

Empowering DINO Representations for Underwater Instance Segmentation via Aligner and Prompter

Zhiyang Chen^{1, 2*}, Chen Zhang^{1, 2*}, Hao Fang^{1, 2}, Runmin Cong^{1, 2†}

¹School of Control Science and Engineering, Shandong University, China

²Key Laboratory of Machine Intelligence and System Control, Ministry of Education, Jinan 250061, China
{zhiyangchen, chen.zhang, fanghaok}@mail.sdu.edu.cn, rmcong@sdu.edu.cn

Abstract

Underwater Instance Segmentation (UIS), integrating pixel-level understanding and instance-level discrimination, is a pivotal technology in marine resource exploration and ecological protection. In recent years, large-scale pretrained visual foundation models, exemplified by DINO, have advanced rapidly and demonstrated remarkable performance on complex downstream tasks. In this paper, we demonstrate that DINO can serve as an effective feature learner for UIS, and we introduce *DiveSeg*, a novel framework built upon two insightful components: (1) The *AquaStyle Aligner*, designed to embed underwater color style features into the DINO fine-tuning process, facilitating better adaptation to the underwater domain. (2) The *ObjectPrior Prompter*, which incorporates binary segmentation-based prompts to deliver object-level priors, provides essential guidance for instance segmentation task that requires both object- and instance-level reasoning. We conduct thorough experiments on the popular UIIS and USIS10K datasets, and the results show that *DiveSeg* achieves the state-of-the-art performance.

code — <https://github.com/ettotof/Diveseg>

Introduction

Understanding underwater scenes is critical for advancing marine exploration and sustainable utilization of marine resources, supporting a broad range of applications such as marine research, ecological monitoring, and ocean resource extraction (Tang et al. 2024; Abdullah et al. 2024). Among the key technologies in this domain, *Underwater Instance Segmentation (UIS)* plays a particularly vital role, as it facilitates both pixel-wise classification and instance-level discrimination of underwater objects for accurate recognition and localization. Compared to underwater semantic segmentation (Islam et al. 2020; Zheng et al. 2024), UIS offers significant advantages in distinguishing overlapping objects (e.g., schools of fish and corals) and accurately delineating object boundaries. These capabilities are essential for enhancing the performance of underwater robots or autonomous underwater vehicles in tasks such as obstacle avoidance, object manipulation, autonomous navigation

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(Cong et al. 2021; Christensen et al. 2022) and underwater video segmentation (Fang et al. 2025a,b).

Unlike natural images, underwater imagery presents unique visual characteristics and significant challenges due to light absorption and scattering, color distortion, low contrast, and limited visibility, all of which substantially degrade image quality (Li et al. 2019). These effects are often non-uniform and depth-dependent, resulting in considerable variation in the appearance of scenes and object instances. A typical underwater imaging system and several degraded images are shown in Figure 1 (a). Early works of UIS (Lian et al. 2023; Jiang et al. 2024; Corrigan, Tay, and Konovessis 2023) use capacity-limited convolutional neural networks to perform end-to-end learning from collected underwater data. Unfortunately, due to their limited representational capacity, these conventional models still exhibit suboptimal performance in this domain, and some visualization results can be seen in Figure 1 (b). Recent advances in visual foundation models (Radford et al. 2021; Kirillov et al. 2023; Oquab et al. 2023; Cong et al. 2025), trained on large-scale datasets, have inspired growing interest in leveraging pretrained visual embeddings for underwater instance segmentation. Li et al. (2025a) introduced the Segment Anything Model (SAM) (Kirillov et al. 2023) to the UIS task, and proposed using the low-rank fine-tuning technique (Hu et al. 2022) to adapt pretrained representations to underwater scenarios. However, the approach still relies heavily on large-scale underwater datasets to alleviate domain misalignment and yields only marginal performance improvements.

DINOv2 (Oquab et al. 2023) is a powerful visual encoder pretrained in a self-supervised manner on the large-scale LVD-142M dataset, enabling it to learn rich and transferable visual representations. It has shown remarkable generalization ability across diverse downstream tasks, such as object detection (Liu et al. 2024), object tracking (Tumanyan et al. 2024), and image segmentation (Li et al. 2023). In contrast to SAM, which relies on task-specific supervised training, DINO learns task-agnostic visual features through self-supervised learning. This generalization capability is especially critical in underwater scenarios, where annotated data is scarce and the visual characteristics often deviate significantly from those of natural images. Nonetheless, due to the substantial domain gap, directly transferring DINOv2 to the



Figure 1: (a) A typical underwater imaging system: direct transmission carries useful scene information, forward scattering causes blurring, and backscattered light reduces visibility. We also present representative underwater images from UIS datasets. (b) Visual comparisons among *Watermask* (CNN-based method), *USIS-SAM* (SAM-based method), and *Ours* (DINO-based method).

underwater task is far from straightforward (see Figure 2).

