

AlignSAM: Aligning Segment Anything Model to Open Context via Reinforcement Learning

Duojun Huang^{1,2} Xinyu Xiong¹ Jie Ma¹ Jichang Li^{1,3} Zequn Jie⁴ Lin Ma⁴ Guanbin Li^{1,2†}

¹School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China

²GuangDong Province Key Laboratory of Information Security Technology

³The University of Hong Kong ⁴Meituan

Abstract

Powered by massive curated training data, Segment Anything Model (SAM) has demonstrated its impressive generalization capabilities in open-world scenarios with the guidance of prompts. However, the vanilla SAM is class-agnostic and heavily relies on user-provided prompts to segment objects of interest. Adapting this method to diverse tasks is crucial for accurate target identification and to avoid suboptimal segmentation results. In this paper, we propose a novel framework, termed AlignSAM, designed for automatic prompting for aligning SAM to an open context through reinforcement learning. Anchored by an agent, AlignSAM enables the generality of the SAM model across diverse downstream tasks while keeping its parameters frozen. Specifically, AlignSAM initiates a prompting agent to iteratively refine segmentation predictions by interacting with the foundational model. It integrates a reinforcement learning policy network to provide informative prompts to the foundational models. Additionally, a semantic recalibration module is introduced to provide fine-grained labels of prompts, enhancing the model’s proficiency in handling tasks encompassing explicit and implicit semantics. Experiments conducted on various challenging segmentation tasks among existing foundation models demonstrate the superiority of the proposed AlignSAM over state-of-the-art approaches. Project page: <https://github.com/Duojun-Huang/AlignSAM-CVPR2024>.

1. Introduction

Compared to traditional computer vision tasks [13, 24, 25, 52, 53], segmentation stands as a fundamental task, playing a pivotal role in visual understanding systems. Accord-

[†]Corresponding author is Guanbin Li. This work was supported in part by the National Natural Science Foundation of China (NO. 62322608), and in part by the Open Project Program of the Key Laboratory of Artificial Intelligence for Perception and Understanding, Liaoning Province (AIPU, No. 20230003).

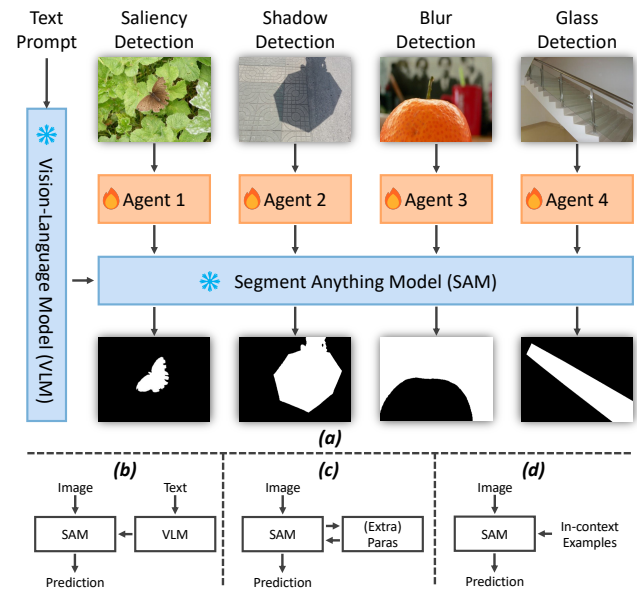


Figure 1. Conceptual comparisons of our method and previous approaches. Frozen and learnable parameters are highlighted in blue and orange, respectively. (a) The proposed method. (b) Text-guided methods [28, 46]. (c) PEFT methods [4, 59]. (d) In-context learning methods [29, 60]. Observed that the proposed agent-based auto-prompting effectively grasps vision and linguistic cues, unleashing the potential of the foundation segmentation model in various contexts, such as saliency detection, shadow detection, blur detection, and glass detection.

ing to different semantic criteria for grouping pixels, various downstream segmentation tasks have emerged, such as saliency detection [34, 50, 64], shadow detection [12, 66], and glass-like object detection [8, 57]. Although significant progress has been achieved, the development of a unified framework that can accommodate the wide variations inherent in the formulations of diverse segmentation tasks continues to pose a challenge.

Recent breakthroughs in Vision Foundation Models (VFM) have demonstrated impressive capabilities in zero-




Figure 2. Segmentation results for CLIP-Surgery [27] using different prompts. (a) and (b) illustrate the prompts of “dog” and “salient object”, respectively. Observed that vision-language models excel in processing explicit semantics but often struggle with implicit semantics.

shot segmentation within open scenarios [16, 48, 49, 68]. Powered by the large-scale high-quality training dataset, Segment Anything Model (SAM) [16] has demonstrated its powerful capacity to be generalized into different tasks and data distributions unseen during training. However, SAM heavily relies on manual prompts such as bounding boxes and points to segment objects of interest since it is pre-trained in a class-agnostic manner. Consequently, achieving target-aware adaptation of SAM for different tasks plays a crucial role in bridging the gap between class-agnostic predictions and industrial needs.

Language-guided segmentation has emerged as a promising approach to customize SAM for specific downstream tasks. By leveraging the cross-modal alignment of the vision-language model [35], referring segmentation can be accomplished by accurately depicting the foreground targets. Nevertheless, there are numerous visual targets that are difficult to accurately express by textual description. For example, the salient object detection task requires visual comparison among different regions within an image, which poses challenges in aligning it with the text encoder as evident in Figure 2.

