related metrics while training an image classifier on the ImageNet dataset based on the different convolution operators (vanCon and DSC), and the different kernel sizes (3×3, and 5×5) on the same convolution operator. Experimental results in Table 2 shows that: 1) When CNNs were constructed with the same kernel size, mobilenetv2 generated much lower energy costs than resnet18, as shown in Table 2 of vanConv♣ and DSC♣. In other words, DSC is more energy friendly than vanConv while achieving similar performance; 2) When CNNs were constructed with different kernel sizes and the same types of operators, the smaller kernel generated much lower energy costs than the larger kernel, as shown in Table 2 of vanConv♣ and vanConv◇. However, the performance gain from the large kernels is not so impressive.

#### 4.1.3 Necks of Detectors

Here we evaluated the effects of different necks (including FPN, PAN, and SharedNeck) in FemtoDet. Experimental results are shown in Table 3, and some observation are summarized as follow: 1) Although FPN enables FemtoDet to achieve preferable results, it not only consumes more energy but also has a large parameter overhead; 2) Due to the limited representation of light-weight models, the top-down and bottom-up feature fusion of PAN is poor; 3) Shared-Neck achieves the best results in the metrics of parameter, energy cost, object detection performance, and factors of

mean energy versus performance tradeoffs by performing adaptive feature fusion.

Table 3: Comparisons of energy-related metrics in FemtoDet with different necks: FPN, PAN, and our proposed SharedNeck, respectively. The detector was trained on the PASCAL VOC.

Necks of Detectors	Param (k)	Power (W)	AP50	mEPT
FPN	176.31	8.31	40.04	4.82
PAN	79.83	7.97	39.91	5.01
SharedNeck	<b>69.77</b>	<b>7.83</b>	<b>42.50</b>	<b>5.43</b>

Table 4: PASCAL VOC object detection results with YOLOX, NanoDet Plus, and FemtoDet.

Methods	Param (k)	Power (W)	AP50	mEPT
YOLOX	70.39	10.42	30.60	2.94
NanoDet Plus	69.60	11.29	38.44	3.40
FemtoDet	<b>68.77</b>	<b>7.83</b>	<b>46.31</b>	<b>5.91</b>
Methods	mAP	mAP-s	mAP-m	mAP-l
YOLOX	13.60	3.00	5.10	16.60
NanoDet Plus	20.78	<b>4.02</b>	8.07	25.48
FemtoDet	<b>22.90</b>	1.10	<b>8.40</b>	<b>28.50</b>

Table 5: Inference speed (FPS) and power of trained detectors (YOLOX, NanoDet Plus, and FemtoDet were trained on PASCAL VOC) on Qualcomm Snapdragon 865 CPU platforms.

Methods	Param (k)	AP50	Inference	
			Power (W)	FPS
YOLOX	70.39	30.60	1.41	38.04
NanoDet Plus	69.60	38.44	1.37	47.80
FemtoDet	<b>68.77</b>	<b>46.31</b>	<b>1.11</b>	<b>64.47</b>

#### 4.2. Validating the Effectiveness of FemtoDet

We verify the effectiveness of FemtoDet on three datasets: PASCAL VOC, COCO, and TJU-DHD, while resizing inputs to 640×640 for training and resizing inputs to 416×416 for validating. Two datasets, PASCAL VOC and TJU-DHD two datasets are converted to the COCO data types for evaluation. In addition, we extract campus data from the TJU-DHD dataset to evaluate the pedestrian detection performance of our proposed extremely light detector, i.e., FemtoDet. Considering extremely light-weight detectors are challenging to fit the complex COCO dataset with poor detection performance, we pull the corresponding results to Appendix E. Also, we compare the detection performance of FemtoDet with YOLOX, and NanoDet Plus [28], in which they are at the same parameter level to ensure fairness. All the detector backbones are pretrained on the Ima-

geNet for 100 epochs, and the metric of mEPT presented in this section is calculated between AP50 and Power.

#### 4.2.1 Results on PASCAL VOC

Experimental results are presented in Table 4, which shows light-weight detectors with eight metrics: parameter (Param), energy cost (Power), AP50, factors of mean energy versus performance tradeoffs (mEPT), mAP, mAP of small size objects (mAP-s), mAP of medium size objects (mAP-m), and mAP of large size of objects (mAP-l) respectively. We can see that although YOLOX has a large number of parameters, it is much lower than FemtoDet in many metrics. NanoDet Plus has the parameter comparable to that of FemtoDet, but only outperforming FemtoDet in the metrics of mAP-s. We believe that for such extremely light-weight detectors, their performances on large-scale objects are the most critical. FemtoDet also presented better results than YOLOX and NanoDet Plus on the metrics of mAP-l. Although the metrics i.e. mEPT, used to evaluate performance and energy balance, FemtoDet achieves the best balance compared to the other two models.

In addition, we evaluated the inference speed (FPS) and power of trained detectors on Qualcomm Snapdragon 865 CPU platforms. As shown in Tab. 5, our FemtoDet also achieves the smallest energy costs and inference speed on edge devices.