In this paper, we focus on leveraging the DINOv2 foundation model to address the task of underwater instance segmentation, and propose a novel framework, called *DiveSeg*. To fully harness the capabilities of DINOv2, it is essential to enhance its adaptation from two perspectives: *underwater scene adaptation* and *underwater objects adaptation*. (1) Underwater scene adaptation aims to mitigate image degradation caused by color distortion and light scattering effects commonly found in underwater environments. To this end, we design a *AquaStyle Aligner*, which captures the unique stylistic features of underwater images through Fourier decomposition, and then injects them into DINOv2 backbone to eliminate the misalignment with the pretrained model through lightweight parameter learning. (2) Underwater objects adaptation focuses on enabling the model to generalize effectively to frequently occurring underwater objects—such as corals, jellyfish, and sea turtles—which are underrepresented in the LVD-142M dataset. To tackle this challenge, we introduce the *ObjectPrior Prompter*, which utilizes a binary mask encompassing all foreground objects to facilitate the learning of instance- and class-agnostic features, and subsequently employs them to prompt instance-specific learning. By decoupling instance-agnostic object perception from fine-grained instance discrimination, this mechanism substantially eases the challenge of adapting to diverse marine instances. We conduct comprehensive experiments on the UIIS and USIS10K datasets, and the results demonstrate that our method achieves superior segmentation performance. The main contributions are summarized as follows:

- We are the first to introduce DINO into the UIS task and propose the *DiveSeg*, which effectively eliminates domain inadaptability. Extensive experiments demonstrate that DINO can serve as a powerful learner for UIS and achieves state-of-the-art (SOTA) performance.
- We design the *AquaStyle Aligner* to embed the color style features of underwater scenes into the DINO fine-tuning process, enabling better adaptation to underwater envi-

ronments.

- We propose the *ObjectPrior Prompter*, which introduces binary object segmentation cues to provide object-level priors for complex instance segmentation, thereby guiding the model to more effectively localize and distinguish underwater target categories and instances.

Related Work

Visual Foundation Models

In the field of computer vision, visual foundation models have become indispensable for improving downstream task performance. Early approaches such as SimCLR (Chen et al. 2020) and MoCo (He et al. 2020) leveraged contrastive learning on tens of millions of images to learn effective feature representations, thereby laying the groundwork for unsupervised visual pretraining. Subsequently, DINO (Caron et al. 2021) introduced a self-distillation mechanism to extract robust representations from unlabeled data. Building upon these developments, MAE (He et al. 2022) further improved model generalization by performing masked image reconstruction on large-scale datasets. Meanwhile, the SAM (Kirillov et al. 2023) has demonstrated zero-shot segmentation capabilities for arbitrary visual objects by jointly training on large-scale manually annotated masks and automatically generated data, thereby introducing a new paradigm for interactive and automated segmentation.

More recently, DINOv2 (Oquab et al. 2023) was pre-trained on over one billion web-sourced images within a multi-modal processing pipeline, yielding substantial gains in classification, detection, and segmentation tasks. Unlike SAM, which focuses on precise boundary delineation, DINOv2 captures rich, high-dimensional semantic features that can be seamlessly integrated into downstream segmentation frameworks, thereby enhancing fine-grained semantic understanding in complex scenarios.

Underwater Instance Segmentation

Underwater instance segmentation aims to accurately identify and segment each individual object in underwater

scenes, offering both pixel-wise classification and instance-level distinction. Lian *et al.* (Lian et al. 2023) formally introduced the underwater instance segmentation task by constructing the USIS dataset and proposing WaterMask—an extension of Mask R-CNN adapted for underwater environments. While WaterMask integrated a dedicated module to mitigate light attenuation and color distortion, its dependence on conventional convolutional backbones limited its representational capacity, leading to suboptimal performance. To advance the field, Lian *et al.* (Lian et al. 2024) subsequently, the USIS10K benchmark dataset was released—this dataset focuses on salient instance segmentation on the basis of underwater salient object detection (Li et al. 2025b; Jin et al. 2024), and improves SAM by introducing USIS-SAM. Through carefully designed adapters and prompt mechanisms, USIS-SAM transfers the pre-trained capabilities of SAM to the underwater domain. However, because the prompt strategy does not fully accommodate the high target density and visual variability that are characteristic of underwater scenes, the segmentation accuracy of USIS-SAM remains limited.

In this work, we propose a novel underwater instance segmentation architecture built upon DINOv2. By harnessing the superior high-dimensional semantic features of DINOv2 and integrating a lightweight, scenario-specific aligner, our network substantially outperforms prior methods. Notably, it achieves these gains with only a modest increase in trainable parameters, demonstrating both effectiveness and efficiency in tackling the unique challenges of underwater instance segmentation.

