iffDetector: Inference-aware Feature Filtering for Object Detection

Mingyuan Mao[1,†], Yuxin Tian[1,†], Baochang Zhang[1,*], Qixiang Ye[2], Wanquan Liu[3], Guodong Guo[4,5], and David Doermann[6]

[1]Beihang University, Beijing, China
[2]University of Chinese Academy of Sciences, Beijing, China
[3]Curtin University, Perth, Australia
[4]Institute of Deep Learning, Baidu Research, Beijing, China
[5]National Engineering Laboratory for Deep Learning Technology and Application
[6]University at Buffalo, Buffalo, USA
[*]Corresponding author, email: bczhang@buaa.edu.cn
[†] Co-first author

Abstract. Modern CNN-based object detectors focus on feature configuration during training but often ignore feature optimization during inference. In this paper, we propose a new feature optimization approach to enhance features and suppress background noise in both the training and inference stages. We introduce a generic Inference-aware Feature Filtering (IFF) module that can easily be combined with modern detectors, resulting in our iffDetector. Unlike conventional open-loop feature calculation approaches without feedback, the IFF module performs closed-loop optimization by leveraging high-level semantics to enhance the convolutional features. By applying Fourier transform analysis, we demonstrate that the IFF module acts as a negative feedback that theoretically guarantees the stability of feature learning. IFF can be fused with CNN-based object detectors in a plug-and-play manner with negligible computational cost overhead. Experiments on the PASCAL VOC and MS COCO datasets demonstrate that our iffDetector consistently outperforms state-of-the-art methods by significant margins[1].

Keywords: Object Detection, iffDetector, IFF, Negative Feedback

## 1 Introduction and Related Work

We have recently witnessed the success of visual object detection thanks to the unprecedented representation capacity of convolutional neural networks (CNNs) [6,5,23,20,22,14,18]. In [16], various taxonomies have been used to categorize the large number of CNN-based object detection methods, including one-stage [23] vs. two-stage [15], single-scale features [23] vs. feature pyramid networks [14], and handcrafted networks [18] vs. network architecture search [4].

---

[1] The test code and model are anonymously available in https://github.com/anonymous2020new/iffDetector

To explore the representation capability of CNNs, these detectors widely adopt a one-path network architecture which focuses on feature configuration during training [3,11,9,30,13,2,25,12] but unfortunately ignore feature optimization during inference. They take the feature maps extracted by the backbone as the input, apply various convolution operations, and acquire the final predictions at the end of the network [8,16,23,17,14]. For instance, the Feature Pyramid Network (FPN) [14] and top-down modulation [24] are designed to fuse multi-scale convolutional features, which improve the standard feature extraction pyramid by adding a second pyramid that takes the high level features from the first pyramid and passes them down to lower layers. This allows features at each level to have access to both higher and lower level features. The feature selective anchor-free (FSAF) [31] approach dynamically selects the most suitable level of feature for each instance based on the instance appearance, so as to address the scale variation problem. To strengthen the deep features learned from lightweight CNN models, RFB Net [17] has a novel module called the Receptive Field Block (RFB) which makes use of a multi-branch pooling with varying kernels corresponding to the receptive fields of different sizes, and reshapes them to generate enforced representations.

To facilitate object localization, many CNN-based detectors leverage anchor boxes at multiple scales and aspect ratios as reference points to address optimal feature-object matching [27,29,26,28]. The MetaAnchor [27] approach learns to optimize anchors from arbitrary customized prior boxes with a sub-network. GuidedAnchoring [26] leverages semantic features to guide the prediction of anchors while replacing dense anchors with predicted anchors. The state-of-the-art FreeAnchor [30] optimizes the anchor matching process by formulating a maximum likelihood estimation (MLE) procedure, which infers the most representative anchor from a "bag" of anchors for each object. Considering that anchor boxes introduce many hyper-parameters and design choices that make the network hard to train, anchor-free frameworks were recently proposed for object detection [13,2,25,12]. CornerNet [13] detects an object as a pair of key points on the top-left and bottom-right corners of a bounding box without anchors, which greatly simplifies the output of the network and eliminates the requirement for designing anchor boxes. Like CornerNet, CenterNet [2] and FoveaBox [12] were introduced with more efficiency for object detection.

Unfortunately, CNN-based detectors are still challenged by complex scenes where the targets and background become ambiguous. This may be caused by the feature representation that is not optimized according to the inference. To the best of our knowledge, there are no systematic approaches for feature optimization during inference, which inhibits the consistent enhancement of detectors. An inference module in [1] was introduced to efficiently distinguish the target objects from the background by adding a cascade structure at inference. FreeAnchor [30] leverages spatial inference to enforce feature representation and feature-object matching. These inference modules however were not investigated to optimize the feature representation in a theoretical framework in these detectors.
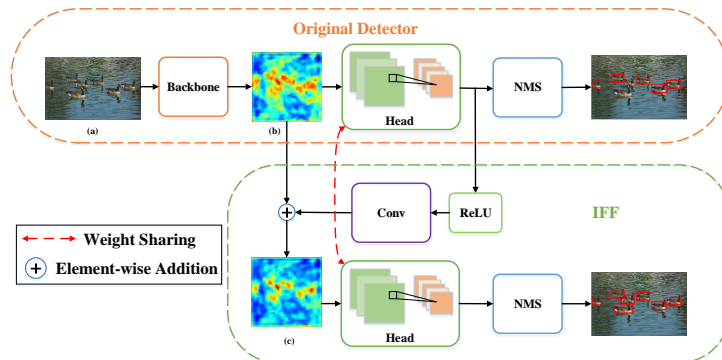
Fig. 1. Feature map denoising: (a) input image; (b) original feature map extracted by the backbone; (c) feature map optimized by IFF. (b) and (c) have more than one channel, and we visualize them by upsampling the feature maps to 300×300, adding all channels together and converting them to heatmaps, where the red regions represent stronger activation.

