

Multi-granularity Interaction Simulation for Unsupervised Interactive Segmentation

Kehan Li^{1,3} Yian Zhao^{1,3} Zhennan Wang² Zesen Cheng^{1,3} Peng Jin^{1,3}
Xiangyang Ji⁴ Li Yuan^{1,2,3} Chang Liu^{4*} Jie Chen^{1,2,3*}

¹ School of Electronic and Computer Engineering, Peking University, Shenzhen, China

² Peng Cheng Laboratory, Shenzhen, China

³ AI for Science (AI4S)-Preferred Program, Peking University, Shenzhen, China

⁴ Department of Automation and BNRist, Tsinghua University, Beijing, China

liuchang2022@tsinghua.edu.cn jiechen2019@pku.edu.cn

Abstract

Interactive segmentation enables users to segment as needed by providing cues of objects, which introduces human-computer interaction for many fields, such as image editing and medical image analysis. Typically, massive and expansive pixel-level annotations are spent to train deep models by object-oriented interactions with manually labeled object masks. In this work, we reveal that informative interactions can be made by simulation with semantic-consistent yet diverse region exploration in an unsupervised paradigm. Concretely, we introduce a **Multi-granularity Interaction Simulation (MIS)** approach to open up a promising direction for unsupervised interactive segmentation. Drawing on the high-quality dense features produced by recent self-supervised models, we propose to gradually merge patches or regions with similar features to form more extensive regions and thus, every merged region serves as a semantic-meaningful multi-granularity proposal. By randomly sampling these proposals and simulating possible interactions based on them, we provide meaningful interaction at multiple granularities to teach the model to understand interactions. Our MIS significantly outperforms non-deep learning unsupervised methods and is even comparable with some previous deep-supervised methods without any annotation.

1. Introduction

Interactive segmentation aims at obtaining high-quality object masks with limited user interaction, allowing users to segment objects as needed. During the development of interactive segmentation, various interactive forms like

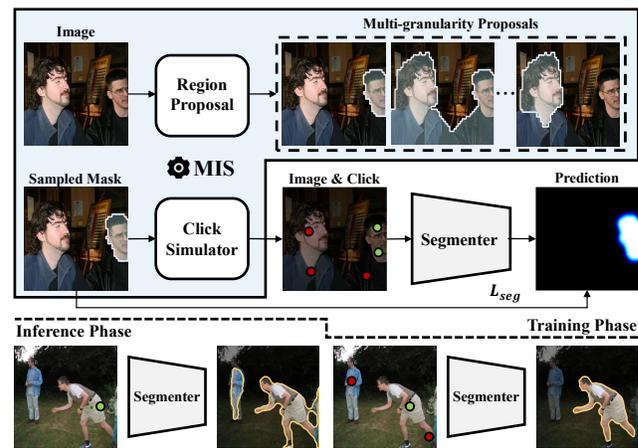


Figure 1. **Top:** To tackle the reliance on object annotations of previous interactive segmentation methods caused by interaction simulation, we propose **Multi-granularity Interaction Simulation (MIS)** that simulates interaction by meaningful regions at multiple granularities. **Down:** The model learns from MIS to successfully understand interactions and produce acceptable segmentation.

bounding boxes [20, 32, 43], scribbles [2, 11, 22], and clicks [44, 17, 25, 35, 8, 36, 24, 9] have been explored. Among them, the click-based interaction becomes the mainstream due to its simplicity and well-established training and evaluation protocols.

State-of-the-art methods for click-based interactive segmentation are based on deep learning, following the basic paradigm proposed by Xu *et al.* [44]. Specifically, these methods encode clicks to a distance map and adapt a semantic segmentation model (e.g., FCN [28]) to take the encoded click map as input, and then train the model with interaction-segmentation pairs. Since it is impractical to collect interaction sequences from real users, previous

*Corresponding author. Project page: <https://lkl.github.io/MIS>

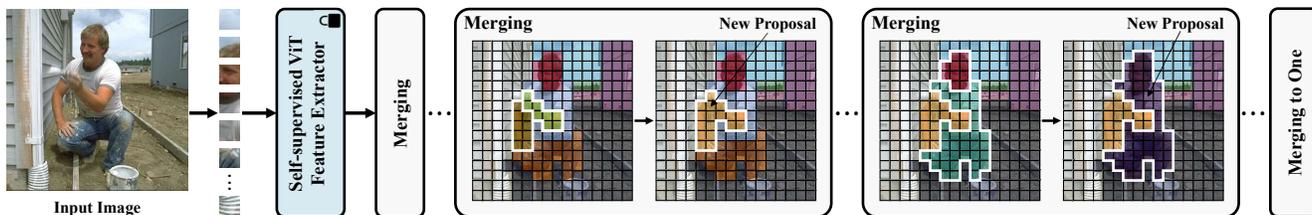


Figure 2. **Illustration of the multi-granularity region proposal generation.** For an input image, We first introduce a self-supervised ViT to produce semantic features for each patch of an image. We then gradually merge two similar patches or regions with similar features to perform a new region. All newly generated regions in the merging process constitute meaningful proposals at multiple granularities.

methods adopt an interaction simulation strategy [44, 36] that randomly samples clicks based on object segmentation annotations. By training with the simulated clicks and corresponding object masks, the model is able to correctly understand user needs and segment objects with a few clicks, significantly outperforming non-deep learning methods. However, the requirement for a large number of expensive pixel-level annotations behind the powerful performance has been ignored by previous methods, which is derived from the object-level interaction simulation.

In this work, we aim at exploring an annotation-free alternation for interaction simulation to eliminate the reliance on object annotations in interactive segmentation, namely unsupervised interactive segmentation. To achieve it, as shown in the top half of Figure 1, our basic idea is to switch previous object-oriented interaction simulation to some semantic-consistent regions, which can be discovered unsupervisedly. In this case, the model requires diverse examples to understand changeable interactions due to the lack of precise information about objects.