Method

From Motivation to Overall Architecture

Benefiting from self-supervised pretraining on the large-scale LVD-142M (Oquab et al. 2023) dataset, DINOv2 demonstrates strong general-purpose visual feature extraction capabilities and excels across various downstream tasks. However, in underwater environments, the absorption and scattering of long-wavelength light by water leave primarily short-wavelength components, causing images to appear predominantly blue-green. In Figure 2, we investigate whether the vanilla DINOv2 can extract effective feature representations from underwater images using the principal component analysis (PCA). The results suggest that while DINOv2 is able to capture most of the primary targets in underwater images, its representations are often affected by background noise compared to natural images and may fail to detect some objects.

Hence, to obtain effective representations for UIS, specialized fine-tuning strategies are needed to facilitate the efficient adaptation of DINOv2 to underwater environments. This paper presents *DiveSeg*, a novel framework whose overall architecture is depicted in Figure 3, comprising two major components: (1) The *AquaStyle Aligner* addresses the domain adaptation challenge from a scene-level perspective by explicitly extracting background color information to model the unique stylistic features of underwater images, thereby mitigating color domain misalignment; (2) The *Ob-*

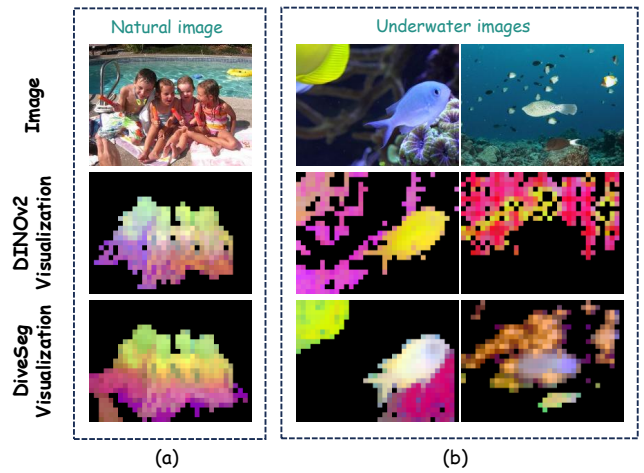


Figure 2: The PCA visualization of DINOv2 and DiveSeg on natural image and underwater images. The background is removed by thresholding the first PCA component.

jectPrior Prompter tackles the adaptation problem from an object-level perspective by leveraging binary object segmentation masks to guide the learning of target categories and instances in underwater scenes. The two adaptation strategies jointly empower DINOv2 to extract more discriminative and accurate representations tailored to underwater scenes, as shown in Figure 2.

AquaStyle Aligner

AquaStyle Aligner contains *Style Extraction* that captures underwater color style information via frequency-domain decomposition, and *Style Injection* that integrates these features into the DINOv2 representation learning process through a cross-attention mechanism.

Style Extraction. A primary distinction between underwater and natural images lies in the color distortion caused by light dispersion in the water medium. This color deviation predominantly manifests in the low-level features of the image. In the frequency domain, the amplitude component of the Fourier spectrum effectively preserves low-level statistical characteristics. Therefore, we adopt the Fourier amplitude as a representation of underwater style features. Specifically, for an image $x \in \mathbb{R}^{H \times W \times 3}$, its Fourier transform \mathcal{F}_x can be expressed as

$$\mathcal{F}_x(u, v) = \sum_{i=0}^{H-1} \sum_{j=0}^{W-1} x(i, j) \cdot e^{-2\pi I \left(\frac{ui}{H} + \frac{vj}{W} \right)}. \quad (1)$$

Here, I represents the imaginary unit, and each channel of the image is calculated independently. The amplitude component A_x and the phase component ϕ_x are respectively expressed as:

$$|A_x(u, v)| = \sqrt{\text{Re}\{F_x(u, v)\}^2 + \text{Im}\{F_x(u, v)\}^2}, \quad (2)$$

$$\phi_x(u, v) = \arctan \left(\frac{\text{Im}\{F_x(u, v)\}}{\text{Re}\{F_x(u, v)\}} \right), \quad (3)$$

