

Zero-Shot Edge Detection with SCESAME: Spectral Clustering-based Ensemble for Segment Anything Model Estimation

Hiroaki Yamagiwa^{1,2} Yusuke Takase¹ Hiroyuki Kambe² Ryosuke Nakamoto^{1,2}
¹Kyoto University ²Rist Inc.

hiroaki.yamagiwa@sys.i.kyoto-u.ac.jp, y.takase@sys.i.kyoto-u.ac.jp,
 {hiroyuki.kambe, ryosuke.nakamoto}@rist.co.jp

Abstract

This paper proposes a novel zero-shot edge detection with SCESAME, which stands for **S**pectral **C**lustering-based **E**nsemble for **S**egment **A**nanything **M**odel **E**stimation, based on the recently proposed Segment Anything Model (SAM). SAM is a foundation model for segmentation tasks, and one of the interesting applications of SAM is Automatic Mask Generation (AMG), which generates zero-shot segmentation masks of an entire image. AMG can be applied to edge detection, but suffers from the problem of over-detecting edges. Edge detection with SCESAME overcomes this problem by three steps: (1) eliminating small generated masks, (2) combining masks by spectral clustering, taking into account mask positions and overlaps, and (3) removing artifacts after edge detection. We performed edge detection experiments on two datasets, BSDS500 and NYUDv2. Although our zero-shot approach is simple, the experimental results on BSDS500 showed almost identical performance to human performance and CNN-based methods from seven years ago. In the NYUDv2 experiments, it performed almost as well as recent CNN-based methods. These results indicate that our method effectively enhances the utility of SAM and can be a new direction in zero-shot edge detection methods.

1. Introduction

Foundation model [4] is the model that is pretrained on large-scale datasets and can be applied directly to downstream tasks, saving significant time and resources by eliminating the need to retrain the model for each specific task.

In the field of computer vision, several foundation models have been proposed for different tasks [13, 23, 35, 39, 40, 60]. The recently proposed Segment Anything Model (SAM) [23] is a foundation model for segmentation tasks,

Our code is available at <https://github.com/ymgw55/SCESAME>.

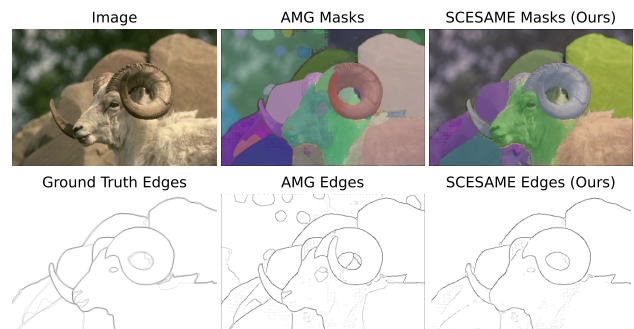


Figure 1. (upper row) Original image and masks generated by AMG and SCESAME. While AMG generates 54 masks, SCESAME combines them into 9 masks after removing smaller ones. (lower row) Ground truth edges and edges from the masks generated by AMG and SCESAME. Unlike AMG, which excessively detects edges from background and shadows, SCESAME restricts such edge detection.

capable of generating segmentation masks from different types of few-shot prompts, including points, bounding boxes, and segmentations. An interesting application of SAM is Automatic Mask Generation (AMG), which generates zero-shot segmentation masks of an entire image. This approach involves providing SAM with a regular grid of points as prompts for an input image, predicting a segmentation mask for each point, and generating the segmentation for the entire image.

One application of AMG is edge detection, one of the most important tasks in image processing and computer vision, which involves identifying the boundaries or other significant features within an image [37]. Edge detection is known to be applicable to downstream tasks such as segmentation [9, 28, 54] and object detection [6, 38, 59]. Therefore, when zero-shot edge detection with AMG shows strong performance, it seems applicable to various downstream tasks. In practice, however, it tends to overdetect edges [23], which is a significant problem.