Therefore, we propose **Multi-granularity Interaction Simulation (MIS)** to achieve diversity. To ensure the semantic consistency of a region, we first parse the semantic of patches in an image by a Vision Transformer (ViT) [10] pre-trained unsupervisedly with DINO [7]. As shown in Figure 2, we then gradually merge two patches or regions with similar features to form a more extensive region for an image, until only one region remains. In this process, the newly generated region in every step makes up semantic-meaningful multi-granularity proposals. The MIS is finally achieved by randomly sampling these proposals and simulating possible interactions based on them when training the model. Moreover, in order to help the model fight against the unsmoothness on the boundary of proposals caused by the feature extractor, we design a smoothness constraint that considers the low-level feature of pixels. With the meaningful interaction simulated at multiple granularities and the smoothness constraint, the model learns to understand clicks by related regions and produce an acceptable segmentation with several user interactions when inference, while completely discarding object annotations, as shown in the bottom half of Figure 1.

We take the standard protocol [44] to evaluate our method on commonly used datasets including Grab-Cut [32], Berkeley [29], SBD [15], and DAVIS [31]. Benefiting from the rich interaction examples from our MIS, the trained model achieves inspiring performance. Specifically, it significantly outperforms non-deep learning methods under the premise of unsupervised setting, and even surpasses some early deep supervised methods. In addition, it can also perform as an unsupervised pre-training to improve the performance when only limited annotations are available. The surprising results support that the interactive segmentation task can be solved in a more label-efficient way. Moreover, we believe the proposed method can reduce the labeling costs in segmentation tasks by training an interactive segmentation model unsupervisedly and using it to efficiently make task-specific annotations as needed.

2. Related Works

Non-deep Learning Interactive Image Segmentation. Interactive segmentation has been an active field for over two decades due to its wide applications. Before the era of deep learning, the solutions are based on low-level features and optimization. One of the best-known methods is proposed by Boykov and Jolly [6], who define a graph to describe user-provided hard constraints and pixel relationship, then solve interactive segmentation by graph cut with max-flow algorithm [5]. After that, Rother *et al.* [32] augment graph cut by more powerful color data modeling and an iterative strategy. Other improved algorithms include random walk [11], geodesic distance [3], and star-convexity [12]. As early unsupervised interactive segmentation methods, these algorithms stuck when processing complex images due to the dependence on low-level features. Unlike their iterative optimization algorithms for input images, we distill user interactions, which are simulated from offline datasets with our method, into a deep neural network. By leveraging the large capabilities of the neural network, our method achieves significantly better performance than them.

Deep Interactive Image Segmentation. Xu *et al.* [44] are the first to introduce deep learning on this task by extending FCN [28] with click map and click sampling strategies.

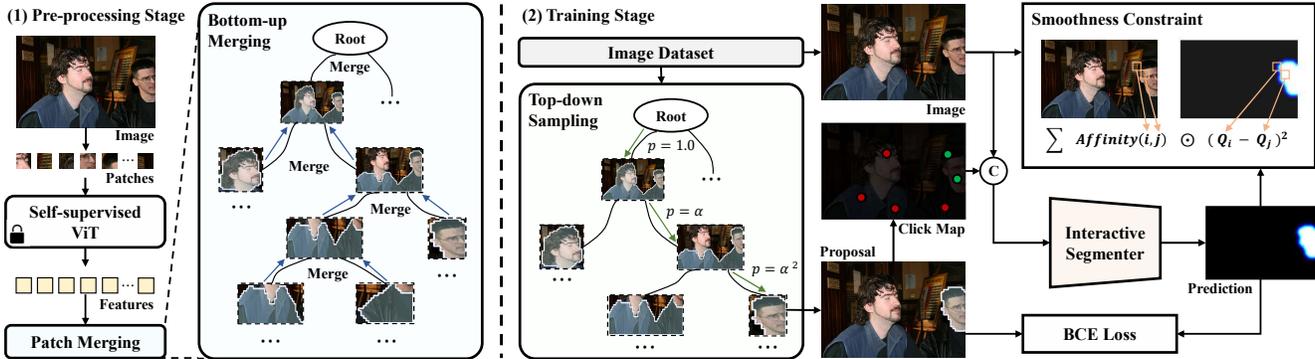


Figure 3. **Illustration of our unsupervised framework to train an interactive segmentation model.** For implementation, the MIS is composed of two stages. In the first pre-processing stage, we first parse the semantic of each patch in the image using a pre-trained ViT and take a gradual merging strategy to produce semantic-consistent region proposals at multiple granularities and save them efficiently with a tree structure. In the training stage, given an image, we randomly select proposals from the corresponding merging tree in a top-down manner and further simulate possible clicks based on them, thus providing the model informative interactions.

Based on this pipeline, the researchers gradually increase the performance ceiling of deep learning methods by emphasizing the first click [25], online optimization [17, 35], diffusing prediction [8], edge-guided flow [14], iterative click sampling [36], powerful feature extractor [26], and Gaussian process posterior [45]. On the other hand, some works [24, 9] focus on efficiency and propose to segment on local crops to accelerate inference. Although the deep learning approaches achieve uplifting performance and efficiency, they require large-scale pixel-level annotations to train, which are expensive and laborious to obtain. Unlike previous works, we focus on the dependency on annotations and show that a fully unsupervised framework can also handle interactive segmentation task well.

Unsupervised Image Segmentation. With the development of self-supervised and unsupervised learning, unsupervised methods for image segmentation task start to emerge. This work is also related to some unsupervised image segmentation methods based on the idea that extracting segments from a pre-trained high-quality dense feature extractor, (e.g., DenseCL [40], DINO [7]). For instance, Simeoni *et al.* [34] proposed a series of hand-made rules to choose pixels belonging to the same object according to their feature similarity. Wang *et al.* [41] introduce normalized cuts [33] on the affinity graph constructed by pixel-level representations from DINO to divide the foreground and background of an image. Hamilton *et al.* [13] train a segmentation head by distilling the feature correspondences from DINO. Melas *et al.* [30] adopt spectral decomposition on the affinity graph to discover meaningful parts in an image. For instance segmentation, FreeSOLO [39] design pseudo instance mask generation based on multi-scale feature correspondences from densely pre-trained models and train an instance segmentation model with these pseudo masks. Although the masks produced by these methods can

serve as pseudo labels to train an interactive segmentation model unsupervisedly, only ordinary results are observed through our experiments in Section 4.1, which demonstrate that these task-specific designs is not apposite for interactive segmentation task. In contrast, we propose a novel strategy to discover semantic-consistent yet diverse regions to fit the interactive segmentation task.

