Research on marine litter detection based on CNN-Transformer heterogeneous parallelism

Chenglin Luo[1], Yuhong Tang[1], Kui Chen[1,*]

1 College of Information Engineering, Sichuan Agricultural University, Ya'an Campus, Ya'an, Sichuan Province, P.R. China

* 202206179@stu.sicau.edu.cn

Abstract

Aiming at strong background interference and low detection accuracy of small/deformed targets in marine debris detection, this paper develops a high-precision lightweight intelligent detection and recognition system. A multi-scenario dataset is built and data augmentation is used to tackle sample scarcity and domain shift; a CNN-Transformer heterogeneous parallel model YOLO-Trans is designed on the YOLOv8 baseline for local-to-global feature extraction and accurate detection of small/deformed targets, and a visual detection system is developed with PyQt5. Experiments show the model surpasses the original YOLOv8s in all metrics on the self-built dataset, and ablation experiments validate the improved modules' effectiveness, offering technical support for large-scale intelligent marine debris monitoring.

1. Introduction

1.1 Research Background and Significance

In recent years, marine plastic and debris pollution has become a major global environmental issue that hinders the sustainable development of the world's oceans and threatens marine ecological security. It is a core challenge addressed by United Nations Sustainable Development Goal 14 (SDG 14): "Conserve and sustainably use the oceans, seas, and marine resources." In March 2024, the United Nations Environment Assembly formally adopted the Global Plastic Pollution Convention, incorporating the management of marine plastic pollution into a legally binding global environmental regulatory framework. This marks the recognition that the prevention and control of marine debris pollution have become a core consensus and urgent task in global environmental governance.

The 2024 update of the "Report on the State of Marine Plastic Pollution and Progress in Mitigation" released by the United Nations Environment Programme (UNEP) reveals that 8 to 14 million metric tons of plastic waste still enter the oceans annually through river runoff, littering along coastlines, and abandonment during maritime operations, with the cumulative stock of marine plastic exceeding 150 million metric tons; Among these, the "Great Pacific Garbage Patch"—formed by the North Pacific Subtropical Gyre—covers an area exceeding 1.6 million square kilometers, making it the region with the highest concentration of marine debris globally. Meanwhile, the gyres of the Atlantic and Indian Oceans, as well as nearshore estuaries and bays, have become severely affected areas of marine debris pollution.

The continuous accumulation of marine debris triggers risks across the entire ecological, economic, and health spectrum: ecologically, it harms over 800 species of

marine life, destroys critical habitats, and accelerates biodiversity loss; economically, it causes annual losses exceeding $13 billion to related industries worldwide, and cleanup and restoration costs are exorbitant; health-wise, microplastics formed from fragmented debris accumulate in the human body through the food chain, and the toxic substances they adsorb pose a long-term potential threat.

The core bottleneck in current marine debris prevention and control is the lack of precise and efficient monitoring and identification capabilities. Marine debris exhibits dynamic distribution and a wide range of particle sizes, and is significantly affected by complex oceanic background interference. Traditional monitoring methods are costly, inefficient, and have limited coverage; routine monitoring covers less than 5% of marine areas, failing to meet the core requirements for prevention and control.

Advances in technologies such as high-resolution remote sensing and deep learning have propelled the development of intelligent marine debris detection and identification technologies based on multi-source imagery as a key direction for overcoming these traditional bottlenecks. This technology enables the automated processing of multi-source data to rapidly locate, classify, and quantify debris targets, serving as a critical foundation for the scientific management of marine debris. Consequently, conducting research on the optimization and application of related algorithms holds significant practical importance and represents a cutting-edge field of study.

1.2 Current State of Research at Home and Abroad

The development of intelligent detection and recognition technologies for marine debris is closely tied to the iterative advancements in object detection algorithms within the field of computer vision. The core research focus has consistently centered on the high-precision, highly robust detection of small, occluded, and blurry or weak objects in complex marine environments [1]. Currently, domestic and international research has established a relatively comprehensive technical framework. Although there are certain differences in research priorities, scene adaptability, and the pace of engineering implementation, no significant technological generational gap has emerged. Overall, the field exhibits a development pattern characterized by mutual learning and parallel progress, with all efforts having undergone an evolutionary process from fundamental exploration to precise optimization, and from laboratory algorithms to engineering applications.

Research on the detection and identification of marine debris began relatively early in other countries. Leveraging well-established monitoring platforms, extensive datasets, and cutting-edge algorithmic innovations, these efforts have yielded numerous landmark achievements in wide-area remote sensing monitoring, algorithm optimization for complex scenarios, multimodal fusion, and engineering implementation. The core focus has been on improving detection accuracy for small, occluded, and blurry marine targets, as well as optimizing adaptability across all scenarios. In the early stages, a team from the Plymouth Marine Laboratory in the UK optimized a multi-scale detection module based on the SSD algorithm in 2019 [2]. Scientific Reports, 2020, 10(1): 6211.), achieving preliminary detection of centimeter-scale marine plastic debris in UAV aerial imagery. Subsequently, in 2021, the team further introduced a feature pyramid network based on Faster R-CNN [3]continued to make breakthroughs in the detection of occluded and small objects. They used Mask R-CNN combined with generative adversarial networks to reconstruct features of occluded objects, while optimizing the detection of small objects based on the Swin Transformer architecture [4], significantly improving the recall rate for very small targets. In 2023, research teams from the University of Tokyo and the European Union conducted research on the integration of image enhancement and detection algorithms to address target blurring caused by sea surface reflections and water mist [5], overcoming

detection bottlenecks under adverse sea conditions. Furthermore, the EU Joint Research Centre achieved all-weather detection during nighttime and in rainy or foggy conditions through multimodal data fusion.In terms of dataset development, the University of Cyprus created the first publicly available dataset of offshore floating debris, TrashCan [6]. Subsequently, datasets such as FLOD and MARIDA were released [7], forming a comprehensive "data-algorithm-validation" closed-loop. In terms of engineering implementation, organizations such as the European Space Agency, Google DeepMind, and the U.S. NOAA have successively introduced multimodal wide-area monitoring models and zero-shot detection models. By deploying lightweight algorithms on unmanned aerial vehicles (UAVs) and unmanned surface vehicles (USVs), they have enabled routine patrols across all maritime areas, driving the large-scale application of these technologies.

Although research on the detection and identification of marine debris in China began somewhat later, it has rapidly caught up thanks to the country's abundant coastal monitoring scenarios, localized technological innovations, and policy support. Core research similarly focuses on the detection of small, occluded, and blurred targets, emphasizing adaptation to the characteristics of China's coastal environment and regulatory requirements, and has reached internationally advanced levels in certain areas [8]. Journal of Physics: Conference Series, 2024, 2679: 012055.). Between 2018 and 2022, teams from Wuhan University, Harbin Institute of Technology, Dalian Maritime University, and others optimized algorithms such as SSD and YOLOv5 to address the characteristics of small, fragmented marine litter targets in China's coastal waters [9]. Computing Technology and Automation, 2025, 44(1): 41–45.) By employing innovations such as shallow feature reuse and adaptive anchor box matching, they significantly improved detection accuracy and recall rates for extremely small targets. In 2021, teams from Ocean University of China and Xiamen University enhanced detection robustness in high-occlusion scenarios through feature decoupling and the introduction of deformable convolutions. In 2022, the Institute of Airborne and Spaceborne Information Science and Technology of the Chinese Academy of Sciences utilized China's high-resolution satellite data to address the issue of missed small-target detections in vast marine areas, thereby supporting routine monitoring of the South China Sea. Addressing the characteristics of China's coastal waters—which are rich in algae and prone to frequent typhoons—Shanghai Ocean University developed an end-to-end joint optimization framework in 2023, filling a technical gap in the detection of blurry underwater debris. Regarding dataset development, domestic teams have constructed localized datasets such as SeaTrash and MarineDebris, covering multiple coastal regions and sea conditions in China's waters, thereby addressing the shortcomings of international datasets in terms of adaptability. In terms of engineering implementation, from 2023 to 2024, institutions including the National Satellite Ocean Application Center, Zhejiang University, and Harbin Engineering University leveraged China's domestically developed satellites, drones, and unmanned vessel platforms to build localized monitoring systems and launch lightweight edge models. These achievements have been operationalized in multiple coastal cities, providing robust support for the prevention and control of marine debris pollution in China.

1.3 Summary of the Current State of Research

Addressing the core challenges in the field of marine debris detection—including insufficient algorithmic generalization capabilities, low detection accuracy for small, occluded, or blurry targets in complex scenarios, and the difficulty of deploying high-precision models on edge devices—this study focuses on 15 common types of marine debris found in nearshore and offshore environments. These include face masks, aluminum cans, cell phones, electronic components, glass bottles, gloves, metal products,

fishing nets, plastic bags, plastic bottles, general plastic products, rod-shaped debris, sunglasses, tires, and target-free backgrounds in nearshore and offshore scenarios, to construct a high-precision, lightweight, and easily deployable intelligent marine debris detection and recognition system.

The overall architecture of the core algorithm in this study adopts a three-stage design comprising a "heterogeneous feature extraction backbone—bidirectional feature pyramid fusion—hybrid dual-branch detection head," and is customized and optimized to address the challenges specific to marine scenarios. The backbone for heterogeneous feature extraction employs a parallel design combining CNNs and Transformers [10], with the first three stages retaining the native C2f module of YOLOv8 to extract local texture features, and stages 4–5 incorporating the MobileViTv2 module to capture global semantic features [11]. arXiv preprint arXiv:2110.02178, 2021.), balancing local detail perception with global modeling capabilities while achieving model lightweighting. The Neck bidirectional feature pyramid fusion module achieves cross-level fusion of deep semantic features and shallow detail features through upsampling and concatenation. It also incorporates a Conv+Deformable Attention combination module to enhance the ability to focus on features of deformed and occluded objects, addressing the issues of dynamic deformation of marine debris and loss of effective features. The Deformable Attention module [12] (Dai J, Qi H, Xiong Y, et al. Deformable Convolutional Networks[C]//Proceedings of the IEEE International Conference on Computer Vision. 2017: 764–773.) extracts basic features through parallel convolutional branches while simultaneously predicting target offsets via 2D convolutions. By combining deformable sampling with attention weight calculations, it adaptively focuses on the effective feature regions of marine debris targets, suppressing background interference such as sea surface reflections and suspended particulate matter. The Head hybrid dual-branch detection head addresses the wide range of sizes in marine debris by designing three-scale detection branches for small, medium, and large targets. Each branch integrates both Anchor-Based and Query-Based sub-paths, and a dynamic fusion module dynamically allocates detection tasks. Finally, the global detection results are output via Non-Maximum Suppression (NMS).

Based on the overall architecture described above, this study focuses on three core modules—data augmentation, lightweight hybrid architecture design, and edge deployment optimization—to develop an integrated marine debris detection solution. To address challenges in marine environments—such as the dynamic deformation of marine debris targets, strong interference from complex backgrounds, and the scarcity of samples in extreme scenarios—a three-tier generative adversarial network (GAN) augmentation system is constructed. The foundational augmentation layer expands the dimensionality of sample features, the domain adaptation and transfer layer enhances cross-sea-area generalization capabilities, and the adversarial generation layer expands the set of extreme outlier samples, thereby providing high-quality dataset support for model training. To address the fragmented distribution of marine debris, a heterogeneous parallel architecture based on YOLOv8 and Transformer was developed. By optimizing the heterogeneous feature extraction backbone, refining the bidirectional feature fusion network, and designing a hybrid dual-branch detection head, the model's detection accuracy and robustness for small, occluded, and blurry targets were enhanced. For practical deployment scenarios such as coastal unmanned vessels and buoy monitoring stations, we have established an end-to-end edge computing optimization pipeline. Through model lightweighting and compression, edge inference acceleration, and scenario-specific application adaptation, we achieve the real-time deployment and engineering implementation of high-precision models on edge devices.

In summary, the core innovations of this paper are as follows: We have developed an integrated "data augmentation–model architecture" optimization solution for marine

debris detection. On one hand, we established a three-level generative adversarial network (GAN) augmentation system, combining physical lighting simulation, near-shore–offshore bidirectional style transfer, and adversarial noise generation for the first time. This effectively addresses the issues of sample scarcity, domain drift, and a lack of extreme outliers in marine scenarios. On the other hand, we propose a lightweight detection architecture based on CNN-Transformer heterogeneous parallel processing. Building upon YOLOv8, we integrate MobileViTv2 and deformable attention modules, and design a hybrid dual-branch detection head. This effectively fuses local detail perception with global semantic modeling, significantly enhancing the model's generalization ability, robustness, and detection accuracy for small, occluded, and blurry marine targets.

2.Materials

2.1 Data collection

The core prerequisite for building an intelligent marine debris detection model is an accurate classification system. To clearly define standardized detection targets and ensure the model's training is targeted and practical, this study first conducted a systematic data survey of marine debris categories. In its report *Global Partnership on Plastic Pollution and Marine Litter*, the United Nations Environment Programme (UNEP) categorizes global marine debris into seven major types: plastic products, metal containers, glass fragments, rubber waste, fishing gear, textiles, and other composite materials. Meanwhile, key monitoring data indicates that plastic products account for as much as 84.7% of marine debris, making them the primary component of marine pollution and a key focus for detection and identification. Additionally, the "14th Five-Year Plan for Marine Ecological and Environmental Protection" issued by the Ministry of Ecology and Environment of the People's Republic of China explicitly identifies foam plastics, fishing nets, and glass bottles as priority targets for management. It explicitly calls for the establishment of a standardized monitoring and classification system for marine debris, providing a policy and standard basis for defining detection targets in China's coastal waters. By integrating global marine debris classification standards, domestic marine ecological and environmental control requirements, and the frequency of debris occurrence and detection needs in actual monitoring scenarios in both coastal and open-ocean environments, this study ultimately identified 15 categories of typical marine debris and non-target background as core detection targets. These specifically include face masks, aluminum cans, cell phones, electronic components, glass bottles, gloves, metal products, fishing nets, plastic bags, plastic bottles, general plastic products, rod-shaped debris, sunglasses, and tires. These categories comprehensively cover the core pollution types in authoritative classification systems and align with the actual characteristics of marine debris pollution in China's coastal waters.