Table 6: TJU-DHD pedestrian detection results with YOLOX, and FemtoDet.

Methods	Param (k)	Power (W)	AP50	mEPT
YOLOX	<b>67.37</b>	9.85	59.40	6.03
FemtoDet	67.54	<b>6.92</b>	<b>66.90</b>	<b>9.67</b>
Methods	mAP	mAP-s	mAP-m	mAP-l
YOLOX	29.00	11.50	37.60	54.20
FemtoDet	<b>33.10</b>	<b>17.60</b>	<b>50.20</b>	<b>61.50</b>

Table 7: Loose metric on TJU-DHD pedestrian detection with YOLOX, and FemtoDet.

Methods	Param (k)	Power (W)	AR20	AP20
YOLOX	<b>67.37</b>	9.85	<b>90.40</b>	42.10
FemtoDet	67.54	<b>6.92</b>	85.80	<b>76.30</b>
Methods	AR20-m	AP20-m	AR20-l	AP20-l
YOLOX	83.70	34.80	96.10	55.60
FemtoDet	<b>94.10</b>	<b>88.80</b>	<b>98.60</b>	<b>95.30</b>

#### 4.2.2 Results on TJU-DHD

Extremely light-weight detectors like FemtoDet have broad application prospects in surveillance scenarios. Hence, we

evaluate the performance of FemtoDet in pedestrian detection using the campus dataset of TJU-DHD, and show the performance of the detectors on the less stringent metric AP20. Table 6 and 7 demonstrate that FemtoDet is competent for pedestrian detection in common surveillance scenarios: 1) FemtoDet has lower energy costs in deployment, which is 29.74% lower than YOLOX; 2) FemtoDet obtains 76.3% AP20, while YOLOX with the same parameters only obtain 71.80% AP20.

Table 8: Ablation study on the effectiveness of our IBE modules and RecWR training strategy. Where FemtoDet\* denotes FemtoDet consists of pure DSC without the enhancement of IBE, 300e and 1200e mean the detector training with the same data augmentation with 300 or 1200 epochs.

\	IBE	RecWR	300e	1200e	AP50
	✗	✗	✓	✗	42.50
	✗	✗	✗	✓	42.11
FemtoDet*	✗	✓	✗	✗	45.13
	✓	✗	✓	✗	45.78
	✓	✓	✗	✗	46.31

#### 4.2.3 Ablation Studies

We conduct a series of ablation studies on PASCAL VOC to demonstrate the effectiveness and efficiency of our IBE modules and RecWR training strategy. As shown in Table 8: 1) The second and third rows present that by training with 300 and 1200 epochs for FemtoDet\*, such a longer epochs do not yield better gains. 2) The fourth and fifth rows present FemtoDet\* was trained with RecWR and 300 epochs with the enhancement of IBE, respectively. Results show that independently equipping IBE and RecWR can improve model performance. 3) The six rows demonstrate that jointly using IBE and RecWR enables FemtoDet to achieve the best performance.

### 5. Conclusion

This paper presented a baseline to encourage the research on energy and performance balance detectors. Our experimental results clearly showed that boosting performance can also cause problem of growth in energy consumption. Conversely, the simple components like ReLU are suitable for building energy-oriented detectors. In addition, we also proposed a novel IBE module and RecWR training strategy to overcome the optimization problem of extremely light-weight detectors. Compared to other state-of-the-art methods under the same parameters setting, IBE and RecWR support these baselines achieving the best performance on VOC, COCO, and TJU-DHD datasets, while consuming the



least energy. In the future, we will continue improving the detectors on the balance of energy and performance.

## References

- [1] Abien Fred Agarap. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018. 1, 6
- [2] Yash Bhalgat, Jinwon Lee, Markus Nagel, Tijmen Blankevoort, and Nojun Kwak. Lsq+: Improving low-bit quantization through learnable offsets and better initialization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 696–697, 2020. 3
- [3] Abhiroop Bhattacharjee, Yeshwanth Venkatesha, Abhishek Moitra, and Priyadarshini Panda. Mime: Adapting a single neural network for multi-task inference with memory-efficient dynamic pruning. *arXiv preprint arXiv:2204.05274*, 2022. 3
- [4] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020. 1, 3
- [5] Yuxuan Cai, Hongjia Li, Geng Yuan, Wei Niu, Yanyu Li, Xulong Tang, Bin Ren, and Yanzhi Wang. Yolobile: Real-time object detection on mobile devices via compression-compilation co-design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 955–963, 2021. 1, 3
- [6] Peng Ding, Huaming Qian, and Shuai Chu. Slimyolov4: lightweight object detector based on yolov4. *Journal of Real-Time Image Processing*, pages 1–12, 2022. 2
- [7] Xiaohan Ding, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11963–11975, 2022. 4
- [8] Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018. 4, 6
- [9] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021. 1, 2, 3, 4, 6
- [10] 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*, pages 7036–7045, 2019. 2
- [11] Raphael Gontijo-Lopes, Sylvia Smullin, Ekin Dogus Cubuk, and Ethan Dyer. Tradeoffs in data augmentation: An empirical study. In *International Conference on Learning Representations*, 2020. 2
- [12] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In *European conference on computer vision*, pages 544–560. Springer, 2020. 3
- [13] Hai Victor Habi, Roy H Jennings, and Arnon Netzer. Hmq: Hardware friendly mixed precision quantization block for cnns. In *European Conference on Computer Vision*, pages 448–463. Springer, 2020. 3
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 3
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [16] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 1, 4, 6
- [17] 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. 4
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 4
- [19] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, pages 734–750, 2018. 3
- [20] 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. *Advances in Neural Information Processing Systems*, 33:21002–21012, 2020. 3
- [21] 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*, pages 2117–2125, 2017. 2, 3, 4
- [22] 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*, pages 8759–8768, 2018. 2, 4, 5
- [23] 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*, pages 21–37. Springer, 2016. 1, 3, 4
- [24] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 4
- [25] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. 2, 4
- [26] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architec-