3. Method

In this section, we introduce the overall framework to solve unsupervised interactive segmentation in detail. Macroscopically, we simulate informative interactions from the proposed **Multi-granularity Interaction Simulation (MIS)** to train an interactive segmentation model. For implementation, the MIS is composed of two stages, as shown in Figure 3. In the first pre-processing stage, we first parse the semantic of each patch in the image using a ViT [10] trained self-supervisedly with DINO [7], which has been proven to produce high-quality dense semantic features [34, 13, 30]. With these features, we take a gradual merging strategy to produce semantic-consistent region proposals at multiple granularities and save them efficiently with a tree structure. In the training stage, given an image, we randomly select proposals from the corresponding merging tree in a top-down manner and further simulate possible clicks based on them, thus providing the model informative interactions for learning to understand interactions during the training process.

For training, the model is optimized by the provided interactions, which include meaningful regions and clicks. Specifically, it takes the image and clicks as input and is optimized to produce similar segmentation as the region. Moreover, we design a smoothness constraint when optimizing, for the purpose of mitigating the misleading brought by the inaccurate boundary of the region, which

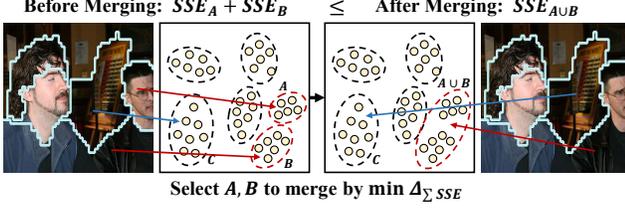


Figure 4. **Illustration of the cost of merging.** The points represent patch features and the circles denote regions that consist of the patches surrounded by it. Merging two regions into one leads to a more divergent cluster and the total SSE increases, thus the cost of merging can be measured as this increment.

is caused by the down-sampling in the ViT. The details of MIS and the training objective are described as follows.

3.1. Multi-granularity Interaction Simulation

Bottom-up Merging. At the core of our MIS is to discover semantic-consistent yet diverse regions, which is achieved efficiently by gradually merging. In the beginning, we produce high-quality semantic features from a ViT [10] trained in a self-supervised manner with DINO [7] to measure the semantic similarity and ensure the semantic-consistency of generated region proposals. In the ViT, an image $\mathbf{I} \in \mathbb{R}^{h \times w \times 3}$ is divided to non-overlapping $n = \frac{h}{s} \times \frac{w}{s}$ patches with stride s , which are then processed by some Transformer [38] encoder layers. The patch features $\mathbf{F} \in \mathbb{R}^{n \times c}$ are obtained from the last layer of the ViT, where c is the feature dimension. With the patch features, we implement the hierarchical merging process by Algorithm 1 and finally get a tree \mathbf{T} which records every merge process.

Initially, n patches form as n leaf nodes of the tree, and a set \mathbb{S} is maintained for saving unmerged nodes, which initially contains all the leaf nodes. In each iteration, the two unmerged nodes with the least cost according to their features will be merged to form a new node, until only one unmerged node remains. Thus, there should be totally $n - 1$ iterations, and the index of the newly generated node in k -th iteration is $n + k$. Moreover, considering the prior that parts belonging to the same object are usually adjacent, we add connectivity constraints when merging to greatly accelerate the search space for finding the minimum cost. Specifically, each patch is only connected to its four adjacent patches and two regions are connected only if at least a pair of patches from them respectively are connected.

We follow Ward’s method [42] to measure the cost of merging, which minimizes the increment in the total within-cluster sum of squared error (SSE) after merging, as shown in Figure 4. Regarding a region A as a cluster that contains some patches, the SSE is defined as

$$e_A = \sum_{i \in A} \|\mathbf{F}_i - \boldsymbol{\mu}_A\|^2, \quad (1)$$

Algorithm 1: Bottom-up Merging

Input : Patch features $\mathbf{F} \in \mathbb{R}^{n \times c}$
Output: Merge tree $\mathbf{T} \in \mathbb{Z}^{(n-1) \times 2}$

- 1 *Initialization*: $\mathbb{S} \leftarrow \{1, 2, \dots, n\}$
- 2 **for** $k \leftarrow 1$ **to** $n - 1$ **do**
- 3 $i, j \leftarrow \arg \min_{i, j \in \mathbb{S}} \text{Cost}(i, j, \mathbf{F})$ s.t. $\text{Connect}(i, j)$;
- 4 $\mathbf{T}_{k, :} \leftarrow (i, j)$;
- 5 $\mathbb{S} \leftarrow \mathbb{C}_{\mathbb{S}}\{i, j\} \cup \{n + k\}$;
- 6 **end**

Algorithm 2: Top-down Sampling

Input : Merge tree \mathbf{T} , decay coefficient α
Output: Node index k

- 1 *Initialization*: $p \leftarrow 1.0, k \leftarrow n - 1$
- 2 **while** $\text{Rand}() < p$ **and** $k > 0$ **do**
- 3 **if** $\text{Rand}() < 0.5$ **then**
- 4 $k \leftarrow \mathbf{T}_{k, 0} - n$;
- 5 **else**
- 6 $k \leftarrow \mathbf{T}_{k, 1} - n$;
- 7 **end**
- 8 $p \leftarrow \alpha \cdot p$;
- 9 **end**
- 10 $k \leftarrow k + n$;

where $\boldsymbol{\mu}_A \in \mathbb{R}^c$ is the center, which can be computed as

$$\boldsymbol{\mu}_A = \frac{1}{s_A} \sum_{i \in A} \mathbf{F}_i, \quad (2)$$

where s_A is the number of patches within A . When merging two regions, the new center is

$$\boldsymbol{\mu}_{A \cup B} = \frac{s_A \boldsymbol{\mu}_A + s_B \boldsymbol{\mu}_B}{s_A + s_B}. \quad (3)$$

Then the cost of merging is defined as the difference of SSE before and after merging. Since the SSE of all clusters except the clusters being merged are unchanged, it can be formulated as

$$\begin{aligned} \text{Cost}(A, B) &= e_{A \cup B} - (e_A + e_B) \\ &= \frac{s_A \cdot s_B}{s_A + s_B} \|\boldsymbol{\mu}_A - \boldsymbol{\mu}_B\|^2. \end{aligned} \quad (4)$$

Detailed derivation is shown in the appendix.

