

# HS-FPN: High Frequency and Spatial Perception FPN for Tiny Object Detection

Zican Shi<sup>1</sup>, Jing Hu<sup>1,2\*</sup>, Jie Ren<sup>1</sup>, Hengkang Ye<sup>1</sup>, Xuyang Yuan<sup>1</sup>, Yan Ouyang<sup>3</sup>, Jia He<sup>4</sup>, Bo Ji<sup>4</sup>, Junyu Guo<sup>4</sup>

<sup>1</sup>Huazhong University of Science and Technology, Wuhan, China

<sup>2</sup>National Key Laboratory of Science and Technology on Multi-Spectral Information Processing, Wuhan, China

<sup>3</sup>Air Force Early Warning Academy, Wuhan, China

<sup>4</sup>Chinese People's Liberation Army 95841 troops, China

{misaka10043, by6040130, jieren, yehengkang, lucasxyyuan}@hust.edu.cn

{ooy\_01, hugh120007, jibo\_2006, henry\_rocky\_1999}@163.com

## Abstract

The introduction of Feature Pyramid Network (FPN) has significantly improved object detection performance. However, substantial challenges remain in detecting tiny objects, as their features occupy only a very small proportion of the feature maps. Although FPN integrates multi-scale features, it does not directly enhance or enrich the features of tiny objects. Furthermore, FPN lacks spatial perception ability. To address these issues, we propose a novel High Frequency and Spatial Perception Feature Pyramid Network (HS-FPN) with two innovative modules. First, we designed a high frequency perception module (HFP) that generates high frequency responses through high pass filters. These high frequency responses are used as mask weights from both spatial and channel perspectives to enrich and highlight the features of tiny objects in the original feature maps. Second, we developed a spatial dependency perception module (SDP) to capture the spatial dependencies that FPN lacks. Our experiments demonstrate that detectors based on HS-FPN exhibit competitive advantages over state-of-the-art models on the AI-TOD dataset for tiny object detection.

## 1 Introduction

Tiny Object Detection (TOD), a subtask of general object detection, focuses on detecting tiny-sized objects that AI-TOD (Wang et al. 2021) defines as being less than  $16 \times 16$  pixels. TOD plays an important role in scenarios such as traffic sign detection, scene monitoring, unmanned aerial vehicle analysis, pedestrian detection, autonomous driving, and maritime rescue.

Despite the development of deep learning bringing forth many outstanding object detection frameworks that have achieved impressive performance on regular-sized objects, the performance of these frameworks drops significantly when it comes to detecting tiny objects. Based on recent reviews and surveys on tiny object detection (Cheng et al. 2023; Li et al. 2022; Tong and Wu 2022), we infer that the performance decline can be attributed to three factors: tiny objects have limited usable features, their features are easily interfered with, and general network architectures do not adequately focus on tiny objects. Moreover, in recent years,

large vision models (Zeng et al. 2023; Zhong et al. 2022) have mostly been studied only on the general COCO dataset (Lin et al. 2014), without giving extra attention to tiny objects. Due to these reasons, tiny object detection has gradually become a challenging research direction in the field of computer vision.

Feature Pyramid Network (FPN) (Lin et al. 2017) has already become an indispensable part of object detection models (Ren et al. 2015). FPN leverages multi-scale feature maps produced by the backbone network (typically ResNet (He et al. 2016)), propagating the strong semantics of deep feature maps to shallow feature maps, significantly enhancing the information content of the shallow features.

Although experiments have shown that FPN significantly improves object detection performance, there are still some challenges when it comes to detecting tiny objects, as described below:

**Tiny objects have limited usable features.** Although FPN integrates deep and shallow features, it does not directly increase the feature content for tiny objects. Due to the small pixel area of tiny objects, frequent downsampling by the backbone network continually compresses the feature size of tiny objects, resulting in only a few pixels of feature representation for tiny objects in the final feature map. Weak feature responses are not enough for precise detection and location.

**Tiny objects lack network attention.** The feature responses of tiny objects are relatively weak and susceptible to interference, requiring extra attention. However, FPN applies the same processing method to features at each layer:  $1 \times 1$  convolutions to reduce the number of channels, top-down features fusion and  $3 \times 3$  convolutions to integrate outputs, without providing special treatment for tiny objects features.

**FPN lacks spatial perception ability.** By merging upper and lower layer features in a pixel-by-pixel addition manner, FPN lacks spatial perception ability around tiny targets. The pixel offset caused by recursive upsampling in FPN results in misalignment of tiny target features between upper and lower layers. Therefore, it is necessary to enhance the spatial perception ability of FPN and enrich the contextual information of tiny target features.

In fact, many traditional methods for detecting tiny objects are based on frequency domain techniques (Achanta

\*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

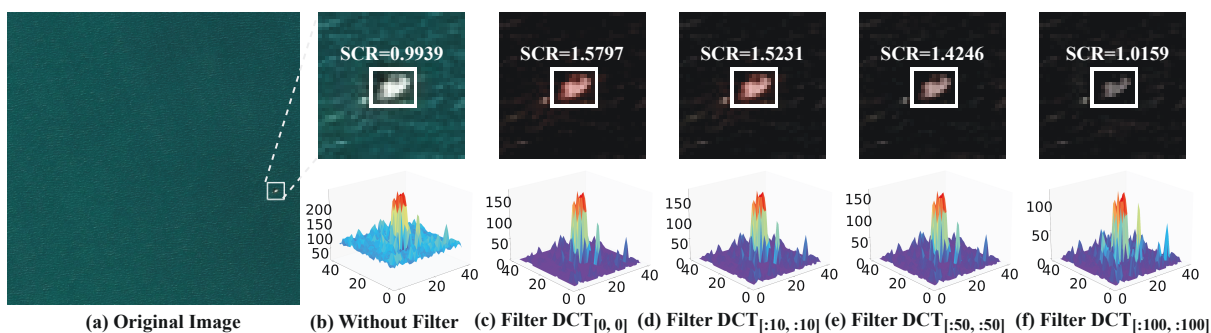


Figure 1: Target region after filtering out low-frequency components across different ranges (top), 3D surface representations (down), and the SCR of the target (marked in red text). (a) The original image containing a tiny ship target; (b) The target neighborhood image without filtering; (c - d) Filtering out low-frequency components in the top-left corner of the DCT results, from the  $0 \times 0$  to  $100 \times 100$  region. It shows that the SCR first increases and then decreases as the filtering range is expanded.