One feasible way to capture contextual information is to utilize sufficient training samples to fine-tune the foundation models. Owing to the substantial number of parameters in foundation models, conducting full fine-tuning for adaptation leads to a significant computational burden. Recent advancements have been directed toward developing Parameter-Efficient Fine-Tuning (PEFT) methods. These methods integrate adapters [4, 10] or LoRA [11] blocks while keeping the backbone network frozen, as indicated in Figure 1(c). They aspire to attain promising performance comparable to the paradigm of full finetuning. Nevertheless, PEFT methods necessitate the calculation of gradients in the intermediate layers of the backbone network, which may be inaccessible due to privacy concerns. Furthermore, these methods remain highly reliant on an abundance of

training data. To achieve sample-efficient adaptation, recent researches [29, 60] explore in-context learning by leveraging similarity matching to generate point prompts for personalized adaptation using limited training instances. However, similarity calculations assume are prone to interference, impeding the capture of implicit semantics.

Inspired by the model-free spirit of reinforcement learning (RL), we introduce a unified framework for adapting SAM to diversified scenarios while keeping parameters of the backbone network frozen. By formulating automatic prompting as a sequential decision-making process, an agent is trained to imitate human annotators to recommend prompting positions. The RL framework is built by an actor branch and a critic branch networks, which collaborate to learn the optimal prompting policy to efficiently refine the segmentation outputs. Moreover, a semantic recalibration module is additionally introduced to provide reliable labels for the selected positions. It contributes to delineating foreground and background areas for downstream tasks of different types. Concretely, for tasks with explicit semantic, a cross-modal attention module integrates visual and linguistic information from the context. In tasks with implicit semantics, a visual-similarity learning branch is presented to capture implicit contextual concepts.

The contributions can be summarized as follows, with Figure 1(a) showcasing its conceptual illustration.

- A general approach, termed AlignSAM, is proposed to optimize the automatic prompting policy for efficiently adapting foundation model to downstream tasks. By constructing universal actions, state and reward signals, it is able to handle various types of downstream tasks within a unified framework.
- A semantic recalibration module is introduced to provide precise prompting labels for adapting the vision foundation model to tasks with explicit and implicit semantics.
- Experiments conducted on various challenging segmentation tasks among existing foundation models demonstrate the superiority of the proposed method over state-of-the-art approaches for efficient adaptation.

2. Related Work

Image segmentation. Different from traditional tasks [18–23, 62], leveraging richly annotated datasets [5, 6] and advanced pretrained feature extractors [9, 43, 63], deep image segmentation has yielded notable results in diverse applications [26, 31, 32, 54, 61], primarily through the development of complex network architectures and training methodologies. For example, DeepLabV3Plus [3] in semantic segmentation utilizes a spatial pyramid pooling structure for multiscale contextual information. In shadow detection, FDRNet [67] enhances training samples by adjusting brightness to differentiate between dark non-shadow and bright shadow areas. Meanwhile, for saliency

detection, MENet [50] proposes a method for progressively aggregating and refining features to address accuracy challenges in cluttered scenes. However, these specialized methods lack versatility, limiting their transferability to different segmentation tasks.

Vision foundation models. Leveraging large-scale datasets such as [38, 39], vision foundation models have made significant advancements in various tasks of computer vision, encompassing image classification [14, 35], segmentation [16, 49, 68], and generation [37]. CLIP [35], trained on massive image-caption pairs, exhibits robust zero-shot learning capabilities in cross-modal task. Segment Anything Model [16] introduces a large-scale segmentation dataset SA-1B and a training framework to enable prompt-driven segmentation in a zero-shot manner. Painter [48, 49] integrates multiple segmentation tasks within an in-context learning framework. Despite their effectiveness, it remains challenging for these models to adapt to diversified downstream tasks while preserving great interpretability as demonstrated in Sec. 1.

Parameter-Efficient Fine-Tuning. Foundation models inherently possess a massive amount of parameters. Therefore, directly fine-tuning all of these parameters for each downstream task results in significant computational costs, which limits the efficiency of specialized applications. Recent efforts have been made to optimize extra learnable parameters on a small-scale, rather than those of the backbone network. Existing approaches for efficient fine-tuning of visual parameters generally fall into three categories: prompt tuning [15], adapter methods [4, 10, 55], and low-rank adaptation (LoRA) [11]. Specifically, VPT [15] injects a small number of learnable parameters into the Transformer’s input space and keeps the backbone frozen during the fine-tuning stage. Adapter-based techniques [4, 10, 55], such as SAM-Adapter [4], integrate tunable modules into pre-trained large models to suit domain-specific tasks, where MLP layers are inserted into each transformer block of the image encoder. Furthermore, LoRA [11] hypothesizes that the change matrix of weights during adaptation has a low rank, and therefore only tunes a pair rank decomposition matrix of the original linear layers weights, while keeping the pre-trained weights frozen. However, most of these algorithms heavily rely on sufficient training data during adaptation and require gradient calculation of the intermediate layers of the foundation model, resulting the lack of sample-efficiency and model-agnostic property.

Reinforcement learning for computer vision. Reinforcement learning has been successfully applied to various vision scenarios including image classification [33, 42], object tracking [30, 58] and semantic segmentation [2, 36].

Notably, RAM [33] formulated attention-based image processing as a sequential control problem and integrated reinforcement learning and recurrent neural network for object recognition task. The reinforced active segmentation method [2] achieves comparable results with significantly less annotated data than other weakly supervised approaches. Inspired by the model-free spirit of reinforcement learning, we train a universal agent to efficiently prompt the vision foundation model, constructing a unified framework that can scale to diversified downstream tasks.

3. Methodology

In this section, we first present a revisit of the Segment Anything Model (SAM) [16] and then introduce the task to be addressed in this paper. Next, we delve into how to achieve the pipeline of the proposed framework and introduce the semantic recalibration module algorithm. Finally, we outline the training and evaluation processes employed in our approach.