Our motivation is to address such issue and propose a more effective zero-shot edge detection method based on AMG. To achieve this, we focused on spectral clustering, a method that uses the spectral information (eigenvalues and eigenvectors) of a graph with affinities between data as edges, considering the data as points in a new space, and then clustering in that space.

In this paper, we demonstrate a performance improvement in zero-shot edge detection with AMG by (1) removing smaller masks generated by AMG, (2) appropriately combing the remaining masks using spectral clustering, taking into account the mask positions and overlaps, and (3) eliminating artifacts that occur when edges are generated from masks. We refer to our method, which defines the affinity between masks generated by AMG and combines these masks using spectral clustering, as SCESAME: **S**pectral Clustering-based **E**nsemble for **S**egment Anything **M**odel **E**stimation. Figure 1 shows an example of masks and edge detection with AMG and SCESAME. While the AMG masks detect excess edges in the background and shadows, SCESAME remove small masks and effectively combine similar masks to reduce such edges.

Through experiments on BSDS500 [1] and NYUDv2 [43], we found that our method exhibits performance nearly equivalent to human performance and CNN-based methods from seven years ago for BSDS500, and nearly equivalent to recent CNN-based methods for NYUDv2, despite being a simple zero-shot technique. While there is still a significant gap between edge detection with SCESAME and the state-of-the-art (SOTA) approaches, these results indicate that our method effectively enhances the utility of SAM and can be a new direction in zero-shot edge detection methods.

2. Related Work

2.1. Edge Detection Method

Edge detection has a long history, with many traditional methods proposed before the advent of deep learning-based methods. In particular, the Sobel filter [24] is one of the earliest edge detection methods, with several advancements including the Canny method [7]. In addition, Felz-Hut [15] achieves refined edge detection by comparing differences between regions using a graph-based representation.

In recent years, deep learning approaches to edge detection have been introduced, including methods using Convolutional Neural Networks (CNN) [25, 50] and the Vision Transformer [37]. Loss functions are also proposed to account for ambiguities in annotations [25, 58].

2.2. SAM-based Model

SAM generates segmentation masks with few prompts, so several segmentation models using SAM have been pro-

posed. PerSAM [56] is a model that can segment specific concepts by one-shot tuning using a pair of an image and a mask. SAA+ [8] is a zero-shot anomaly detection model that uses Grounding DINO [29] to generate bounding boxes from text and then provides them as prompts to SAM. Track Anything [52] is a model that can track objects in a video with just a few clicks. HQ-SAM [22] is a model that performs additional learning for SAM parameters to generate more accurate masks.

2.3. Segmentation Method by Spectral Clustering

Many methods have been proposed for segmentation using spectral clustering. The method that combines a block-wise segmentation strategy with stochastic ensemble consensus [47] considers segment-level clustering and is related to our proposed approach. Linear spectral clustering [26] is a superpixel segmentation algorithm based on k -means clustering. The method of coupling local brightness, color, and texture cues using spectral clustering to detect contours has been proposed [1]. For unsupervised semantic segmentation, a parametric approach has been proposed that employs neural network-based eigenfunctions to generate embeddings for spectral clustering [11]. In the field of medical image segmentation, spectral clustering-based methods have been proposed by using prior information [49] or by identifying the tumor region [33].

3. Zero-Shot Edge Detection with AMG

In this section, we introduce the zero-shot edge detection pipeline using AMG, based on the original SAM paper [56]. Note that AMG for edge detection differs from standard AMG in terms of the number of points provided as prompts and the mask removal process, but we simply refer to AMG for edge detection as AMG throughout this paper unless there is confusion. For details on standard AMG, see the original SAM paper [56].

First, we explain AMG for edge detection. A 16×16 regular grid of points is given to SAM as prompts, which predicts three different scale masks at each point, generating a total of 768 masks. Non-Maximum Suppression (NMS) is then applied to the masks to remove redundant masks.