In this paper, we propose an Inference-aware Feature Filtering (IFF) approach, which optimizes feature learning in a theoretical framework by introducing a feedback architecture in detection networks. IFF differs from the existing inference frameworks in that it incorporates the conventional inference module into our filtering module to refine the feature map, instead doing it just for the final detection. Prior research shows that activated regions on a feature map are more likely to contain targets. In reality, however, the feature map extracted by the backbone may not be accurate. On the original feature map, some parts of the background are falsely activated (see Fig. 1(b)), which may lead to detection errors. By applying IFF to the original feature map, we can enhance the correctly activated regions, making them more intense, while at the same time depressing the background (see Fig. 1(c)). This results in more accurate object predictions. The IFF procedure, as shown in Fig. 1 and detailed in Sections 2 and 3, can be performed repeatedly.

Approximating a negative feedback system in the detection network, we prove the stability of the IFF-based object detection system. This leads to a new and generic inference model based on light-weight convolutional layers to improve the detection performance for both anchor-based and anchor-free detectors. By evaluating various CNN-based object detectors including SSD [18], YOLO [21], FreeAnchor [30], FoveaBox [12], and CenterNet [2], on PASCAL VOC2007+2012 and COCO 2017, we demonstrate the general applicability and effectiveness of IFF. The contributions of this paper are summarized as follows:

1. We present a generic Inference-aware Feature Filtering (IFF) module, which can be easily mounted on both anchor-based and anchor-free detectors, to optimize CNN features during inference.

2. We provide an explainable investigation into the proposed IFF, which theoretically guarantees the feature learning stability of the CNN-based detectors.
3. We achieve a significant performance gain over various existing detectors (including the state-of-the-art) on average precision (AP) on the COCO 2017 data set, with a negligible increase in computational cost.
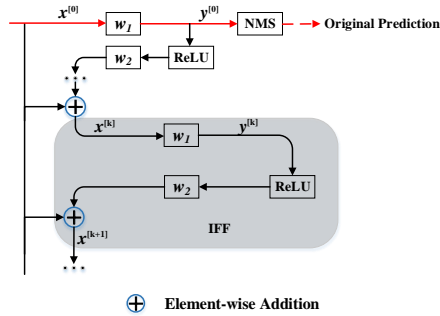
Fig. 2. Inference-aware feature filtering module: w_1 denotes the convolutional filters in the head of the original detector; w_2 denotes the convolutional filters we add to the system; $x^{[k]}$ and $y^{[k]}$ represent the feature map and the output in the k^{th} forward propagation respectively; the addition operation introduces a feedback from the detection output into the feature learning to achieve a negative feedback (control theory [10]).

2 Inference-aware Feature Filtering

In this section, we first detail the mechanism of the proposed filtering module and its advantages, then provide the framework for integrating IFF into a detector and explain how it works.

2.1 The Filtering Module

Our IFF is designed to highlight the features that can actively respond to the detection output (location and classification) and it provides a new and generic inference method for object detection. Our filtering module introduces a feedback link from the detection output to filter undesirable features. Suppose that for each input feature map $x^{[0]}$, there exists an ideal feature map n that can lead to the ideal detection results, and additive noise δ, such that

$$x^{[0]} = n + \delta. \tag{1}$$

To mitigate the noise effect on the features (Fig. 2), we incorporate the output of the detection into our inference-aware feature filtering (IFF) module and have

$$x^{[k+1]} = x^{[0]} + w_2 \otimes ReLU(y^{[k]}), \tag{2}$$

and

$$y^{[k]} = w_1 \otimes x^{[k]}, \tag{3}$$

where $x^{[k]}$ and $y^{[k]}$ represent the feature map and the output in the k^{th} forward propagation; w_1 denotes the forward convolution filters, w_2 are the new convolutional filters added, and $\otimes$ denotes the convolution operation. All the ReLU functions in this paper are Leaky ReLU.

Existing detectors without feedback only implement the forward propagation once at inference, which is the red path in Fig. 2. The corresponding output[2] is

$$y^{[0]} = w_1 \otimes x^{[0]}. \tag{4}$$

Note that we not only regard $y^{[0]}$ as heuristic information in IFF, but also use it for the final prediction. Thus, we introduce negative feedback architectures and obtain $\{(x^{[0]}, y^{[0]}), (x^{[1]}, y^{[1]}), \cdots (x^{[k]}, y^{[k]}), \cdots (x^{[M_I]}, y^{[M_I]})\}$ based on Fig. 2. The maximum value M_I is flexible and can be adjusted as necessary. $y^{[M_I]}$ is used to make the final prediction.