3.1. Preliminary

Revisit of Segment Anything Model. The Segment Anything Model (SAM) is proposed as a vision foundation model for achieving a unified capacity of segmentation. In this case, given a set of prompts, i.e., foreground and background points, bounding boxes, or a mask, SAM is capable of segmenting the prompted target. In detail, the SAM model is made up of an image encoder $\mathcal{E}_I(\cdot)$, a prompt encoder $\mathcal{E}_P(\cdot)$ and a mask decoder $\mathcal{D}_M(\cdot)$. During model inference, the image encoder and the prompt encoder are first utilized to encode the given image x and a prompt set P respectively, which can be formulated as follows:

$$F_I = \mathcal{E}_I(x), \quad P_t = \mathcal{E}_P(P), \quad (1)$$

where $F_I \in R^{h \times w \times c}$ and $P_t \in R^{k \times c}$, with k denoting the amount of prompts. Further, the image feature F_I and prompt embedding P_t will be fed into $\mathcal{D}_M(\cdot)$ to generate the prediction mask M as follows:

$$M = \mathcal{D}_M(F_I, P_t). \quad (2)$$

Overview. In this paper, we present a novel framework termed AlignSAM for adapting SAM to diversified scenarios while keeping its parameters frozen. First, we formulate automatic prompting as a sequential decision-making process and introduce Target-aware Reinforcement Learning to construct a unified pipeline to handle tasks with diversified objects. Specifically, we train a reinforcement agent to imitate human annotators to recommend prompting positions, which executes a series of prompting actions to refine the segmentation progressively. Moreover, a Semantic Recalibration Module is additionally proposed to delineate foreground and background areas, providing precise label for the selected prompts.

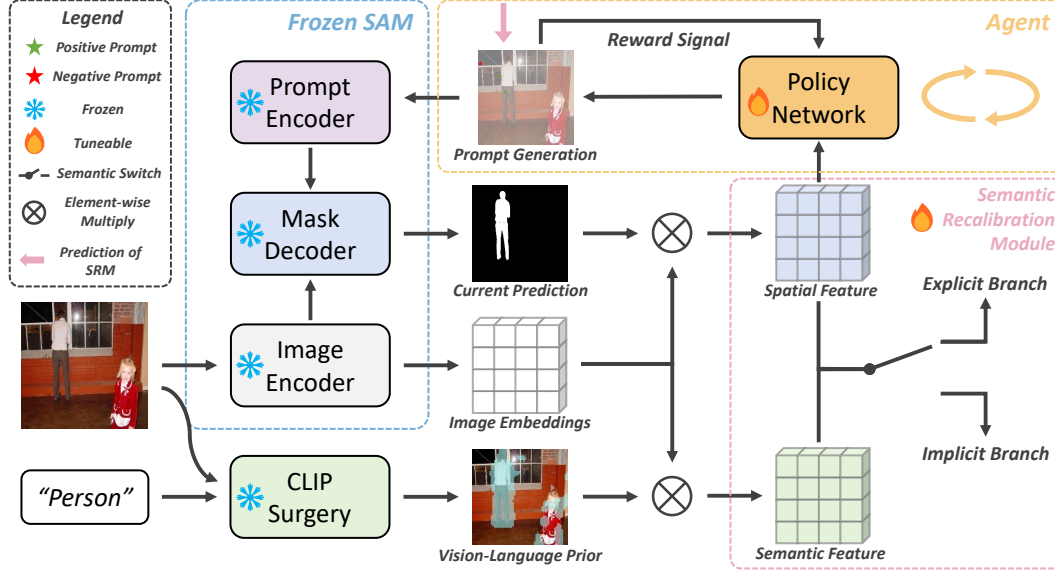


Figure 3. An overview of the proposed AlignSAM, which consists of a Segment Anything Model (SAM), a vision-language model (CLIP-Surgery [27]), a reinforcement learning agent, and a semantic recalibration module. The frozen SAM receives the point prompts generated by the agent and semantic recalibration module, dealing with various downstream scenarios without relying on manual prompting.

3.2. Target-aware Reinforcement Learning

We train a target-aware reinforcement agent to interact with the environment to maximize the accumulated rewards. The objective of the agent is to recommend optimal prompting positions to facilitate progressive refinement of the mask prediction by the segmentation model. To be specific, during each interaction loop, the agent will first select an action according to current state observation and subsequently receive a signal of reward from the environment. Here, we first present how to construct the state space, action space, and the reward function separately in our RL framework, and then illustrate the process of model training.

Action space. To construct the action space, a naive approach is to regard each pixel as a candidate action in the image. However, it is cost intensive and in-efficient. To relieve the computation burden while maintaining the local details of image, we propose to construct action space in the patch-level rather than pixel-level for each image. Specifically, we propose to divide each image into different regions to build the action space for reinforced learning. Given an input image $x \in R^{H \times W \times C}$, we divide it into image patches $x_p \in R^{H' \times W' \times C}$. The candidate action set is the center points of each patch in an image, which can be formulated as follows:

$$A = \{(h_i, w_i) | (h_i, w_i) = center(x_p)\}. \quad (3)$$

Therefore, the action set consists of $N = \frac{H}{H'} \cdot \frac{W}{W'}$ candidate points in an image. For each image, the goal is to select the optimal action among the N candidates at each time-

step, and the process will be terminated after T time-steps in total.

State space. At each timestep, the agent determines which action to execute based on the current state. In order to facilitate the agent to perform gainful decisions, the state must encompass comprehensive information about the environment. Therefore, we combine the visual feature embedding and the iterative prediction to form the state representation. For an image sample x at the timestep t , the state representation can be formulated as:

$$s_t = \mathcal{E}_{\mathcal{I}}(x) \cdot M_{t-1}, \quad (4)$$

where M_{t-1} denotes the mask predicted by the mask decoder of SAM at last step. The feature extracted by the image encoder represents the inherent spatial feature of the sample, while the information derived from the previous prediction maintains the contextual continuity across the sequential dimension. The multiplication operation highlights the features of regions that are more likely to be recognized as foreground in the historical iteration round and diminishes the features of regions that are prone to be classified as background. In this way, the policy network maps the state to a probability distribution over the actions, enabling the optimization of the maximum expected reward in the full training process.

Reward function. Previous RL-based methods for computer vision tasks, such as [2, 17], defined the reward function to improve the task performance between iterations.