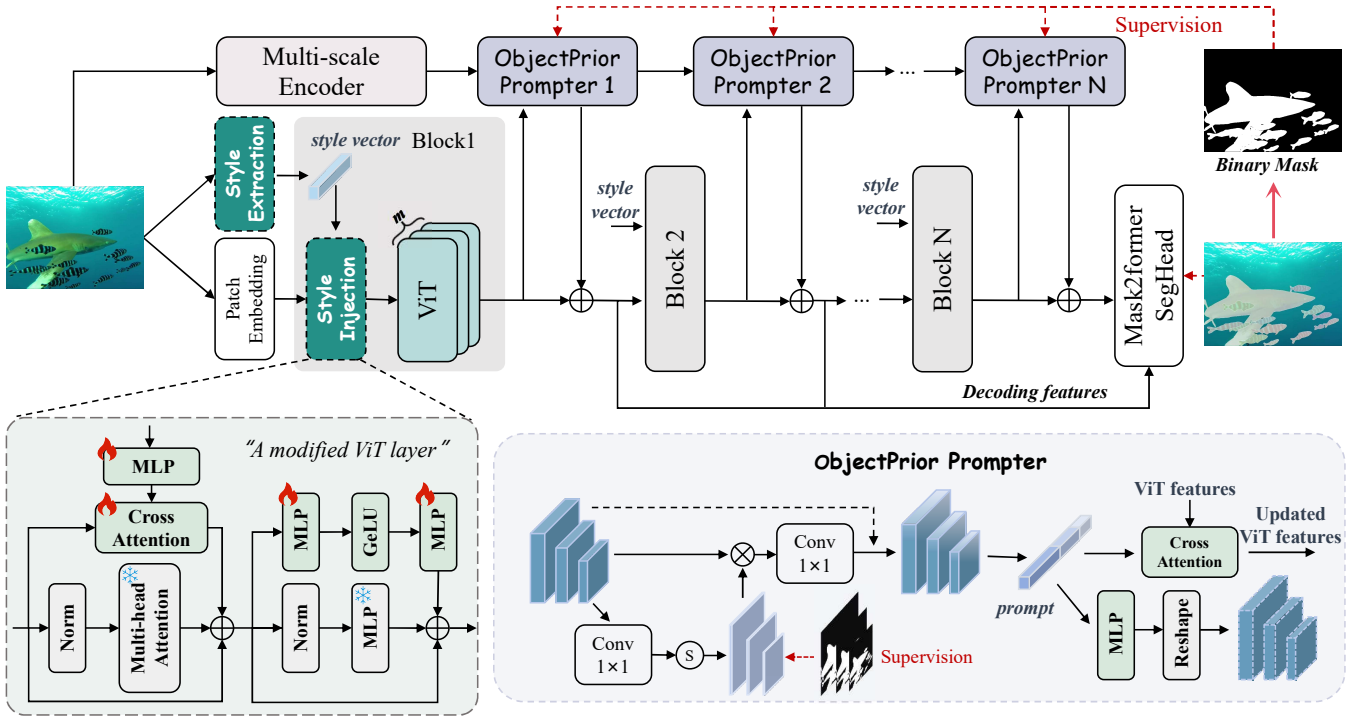


Figure 3: The overall framework of the proposed *DiveSeg* is illustrated as follows. First, we employ a *Style Extraction* module to obtain an underwater style vector. This vector is subsequently injected into the frozen DINOv2 backbone via the *Style Injection* module, enabling rapid adaptation to the underwater domain. Together, these two modules constitute the *AquaStyle Aligner*. In addition, the *ObjectPrior Prompter* leverages binary masks to learn object-level priors, which guide the network to focus on underwater objects and ease the challenge of directly segmenting specific instances.




Figure 4: Underwater images and the corresponding style images, FT and iFT represents Fourier transform and inverse Fourier transform.

where $\text{Re}(\cdot)$ and $\text{Im}(\cdot)$ are the real part and the imaginary part respectively. To represent the image style information, we fix the phase at the average value $\bar{\phi}_x$ and use the amplitude information to reconstruct the style image \hat{x} via the inverse Fourier transform \mathcal{F}^{-1} :

$$\hat{x}(i, j) = \mathcal{F}^{-1} \left\{ |A_x(u, v)| \cdot e^{i\bar{\phi}_x} \right\}. \quad (4)$$

As illustrated in Figure 4, averaging the phase while retaining the amplitude in the frequency domain effectively removes object-related content from the image, preserving the distinctive color characteristics of underwater scenes. Based on this, we employ multi-layer convolution and global aver-

age pooling to encode the style image into a compact style vector p_x .

Style Injection. Inspired by the adapter design proposed by Houlsby *et al.* (Houlsby *et al.* 2019), we develop a style injection module that employs cross-attention rather than relying solely on the standard multilayer perceptron (MLP), enabling more effective interaction between style representations and image features. Specifically, implemented as a parallel branch to the Multi-head Attention with normalization layer (MHA) in ViT, our style injection module shares the same input but integrates the style vector p_x through a cross-attention mechanism. In this design, the query is obtained from ViT features, while the key and value are generated from the style vector, enabling effective fusion of visual content and style information. This injection process can be expressed as:

$$\omega_1 = \text{MHA}(V_{in}) + \text{CrossAttn}(V_{in}, \text{MLP}(p_x)), \quad (5)$$

where V_{in} represents the input features of the ViT block, and all parameters of the *MHA* are frozen. Essentially, the features produced by the cross-attention can be regarded as a complementary representation to those from the original *MHA*, helping to prevent the degradation of the pretrained model. The output of this step is denoted as ω_1 .