- ture design. In *Proceedings of the European conference on computer vision (ECCV)*, pages 116–131, 2018. 3
- [27] Yanwei Pang, Jiale Cao, Yazhao Li, Jin Xie, Hanqing Sun, and Jinfeng Gong. Tju-dhd: A diverse high-resolution dataset for object detection. *IEEE Transactions on Image Processing*, 30:207–219, 2020. 3
- [28] RangiLyu. Nanodet-plus: Super fast and high accuracy lightweight anchor-free object detection model. <https://github.com/RangiLyu/nanodet>, 2021. 7
- [29] 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*, pages 779–788, 2016. 1, 3, 4
- [30] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017. 1, 3
- [31] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 1, 2, 3, 6
- [32] 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, 2015. 1, 3
- [33] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 6
- [34] Mohammad Javad Shafiee, Brendan Chywl, Francis Li, and Alexander Wong. Fast yolo: A fast you only look once system for real-time embedded object detection in video. *arXiv preprint arXiv:1709.05943*, 2017. 1, 3
- [35] Zhuo Su, Wenzhe Liu, Zitong Yu, Dewen Hu, Qing Liao, Qi Tian, Matti Pietikäinen, and Li Liu. Pixel difference networks for efficient edge detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5117–5127, 2021. 5
- [36] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14454–14463, 2021. 1, 3
- [37] 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*, pages 10781–10790, 2020. 2
- [38] 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*, pages 9627–9636, 2019. 1, 3
- [39] Jin Xie, Yanwei Pang, Jing Nie, Jiale Cao, and Jungong Han. Latent feature pyramid network for object detection. *IEEE Transactions on Multimedia*, 2022. 2, 4
- [40] Jin Xu, Xu Tan, Kaitao Song, Renqian Luo, Yichong Leng, Tao Qin, Tie-Yan Liu, and Jian Li. Analyzing and mitigating interference in neural architecture search. In *International Conference on Machine Learning*, pages 24646–24662. PMLR, 2022. 3
- [41] Shangliang Xu, Xinxin Wang, Wenyu Lv, Qinyao Chang, Cheng Cui, Kaipeng Deng, Guanzhong Wang, Qingqing Dang, Shengyu Wei, Yuning Du, et al. Pp-yoloe: An evolved version of yolo. *arXiv preprint arXiv:2203.16250*, 2022. 1
- [42] Haichuan Yang, Yuhao Zhu, and Ji Liu. Ecc: Platform-independent energy-constrained deep neural network compression via a bilinear regression model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11206–11215, 2019. 1, 3, 4
- [43] Xinyi Ying, Yingqian Wang, Longguang Wang, Weidong Sheng, Li Liu, Zaipin Lin, and Shilin Zho. Mocopnet: Exploring local motion and contrast priors for infrared small target super-resolution. *arXiv preprint arXiv:2201.01014*, 2022. 5
- [44] Guanghua Yu, Qinyao Chang, Wenyu Lv, Chang Xu, Cheng Cui, Wei Ji, Qingqing Dang, Kaipeng Deng, Guanzhong Wang, Yuning Du, et al. Pp-picodet: A better real-time object detector on mobile devices. *arXiv preprint arXiv:2111.00902*, 2021. 1, 3, 6
- [45] Jiahui Yu and Thomas Huang. Autoslim: Towards one-shot architecture search for channel numbers. *arXiv preprint arXiv:1903.11728*, 2019. 3, 5
- [46] Haibo Zhang, Shulin Zhao, Ashutosh Pattnaik, Mahmut T Kandemir, Anand Sivasubramaniam, and Chita R Das. Distilling the essence of raw video to reduce memory usage and energy at edge devices. In *Proceedings of the 52nd Annual IEEE/ACM international symposium on microarchitecture*, pages 657–669, 2019. 3
- [47] 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*, pages 9759–9768, 2020. 3