Working as a binary segmentation task, the performance metrics are diversified across segmentation scenarios. To relieve the burden of hand-crafted function engineering, we formulate the reward function as the score of querying the target foreground. At each timestep, if the agent’s action points to the foreground region, it earns a reward of **+1**; otherwise, it incurs a penalty of **-1** as feedback.

Training RL model. We employ a widely-used proximal proxy optimization (PPO) algorithm [40] to instantiate the RL framework, which consists of an actor network π_θ and a critic network \mathcal{V}_θ . The actor explores and improves its policy by interacting with the environment by mapping the state to a probability distribution over all the actions. Meanwhile, the critic is trained in a supervised manner to predict the state value of the current policy. The actor and critic work in a cooperative manner to learn the optimal policy.

During the training phase, the RL networks will be trained by E episodes with each episode containing T prompting steps. At the beginning of each episode, the prompt set P_0 for each image is initialized as a pair of positive and negative prompt points randomly selected from the sample itself. After the first inference prompted by P_0 , the initial state s_0 can be calculated according to Eq. 4. At each prompting step t , the agent selects an action from the probability distribution predicted by the actor network, which can be formulated as follows:

$$a_t \sim p_t(\cdot) = \pi_\theta(\cdot | s_t), \quad (5)$$

where p_t represents the probability distribution across the action space. After action a_t is executed, the agent receives a scalar reward r_t from the environment. At each iteration, the tuple comprising the current state, action, reward, and next state is stored in a memory buffer. When an episode is completed, the parameters of π_θ and \mathcal{V}_θ will be updated using stochastic gradient descent (SGD) for K epochs. To form the optimization target, the action value function is first calculated as follows:

$$Q(s_t, a_t) = r_t + \gamma r_{t+1} + \dots + \gamma^{T-t} \mathcal{V}_\theta(s_T), \quad (6)$$

where γ denotes a discount factor to balance between the current reward and future expected values. The action value represents the expected cumulative reward after taking action a_t under state s_t . Furthermore, the advantage function can be formulated as follows:

$$A_t = Q(s_t, a_t) - \mathcal{V}_\theta(s_t), \quad (7)$$

where $\mathcal{V}_\theta(s_t)$ denotes the state value estimated by the critic network. The advantage function measures the expected value of taking specific action under a given state over the average performance. To improve the training stability, importance resampling is utilized during the optimization of

Algorithm 1: Training procedure of the RL model.

Input: Input image x and its ground truth mask y , SAM’s image encoder $\mathcal{E}_I(\cdot)$, prompt encoder $\mathcal{E}_P(\cdot)$, mask decoder $\mathcal{D}_M(\cdot)$

Output: Optimal parameters of actor network π_θ and critic network \mathcal{V}_θ

```

1 for  $e = 1, 2, \dots, E$  do
2    $P_0 = \{a_p, a_n\} \sim y$ 
3    $M_0 = \mathcal{D}_M(\mathcal{E}_I(x), P_0)$ 
4    $s_1 = \mathcal{E}_I(x) \cdot M_0$ 
5   for  $t = 1, 2, \dots, T$  do
6      $a_t \sim \pi_\theta(\cdot | s_t)$ 
7      $P_t = P_{t-1} \cup \{a_t\}$ 
8      $M_t = \mathcal{D}_M(\mathcal{E}_I(x), P_t)$ 
9      $s_t = \mathcal{E}_I(x) \cdot M_t$ 
10     $r_t = +1$  if  $y(a_t)=1$  else  $-1$ 
11  // Compute advantage estimates
12  for  $t = T, T-1, \dots, 1$  do
13     $\hat{A}_t = Q(s_t, a_t) - \mathcal{V}_\theta(s_t)$ 
14  // Update  $\pi_\theta$  and  $\mathcal{V}_\theta$  with  $K$  epochs
15  for  $k = 1, 2, \dots, K$  do
16     $\pi_\theta \leftarrow \pi_\theta + \nabla_{\pi_\theta} L_{act}$ 
17     $\mathcal{V}_\theta \leftarrow \mathcal{V}_\theta - \nabla_{\mathcal{V}_\theta} L_{cri}$ 

```

policy model. Specifically, the probability ratio $\gamma_\theta(t)$ between the previous and current actor models is formulated as follows:

$$\gamma_\theta(t) = \frac{\pi_\theta(a_t | s_t)}{\pi_\theta(a_{t-1} | s_{t-1})}. \quad (8)$$

With the advantage function denoted by A_t at timestep t , the objective function of the actor network is formulated as follows:

$$L_{act} = \mathbb{E}_{t \in [1, T]} [\min(\gamma_\theta(t) A_t, \text{clip}(\gamma_\theta(t), 1-\epsilon, 1+\epsilon) A_t)], \quad (9)$$

where $\text{clip}(\gamma_\theta(t), 1-\epsilon, 1+\epsilon)$ is a clipping function to constrain the value of $\gamma_\theta(t)$ to the interval between $(-\epsilon, 1+\epsilon)$, and ϵ is a hyperparameter to control the magnitude of model update. Concurrently, the critic network is optimized to minimize the discrepancy between the estimated and the actual values. The training objective of the critic network can be formulated as follows:

$$L_{cri} = \mathbb{E}_{t \in [1, T]} [(Q(s_t, a_t) - \mathcal{V}_\theta(s_t))^2]. \quad (10)$$

The training process of the RL model is summarized in Algorithm 1.

3.3. Semantic Recalibration Module

After tuning with reward, the reinforcement learning agent tends to prompt positions of foreground target for the segmentation model. A naive prompting policy is to regard all the selected points as target foreground. However, it is unrealistic and overly idealistic due to the limited capacity of the action space. To generate prompts with reliable label information, we design a semantic recalibration module (SRM) to transition the prompting from a coarse-grained policy to a fine-grained one. This module can generalize to tasks with explicit and implicit semantics by switchable branches. Eventually, the SRM output serves as the reference mask for the selected action to construct the complete prompt information.