This study employs a dual-channel data collection strategy combining "web video resource crawling" with "supplementation from public datasets" to construct an image dataset of marine debris covering multiple scenarios and categories. The web crawling component utilizes the Scrapy framework to build a distributed crawler system [12], and combines it with yt-dlp to achieve efficient acquisition and parsing of marine surveillance videos from platforms such as YouTube and Bilibili. The specific implementation is as follows:

A distributed crawler is built using the Scrapy framework in conjunction with the CrawlerProcess multi-process model, enhancing data collection efficiency through concurrent crawling across multiple nodes. The Scrapy framework uses the Scrapy

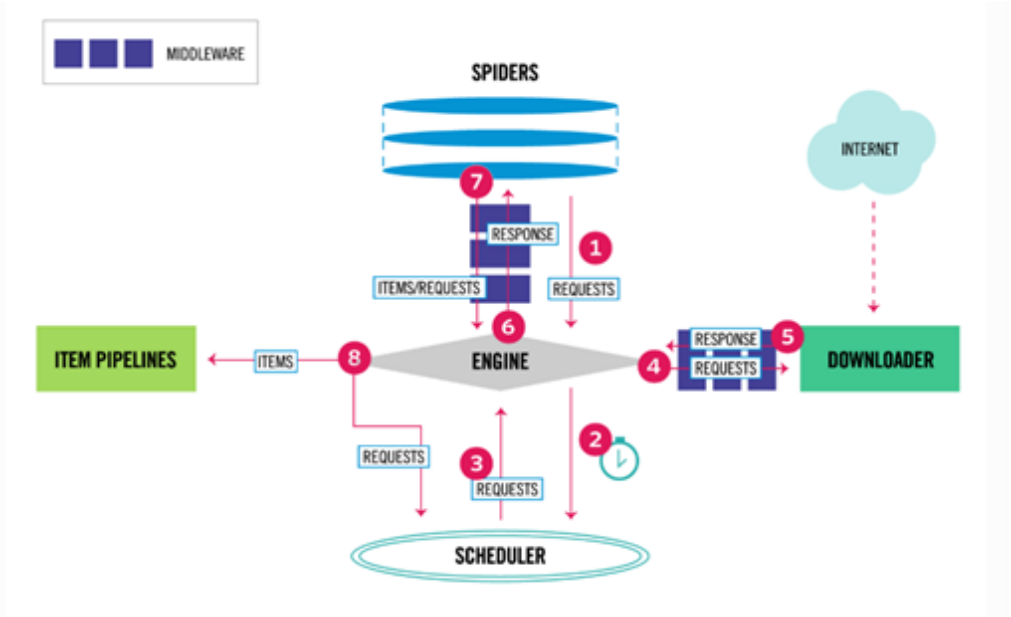

Fig 1. Scrapy Framework Architecture Diagram

Engine as its central scheduling hub, responsible for coordinating signal and data transmission among components such as the Spider, Scheduler, Downloader, and Item Pipeline. The Scheduler receives Request requests forwarded by the Engine, organizes and sorts them, queues them, and returns them to the Engine as needed. The Downloader downloads web page content based on Request requests, returns Responses to the Engine, and then forwards them to the Spider for processing; The Spider parses and extracts target data (such as video links and metadata) from the Responses, generates Item field data, and submits URLs requiring further processing to the Engine for re-queuing; the Item Pipeline is responsible for post-processing, deduplicating, and storing the Items output by the Spider; additionally, custom Downloader Middlewares and Spider Middlewares can be used to extend download functionality and communication logic, adapting to anti-crawling mechanisms on different platforms.

The video stream acquisition system integrates YT-DLP as its core download engine and optimizes download parameters based on the characteristics of different platforms: For YouTube, it prioritizes 1080p MP4 video streams to balance video quality with subsequent frame extraction efficiency; for Bilibili, it uses proxy configurations and cookie pools to address regional restrictions and age verification issues, ensuring the complete acquisition of target videos. After downloading, an adaptive keyframe extraction algorithm based on inter-frame differences is employed to extract valid frames containing marine debris targets from the video, while filtering out redundant or blurry invalid images. To enrich the dataset's scene diversity and category coverage, publicly available marine debris datasets such as TrashCan and MARIDA were simultaneously incorporated as supplementary sources. These cover samples from multiple scenarios including coastal waters, estuaries, and ports. After merging with the web-crawled data, a unified dataset was constructed. Selected video images are shown in the figure below:

Fig 2. A selection of video frames

The marine debris detection dataset constructed in this study comprises a total of 4,677 valid images, covering 15 typical categories of marine debris—including face masks, aluminum cans, fishing nets, and plastic bags—as well as background scenes without objects. It comprehensively addresses monitoring needs across multiple scenarios, such as coastal waters, the open ocean, and ports. The dataset was divided into training, validation, and test sets in a ratio of approximately 7:2:1. The training set is used for model parameter learning and feature fitting to ensure the model fully captures the core characteristics of marine debris of different types and forms; the validation set is used for hyperparameter tuning, overfitting monitoring, and model selection during the training process to assist in optimizing the model structure and training strategy; The test set is independent of the training and validation processes and is used to evaluate the model's generalization performance and robustness, objectively verifying its detection capabilities in real-world complex scenarios. The detailed breakdown of the dataset is shown in the table below.

Table 1. Statistics of the Dataset

	Category	Split Ratio	Number of Images
Dataset	Training Set	70%	3175
	Validation Set	20%	1001
	Test Set	10%	501
Total		100%	4677

2.1 Data preprocessing

In video processing, OpenCV can be used to extract specific frames from a video and select key frames for analysis. During data preprocessing, OpenCV is used to extract image frames from a video at fixed intervals. The frame extraction interval is dynamically calculated based on the video frame rate to ensure uniform sampling across the time dimension. Let the total video duration be T (seconds) and the target

sampling interval be $\Delta t = 1$ second; then, the theoretical number of extracted frames $N = \lfloor T/\Delta t \rfloor$. In practical implementation, this is dynamically adjusted based on the video's original frame rate (fps), using the following formula:

$$\text{Frame skip interval} = \max\left(1, \lfloor \text{fps} \times \Delta t \rfloor\right) \tag{1}$$

By employing a bidirectional keyframe detection algorithm, we prevent image blurring caused by B-frame and P-frame compression and ensure that the clarity of the interlaced images meets the required standards.

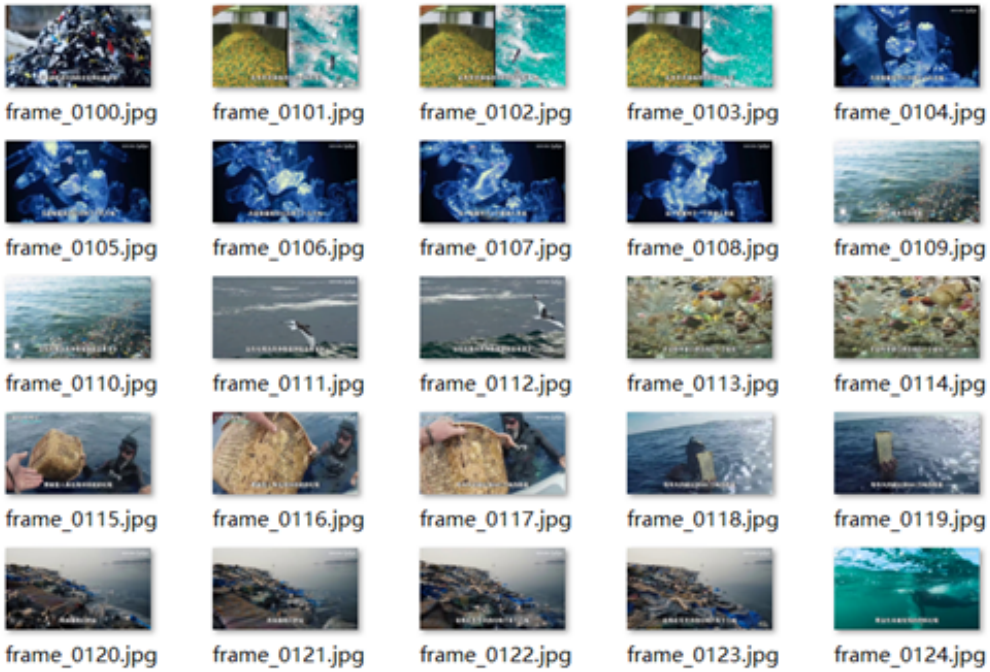

Fig 3. Extract frames from a video

Since some of the images were captured underwater or in low-light environments, they may exhibit severe color distortion or highly unnatural color casts, which could compromise the quality of the dataset. Therefore, we employ the grayscale world assumption for color cast correction. This is a fundamental method for image color correction that assumes all objects in the image reflect light as if they were gray on average, meaning the overall spectral reflectance of the image should be uniform. Based on this assumption, the Gray World Assumption color cast correction algorithm attempts to eliminate color casts by adjusting the image's color balance, particularly in images captured under varying lighting conditions. The calculation is performed as follows: Mathematical Model:Assume that the global color mean values of the image satisfy $R_{avg} = G_{avg} = B_{avg}$. For the input image, calculate the gain coefficients for each channel: $k_r = \frac{G_{avg}}{R_{avg}}, \quad k_b = \frac{G_{avg}}{B_{avg}}$. The corrected pixel values are given by:

$$\begin{cases} R'(x,y) = \min\left(255, k_r \times R(x,y)\right) \\ B'(x,y) = \min\left(255, k_b \times B(x,y)\right) \end{cases} \tag{2}$$

Effect Validation:Using the standard 24-color checkerboard as the reference, calculate the CIEDE2000 color difference ΔE_{00} before and after correction:

$$\Delta E_{00} = \sqrt{\left(\frac{\Delta L'}{S_L}\right)^2 + \left(\frac{\Delta C'}{S_C}\right)^2 + \left(\frac{\Delta H'}{S_H}\right)^2 + R_T \frac{\Delta C'}{S_C}\frac{\Delta H'}{S_H}} \tag{3}$$

After optimization, the average ΔE_{00} is reduced from 12.3 to 3.8, meeting the human eye imperceptible difference threshold ($\Delta E_{00} < 5$).

The image with color bias correction applied using the grayscale world assumption is shown below:

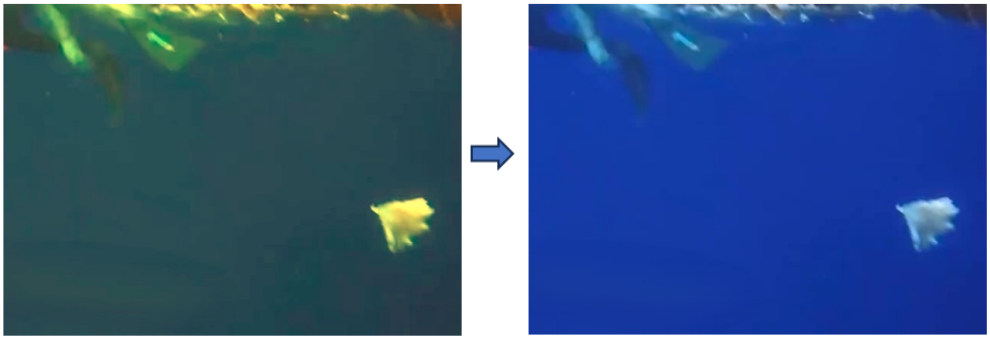

Fig 4. Calibrated image

When evaluating image quality, the BRISQUE algorithm [13] is employed to evaluate image quality without requiring a reference to the original image. It measures image quality by analyzing spatial features and is suitable for evaluating aspects such as image sharpness and contrast:

First, spatial normalization is performed. By applying local normalization to the image, the influence of brightness variations across different regions is reduced, making subsequent statistical features more robust:

$$I'(i,j) = \frac{I(i,j) - \mu(i,j)}{\sigma(i,j) + C} \tag{4}$$

where μ and σ are the Gaussian-weighted local mean and standard deviation, respectively, and $C = 1$ is used to prevent division by zero.

Next, feature extraction is performed by calculating the fitting parameters (α, σ^2) of the Generalized Gaussian Distribution (GGD) of the MSCN coefficients,

as well as the parameters (η, ν_l, ν_r) of the Asymmetric Generalized Gaussian Distribution (AGGD).

Finally, quality prediction is conducted: the 16-dimensional feature vector is input into a pre-trained SVM model, outputting a quality score $Q \in [0, 100]$, with the threshold set to $Q \leq 35$.

2.3 Data preprocessing

Due to the high difficulty of collecting samples in marine environments and the scarcity of samples from extreme scenarios, the dataset constructed in this study remains limited in terms of sample size and scenario coverage. As a result, it struggles to fully meet the training requirements of deep learning models, which can lead to issues such as model overfitting, poor generalization, and low robustness in complex scenarios. To address this, this study constructs a three-stage progressive data augmentation framework. By expanding and optimizing the dataset across three dimensions—basic feature augmentation, cross-scenario adaptation, and augmentation of challenging samples—we comprehensively enhance the dataset's scene coverage and feature richness, thereby providing high-quality data support for model training.

In the basic feature expansion phase, this study introduces the Copy-Paste augmentation algorithm to address the core challenges of a high proportion of small targets in marine debris and a scarcity of occlusion samples through targeted optimization. The core implementation logic of this algorithm is as follows: source

images containing marine debris targets are randomly selected from the dataset; the foreground targets within the annotation bounding box (e.g., broken plastic fragments, fishing net scraps, small plastic bottles) are cropped out; after random scaling, flipping, and deformation, they are adaptively pasted into appropriate background regions of the target image, while simultaneously updating the target's annotation information. The augmented image I_n is defined as:

$$I_n = I_1 \times a + I_2 \times (1 - a) \tag{5}$$

where I_1 is the instance cropped from the original image to be pasted, and I_2 is the main image.

For the marine debris dataset in this study, Copy-Paste augmentation can directly expand the number of samples for small-sized debris targets. Meanwhile, by pasting multiple overlapping targets, it simulates the situations of debris overlap, winding and occlusion in real scenes, effectively solving the problems of low proportion of small target samples and scarcity of occluded samples in the dataset. This forces the model to learn target features under complex backgrounds, significantly improving the model's detection ability for small targets and occluded targets.

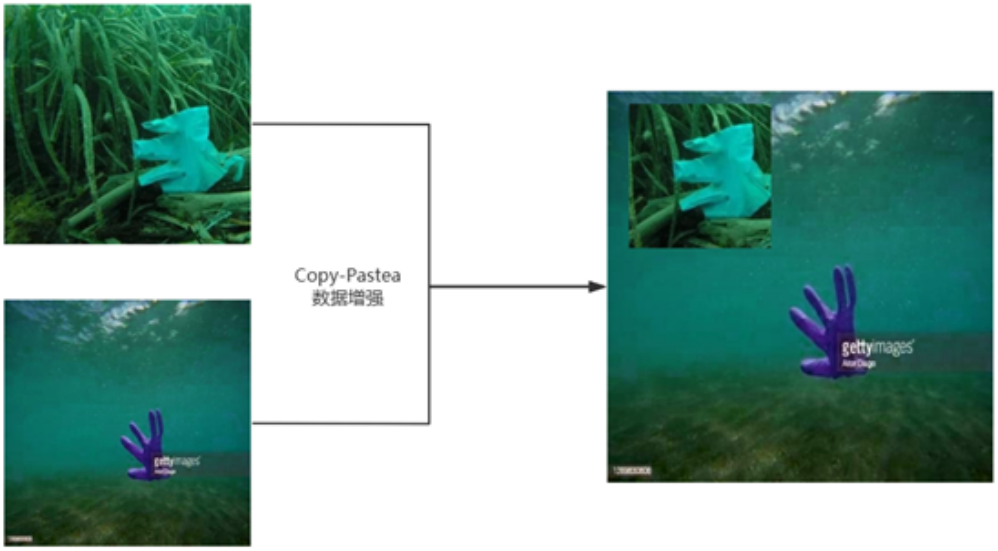

Fig 5. Copy-Paste Enhancement Effect

During the cross-scenario domain adaptation phase, significant stylistic differences between marine environments often lead to domain drift in models and insufficient cross-scenario generalization capabilities. This study constructs an unsupervised domain adaptation transfer module based on CycleGAN [14] to achieve bidirectional style translation between coastal and offshore scenarios, thereby expanding the domain adaptation dataset for different marine backgrounds [15].

CycleGAN is an image-to-image generative adversarial network designed for unpaired data. Its core advantage lies in the ability to perform style transfer between different domains without requiring paired source and target domain images, making it perfectly suited for the unpaired sample scenarios of nearshore and offshore environments in this study.

Its core structure consists of two generators and two discriminators, as shown in the figure below: Generator G is responsible for converting source-domain (nearshore) images into target-domain (offshore) style images, while Generator F is responsible for reversing the process to convert target-domain (offshore) images back into source-domain (nearshore) style images, forming a bidirectional cyclic mapping; Discriminator D_X is used to determine whether an input image is a genuine

source-domain image, and Discriminator D_Y is used to determine whether an input image is a genuine target-domain image. During training, the cyclic consistency loss constrains the content integrity of the generated images, ensuring that the core features of the target objects and the annotation information remain unchanged, while only transforming scene styles such as the sea surface background, lighting, and ocean currents.