et al. 2009; Zhang et al. 2011). For example, wavelet transform can be used to separate the original image into low-frequency and high-frequency components, and then thresholding can be applied to the high-frequency components to detect tiny objects (Zhao, Liu, and Mo 2012). Because low-frequency components represent the overall contours and large, smooth areas of the image, while tiny objects typically appear as details and edges, which correspond to high-frequency components in the image’s frequency spectrum. By filtering out the low-frequency components, tiny objects can be highlighted. This can be achieved by using discrete cosine transform (DCT) and a high-pass filter. The effectiveness of this filtering can be quantified using the Signal to Clutter Ratio (SCR), which reflects the degree of discrimination between the target and its surrounding background in the local image region. A higher SCR value generally indicates a more salient target. The SCR is mathematically defined as:

$$SCR = \frac{|\mu_t - \mu_b|}{\sigma_b} \quad (1)$$

where  $\mu_t$  is the average gray value of the target region (a local image region of  $40 \times 40$  pixels centered around the target in this work), and  $\mu_b$  and  $\sigma_b$  are the average and standard deviation of the gray value in the target’s neighborhood, respectively. Figure 1 shows the target regions, their 3D surface representations, and the SCR values after filtering out low-frequency components across different ranges. The results indicate that moderate filtering of low-frequency components can significantly improve SCR, making tiny objects more prominent.

Inspired by the above, we firstly design a high frequency perception module (HFP) to enhance the features of tiny objects by filtering out low-frequency components in the FPN feature maps. HFP begins with a predefined high-pass filter to extract the high-frequency response of the input feature map. This high-frequency response is used to refine the original feature map through both channel and spatial branches. The channel path dynamically allocates weights to each channel of the original feature map based on the high-frequency response, highlighting channels with more tiny object features. The spatial path allocates weights to

each pixel in the original feature map based on the high-frequency response, generating a spatial mask that directs attention on the areas with tiny object features. Secondly, we design a spatial dependency perception module (SDP) operated between adjacent upper and lower feature maps. SDP calculates similarity at the pixel level and learns the spatial dependencies between adjacent pixel in the upper and lower feature maps, enriching tiny object features by leveraging valuable spatial dependencies.

Experiments show that using ResNet50 as the backbone, Faster R-CNN based on HS-FPN achieves an average precision (*AP*) of 20.3 on the AI-TOD dataset, compared to 18.3 with FPN. Similarly, Cascade R-CNN based on HS-FPN improves *AP* from 20.2 to 23.6. Furthermore, HS-FPN has a similar overall structure to FPN and can easily embedded into any model that requires FPN.

The main contributions of this paper are as follows:

- We reveal three issues faced by current FPN-based object detection models in detecting tiny objects.
- We propose a novel feature pyramid network, HS-FPN, which addresses these issues mentioned above through high frequency perception module and spatial dependency perception module.
- We integrate HS-FPN into various detection models to replace FPN, and evaluate its performance on multiple tiny object detection datasets. Experiments indicate that HS-FPN significantly improves performance compared to FPN.

## 2 Related Work

### 2.1 Tiny Object Detection

Most of the research on TOD focuses on three main aspects: benchmark construction, better label assignment strategy, and multi-scale learning.

**Benchmark construction.** Generally used datasets like COCO (Lin et al. 2014) lack sufficient tiny object instances. To address this, new datasets such as AI-TOD (Wang et al. 2021), SODA (Cheng et al. 2023), and AI-TOD-v2 (Xu




Figure 2: The overall structure of HS-FPN and the details of the lateral connection. HFP means high frequency perception module while SDP means spatial dependency perception module. For better visualization, convolution layers are omitted. Please note that SDP does not appear in  $P_5$ .

et al. 2022a) have been developed, including more tiny object samples and effective evaluation metrics, advancing detection techniques.

**Better label assignment strategy.** The Intersection over Union (IoU) metric, commonly used for matching predicted boxes with ground truth, is highly sensitive to location deviations in tiny objects, significantly compromising label assignment quality in anchor-based detectors. There have already been some improvements to the IoU metric, methods such as Dot Distance (Xu et al. 2021), NWD (Xu et al. 2022a), and RFLA (Xu et al. 2022b), have demonstrated effectiveness in TOD.

**Multi-scale learning.** As a classic multi-scale learning method, FPNs (Lin et al. 2017; Liu et al. 2018; Luo et al. 2022) effectively integrates high-level semantics from deeper feature maps with detailed information from shallower ones through a top-down structure and lateral connections, significantly improving tiny object detection performance.