Specifically, to enhance the RL agent’s comprehension of contextual information, we propose to recalibrate the state representation by considering both the semantic-dominated and spatial-dominated aspects of the environment. Following prior research [27], we adopt an approach to encode explicit semantic context using the text-guided attention map derived from the visual-language model CLIP [35], achieved without requiring additional training. The attention map, denoted as M_c , is derived from the pairwise similarity between the image feature and the corresponding text feature extracted by the visual-language model. As for the spatial-dominated state, we employ the previous probability map predicted by the mask decoder of SAM to recalibrate the state, as it contains historical information learned from the visual samples. Finally, the semantic-dominated and the spatial-dominated states are formulated as the multiplication between their corresponding attention masks and the inherent feature of the sample, which can be formulated as follows:

$$s_c = \mathcal{E}_{\mathcal{I}}(x) \cdot M_c, s_t = \mathcal{E}_{\mathcal{I}}(I) \cdot M_{t-1}. \quad (11)$$

Subsequently, the SRM receives states representation as input to predict the prompting mask for each candidate position, which can be formulated as:

$$y_r = K(V(s_c, s_t)), \quad (12)$$

where $K(\cdot)$ represents the classification head and $V(\cdot, \cdot)$ is a semantic switch which chooses the implicit branch or explicit branch to execute, depending on the type of downstream task.

Implicit branch. The implicit branch is designed for segmentation tasks with abstract concept such as saliency detection, where the text feature is unavailable, since the prompt “salient object” is implicit semantics, as observed in Figure 2. Concretely, the $V(\emptyset, s_t)$ is utilized in Eq. 12 and formulated as two convolution blocks to extract the targeted spatial feature without noisy attention from the text feature.

Name	Task	Branch
CUHK [41]	Blur Detection	I
SBU [45]	Shadow Detection	E
MSD [56]	Glass Detection	I
DUTS [47]	Saliency Detection	I
PASCAL [6]	Semantic Segmentation	E

Table 1. Summary of datasets used in our benchmark. The term “Branch” refers to the semantic branch allocated for this dataset, with “I” representing the implicit branch and “E” denoting the explicit branch.

Explicit branch. The explicit branch is designed for concrete object segmentation. Since textual information allows for a more comprehensive understanding of the concrete content in the image, we utilize $V(s_c, s_t)$ and formulate it as three convolution layers to aggregate the two states followed by a self-attention layer to capture the relationships and dependencies between different elements in an image. Finally, the feature is sent to two convolution layers to conduct the mask decoding.

Training objective. The training objective to optimize the proposed implicit and explicit branches is a combination of the Dice loss [65] and the binary cross-entropy loss [3], which are both widely adopted in image segmentation:

$$L_{seg} = L_{dice}(y_r, y') + L_{bce}(y_r, y'), \quad (13)$$

where y' denotes the ground truth mask of such a sample x downsampled to the same size of y_r . During training, the SRM module is updated over q iterations after the execution of the action. During evaluation, the prompt set should be initialized as a positive point and a negative point respectively responsible for the predicted probabilities with the highest and the lowest scores by the SRM module. During evaluation, the prompt set will be initialized as a positive point and a negative point with the highest and the lowest scores predicted as foreground areas by the SRM module.

4. Experiments

4.1. Experimental Setups

Datasets. We validate AlignSAM in various challenging benchmark datasets. The detailed dataset partition is shown in Table 1. For quantitative comparison, we follow the previously used evaluation metrics in each task, including mean intersection over union (mIoU), mean absolute error (MAE), balance error rate (BER) [44], F-measure (F_β) [1], and E-measure (E_ϕ) [7]. For each segmentation task, we randomly choose select 50 samples as the training set and utilize the original testing set for evaluation.

Implementation details. All the experiments are performed on a single NVIDIA A100 GPU with 80 GB mem-

Method	Blur CUHK [41]		Shadow SBU [45]		Glass MSD [56]		Saliency DUTS [47]	
	mIoU \uparrow	F_β \uparrow	mIoU \uparrow	BER \downarrow	mIoU \uparrow	F_β \uparrow	E_ϕ \uparrow	MAE \downarrow
	SAMed [59]	55.44	71.68	17.24	42.84	41.35	52.67	76.41
SEEM [68]	60.84	67.45	19.52	48.01	32.35	37.44	59.88	0.326
Painter [48]	18.61	27.25	9.31	47.89	8.62	14.85	81.06	0.113
PerSAM [60]	55.84	71.59	18.68	49.51	31.18	38.00	64.13	0.257
Ours	68.47	76.89	30.78	34.62	45.44	57.28	78.21	0.082
Ours-w/o RL	59.75	70.98	25.62	37.57	33.41	46.91	74.19	0.086
Ours-w/o SRM	66.89	73.15	21.29	42.72	31.92	36.76	30.52	0.587
Ours-w/o MSI	46.92	67.02	18.65	42.91	35.86	47.95	68.59	0.134

Table 2. Comparison results with existing SOTA approaches and ablation study results on four different segmentation tasks. The best performance among all approaches is highlighted in **bold**. “RL”, “SRM”, “MSI” represent reinforcement learning, semantic recalibration module, multi-step interaction, respectively.