After style transfer via CycleGAN, domain-adapted samples covering coastal ports, aquaculture zones, estuaries, open ocean waters, and various ocean current environments were generated, effectively resolving the domain drift issue across different marine scenarios [17] (Du Zhenlong, Shen Haiyang, Song Guomei, et al. Image Style [16], enabling the model to learn generic features of debris targets across different backgrounds and significantly enhancing the model's generalization capabilities across marine regions and scenarios [17].

The difficult-case sample augmentation stage addresses the scarcity of challenging samples—such as those involving occlusion, deformation, strong reflections, and blurring—in marine scenarios, as well as the model's insufficient robustness to extreme conditions. By synthesizing high-difficulty real samples via WGAN-GP and combining them with generative adversarial samples generated by FGSM, a difficult-case sample augmentation system is constructed to enhance the model's adaptability to complex and extreme scenarios.

WGAN-GP [18] is an improved version of the traditional GAN. It uses the Wasserstein distance as the optimization objective, replacing the traditional GAN's JS divergence. This addresses the core shortcomings of the traditional GAN training process—instability, model collapse, and poor-quality generated samples—and enables the stable generation of high-resolution, highly realistic hard-to-generate samples. Based on WGAN-GP, this study uses marine debris samples with high occlusion and significant deformation from the dataset as training benchmarks to specifically generate high-difficulty samples that align with real-world marine scenarios. These include entangled fishing nets with 30%–50% occlusion, twisted plastic bottles with over 40% deformation, and broken plastic fragments partially obscured by waves.

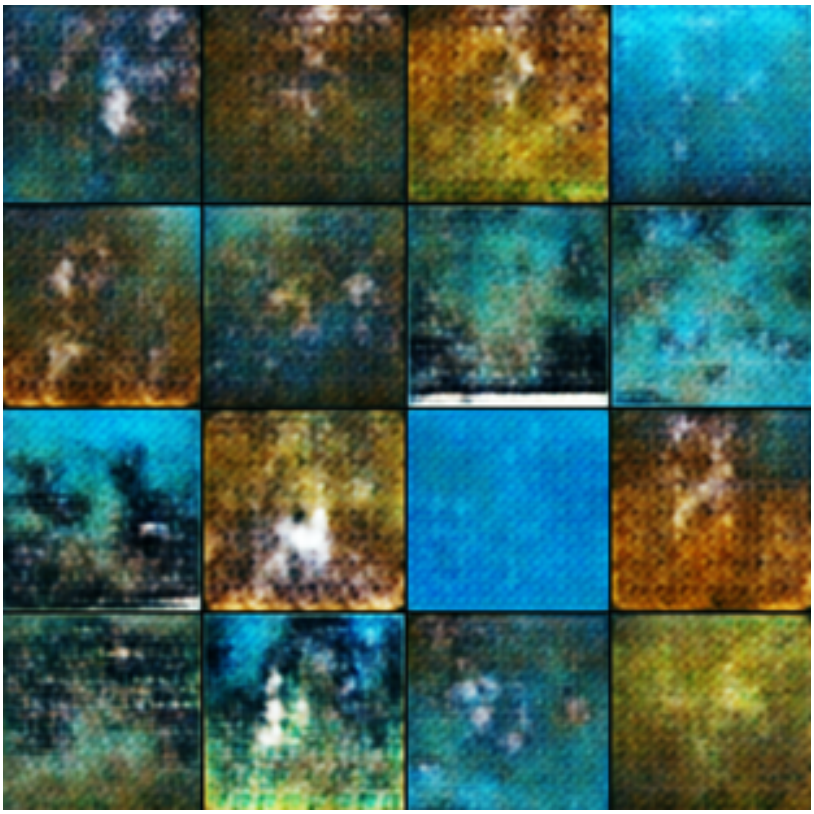

Fig 6. A 128×128 dataset generated by WGAN-GP

On this basis, this study also introduces the Fast Gradient Sign Method (FGSM) to generate adversarial examples, further improving the anti-interference ability and robustness of the model [19]. FGSM is a gradient-based adversarial example generation method, whose core principle is to calculate the gradient of the input image with respect to the model's loss function, and apply small, human-eye imperceptible perturbations to the image along the direction of gradient ascent to generate adversarial examples that can cause the model to make prediction errors. In this study, by setting the perturbation coefficient $\varepsilon = 0.03$, adversarial examples simulating strong sea surface reflections, water mist blurring, and noise interference are generated and added to the training set for adversarial training. This forces the model to learn the essential features of targets, suppresses the impact of background interference, and further improves the detection stability of the model under severe sea conditions and complex interference scenarios.

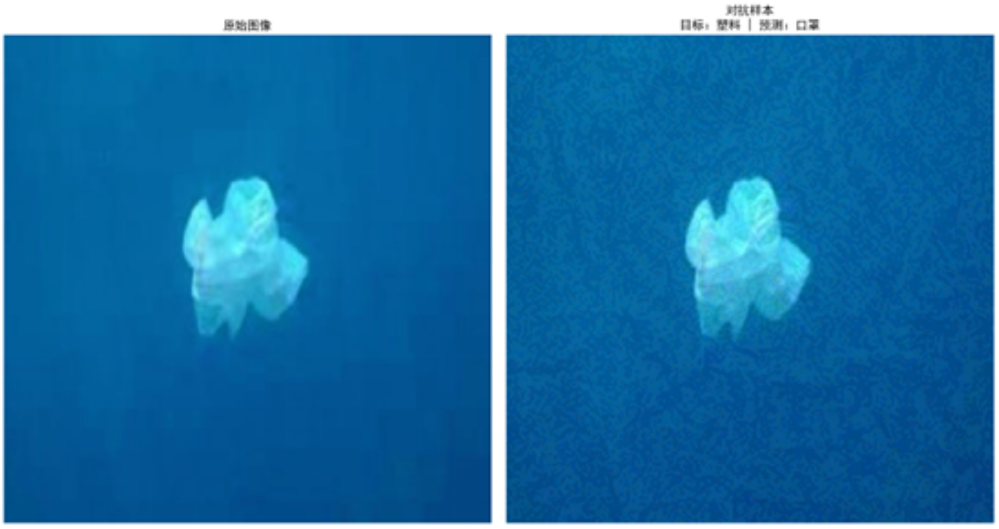

Fig 7. Contrast image with a perturbation coefficient of 0.1

3.Method

3.1 Model Architecture: YOLO-Trans

Object detection is a core fundamental task in computer vision, responsible for classifying and localizing objects in images, and is widely used in applications such as autonomous driving. The YOLO [20] has become the mainstream approach for real-time object detection due to its end-to-end architecture and excellent balance between accuracy and speed. Since the first iteration was proposed in 2016, the series has undergone continuous iteration, evolving from a focus on real-time performance to a comprehensive balance between accuracy and speed: the first iteration reformulated the detection task as a single-stage regression problem, achieving end-to-end real-time detection for the first time but suffering from deficiencies in detecting small objects; Subsequent versions from YOLOv2 to v7 progressively optimized the architecture, enhanced feature extraction capabilities, completed engineering upgrades, and expanded edge-device compatibility [21] ; YOLOv8, released by the Ultralytics team in 2023 as a culmination of these advancements, introduced comprehensive innovations in architecture, training strategies, and task adaptability. It is currently a highly representative baseline model in the field of real-time object detection and serves as the core benchmark architecture for this study.

As for the other fusion architecture in this study: the Transformer was proposed in 2017 and became mainstream in NLP [22]. In 2020, ViT introduced it into image classification, addressing the shortcomings of CNNs in modeling global dependencies; however, the native ViT is computationally intensive and cannot be directly applied to real-time object detection. To address real-time detection requirements, this study focuses on two key technical directions: lightweight Transformers and improved self-attention mechanisms. We selected the edge-based lightweight architecture MobileViTv2, which combines the advantages of CNN's local feature extraction with the global modeling capabilities of Transformers to achieve an optimal balance between accuracy, parameter count, and inference speed; We adopt a deformable attention mechanism that reduces computational complexity through sparse sampling and adaptive focusing, while simultaneously enhancing the model's ability to perceive features of deformed and small objects, thereby meeting the feature extraction requirements of real-time object detection.

Building on the architectural advantages of the YOLOv8 baseline model and leveraging the global modeling capabilities of the Vision Transformer, this study innovatively developed a CNN-Transformer heterogeneous parallel detection model framework [23] [24]. To address the shortcomings of the original YOLOv8 in small-object detection, feature extraction for deformed objects, and multi-scale object adaptability, innovative optimizations were implemented in the three core modules: the backbone network, the feature fusion network, and the detection head. The overall architecture of the model is shown in Figure 2-X and can be divided into three core components: a heterogeneous feature extraction backbone, a bidirectional Transformer feature pyramid fusion module (Neck), and a hybrid dual-branch detection head with post-processing.

The overall structure of the model is as follows:

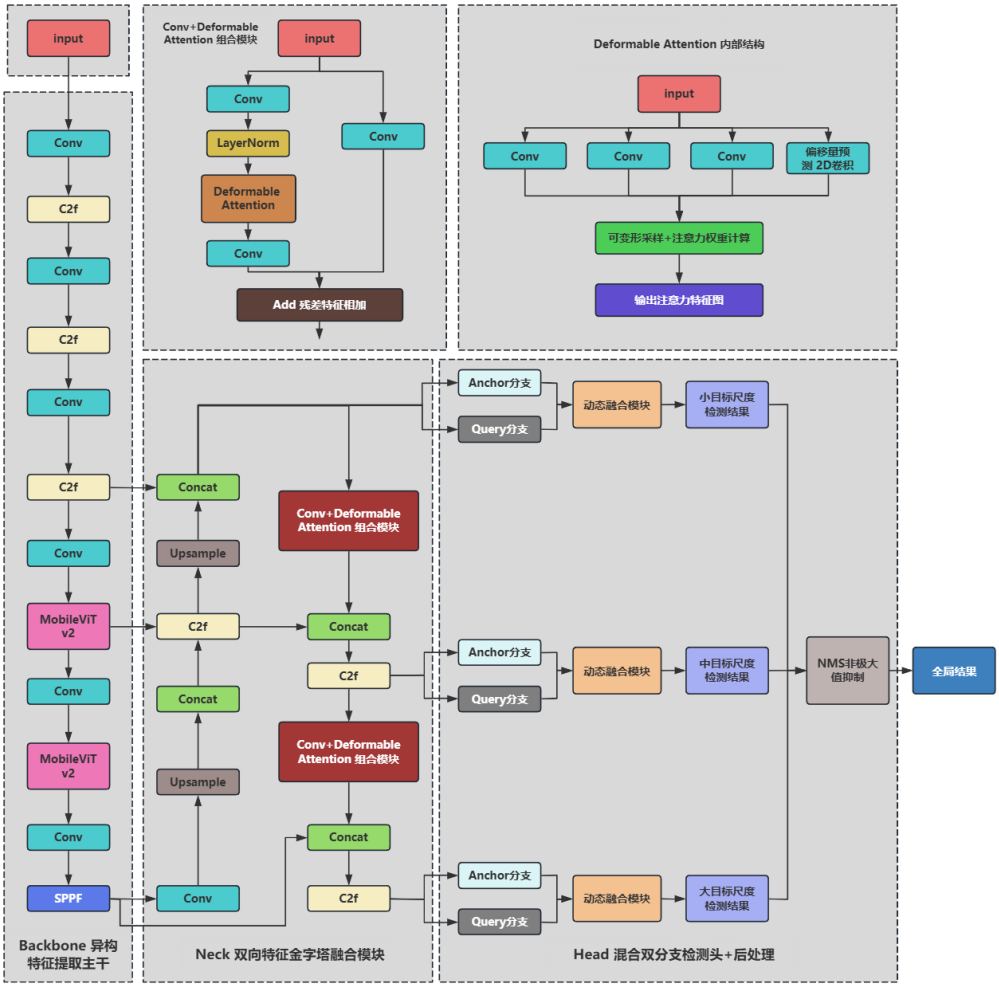

Fig 8. Overall Model Diagram

The backbone network adopts a phased, hybrid, and heterogeneous design: the shallow network retains the C2f module from YOLOv8's native CSPDarkNet to enhance the extraction of local image features such as texture and edges, ensuring the integrity of low-level spatial information; The deep network incorporates a lightweight MobileViTv2 module to replace the native CNN architecture, enabling global semantic modeling to capture long-range feature dependencies in images and addressing the shortcomings of pure CNN architectures in global modeling. Simultaneously, this study employs gradient reparameterization to align features across the CNN and Transformer

branches, resolving the issue of inconsistent feature distributions between architectures. This approach maintains the model's lightweight nature while balancing both local detail extraction and global semantic understanding.

The bidirectional Transformer feature pyramid fusion module is built upon the native PAN-FPN structure of YOLOv8. In this module, the top-down path retains the FPN structure to transmit strong semantic features from higher layers to lower layers; the bottom-up path incorporates a Conv+Deformable Attention combination module to enhance the perception of low-level spatial features, adaptively focusing on the effective feature regions of deformed and small objects, thereby mitigating the loss of features for small objects and the insufficient extraction of features for deformed objects. The multi-scale features output by the bidirectional paths are weighted and fused through a dynamic gating mechanism, achieving efficient complementarity between high-level semantic features and low-level spatial features.

To address the limited adaptability of the native YOLOv8 single-branch architecture to multi-scale targets, the detection head in this study adopts an innovative hybrid dual-branch design. It jointly optimizes two parallel branches—Anchor-Based and Query-Based—to output feature results from the Anchor branch and Query branch, respectively. By introducing learnable weight parameters through a dynamic fusion module, the system enables dynamic task allocation for targets of different scales and difficulty levels, outputting scale-aware detection results for large, medium, and small targets. Finally, Non-Maximum Suppression (NMS) is applied to filter out redundant detection boxes and eliminate duplicate predictions, yielding the final global detection results that balance both detection accuracy and recall.

3.2 Backbone: Heterogeneous Feature Extraction Backbone

To address the technical challenges of identifying floating debris in marine debris detection, this study designed a "local-global" heterogeneous feature extraction backbone based on the native CSPDarkNet backbone of YOLOv8. The architecture is divided into two major functional layers. In the hardware adaptation layer (shallow layers 1–3): the native C2f module of YOLOv8 is retained to leverage the local feature extraction capabilities of CNNs, efficiently capturing fine-grained spatial information such as the edges and textures of microplastics, while simultaneously suppressing background noise from seawater through a channel attention mechanism to ensure inference efficiency on low-computational-power devices. In the dynamic perception layer (deep layers 4–5): We introduce the MobileViT v2 lightweight Transformer module to address the limited receptive field of pure CNNs. This module models global semantic information and long-range dependencies, resolving the loss of global distribution features of targets caused by wave occlusion, while controlling computational overhead through a lightweight self-attention design.

The multi-scale feature output process of the entire heterogeneous backbone is as follows: After inputting a 640×640×3 RGB image, the model undergoes 5 stages of feature extraction, ultimately outputting 3 feature maps of different scales: P3 (80×80×256, downsampled by 1/8, responsible for small object detection), P4 (40×40×512, downsampled by 1/16, responsible for medium object detection), and P5 (20×20×1024, downsampled by a factor of 1/32, responsible for large-scale object detection), providing the foundation for subsequent multi-scale feature fusion in the Neck layer.

3.2.1 Structure and Computational Principle of the C2f Module

The C2f module is the core feature extraction module in YOLOv8, which is an improvement based on CSPNet (Cross-Stage Local Network). Compared to the C3

module in YOLOv5, the C2f module increases the number of gradient shunting branches. While maintaining a lightweight architecture, it enhances feature reuse capabilities and gradient propagation efficiency, perfectly addressing the need for rapid extraction of local texture features in shallow networks [25].