During the merging process, the center $\boldsymbol{\mu}$ and the size s of each region are recorded. Initially, the size is set to 1 for each patch and the center is its feature. The detailed calculation after each merging is as follows: **(1)** Computing the center of the newly generated region according to Equation (3). **(2)** Computing the size of the newly generated region as $s_{A \cup B} = s_A + s_B$. **(3)** Computing the cost of



Figure 5. **Qualitative results of discovered regions for the baselines and our MIS.** While the baselines only focus on some salient regions in a fixed granularity, our MIS produces semantic-consistent regions at multiple granularities thus catching more objects and parts.

merging newly generated regions with other regions according to Equation (4). Since the other regions are not changed, the cost between them remains the same.

After $n-1$ iterations, we get a tree structure that is stored by a matrix $T \in \mathbb{Z}^{(n-1) \times 2}$, where the i -th row is the indices of the two children of the $i+n$ node. In the tree, each node represents a merge and also represents a region proposal which consists of all patches within the subtree rooted at this node. Hence, with this hierarchical relationship, we obtain $n-1$ region proposals at multiple granularities for each image, which help the interactive segmentation model to learn various interactions. For storage, there is no need to store all the $n-1$ masks but only the tree structure, that is, the matrix T , which makes it more space-efficient.

Top-down Sampling. When training, the MIS is worked based on the stored merging tree. Specifically, when sampling an image for training, we select a mask proposal from its corresponding merging tree for simulating an interaction. In the merging tree, shallow nodes correspond to relatively complete parts or their combinations and deeper nodes are more relevant to finely divided parts. Intuitively, we prefer the complete one while maintaining diversity. Therefore, we design a sampling strategy in a top-down manner and gradually decrease the probability to go deeper, as shown in Algorithm 2. Concretely, we start at the root node (*i.e.*, the node with an index of $2n-1$) and the probability $p=1.0$. Then in each iteration, we choose to go deeper or stop with probability p , and decrease p with a coefficient α . The loop will terminate when choosing to stop or reach a leaf node (*i.e.*, nodes whose index is less than or equal to n). If it is decided to go deeper, we randomly move the current node to one of its children. Finally, we get the index of a node k corresponding to a region proposal, where the patches in the sub-tree rooted at k are annotated as foreground.

Click Sampling. With the sampled region, the last step is to simulate possible clicks according to the region. Drawing on the experience of previous methods, we use a combination of random and iterative click simulation strategies to produce a click map following RITM [36]. Specifically, some positive and negative clicks are first generated in parallel according to the foreground and background defined by the region. Then some additional clicks are appended by checking the erroneous region of the prediction which is produced by previous clicks. Finally, the MIS generates complete interactions which consist of a segmentation target and corresponding clicks.

3.2. Training

Interactive Segmentation Model. Typical interactive segmentation models are built based on semantic segmentation models (*e.g.*, FCN [28]) by taking an additional click map as input, which is transformed from the coordinate of clicks. Formulaically, given an image $I \in \mathbb{R}^{h \times w \times 3}$, a click map $D \in \mathbb{R}^{h \times w \times 2}$ which presents positive and negative clicks respectively, and previous prediction $B \in \{0, 1\}^{h \times w}$, the model $\mathcal{F}(\cdot)$ outputs the probability of being foreground for each pixel

$$Q = \mathcal{F}(\text{Fusion}(I, D, B)), Q \in [0, 1]^{h \times w}, \quad (5)$$

where $\text{Fusion}(\cdot)$ is some kind of fusion operation such as concatenate. Through training, the model learns to map the image and the user clicks to foreground and background segmentation thus can be used interactively when inference.

Training with Proposals. Under our unsupervised setting, the model learns this mapping by the simulated interactions from the MIS. The model is first learned by the supervision from the multi-granularity region proposals M by optimiz-

Method	Backbone	GrabCut		Berkeley		SBD		DAVIS	
		NoC@85 ↓	NoC@90 ↓	NoC@85 ↓	NoC@90 ↓	NoC@85 ↓	NoC@90 ↓	NoC@85 ↓	NoC@90 ↓
<i>Deep Supervised Methods</i>									
DIOS w/o GC [44] _{CVPR16}	FCN	<u>8.02</u>	<u>12.59</u>	-	-	<u>14.30</u>	<u>16.79</u>	<u>12.52</u>	<u>17.11</u>
DIOS w/ GC [44] _{CVPR16}	FCN	<u>5.08</u>	<u>6.08</u>	-	-	<u>9.22</u>	<u>12.80</u>	<u>9.03</u>	<u>12.58</u>
LD [23] _{CVPR18}	VGG-19	<u>3.20</u>	<u>4.79</u>	-	-	<u>7.41</u>	<u>10.78</u>	5.05	<u>9.57</u>
BRS [17] _{CVPR19}	DenseNet	<u>2.60</u>	<u>3.60</u>	-	<u>5.08</u>	6.59	<u>9.78</u>	5.58	8.24
f-BRS [35] _{CVPR20}	ResNet-101	<u>2.30</u>	<u>2.72</u>	-	4.57	4.81	7.73	5.04	7.41
FCA-Net [25] _{CVPR20}	ResNet-101	-	2.08	-	3.92	-	-	-	7.57
IA+SA [19] _{ECCV20}	ConvNet	-	<u>3.07</u>	-	<u>4.94</u>	-	-	5.16	-
CDNet [8] _{ICCV21}	ResNet-34	1.86	2.18	1.95	3.27	5.18	7.89	5.00	6.89
RITM [36] _{ICIP22}	HRNet-18	1.76	2.04	1.87	3.22	3.39	5.43	4.94	6.71
FocalClick [9] _{CVPR22}	MiT-B0	1.66	1.90	-	3.14	4.34	6.51	5.02	7.06
FocusCut [24] _{CVPR22}	ResNet-101	1.46	1.64	1.81	3.01	3.40	5.31	4.85	6.22
PseudoClick [27] _{ECCV22}	HRNet-32	-	1.84	-	2.98	-	5.61	4.74	6.16
SimpleClick [26] _{arXiv22}	ViT-B	1.40	1.54	1.44	2.46	3.28	5.24	4.10	5.48
<i>Non-deep Learning Methods</i>									
GraphCut [6] _{ICCV01}	N/A	7.98	10.00	-	14.22	13.60	15.96	15.13	17.41
Random Walk [11] _{TPAMI06}	N/A	11.36	13.77	-	14.02	12.22	15.04	16.71	18.31
Geodesic Matting [3] _{ICV09}	N/A	13.32	14.57	-	15.96	15.36	17.60	18.59	19.50
GSC [12] _{CVPR10}	N/A	7.10	9.12	-	12.57	12.69	15.31	15.35	17.52
ESC [12] _{CVPR10}	N/A	7.24	9.20	-	12.11	12.21	14.86	15.41	17.70
<i>Deep Unsupervised Methods</i>									
TokenCut* [41] _{CVPR22}	ViT-B	3.04	5.74	6.30	9.97	10.82	13.16	10.56	15.01
FreeMask* [39] _{CVPR22}	ViT-B	4.06	6.10	5.55	9.02	8.61	12.01	8.26	13.05
DSM* [30] _{CVPR22}	ViT-B	3.64	4.64	5.49	7.75	8.59	11.57	7.08	10.11
MIS w/o \mathcal{L}_{smooth} (Ours)	ViT-B	2.08	2.56	3.40	5.51	7.38	9.92	6.78	10.08
MIS (Ours)	ViT-B	1.94	2.32	3.09	4.58	6.91	9.51	6.33	8.44

