

SAM-PM: Enhancing Video Camouflaged Object Detection using Spatio-Temporal Attention

Muhammad Nawfal Meeran Gokul Adethya T[†] Bhanu Pratyush Mantha[†]
 National Institute of Technology, Tiruchirappalli




Figure 1. Comparison of mask predictions between ground truth (GT), SAM-PM (Ours), and SLT-Net.

Abstract

In the domain of large foundation models, the Segment Anything Model (SAM) has gained notable recognition for its exceptional performance in image segmentation. However, tackling the video camouflage object detection (VCOD) task presents a unique challenge. Camouflaged objects typically blend into the background, making them difficult to distinguish in still images. Additionally, ensuring temporal consistency in this context is a challenging problem. As a result, SAM encounters limitations and falls short when applied to the VCOD task. To overcome these challenges, we propose a new method called the SAM Propagation Module (SAM-PM). Our propagation module enforces temporal consistency within SAM by employing

spatio-temporal cross-attention mechanisms. Moreover, we exclusively train the propagation module while keeping the SAM network weights frozen, allowing us to integrate task-specific insights with the vast knowledge accumulated by the large model. Our method effectively incorporates temporal consistency and domain-specific expertise into the segmentation network with an addition of less than 1% of SAM’s parameters. Extensive experimentation reveals a substantial performance improvement in the VCOD benchmark when compared to the most recent state-of-the-art techniques. Code and pre-trained weights are open-sourced at <https://github.com/SpiderNitt/SAM-PM>

[†]Co-authors who contributed significantly to pivotal ideation.

1. Introduction

The task of detecting objects that seamlessly blend into the background is crucial for various applications, including surveillance and security [30], autonomous driving [9, 47], robotics [38], and medical image segmentation [17, 53]. This challenging task is addressed by Video Camouflaged Object Detection (VCOD) and Camouflage Object Detection (COD). Despite their broad practical utility, these tasks are daunting as camouflaged objects are often indistinguishable to the naked-eyes. Consequently, in the video domain, related problems like Video Object Segmentation (VOS) [2, 59], and Video Motion Segmentation (VMS) [23, 57] have received considerable attention in computer vision, but VCOD remains relatively under-explored.

Recently, the Segment Anything Model (SAM) [25] has gained a lot of attention among foundation (i.e., models trained on vast amounts of data) models like CLIP [44], BLIP [27, 28], and DALL-E [45, 46] in the field of computer vision. SAM is capable of producing high-quality segmentation masks in diverse scenarios. However, it performs poorly in VCOD and COD, as indicated by our experiments and the findings of this study [5] respectively. In VCOD, SAM confronts three key challenges. Firstly, the extensive visual training corpus primarily comprises objects with well-defined boundaries, neglecting the representation of camouflaged objects characterized by ambiguous and indistinguishable boundaries. Secondly, SAM is trained on static image data and struggles to capture motion and maintain temporal consistency across consecutive video frames. Thirdly, the ambiguous boundary means the appearance of the camouflaged object resembles the background. This gives rise to two fundamental problems: **1)** the object’s boundaries seamlessly merge with the background, becoming perceptible only during movement; **2)** the object typically exhibits repetitive textures akin to its surroundings. As a result, the determination of pixel movement across frames for motion estimation (e.g., using optical flow) becomes erratic and prone to errors. Consequently, using VOS or VMS methods to address the issue encounters significant failures. In addressing this challenge, the prior approach SLT-Net[10] has employed two distinct modules to implicitly capture motion and maintain temporal consistency. However, this approach proves to be highly resource-intensive in terms of training, presenting a significant drawback.

In this study, we introduce the **SAM-Propagation Module (SAM-PM)**, a novel approach to tackle VCOD. SAM-PM incorporates a Temporal Fusion Mask Module (TFMM) and a Memory Prior Affinity Module (MPAM) into the Propagation Module (PM) to ensure the temporal consistency of masks through SAM. Utilizing TFMM, we extract mask embeddings that are both spatially and temporally informed. These embeddings play a crucial role in enhancing

temporal information within MPAM. Furthermore, to harness the capabilities acquired by SAM from the massive data, we opt to freeze SAM weights and only train our PM to gain the domain-specific information required for VCOD. In this work, our contributions can be summarized as follows:

- We introduce a novel VCOD framework and demonstrate its effectiveness by deploying it alongside SAM. This approach ensures temporal consistency and imparts domain knowledge through a unified optimization target.
- Achieving these results is possible by training fewer than 1 million parameters while keeping the SAM weights frozen. This renders our SAM-PM exceptionally parameter-efficient during training.
- Establishing a new state-of-the-art in the VCOD task, we surpass the performance of the previous state-of-the-art (SOTA) SLT-Net [10] by an impressive margin.

2. Related Works

COD. Even with well-trained eyes, humans can surprisingly overlook camouflaged objects, highlighting the remarkable effectiveness of camouflage in deceiving our visual perception. However, biological studies [19] reveal that predatory animals follow a two-step process: searching before identifying their prey. Building on this insight, SINet [16] introduces two modules: one for the initial search of camouflaged objects and another for their precise detection. Motivated by the same concept, PFNet [36] and MirrorNet [55] adopt a strategy of first determining the coarse location of camouflage objects and then refining the process through segmentation to generate intricate masks. MGL [65] incorporated edge details into the segmentation stream via two graph-based modules. More recently, vision transformer-based models like SINet-v2 [18], ZoomNet [41], and FSPNet [21] have shown strong global and local context modeling capabilities in camouflaged object detection.