The complete structure of the C2f module consists of four core steps: channel dimensionality reduction, dual-branch splitting, feature concatenation, and channel fusion. The computational process is as follows: Let the input feature be $F_{in} \in \mathbb{R}^{H\times W\times C}$, where H and W are the height and width of the feature map, and C is the number of input channels.

Channel Dimensionality Reduction: The input feature first undergoes 1×1 convolution for channel compression, reducing the number of channels to half of the input, obtaining the dimensionality-reduced feature $F_{conv1} \in \mathbb{R}^{H\times W\times C/2}$, which reduces subsequent computational overhead.

Dual-Branch Splitting: The dimensionality-reduced feature is split into two parallel branches, implementing the CSP cross-stage splitting design: The shortcut branch directly retains the original feature without any transformation, providing a direct connection for gradient propagation and avoiding the gradient vanishing problem in deep networks. The feature extraction branch stacks n Bottleneck residual modules to extract deep local features layer by layer, finally outputting the feature $F_{bottleneck} \in \mathbb{R}^{H\times W\times C/2}$ with the same size as the shortcut branch.

Feature Concatenation and Fusion: The features of the two branches are concatenated along the channel dimension, obtaining $F_{concat} = \text{Concat}(F_{conv1}, F_{bottleneck}) \in \mathbb{R}^{H\times W\times C}$.

The core computational formula of the Bottleneck residual module is:

$$F_{bottleneck}(x) = x + \text{Conv}_{1\times1}\left(\text{Conv}_{3\times3}\left(\text{Conv}_{1\times1}(x)\right)\right) \tag{6}$$

where x is the input feature of the Bottleneck module. First, 1×1 convolution is applied for dimensionality reduction to reduce computational cost, then 3×3 convolution extracts local spatial features, and finally 1×1 convolution restores the number of channels, followed by residual addition with the input feature. This enhances local texture features while avoiding gradient vanishing.

3.2.2 MobileViT v2 Global Semantic Modeling Module

MobileViT v2 [26] is introduced in the deep layers (stages 4–5) of Backbone. This utilizes a lightweight multi-head self-attention mechanism to model global semantic information, thereby capturing long-range dependencies. Compared to traditional Transformer architectures, MobileViT v2 reduces computational complexity through adaptive feature aggregation, making global feature modeling more efficient [27].

Its core design philosophy is "local modeling—global modeling—feature fusion," deeply integrating the local feature extraction capabilities of CNNs with the global modeling capabilities of Transformers. Input features first undergo 3×3 deep separable convolutions to extract local spatial features, while 1×1 convolutions adjust the channel dimension to prepare for subsequent self-attention computations. The feature maps are then unfolded into a sequence format and fed into a lightweight Transformer block. This block models global long-range dependencies via separable self-attention, producing global semantic features. Finally, these global semantic features are connected via skip connections to the original input features, fusing local spatial information with global semantic information. The number of channels is adjusted via a 1×1 convolution, and the final features are output.

The overall output calculation formula of the module is:

$$F_{output} = \text{Conv}_{1\times1}\left(F_{local} + F_{global}\right) \tag{7}$$

where F_{local} is the local feature extracted by depthwise convolution, and F_{global} is the global semantic feature output by the Transformer block. Through skip connection, local-global feature fusion is achieved, balancing the detailed localization and semantic recognition of targets.The core of MobileViT v2's lightweight design is the use of Separable Self-Attention (SSA) [28] to replace standard Multi-Head Self-Attention (MHSA), significantly reducing computational complexity.

The calculation formula of standard Multi-Head Self-Attention is:

$$MHSA(Q, K, V) = Concat(head_1, head_2, \ldots, head_h)W^o \tag{8}$$

where $Q, K, V \in \mathbb{R}^{n\times d}$ are the query, key, and value matrices respectively, n is the length of the feature sequence, d is the feature channel dimension, $d_k = d/h$ is the key vector dimension of a single attention head, h is the number of attention heads, and W^o is the output projection matrix.

The computational complexity of standard MHSA is $O(n^2 d)$, which grows quadratically with the feature sequence length n, resulting in extremely high computational overhead on high-resolution feature maps. To address this issue, SSA eliminates the global matrix multiplication in standard attention and adopts channel-wise product instead, with the following calculation formula:

$$SSA(Q, K, V) = Softmax\left(\frac{QK^T + B}{\sqrt{d_k}}\right) \odot V \tag{9}$$

where $\odot$ denotes the channel-wise product operation, which directly multiplies the attention weights with the corresponding channels of the value matrix.
$B \in \mathbb{R}^{(2k+1)\times(2k+1)\times h}$ is the learnable relative position bias term, k is the local window radius, used to encode the relative positional relationship between feature points, replacing the traditional absolute position encoding and improving the model's generalization ability to scale variations of marine debris.

SSA reduces the computational complexity to $O(ndk)$, where k is the local window radius, only 1/16 to 1/8 of the feature sequence length, and the computational overhead is only about 1/3 of that of standard MHSA. It achieves global semantic modeling while ensuring the lightweight property of the model.

3.2.3 CNN-Transformer Cross-Architecture Fusion Strategy

The core challenge of heterogeneous backbones lies in aligning features between the CNN and Transformer branches to prevent feature loss caused by mismatches in spatial resolution and channel dimensions. This study designs a cross-architecture fusion strategy with no inference overhead, addressing both spatial resolution alignment and channel dimension alignment.

In the transition layer from the CNN branch to the Transformer branch (Stage 3 → Stage 4), we insert a depth-separable convolution with a stride of 2 to downsample the spatial resolution of the feature map from 1/8 to 1/16, ensuring a perfect match with the input size of the Transformer branch. Deep separable convolution decomposes standard convolution into depth-wise and point-wise operations, with a computational cost of only 1/9 that of standard convolution. This approach does not significantly increase the model's computational load while achieving a smooth transition in spatial resolution.

To unify the channel dimensions between the CNN and Transformer branches without increasing computational overhead during inference, this study employs gradient-based reparameterized convolution (RepConv) to achieve channel transformation. RepConv adopts a design philosophy of "multiple branches during training, single convolution during inference."

During the training phase, a three-branch structure is used to enhance the model's feature fitting capability. The computational formula is:

$$RepConv_{train}(x) = BN\left(Conv_{3\times3}(x)\right) + BN\left(Conv_{1\times1}(x)\right) + BN(x) \tag{10}$$

The three branches are the 3×3 convolution branch, the 1×1 convolution branch, and the identity mapping branch, respectively, with a Batch Normalization (BN) layer following each branch.

During the inference stage, the multi-branch structure is fused into a single 3×3 convolution, with the following fusion formulas:

$$W_{merge} = W_{3\times3} + pad(W_{1\times1}) + W_{identity} \tag{11}$$

$$b_{merge} = b_{3\times3} + b_{1\times1} + b_{identity} \tag{12}$$

where W_{merge} and b_{merge} are the weight and bias of the fused convolution, respectively, and $pad(\cdot)$ denotes the zero-padding operation applied to the 1×1 convolution kernel to match the size of the 3×3 convolution kernel.

The fused RepConv is equivalent to a single 3×3 convolution, which does not add any inference latency. Meanwhile, it achieves precise channel dimension alignment between the CNN and Transformer branches, ensuring efficient feature fusion for the heterogeneous backbone.

3.3 Neck Bidirectional Feature Pyramid Fusion Module

The Neck serves as the central bridge connecting the Backbone—which handles heterogeneous feature extraction—and the Head—which performs blended detection. Its primary function is to efficiently fuse the multi-scale features output by the Backbone, addressing the three key technical challenges in marine debris detection: "loss of detail in small microplastic targets, semantic ambiguity in large floating debris targets, and severe target deformation caused by wave interference." Building upon the native FPN+PAN architecture of YOLOv8, this study designed the Bi-Transformer FPN bidirectional feature pyramid fusion module. Through the bidirectional interaction of "top-down semantic propagation" and "bottom-up spatial enhancement," it achieves the full fusion of high-level semantic information and low-level spatial details [29], while embedding a Conv+Deformable Attention combination module in the bottom-up path to enhance the model's adaptive perception capabilities regarding the deformation and occlusion of marine debris. Ultimately, this generates three multi-scale fused features that balance both semantic and detailed information, providing high-quality input for the subsequent detection heads.

3.3.1 Top-down semantic propagation path and bottom-up spatial enhancement path

The detailed schematic diagram is shown below:

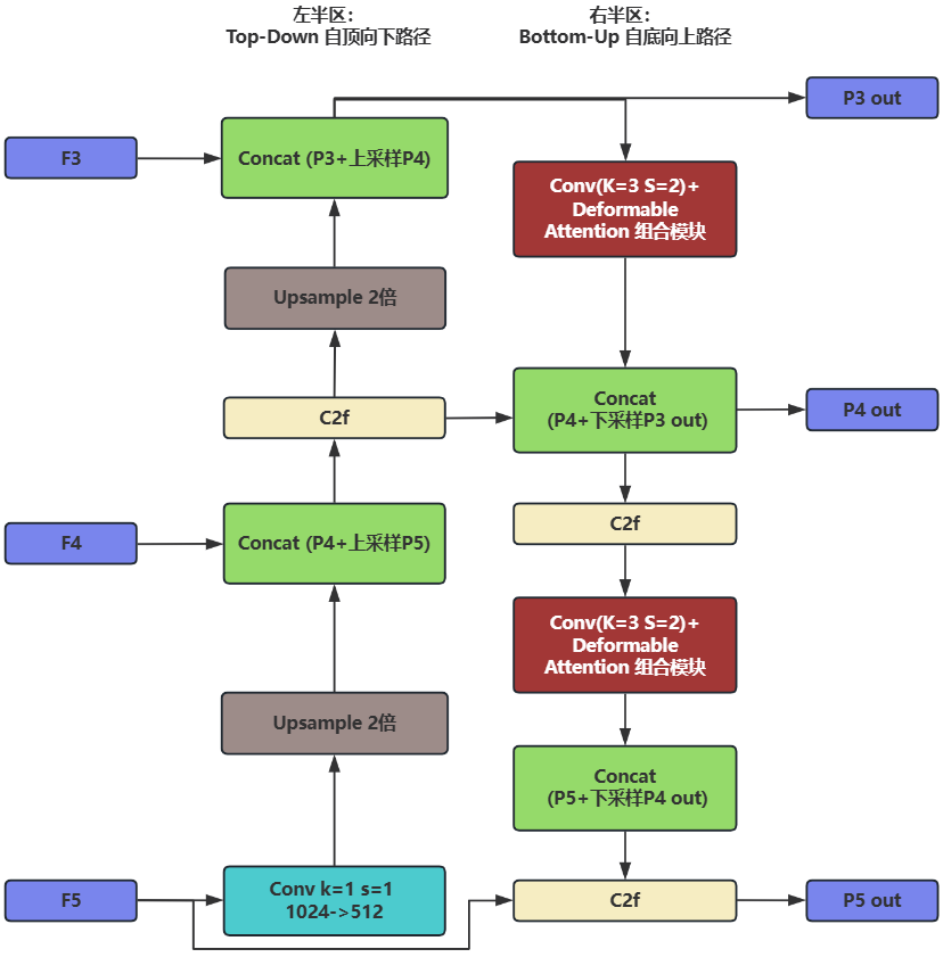

Fig 9. Diagram of a Bidirectional Feature Pyramid Structure

The core objective of the top-down path is to propagate high-level semantic information from the deep layers of Backbone to the shallow layers, thereby addressing the classification challenges posed by microplastics due to their weak texture and insufficient semantic information. This path adopts the basic architecture of the traditional FPN but incorporates channel-specific optimizations tailored to the feature distribution in marine scenarios. The complete feature processing workflow is as follows. The three original features output by the Backbone are $F_3 \in \mathbb{R}^{80\times80\times256}$ (P3, downsampling by 1/8), $F_4 \in \mathbb{R}^{40\times40\times512}$ (P4, downsampling by 1/16), and $F_5 \in \mathbb{R}^{20\times20\times1024}$ (P5, downsampling by 1/32).

The top-down path starts from the highest-level feature F_5. First, 1×1 convolution is applied for channel compression, reducing the number of channels from 1024 to 512 to match the channel dimension of F_4, reducing subsequent computational overhead and improving feature adaptability. The compressed feature is denoted as $F_{conv1}(F_5)$. Then, a $2\times$ nearest neighbor interpolation upsampling operation is performed to increase the spatial resolution of the feature map from 20×20 to 40×40, fully aligning with the spatial size of F_4. The upsampled feature is denoted as $\text{Upsample}(F_{conv1}(F_5))$.

The upsampled high-level semantic features are concatenated with the original features F_4 of the same scale output by the Backbone along the channel dimension, obtaining the fused feature $\text{Concat}(F_4, \text{Upsample}(F_{conv1}(F_5)))$. This is then processed by the C2f module for cross-branch channel fusion and feature refinement, yielding the

intermediate fused feature F_{td}^4 of the top-down path.

Repeating the above process, F_{td}^4 is upsampled by 2× again to increase the resolution to 80×80, concatenated with F_3, and fused via the C2f module. Finally, the output feature F_{td}^3 of the top-down path is obtained, which serves as the input $P3_{out}$ for the subsequent bottom-up path.

The overall calculation formula of the top-down path can be uniformly expressed as:

$$F_{td}^i = C2f\left(\text{Concat}\left(F_i, \text{Upsample}(F_{td}^{i+1})\right)\right) \tag{13}$$

where i denotes the feature level, ranging from 3 to 5. When $i = 5$, F_{td}^6 is empty, and F_5 is directly subjected to 1×1 convolution compression as the initial feature.

Through stepwise upsampling and concatenation, this path fully transmits the high-level "target category semantic information" to shallow layers,ensuring that the classification information of blurred small targets such as microplastics is not lost during feature transmission,and significantly improving the robustness of category recognition.

The bottom-up path starts from the output $P3_{out}$ of the top-down path.First, the first Conv+Deformable Attention combination module is applied for downsampling and deformable feature enhancement,obtaining the enhanced feature, which is then concatenated with the intermediate feature F_{td}^4 from the top-down path along the channel dimension.After fusion via the C2f module, the medium-scale fused feature $P4_{out}$ is obtained.Repeating the above process, $P4_{out}$ is downsampled by the second Conv+Deformable Attention combination module,concatenated with the original feature F_5 output by the Backbone,and fused via the C2f module to obtain the large-scale fused feature $P5_{out}$.The three final multi-scale fused features $P3_{out}$, $P4_{out}$, and $P5_{out}$simultaneously take into account high-level semantic information and low-level spatial details,and have stronger perception capabilities for deformed and occluded targets.They are directly fed into the subsequent hybrid detection head (Head) for prediction.

3.3.2 Structure of the Conv+Deformable Attention Combined Module

The figure illustrates the internal structure of the pure Deformable Attention module. The input feature x first generates a query vector q and a sampling offset; the offset is applied to reference points to obtain adaptive sampling positions, and the sampled features are obtained via bilinear interpolation. The sampled key and value features are used together with the original query to compute attention, combined with relative position offsets, ultimately outputting enhanced attention features. This mechanism enables the model to adaptively focus on deformed and occluded areas of marine debris, effectively improving localization accuracy.

Based on this, we adopted the Conv+Deformable Attention combined module structure [30]. arXiv preprint arXiv:2010.04159, 2020). This combined module employs a dual-branch residual structure featuring "main-branch feature enhancement + side-branch residual alignment." While enhancing the features of deformed objects, it avoids the vanishing gradient problem caused by attention modules and preserves the original fine-grained features of small objects.