## 2.2 Learning in Frequency Domain

Frequency domain analysis has long been a crucial tool in image processing. With the breakthroughs in deep learning, researchers have increasingly integrated frequency domain information into deep learning methods (Xu et al. 2020). For instance, Rao et al. (2023) used 2D discrete Fourier transform and a set of learnable frequency filters to enable frequency domain interactions between tokens, learning long-term spatial dependencies in the frequency domain with logarithmic complexity. Qin et al. (2021) mathematically demonstrated that the commonly used Global Average Pooling is equivalent to the lowest frequency component in the DCT and assigned weights to each channel of the feature map

based on DCT. While these studies were not all specifically focused on tiny object detection, they have inspired our research.

## 3 Method

In this section, we detail the HS-FPN framework and its components. As shown in Figure 2, the structure of HS-FPN is similar to FPN, collecting four feature maps from the backbone network and reducing their channel dimensions to 256 using  $1 \times 1$  convolutions. These reduced-channel features are named as  $\{C_2, C_3, C_4, C_5\}$  which have strides of  $\{4, 8, 16, 32\}$  pixels relative to the input image. The feature pyramid  $\{P_2, P_3, P_4, P_5\}$  is generated by the top-down pathway in HS-FPN. Each lateral connection in HS-FPN contains two modules: high frequency perception module (HFP) and spatial dependency perception (SDP) module. HFP generates high frequency responses to enrich the information of tiny objects in  $C_i$  while SDP takes  $\{C_i, P_{i+1}\}$  as inputs using pixel-level cross attention mechanism to learn the relationships between pixels in  $\{C_i, P_{i+1}\}$  and then using valuable spatial dependencies to enrich the features of tiny targets. Please note that all laterals of HS-FPN contain the HFP module, while only  $\{P_2, P_3, P_4\}$  layers contain the SDP module. Finally, the output features of HS-FPN at each layer are obtained from  $\{P_2, P_3, P_4, P_5\}$  through separate  $3 \times 3$  convolutions for subsequent detection tasks.

### 3.1 High Frequency Perception Module

To address the issues mentioned in Section 1 that tiny objects have limited usable features and FPN does not pay extra attention to tiny objects, we propose the high frequency perception module (HFP). As shown in Figure 3, HFP consists of a high-frequency feature generator, a channel path (CP), and a spatial path (SP). Assuming the input feature  $C_i$  has a size of  $R^{(C \times H_i \times W_i)}$ , HFP first applies a high-pass filter to  $C_i$  to obtain the high frequency responses  $F_i$  which has the same size of  $C_i$ . Then  $F_i$  is used respectively as input for CP and SP to generate channel attention weights  $u^{CP} \in R^{(C \times 1 \times 1)}$  and spatial attention weights  $u^{SP} \in R^{(1 \times H_i \times W_i)}$ . These weights are broadcast along the spatial and channel dimensions, respectively, and element-wise multiplied with  $C_i$ . Finally, the results from CP and SP are added and then passed through a  $3 \times 3$  convolution to produce the output of HFP, denoted as  $C_i^{HFP}$ .

**Channel Path.** It is undeniable that each channel of  $C_i$  contributes differently to the representation of tiny objects. Therefore, we aim to identify the channels within  $C_i$  that contain more tiny object features and assign them distinct weights. General channel attention (Roy, Navab, and Wachinger 2018; Woo et al. 2018) is typically calculated across the entire feature map. However, since tiny object features occupy a relatively small proportion of the original feature map, the channel weights computed across the whole map are easily influenced by low-frequency homogeneous background. By applying a high-pass filter to  $C_i$ , most low-frequency components and homogeneous background are filtered out, increasing the proportion of tiny object fea-




Figure 3: The design ideas of HFP, which consists of a high frequency feature generator, a channel path (CP), and a spatial path (SP). Please note that the high frequency feature extractor within the dashed box is implemented only in the  $\{P_2, P_3\}$  and the two high frequency features in the diagram are identical.  $\odot$  means Hadamard product,  $\oplus$  means feature cognate and  $\oplus$  denotes pixel-by-pixel summation.  $\sigma$  refers ReLU function.

tures in the high-frequency responses  $F_i$ . As a result, calculating channel attention at  $F_i$  becomes more accurate.

To more effectively extract information from each channel, we chose to compress the spatial dimensions of  $F_i$ , ensuring that each channel of  $F_i$  is representative. Global Average Pooling (GAP) can integrate the spatial information of the feature map to obtain an overall representation, while Global Max Pooling (GMP) captures the maximum activation value in each channel. Therefore, we use both GAP and GMP to process  $F_i$ , integrating its feature information to facilitate the subsequent calculation of each channel's importance to tiny objects. Firstly, as shown in Figure 3, we perform a GAP and a GMP on  $F_i$  separately, resulting in  $F_i^{GAP}$  and  $F_i^{GMP}$ . It is important to note that we directly pool  $F_i$  into  $R^{(C \times k \times k)}$  instead of pooling it to the same size of  $u^{CP}$ , as that results in a significant loss of information. In this paper,  $k$  is generally set to 16, as when  $k$  is too small, key information is lost in both GAP and GMP; when  $k$  is too large, the effectiveness of both operations diminishes. Next, we apply the ReLU function to both  $F_i^{GAP}$  and  $F_i^{GMP}$  to retain positive values, and then sum across each channel to generate two one-dimensional feature vectors. These vectors are then passed through separate  $1 \times 1$  group convolutions to generate different channel scores. These scores are then concatenated and passed through another  $1 \times 1$  group convolution, producing the final channel attention weights  $u^{CP}$ .