VCOD. Prior works [3, 4, 57] in this field have mainly relied on optical flow-based methods. [37] proposed to segment multiple motions in a rapid and non-iterative way using optical flows. In [57], the authors utilized a network to segment the angle field instead of the raw optical flow. [26] proposed a framework consisting of two components: a registration module aligning background across consecutive frames and a motion segmentation module with memory for detecting moving objects. They also proposed a larger camouflaged dataset (MoCA) with bounding boxes labeled for every five frames. SLT-Net presented a framework that models both short and long-term temporal consistency in camouflaged videos. Furthermore, they proposed a large-scale MoCA-Mask dataset containing 87 video sequences with pixel-wise ground truth masks.

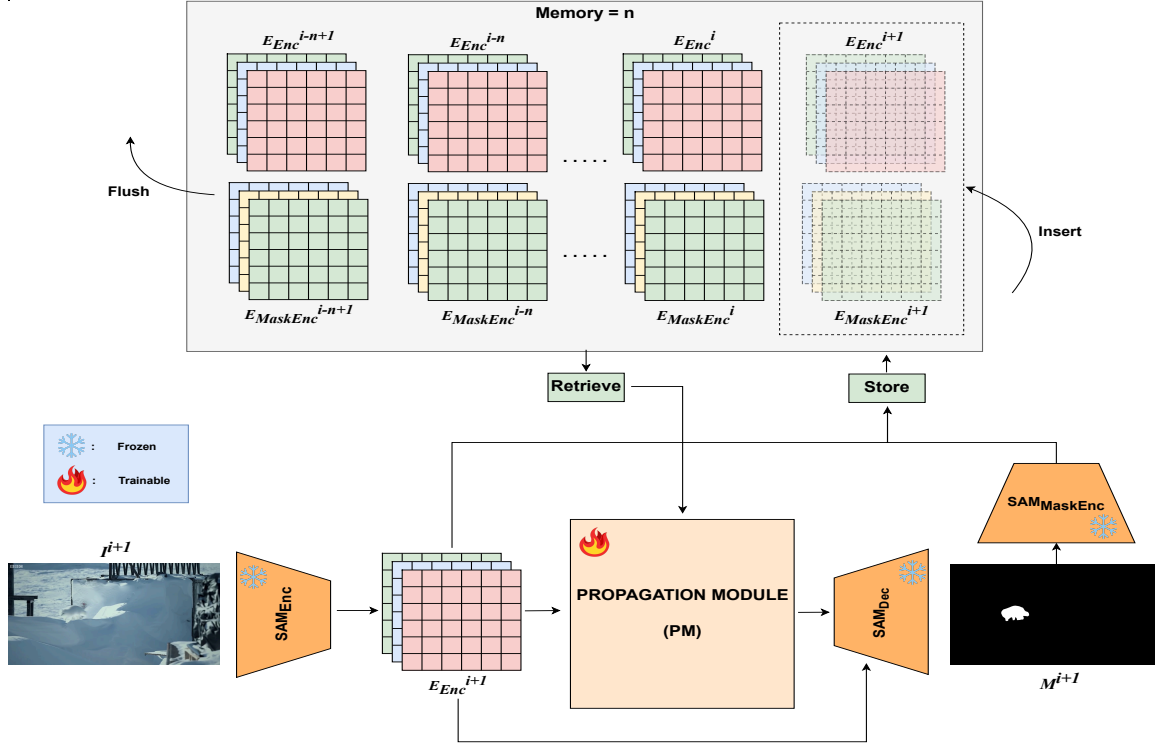


Figure 2. Overall framework of SAM-PM consisting of SAM(frozen), PM(trainable) and Memory. Instead of adding the image and mask directly, we incorporate their embeddings into Memory. This prevents redundant encoding when predicting future masks. One significant advantage of this architecture is that we only need to encode the input frame once, regardless of the number of objects we aim to track.

VOS. This task requires localizing objects in videos with temporally consistent pixel-wise masks. Video Object Segmentation (VOS) involves distinct subtasks, including unsupervised (automatic) video segmentation [11, 33, 49–51], semi-supervised (mask-guided semi-automatic) video segmentation [34, 58, 60–63], and interactive (scribble or click-based) video segmentation [7]. In the realm of semi-supervised VOS, maintaining temporal consistency has led to the development of several spatiotemporal memory-based methods, such as XMem [6], QDMN [31], and Space-Time Memory Network (STM) [40]. Notably, Space-Time Correspondence Network (STCN) [8] introduces a direct image-to-image correspondence approach, which stands out for its simplicity, efficiency, and effectiveness compared to STM. XMem, drawing inspiration from the Atkinson-Shiffrin memory model [1], proposes a feature memory tailored for long video VOD, addressing the need for temporal coherence. There also has been research done to make VOS more applicable to real-world setting using more complex setups. MOSE [12] introduces a diverse dataset featuring numerous crowds, occlusions, and frequent object disappearance-reappearance instances in extended videos. This poses a significant challenge to the performance of existing VOS models. Similarly, other datasets such as OVIS

[42] focus on heavily occluded objects, UVO [52] contributes towards open-world dense objects and VOST [48] is directed toward complex object transformations.

3. Proposed Framework

We propose SAM-PM to adapt SAM for Video Camouflaged Object Detection. In Sec 3.1, we briefly review the architecture of SAM on which our SAM-PM is built. Subsequently, in Sec 3.2, we present SAM-PM alongside its key components: the Temporal Fusion Mask Module (TFMM) and the Memory Prior Affinity Module (MPAM), which are the key components to achieve better segmentation quality on videos containing camouflaged objects.