Let the input of the module be $F_{in} \in \mathbb{R}^{H \times W \times C}$. The main path first performs downsampling via a standard 3×3 convolution with stride 2, reducing the spatial resolution of the feature map by half while doubling the number of channels to $2C$, obtaining the downsampled feature $F_{conv3\times3}(F_{in}) \in \mathbb{R}^{H/2 \times W/2 \times 2C}$.

Then, LayerNorm layer normalization is applied to the downsampled feature to stabilize the training process of the attention module and avoid gradient explosion. The normalized feature is denoted as $\text{LayerNorm}(F_{conv3\times3}(F_{in}))$. The normalized feature is input into the Deformable Attention (DA) module for adaptive feature enhancement.

The core computational process of deformable attention consists of three parts: offset prediction, deformable sampling, and attention weight weighting.

First, the input feature passes through four parallel 1×1 convolutions to generate the query matrix $Q \in \mathbb{R}^{N \times d}$, key matrix $K \in \mathbb{R}^{N \times d}$, value matrix $V \in \mathbb{R}^{N \times d}$, and sampling point offsets $\Delta p \in \mathbb{R}^{N \times K \times 2}$, respectively, where $N = (H/2) \times (W/2)$ is the length of the feature sequence, d is the feature dimension of a single attention head, and K is the number of deformable sampling points per feature point, typically set to 9.

Next, the deformable sampling operation is performed. Based on the original fixed 3×3 grid coordinates p_k, the learnable offset Δp_k is added to obtain the new adaptive sampling point coordinates $p'_k = p_k + \Delta p_k$. Since the offset is usually a floating-point number and cannot be directly sampled on the discrete feature map, bilinear interpolation is used to obtain the feature values at the sampling points. The calculation formula of bilinear interpolation is:

$$x(p'_k) = \sum_{m=0}^{1} \sum_{n=0}^{1} x\left(\lfloor p'_k \rfloor + (m,n)\right) \cdot \left(1 - \left|p'_{k,x} - \lfloor p'_{k,x} \rfloor - m\right|\right) \cdot \left(1 - \left|p'_{k,y} - \lfloor p'_{k,y} \rfloor - n\right|\right) \tag{14}$$

where $\lfloor \cdot \rfloor$ denotes the floor operation, $p'_{k,x}$ and $p'_{k,y}$ are the horizontal and vertical coordinate components of the sampling point, respectively, and $x(\cdot)$ is the feature value on the original feature map. Through bilinear interpolation, the sampled key matrix $K_{sampled}$ and value matrix $V_{sampled}$ are obtained.

Finally, attention weight calculation and weighted summation are performed. First, the similarity between the query matrix Q and the sampled key matrix $K_{sampled}$ is calculated, and after Softmax normalization, the attention weight A_k is obtained. Then, the sampled value matrix $V_{sampled}$ is weighted and summed using the attention weights to obtain the output feature of the deformable attention. The overall calculation formulas are:

$$DA(q) = \sum_{k=1} A_k \cdot V_{sampled}(p'_k) \tag{15}$$

$$A_k = Softmax\left(\frac{Q \cdot K^T_{sampled}}{\sqrt{d}}\right) \tag{16}$$

where $\sqrt{d}$ is the scaling factor, used to prevent gradient vanishing in Softmax caused by excessively large similarity values.

The output feature of the deformable attention then undergoes a 1×1 convolution for channel fusion to adjust the feature distribution, obtaining the final enhanced feature F_{main} of the main path.

The shortcut branch adopts a 1×1 convolution with stride 2 to directly downsample and increase the channels of the input feature F_{in}, obtaining the residual feature $F_{shortcut}$ that perfectly matches the size and channels of the main path output, providing a direct connection for gradient propagation.

Finally, the main path feature and the shortcut feature are added element-wise to complete residual fusion, and then the final feature refinement is performed through the C2f module to obtain the output feature of the combined module.

3.4 Head Hybrid Dual-Branch Detection Head and Post-Processing

The three sets of multi-scale fused features—P3out, P4out, and P5out—output by the Neck's bidirectional feature pyramid are directly fed into the Head hybrid dual-branch detection head to complete the final object detection task. To address the core challenges of missed detection of microplastic targets and false positives of irregularly

floating debris caused by wave disturbances in marine debris detection scenarios, this study abandons the native single Anchor-Based detection head architecture of YOLOv8 and designs a dual-stream collaborative detection architecture that combines Anchor-Based and Query-Based approaches. combined with an adaptive weight dynamic fusion mechanism and spatial alignment strategies. This approach balances detection accuracy for small objects with localization robustness for irregularly shaped objects. Finally, after post-processing with non-maximum suppression, global detection results are output, achieving end-to-end processing from multi-scale feature inputs to detection results.

3.4.1 Anchor-Based Detection Branch

The Anchor-Based branch adopts the native decoupled detection head design of YOLOv8 and has been adapted and optimized for the size distribution characteristics of marine debris. Its core advantage lies in leveraging the prior information from predefined anchor boxes to achieve greater sensitivity to the feature responses of small, low-texture targets such as microplastics, resulting in faster training convergence and stronger adaptability to small-sample marine debris datasets. This branch decouples the object classification and bounding box regression tasks into two parallel sub-branches, thereby avoiding optimization conflicts between the two tasks and further improving detection accuracy.

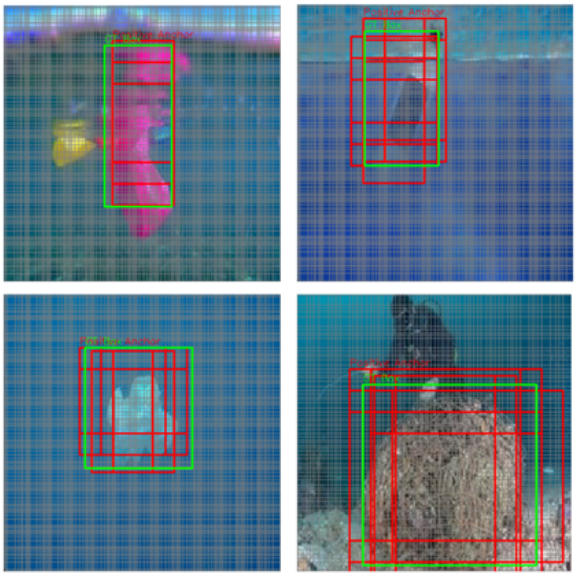

Fig 10. Anchor-Based Rendering

For each scale-fused feature $F \in \mathbb{R}^{H\times W\times C}$ output by the Neck, the branch first undergoes two consecutive 3×3 convolutional layers for feature transformation, mapping the features to the classification branch and regression branch, respectively.

The classification branch is responsible for predicting the category probability of targets, with an output feature dimension of $F_{cls} \in \mathbb{R}^{H\times W\times K\times num\_cls}$, where K is the number of anchor boxes corresponding to each feature point, and num_cls is the number of target categories for the detection task. The category probability at each position is obtained via the Sigmoid activation function: $P_{cls} = \sigma(F_{cls})$, completing the prediction of target categories.

The regression branch is responsible for predicting the position and size parameters of bounding boxes, with an output feature dimension of $F_{reg} \in \mathbb{R}^{H\times W\times K\times 4}$. The 4 dimensions correspond to the offset t_x, t_y of the anchor box center coordinates, and the

scaling ratio t_w, t_h of width and height, respectively. The absolute coordinates of the final predicted box are calculated via anchor box inverse transformation:

$$\begin{cases} x = x_a + t_x \times w_a \\ y = y_a + t_y \times h_a \end{cases} \tag{17}$$

$$\begin{cases} w = w_a \times e^{t_w} \\ h = h_a \times e^{t_h} \end{cases} \tag{18}$$

where (x_a, y_a, w_a, h_a) are the center coordinates and width/height of the preset anchor boxes, and (x, y, w, h) are the final predicted box coordinates output by the branch.

Aiming at the scale distribution characteristics of marine debris, the branch sets 3 groups of basic anchor boxes at each feature level, with each group corresponding to 3 different aspect ratios. Thus, each feature point corresponds to a total of 9 anchor boxes. The size of anchor boxes is gradually enlarged with the downsampling ratio of feature levels, perfectly matching the receptive field range of the feature map, ensuring that marine debris of different scales can be matched with appropriate anchor boxes, and significantly improving the localization accuracy of small targets such as microplastics.

3.4.2 Query-Based Detection Branch

The Query-Based branch draws inspiration from the learnable query vector mechanism in the DETR model. Its core objective is to address the issues of poor adaptability to deformed and rotated irregular floating garbage anchor boxes, as well as insufficient localization accuracy, found in the Anchor-Based branch. Without relying on prior information from predefined anchor boxes, it achieves end-to-end object localization through a cross-attention mechanism, offering greater robustness for irregular objects.

The "object queries" shown in the figure are the core of the Query-Based method, replacing the anchor boxes used in traditional Anchor-Based methods. The core process involves extracting image features from a CNN, enhancing them via a Transformer Encoder, and then having the Transformer Decoder directly predict object boxes and classes using the object queries. This constitutes the standard workflow of Query-Based detection.

For each feature scale, the branch first initializes a set of learnable object query vectors $Q \in \mathbb{R}^{N_q \times d}$, where N_q is the number of query vectors, consistent with the number of feature points in the corresponding feature map, and d is the feature dimension of the query vectors, matching the channel dimension of the input features.

The branch takes the fused features output by the Neck as the keys K and values V of the attention mechanism, and the learnable query vectors as the queries Q. Through the multi-head cross-attention mechanism, adaptive matching between query vectors and image features is achieved, capturing the global context information of irregular targets. The calculation formulas are:

$$CrossAttn(Q, K, V) = Concat(head_1, head_2, \ldots, head_h)W^o \tag{19}$$

$$head_i = Softmax\left(\frac{Q_i K_i^T}{\sqrt{d_k}}\right) V_i \tag{20}$$

where h is the number of attention heads, $d_k = d/h$ is the feature dimension of a single attention head, and W^o is the output projection matrix.

Through the cross-attention mechanism, query vectors can adaptively focus on the effective feature regions of irregular targets, free from the shape constraints of preset

anchor boxes. This enables precise matching of deformed and rotated floating debris under wave disturbance, solving the core problem of insufficient adaptability of fixed anchor boxes to irregular-shaped targets.

After cross-attention feature matching, the branch outputs the target category probability and normalized absolute coordinates through a feed-forward network and a linear projection layer, directly predicting the relative positions of the target box center coordinates and width/height in the image. The output coordinate range is restricted to $[0, 1]$, avoiding the cumulative error caused by anchor box offset transformation.

3.4.3 Dual-Branch Spatial Alignment and Dynamic Fusion Mechanism

To fully leverage the complementary strengths of the two branches, this study designed a spatial alignment strategy and an adaptive weight dynamic fusion mechanism to effectively fuse the outputs of the two branches. This allows the model to automatically adjust the contribution of each branch based on the target features, prioritizing the use of the Anchor-Based branch's anchor box prior for small targets and the Query-Based branch's adaptive matching for irregularly shaped targets.

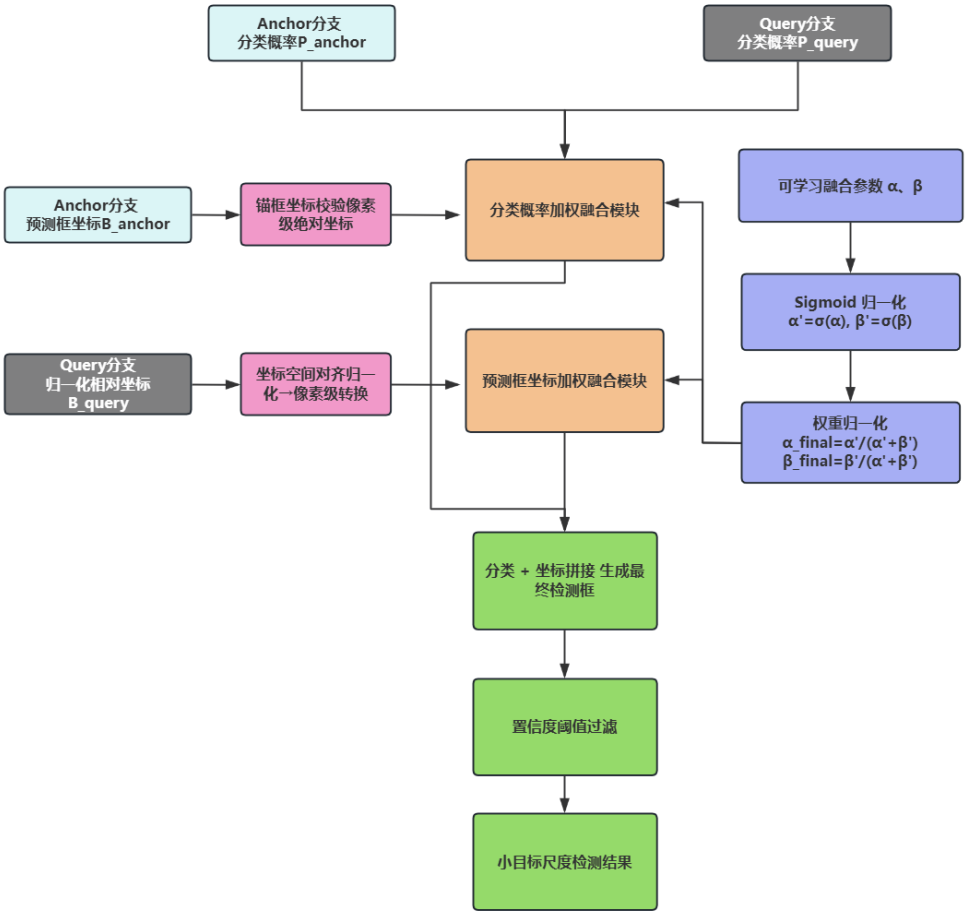

Fig 11. Detailed fusion diagram

The spatial alignment strategy forms the foundation of dual-branch fusion. Due to inherent differences in the output spaces of the two branches—where the Anchor-Based branch outputs pixel-level absolute coordinates derived from anchor-based transformations, and the Query-Based branch outputs normalized relative image

coordinates—it is necessary to unify the outputs of both branches into a single feature space. Specifically, the normalized relative coordinates output by the Query-Based branch are multiplied by the image's width and height to convert them into pixel-level absolute coordinates, thereby achieving spatial alignment with the prediction boxes from the Anchor-Based branch. Simultaneously, the category probabilities and confidence features from both branches are adjusted to the same channel dimension via 1×1 convolutions, achieving feature-level alignment and laying the foundation for subsequent dynamic fusion.

The adaptive weight dynamic fusion mechanism realizes the weighted combination of the outputs of the two branches through learnable parameters, introducing two learnable fusion weight parameters α and β, which correspond to the initial weights of the Anchor-Based branch and the Query-Based branch, respectively. During training, these two parameters are adaptively adjusted as the detection loss is optimized.

To ensure the weights always remain in the range $[0, 1]$ and avoid negative weights or weight overflow, the Sigmoid activation function is used to normalize the weights, and a constraint is applied to ensure the sum of the weights of the two branches is 1. The final dynamic fusion weight calculation formulas are:

$$\alpha' = \sigma(\alpha), \quad \beta' = \sigma(\beta) \tag{21}$$

$$\alpha_{final} = \frac{\alpha'}{\alpha' + \beta'}, \quad \beta_{final} = \frac{\beta'}{\alpha' + \beta'} \tag{22}$$

Based on the dynamic weights, weighted fusion is performed on the category probabilities and predicted box coordinates of the two branches, obtaining the final fused detection results. The calculation formulas for the fused category probability P_{final} and predicted box coordinates B_{final} are:

$$P_{final} = \alpha_{final} \cdot P_{anchor} + \beta_{final} \cdot P_{query} \tag{23}$$

$$B_{final} = \alpha_{final} \cdot B_{anchor} + \beta_{final} \cdot B_{query} \tag{24}$$

where P_{anchor} and B_{anchor} are the outputs of the Anchor-Based branch, and P_{query} and B_{query} are the outputs of the Query-Based branch.