**Spatial Path.** The SP of HFP is designed to help the model better focus on tiny object features within the spatial range, achieving an effect similar to self-attention (Vaswani et al. 2017). However, self-attention relies on similarity calculations between feature pixels to emphasize important information. Given that tiny object features are weak and FPN's shallow feature maps contain noise, self-attention may struggle to effectively highlight tiny object features and could even amplify noise and background responses. With the help of the high-pass filter, the low-frequency homogeneous background in  $F_i$  is filtered out and tiny object fea-




Figure 4: Illustration of High-Pass Filter.

tures are enhanced. Therefore, we use  $F_i$  as a spatial mask. Specifically, a  $1 \times 1$  convolution aggregates the channel-wise information of  $F_i$  to generate the spatial attention mask  $u^{SP}$ .

**High-Pass Filter.** To highlight tiny object features in the input feature map, we designed a high-pass filter in HFP, as shown in Figure 4. The filter  $z$  has the same size of  $F_i$ , assumed to be  $R^{(C \times H_i \times W_i)}$ . Since the low-frequency components of the DCT spectrum are concentrated in the top-left corner and the high-frequency components are distributed towards the bottom-right corner, we introduce a hyper-parameter  $\alpha$  to control the frequency band of the filter. Specifically, the values of filter  $z$  are determined as follows:

$$z(u, v) = \begin{cases} 0 & u < \alpha H, v < \alpha W \\ 1 & \text{else} \end{cases}, \alpha \in [0, 1] \quad (2)$$

When  $\alpha = 0$ , all values of filter  $z$  are 1, making  $F_i$  identical to  $C_i$ . When  $\alpha = 1$ , all values of filter  $z$  are 0, blocking all frequencies. The high-pass filter  $z$  only works when  $\alpha$  is between 0 and 1, filtering out low-frequency components from  $C_i$  to generate high-frequency responses  $F_i$ .

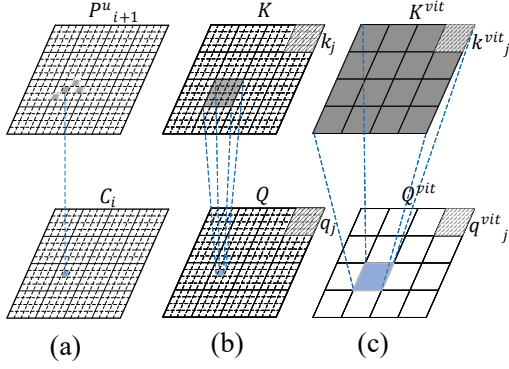


Figure 5: Computing details between different attention methods. (a) Lack of spatial dependency in FPN; (b) Attention between pixels used in SDP. (c) Attention between feature blocks used in ViT.

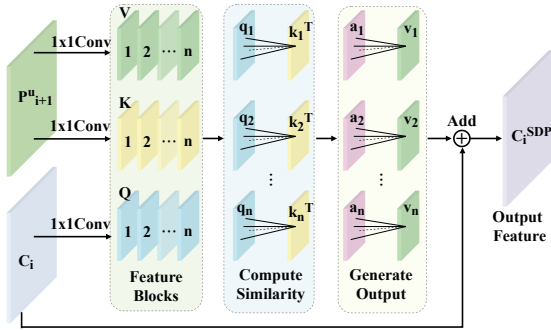


Figure 6: An over review of SDP.  $P_{i+1}^u$  is upper feature and has already been upsampled.

### 3.2 Spatial Dependency Perception Module

By merging upper and lower layer features in a pixel-by-pixel addition manner, FPN lacks spatial perception ability around tiny targets, as shown in Figure 5 (a). However, enhancing spatial perception is crucial. For instance, the pixel offset caused by recursive upsampling in FPN results in misalignment between features  $\{C_i, P_{i+1}\}$ , which greatly affects the accuracy of tiny object features. Inspired by the attention mechanism in Vision Transformer (ViT) (Dosovitskiy et al. 2020), we designed the spatial dependency perception module (SDP) to capture the mapping between pixels in  $\{C_i, P_{i+1}\}$ . The attention mechanism, proficient at capturing long-range dependencies, enables pixels from lower-level features  $C_i$  to interact not only with pixels at the same spatial location in  $P_{i+1}$ , but also with a wider range of pixels. Ultimately, we leverage the valuable spatial dependencies captured before to enrich the features of tiny objects.

As shown in Figure 6, we first upsample  $P_{i+1}$  to match the size of  $C_i$ , and name the upsampled feature  $P_{i+1}^u$ . Subsequently, we compute the Query (Q), Key (K), and Value (V) matrices through three separate  $1 \times 1$  convolutions, where Q is derived from  $C_i$  and K, V are derived from  $P_{i+1}^u$ . Next, we divide the  $\{Q, K, V\}$  matrices into multiple feature blocks.

For simplicity, let’s assume the dimensions of  $\{Q, K, V\}$  are  $R^{(C \times H_i \times W_i)}$ . Since the detection framework doesn’t constrain input image size, the feature block size must be adjusted based on the actual input image to avoid issues with non-divisible dimensions. Please note that the input feature maps in the feature pyramid  $\{C_2, C_3, C_4, C_5\}$  are multiples of each other in size, so making the feature block have the same size with  $C_5$  could ensure a perfect division. Based on this assumption, we partition  $\{Q, K, V\}$  into feature blocks and reshape them into  $R^{(n \times (H_5 \times W_5) \times C)}$ , where  $n = \frac{H_i}{H_5} \times \frac{W_i}{W_5}$ ,  $C$  is the number of channels in  $C_i$ ,  $H_5 \times W_5$  is the resolution of  $C_5$ . It can be explained as:  $C_i$  can be divided into  $n$  feature blocks with each block containing  $H_5 \times W_5$  feature points, and each point is represented as a  $C$  dimensional column vector. Then we extract corresponding feature blocks  $q_j$  and  $k_j$  ( $1 \leq j \leq n, q_j, k_j \in R^{((H_5 \times W_5) \times C)}$ ) from  $\{Q, K\}$  and calculate their similarity and obtain the similarity matrix  $a_j \in R^{((H_5 \times W_5) \times (H_5 \times W_5))}$ . Finally, each feature point in the corresponding block  $v_j$  from V is weighted by  $a_j$  to generate feature blocks rich in spatial dependencies and high-level semantics.