Method	mIoU	Bottle	Car	Sheep	Cat	Chair	Dog	Person	Sofa	Cow	Horse
MSA [51]	47.12	31.56	41.05	56.47	60.65	19.79	55.95	38.81	33.97	60.38	61.99
SAMed [59]	51.42	36.13	48.72	60.59	58.46	16.22	71.48	49.08	47.65	74.33	71.72
SEEM [68]	52.30	36.95	44.95	63.93	84.46	19.67	66.40	63.87	47.89	73.74	74.10
Painter [48]	59.27	35.33	61.93	77.29	63.16	29.90	59.82	46.40	55.58	78.71	72.82
PerSAM [60]	53.02	36.03	46.94	64.42	69.39	22.28	67.25	49.07	36.84	68.79	65.56
Ours	62.09	48.09	66.13	73.12	85.99	31.68	79.84	64.72	61.63	72.97	75.10
Ours-w/o RL	54.06	41.69	56.28	69.30	80.80	26.84	69.16	50.83	48.20	62.74	72.91
Ours-w/o SRM	27.73	19.78	27.42	30.33	39.29	14.08	29.16	24.91	30.39	31.32	29.02

Table 3. Comparison results with existing SOTA approaches and ablation study results on Pascal-VOC 2012. We report the mean IoU of all 20 categories and the IoU of 10 randomly selected categories. The best performance among all approaches is highlighted in **bold**.

ory. We export Adam as the optimizer conducted on all experiments, with a learning rate of 1.0×10^{-4} . To construct the action space, we set both height and width of each image to 800, and each patch within an image is assigned dimensions of 80 for both height and width. In the training of the RL agent, the hyperparameters γ and ϵ are assigned the values of 0.99 and 0.20, respectively. Additionally, we set the episode E to 50 and set the epoch K to 20. Furthermore, we fix the number of interaction rounds T to 15. For the implicit branch, Q is set to 1, whereas for the explicit branch, it is set to 5.

Baselines. We compare our proposed AlignSAM with following state-of-the-art (SOTA) approaches based on foundation models. Specifically, MSA [51] and SAMed [59] are parameter-efficient fine-tuning methods based on adapter [10] and LoRA [11]; Painter [48] and SEEM [68] are other foundation segmentation models; PerSAM [60] is an in-context learning variant of SAM.

4.2. Results

Quantitative results. Table 2 showcases the comparative performance of the proposed AlignSAM and other SOTA efficient tuning methods on four distinct segmentation tasks: blur detection, shadow detection, glass detection, and salient object detection. Notably, AlignSAM outperforms other SOTA methods in the majority of reported




Figure 4. Qualitative comparisons between our and other methods. (a) Blur detection. (b) Shadow detection. (c) Glass detection. (d) Semantic segmentation.

benchmarks across various evaluation metrics.

To verify the effectiveness of our method in explicit tasks, we further illustrate the comparison results for 20 classes on the Pascal-VOC dataset as shown in Table 3. Notably, AlignSAM exhibits superior performance across almost all

	Blur mIoU	Shadow mIoU	Glass mIoU	Saliency E_ϕ	Semantic mIoU
Implicit Branch	68.47	19.40	45.44	78.21	46.03
Explicit Branch	47.98	30.78	25.72	69.31	62.09

Table 4. Ablation study results of different branches on SAM. Concerning the explicit branch, the text prompts for segmentation tasks are arranged from left to right as follows: “defocus background”, “shadow”, “glass”, “salient objects” and the category names (such as bottle, car) respectively.

T	1	5	10	15	20
Bottle	32.59	40.74	46.42	48.09	48.70
Car	42.55	60.15	65.04	66.13	66.54
Sheep	54.69	67.52	71.16	73.12	74.00
Cat	74.59	81.37	83.86	85.99	86.13

Table 5. Ablation study results w.r.t. iterative numbers.

the scenarios, achieving comparable results to the best-performing competitor for only a few categories.

Qualitative results. The comparisons of qualitative results between ours and previous SOTA algorithms on various segmentation tasks are shown in Figure 4. It can be observed that our AlignSAM is able to handle diverse challenging scenarios and produce more accurate results.

4.3. Ablation Study

Effect of each individual component. Table 2 presents a detailed ablation result of each component of AlignSAM, showcasing their individual efficacy. For example, for the blur detection task, the removal of reinforcement learning, relying solely on random action selection, resulted in an 8.72% reduction in mIoU. Additionally, excluding the semantic recalibration module caused a 1.58% decrease in mIoU, underscoring its role in generating precise point prompt labels. Notably, there is a significant 21.55% decrease in mIoU when SAM failed to engage in multi-step interaction during inference, emphasizing the pivotal contribution of iterative progress. The multi-step iterative point prompt selection significantly improves mask accuracy compared to using single point prompt. Furthermore, the ablation experiments conducted on Pascal-VOC also validate the effect of AlignSAM as illustrated in Table 3.

Effect of semantic switch. We also examine performance differences between the implicit and explicit branches in the semantic switch in Table 4. Abstract concept segmentation, such as saliency detection, shows superior performance when utilizing image state features only. This is due to the implicit semantic acquired from the text model is unreliable and can be detrimental to visual state features. Conversely, in concrete object segmentation, such as semantic segmentation, the utilization of leveraging hybrid




Figure 5. Several examples to illustrate the iterative point selection and the corresponding segmentation results. The sequence progresses from left to right, showing a gradual increase in the number of point prompts.

vision-language features achieves better performance. The text prompt containing the category name provides valuable guidance for segmentation.

Effect of multi-step iterative point selection. We visualize the selection of iterative point prompts and report their corresponding segmentation results in Figure 5. In simple object segmentation scenarios (row 1), the model demonstrates effectiveness with only a few point prompts and more prompt points can lead to improved accuracy in the boundary of objects. However, in scenarios involving multiple objects, it requires more point prompts to segment the complete foreground areas (row 2). This limitation is also evident in low-level segmentation tasks, such as defocus blur detection (row 3). Feeding the foundation model with an adequate number of precise point prompts can significantly enhance the segmentation precision. Table 5 further substantiates these observations, confirming the benefits of the multi-step iterative point selection process.

5. Conclusion

In this paper, we propose a novel framework for automatic prompting to align SAM to open context via reinforcement learning. This unveiling of the SAM unleashes its potential across diverse downstream tasks while maintaining its adaptability intact. Initially, we devise a reinforcement learning agent to discern informative points for prompting. Subsequently, a semantic recalibration module is introduced to ensure the precise binary classification of the selected prompts. The RL agent and the recalibration module flexibly explore visual and linguistic knowledge to address the implicit and explicit semantics tasks in a unified framework. Extensive experiments conducted on diverse benchmarks confirm the efficacy of the proposed method.