Table 1. Comparison with previous methods and baselines on four interactive segmentation benchmarks. \star denotes baselines that we introduce from other unsupervised tasks. From top to bottom we show the previous deep supervised method, unsupervised non-deep learning method, and our newly proposed deep unsupervised method. All models listed here are trained on SBD dataset. Our MIS achieves superior performance within unsupervised methods and baselines and even exceeds early deep supervised methods (denoted by underline).

ing the objective

$$\mathcal{L}_{pseudo} = \text{BCE}(\mathbf{M}, \mathbf{Q}), \quad (6)$$

where $\text{BCE}(\cdot)$ denotes the binary cross entropy loss. Although the region proposals generated by the MIS contains semantics, the possible inaccuracy of the self-supervised feature extractor and the patching process in the ViT will bring some inaccurate part. Fortunately, as discovered by previous research [1] that the model tends to learn effect patterns rather than memorize the noise, we find that the model can learn to correctly map clicks to segmentation after training with these simulated interactions.

Smoothness Constraint. In order to help the model correct the error caused by patching, we propose a smoothness constraint. We first introduce the bilateral affinity matrix [37, 4] \mathbf{W} which presents the affinity of pixels based on their low-level feature and the item $\mathbf{W}_{i,j}$ is

$$\exp\left(-\frac{\|\mathbf{a}_i^{xy} - \mathbf{a}_j^{xy}\|^2}{2\sigma_{xy}^2} - \frac{(a_i^l - a_j^l)^2}{2\sigma_l^2} - \frac{\|\mathbf{a}_i^{uv} - \mathbf{a}_j^{uv}\|^2}{2\sigma_{uv}^2}\right),$$

where \mathbf{a}_i is a pixel with a spatial position (p_i^x, p_i^y) and color (p_i^l, p_i^u, p_i^v) and $\sigma_{xy}, \sigma_l, \sigma_{uv}$ control the extent of the spatial, luma, and chroma support of the filter. Motivated by the bilateral solver [4], the smooth constraint is implemented by restricting the model to produce consistent predictions on pixels with high bilateral affinity, which is formulated as

$$\mathcal{L}_{smooth} = \sum_i \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} \mathbf{W}_{ij} \cdot (\mathbf{Q}_i - \mathbf{Q}_j)^2. \quad (7)$$

As we mainly concern with the predictions on the boundary, we only perform the loss locally for fast computation and memory efficiency, *i.e.*, for pixel i the loss is calculated in its 5×5 neighborhoods \mathcal{N}_i . Finally, the overall training objective is

$$\mathcal{L} = \mathcal{L}_{pseudo} + \lambda \cdot \mathcal{L}_{smooth}, \quad (8)$$

where λ controls the strength of the smoothness constraint.

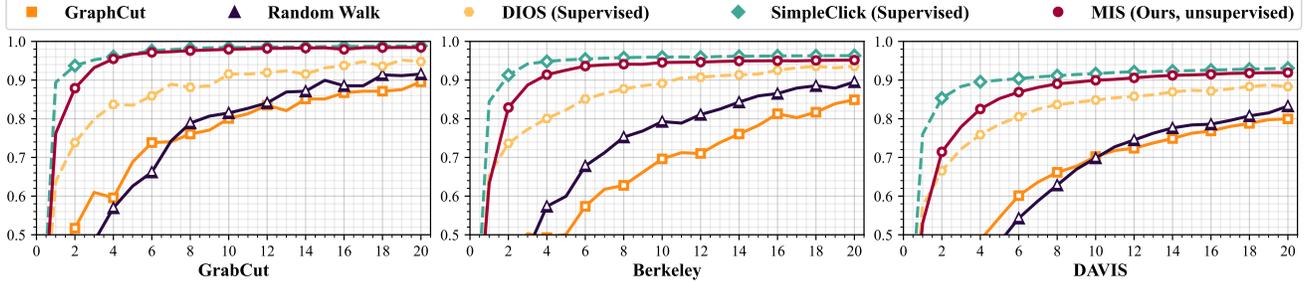


Figure 6. **Comparison of IoU-NoC curve.** The two dotted lines indicate the performance of supervised approaches. The proposed method achieves surprising performance that significantly exceeds traditional methods and is close to supervised method when the NoC increases.