At the subsequent Feed-Forward (FF) in ViT block, which consists of an MLP followed by a normalization layer, following the standard adapter design paradigm, we employ

Method	Backbone	Params	mAP	AP_{50}	AP_{75}
CNN-Based Methods					
Mask R-CNN	ResNet-101	63M	23.4	40.9	25.3
Mask Scoring R-CNN	ResNet-101	79M	24.6	41.9	26.5
Cascade Mask R-CNN	ResNet-101	88M	25.5	42.8	27.8
BMask R-CNN	ResNet-101	66M	22.1	36.2	24.4
Point Rend	ResNet-101	63M	25.9	43.4	27.6
R ³ -CNN	ResNet-101	77M	24.9	40.5	27.8
Mask Transfiner	ResNet-101	63M	24.6	42.1	26.0
Mask2Former	ResNet-101	63M	25.7	38.0	27.7
WaterMask	ResNet-101	67M	27.2	43.7	29.3
Transformer-Based Methods					
USIS-SAM	ViT-H	701M	29.4	<u>45.0</u>	<u>32.3</u>
<i>DiveSeg</i> (Ours)	ViT-L	390M	35.6	52.0	38.5

Table 1: Quantitative comparisons with state-of-the-arts on the UIIS dataset. The best results are highlighted in bold, while the second-best results are indicated with underlines.

cascaded multilayer perceptrons to integrate the features from the previous step at a deeper level. This enables DINOv2 to better capture the distinctive color characteristics of underwater images and acquire more discriminative visual representations. This is formulated as:

$$\omega_2 = FF(\omega_1) + MLP(GeLU(MLP(\omega_1))). \quad (6)$$

Likewise, the parameters of the FF remain frozen. Two additional MLP layers are employed to reduce and subsequently restore the feature dimensionality, forming a bottleneck structure. The output of this step, denoted as ω_2 , is taken as the output of the ViT block.

Notably, we evenly divide all ViT layers in the original DINOv2 architecture into four sequential blocks to adapt to the instance segmentation decoder. To minimize the computational overhead and parameter growth, the *AquaStyle Aligner* is only applied to the first ViT layer in each block, resulting in a total of four *AquaStyle Aligners* across the feature extraction process.

ObjectPrior Prompter

Along with the color adaptation for underwater scenes largely resolved, the instance segmentation of underwater-specific objects continues to pose significant challenges. This is primarily due to the rarity of these object categories in the DINOv2 pretraining dataset, rendering direct learning of instance-level segmentation in underwater contexts a nontrivial task. Driven by this insight, we introduce the *ObjectPrior Prompter* that initially applies a binary mask to constrain class-agnostic and instance-agnostic feature learning, subsequently embedding these features as prior prompts within the DINOv2 backbone to ease the learning of instance-level underwater objects.

As illustrated in the top section of Figure 3, we initially extract image features at multiple resolutions through a multi-scale encoder. Subsequently, the *ObjectPrior Prompter* interacts with the backbone to progressively incorporate underwater object priors. Specifically, the multi-scale encoder follows a simple design, consisting of three

convolutional layers for feature extraction, stride-2 convolutions for downsampling, and 1×1 convolutions for dimensionality reduction. The encoder outputs a three-scale pyramid of features, denoted as $\{f_M^1, f_M^2, f_M^3\}$, with corresponding resolutions $\{\frac{1}{8^2}, \frac{1}{16^2}, \frac{1}{32^2}\}$ of the input image.

In the *ObjectPrior Prompter*, the multi-scale features are first used to generate pseudo masks via 1×1 convolution layers followed by Sigmoid functions, supervised by a binary mask. The binary mask is a single-channel image derived from the segmentation ground truth, covering all target objects. During both training and inference, pseudo masks P_{mask} are dynamically generated for each image to represent all foreground instances. It can be formulated as:

$$P_{mask}^k = \sigma\left(\text{Conv}_{1 \times 1}^k(f_M^k)\right), \quad (7)$$

where $k \in \{1, 2, 3\}$. Then, we perform element-wise multiplication between the pseudo masks and the corresponding original features to preserve object-related information. The resulting features are further integrated with the original ones via a convolutional layer and residual connection, enhancing instance-level representations without compromising the original semantics. Formally defined as:

$$f_{MT}^k = \text{Conv}_{1 \times 1}(P_{mask}^k \cdot f_M^k) + f_M^k, \quad (8)$$

Subsequently, the features f_{MT}^k are flattened and concatenated, termed as O_{prompt} , and used as an object-level prior prompt to interact with ViT features. In the cross-attention mechanism, the prompt serves as the key and value, while the ViT features act as the query:

$$f_{ViT}^{opp} = \text{CrossAttn}(f_{ViT}, O_{prompt}). \quad (9)$$

Here, f_{ViT} and f_{ViT}^{opp} denote the ViT features before and after prompt-based interaction, respectively. Finally, we update the features O_{prompt} through multiple linear layers, and then reshape them back to the original pyramid feature dimensions, allowing supervision from the binary mask. We embed the *ObjectPrior Prompter* module after each block and use the sum of its output f_{ViT}^{opp} and the original ViT feature f_{ViT} as the decoder input.