SAM-PM takes as input a video clip featuring camouflaged objects and produces a series of pixel-wise binary masks corresponding to each frame in the video. Let the video clip consisting of T frames be denoted as $\{I^i\}_{i=1}^T$ where $I^i \in R^{(3 \times H \times W)}$. Our objective is to assign binary masks $M^i \in \{0, 1\}^{(H \times W)}$ to these frames. We encode the frames and generated masks using the SAM image and prompt encoder. We denote the image embeddings and mask embeddings for the i^{th} frame as E_{Enc}^i and $E_{MaskEnc}^i$ respectively. The initial input to our model is the first frame and the corresponding ground truth.

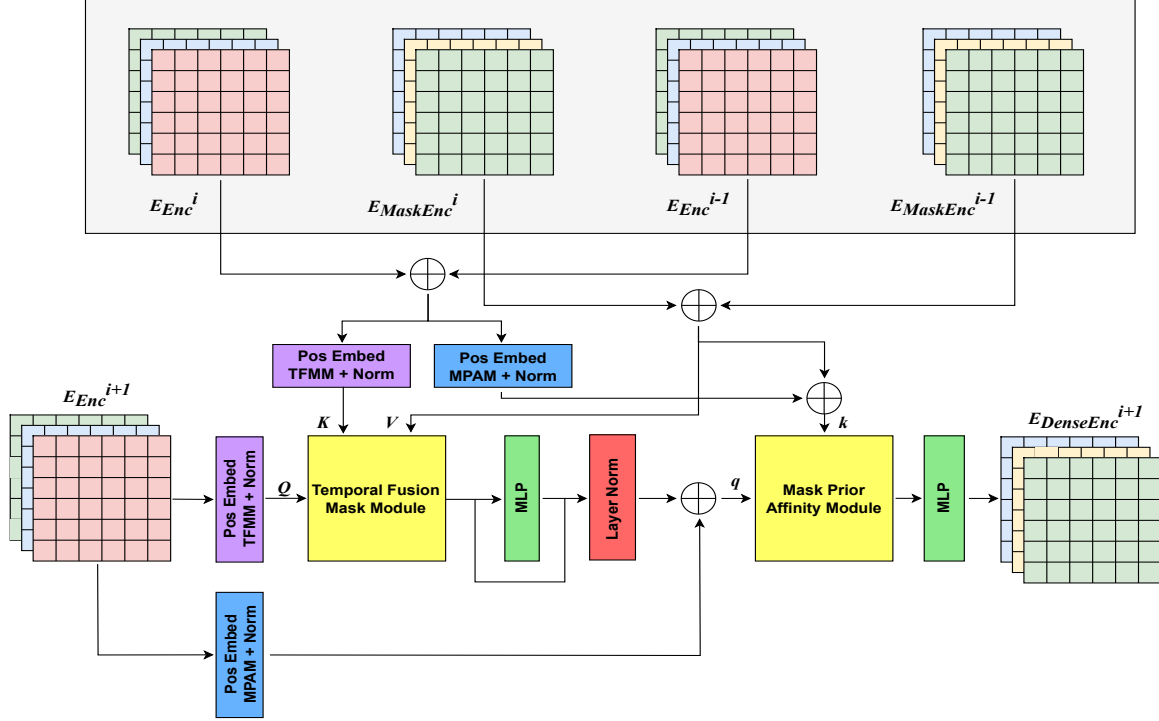


Figure 3. Overview of our Propagation Module consisting of TFMM and MPAM. It’s important to observe that we utilize separate positional encoding for each module to provide greater flexibility within the model. Both positional encodings come with trainable parameters, enabling the model to regulate the extent of positional encoding applied to them.

3.1. Preliminaries: SAM

SAM [25] comprises of three key modules: **(1)** Image encoder SAM_{Enc} : Utilizes a Masked Autoencoder (MAE) [20] pre-trained Vision Transformer [13] for image feature extraction. The image encoder’s output consists of image embeddings 16 times smaller than the original image. **(2)** Prompt encoder: It encodes the positional information from the input points/boxes to generate sparse embeddings, while input mask prompts are encoded into dense embeddings using the mask encoder $SAM_{MaskEnc}$. **(3)** Mask decoder SAM_{Dec} : It efficiently maps image embeddings and prompt embeddings (i.e. sparse and dense embeddings) to an output mask. The SAM model is trained on the SA-1B dataset, which consists of over 1 billion masks and 11 million images. For more details regarding SAM, we refer readers to [25].

3.2. Ours: SAM-PM

The overall framework of SAM-PM is shown in Figure 2. SAM-PM consists of a Propagation Module (PM) and a simple yet effective Memory alongside the frozen SAM network. SAM-PM addresses the VCOD task in a semi-supervised way, utilizing the first frame’s ground truth mask along with video frames as input to prompt encoder and image encoder, respectively. We store the sequence of frames

and their corresponding predicted masks in the memory to serve as the inputs to several components within the Propagation Module (PM). The PM is composed of the TFMM and MPAM as shown in Figure 3. The TFMM takes as input the current frame’s image embedding, concatenated image embeddings of previous frames, and concatenated predicted mask embeddings from memory. It performs spatio-temporal cross attention, resulting in a temporally infused mask embedding. The MPAM concatenates the temporally infused mask embedding with the current frame’s image embedding and the previous frame predicted mask embeddings with their corresponding image embeddings from the memory to perform affinity, consequently enhancing the temporal information. During training, we consider each sample training point to consist of frames extracted from a fixed-length memory n , representing subsequences from the full video. During inference, we store and retrieve frames for the whole sequence of the video using a simple strategy. Furthermore, we use a two-stage training method to maintain stability during the training process.