Data Ratio	Initialization	GrabCut		Berkeley		SBD		DAVIS	
		NoC@85 ↓	NoC@90 ↓						
1%	Default (MAE [16])	2.32	2.52	2.36	4.44	8.01	10.43	5.81	8.28
	MIS (<i>Ours</i>)	1.54 (↓ 0.78)	1.90 (↓ 0.62)	1.95 (↓ 0.41)	3.73 (↓ 0.71)	5.57 (↓ 2.44)	8.15 (↓ 2.28)	5.36 (↓ 0.45)	7.40 (↓ 0.88)
5%	Default (MAE [16])	1.68	2.22	1.97	3.64	5.48	8.13	5.14	7.23
	(<i>Ours</i>)	1.62 (↓ 0.06)	1.88 (↓ 0.34)	1.95 (↓ 0.02)	3.32 (↓ 0.32)	5.21 (↓ 0.27)	7.49 (↓ 0.64)	4.54 (↓ 0.60)	6.11 (↓ 1.12)
10%	Default (MAE [16])	1.64	1.84	1.88	3.21	5.33	7.89	4.88	6.77
	MIS (<i>Ours</i>)	1.40 (↓ 0.24)	1.50 (↓ 0.34)	1.64 (↓ 0.24)	2.73 (↓ 0.48)	4.61 (↓ 0.72)	6.95 (↓ 0.94)	4.25 (↓ 0.63)	5.94 (↓ 0.83)

Table 2. **Performance with limited annotations.** When training the model in a supervised manner with a limited number of annotations (1%, 5%, 10%), the performance can be significantly improved by unsupervised pre-training using the proposed method.

4. Experiments

Implementation Details. We follow SimpleClick [26] to build the interactive segmentation model, which consists of two patch embedding modules for image and click map respectively, a ViT [10] backbone initialized with MAE [16], a simple feature pyramid [21], and an MLP segmentation head. For training, we optimize the model for 55 epochs using the Adam [18] optimizer with a learning rate of $5e-5$, which decays by a factor of 10 at 50 epochs. The decay coefficient α in Algorithm 2 and the weight λ in Equation (8) are 0.9 and 10.0, respectively. In addition, we set a threshold of 0.05 to filter small proposals. That is to say, we ignore the proposals whose areas are less than 0.05 times the image area. For inference, we set the threshold of binarizing the prediction to 0.5 as normal binary segmentation and use the same pre-processing as [26]. To ensure the fully unsupervised setting, we employ a self-supervisedly pre-trained model as the feature extractor. Specifically, a Vision Transformer (ViT) [10] trained with DINO [7] is adopted since the property of representing patch-level semantics [7, 13, 34]. We use ViT-Small [34] with a patch size of 8. The results with other models are included in the appendix. Before feeding an image into the ViT, we resize the image to make its height and width divisible by the patch size (*i.e.*, the target height is $(h + p - h \bmod p)$ when the original height h is not divisible by the patch size p), and the position embedding is interpolated to fit the image size using bilinear interpolation. We finally use all output tokens of the last block except the *cls* token as the patch features.

Evaluation. The evaluation is done under the standard evaluation protocol of click-based interactive segmentation. Specifically, we adopt the same click simulator as previous work [44, 17, 25, 35, 8, 36, 24, 9] to sample clicks when evaluation. Roughly, the next click will be placed at the center of the largest error region by comparing the ground-truth and prediction. We adopt the Number of Click (NoC) as the evaluation metric, which counts the average number of clicks needed to achieve a fixed Intersection over Union (IoU), the smaller the better. We use two commonly used target intersection-over-union (IoU) thresholds 85% and 90%, denoted as NoC@85 and NoC@90 respectively. Additionally, the IoU-NoC curves are also adopted to represent the convergence trend of IoU when the NoC increases, as well as compare the IoU metric under the same NoC.

Dataset. We evaluate the performance of our MIS on the following four benchmarks. (1) **GrabCut** [32]: The dataset contains 50 relatively easy images whose background and foreground have a clear difference. (2) **Berkeley** [29]: The dataset contains 96 images with 100 instances and some of them are more challenging than GrabCut. (3) **SBD** [15]: The dataset contains 2,857 images with 6,671 challenging instances, which does not overlap with our training set. (4) **DAVIS** [31]: We follow previous works [24, 9, 27, 36, 26] and use the same 345 frames for evaluation.

4.1. Quantitative Results

Baselines. Since the fully unsupervised interactive segmentation has not been explored by previous research, we in-



Figure 7. **Visualization of the interaction and segmentation.** The green and red points indicate positive and negative clicks respectively.

Value	GrabCut		Berkeley		SBD		DAVIS	
	NoC@85 ↓	NoC@90 ↓						
<i>Strength of the Smoothness Constraint λ</i>								
0.0	2.08	2.56	3.40	5.51	7.38	9.92	6.78	10.08
3.0	2.02 (↓ 0.06)	2.50 (↓ 0.06)	3.07 (↓ 0.33)	5.04 (↓ 0.47)	6.84 (↓ 0.54)	9.33 (↓ 0.59)	6.32 (↓ 0.46)	8.78 (↓ 1.30)
5.0	1.94 (↓ 0.14)	2.34 (↓ 0.22)	3.01 (↓ 0.39)	4.71 (↓ 0.80)	6.79 (↓ 0.59)	9.39 (↓ 0.53)	6.39 (↓ 0.39)	8.55 (↓ 1.53)
10.0	1.94 (↓ 0.14)	2.32 (↓ 0.24)	3.09 (↓ 0.31)	4.58 (↓ 0.93)	6.91 (↓ 0.47)	9.51 (↓ 0.41)	6.33 (↓ 0.45)	8.44 (↓ 1.64)
15.0	2.12 (↑ 0.04)	2.50 (↓ 0.06)	3.25 (↓ 0.15)	5.13 (↓ 0.38)	7.12 (↓ 0.26)	9.85 (↓ 0.07)	6.71 (↓ 0.07)	9.10 (↓ 0.98)

Table 3. **Ablation on the smoothness constraint.** The results demonstrate its effectiveness and robustness.

introduce some baselines which are originally designed for unsupervised salient detection (TokenCut [41]), instance segmentation (FreeMask [39]), and semantic segmentation (DSM [30]) to demonstrate the effectiveness of our method. We follow the official implementation of these baselines to extract mask proposals and use them to train the same model as ours. We denote these baselines by “*” in Table 1.

Comparison with Previous Work. For comparison, we present experimental results of previous supervised deep learning methods, non-deep learning methods, baselines, and the proposed MIS on four aforementioned datasets. Under the unsupervised setting, our method significantly outperforms all other methods as shown in the second part of Table 1. It is also worth mentioning that the proposed MIS achieves inspiring performance that is even comparable with some previous supervised approaches. This supports that the interactive segmentation task can be solved in a more label-efficient way. Moreover, we show the IoU-NoC curve of previous interactive segmentation methods and MIS in Figure 6, where the dotted lines indicate the upper and lower performance of previous supervised methods. It can be found that our method reaches obviously higher

IoU compared to the non-deep learning methods, especially in the first few clicks, and gradually converge to the supervised counterpart when NoC increases.