Next, we explain the edge detection process using the AMG masks. The logits of the masks are converted to probability values using an element-wise sigmoid function, and then a Sobel filter is applied for edge detection. During this process, all values except those at the boundaries of the masks are set to 0. Using the probabilities obtained for each mask, the maximum probability for each pixel over all the probabilities is determined, followed by min-max normalization over the entire image.

Finally, a Gaussian blur is applied, and then edge NMS [7, 12] is used to thin the edges, although the Gaussian blur for improving edge NMS is not mentioned in the

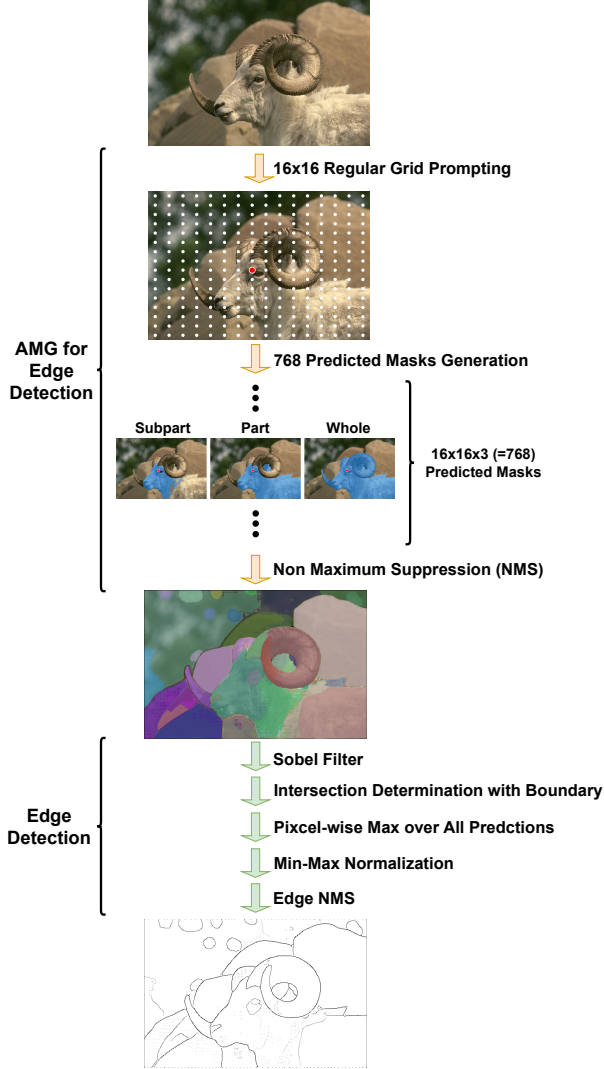


Figure 2. The zero-shot edge detection pipeline using AMG.

original paper.

Figure 2 illustrates the generation of masks by AMG and the edge detection process. It can be seen that in AMG, masks are generated at three scales: subpart, part, and whole, using a one point prompt. By performing edge detection on the masks remaining after NMS, edges are generated that reflect the contours of the masks. For details on the implementation, see § 6.2.

4. Spectral Clustering

In this section, based on the well-known tutorial on spectral clustering [48], we explain the algorithm for spectral clustering [34] that is used in our proposed method.

Consider an undirected graph G with vertex set $V = \{v_1, \dots, v_n\}$, and define an affinity matrix between vertices $\mathbf{W} = (w_{ij}) \in \mathbb{R}_{\geq 0}^{n \times n}$. Also, define the degree matrix $\mathbf{D} =$

Algorithm 1 Spectral Clustering

Input: Affinity matrix $\mathbf{W} = (w_{ij}) \in \mathbb{R}_{\geq 0}^{n \times n}$, number of clusters k .

Output: Clusters A_1, \dots, A_k , $A_i = \{j \mid y_j \in C_i\}$.