For a given frame j we calculate E_{Enc}^j and $E_{MaskEnc}^j$ as follows

$$E_{Enc}^j = SAM_{Enc}(I^j) \quad (1)$$

$$E_{MaskEnc}^j = SAM_{MaskEnc}(M^j) \quad (2)$$

To predict the mask M^{i+1} for the $i + 1^{th}$ frame, we first pass the current frame’s image embedding E_{Enc}^{i+1} , along with the image embeddings $E_{Enc}^{i-n+1:i}$ and the corresponding mask embeddings $E_{MaskEnc}^{i-n+1:i}$ for the previous n frames, to the PM module. This yields the dense embedding $E_{DenseEnc}^{i+1}$ for the current frame as described in Eqn. 4

$$E_{Enc}^{i+1} = SAM_{Enc}(I^{i+1}) \quad (3)$$

$$E_{DenseEnc}^{i+1} = PM(E_{Enc}^{i-n+1:i+1}, E_{MaskEnc}^{i-n+1:i}) \quad (4)$$

$$M^{i+1} = SAM_{Dec}(E_{Enc}^{i+1}, E_{DenseEnc}^{i+1}) \quad (5)$$

In Eqn. 5, the dense and image embeddings of the current frame are passed to the decoder to obtain the current mask prediction. This process is repeated for the entire video, as described in Eqn. 3-5 during which the memory is updated by adding new image and dense embeddings while removing outdated ones when the memory limit is reached.

3.2.1 Temporal Fusion Mask Module

Building upon prior VOS research [6, 8, 31] emphasizing the incorporation of cross-attention mechanisms for collecting spatio-temporal information within the mask embedding space, we introduce the Temporal Fusion Mask Module (TFMM). The goal of TFMM is to produce mask embeddings that are infused with temporal information from the previous frames.

In Eqn. 6, we start by adding positional embedding (PE) to the current frame image embedding. Then, we apply LayerNorm (LN) and use the query head of TFMM to obtain Q . Similarly, in Eqn. 7, we add positional embedding to the concatenated image embeddings from memory, apply LayerNorm, and use the key head of TFMM to obtain K . In Eqn. 8, we apply LayerNorm to the concatenated mask embeddings and use the value head of TFMM to obtain V .

$$Q = Q_{TFMM}(LN(PE_{TFMM} + E_{Enc}^{i+1})) \quad (6)$$

$$K = K_{TFMM}(LN(PE_{TFMM} + E_{Enc}^{(i-n+1:i)})) \quad (7)$$

$$V = V_{TFMM}(LN(E_{MaskEnc}^{(i-n+1:i)})) \quad (8)$$

We utilize the standard cross-attention mechanism with Q , K , and V , followed by an MLP (with skip connection) and LayerNorm, to obtain the mask embeddings $Output_{TFMM}$, as described in Eqn. 9-11.

$$O_{TFMM} = Attention_{TFMM}(Q, K, V) \quad (9)$$

$$O_{TFMM} = MLP(O_{TFMM}) + O_{TFMM} \quad (10)$$

$$Output_{TFMM} = LN(O_{TFMM}) \quad (11)$$

3.2.2 Memory Prior Affinity Module

Drawing inspiration from [31, 40], we introduce the MPAM module as an extension. Previous approaches involve using

either the sole current image post-image encoder processing [40] or concatenating the current frame with the prior [31] mask before undergoing encoder processing. Instead of relying on preceding masks, we extract values from our Temporal Fusion Mask Module $Output_{TFMM}$. These values offer a generalized representation of mask embeddings that are subsequently concatenated (\oplus) with the current frame’s image embedding E_{Enc}^{i+1} . This enhances information flow compared to utilizing the previous frame in its raw form. This method leverages existing SAM encoders, eliminating the necessity to train an additional encoder as typically needed when concatenating images and masks within the original image space. We use an embedding dimension of 128 for q_k , q_v , m_k and m_v and define them as follows

$$q = LN(PE_{MPAM} + E_{Enc}^{i+1} \oplus Output_{TFMM}) \quad (12)$$

$$k = LN(PE_{MPAM} + E_{Enc}^{(i:i-n+1)}) \quad (13)$$

$$q_k, q_v = Linear_1(q), Linear_2(q) \quad (14)$$

$$m_k, m_v = Linear_3(k), Linear_4(k) \quad (15)$$

Using the values defined in Eqn. 14-15, we compute V_{MPAM} and the dense embedding from our module as follows

$$V_{MPAM} = q_v \oplus Attention(q_k, m_k, m_v) \quad (16)$$

$$E_{DenseEnc}^{i+1} = MLP(V_{MPAM}) \quad (17)$$

Finally, we pass both $E_{DenseEnc}^{i+1}$ and E_{Enc}^{i+1} to the decoder SAM_{Dec} to get the predicted mask M^{i+1} as described in Eqn. 5.

4. Experiments

Here, we analyze our proposed framework on the CAD and MoCA-Mask datasets.

4.1. Datasets

We use three publicly available camouflage datasets: COD10K [18], CAD [3], and MoCA-Mask [10].