$$a_j = \text{Softmax}\left(\frac{q_j \times k_j^T}{\sqrt{C}}\right), 1 \leq j \leq n \quad (3)$$

After applying the cross-attention mechanism to each feature block of Q, the blocks are aggregated according to their spatial locations to form a new feature rich in spatial dependencies, which is then added to  $C_i$ . This process effectively enriches the feature content of tiny objects and infuses them with the rich semantics of high-level features.

Please note that the attention calculation process in SDP differs from that in ViT. First, ViT reshapes Q into  $R^{(n \times (H_5 \times W_5) \times C)}$ , which means Q consists of  $n$  feature points with each being a vector of dimensions  $(H_5 \times W_5 \times C)$ . Second, there is only one similarity matrix of size  $R^{(n \times n)}$  in ViT instead of  $n$  similarity matrices in SDP. In summary, SDP computes cross-attention between pixels within feature blocks, while ViT computes attention across feature blocks. These differences can be intuitively understood from Figure 5 (b, c).

## 4 Experiments and Analysis

### 4.1 Datasets and Metrics

Experiments are performed on two TOD datasets. The main experiments are conducted on the challenging AI-TOD (Wang et al. 2021), which provides a total of 700,621 instances in 8 categories in 28,036 aerial images. The average absolute instance size of AI-TOD is 12.8 pixels. Furthermore, we build a subset consisting of 10 categories named  $DOTA_{mini10}$  from DOTA (Xia et al. 2018) and test our method on it. This subset contains 410,305 instances, with 17.3% of the instances smaller than 16 pixels and 38.5% smaller than 32 pixels. All metrics used in the experiments follow the AI-TOD benchmark.

### 4.2 Experiments Settings

All of our experiments are implemented based on MMDetection (Chen et al. 2019) and PyTorch (Paszke et al. 2019).

HFP			SDP	$AP$	$AP_{50}$	$AP_{75}$	$AP_{vt}$	$AP_t$	$AP_s$	$AP_m$
CP	SP									
×	×	×		20.2	47.4	13.8	9.9	21.3	24.1	30.3
✓	×	×		21.6	49.2	15.2	9.6	23.0	25.7	30.8
×	✓	×		21.6	49.6	15.3	10.0	23.1	25.4	<b>31.7</b>
✓	✓	×		<u>22.4</u>	<u>50.4</u>	<u>16.1</u>	<b>12.1</b>	<u>23.9</u>	<u>26.1</u>	31.5
×	×	✓		21.3	49.1	15.1	8.3	23.1	25.3	30.3
✓	✓	✓		<b>23.6</b>	<b>52.3</b>	<b>17.4</b>	<u>11.6</u>	<b>25.2</b>	<b>27.0</b>	31.6

Table 1: Effect of each component on AI-TOD test set.

$\alpha$	$AP$	$AP_{50}$	$AP_{75}$	$AP_{vt}$	$AP_t$	$AP_s$	$AP_m$
0.00	21.7	49.3	15.5	9.7	23.2	<b>26.4</b>	29.9
0.25	<b>22.4</b>	<u>50.4</u>	<b>16.1</b>	<b>12.1</b>	<b>23.9</b>	<u>26.1</u>	<b>31.5</b>
0.50	<u>22.1</u>	49.9	<u>15.9</u>	10.3	<u>23.6</u>	25.8	<u>31.3</u>
0.75	21.9	<b>50.7</b>	15.4	<u>10.7</u>	23.3	25.9	31.0
1.00	21.5	49.0	15.2	10.2	22.7	25.7	30.1

Table 2: Optimal studies of  $\alpha$  on AI-TOD test set.

We train and evaluate the detectors on two NVIDIA 3080 Ti GPUs (one image per GPU). All detectors are trained using the Stochastic Gradient Descent (SGD) optimizer for 12 epochs with a momentum of 0.9 and weight decay of 0.0001. The initial learning rate of two-stage models like Faster R-CNN (Ren et al. 2015) and one-stage models like RetinaNet (Lin et al. 2020) is set to 0.005 and 0.001 respectively, and decays at the 8<sup>th</sup> and 11<sup>th</sup> epochs. Besides, the basic anchor size is set to 2 for all anchor-based models.

### 4.3 Ablation Study

We analyze the effect of each proposed component of HS-FPN on AI-TOD test subset. Results are reported in Table 1. We progressively incorporate the channel path (CP) and spatial path (SP) of HFP and SDP into FPN. Baseline detector is Cascade R-CNN (Cai and Vasconcelos 2018) with ResNet50 (He et al. 2016) as backbone.

**The effect of High Frequency Perception Module.** We implement HFP independently on each lateral layer of FPN. Results show that naively added HFP to FPN can bring 2.2  $AP$  points higher than baseline detectors, as well as  $AP_t$  and  $AP_s$  achieve 2.6 points and 2.0 points higher than baseline, respectively. This indicates that HFP has a positive effect on improving the performance of TOD. As shown in Figure 7 (a, b), the tiny object features output by HFP are significantly enhanced compared to the input features.