We explore the relationship between $x^{[k]}$ and $x^{[k+1]}$ for feature filtering, and show that the feature capacity is visually enhanced by our generic module in Fig. 2. In addition, the efficacy is theoretically validated based on the feedback mechanism in control theory, where the addition operation introduces feedback from the detection output into the feature learning to achieve a negative feedback (control theory [10]). The detailed proof is provided in Section 3, verifying that such a feedback system is stable in terms of filtering the noise δ. Briefly, the noise effect on $x^{[k+1]}$ is reduced in comparison to the effect on $x^{[k]}$. The framework of integrating IFF into a detector is provided in Section 2.2. This feedback structure has two main advantages, denoising and exploring high level semantics:

1. **Denoising.** Input images often have large variations of lighting and object occlusion, which leads to a "polluted" feature map generated by the CNN backbone. As a result, a prediction based on such a feature map can be somewhat ambiguous. Introducing a feedback structure to the CNN updates the open-loop feature calculation to a closed-loop feature optimization. Based on control theory, such an architecture makes the feature map more robust to noise. This negative feedback reuses the heuristic information in the original prediction, denoises the feature map, and makes the correctly activated regions on it more intensive. More details can be found in Section 3.2.
2. **Exploring high-level semantics.** It is rational to infer that classification and bounding boxes have a coupling relationship. For example, the bounding boxes for cars and people are generally different, in particular with respect to their aspect ratios. To this end, the feedback structure helps merge the location and classification, promoting a detector to learn the coupling relationships.

---

[2] $Prediction = NMS(y^{[0]})$, where NMS is short for Non-Maximum Suppression.

2.2   Our iffDetector

The IFF module is based on a set of light-weight convolution operations that can lead to enhanced features. To clarify how IFF works, we update YOLOv2 [21] to iffDetector (see Fig. 3). IFF introduces a unique feedback strategy for the inference, forming a cascade architecture. Existing detectors without feedback use the one-path architecture at inference, which is the red path in Fig. 3. For IFF an extra convolution layer w_2 is added to merge the prediction and the feature map. The original object scores are not regarded as the final predictions.
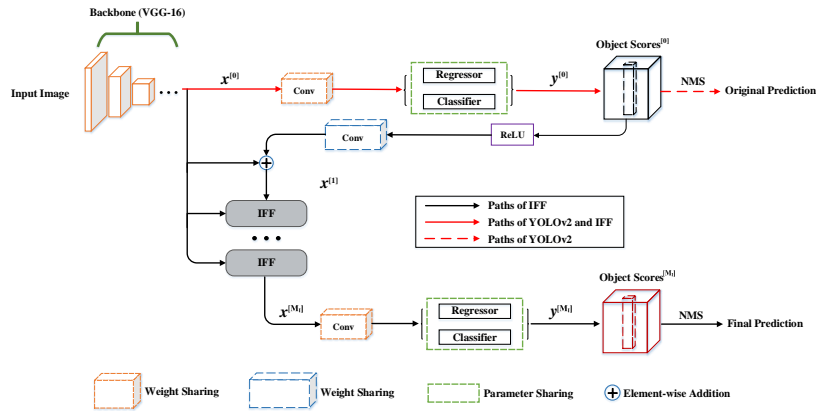
Fig. 3. The architecture of iffDetector: $*^{[k]}$ represents the feature in the k^{th} filtering module, and each IFF module is shown in Fig. 2; M_I is the number of the filtering modules we use; $x^{[0]}$ is the original feature map. Note that the head of YOLOv2 [21] only has one convolution layer.

Instead we apply the convolution to the object scores, and then add the convolution results to the original feature map $x^{[0]}$, generating a new feature map with better features. As mentioned in Section 2.1, this process can be implemented several times, which means the depth of the cascade architecture is flexible. The final feature map we acquire generates the final prediction. Unlike YOLOv2 [21], most existing detectors predict bounding boxes at different scales, which can also be updated to iffDetectors by simply adding the filtering modules independently across every scale similar to Fig. 3. M_I is the number of the filtering modules used. We then apply NMS on $y^{[M_I]}$ to obtain the final prediction. Note that a larger M_I does not always guarantee a better result because of training inefficiency. Our experiments show that setting M_I to 1 or 2 is appropriate.