COD10K Currently, it is the largest COD dataset available, which consists of 5066 camouflaged images (3040 for training, 2026 for testing), 1934 non-camouflaged images, and 3000 background images divided into 10 superclasses and 78 sub-classes (69 camouflaged, 9 non-camouflaged).

CAD CamouflagedAnimalDataset consists of 9 video sequences that were extracted from YouTube videos and contain hand-labeled ground-truth masks on every 5th frame.

MoCA-Mask MoCA-Mask dataset consists of 87 video sequences (71 for training, 16 for testing) which extend the bounding box annotations provided by the original Moving Camouflaged Animal (MoCA) dataset [26] to dense segmentation masks

Table 1. Quantitative results on different VCOD Datasets. The best results are highlighted in **bold**

Model	MoCA-Mask w/o pseudo labels						MoCA-Mask with pseudo labels						CAD					
	$S_\alpha \uparrow$	$F_\beta^w \uparrow$	$E_\phi \uparrow$	$M \downarrow$	mDic	mIoU	$S_\alpha \uparrow$	$F_\beta^w \uparrow$	$E_\phi \uparrow$	$M \downarrow$	mDic	mIoU	$S_\alpha \uparrow$	$F_\beta^w \uparrow$	$E_\phi \uparrow$	$M \downarrow$	mDic	mIoU
EGNet [66]	0.547	0.110	0.574	0.035	0.143	0.096	0.546	0.105	0.573	0.034	0.135	0.090	0.619	0.298	0.666	0.044	0.324	0.243
BASnet [43]	0.561	0.154	0.598	0.042	0.190	0.137	0.537	0.114	0.579	0.045	0.135	0.100	0.639	0.349	0.773	0.054	0.393	0.293
CPD [54]	0.561	0.121	0.613	0.041	0.162	0.113	0.550	0.117	0.613	0.038	0.147	0.104	0.622	0.289	0.667	0.049	0.330	0.239
PraNet [17]	0.614	0.266	0.674	0.030	0.311	0.234	0.568	0.171	0.576	0.045	0.211	0.152	0.629	0.352	0.763	0.042	0.378	0.290
SINet [16]	0.598	0.231	0.699	0.028	0.276	0.202	0.574	0.185	0.655	0.030	0.221	0.156	0.636	0.346	0.775	0.041	0.381	0.283
SINet-v2 [18]	0.588	0.204	0.642	0.031	0.245	0.180	0.571	0.175	0.608	0.035	0.211	0.153	0.653	0.382	0.762	0.039	0.413	0.318
PNS-Net [22]	0.544	0.097	0.510	0.033	0.121	0.101	0.576	0.134	0.562	0.038	0.189	0.133	0.655	0.325	0.673	0.048	0.384	0.290
RCRNet [56]	0.555	0.138	0.527	0.033	0.171	0.116	0.597	0.174	0.583	0.025	0.194	0.137	0.627	0.287	0.666	0.048	0.309	0.229
MG [57]	0.530	0.168	0.561	0.067	0.181	0.127	0.547	0.165	0.537	0.095	0.197	0.141	0.594	0.336	0.691	0.059	0.368	0.268
SLT-Net [10]	0.631	0.311	0.759	0.027	0.360	0.272	0.656	0.357	0.785	0.021	0.397	0.310	0.696	0.481	0.845	0.030	0.493	0.401
SAM [25]	0.667	0.547	0.733	0.138	0.559	0.469	0.650	0.517	0.722	0.140	0.537	0.441	0.490	0.211	0.460	0.184	0.214	0.152
SAM-PM (Ours)	0.728	0.567	0.813	0.009	0.594	0.502	0.695	0.464	0.732	0.011	0.497	0.416	0.729	0.602	0.746	0.018	0.594	0.493




Figure 4. Visual comparison of some recent VCOD models with ours on MoCA-Mask




Figure 5. Visual comparison of some recent VCOD models with ours on CAD Dataset.

4.2. Evaluation Metrics

Six commonly used metrics are employed for VCOD tasks, which include mean absolute error M which evaluates the pixel level accuracy between prediction and labeled masks, Weighted F-measure F_β^w [35], Enhanced-alignment measure E_ϕ [15] which simultaneously assesses pixel-level correspondence and overall image statistical alignment, structure measure S_α [14] which evaluates region-aware and object-aware structural similarity, mean Dice and mean IoU measures the similarity and overlapping between two sets of data and masks respectively. Larger F_β^w , E_ϕ , S_α , mDice, mIoU and smaller M indicate better segmentation performance.

4.3. Implementation Details

The images are resized to 1024 x 1024 for static image pre-training and main training after applying RandomAffine,

Gaussian Blur, ColorJitter, Random Horizontal Flip, and Random GrayScale augmentation. We use SAM-L with a frozen architecture, where only the PM is trainable. We use a combination of Focal loss [29], Dice Loss [39], and MSE Loss between the IoU prediction of SAM and the IoU of the predicted mask with the ground truth mask in a ratio of 20:1:1 following [25].

We adopt a two-stage training approach to ensure stability throughout the training process. First, we pre-train the model on the COD10K dataset, where we randomly pick three frames following [31] to form a training sample for the main training of the propagation module. All the training was done with A40 GPU. For optimization, we use AdamW [24, 32] with a weight decay of 1e-2 and a batch size of 4 for both stages. The pre-training is done for 7.6k iterations with an initial learning rate of 4e-4, and the main training is performed for 2.5k iterations with an initial learning rate

Table 2. Ablation results on different VCOD Datasets.