- 1: Compute the graph Laplacian \mathbf{L} .
- 2: Select the k smallest eigenvalues of \mathbf{L} and denote their corresponding eigenvectors as $\mathbf{u}_1, \dots, \mathbf{u}_k$.
- 3: Define the matrix $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_k] \in \mathbb{R}^{n \times k}$, and let the row vectors of \mathbf{U} be $\mathbf{y}_i \in \mathbb{R}^k$. Then, $\mathbf{U} = [\mathbf{y}_1, \dots, \mathbf{y}_n]^\top$.
- 4: Treat the vectors $(\mathbf{y}_i)_{i=1}^n$ as corresponding to each vertex, and use the k -means clustering to classify them into clusters C_1, \dots, C_k .




Figure 3. The utility of spectral clustering: (left) Inspired by a simple example [55], 100 points are randomly sampled from circles of radius 0.1, 0.5, and 1.0, and Gaussian noise is added. (middle) An affinity matrix is computed from a 10-nearest neighbor graph, and normalized spectral clustering with $k = 3$ classifies the three circles separately. (right) With k -means clustering, where $k = 3$, the points are classified based on their distance proximity only, and the three circles cannot be classified separately.

$\text{diag}(d_1, \dots, d_n) \in \mathbb{R}^{n \times n}$ where $d_i := \sum_{j=1}^n w_{ij}$.

The graph Laplacian for \mathbf{D}, \mathbf{W} is defined as:

$$\mathbf{L} := \mathbf{D} - \mathbf{W} \in \mathbb{R}^{n \times n}. \quad (1)$$

Since \mathbf{D}, \mathbf{W} are symmetric matrices, \mathbf{L} is also symmetric. For any $\mathbf{f} = (f_1, \dots, f_n)^\top \in \mathbb{R}^n$,

$$\mathbf{f}^\top \mathbf{L} \mathbf{f} = \frac{1}{2} \sum_{i,j=1}^n w_{ij} (f_i - f_j)^2 \geq 0 \quad (2)$$

holds, so \mathbf{L} is a positive semidefinite matrix [48]. By defining a constant vector with all components equal to 1 as $\mathbf{1} \in \mathbb{R}^n$, it follows from the definitions of \mathbf{D}, \mathbf{W} that $\mathbf{L}\mathbf{1} = (\mathbf{D} - \mathbf{W})\mathbf{1} = \mathbf{0}$. Thus, the smallest eigenvalue of \mathbf{L} is 0, and its corresponding eigenvector is $\mathbf{1}$.

The multiplicity k of the eigenvalue 0 of \mathbf{L} corresponds to the number of connected components in G , and if we denote their index sets as A_1, \dots, A_k , the eigenvectors are given by the indicator vectors $\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_k} \in \mathbb{R}^n$ [48]. Here, for $\mathbf{1}_{A_l} = (g_1, \dots, g_n)^\top$, $g_i = 1$ if $i \in A_l$, and $g_i = 0$ otherwise.

Next, we define a matrix $\mathbf{U} = [\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_k}] \in \mathbb{R}^{n \times k}$ with column vectors $\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_k}$. We denote the row




Figure 4. (upper row) Ground truth edges. (lower row) The edges from the SCESAME masks, with edges within 5 pixels of the image boundary highlighted in red.

vectors of \mathbf{U} as $\mathbf{y}_i \in \mathbb{R}^k$, that is $\mathbf{U} = [\mathbf{y}_1, \dots, \mathbf{y}_n]^\top$. Here, vertex v_i belongs to the connected component corresponding to the index where the component of \mathbf{y}_i is equal to 1. However, the graph does not necessarily have k connected components. In such cases, by considering k smallest eigenvalues of \mathbf{L} and the eigenvectors instead of $\mathbb{1}_{A_1}, \dots, \mathbb{1}_{A_k}$, we can redefine a matrix \mathbf{U} with these eigenvectors and use k -means clustering to determine the cluster to which the row vector \mathbf{y}_i corresponding to the vertex v_i belongs. The algorithm is described in Algorithm 1. For details, see the tutorial [48].