Generalizability of Inference-aware Feature Filtering. The filtering module can be applied to various CNN-based detectors including anchor-based and anchor-free frameworks. Without loss of generality our filtering module is simple and effective. The size of the convolution filters we design is flexible. We

---

Algorithm 1 Training iffDetector based on IFF

Input:
X: Input images.
G: Labels of the training data.
M_I: The number of the filtering modules.
Output: θ (e.g., w_1, w_2); Detection parameters.
$\theta \leftarrow$ initialize network parameters.
Obtain original feature map $x^{[0]}$.
repeat
    for $k = 0 : M_I - 1$ **do**
        Forward propagation: Obtain $y^{[k]}$: $y^{[k]} \leftarrow w_1 \otimes x^{[k]}$.
        Denoise the feature map: $x^{[k+1]} \leftarrow x^{[0]} + w_2 \otimes ReLU(y^{[k]})]$.
    end for
    Calculate the loss: $L(\theta; G)$.
    Backward propagation: $\theta \leftarrow \theta - \bigtriangledown L(\theta; G)$
until The convergence of the loss
Return θ

---

use 1×1 filters for the backbones darknet19 and VGG-16 and 3×3 filters for ResNet-50-FPN and ResNeXt-64x4d-101-FPN, with less than a 2% increase in the number of parameters.

Based on the above description of IFF, we develop a universal detector training procedure as described in Algorithm 1.

3   Theoretical Investigation

To analyze our iffDetector, we first apply the Fourier transform on the inference module. We then explain how our module affects feature learning from a theoretical point of view.

3.1   Properties of the Detector Based on the Fourier Transform

To discuss the properties of the detector we apply the Fourier transform on Eqs. 2 and 3, and have

$$Y^{[k]} = W_1 * X^{[k]}, \tag{5}$$

which represents the linear case for the Fourier transform, and

$$\begin{aligned} X^{[k+1]} &= X^{[0]} + W_2 * \mathcal{F}(h(y^{[k]})) \\ &= X^{[0]} + W_2 * H(Y^{[k]}), \end{aligned} \tag{6}$$

where $*$ is the element-wise multiplication operator, $h(x)$ is the ReLU function, $\mathcal{F}$ denotes the Fourier transform, and each capitalized letter denotes the Fourier transform of the corresponding variable.

Eq. 6 represents the nonlinear case when using Fourier transform to analyze the detector. This remains a challenge in terms of interpretability for deep

learning, because of the nonlinear ReLU function. Theorem 1 shows that the Leaky ReLU function can reduce the energy of a signal's Fourier transform and at the same time can be further used to prove the feature learning stability of IFF (Section 3.2).

Theorem 1. *For any input feature map x, it satisfies*

$$\|\mathcal{F}(h(x))\| \leq \|\mathcal{F}(x)\|, \tag{7}$$

where $\| \cdot \|$ denotes the Frobenius norm of a matrix.

Proof. The Leaky ReLU is defined as

$$h(x) = \begin{cases} x, & x > 0, \\ \lambda x, & x \leq 0, \end{cases}$$

where $\lambda \in (0, 1)$. If x is a matrix of size $W \times H$, we use x_i to denote its elements, where $i \in [0, WH - 1]$. According to the definition of the Leaky ReLU function, it is obvious that

$$\sum_{i=0}^{WH-1} |h(x_i)|^2 \leq \sum_{i=0}^{WH-1} |x_i|^2. \tag{8}$$

From Lemma 1 in the appendix, we can obtain

$$\sum_{i=0}^{WH-1} |h(x_i)|^2 = \frac{1}{WH} \sum_{k=0}^{WH-1} |H(X_k)|^2, \tag{9}$$

and

$$\sum_{i=0}^{WH-1} |x_i|^2 = \frac{1}{WH} \sum_{k=0}^{WH-1} |X_k|^2, \tag{10}$$

where X_k denotes the element of X, which is the Fourier transform of x. Then we have

$$\|\mathcal{F}(h(x))\| \leq \|\mathcal{F}(x)\|, \tag{11}$$

which proves Theorem 1.

3.2   Feature Learning based on Feedback

In this section, we prove the robustness of iffDetector based on the Fourier transform. As mentioned previously, we assume there exists an ideal feature map n for an input image, with N as its corresponding Fourier transform. However, for any input image, we often have a "polluted" feature map $X^{[0]}$ generated by the CNN backbone defined as

$$X^{[0]} = N + \Delta, \tag{12}$$

where Δ denotes the Fourier transform of the noise δ. Correspondingly, we also have an ideal output y_n defined as

$$
\begin{aligned}
y_n &= w_1 \otimes n, \\
Y_N &= W_1 * N.
\end{aligned}
\tag{13}
$$

We define the optimization objective function as

$$
V(Y) = \|Y - Y_N\|^2,
\tag{14}
$$

where $\|\cdot\|$ denotes the Frobenius norm, V is a positive definite function and $V(Y) = 0$ only when $Y = Y_N$.

Eq. 14 represents the deviation between the actual output Y and the ideal output Y_N. We develop Theorem 2 to show that as the distance between $X^{[k]}$ and N increases, the distance between $Y^{[k+1]}$ and Y_N is reduced by IFF, which guarantees the stability of our iffDetector.

Theorem 2. *In iffDetector, there always exists $\epsilon > 0$, such that when $\|X^{[k]} - N\| > \epsilon$, it follows that*

$$
V(Y^{[k+1]}) \le V(Y^{[k]}),
\tag{15}
$$

where $V(\cdot)$ is defined in Eq. 14.