Comparison with Baselines. MIS also gets advanced results compared to the baselines which only focus on the specific granularity such as the most salient object or fixed number of objects, demonstrating the effectiveness of diverse interactions under the unsupervised setting.

Supervised Fine-tuning. Besides the fully unsupervised setting, we further fine-tune our model in a supervised manner with only a limited number of annotations (1%, 5%, 10%). As shown in Table 2, the model initialized with MIS pre-training the MAE [16] counterpart especially on the relatively hard dataset (*i.e.*, SBD and DAVIS). These results demonstrate that MIS can serve as a strong unsupervised pre-training method for interactive segmentation when the available annotations are rare.

4.2. Qualitative Results

Multi-granularity Region Proposal. In Figure 5 we show the qualitative results of the region proposals generated by the baselines and our MIS, where different colors mean dif-

Value	GrabCut		Berkeley		SBD		DAVIS	
	NoC@85 ↓	NoC@90 ↓						
<i>Decay coefficient α</i>								
0.0	2.12	2.48	3.89	5.52	8.71	10.87	7.84	9.92
0.8	1.94 (↓ 0.18)	2.40 (↓ 0.08)	3.08 (↓ 0.81)	4.99 (↓ 0.53)	7.03 (↓ 1.68)	9.62 (↓ 1.25)	6.36 (↓ 1.48)	8.57 (↓ 1.35)
0.9	1.94 (↓ 0.18)	2.32 (↓ 0.16)	3.09 (↓ 0.80)	4.58 (↓ 0.94)	6.91 (↓ 1.80)	9.51 (↓ 1.36)	6.33 (↓ 1.51)	8.44 (↓ 1.48)
0.95	2.22 (↑ 0.10)	2.62 (↑ 0.14)	3.19 (↓ 0.70)	4.89 (↓ 0.63)	6.91 (↓ 1.80)	9.52 (↓ 1.35)	6.09 (↓ 1.75)	8.54 (↓ 1.38)

Table 4. **Ablation on top-down sampling.** In the case of $\alpha = 0$, only the two children of the root node are sampled and thus multi-granularity is not allowed. The results show the effectiveness of multi-granularity and the robustness of α .

Sampling	GrabCut		Berkeley		SBD		DAVIS	
	NoC@85 ↓	NoC@90 ↓						
Uniform	3.84	4.32	4.42	6.20	7.11	9.95	7.71	10.10
Top-down	1.94 (↓ 1.90)	2.32 (↓ 2.00)	3.09 (↓ 1.33)	4.58 (↓ 1.62)	6.91 (↓ 0.20)	9.51 (↓ 0.44)	6.33 (↓ 1.38)	8.44 (↓ 1.66)

Table 5. **Ablation on sampling strategy.** The proposed top-down sampling strategy is effective because it reasonably assigns different weights to fragmented and complete components, so as to sample complete objects as much as possible while maintaining diversity.

ferent mask proposals and our mask proposals. Compared with them, our multi-granularity proposals capture more informative regions and are more diverse. Through the top-down random sampling strategy, our proposals provide the model with rich interactions, which makes the model learn a better mapping from clicks to segmentation under the unsupervised setting.

Interactive Segmentation. Figure 7 show some interactive segmentation examples of our method. The trained interactive segmentation model is able to capture the outline of the needed object by the first one or two clicks and achieves high IoU with a few clicks, which demonstrates the feasibility of our newly proposed unsupervised setting and the effectiveness of the MIS. For some simple objects, our model can produce satisfactory predictions within the first few clicks. It can also handle complex scenes such as partially occluded objects and adjacent instances.

4.3. Ablation Study

Smoothness Constraint. For the smoothness constraint, we switch the weight λ in Equation (8) to show its effectiveness. The experimental results in Table 3 demonstrate that the smoothness constraint is helpful for improving the performance through the principle that helps the model correct the boundary errors in the proposals. The effectiveness is robust when λ is within a reason range.

Top-down Sampling. The top-down sampling (Algorithm 2) is the core part of our MIS to achieve diversity. We evaluate its effectiveness by comparing with a fixed-granularity strategy, and a fully random strategy that randomly sample nodes on the merging tree from a uniform distribution. In Table 4, we take $\alpha = 0.0$ as the baseline. In this case, we always use one of the two children of the root node in the merging tree, thus multi-granularity is not allowed. Comparing the results with $\alpha = 0.8, 0.9, 0.95$ we

can find that the introduction of multi-granularity brings noticeable improvement, especially in the hard dataset. It is also obvious that this hyper-parameter is relatively robust. In addition, the results compared to the fully random sampling strategy in Table 5 also show clear advantages of our top-down sampling strategy. When sampling uniformly, the finely divided object components are more likely to be selected due to the large number, resulting in local perception and thus requiring more clicks to localize objects. In contrast, our top-down sampling strategy reasonably assigns different weights to fragmented and complete components, so as to sample complete objects as much as possible while maintaining diversity.

5. Conclusion

In this work, we open up a promising direction for deep interactive segmentation that completely discards the dependence on manual annotations. While previous methods learn the interaction through object-oriented simulation with pixel-level annotations, we propose MIS and demonstrate the feasibility to simulate informative interaction with semantic-consistent regions discovered in an unsupervised manner. Inspiring experimental results show the effectiveness of our method and the potential to reduce the labor of creating annotation without cost via training an interactive segmentation model unsupervisedly. Our experiment also supports that the interactive segmentation task can be solved in a more label-efficient way to trade off the performance and annotations. We hope our exploration can promote the attendance of label efficiency on interactive segmentation.

Acknowledgements. This work was supported by the National Key R&D Program of China (2022ZD0118101) and the Natural Science Foundation of China (No. 61972217, 32071459, 62176249, 62006133, 62271465, 62202014). Li Yuan is sponsored by CCF Tencent Open Research Fund.