Method	Backbone	Params	Class-Agnostic			Multi-Class		
			mAP	AP_{50}	AP_{75}	mAP	AP_{50}	AP_{75}
CNN-Based Methods								
S4Net	ResNet-50	47M	32.8	64.1	27.3	23.9	43.5	24.4
RDPNet	ResNet-101	66M	54.7	78.3	63.0	39.3	55.9	45.4
OQTR	ResNet-50	50M	56.6	79.3	62.6	19.7	30.6	21.9
WaterMask	ResNet-101	67M	59.0	80.6	67.2	38.7	54.9	43.2
Transformer-Based Methods								
SAM+BBBox	ViT-H	641M	45.9	65.9	52.1	26.4	38.9	29.0
SAM+Mask	ViT-H	641M	55.1	80.2	62.8	38.5	56.3	44.0
RSPrompter	ViT-H	632M	58.2	79.9	65.9	40.2	55.3	44.8
USIS-SAM	ViT-H	701M	<u>59.7</u>	<u>81.6</u>	<u>67.7</u>	<u>43.1</u>	<u>59.0</u>	<u>48.5</u>
<i>DiveSeg</i> (Ours)	ViT-L	390M	64.1	82.8	72.2	48.4	62.3	54.4

Table 2: Quantitative comparisons with state-of-the-arts on the USIS10K dataset. The best results are highlighted in bold, while the second-best results are indicated with underlines.




Figure 5: Qualitative comparisons of *DiveSeg* with SOTA UIS methods on the USIS10K and UIIS datasets.

Experiments

We conduct extensive experiments on UIIS (Lian et al. 2023) and USIS10K (Lian et al. 2024) datasets to verify the effectiveness of the proposed *DiveSeg*. Experimental results show that *DiveSeg* achieves better performance than the existing SOTA methods. In addition, we also performed ablation experiments on the proposed *AquaStyle Aligner* and *ObjectPrior Prompter* to verify the effectiveness of individual modules.

Experimental Setup

Datasets. The proposed method is evaluated on UIIS and USIS10K datasets. The UIIS dataset contains 3937 training images and 691 testing images, covering 7 instance categories, as follows: Fish, Reefs, Aquatic plants, Wrecks/Ruins, Human divers, Robots, and Sea-floor. The USIS10K dataset contains 10632 images, which are divided into training set, validation set and test set according to the ratio of 7:1.5:1.5. The dataset contains not only the seven instance categories mentioned above, but also class-agnostic labels

to support class-agnostic instance detection tasks.

Evaluation Metrics. To ensure a fair and comprehensive comparison, we adopt the standard mask AP evaluation protocol commonly used in instance segmentation benchmarks. This includes metrics such as mAP , AP_{50} , and AP_{75} , which correspond to average precision computed at different IoU thresholds. These metrics collectively capture both coarse and fine-grained segmentation accuracy, allowing for a thorough assessment of the model’s ability to detect and localize instance masks with varying levels of precision.

Implementation Details. We implemented the proposed approach using PyTorch and the Detectron2 framework. Training was conducted on an NVIDIA A100 GPU with a batch size of 8, using AdamW as the optimizer with a weight decay of 0.05. The initial learning rate was set to $1e-4$, and the warm-up strategy was employed. The training lasted for 30,000 iterations, during which the learning rate was decayed to one-tenth of its initial value at the 23,000th and 27,000th iterations. The segmentation head adopted the Mask2Former (Cheng et al. 2022) architecture. We applied

Methods	mAP	AP_{50}	AP_{75}
Full Model	35.6	52.0	38.5
w/o AquaStyle Aligner	34.8	50.6	37.6
w/o ObjectPrior Prompter	34.1	50.8	37.8
DINOv2+Mask2Former	30.9	44.6	32.2

Table 3: Ablation study of individual components.

classification loss and mask loss to constrain the predictions of Mask2Former, while binary cross-entropy (BCE) loss, IoU loss, and L1 loss were used to guide the learning of pseudo masks.