Proof. For simplicity, we first define

$$
\begin{aligned}
V'(Y^{[k]}) &= V(Y^{[k+1]}) - V(Y^{[k]}) \\
&= \|Y^{[k+1]} - Y_N\|^2 - \|Y^{[k]} - Y_N\|^2.
\end{aligned}
\tag{16}
$$

Plugging Eqs. 5, 6 and 13 into $\|Y^{[k+1]} - Y_N\|^2$ results in

$$
\begin{aligned}
V'(Y^{[k]}) &= \|W_1 * (X^{[0]} - N) + W_1 * W_2 * H(W_1 * X^{[k]})\|^2 \\
&\quad - \|Y^{[k]} - Y_N\|^2.
\end{aligned}
\tag{17}
$$

Besides, it follows that

$$
\begin{aligned}
-\|Y^{[k]} - Y_N\|^2 &\le -(\|Y^{[k]}\| - \|Y_N\|)^2 \\
&= -(\|Y^{[k]}\|^2 - 2\|Y^{[k]}\|\|Y_N\| + \|Y_N\|^2).
\end{aligned}
\tag{18}
$$

Based on Lemma 2, Eqs. 12, 13, 18 and Theorem 1, we have

$$
\begin{aligned}
V'(Y^{[k]}) &\le \|W_1\|^2\|\Delta\|^2 \quad + 2\|W_1\|^2\|\Delta\|\|W_2\|\|H(W_1 * X^{[k]})\| \\
&\quad + \|W_1\|^2\|W_2\|^2\|H(W_1 * X^{[k]})\|^2 - \|W_1\|^2\|X^{[k]}\|^2 \\
&\quad + 2\|W_1\|^2\|X^{[k]}\|\|N\| - \|W_1\|^2\|N\|^2 \\
&\le \|W_1\|^2\|\Delta\|^2 + 2\|W_1\|^2\|\Delta\|\|W_2\|\|W_1\|\|X^{[k]}\| \\
&\quad + \|W_1\|^4\|W_2\|^2\|X^{[k]}\|^2 - \|W_1\|^2\|X^{[k]}\|^2 \\
&\quad + 2\|W_1\|^2\|X^{[k]}\|\|N\| - \|W_1\|^2\|N\|^2 \\
&= \|W_1\|^2(\|W_1\|^2\|W_2\|^2 - 1)\|X^{[k]}\|^2 \\
&\quad + 2(\|W_1\|^2\|N\| + \|W_1\|^3\|\Delta\|\|W_2\|)\|X^{[k]}\| \\
&\quad + \|W_1\|^2(\|\Delta\|^2 - \|N\|^2).
\end{aligned}
\tag{19}
$$

We note that W_1 and W_2 are 3D filters, however, the convolution operation in CNNs is actually the 2D convolution. We can only choose the 2D matrix of the maximum norm and then directly use Theorem 1. By setting

$$\begin{cases} A = \|W_1\|^2(\|W_1\|^2\|W_2\|^2 - 1\|), \\ B = 2(\|W_1\|^2\|N\| + \|W_1\|^3\|\Delta\|\|W_2\|), \\ C = \|W_1\|^2(\|\Delta\|^2 - \|N\|^2), \end{cases} \tag{20}$$

we obtain

$$V^{'}(Y^{[k]}) \leq A\|X^{[k]}\|^2 + B\|X^{[k]}\| + C, \tag{21}$$

where A, B and C are constants[3]. By inserting Eq. 12 into Eq. 6, we have

$$\begin{aligned} X^{[k]} &= N + [\Delta + W_2 * H(Y^{[k-1]})] \\ &= N + \Delta^{[k]}, \end{aligned} \tag{22}$$

where $\Delta^{[k]}$ is a variable and equals $X^{[k]} - N$. Inserting Eq. 22 into Eq.21, we obtain

$$V^{'}(Y^{[k]}) \leq A\|N + \Delta^{[k]}\|^2 + B\|N + \Delta^{[k]}\| + C, \tag{23}$$

where $A \leq 0$ is easily satisfied because $\|W_1\|^2\|W_2\|^2 < 1$ can be achieved by controlling by W_2. According to the property of a quadratic function, there always exists ϵ, such that when $\|\Delta^{[k]}\| = \|X^{[k]} - N\| > \epsilon$, we have $A\|X^{[k]}\|^2 + B\|X^{[k]}\| + C < 0$. That is to say, $V(Y^{[k+1]}) \leq V(Y^{[k]})$, which proves Theorem 2.

Theorem 2 indicates that once the deviation of $X^{[k]}$ from N is too large, $Y^{[k+1]}$ will move back to Y_N (the ideal output), which proves the stability of the filtering module. Eq. 14 is a classical Lyapunov function [19], which links the CNN-based detectors to the control theory that explains the feedback strategy can bring robustness to our iffDetector.

4   Experiments

In this section, we present the implementation of IFF for detectors to evaluate the effect of the proposed filtering module. We also compare iffDetector with the counterpart and the state-of-the-art approaches. Experiments are carried out on several datasets, including PASCAL VOC2007, PASCAL VOC2012, and MSCOCO 2017.