References

- [1] Radhakrishna Achanta, Sheila Hemami, Francisco Estrada, and Sabine Susstrunk. Frequency-tuned salient region detection. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1597–1604. IEEE, 2009. 6
- [2] Arantxa Casanova, Pedro O Pinheiro, Negar Rostamzadeh, and Christopher J Pal. Reinforced active learning for image segmentation. *arXiv preprint arXiv:2002.06583*, 2020. 3, 4
- [3] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 2, 6
- [4] Tianrun Chen, Lanyun Zhu, Chaotao Deng, Runlong Cao, Yan Wang, Shangzhan Zhang, Zejian Li, Lingyun Sun, Ying Zang, and Papa Mao. Sam-adapter: Adapting segment anything in underperformed scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3367–3375, 2023. 1, 2, 3
- [5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 2
- [6] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010. 2, 6
- [7] Deng-Ping Fan, Cheng Gong, Yang Cao, Bo Ren, Ming-Ming Cheng, and Ali Borji. Enhanced-alignment measure for binary foreground map evaluation. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 698–704. International Joint Conferences on Artificial Intelligence Organization, 2018. 6
- [8] Hao He, Xiangtai Li, Guangliang Cheng, Jianping Shi, Yunhai Tong, Gaofeng Meng, Véronique Prinet, and LuBin Weng. Enhanced boundary learning for glass-like object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15859–15868, 2021. 1
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [10] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019. 2, 3, 7
- [11] Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 2, 3, 7
- [12] Xiaowei Hu, Tianyu Wang, Chi-Wing Fu, Yitong Jiang, Qiong Wang, and Pheng-Ann Heng. Revisiting shadow detection: A new benchmark dataset for complex world. *IEEE Transactions on Image Processing*, 30:1925–1934, 2021. 1
- [13] DuoJun Huang, Jichang Li, Weikai Chen, Junshi Huang, Zhenhua Chai, and Guanbin Li. Divide and adapt: Active domain adaptation via customized learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7651–7660, 2023. 1
- [14] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 3
- [15] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision*, pages 709–727. Springer, 2022. 3
- [16] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollar, and Ross Girshick. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4015–4026, 2023. 2, 3
- [17] Debang Li, Huikai Wu, Junge Zhang, and Kaiqi Huang. Fast a3rl: Aesthetics-aware adversarial reinforcement learning for image cropping. *IEEE Transactions on Image Processing*, 28(10):5105–5120, 2019. 4
- [18] Jichang Li, Si Wu, Cheng Liu, Zhiwen Yu, and Hau-San Wong. Semi-supervised deep coupled ensemble learning with classification landmark exploration. *IEEE Transactions on Image Processing*, 29:538–550, 2019. 2
- [19] Jichang Li, Guanbin Li, Yemin Shi, and Yizhou Yu. Cross-domain adaptive clustering for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2505–2514, 2021.
- [20] Jichang Li, Guanbin Li, Feng Liu, and Yizhou Yu. Neighborhood collective estimation for noisy label identification and correction. In *European Conference on Computer Vision*, pages 128–145. Springer, 2022.
- [21] Jichang Li, Guanbin Li, Hui Cheng, Zicheng Liao, and Yizhou Yu. Feddiv: Collaborative noise filtering for federated learning with noisy labels. *arXiv preprint arXiv:2312.12263*, 2023.
- [22] Jichang Li, Guanbin Li, and Yizhou Yu. Adaptive betweenness clustering for semi-supervised domain adaptation. *IEEE Transactions on Image Processing*, 2023.
- [23] Jichang Li, Guanbin Li, and Yizhou Yu. Inter-domain mixup for semi-supervised domain adaptation. *Pattern Recognition*, 2023. 2
- [24] Jiaming Li, Jiacheng Zhang, Jichang Li, Ge Li, Si Liu, Liang Lin, and Guanbin Li. Learning background prompts to discover implicit knowledge for open vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1
- [25] Luoqin Li, Jiabing Wang, Jichang Li, Qianli Ma, and Jia Wei. Relation classification via keyword-attentive sentence mech-