Model	MoCA-Mask w/o pseudo labels						CAD					
	$S_\alpha \uparrow$	$F_\beta^w \uparrow$	$E_\phi \uparrow$	$M \downarrow$	mDic	mIoU	$S_\alpha \uparrow$	$F_\beta^w \uparrow$	$E_\phi \uparrow$	$M \downarrow$	mDic	mIoU
SAM-PM	0.728	0.567	0.813	0.009	0.594	0.502	0.729	0.602	0.746	0.018	0.594	0.493
SAM-PM without positional embeddings	0.690	0.476	0.776	0.012	0.502	0.420	0.712	0.572	0.738	0.022	0.560	0.457
SAM-PM without MPAM	0.661	0.424	0.673	0.011	0.440	0.367	0.673	0.491	0.627	0.023	0.472	0.382

of 5e-5. We employ a multi-step learning rate decay schedule, reducing the learning rate by half after 3.8k iterations during the pre-training stage.

We initiate the process during inference with the first frame and its corresponding ground truth mask. Subsequent masks are predicted by concatenating all frames and masks in the memory with a set memory length of 2. To manage memory, we remove the oldest frame whenever the memory reaches full capacity.

4.4. Results

Table 1 illustrates our model performance with recent state-of-the-art VCOD methods. For the evaluation of SAM, we provide the geometric center of the animal as a point prompt, which is sampled from the ground truth. Since SAM-PM utilizes the ground truth of the initial frame as input to the model, we exclude it from the metric calculations.

We provide qualitative results on the MoCA-Mask and CAD datasets in Figure 4 and Figure 5, respectively. When compared with the previous SOTA SLT-Net, our method achieves average performance gains of 82.31%, 15.37%, 7.11%, 200%, 65%, and 84.55% in terms of F_β^w , S_α , E_ϕ , M , mDic, mIoU respectively on MoCA-Mask dataset.

This demonstrates that our model can segment out accurately and locate camouflaged objects in various complex and challenging scenarios. Our method outperforms all the previous methods on the CAD dataset. It surpasses the previous SOTA SLT-Net by 25.16%, 4.74%, 66.67%, 20.48% and 22.94% in terms of F_β^w , S_α , M , mDic, mIoU respectively.

4.5. Ablation

We perform ablation on the MoCA-Mask and CAD dataset to analyze the performance contribution of positional embedding and affinity in the propagation framework. The results are tabulated in table 2.

Positional Embedding. Incorporating positional embedding contributes to a performance boost of 5.5% on S_α , 19.11% on F_β^w , 4.77% on E_ϕ , 33.33% on M , 18.32% on mDic, and 19.5% on mIoU for MoCA-Mask Dataset and a performance boost of 2.39% on S_α , 5.24% on F_β^w , 1.08%

on E_ϕ , 22.22% on M , 6.07% on mDic and 7.87% on mIoU for CAD dataset. The positional embedding is introduced to preserve spatial information in embedding generated by the SAM encoder. Camouflaged images often involve intricate patterns and subtle variations in colour and texture. Preserving the spatial arrangement of these patterns is crucial for distinguishing between the camouflage and the background. The patterns are designed to disrupt the perception of shapes and structures. If a model treats different spatial arrangements as equivalent (symmetric), it may struggle to differentiate between the camouflage and the background, which is now solved by using positional embedding.

MPAM. Training the model with MPAM results in a performance increase of 10.13% on S_α , 33.72% on F_β^w , 20.8% on E_ϕ , 22.22% on M , 35% on mDic and 36.78% on mIoU for MoCA-Mask dataset and a performance increase of 8.32% on S_α , 22.61% on F_β^w , 18.97% on E_ϕ , 27.78% on M , 25.85% on mDic and 29.05% on mIoU for CAD dataset. It is a crucial module in refining the superficial predictions from the previous module. Camouflaged images tend to produce low-confidence fuzzy features, which are now rectified with the help of the affinity model, leading to improvement in performance.

5. Limitations

Even though the PM is extremely small and efficient in terms of parameters, the throughput and latency of SAM-PM are heavily bottlenecked due to the large encoder (contributing to more than 98% parameters in SAM-PM) that is being used. We have also observed the model failing sometimes for smaller targets. Since the encoder is one of the crucial components in effectively getting the essential target representation, our PM is limited by the encoder’s performance in COD tasks. Our SAM-PM approach is directed towards parameter-efficient fine-tuning therefore, further encoder fine-tuning might be essential for better performance. Currently, we employ a naive queue-based memory bank, which can be improved further by using complex strategies, such as memory updation and contrastive items in CTVIS [64].

6. Conclusion

In summary, our novel SAM Propagation Module (SAM-PM) presents a noteworthy improvement in Video Camouflaged Object Detection, outperforming previous state-of-the-art models. Our Propagation Module seamlessly integrates with the Segment Anything Model (SAM), requiring less than 1% parameters for training. Its successful integration proves its effectiveness in tackling camouflaged objects in videos and suggests its potential to enhance large foundation models for real-world applications. The demonstrated performance improvements and minimal overhead of SAM-PM offer a promising avenue for advancing the field, opening possibilities for more efficient and accurate video-based segmentation tasks.