4.1   Implementation Details

IFF is implemented on FreeAnchor [30] and CenterNet [2], which are the state-of-the-art one-stage anchor-based and anchor-free detectors, respectively. We

---

[3] For a certain input image and forward propagation, the original feature map $x^{[0]}$, the ideal feature map N, filters W_1 and W_2 are all constants.

also test the filtering module on other mainstream detectors including SSD [18], YOLOv2 [21], and FoveaBox [12]. Since we only modify the head of the detector, the pretrained weights of the backbone can be used which saves training time. The training and test times of iffDetectors are almost the same as their baselines. To select the best number of filtering modules (M_I), we test our iffDetector based on FreeAnchor with different M_I. To demonstrate the generality of IFF, we carry out experiments on three datasets: COCO 2017, and PASCAL VOC2007 and 2012. We use 8 Tesla V100 GPUs for the experiments on FreeAnchor (with ResNeXt-101 as the backbone) and 4 TITAN Xp GPUs for other experiments. The training schedule and batch size are adjusted according to the linear scaling rule [7].

4.2   The Choice of M_I

As we mentioned in Section 2, the number of filtering modules is flexible. To select an appropriate value of M_I, we test iffDetector using FreeAnchor (with ResNet-50 as the backbone) on COCO 2017 validation set with different M_I, and the results are shown in Table 1. We find that having more filtering modules does not guarantee better performance. Furthermore, having more filtering modules means more training and inference time. Thus, we set M_I to 1 in all our experiments. Note that $M_I = 0$ corresponds to the original detector without feedback.

Table 1. iffDetector performance with different M_I.

| M_I	0	1	2	3
mAP	38.5	39.3	39.2	39.2

4.3   Model Effect

Feature maps generated by the backbones may have noise such that some regions without object are falsely activated. We use SSD with or without IFF trained with VOC2007 to illustrate the advantages of effect of the proposed module. Our filtering module at inference reuses the heuristic information in the original predictions, successfully denoising the feature maps and making the correctly activated regions more intensive, thus leading to more accurate predictions (see Fig. 4). When multiple objects of the same class appear very close to each other, the inaccurate activated regions on feature maps usually lead to missing detections. We have 6 feature maps for prediction, and choose the corresponding kernels of the maximum norm. We calculate A in Eq. 23 and obtain $A = -0.361, -0.301, -0.226, -0.185, -0.201, -0.149$ in a descent order, which

are consistent with Theorem 2. Furthermore, iffDetector achieves a better performance with almost the same inference cost, *e.g.*, SSD operates at 45.45 frames per second (FPS), while iffDetector achieves 44.62 FPS.
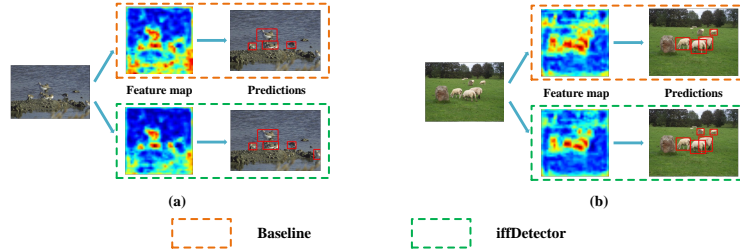
Fig. 4. Visualization of features w/o IFF: (a) and (b) are the detection results of two input images. To achieve more detailed and direct illustration, we upsample the feature maps to 300×300, add all the channels, and convert them to heatmaps, where the red regions represent stronger activations.
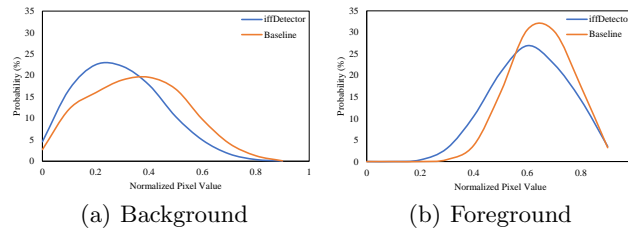
Fig. 5. The energy distributions on the feature maps. The baseline is SSD [18].

IFF is designed to filter undesired features enforcing the foreground while depressing the background. To verify this, we show the distribution of pixel values in the background of the feature map, which can be regarded as the energy distribution (see Fig. 5). To do this, we randomly select 60 images from VOC2007 to obtain two groups of feature maps (original feature maps and feature maps optimized by IFF). For every feature map, we add all the channels together and normalize the pixel values. We then compute the distribution of the pixel values of each group, resulting in two curves corresponding to the two groups. Fig. 5(a) shows that the energy of the background obtained by iffDetector is clearly lower than that of the original detector which means IFF successfully depresses the background. Meanwhile, considering the strongly activated regions of the foreground obtained by iffDetector is more concentrated (see Fig. 4), the energy of the foreground obtained by iffDetector is also a little bit lower than

that of the original detector, which is beneficial to object detection, because in most cases the center points of the objects generate the best bounding boxes.

Table 2. Detection performance comparison of our iffDetectors and the baselines on PASCAL VOC.

Backbone	Detector	Training	mAP
Darknet-19	YOLOv2 416×416	VOC07	59.3
	iffDetector		**60.4**
VGG-16	SSD300	VOC07+12	77.1
	iffDetector		**78.0**

Table 3. Detection performance comparison of our iffDetectors and the baselines on COCO 2017.