Comparison with State-of-the-arts

Quantitative Results. (1) On the UIIS dataset. As shown in Table 1, we compare our method against the SOTA algorithms. Our *DiveSeg*, which employs ViT-L as the backbone, achieves mAP , AP_{50} , and AP_{75} of 35.6, 52.0, and 38.5, respectively, representing improvements of 21.1%, 15.6%, and 19.2% over USIS-SAM with a ViT-H backbone. Furthermore, *DiveSeg* substantially outperforms WaterMask and other competitors across all metrics, demonstrating its superior capability in underwater instance segmentation. **(2) On the USIS10K dataset.** As summarized in Table 2, our approach demonstrates clear superiority over existing methods. In the class-agnostic setting, it achieves mAP , AP_{50} , AP_{75} of 64.1, 82.8, 72.2, outperforming USIS-SAM by 4.4, 1.2, and 4.5 percentage points, respectively. Under the multi-class setting, it attains mAP , AP_{50} , AP_{75} of 48.4, 62.3, 54.4—substantially higher than USIS-SAM—while using only 55.6% of its parameters. These results compellingly demonstrate that our model not only excels at class-agnostic instance segmentation but also delivers precise class-level and instance-level predictions in challenging underwater environments.

Qualitative Results. As shown in Figure 5, we qualitatively compare our method with WaterMask and USIS-SAM on the UIIS and USIS10K test sets. Our approach more accurately delineates object boundaries for instance, successfully segmenting fish in shadows (Column 1) and underwater ruins (Column 6), and detecting coral missed by others (Column 4). In terms of classification accuracy, our model also outperforms the competitors. For example, in Column 2, other methods misclassify fish as coral due to complex environmental interference, whereas our model produces correct predictions. Moreover, in cases of instance overlap such as fish and coral (Column 3) or divers with equipment (Column 5) our model maintains clear instance boundaries, while others often produce boundary confusion.

Ablation Study

Effectiveness of the Individual Modules. Table 3 presents an ablation study evaluating the contributions of the *AquaStyle Aligner* and the *ObjectPrior Prompter*, based on a baseline that comprises a frozen DINOv2 backbone and a trainable Mask2Former decoder. Introducing the *AquaStyle Aligner* increases mAP by 3.9 points, AP_{50} by 6.0 points,

Methods	mAP	AP_{50}	AP_{75}
Frozen	30.9	44.6	32.2
Full Fine-tuning	31.1	45.4	35.2
LoRA	31.8	47.4	34.6
Adapter	32.7	48.1	36.4
<i>AquaStyle Aligner</i>	34.1	50.8	37.8

Table 4: Ablation study of *AquaStyle Aligner*.

and AP_{75} by 5.4 points, while adding the *ObjectPrior Prompter* yields improvements of 3.2, 6.2, and 5.6 points on these metrics, respectively. When both modules are combined, mAP reaches 35.6, with corresponding further gains in AP_{50} and AP_{75} , which demonstrates that the *AquaStyle Aligner* and the *ObjectPrior Prompter* play significant roles in adapting DINOv2 for underwater instance segmentation.

Ablation Study on the AquaStyle Aligner. Table 4 reports an analysis of alternative adaptation strategies for DINOv2 to validate the efficacy of our *AquaStyle Aligner*. "Frozen" denotes that all parameters of DINOv2 are kept fixed, and no additional learnable parameters are introduced in the encoder. First, we adopt the full fine-tuning strategy for DINOv2 (Row 2), which enables gradient updates for all parameters of the DINOv2 model that are kept frozen in the "Frozen" setting. However, this approach yields only marginal improvements over the frozen backbone, possibly due to catastrophic forgetting of the pre-trained knowledge caused by updating a large number of parameters. Moreover, inspired by recent studies demonstrating that low-rank adaptation (LoRA) and lightweight Adapters can effectively tailor pre-trained ViTs with minimal parameter overhead, we apply these techniques in Rows 3 and 4. Both approaches lead to substantial performance improvements, confirming the effectiveness of targeted parameter fine-tuning. The proposed *AquaStyle Aligner* scheme (Row 5), which incorporates specialized underwater style extraction and effective style injection, surpasses these methods. This demonstrates that integrating underwater style features into DINOv2 provides more pronounced benefits for understanding and segmenting underwater scenes.

Conclusion

This paper introduces *DiveSeg*, a novel Underwater Instance Segmentation (UIS) framework built on the DINOv2 visual foundation model. It tackles two key challenges of UIS: underwater scene adaptation and underwater object adaptation. *DiveSeg* has two core components: *AquaStyle Aligner*, which embeds underwater color style features into DINOv2 fine-tuning via Fourier decomposition and cross-attention to boost domain adaptation; and *ObjectPrior Prompter*, which uses binary segmentation prompts to provide object-level priors, enabling instance segmentation through dual object- and instance-level reasoning. Extensive experiments on UIIS and USIS10K datasets show *DiveSeg* achieves state-of-the-art performance, significantly outperforming existing methods in both quantitative and qualitative results.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62471278, and in part by the Taishan Scholar Project of Shandong Province under Grant tsqn202306079.