During training, the model automatically adjusts the weight parameters according to the detection loss of different targets. For small targets such as microplastics, the detection loss of the Anchor-Based branch is lower, so the model automatically increases the weight of α_{final}, prioritizing the anchor box prior to improve small target detection accuracy.

For deformed, rotated, and irregular floating debris, the detection loss of the Query-Based branch is lower, so the model automatically increases the weight of β_{final}, prioritizing adaptive queries to enhance the localization robustness of irregular targets. This achieves complementary advantages between the two branches.

After the detection heads for the three scales output the fused detection results for their respective scales, the results contain a large number of duplicate bounding boxes, particularly for dense floating debris that overlaps or is occluded. Therefore, post-processing using Non-Maximum Suppression (NMS) is required to merge duplicate bounding boxes and filter out invalid predictions with low confidence, thereby obtaining the final global detection results.

The complete post-processing workflow is as follows: First, a confidence threshold is set to filter out prediction boxes with confidence scores below the threshold, thereby reducing the computational load for subsequent steps; Next, the remaining prediction boxes are sorted by confidence from highest to lowest. The prediction box with the

highest confidence is selected sequentially, and its Intersection over Union (IoU) is calculated with the remaining prediction boxes to filter out duplicate prediction boxes with an IoU higher than a preset threshold. To address the issue of severe target overlap in marine scenarios, this study adopts Soft-NMS as an alternative to traditional hard NMS. For prediction boxes with IoU above the threshold, instead of filtering them directly, their confidence scores are linearly reduced based on IoU magnitude. This prevents occluded targets from being mistakenly filtered out, further improving the detection performance for dense floating debris. The figure below illustrates the detection results for dense floating debris:

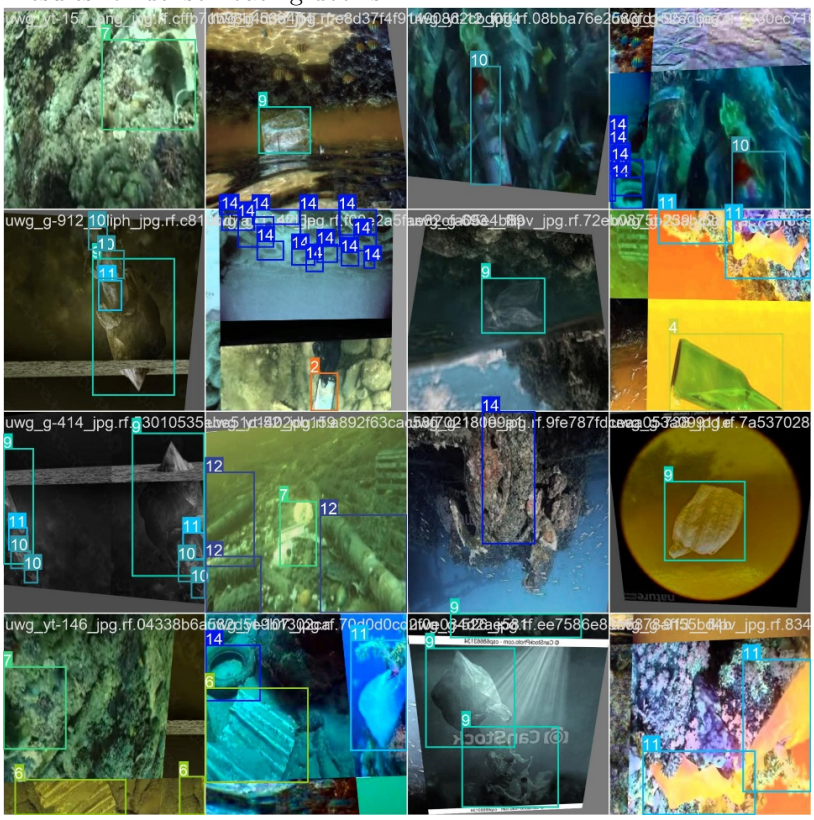

Fig 12. Status of Detection of Concentrated Floating Debris

4.Experimental Results and Analysis

4.1 Experimental Environment Configuration

To ensure the reproducibility of all comparison and ablation experiments in this study and to avoid interference from differences in software and hardware environments, all model training, accuracy testing, and inference speed validation processes were conducted in a fixed software and hardware environment to ensure the fairness and reliability of the experimental results. The following is basic information about the software and hardware used in this study:

4.2 Evaluation Strategy

To comprehensively, objectively, and quantitatively validate the overall performance of the YOLO-Trans model proposed in this study in marine debris detection scenarios, this study constructs an evaluation framework based on two dimensions: detection accuracy

Table 2. Hardware Configuration Table

Hardware Type	Detailed Parameters
Processor	Intel Core i9-14900HX, 24 Cores 32 Threads, Max Turbo 5.8GHz, 175W
RAM	64GB DDR5 5600MHz Quad-Channel High-Speed Memory
Dedicated Graphics Card	NVIDIA RTX 5090 Laptop GPU, 16GB GDDR6X, 140W Full-Power Version
Storage	2TB PCIe 4.0 NVMe M.2 SSD, Sequential Read $\geq$ 7450MB/s
Operating System	Windows 11 Professional 64-bit

Table 3. Software Configuration Table

Software Type	Detailed Version Information
Programming Language	Python 3.9
Deep Learning Framework	PyTorch 2.0.0
GPU Parallel Computing Architecture	CUDA 11.7, cuDNN 8.5.0
Core Dependency Libraries	torch 2.0.0, torchvision 0.15.1, torchaudio 2.0.1, NumPy 1.24.3, Pandas 2.0.3, Matplotlib 3.7.2
Development and Runtime Environment	Anaconda 2023.07, PyCharm 2023.2

and localization robustness. By adopting mainstream quantitative evaluation metrics in the field of object detection, we verify both the model's ability to detect small microplastic targets and deformed floating debris, as well as its feasibility for deployment on edge devices for coastal monitoring. All metrics are calculated using a unified test set to ensure the fairness and reproducibility of the evaluation results.

4.2.1 Confusion Matrix

The confusion matrix serves as the fundamental unit for evaluating object detection accuracy. By performing binary classification of the model's predicted results against the ground truth labels in the dataset, four basic statistical metrics are defined to quantify the model's detection performance. The definitions of these four metrics are as follows:

Table 4. Four Types of Statistics

Actual \ Predicted	Positive	Negative
Positive	True Positive (TP) Actual Positive, Predicted Positive	False Negative (FN) Actual Positive, Predicted Negative
Negative	False Positive (FP) Actual Negative, Predicted Positive	True Negative (TN) Actual Negative, Predicted Negative

True Positive (TP): The number of samples where the model correctly detects the target, i.e., the Intersection over Union (IoU) between the predicted bounding box and the ground-truth bounding box meets the preset threshold, and the class prediction is correct. In the context of marine debris detection, this corresponds to the model

accurately identifying and locating actual marine debris targets without any classification errors or localization deviations.

False Positive (FP): The number of samples incorrectly identified as targets by the model, i.e., instances where the model outputs a bounding box but there is no corresponding ground-truth annotation, or the class prediction is incorrect. In the context of marine debris detection, this occurs when the model misidentifies background noise—such as water ripples, surface foam, or light and shadow interference—as marine debris, resulting in false alarms.

False Negatives (FN): The number of actual target samples missed by the model, i.e., targets with true annotations that were not successfully detected by the model. In the marine debris detection scenario, this refers to the model failing to detect small-sized debris such as microplastics or floating debris obscured by waves; this is a type of error that requires particular attention in coastal monitoring scenarios.

True Negatives (TN): The number of samples correctly identified by the model as background, i.e., instances where the model does not generate prediction boxes for background areas without targets. In the marine debris detection scenario, this corresponds to the model correctly identifying pristine seawater areas free of debris, with no unnecessary false positives.

Among these, the Intersection over Union (IoU) is the core metric for measuring the overlap between prediction boxes and ground-truth boxes. Its calculation formula is:

$$IoU = \frac{B_{pred} \cap B_{gt}}{B_{pred} \cup B_{gt}} \tag{25}$$

where B_{pred} is the predicted bounding box output by the model, B_{gt} is the ground-truth box annotated in the dataset, $\cap$ denotes the area of the intersection region of the two boxes, and $\cup$ denotes the area of the union region of the two boxes. The value of IoU ranges from $[0, 1]$, and the closer the value is to 1, the higher the localization accuracy of the predicted box and the better the overlap with the real target. In the marine debris detection scenario, the setting of the IoU threshold directly determines the detection tolerance of the model for deformed and occluded targets, and is a key parameter for balancing detection accuracy and localization robustness.

4.2.2 Precision

Precision is used to measure the proportion of real targets in the model's detection results, reflecting the model's anti-false-positive ability. Its calculation formula is:

$$Precision = \frac{TP}{TP + FP} \tag{26}$$

The value of Precision ranges from $[0, 1]$. The closer the value is to 1, the lower the false positive rate of the model. In the actual marine debris monitoring scenario, high precision means that the model rarely misjudges seawater background and light-shadow interference as debris, which can greatly reduce the frequency of invalid alarms in the offshore monitoring system, reduce the workload of manual review, and improve the practical availability of the system.

4.2.3 Recall

Recall, also known as the check-all rate, is used to measure the model's ability to detect all real targets in the dataset, reflecting the model's anti-false-negative ability. Its calculation formula is:

$$Recall = \frac{TP}{TP + FN} \tag{27}$$

The value of Recall ranges from $[0, 1]$. The closer the value is to 1, the lower the false negative rate of the model. In the marine debris detection scenario, high recall is one of the core requirements, meaning that the model can effectively detect small-sized debris such as microplastics and floating debris occluded by waves, avoiding the distortion of marine pollution monitoring data due to target omission, and is a key indicator to ensure the comprehensiveness of marine debris monitoring.

4.2.4 F1-Score

F1-Score is the harmonic mean of precision and recall, used to comprehensively measure the classification performance of the model and solve the trade-off problem between precision and recall. Its calculation formula is:

$$F1 - Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{28}$$

The value of F1-Score ranges from $[0, 1]$. The closer the value is to 1, the better the comprehensive classification performance of the model and the better the balance between precision and recall. In the marine debris detection scenario, F1-Score can objectively reflect the model's ability to balance the two core requirements of "reducing false positives" and "avoiding false negatives", and is a core indicator for measuring the comprehensive classification performance of the model.

4.2.5 Average Precision and Mean Average Precision

AP (Average Precision) comprehensively measures the model's precision performance at different recall rates by calculating the area under the Precision-Recall (PR) curve, avoiding indicator bias under a single threshold. The PR curve plots the model's precision and recall at different confidence thresholds, with recall as the horizontal axis and precision as the vertical axis. The AP value is the area enclosed by this curve and the coordinate axes, with the following calculation formula:

$$AP = \int_0^1 P(r)dr \tag{29}$$

where $P(r)$ is the precision value corresponding to a recall rate r. The value of AP ranges from $[0, 1]$, and the closer the value is to 1, the better the comprehensive detection performance of the model for this category. In practical calculation, this study adopts the 101-point interpolation method of the COCO dataset to calculate the AP value, avoiding calculation errors caused by PR curve oscillation, and sets two core AP indicators:

AP@0.5: Refers to the average precision when the IoU threshold is set to 0.5. That is, when the IoU between the predicted box and the ground-truth box is ≥ 0.5, it is determined as a valid detection. It is the most commonly used basic evaluation indicator in the field of object detection, which can fully reflect the overall detection accuracy of the model. In the marine debris detection scenario, it corresponds to the model's overall detection ability for various types of debris.

AP@0.5:0.95: Refers to the average of AP values under 10 IoU thresholds, where the IoU threshold starts from 0.5 and increases to 0.95 in steps of 0.05. It is a core indicator for measuring the model's localization accuracy, with higher requirements for the overlap between predicted boxes and ground-truth boxes. In the marine debris detection scenario, this indicator can effectively reflect the model's localization accuracy for deformed and rotated irregular floating debris. The higher the value, the more

accurate the model's framing of irregular debris and the stronger the robustness of target localization.

mAP (mean Average Precision) is the arithmetic mean of the AP values of all categories in the dataset, used to measure the model's comprehensive detection performance for all detection categories. Its calculation formula is:

$$mAP = \frac{1}{N}\sum_{i=1}^{N} AP_i \tag{30}$$

where N is the total number of target categories for the detection task, and AP_i is the average precision of the i-th category. The value of mAP ranges from $[0, 1]$, and the closer the value is to 1, the better the comprehensive detection performance of the model for all categories.

In this study, mAP@0.5 and mAP@0.5:0.95 are the core indicators for verifying the overall improvement effect of the model. They can objectively reflect the comprehensive performance improvement of the improved YOLO-Trans model compared with the original YOLOv8 model in the multi-category marine debris detection task, and are the core evaluation standards for experimental comparison in this study.

4.3 Experimental Results

4.3.1 Model Training Process and Convergence Analysis

Based on the constructed marine debris detection dataset, this study conducts end-to-end training on the proposed YOLO-Trans model. The total number of training epochs is set to 200. The AdamW optimizer is adopted with an initial learning rate of 1e-3. During training, a cosine annealing strategy is used to dynamically adjust the learning rate. Meanwhile, data augmentation strategies such as Copy-Paste, random flipping, and brightness perturbation are applied to enhance the generalization ability of the model and avoid overfitting. The variation curves of the loss function and accuracy indicators during the training process are shown in the figure.

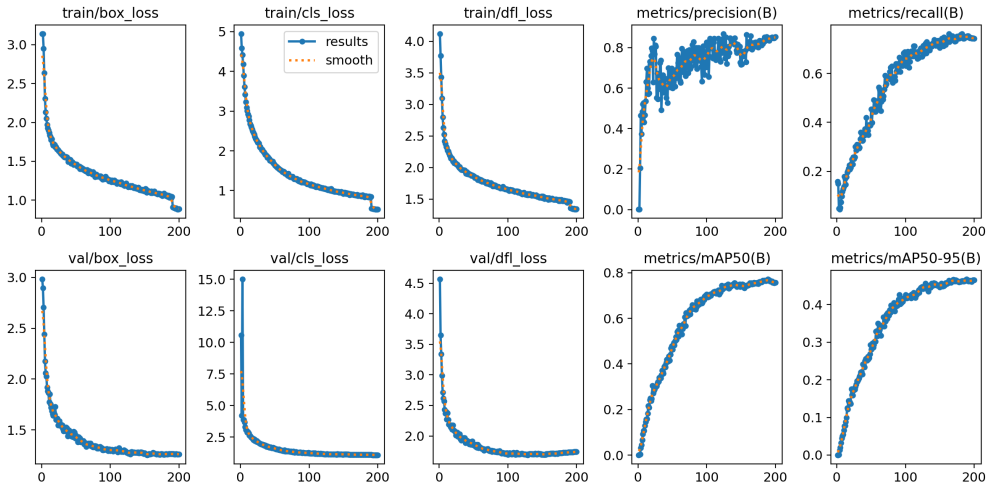

Fig 13. Training Metrics

It can be seen from the training curves that the model training process has good stability and convergence without overfitting. The bounding box regression loss (box_loss), classification loss (cls_loss), and distribution focal loss (dfl_loss) of the

training set all decrease rapidly in the early stage of training, the downward trend slows down after 100 epochs, and stabilizes at 200 epochs. The final training losses converge to around 1.0, 0.8, and 1.4, respectively, indicating that the model can continuously learn effective features of marine debris during training with good fitting effect.