Furthermore, we also report the contribution of CP and SP within HFP. As shown in Table 1, when adding CP and SP individually to FPN,  $AP$  both achieve 1.4 points improvement as well as  $AP_t$  and  $AP_s$  are also improved. Moreover, when CP and SP are combined in HFP,  $AP$  is improved by 2.2 compared to FPN, and by 0.8 compared to using only SP or CP, demonstrating that the combination of CP and SP has a synergistic positive effect.

In HFP, we use the parameter  $\alpha$  to control the range of the high-pass filter. To investigate the impact of different filter-

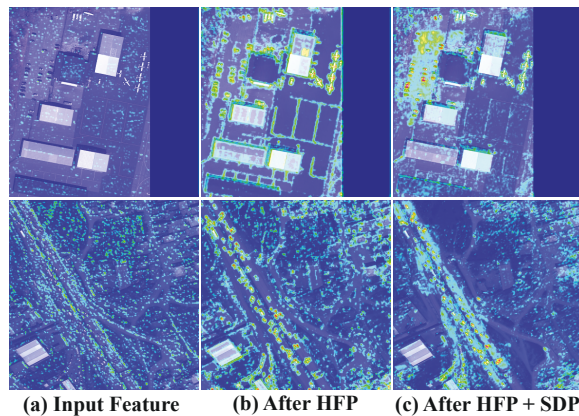


Figure 7: Visualization of the feature at the  $P_2$  level after the gradual addition of HFP and SDP in FPN.

ing rates on tiny object detection, we conducted optimization experiments for  $\alpha$ . The results show that as  $\alpha$  increases from 0 to 1, detection performance first improves and then declines, following a similar trend to the SCR depicted in Figure 1. This confirms the effectiveness of the HFP module and indicates that filtering out low-frequency information within an appropriate range can enhance the detection performance of tiny objects. Results are summarized in Table 2. Finally, in this paper,  $\alpha$  is set to 0.25 for both AI-TOD and  $DOTA_{mini10}$ .

### The effect of Spatial Dependency Perception Module.

We implement SDP on each lateral layer of FPN independently to evaluate its performance, resulting in an improvement of  $AP$  by 1.1 points. When combined with HFP, the  $AP$  increased by 3.4 points compared to FPN, 1.2 points over HFP alone, and 2.3 points over SDP alone, showing the benefits of their combination. As shown in Figure 7 (c), the combination of HFP and SDP effectively enriches and enhances the features of tiny objects, while also suppressing some of the high-frequency noise amplified by HFP.

### 4.4 Main Result

We evaluate HS-FPN on the AI-TOD and  $DOTA_{mini10}$  dataset, comparing its performance with the baseline model and other state-of-the-art models. The results are presented in Tables 3 and 4. By simply replacing FPN with HS-FPN, nearly all baseline models show significant improvements in  $AP$ . Even with a more powerful backbone, like ResNet101, or a lightweight network like MobileNetV2, the  $AP$  on Faster R-CNN still improved by 1.7 and 1.9 on AI-TOD, respectively. Additionally, using advanced methods like NWD and RLFA to optimize label assignment in anchor-based detectors further enhanced HS-FPN’s performance, yielding notable  $AP$  improvements in Faster R-CNN, Cascade R-CNN, and DetectoRS.

### 4.5 Computational Costs

Table 5 demonstrates the computation increase of each component of HS-FPN within the entire TOD framework. The entire detection model is Cascade R-CNN with ResNet50 as

Method	Backbone	$AP$	$AP_{50}$	$AP_{75}$	$AP_{vt}$	$AP_t$	$AP_s$	$AP_m$
ATSS (Zhang et al. 2019)	R-50 + FPN	15.5	38.2	10.1	4.0	14.5	21.5	31.9
MENet (Zhang et al. 2024)	R-50 + FPN	20.4	50.0	12.9	8.9	21.4	23.2	31.0
RetinaNet* (Lin et al. 2020)	R-50 + FPN	13.7	31.5	9.4	5.8	15.7	15.0	16.6
Cascade RPN (Vu et al. 2019)	R-50 + FPN	18.1	44.9	11.2	7.1	18.8	22.3	27.9
Grid R-CNN (Lu et al. 2019)	R-50 + FPN	18.4	38.7	14.6	12.1	19.0	21.8	25.5
Faster R-CNN (Ren et al. 2015)	R-50 + FPN	18.3	44.7	11.3	9.4	19.4	22.4	27.0
Faster R-CNN w/NWD (Xu et al. 2022a)	R-50 + FPN	19.4	47.3	12.5	8.0	20.4	22.7	28.6
Faster R-CNN w/RFLA (Xu et al. 2022b)	R-50 + FPN	20.9	50.8	13.0	8.1	21.2	25.9	32.7
Faster R-CNN (Sandler et al. 2018)	MobileNetV2 + FPN	15.6	39.2	9.2	9.8	16.9	18.7	20.5
Faster R-CNN	R-101 + FPN	18.8	45.8	11.6	11.1	19.7	22.9	27.9
Cascade R-CNN	R-50 + FPN	20.2	47.4	13.8	9.9	21.3	24.1	30.3
Cascade R-CNN w/NWD	R-50 + FPN	21.4	49.1	15.1	9.4	22.4	25.1	31.5
Cascade R-CNN w/RFLA	R-50 + FPN	21.9	50.6	15.9	8.3	21.9	26.9	34.7
DetectoRS (Qiao, Chen, and Yuille 2021)	R-50 w/SAC + RFP	22.8	51.4	16.9	10.7	23.7	27.4	32.3
DetectoRS w/RFLA	R-50 w/SAC + RFP	23.9	54.6	17.8	9.8	24.0	<u>28.9</u>	<b>37.4</b>
RetinaNet*	R-50 + HS-FPN	15.0[+1.3]	33.5	10.8	7.4	17.4	15.8	16.7
Cascade RPN	R-50 + HS-FPN	19.3[+1.2]	46.3	12.7	10.5	20.4	23.9	30.1
Grid R-CNN	R-50 + HS-FPN	19.3[+0.9]	40.6	15.7	11.3	21.0	21.6	26.6
Faster R-CNN	R-50 + HS-FPN	20.3[+2.0]	48.8	13.3	11.6	22.0	25.5	27.8
Faster R-CNN w/NWD	R-50 + HS-FPN	21.7[+2.3]	51.5	14.7	11.0	23.7	23.8	28.8
Faster R-CNN w/RFLA	R-50 + HS-FPN	23.1[+2.2]	54.3	15.5	9.9	23.8	27.4	33.6
Faster R-CNN	MobileNetV2 + HS-FPN	17.5[+1.9]	42.0	11.3	9.9	19.2	20.9	22.9
Faster R-CNN	R-101 + HS-FPN	20.5[+1.7]	48.9	13.4	12.2	22.1	23.9	27.4
Cascade R-CNN	R-50 + HS-FPN	23.6[+3.4]	52.3	17.4	11.6	25.2	27.0	31.6
Cascade R-CNN w/NWD	R-50 + HS-FPN	24.3[+2.9]	<u>54.8</u>	17.9	<u>12.7</u>	25.2	28.8	32.9
Cascade R-CNN w/RFLA	R-50 + HS-FPN	<u>24.6[+2.7]</u>	54.6	<u>18.6</u>	9.2	25.0	28.6	35.9
DetectoRS	R-50 w/SAC + HS-FPN	24.2[+1.4]	54.5	17.8	<b>13.1</b>	<b>25.8</b>	27.6	33.2
DetectoRS w/RFLA	R-50 w/SAC + HS-FPN	<b>25.1[+1.2]</b>	<b>55.7</b>	<b>19.1</b>	12.1	<u>25.3</u>	<b>29.9</b>	<u>36.9</u>