References

- [1] Richard C Atkinson and Richard M Shiffrin. Human memory: A proposed system and its control processes. In *Psychology of learning and motivation*, pages 89–195. Elsevier, 1968. 3
- [2] Sara Beery, Guanhang Wu, Vivek Rathod, Ronny Votel, and Jonathan Huang. Context r-cnn: Long term temporal context for per-camera object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13075–13085, 2020. 2
- [3] Pia Bideau and Erik Learned-Miller. It’s moving! a probabilistic model for causal motion segmentation in moving camera videos. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 433–449. Springer, 2016. 2, 5
- [4] Pia Bideau, Rakesh R Menon, and Erik Learned-Miller. Moa-net: self-supervised motion segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 2
- [5] Tianrun Chen, Lanyun Zhu, Chaotao Ding, Runlong Cao, Shangzhan Zhang, Yan Wang, Zejian Li, Lingyun Sun, Papa Mao, and Ying Zang. Sam fails to segment anything?—sam-adapter: Adapting sam in underperformed scenes: Camouflage, shadow, and more. *arXiv preprint arXiv:2304.09148*, 2023. 2
- [6] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European Conference on Computer Vision*, pages 640–658. Springer, 2022. 3, 5
- [7] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Modular interactive video object segmentation: Interaction-to-mask, propagation and difference-aware fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5559–5568, 2021. 3
- [8] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *Advances in Neural Information Processing Systems*, 34:11781–11794, 2021. 3, 5
- [9] Xuelian Cheng, Yiran Zhong, Yuchao Dai, Pan Ji, and Hongdong Li. Noise-aware unsupervised deep lidar-stereo fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6339–6348, 2019. 2
- [10] Xuelian Cheng, Huan Xiong, Deng-Ping Fan, Yiran Zhong, Mehrtash Harandi, Tom Drummond, and Zongyuan Ge. Implicit motion handling for video camouflaged object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13864–13873, 2022. 2, 5, 6
- [11] Suhwan Cho, Minhyeok Lee, Seunghoon Lee, Chaewon Park, Donghyeong Kim, and Sangyoun Lee. Treating motion as option to reduce motion dependency in unsupervised video object segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5140–5149, 2023. 3
- [12] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, Philip HS Torr, and Song Bai. Mose: A new dataset for video object segmentation in complex scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20224–20234, 2023. 3
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 4
- [14] Deng-Ping Fan, Ming-Ming Cheng, Yun Liu, Tao Li, and Ali Borji. Structure-measure: A new way to evaluate foreground maps. In *Proceedings of the IEEE international conference on computer vision*, pages 4548–4557, 2017. 6
- [15] Deng-Ping Fan, Cheng Gong, Yang Cao, Bo Ren, Ming-Ming Cheng, and Ali Borji. Enhanced-alignment measure for binary foreground map evaluation. *arXiv preprint arXiv:1805.10421*, 2018. 6
- [16] Deng-Ping Fan, Ge-Peng Ji, Guolei Sun, Ming-Ming Cheng, Jianbing Shen, and Ling Shao. Camouflaged object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2777–2787, 2020. 2, 6
- [17] Deng-Ping Fan, Ge-Peng Ji, Tao Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Prant: Parallel reverse attention network for polyp segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 263–273. Springer, 2020. 2, 6
- [18] Deng-Ping Fan, Ge-Peng Ji, Ming-Ming Cheng, and Ling Shao. Concealed object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2, 5, 6
- [19] Joanna R Hall, Innes C Cuthill, Roland Baddeley, Adam J Shohet, and Nicholas E Scott-Samuel. Camouflage, detection and identification of moving targets. *Proceedings of the Royal Society B: Biological Sciences*, 280(1758):20130064, 2013. 2
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 4

- [21] Zhou Huang, Hang Dai, Tian-Zhu Xiang, Shuo Wang, Huai-Xin Chen, Jie Qin, and Huan Xiong. Feature shrinkage pyramid for camouflaged object detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5557–5566, 2023. 2
- [22] Ge-Peng Ji, Yu-Cheng Chou, Deng-Ping Fan, Geng Chen, Huazhu Fu, Debesh Jha, and Ling Shao. Progressively normalized self-attention network for video polyp segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 142–152. Springer, 2021. 6
- [23] Pan Ji, Yiran Zhong, Hongdong Li, and Mathieu Salzmann. Null space clustering with applications to motion segmentation and face clustering. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 283–287. IEEE, 2014. 2
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [25] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 2, 4, 6
- [26] Hala Lamdouar, Charig Yang, Weidi Xie, and Andrew Zisserman. Betrayed by motion: Camouflaged object discovery via motion segmentation. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 2, 5
- [27] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2
- [28] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 2
- [29] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 6
- [30] Ting Liu, Yao Zhao, Yunchao Wei, Yufeng Zhao, and Shikui Wei. Concealed object detection for activate millimeter wave image. *IEEE Transactions on Industrial Electronics*, 66(12): 9909–9917, 2019. 2
- [31] Yong Liu, Ran Yu, Fei Yin, Xinyuan Zhao, Wei Zhao, Weihao Xia, and Yujiu Yang. Learning quality-aware dynamic memory for video object segmentation. In *European Conference on Computer Vision*, pages 468–486. Springer, 2022. 3, 5, 6
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [33] Xiankai Lu, Wenguan Wang, Chao Ma, Jianbing Shen, Ling Shao, and Fatih Porikli. See more, know more: Unsupervised video object segmentation with co-attention siamese networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3623–3632, 2019. 3
- [34] Xiankai Lu, Wenguan Wang, Martin Danelljan, Tianfei Zhou, Jianbing Shen, and Luc Van Gool. Video object segmentation with episodic graph memory networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 661–679. Springer, 2020. 3
- [35] Ran Margolin, Lihi Zelnik-Manor, and Ayellet Tal. How to evaluate foreground maps? In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 248–255, 2014. 6
- [36] Haiyang Mei, Ge-Peng Ji, Ziqi Wei, Xin Yang, Xiaopeng Wei, and Deng-Ping Fan. Camouflaged object segmentation with distraction mining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8772–8781, 2021. 2
- [37] Etienne Meunier, Anaïs Badoual, and Patrick Bouthemy. Em-driven unsupervised learning for efficient motion segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4462–4473, 2022. 2
- [38] Jeff Michels, Ashutosh Saxena, and Andrew Y Ng. High speed obstacle avoidance using monocular vision and reinforcement learning. In *Proceedings of the 22nd international conference on Machine learning*, pages 593–600, 2005. 2
- [39] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016. 6
- [40] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9226–9235, 2019. 3, 5
- [41] Youwei Pang, Xiaoqi Zhao, Tian-Zhu Xiang, Lihe Zhang, and Huchuan Lu. Zoom in and out: A mixed-scale triplet network for camouflaged object detection. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 2160–2170, 2022. 2
- [42] Jiyang Qi, Yan Gao, Yao Hu, Xinggang Wang, Xiaoyu Liu, Xiang Bai, Serge Belongie, Alan Yuille, Philip HS Torr, and Song Bai. Occluded video instance segmentation: A benchmark. *International Journal of Computer Vision*, 130(8): 2022–2039, 2022. 3
- [43] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. Basnet: Boundary-aware salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7479–7489, 2019. 6
- [44] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [45] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Confer-*