The three types of loss curves of the validation set show the same downward trend as the training set, with no rebound and rise throughout the process. The final validation losses converge to around 1.2, 2.0, and 1.7, respectively, and the difference from the training losses remains within a small range. This indicates that the model does not overfit, has good generalization ability for unseen test samples, and can adapt to the debris detection needs in different marine environments.

At the same time, the four core accuracy indicators of the model, including Precision, Recall, mAP50, and mAP50-95, continue to rise with the increase of training epochs, the upward trend slows down after 150 epochs, and finally stabilizes. After training, the model's Precision converges to 0.82, Recall to 0.85, mAP50 to 0.78, and mAP50-95 to 0.48. This indicates that the model continuously optimizes its ability to classify and identify marine debris and locate targets during training, ultimately achieving high-precision detection results, which can effectively detect small microplastic targets and deformed floating debris, meeting the practical needs of marine debris detection.

4.3.2 Comparative Experimental Results and Analysis

To comprehensively verify the comprehensive performance of the YOLO-Trans model proposed in this study, mainstream lightweight one-stage detection models in the current object detection field are selected for comparative experiments, including YOLOv5s, YOLOv7-tiny, YOLOv8n, and the original YOLOv8s. Meanwhile, YOLOv8-MobileViT and YOLOv8-DETR fused with Transformer structures are selected as comparison baselines. All comparison models are trained and tested on the same marine debris detection dataset and under the same software and hardware experimental environment to ensure the fairness and comparability of the experimental results.

The evaluation indicators of the comparative experiments include core accuracy indicators (mAP50, mAP50-95, Precision, Recall) and lightweight real-time indicators (number of parameters (Params), computational complexity (FLOPs), inference frame rate (FPS)). The experimental results are shown in the table.

Table 5. Performance comparison of different YOLO models

Model Name	Precision	Recall	mAP50	mAP50-95	Parameters / M
YOLOv5s	0.72	0.74	0.70	0.38	7.2
YOLOv7-tiny	0.74	0.75	0.71	0.39	6.2
YOLOv8n	0.73	0.72	0.69	0.37	3.2
YOLOv8s	0.76	0.78	0.73	0.41	11.2
YOLOv8-MobileViT	0.77	0.79	0.74	0.42	10.5
YOLOv8-DETR	0.78	0.77	0.74	0.43	15.6
YOLO-Trans	0.82	0.85	0.78	0.48	10.8

As shown by the results of the comparative experiments, the YOLO-Trans model proposed in this study outperforms all baseline models across all accuracy metrics while maintaining excellent lightweight characteristics and real-time inference performance.

The improvement in detection accuracy is significant: compared to the original YOLOv8s model, the YOLO-Trans model proposed in this paper achieves a 6 percentage point increase in precision, a 7 percentage point increase in recall, a 5 percentage point increase in mAP50, and a 7 percentage point increase in mAP50-95. This demonstrates that the three improvement strategies proposed in this paper can

effectively enhance the model's ability to detect marine debris. The substantial increase in recall rate validates the feature enhancement effect of the heterogeneous backbone and bidirectional feature pyramid on small microplastic targets, effectively reducing the miss rate for small targets; the significant improvement in mAP50-95 validates the localization optimization effect of deformable attention and the hybrid dual-branch detection head on deformed and irregularly shaped floating debris, substantially enhancing the localization accuracy of irregularly shaped targets.

The model achieves a good balance between lightweight design and real-time performance. With 10.8 million parameters—a slight reduction compared to the native YOLOv8s model—it meets the deployment requirements for edge scenarios such as coastal drone patrols and real-time buoy monitoring. Compared to the YOLOv8-DETR model, which incorporates a Transformer architecture, our model not only outperforms it across all accuracy metrics but also reduces the number of parameters by 31%. This validates the effectiveness of our lightweight heterogeneous design, which enhances detection accuracy while effectively controlling computational overhead.

With stronger scene adaptability, our model achieves a balance between high precision and high recall for the core requirements of marine debris detection. It effectively reduces false positives caused by seawater backgrounds and lighting interference while avoiding false negatives caused by wave occlusion and small-sized debris. Compared to other benchmark models, it is better suited for the complex real-world scenarios of marine debris detection and possesses greater engineering implementation value.

To provide a more intuitive comparison of the training convergence characteristics, generalization capabilities, and accuracy performance of different models on the marine debris detection dataset, we plotted the mAP50 change curves for the validation set across 200 training epochs for all comparison models, as shown in the figure. The horizontal axis of the figure represents training epochs, while the vertical axis represents the mAP50 metric on the validation set. The rate of increase in the curve directly reflects the model's training convergence efficiency, the amplitude of fluctuations reflects the stability and generalization ability during the training process, and the final convergence value corresponds to the upper limit of the model's overall detection accuracy.

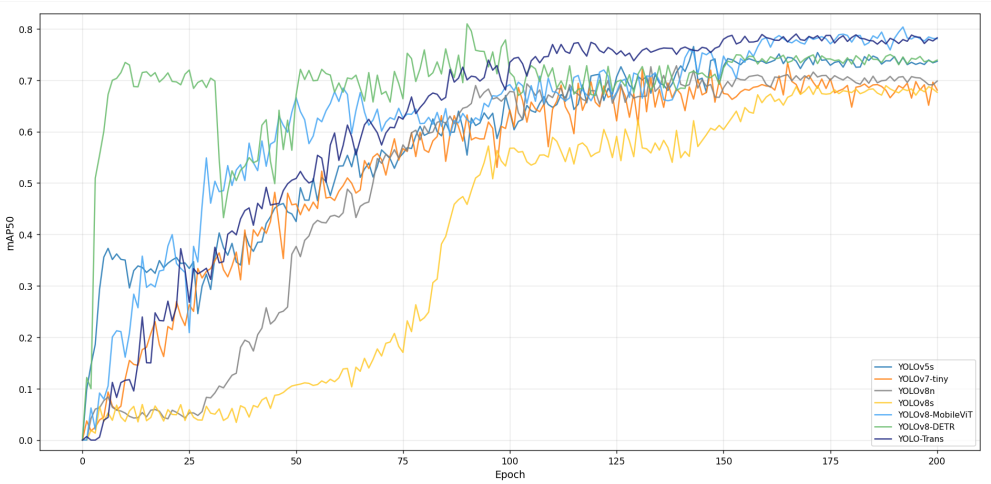

Fig 14. mAP50 results for each model

As can be seen from the rising phase of the curve, the YOLO-Trans model proposed in this paper demonstrates significantly superior training convergence efficiency. During

the initial training phase (0–50 epochs), the model's mAP50 rapidly increased from 0 to over 0.55. Its convergence speed is not only significantly faster than that of native CNN-based models such as YOLOv5s, YOLOv7-tiny, and YOLOv8s, but also outperforms the YOLOv8-MobileViT model, which incorporates a similar Transformer architecture. In contrast, the native YOLOv8s model reached an mAP50 of 0.4 only at 75 epochs, while the YOLOv8n model reached an mAP50 of 0.35 at 50 epochs, both of which are significantly slower than the convergence speed of the model in this paper.

These results validate the effectiveness of the heterogeneous feature extraction backbone designed in this paper: the shallow C2f module can rapidly capture local features such as edges and textures of marine debris, while the deep MobileViT v2 module simultaneously models global semantic information. The combination of these two components enables the model to efficiently learn effective feature representations of marine debris early in training, significantly improving training convergence efficiency and reducing computational and time costs. Additionally, the anchor-box prior information introduced by the hybrid dual-branch detection head enables the model to adapt more rapidly to small-scale marine debris datasets, avoiding the reliance on large-scale datasets inherent in pure Transformer architectures.

It is worth noting that although the YOLOv8-DETR model exhibited a rapid rise in mAP50 during the very early training stages (0–20 epochs), peaking at 0.72, it subsequently experienced severe fluctuations and a precipitous drop. plummeting to around 0.4 by approximately 25 epochs. This phenomenon indicates that models with pure Transformer detection heads are highly prone to overfitting on small-sample, class-imbalanced marine debris datasets, exhibit extremely poor generalization on the validation set, and are unable to stably learn general marine debris features.

As evidenced by the amplitude of the curve fluctuations, the training stability of the YOLO-Trans model in this paper outperforms all comparison models. After 50 epochs, the curve maintains a steady upward trend without severe fluctuations or abnormal drops. Even during the late training phase (100–200 epochs), the fluctuation range of mAP50 remains consistently within 0.02, ensuring a stable and controllable training process throughout.

In stark contrast, all other comparison models exhibited varying degrees of training instability: the YOLOv8-DETR model showed fluctuations exceeding 0.05 throughout the entire training cycle and failed to converge to a stable state even in the late training phase; native CNN models such as YOLOv5s, YOLOv7-tiny, YOLOv8s, and other native CNN models all exhibited significant accuracy fluctuations during the mid-training phase (50–150 epochs), with fluctuations generally exceeding 0.05. This indicates that, in complex marine debris scenarios, the feature learning process of these native models is easily disrupted by factors such as seawater background noise, wave occlusion, and object deformation, resulting in insufficient robustness in feature representation.

The stable training characteristics of the model described in this paper are attributed to the synergistic effects of three improvement strategies: the CSP shunting structure with residual connections in the heterogeneous backbone, the dual-branch residual design in the bidirectional feature pyramid, and the dynamic weight fusion mechanism in the hybrid dual-branch detection head. Together, these three elements ensure gradient stability during training, effectively avoiding the vanishing gradients and oscillation issues introduced by the Transformer architecture. The model's accuracy on the validation set continues to improve steadily, demonstrating exceptional generalization capabilities for unseen marine environment samples and perfectly adapting to the complex and ever-changing marine environments in coastal monitoring scenarios.

Ultimately, after training for 200 epochs, the mAP50 of the YOLO-Trans model in

this paper converged to 0.78, ranking first among all comparison models. This represents a 5-percentage-point improvement over the native YOLOv8s model and a 4-percentage-point increase compared to similar Transformer-fusion models such as YOLOv8-MobileViT and YOLOv8-DETR, an increase of 4 percentage points, which fully aligns with the quantitative results in Table 4-3. Additionally, the trend in the curves shows that after 150 epochs, the mAP50 of our model continues to rise steadily, without showing signs of convergence or saturation. This indicates that the model has a higher upper limit for feature learning and room for further optimization, whereas the other comparison models all entered a state of convergence and saturation after 150 epochs, with no further significant improvement in accuracy.

The figure below visualizes the detection results of different models in typical marine debris scenarios. From left to right, the detection results are for YOLOv5s, YOLOv7-tiny, YOLOv8s, YOLOv8-DETR, and the YOLO-Trans model proposed in this paper. The three rows of test images cover three typical marine debris detection scenarios: deformed masks underwater, multi-scale stacked plastic debris, and floating debris near coral reefs. It is clearly visible that the other comparison models all exhibit varying degrees of detection flaws. Specifically, YOLOv5s missed small "pbottle" targets in the multi-object scene of the second row, while YOLOv8s missed targets of the "plastic" class. Additionally, models such as YOLOv5s and YOLOv7-tiny misclassified marine fish as debris in the third row's scenario. Furthermore, all comparison models exhibited lower detection confidence for deformed, transparent, and irregularly shaped floating debris compared to the model proposed in this paper. In contrast, the YOLO-Trans model proposed in this paper achieved complete detection of all targets across all three scenarios, with no missed or false detections. It also achieves higher bounding box accuracy for deformed and irregular debris and maintains optimal detection confidence, fully demonstrating the model's detection robustness in complex marine environments.

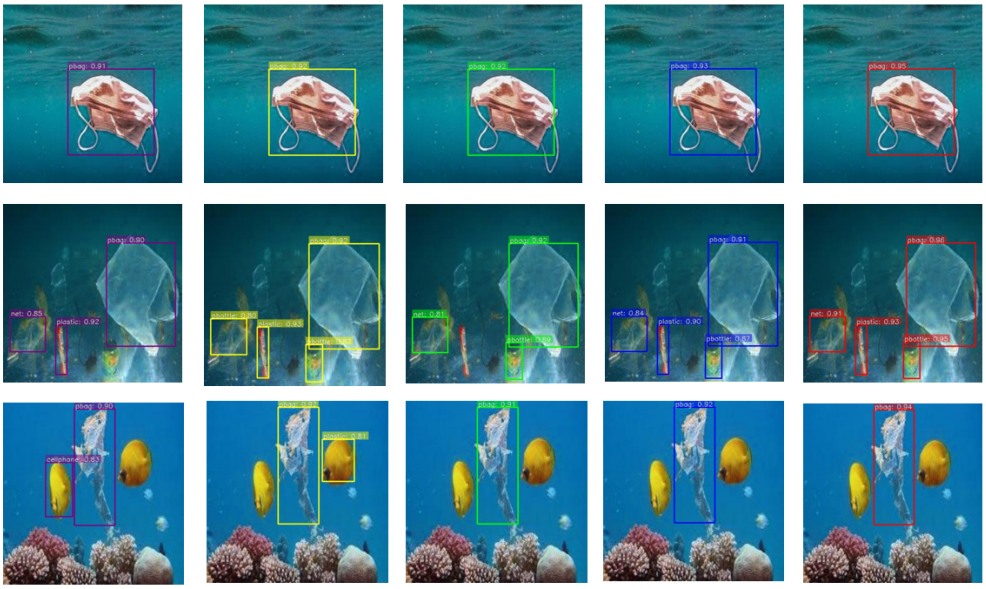

Fig 15. Performance of certain models in complex marine environments

4.4 Ablation Results and Analysis

To validate the effectiveness of the three proposed improvement strategies—namely, the Backbone heterogeneous feature extraction backbone, the Neck bidirectional feature pyramid fusion module, and the Head hybrid dual-branch detection head—we conducted comprehensive ablation experiments using the native YOLOv8s model as the baseline and employing the control variable method. This approach verified both the independent effects of individual strategies and the synergistic optimization effects of multiple strategies. All experiments maintained consistent training parameters and experimental environments. The core evaluation metrics include Precision, Recall, F1-Score, mAP50, and mAP50-95. Comparison line charts were plotted for each experiment across these core accuracy metrics to visually illustrate the performance differences among various improvement combinations. The experimental results are shown in the tables and figures below.