Table 3: Main results on AI-TOD. Models are trained on the train-val set and evaluated on the test set. Note that RetinaNet\* means using P2-P6 of FPN. The relative improvements of  $AP$  are shown in parenthesis.  $\alpha$  is set to 0.25. "w/" means with.

Method	$AP$	$AP_{50}$	$AP_t$	$AP_s$	$AP_m$
RetinaNet* + FPN	38.2	68.1	13.7	31.8	44.9
Faster R-CNN + FPN	46.9	74.2	18.6	38.5	54.3
Cascade R-CNN + FPN	<u>49.4</u>	74.2	18.1	40.0	<u>58.1</u>
RetinaNet* + HS-FPN	40.1	71.1	15.9	32.9	47.0
Faster R-CNN + HS-FPN	48.4	<u>75.7</u>	<u>21.7</u>	<u>40.3</u>	55.9
Cascade R-CNN + HS-FPN	<b>50.9</b>	<b>76.6</b>	<b>22.2</b>	<b>41.8</b>	<b>59.0</b>

Table 4: Results on  $DOTA_{mini10}$ .  $\alpha$  is set to 0.25.

Method	$AP$	$AP_t$	$AP_s$	FLOPs(G)	Params(M)
Baseline (FPN)	20.2	21.3	24.1	162.34	68.95
only CP	21.6	23.0	25.7	193.67	71.34
only SP	21.6	23.1	25.4	193.69	71.31
only HFP	<u>22.4</u>	<u>23.9</u>	<u>26.1</u>	193.69	71.34
only SDP	21.3	23.1	25.3	169.32	69.35
HS-FPN (ALL)	<b>23.6</b>	<b>25.2</b>	<b>27.0</b>	200.65	71.73

Table 5: FLOPs and parameters comparison for HS-FPN components.

the backbone and an input size of (800, 800). The Params and FLOPs for HFP mainly come from the  $3 \times 3$  convolutions in Figure 2, which results in similar computational costs for CP and SP. Given the context of the entire detection framework, HS-FPN slightly increases computational cost but significantly improves performance compared to FPN.

## 5 Conclusion

In this paper, we propose a high frequency and spatial perception feature pyramid network (HS-FPN) to enhance the detection performance of tiny objects. Given the small scale and low-quality feature representation of tiny objects, we designed a high frequency perception module (HFP) to improve feature representation from the frequency domain. Additionally, we introduced a spatial dependency perception module (SDP) to capture the spatial dependencies between adjacent pixel points in the upper and lower feature maps, enriching the features of tiny objects. Extensive experiments on two tiny object detection datasets validate the effectiveness of our proposed method.

**Extended version.** Additional motivation and details are provided in the extended version<sup>1</sup>.

<sup>1</sup><https://arxiv.org/abs/2412.10116>

## Acknowledgments

This work was partly supported by the National Key Laboratory of Science and Technology on Multi-Spectral Information Processing.