Backbone	Detector	Parameter Size (M)	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
ResNet-50	FoveaBox	36.5	36.7	56.6	39.2	20.3	39.9	45.3
	iffDetector	36.6 ($\uparrow 0.5\%$)	**37.6**	**57.4**	**40.4**	**21.2**	**40.8**	**46.7**
ResNet-50	FreeAnchor	33.8	38.7	57.3	41.6	20.2	41.3	50.1
	iffDetector	34.1 ($\uparrow 0.8\%$)	**39.6**	**58.5**	**42.7**	**21.3**	**42.2**	**50.6**
Hourglass-104	CenterNet	191.3	45.1	63.9	49.3	26.6	47.1	57.7
	iffDetector	191.7 ($\uparrow 0.2\%$)	**45.8**	**64.8**	**50.0**	**27.7**	**47.8**	**58.6**
ResNeXt-101	FreeAnchor	91.5	47.3	66.3	51.5	30.6	50.4	59.0
	iffDetector	93.2 ($\uparrow 1.8\%$)	**47.9**	**66.7**	**52.2**	**30.7**	**50.7**	**59.9**

4.4   Detection Performance

We compare the proposed iffDetector with multiple baselines, including anchor-based approaches, YOLOv2 [21], SSD [18], and Freeanchor [30], and anchor-free approaches, Foveabox [12] and CenterNet [2]. More experimental results are also provided in the supplementary material, including training curves, result analysis and other implementation details.

For YOLOv2 and SSD, we upgrade each with our IFF by setting the size of the filters (w_2 in Fig. 2) to 1×1 and test if the iffDetectors outperform the baselines. As shown in Table 2, the IFF module works very well on the PASCAL VOC datasets. For other baselines, considering their backbones are large, we set the size of the filters to 3×3 and test the iffDetectors on COCO 2017 (see Table 3). Again, IFF consistently improves the baselines on all evaluation metrics with less than a 2% increase in the number of parameters. Note that

all the performance gains are achieved with negligible extra cost of training and inference time. More importantly, compared to the most recent FreeAnchor (on ResNeXt-101), iffDetector respectively improves the AP and AP_L up to 0.6% and 0.9%, which are significant margins in terms of the challenging object detection task.

iffDetector is also compared with other state-of-the-art detectors in Table 4 using standard settings, ResNet as the backbone. Note that iffDetector* is acquired by updating the corresponding FreeAnchor detectors. Besides, Advanced settings use the jitter over scales of $\{640, 672, 704, 736, 768, 800\}$ during training on the ResNeXt-32x8d-101 backbone. The experiments show that iffDetector outperform both the anchor-based counterparts (RetinaNet [15] and FreeAnchor [30]) and the anchor-free approaches (FoveaBox [12], CornerNet [13] and CenterNet [2]). With less than a 2% increase in the number of parameters and almost the same training and inference time, IFF upgrades the state-of-the-art detectors with significant margins.

Table 4. Performance comparison with state-of-the-art detectors.

Detector	Backbone	Iter.	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
RetinaNet [15]	ResNet-50	70k	35.7	55.0	38.5	19.9	38.9	46.3
FoveaBox [12]	ResNet-50	70k	36.7	56.6	39.2	20.3	39.9	45.3
FreeAnchor [30]	ResNet-50	90k	38.7	57.3	41.6	20.2	41.3	50.1
iffDetector*(ours)	ResNet-50	90k	**39.6**	**58.5**	**42.7**	**21.3**	**42.2**	**50.6**
CornerNet [13]	Hourglass-104	500k	40.6	56.4	43.2	19.1	42.8	54.3
RetinaNet [15]	ResNeXt-101	135k	40.8	61.1	44.1	24.1	44.2	51.2
FoveaBox [12]	ResNeXt-101	135k	42.1	61.9	45.2	24.9	46.8	55.6
FSAF [31]	ResNeXt-101	180k	44.6	65.2	48.6	29.7	47.1	54.6
Cascade R-CNN [1]	ResNeXt-101	180k	44.9	63.7	48.9	25.9	47.7	57.1
CenterNet [2]	Hourglass-104	480k	45.1	63.9	49.3	26.6	47.1	57.7
FreeAnchor [30]	ResNeXt-101	180k	47.3	66.3	51.5	30.6	50.4	59.0
iffDetector*(ours)	ResNeXt-101	180k	**47.9**	**66.7**	**52.2**	**30.7**	**50.7**	**59.9**

5   Conclusion

We have proposed an elegant and effective approach referred to as Inference-aware Feature Filtering (IFF), for visual object detection. The proposed IFF upgrades both anchor-based and anchor-free frameworks to new detectors with a filtering ability by introducing a feedback architecture to the inference procedure and achieves more accurate and robust predictions. With IFF implemented, we have significantly improved the performance of object detection over baseline detectors. The underlying effect is that IFF denoises the feature maps and facilitates generating new feature maps with more precise and intensive activated regions, which lead to improved detection results. Our IFF approach provides a fresh insight to the visual object detection problem.