- ence on Machine Learning, pages 8821–8831. PMLR, 2021. [2](#)
- [46] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. [2](#)
- [47] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J Black. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12240–12249, 2019. [2](#)
- [48] Pavel Tokmakov, Jie Li, and Adrien Gaidon. Breaking the” object” in video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22836–22845, 2023. [3](#)
- [49] Wenguan Wang, Jianbing Shen, Ruigang Yang, and Fatih Porikli. Saliency-aware video object segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 40(1):20–33, 2017. [3](#)
- [50] Wenguan Wang, Xiankai Lu, Jianbing Shen, David J Crandall, and Ling Shao. Zero-shot video object segmentation via attentive graph neural networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9236–9245, 2019.
- [51] Wenguan Wang, Hongmei Song, Shuyang Zhao, Jianbing Shen, Sanyuan Zhao, Steven CH Hoi, and Haibin Ling. Learning unsupervised video object segmentation through visual attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3064–3074, 2019. [3](#)
- [52] Weiyao Wang, Matt Feiszli, Heng Wang, and Du Tran. Unidentified video objects: A benchmark for dense, open-world segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10776–10785, 2021. [3](#)
- [53] Yu-Huan Wu, Shang-Hua Gao, Jie Mei, Jun Xu, Deng-Ping Fan, Rong-Guo Zhang, and Ming-Ming Cheng. Jcs: An explainable covid-19 diagnosis system by joint classification and segmentation. *IEEE Transactions on Image Processing*, 30:3113–3126, 2021. [2](#)
- [54] Zhe Wu, Li Su, and Qingming Huang. Cascaded partial decoder for fast and accurate salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3907–3916, 2019. [6](#)
- [55] Jinnan Yan, Trung-Nghia Le, Khanh-Duy Nguyen, Minh-Triet Tran, Thanh-Toan Do, and Tam V Nguyen. Mirror-net: Bio-inspired camouflaged object segmentation. *IEEE Access*, 9:43290–43300, 2021. [2](#)
- [56] Pengxiang Yan, Guanbin Li, Yuan Xie, Zhen Li, Chuan Wang, Tianshui Chen, and Liang Lin. Semi-supervised video salient object detection using pseudo-labels. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7284–7293, 2019. [6](#)
- [57] Charig Yang, Hala Lamdouar, Erika Lu, Andrew Zisserman, and Weidi Xie. Self-supervised video object segmentation by motion grouping. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7177–7188, 2021. [2, 6](#)
- [58] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. *Advances in Neural Information Processing Systems*, 35:36324–36336, 2022. [3](#)
- [59] Zhao Yang, Qiang Wang, Luca Bertinetto, Weiming Hu, Song Bai, and Philip HS Torr. Anchor diffusion for unsupervised video object segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 931–940, 2019. [2](#)
- [60] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. In *European Conference on Computer Vision*, pages 332–348. Springer, 2020. [3](#)
- [61] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. *Advances in Neural Information Processing Systems*, 34:2491–2502, 2021.
- [62] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by multi-scale foreground-background integration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):4701–4712, 2021.
- [63] Zongxin Yang, Jiaxu Miao, Xiaohan Wang, Yunchao Wei, and Yi Yang. Scalable multi-object identification for video object segmentation. *arXiv preprint arXiv:2203.11442*, 2022. [3](#)
- [64] Kaining Ying, Qing Zhong, Weian Mao, Zhenhua Wang, Hao Chen, Lin Yuanbo Wu, Yifan Liu, Chengxiang Fan, Yunzhi Zhuge, and Chunhua Shen. Ctvis: Consistent training for online video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 899–908, 2023. [7](#)
- [65] Qiang Zhai, Xin Li, Fan Yang, Chenglizhao Chen, Hong Cheng, and Deng-Ping Fan. Mutual graph learning for camouflaged object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12997–13007, 2021. [2](#)
- [66] Jia-Xing Zhao, Jiang-Jiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. Egnet: Edge guidance network for salient object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8779–8788, 2019. [6](#)