References

- Abdullah, A.; Barua, T.; Tibbetts, R.; Chen, Z.; Islam, M. J.; and Rekleitis, I. 2024. Caveseg: Deep semantic segmentation and scene parsing for autonomous underwater cave exploration. In *IEEE International Conference on Robotics and Automation*, 3781–3788.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9650–9660.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, 1597–1607.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girdhar, R. 2022. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1290–1299.
- Christensen, L.; de Gea Fernández, J.; Hildebrandt, M.; Koch, C. E. S.; and Wehbe, B. 2022. Recent advances in ai for navigation and control of underwater robots. *Current Robotics Reports*, 3(4): 165–175.
- Cong, R.; Yu, Z.; Fang, H.; Sun, H.; and Kwong, S. 2025. UIS-Mamba: Exploring mamba for underwater instance segmentation via dynamic tree scan and hidden state weaken. In *Proceedings of the ACM International Conference on Multimedia*, 343–352.
- Cong, Y.; Gu, C.; Zhang, T.; and Gao, Y. 2021. Underwater robot sensing technology: A survey. *Fundamental Research*, 1(3): 337–345.
- Corrigan, B. C.; Tay, Z. Y.; and Konovovsis, D. 2023. Real-time instance segmentation for detection of underwater litter as a plastic source. *Journal of Marine Science and Engineering*, 11(8): 1532.
- Fang, H.; Cong, R.; Lu, X.; Zhou, X.; Kwong, S.; and Zhang, W. 2025a. Decoupled motion expression video Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13821–13831.
- Fang, H.; Zhang, T.; Zhou, X.; and Zhang, X. 2025b. Learning better video query with SAM for video instance segmentation. *IEEE Transactions on Circuits and Systems for Video Technology*, 35(4): 2963–2974.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16000–16009.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning*, 2790–2799.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *International Conference on Learning Representations*, 1(2): 3.
- Islam, M. J.; Edge, C.; Xiao, Y.; Luo, P.; Mehtaz, M.; Morse, C.; Enan, S. S.; and Sattar, J. 2020. Semantic segmentation of underwater imagery: dataset and benchmark. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Jiang, J.; Zuo, X.; Shu, X.; Xu, D.; and Qian, P. 2024. UW-DETR: Underwater image instance segmentation with Co-DETR. In *International Conference on Pattern Recognition and Artificial Intelligence*, 125–130.
- Jin, J.; Jiang, Q.; Wu, Q.; Xu, B.; and Cong, R. 2024. Underwater salient object detection via dual-stage self-paced learning and depth emphasis. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, P. D.; Wan-Yen; and Girshick, R. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Li, C.; Guo, C.; Ren, W.; Cong, R.; Hou, J.; Kwong, S.; and Tao, D. 2019. An underwater image enhancement benchmark dataset and beyond. *IEEE Transactions on Image Processing*, 29: 4376–4389.
- Li, F.; Zhang, H.; Xu, H.; Liu, S.; Zhang, L.; Ni, L. M.; and Shum, H.-Y. 2023. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3041–3050.
- Li, H.; Lian, S.; Li, Z.; Cong, R.; and Kwong, S. 2025a. UWSAM: Segment anything model guided underwater instance segmentation and a large-scale benchmark dataset. *arXiv preprint arXiv:2505.15581*.
- Li, H.; Lin, G.; Li, Z.; Kwong, S.; and Cong, R. 2025b. FSCDiff: Frequency-spatial entangled conditional diffusion model for underwater salient object detection. In *Proceedings of the ACM International Conference on Multimedia*, 8379–8388.
- Lian, S.; Li, H.; Cong, R.; Li, S.; Zhang, W.; and Kwong, S. 2023. Watermask: Instance segmentation for underwater imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1305–1315.
- Lian, S.; Zhang, Z.; Li, H.; Li, W.; Yang, L. T.; Kwong, S.; and Cong, R. 2024. Diving into underwater: Segment anything model guided underwater salient instance

segmentation and a large-scale dataset. *arXiv preprint arXiv:2406.06039*.

Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, 38–55.

Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763.

Tang, Y.; Zhu, C.; Wan, R.; Xu, C.; and Shi, B. 2024. Neural underwater scene representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11780–11789.

Tumanyan, N.; Singer, A.; Bagon, S.; and Dekel, T. 2024. Dino-tracker: Taming dino for self-supervised point tracking in a single video. In *European Conference on Computer Vision*, 367–385.

Zheng, Z.; Liang, H.; Hua, B.-S.; Wong, Y. H.; Ang, P.; Chui, A. P. Y.; and Yeung, S.-K. 2024. Coralscop: Segment any coral image on this planet. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28170–28180.