## References

- Achanta, R.; Hemami, S.; Estrada, F.; and Susstrunk, S. 2009. Frequency-tuned salient region detection. In *2009 IEEE conference on computer vision and pattern recognition*, 1597–1604. IEEE.
- Cai, Z.; and Vasconcelos, N. 2018. Cascade R-CNN: Delving Into High Quality Object Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chen, K.; Wang, J.; Pang, J.; Cao, Y.; Xiong, Y.; Li, X.; Sun, S.; Feng, W.; Liu, Z.; Xu, J.; Zhang, Z.; Cheng, D.; Zhu, C.; Cheng, T.; Zhao, Q.; Li, B.; Lu, X.; Zhu, R.; Wu, Y.; Dai, J.; Wang, J.; Shi, J.; Ouyang, W.; Loy, C. C.; and Lin, D. 2019. MMDetection: Open MMLab Detection Toolbox and Benchmark. *ArXiv*, abs/1906.07155.
- Cheng, G.; Yuan, X.; Yao, X.; Yan, K.; Zeng, Q.; Xie, X.; and Han, J. 2023. Towards Large-Scale Small Object Detection: Survey and Benchmarks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11): 13467–13488.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Li, Z.; Wang, Y.; Zhang, N.; Zhang, Y.; Zhao, Z.; Xu, D.; Ben, G.; and Gao, Y. 2022. Deep Learning-Based Object Detection Techniques for Remote Sensing Images: A Survey. *Remote Sensing*, 14(10).
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature Pyramid Networks for Object Detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 936–944.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2020. Focal Loss for Dense Object Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2): 318–327.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In Fleet, D.; Pajdla, T.; Schiele, B.; and Tuytelaars, T., eds., *Computer Vision – ECCV 2014*, 740–755. Cham: Springer International Publishing. ISBN 978-3-319-10602-1.
- Liu, S.; Qi, L.; Qin, H.; Shi, J.; and Jia, J. 2018. Path aggregation network for instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8759–8768.
- Lu, X.; Li, B.; Yue, Y.; Li, Q.; and Yan, J. 2019. Grid R-CNN. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7355–7364.
- Luo, Y.; Cao, X.; Zhang, J.; Guo, J.; Shen, H.; Wang, T.; and Feng, Q. 2022. CE-FPN: enhancing channel information for object detection. *Multimedia Tools and Applications*, 81(21): 30685–30704.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Qiao, S.; Chen, L.-C.; and Yuille, A. 2021. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10213–10224.
- Qin, Z.; Zhang, P.; Wu, F.; and Li, X. 2021. FcaNet: Frequency Channel Attention Networks. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 763–772.
- Rao, Y.; Zhao, W.; Zhu, Z.; Zhou, J.; and Lu, J. 2023. GFNet: Global Filter Networks for Visual Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9): 10960–10973.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In Cortes, C.; Lawrence, N.; Lee, D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Roy, A. G.; Navab, N.; and Wachinger, C. 2018. Concurrent spatial and channel ‘squeeze & excitation’ in fully convolutional networks. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I*, 421–429. Springer.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4510–4520.
- Tong, K.; and Wu, Y. 2022. Deep learning-based detection from the perspective of small or tiny objects: A survey. *Image and Vision Computing*, 123: 104471.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Vu, T.; Jang, H.; Pham, T. X.; and Yoo, C. 2019. Cascade rpn: Delving into high-quality region proposal network with adaptive convolution. *Advances in neural information processing systems*, 32.
- Wang, J.; Yang, W.; Guo, H.; Zhang, R.; and Xia, G.-S. 2021. Tiny Object Detection in Aerial Images. In *2020 25th*

*International Conference on Pattern Recognition (ICPR)*, 3791–3798.

Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. CBAM: Convolutional Block Attention Module. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

Xia, G.; Bai, X.; Ding, J.; Zhu, Z.; Belongie, S.; Luo, J.; Datcu, M.; Pelillo, M.; and Zhang, L. 2018. DOTA: A Large-Scale Dataset for Object Detection in Aerial Images. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3974–3983. Los Alamitos, CA, USA: IEEE Computer Society.

Xu, C.; Wang, J.; Yang, W.; Yu, H.; Yu, L.; and Xia, G.-S. 2022a. Detecting tiny objects in aerial images: A normalized Wasserstein distance and a new benchmark. *ISPRS Journal of Photogrammetry and Remote Sensing*, 190: 79–93.

Xu, C.; Wang, J.; Yang, W.; Yu, H.; Yu, L.; and Xia, G.-S. 2022b. RFLA: Gaussian receptive field based label assignment for tiny object detection. In *European conference on computer vision*, 526–543. Springer.

Xu, C.; Wang, J.; Yang, W.; and Yu, L. 2021. Dot distance for tiny object detection in aerial images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1192–1201.

Xu, K.; Qin, M.; Sun, F.; Wang, Y.; Chen, Y.-K.; and Ren, F. 2020. Learning in the Frequency Domain. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1737–1746.

Zeng, Y.; Jiang, C.; Mao, J.; Han, J.; Ye, C.; Huang, Q.; Yeung, D.-Y.; Yang, Z.; Liang, X.; and Xu, H. 2023. CLIP2: Contrastive Language-Image-Point Pretraining from Real-World Point Cloud Data. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15244–15253.

Zhang, S.; Chi, C.; Yao, Y.; Lei, Z.; and Li, S. Z. 2019. Bridging the Gap Between Anchor-Based and Anchor-Free Detection via Adaptive Training Sample Selection. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9756–9765.

Zhang, T.; Zhang, X.; Zhu, X.; Wang, G.; Han, X.; Tang, X.; and Jiao, L. 2024. Multistage enhancement network for tiny object detection in remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–12.

Zhang, Y.; Zhang, J.; Wang, D.; and Chen, C. 2011. Infrared small target detection based on morphology and wavelet transform. In *2011 2nd International Conference on Artificial Intelligence, Management Science and Electronic Commerce (AIMSEC)*, 4033–4036.

Zhao, J.; Liu, F.; and Mo, B. 2012. An Algorithm of Dim and Small Target Detection Based on Wavelet Transform and Image Fusion. In *2012 Fifth International Symposium on Computational Intelligence and Design*, volume 2, 43–45.

Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; and Gao, J. 2022. RegionCLIP: Region-based Language-Image Pretraining. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16772–16782.