References

- [1] Devansh Arpit, Stanislaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, pages 233–242. PMLR, 2017. 6
- [2] Junjie Bai and Xiaodong Wu. Error-tolerant scribbles based interactive image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 392–399, 2014. 1
- [3] Xue Bai and Guillermo Sapiro. A geodesic framework for fast interactive image and video segmentation and matting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1–8, 2007. 2, 6
- [4] Jonathan T Barron and Ben Poole. The fast bilateral solver. In *European Conference on Computer Vision*, pages 617–632. Springer, 2016. 6
- [5] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(9):1124–1137, 2004. 2
- [6] Yuri Y Boykov and M-P Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in ND images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 105–112, 2001. 2, 6
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. 2, 3, 4, 7
- [8] Xi Chen, Zhiyan Zhao, Feiwu Yu, Yilei Zhang, and Manni Duan. Conditional diffusion for interactive segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7345–7354, 2021. 1, 3, 6, 7
- [9] Xi Chen, Zhiyan Zhao, Yilei Zhang, Manni Duan, Donglian Qi, and Hengshuang Zhao. FocalClick: Towards practical interactive image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1300–1309, 2022. 1, 3, 6, 7
- [10] 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. In *International Conference on Learning Representations*, 2020. 2, 3, 4, 7
- [11] Leo Grady. Random walks for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1768–1783, 2006. 1, 2, 6
- [12] Varun Gulshan, Carsten Rother, Antonio Criminisi, Andrew Blake, and Andrew Zisserman. Geodesic star convexity for interactive image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3129–3136, 2010. 2, 6
- [13] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T Freeman. Unsupervised semantic segmentation by distilling feature correspondences. In *International Conference on Learning Representations*, 2022. 3, 7
- [14] Yuying Hao, Yi Liu, Zewu Wu, Lin Han, Yizhou Chen, Guowei Chen, Lutao Chu, Shiyu Tang, Zhiliang Yu, Zeyu Chen, et al. Edgeflow: Achieving practical interactive segmentation with edge-guided flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1551–1560, 2021. 3
- [15] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhansu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 991–998. IEEE, 2011. 2, 7
- [16] 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. 7, 8
- [17] Won-Dong Jang and Chang-Su Kim. Interactive image segmentation via backpropagating refinement scheme. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5297–5306, 2019. 1, 3, 6, 7
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [19] Theodora Kontogianni, Michael Gygli, Jasper Uijlings, and Vittorio Ferrari. Continuous adaptation for interactive object segmentation by learning from corrections. In *European Conference on Computer Vision*, pages 579–596. Springer, 2020. 6
- [20] Victor Lempitsky, Pushmeet Kohli, Carsten Rother, and Toby Sharp. Image segmentation with a bounding box prior. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 277–284, 2009. 1
- [21] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *European Conference on Computer Vision*, pages 280–296. Springer, 2022. 7
- [22] Yin Li, Jian Sun, Chi-Keung Tang, and Heung-Yeung Shum. Lazy snapping. *ACM Transactions on Graphics*, 23(3):303–308, 2004. 1
- [23] Zhuwen Li, Qifeng Chen, and Vladlen Koltun. Interactive image segmentation with latent diversity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 577–585, 2018. 6
- [24] Zheng Lin, Zheng-Peng Duan, Zhao Zhang, Chun-Le Guo, and Ming-Ming Cheng. FocusCut: Diving into a focus view in interactive segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2637–2646, 2022. 1, 3, 6, 7
- [25] Zheng Lin, Zhao Zhang, Lin-Zhuo Chen, Ming-Ming Cheng, and Shao-Ping Lu. Interactive image segmentation with first click attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13339–13348, 2020. 1, 3, 6, 7

- [26] Qin Liu, Zhenlin Xu, Gedas Bertasius, and Marc Niethammer. Simpleclick: Interactive image segmentation with simple vision transformers. *arXiv preprint arXiv:2210.11006*, 2022. 3, 6, 7
- [27] Qin Liu, Meng Zheng, Benjamin Planche, Srikrishna Karanam, Terrence Chen, Marc Niethammer, and Ziyang Wu. Pseudoclick: Interactive image segmentation with click imitation. In *European Conference on Computer Vision*, pages 728–745. Springer, 2022. 6, 7
- [28] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. 1, 2, 5
- [29] Kevin McGuinness and Noel E O’connor. A comparative evaluation of interactive segmentation algorithms. *Pattern Recognition*, 43(2):434–444, 2010. 2, 7
- [30] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Deep spectral methods: A surprisingly strong baseline for unsupervised semantic segmentation and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8364–8375, 2022. 3, 6, 8
- [31] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 724–732, 2016. 2, 7
- [32] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. “grabcut” interactive foreground extraction using iterated graph cuts. *ACM Transactions on Graphics*, 23(3):309–314, 2004. 1, 2, 7
- [33] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000. 3
- [34] Oriane Siméoni, Gilles Puy, Huy V Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. In *British Machine Vision Conference*, 2021. 3, 7
- [35] Konstantin Sofiiuk, Iliia Petrov, Olga Barinova, and Anton Konushin. f-brs: Rethinking backpropagating refinement for interactive segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8623–8632, 2020. 1, 3, 6, 7
- [36] Konstantin Sofiiuk, Ilya A Petrov, and Anton Konushin. Reviving iterative training with mask guidance for interactive segmentation. In *IEEE International Conference on Image Processing*, pages 3141–3145, 2022. 1, 2, 3, 5, 6, 7
- [37] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 839–846, 1998. 6
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017. 4
- [39] Xinlong Wang, Zhiding Yu, Shalini De Mello, Jan Kautz, Anima Anandkumar, Chunhua Shen, and Jose M Alvarez. Freesolo: Learning to segment objects without annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14176–14186, 2022. 3, 6, 8
- [40] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3024–3033, 2021. 3
- [41] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14543–14553, 2022. 3, 6, 8
- [42] Joe H Ward Jr. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301):236–244, 1963. 4
- [43] Jiajun Wu, Yibiao Zhao, Jun-Yan Zhu, Siwei Luo, and Zhuowen Tu. Milcut: A sweeping line multiple instance learning paradigm for interactive image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 256–263, 2014. 1
- [44] Ning Xu, Brian Price, Scott Cohen, Jimei Yang, and Thomas S Huang. Deep interactive object selection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 373–381, 2016. 1, 2, 6, 7
- [45] Minghao Zhou, Hong Wang, Qian Zhao, Yuexiang Li, Yawen Huang, Deyu Meng, and Yefeng Zheng. Interactive segmentation as gaussian process classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5297–5306, 2019. 3