References

1. Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. In: IEEE CVPR. pp. 6154–6162 (2018)
2. Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., Tian, Q.: Centernet: Object detection with keypoint triplets. In: IEEE CVPR (2019)
3. Fu, C.Y., Liu, W., Ranga, A., Tyagi, A., Berg, A.C.: DSSD: Deconvolutional single shot detector. In: arXiv preprint arXiv:1701.06659
4. Ghiasi, G., Lin, T., Le, Q.V.: NAS-FPN: learning scalable feature pyramid architecture for object detection. In: IEEE CVPR. pp. 7036–7045 (2019)
5. Girshick, R.B.: Fast R-CNN. In: IEEE ICCV. pp. 1440–1448 (2015)
6. Girshick, R.B., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: IEEE CVPR. pp. 580–587 (2014)
7. Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch sgd: training imagenet in 1 hour. arXiv: 1706.02677 (2017)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE CVPR. pp. 770–778 (2016)
9. He, Y., Zhu, C., Wang, J., Savvides, M., Zhang, X.: Bounding box regression with uncertainty for accurate object detection. arXiv:1809.08545 (2018)
10. Horowitz, I.: Fundamental theory of automatic linear feedback control systems. IRE Transactions on Automatic Control **4**(3), 5–19 (1959)
11. Jiang, B., Luo, R., Mao, J., Xiao, T., Jiang, Y.: Acquisition of localization confidence for accurate object detection. In: ECCV. pp. 784–799 (2018)
12. Kong, T., Sun, F., Liu, H., Jiang, Y., Shi, J.: Foveabox: Beyond anchor-based object detector. arXiv:1904.03797 (2019)
13. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: ECCV. pp. 765–781 (2018)
14. Lin, T., Dollár, P., Girshick, R.B., He, K., Hariharan, B., Belongie, S.J.: Feature pyramid networks for object detection. In: IEEE CVPR. pp. 936–944 (2017)
15. Lin, T., Goyal, P., Girshick, R.B., He, K., Dollár, P.: Focal loss for dense object detection. In: IEEE ICCV. pp. 2999–3007 (2017)
16. Liu, L., Ouyang, W., Wang, X., Fieguth, P., Chen, J., Liu, X., Pietikainen, M.: Deep learning for generic object detection: A survey. Int. J. Comp. Vis. (2019)
17. Liu, S., Huang, D., Wang, Y.: Receptive field block net for accurate and fast object detection. In: ECCV. pp. 385–400 (2018)
18. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S.E., Fu, C., Berg, A.C.: SSD: single shot multibox detector. In: ECCV. pp. 21–37 (2016)
19. Lyapunov, A.M., Walker, J.A.: The general problem of the stability of motion. Journal of Applied Mechanics **61**(1), 226 (1892)
20. Redmon, J., Divvala, S.K., Girshick, R.B., Farhadi, A.: You only look once: Unified, real-time object detection. In: IEEE CVPR. pp. 779–788 (2016)
21. Redmon, J., Farhadi, A.: YOLO9000: better, faster, stronger. In: IEEE CVPR. pp. 6517–6525 (2017)
22. Redmon, J., Farhadi, A.: Yolov3: An incremental improvement. arXiv: 1804.02767 (2018)
23. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: NeurIPS. pp. 91–99 (2015)

24. Shrivastava, A., Sukthankar, R., Malik, J., Gupta, A.: Beyond skip connections: Top-down modulation for object detection. arXiv:1612.06851 (2016)
25. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. arXiv:1904.01355 (2019)
26. Wang, J., Chen, K., Yang, S., Loy, C.C., Lin, D.: Region proposal by guided anchoring. In: IEEE CVPR. pp. 2965–2974 (2019)
27. Yang, T., Zhang, X., Li, Z., Zhang, W., Sun, J.: Metaanchor: Learning to detect objects with customized anchors. In: NeurIPS. pp. 320–330 (2018)
28. Yu, J., Jiang, Y., Wang, Z., Cao, Z., Huang, T.S.: Unitbox: An advanced object detection network. In: ACM Multimedia Conference. pp. 516–520 (2016)
29. Zhang, S., Wen, L., Bian, X., Lei, Z., Li, S.Z.: Single-shot refinement neural network for object detection. In: IEEE CVPR. pp. 4203–4212 (2018)
30. Zhang, X., Wan, F., Liu, C., Ji, R., Ye, Q.: FreeAnchor: Learning to match anchors for visual object detection. In: NeurIPS (2019)
31. Zhu, C., He, Y., Savvides, M.: Feature selective anchor-free module for single-shot object detection. In: IEEE CVPR. pp. 840–849 (2019)

Appendix

Lemma 1. *According to Parseval's Theorem, we have*

$$\sum_{n=0}^{N-1} |x_n|^2 = \frac{1}{N} \sum_{k=0}^{N-1} |X_k|^2, \tag{24}$$

where x_n is the signal in the spatial domain, and $X_k = \mathcal{F}(x_n)$.

Lemma 2. *Any complex matrices X and Y satisfy the following inequality*

$$\|X \pm Y\|^2 \leq \|X\|^2 + 2\|X\|\|Y\| + \|Y\|^2, \tag{25}$$

where $\|\cdot\|$ denotes any matrix norm.