- anism and synthetic stimulation loss. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(9):1392–1404, 2019. 1
- [26] Wenxue Li, Xinyu Xiong, Siying Li, and Fugui Fan. Hybridvps: Hybrid-supervised video polyp segmentation under low-cost labels. *IEEE Signal Processing Letters*, 2023. 2
- [27] Yi Li, Hualiang Wang, Yiqun Duan, and Xiaomeng Li. Clip surgery for better explainability with enhancement in open-vocabulary tasks, 2023. 2, 4, 6
- [28] Yonglin Li, Jing Zhang, Xiao Teng, and Long Lan. Refsam: Efficiently adapting segmenting anything model for referring video object segmentation. *arXiv preprint arXiv:2307.00997*, 2023. 1
- [29] Yang Liu, Muzhi Zhu, Hengtao Li, Hao Chen, Xinlong Wang, and Chunhua Shen. Matcher: Segment anything with one shot using all-purpose feature matching. *arXiv preprint arXiv:2305.13310*, 2023. 1, 2
- [30] Wenhan Luo, Peng Sun, Fangwei Zhong, Wei Liu, Tong Zhang, and Yizhou Wang. End-to-end active object tracking via reinforcement learning. In *International conference on machine learning*, pages 3286–3295. PMLR, 2018. 3
- [31] Jie Ma, Chuan Wang, Yang Liu, Liang Lin, and Guanbin Li. Enhanced soft label for semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1185–1195, 2023. 2
- [32] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3523–3542, 2021. 2
- [33] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. *Advances in neural information processing systems*, 27, 2014. 3
- [34] Youwei Pang, Xiaoqi Zhao, Lihe Zhang, and Huchuan Lu. Multi-scale interactive network for salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9413–9422, 2020. 1
- [35] 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, 3, 6
- [36] Md Alimoor Reza and Jana Kosecka. Reinforcement learning for semantic segmentation in indoor scenes. *arXiv preprint arXiv:1606.01178*, 2016. 3
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [38] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 3
- [39] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 3
- [40] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 5
- [41] Jianping Shi, Li Xu, and Jiaya Jia. Discriminative blur detection features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2965–2972, 2014. 6, 7
- [42] Dongseok Shim and H Jin Kim. Gaussian ram: Lightweight image classification via stochastic retina-inspired glimpse and reinforcement learning. In *2020 20th International Conference on Control, Automation and Systems (ICCAS)*, pages 155–160. IEEE, 2020. 3
- [43] K Simonyan and A Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. 2
- [44] Tomás F Yago Vicente, Minh Hoai, and Dimitris Samaras. Leave-one-out kernel optimization for shadow detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3388–3396, 2015. 6
- [45] Tomás F Yago Vicente, Le Hou, Chen-Ping Yu, Minh Hoai, and Dimitris Samaras. Large-scale training of shadow detectors with noisily-annotated shadow examples. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VI 14*, pages 816–832. Springer, 2016. 6, 7
- [46] Haoxiang Wang, Pavan Kumar Anasosalu Vasu, Fartash Faghri, Raviteja Vemulapalli, Mehrdad Farajtabar, Sachin Mehta, Mohammad Rastegari, Oncel Tuzel, and Hadi Pouransari. Sam-clip: Merging vision foundation models towards semantic and spatial understanding. *arXiv preprint arXiv:2310.15308*, 2023. 1
- [47] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 6, 7
- [48] Xinlong Wang, Wen Wang, Yue Cao, Chunhua Shen, and Tiejun Huang. Images speak in images: A generalist painter for in-context visual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6830–6839, 2023. 2, 3, 7
- [49] Xinlong Wang, Xiaosong Zhang, Yue Cao, Wen Wang, Chunhua Shen, and Tiejun Huang. Seggpt: Towards segmenting everything in context. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1130–1140, 2023. 2, 3
- [50] Yi Wang, Ruili Wang, Xin Fan, Tianzhu Wang, and Xi-angjian He. Pixels, regions, and objects: Multiple enhancement for salient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10031–10040, 2023. 1, 3
- [51] Junde Wu, Rao Fu, Huihui Fang, Yuanpei Liu, Zhaowei Wang, Yanwu Xu, Yueming Jin, and Tal Arbel. Medical sam

- adapter: Adapting segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.12620*, 2023. [7](#)
- [52] Si Wu, Guangchang Deng, Jichang Li, Rui Li, Zhiwen Yu, and Hau-San Wong. Enhancing triplegan for semi-supervised conditional instance synthesis and classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10091–10100, 2019. [1](#)
- [53] Si Wu, Jichang Li, Cheng Liu, Zhiwen Yu, and Hau-San Wong. Mutual learning of complementary networks via residual correction for improving semi-supervised classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6500–6509, 2019. [1](#)
- [54] Xinyu Xiong, Siying Li, and Guanbin Li. Unpaired image-to-image translation based domain adaptation for polyp segmentation. In *International Symposium on Biomedical Imaging*, pages 1–5. IEEE, 2023. [2](#)
- [55] Xinyu Xiong, Churan Wang, Wenxue Li, and Guanbin Li. Mammo-sam: Adapting foundation segment anything model for automatic breast mass segmentation in whole mammograms. In *International Workshop on Machine Learning in Medical Imaging*, pages 176–185. Springer, 2023. [3](#)
- [56] Xin Yang, Haiyang Mei, Ke Xu, Xiaopeng Wei, Baocai Yin, and Rynson WH Lau. Where is my mirror? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8809–8818, 2019. [6, 7](#)
- [57] Letian Yu, Haiyang Mei, Wen Dong, Ziqi Wei, Li Zhu, Yuxin Wang, and Xin Yang. Progressive glass segmentation. *IEEE Transactions on Image Processing*, 31:2920–2933, 2022. [1](#)
- [58] Da Zhang, Hamid Maei, Xin Wang, and Yuan-Fang Wang. Deep reinforcement learning for visual object tracking in videos. *arXiv preprint arXiv:1701.08936*, 2017. [3](#)
- [59] Kaidong Zhang and Dong Liu. Customized segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.13785*, 2023. [1, 7](#)
- [60] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*, 2023. [1, 2, 7](#)
- [61] Zhicheng Zhang, Song Chen, Zichuan Wang, and Jufeng Yang. Planeseg: Building a plug-in for boosting planar region segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 1(1):1–15, 2024. [2](#)
- [62] Zhicheng Zhang, Junyao Hu, Wentao Cheng, Danda Paudel, and Jufeng Yang. Extdm: Distribution extrapolation diffusion model for video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [2](#)
- [63] Zhicheng Zhang, Pancheng Zhao, Eunil Park, and Jufeng Yang. Mart: Masked affective representation learning via masked temporal distribution distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [2](#)
- [64] 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. [1](#)
- [65] Rongjian Zhao, Buyue Qian, Xianli Zhang, Yang Li, Rong Wei, Yang Liu, and Yinggang Pan. Rethinking dice loss for medical image segmentation. In *2020 IEEE International Conference on Data Mining (ICDM)*, pages 851–860. IEEE, 2020. [6](#)
- [66] Quanlong Zheng, Xiaotian Qiao, Ying Cao, and Rynson WH Lau. Distraction-aware shadow detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5167–5176, 2019. [1](#)
- [67] Lei Zhu, Ke Xu, Zhanghan Ke, and Rynson WH Lau. Mitigating intensity bias in shadow detection via feature decomposition and reweighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4702–4711, 2021. [2](#)
- [68] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *arXiv preprint arXiv:2304.06718*, 2023. [2, 3, 7](#)