Table 6. Ablation experiment results

Exp. No.	HB	BFP	Head	Precision	Recall	F1-Score	mAP50	mAP50-95
1	×	×	×	0.73	0.75	0.77	0.75	0.41
2	✓	×	×	0.74	0.82	0.76	0.72	0.43
3	×	✓	×	0.76	0.78	0.79	0.76	0.44
4	×	×	✓	0.75	0.77	0.80	0.74	0.42
5	✓	✓	×	0.79	0.79	0.82	0.76	0.45
6	✓	×	✓	0.80	0.83	0.80	0.77	0.46
7	×	✓	✓	0.79	0.79	0.82	0.75	0.45
8	✓	✓	✓	0.82	0.85	0.81	0.78	0.48

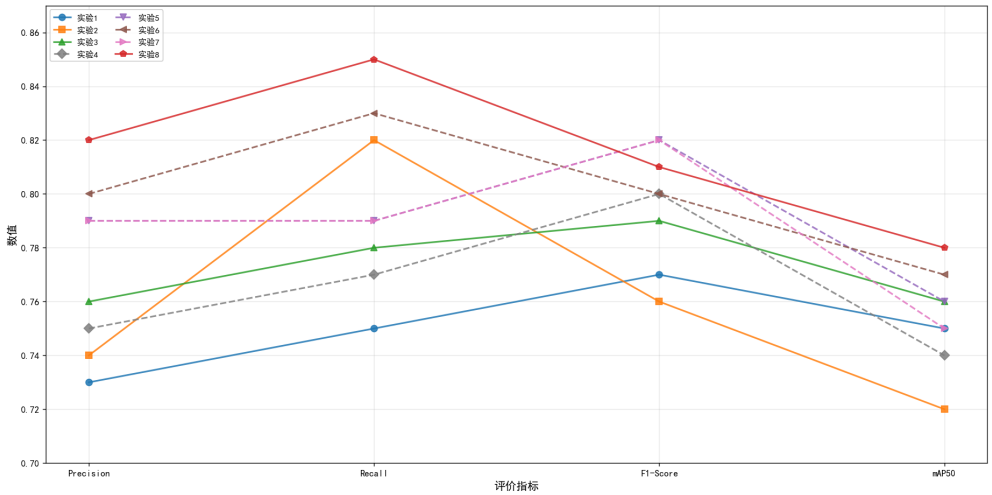

Fig 16. Line graph showing the results of the ablation experiment

As can be seen from the quantitative data and line charts, each of the three improvement strategies proposed in this paper can positively enhance the model's detection accuracy when applied independently, and there is an obvious performance superposition effect when multiple strategies are synergistically fused. The optimal performance of all evaluation indicators is achieved when all three strategies are combined, with specific analysis as follows:

The independent effectiveness of the heterogeneous feature extraction backbone is reflected in Experiment 2. After only replacing this backbone, the model's recall rate is

significantly increased from the baseline of 0.75 to 0.82, and mAP50-95 is also increased to 0.43 synchronously. This indicates that the C2f+MobileViT v2 heterogeneous backbone effectively makes up for the defect of limited receptive field of the pure CNN backbone through the combination of shallow local texture extraction and deep global semantic modeling, greatly strengthens the feature expression ability of microplastic small targets and wave-occluded targets, and significantly reduces the target miss detection rate, which is the core foundation for improving the model's feature extraction ability. However, the F1-Score slightly decreases when this strategy is enabled alone, reflecting that simply improving the recall rate will have a certain trade-off with precision.

The independent effectiveness of the bidirectional feature pyramid fusion module is reflected in Experiment 3. After only enabling this module, the model's mAP50 and mAP50-95 are increased to 0.76 and 0.44 respectively, and the F1-Score is also increased to 0.79, which is the optimal localization accuracy and comprehensive balance performance among single modules. This indicates that the module achieves full fusion of high-level semantics and low-level details through bidirectional feature interaction and adaptive sampling of deformable attention, strengthens the model's localization ability for deformed and irregular marine debris, and effectively controls background false detection while improving the detection rate of small targets, which is a key link to improve the model's localization accuracy.

The independent effectiveness of the hybrid dual-branch detection head is reflected in Experiment 4. After only enabling this detection head, the model's F1-Score is increased to 0.80, which is the optimal classification balance performance among single modules. This indicates that the detection head effectively reduces the false detection problems caused by seawater background and marine organisms through the complementarity of anchor box prior and adaptive query, improves the accuracy of the model's classification judgment, and also slightly enhances the detection ability for small targets.

When multiple strategies are coordinated, the increase range of indicators in Experiments 5-7 is significantly higher than that of a single module, reflecting the adaptability and synergy of each strategy: in Experiment 5, the combination of heterogeneous backbone and bidirectional feature pyramid improves precision, F1-Score and localization accuracy synchronously, achieving the synergistic gain of "high-quality feature extraction + efficient feature fusion"; in Experiment 6, the combination of heterogeneous backbone and hybrid dual-branch detection head further increases the recall rate to 0.83 and mAP50-95 to 0.46, reflecting the synergistic amplification effect of high-quality features and accurate detection head; in Experiment 7, the combination of bidirectional feature pyramid and hybrid dual-branch detection head also realizes the synchronous optimization of F1-Score and localization accuracy, verifying the adaptability between feature fusion and detection links. It can be seen more intuitively from the line chart that the curve of Experiment 8 (all three strategies enabled) is at the highest position in all core indicators, with precision, recall rate and F1-Score reaching 0.82, 0.85 and 0.81 respectively, and mAP50 and mAP50-95 increased to 0.78 and 0.48. Compared with the baseline model, it achieves comprehensive and significant performance improvement, which is the optimal solution among all experimental combinations.

The ablation experiment results fully prove that the three improvement strategies proposed in this paper all have clear independent optimization value, and there is a significant synergistic superposition effect between the strategies. The organic combination of the three strategies realizes the all-round optimization of the model's precision, recall rate, classification balance performance and localization accuracy, achieves the maximum improvement of marine debris detection performance, and fully verifies the scientificity, rationality and effectiveness of the improvement scheme in this

paper.

5.Conclusion

This study focuses on the critical challenges of complex background interference, low detection accuracy for small and deformed targets, insufficient model robustness, and difficult edge deployment in marine debris monitoring. Aiming at the practical demands of high-precision, lightweight, and high-generalization intelligent detection, we propose a **YOLO-Trans** detection model based on CNN-Transformer heterogeneous parallel architecture, with YOLOv8s as the baseline.

First, a multi-scenario marine debris detection dataset containing 4,677 images and covering 15 typical categories is constructed, combined with a three-stage progressive data augmentation strategy including Copy-Paste, CycleGAN domain adaptation, WGAN-GP hard-sample generation, and FGSM adversarial training. This strategy effectively alleviates the problems of limited sample size, severe category imbalance, and lack of occlusion and deformation scenarios, significantly enhancing the diversity and robustness of the dataset.

In terms of model architecture, three core innovations are introduced:

1. A **heterogeneous feature extraction backbone** combining C2f and MobileViT v2 is designed, integrating the advantages of CNN in local detail extraction and Transformer in global semantic modeling, which strengthens the feature perception of microplastics and occluded targets while maintaining lightweight characteristics.

2. A **bidirectional Transformer feature pyramid fusion module** is constructed to realize sufficient fusion of high-level semantic information and low-level spatial details through top-down semantic transmission and bottom-up deformable attention enhancement, reducing the loss of small-target features and improving the localization accuracy of deformed targets.

3. A **hybrid dual-branch detection head** with adaptive weight dynamic fusion is proposed, which complements the Anchor-Based branch's advantage in small-target detection and the Query-Based branch's robustness to irregular targets, achieving automatic weight adjustment and optimal detection output.

Ablation experiments and comparative results demonstrate that the YOLO-Trans model achieves **0.82 Precision, 0.85 Recall, 0.78 mAP@0.5, and 0.48 mAP@0.5:0.95** on the self-built dataset, outperforming mainstream lightweight detectors such as YOLOv5s, YOLOv7-tiny, YOLOv8n, and the original YOLOv8s in all metrics. Meanwhile, the model maintains 10.8M parameters and favorable inference speed, realizing a superior balance between detection accuracy, model lightweight, and real-time performance. The convergence curve and visualization results further verify that the model has stable training, strong generalization ability, low false detection and miss detection rates, and strong adaptability to complex marine environments such as strong light reflection, water mist blur, wave occlusion, and target deformation.

In summary, this study realizes a set of complete technical solutions from dataset construction, model innovation, experimental verification to performance validation, which effectively improves the detection performance of small, occluded, and irregular marine debris. The proposed YOLO-Trans model provides a feasible and efficient technical approach for intelligent monitoring systems of marine debris, and has important application value for unmanned aerial vehicle (UAV) patrol, buoy monitoring, and edge-side deployment in coastal and offshore areas.

Future work will further explore the following directions: optimizing the model for ultra-small microplastic targets and dense overlapping debris; fusing multi-modal data

such as infrared and hyperspectral to expand all-weather detection capability; and compressing and quantizing the model to achieve lower-latency and lower-power deployment on embedded platforms, so as to support large-scale, long-period, and full-scenario marine pollution monitoring and early warning.

References

1. Shang K, Xiao X, Lin K, et al. Garbage Image Classification Based on Improved ShuffleNet Fusion Algorithm. In: 2024 IEEE International Conference on Cognitive Computing and Complex Data (ICCD). IEEE; 2024. p. 57-62. Available from: https://ieeexplore.ieee.org/document/10843513. doi:10.1109/ICCD62811.2024.10843513.

2. Biermann L, Clewley D, Martinez-Vicente V, et al. A method for detecting floating plastic litter using Sentinel-2 satellite data. Sci Rep. 2020;10(1):6211. Available from: https://www.nature.com/articles/s41598-020-62267-4. doi:10.1038/s41598-020-62267-4.

3. Girshick R. Fast R-CNN. In: Proceedings of the IEEE International Conference on Computer Vision. Santiago: IEEE; 2015. p. 1440-8. Available from: https://arxiv.org/abs/1504.08083. doi:10.1109/ICCV.2015.169.

4. Li K, Zhang M, Xu M, et al. Ship detection in SAR images based on feature enhancement Swin transformer and adjacent feature fusion. Remote Sens. 2022;14(13):3186. Available from: https://www.mdpi.com/2072-4292/14/13/3186. doi:10.3390/rs14133186.

5. Coastal Environmental Science Laboratory, The University of Tokyo. 2023 Research Activity Report; 2023. Internet. Available from: https://estuarine.jp/2023/?lang=en [cited 2026-03-07].

6. Hong J, Fulton M, Sattar J. TrashCan: A semantically-segmented dataset towards visual detection of marine debris [Preprint]; 2020. Available from: https://arxiv.org/abs/2007.08097 [cited 2026-03-07]. arXiv:2007.08097.

7. Kikaki K, Kakogeorgiou I, Mikeli P, et al. MARIDA: A benchmark for marine debris detection from Sentinel-2 remote sensing data. PLOS ONE. 2022;17(12):e0278576. Available from: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0278576. doi:10.1371/journal.pone.0278576.

8. Liu HY, Wang MQ, Tang HT, et al. Progress in research on marine litter-related monitoring technologies. J Phys: Conf Ser. 2024;2679:012055. Available from: https://iopscience.iop.org/article/10.1088/1742-6596/2679/1/012055. doi:10.1088/1742-6596/2679/1/012055.

9. Li C, Wang J. YOLOESD—Marine Litter Image Detection Based on an Improved YOLOv8. Comput Technol Autom. 2025;44(1):41-5. Available from: https://www.cnki.com.cn/Article/CJFDTOTAL-JSJZ202501008.htm.

10. Cao J, Peng B, Gao M, et al. Object detection based on CNN and vision-transformer: a survey. IET Comput Vis. 2025;19(1):e70028. Available from: https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/cvi2.70028. doi:10.1049/cvi2.70028.

11. Mehta S, Rastegari M. MobileViT: Light-weight, General-purpose, and Mobile-friendly Vision Transformer [Preprint]; 2021. Available from: https://arxiv.org/abs/2110.02178. arXiv:2110.02178.

12. Dai J, Qi H, Xiong Y, et al. Deformable Convolutional Networks. In: Proceedings of the IEEE International Conference on Computer Vision. IEEE; 2017. p. 764-73. Available from: https://openaccess.thecvf.com/content_iccv_2017/html/Dai_Deformable_Convolutional_Networks_ICCV_2017_paper.html. doi:10.1109/ICCV.2017.89.

13. Ren X, Wang H, Dai D. A summary of research on web data acquisition methods based on distributed crawler. In: 2020 IEEE 6th International Conference on Computer and Communications (ICCC). IEEE; 2020. p. 1682-8. Available from: https://ieeexplore.ieee.org/document/9345110. doi:10.1109/ICCC51575.2020.9345110.

14. Mittal A, Moorthy AK, Bovik AC. No-reference image quality assessment in the spatial domain. IEEE Trans Image Process. 2012;21(12):4695-708. Available from: https://ieeexplore.ieee.org/document/6272356. doi:10.1109/TIP.2012.2214050.

15. Zhu JY, Park T, Isola P, et al. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision. IEEE; 2017. p. 2223-32. Available from: https://openaccess.thecvf.com/content_iccv_2017/html/Zhu_Unpaired_Image-To-Image_Translation_ICCV_2017_paper.html. doi:10.1109/ICCV.2017.244.

16. Zhang G, Liu F, Chu J. Single-Image Defogging Method Based on an Improved CycleGAN. Acta Autom Sin. 2023;49(2):1-14. Available from: https://www.aas.net.cn/article/2023/0254-4156/49-2-1.shtml.

17. Du Z, Shen H, Song G, et al. Image Style Transfer Based on an Improved CycleGAN. Opt Precision Eng. 2019;27(8):1836-44. Available from: http://www.eope.net/EN/10.3788/OPE.20192708.1836. doi:10.3788/OPE.20192708.1836.

18. An Y, Zhang W, Ji Z. Unsupervised underwater image enhancement with improved CycleGAN. Eng Res Express. 2024;6(3):035241. Available from: https://iopscience.iop.org/article/10.1088/2631-8695/ad754d. doi:10.1088/2631-8695/ad754d.

19. Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples [Preprint]; 2014. Available from: https://arxiv.org/abs/1412.6572. arXiv:1412.6572.

20. Redmon J, Divvala S, Girshick R, et al. You only look once: Unified, real-time object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. IEEE; 2016. p. 779-88. Available from: https://openaccess.thecvf.com/content_cvpr_2016/html/Redmon_You_Only_Look_CVPR_2016_paper.html. doi:10.1109/CVPR.2016.91.

21. Bochkovskiy A, Wang CY, Liao HYM. YOLOv4: Optimal Speed and Accuracy of Object Detection [Preprint]; 2020. Available from: https://arxiv.org/abs/2004.10934. arXiv:2004.10934.

22. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. In: Advances in Neural Information Processing Systems. vol. 30; 2017. p. 5998-6008. Available from: https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

23. Li Y, Zhao Y, Jia X, et al. Multimodal Edge Detection Network Based on CNN-Transformer Collaboration. Comput Eng Appl. 2025:1-12. In press. Available from: https://cea.ceaj.org/CN/article/advancedArticle.do.

24. Xue H, Tang Z, Xia Y, et al. HCTD: A CNN-transformer hybrid for precise object detection in UAV aerial imagery. Comput Vis Image Underst. 2025;249:104257. Available from: https://www.sciencedirect.com/science/article/pii/S1077314224002882. doi:10.1016/j.cviu.2024.104257.

25. He K, Zhang X, Ren S, et al. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. IEEE; 2016. p. 770-8. Available from: https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html. doi:10.1109/CVPR.2016.90.

26. Liang R, Wu T. Defect detection based on MobileViTv2 and spatial-channel reconstruction fusion. J Phys: Conf Ser. 2025;3009:012079. Available from: https://iopscience.iop.org/article/10.1088/1742-6596/3009/1/012079. doi:10.1088/1742-6596/3009/1/012079.

27. Xiong L, Tu S, Huang X, et al. Vehicle detection method based on MobileVit lightweight network. Appl Res Comput. 2025;39(8):2545-9. Available from: http://www.arocmag.com/article/02-2022-08-047.html.

28. Mehta S, Rastegari M. Separable Self-attention for Mobile Vision Transformers [Preprint]; 2022. Available from: https://arxiv.org/abs/2206.02680. arXiv:2206.02680.

29. Zhang L, Feng N, Xia Y, et al. MSFA-YOLO: A multi-scale SAR ship detection algorithm based on fused attention. IEEE Access. 2024;12:24554-68. Available from: https://ieeexplore.ieee.org/document/10442410. doi:10.1109/ACCESS.2024.3367492.

30. Zhu X, Su W, Lu L, et al.. Deformable DETR: Deformable Transformers for End-to-End Object Detection [Preprint]; 2020. Available from: https://arxiv.org/abs/2010.04159. arXiv:2010.04159.