The matrix \mathbf{L} is precisely referred to as the unnormalized graph Laplacian, and the normalized graph Laplacian \mathbf{L}_{sym} is defined using the identity matrix $\mathbf{I} \in \mathbb{R}^{n \times n}$ as follows:

$$\mathbf{L}_{\text{sym}} := \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2} \in \mathbb{R}^{n \times n} \quad (3)$$

Normalized spectral clustering can be considered using the same procedure as in the unnormalized case.

A comparison between normalized spectral clustering and k -means clustering for two-dimensional points is shown in Figure 3. Normalized spectral clustering can classify the three circles separately, while k -means clustering can not do so. Since spectral clustering also uses k -means clustering, this example illustrates that the row vectors \mathbf{y}_i of the matrix \mathbf{U} provide suitable embeddings.

5. Proposed Method

This section describes three steps for edge detection with SCESAME: Spectral Clustering-based Ensemble for Segment Anything Model Estimation, based on AMG.

5.1. Removal of Small Noise Masks

Edge detection with AMG tends to be overly sensitive to minor changes that humans would not notice. For example, comparing the original image with the AMG masks in Figure 1, we can see that detail edges are detected in response to background light because AMG generates small masks.

However, since humans do not detect edges for such small changes, edge detection with AMG results in excessive detection of unnecessary edges.

Based on this observation, we propose a preprocessing step to remove small masks that would act as noise during edge detection. We sort the AMG masks by size and select only the top $1/t$ masks, where $t \in \mathbb{N}$. We call this operation Top Mask Selection (TMS). Despite its simplicity, edge detection with TMS improves performance over edge detection with AMG. See § 6.5 for details.

In the next section, we show a method to achieve higher performance edge detection based on the TMS masks.

5.2. Mask Ensemble Using Spectral Clustering

TMS selects masks based solely on their size, without taking into account their positions or overlaps. This may lead to the over-detection of unnecessary edges. To manage this issue, we propose merging the masks obtained by TMS. We define an affinity based on mask positions and overlaps, and use Spectral Clustering (SC) for this merging process.

Let the masks obtained by TMS be $\{M_i\}_{i=1}^n$. For each mask M_i , the position of the mask \mathbf{x}_i is determined by the center point of its bounding box. Let S_i denote the area of M_i . If we define the overlapping area between M_i and M_j as $S_i \cap S_j$, then the ratio of this overlap to the smaller mask area is $r_{ij} := S_i \cap S_j / \min\{S_i, S_j\} \in [0, 1]$.

Using the ratio of the overlapping area r_{ij} and the distance between the masks $\|\mathbf{x}_i - \mathbf{x}_j\|$, we model the affinity w_{ij} between M_i, M_j as follows:

$$w_{ij} := \exp\left(\frac{r_{ij}}{\tau}\right) \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{\sigma_i \sigma_j}\right) \quad (4)$$

In the above, $\tau \in \mathbb{R}_{>0}$ is the temperature hyperparameter, and $\sigma_i \in \mathbb{R}_{>0}$ is a normalization constant specific to \mathbf{x}_i , determined by the distance to the seventh closest point [55]. According to the definition of (4), w_{ij} increases as r_{ij} increases and as $\|\mathbf{x}_i - \mathbf{x}_j\|$ decreases. When $\tau < 1$, the affinity emphasize the ratio of the overlapping area rather than the distance between the masks. Note that, if $r_{ij} = 0$ in equation (4), it aligns with the local scaling affinity presented in [55].

From (4), a similarity matrix $\mathbf{W} = (w_{ij}) \in \mathbb{R}_{\geq 0}^{n \times n}$ can be derived. First, we set the number of clusters to $k = \max\{\lfloor n/c \rfloor, 2\} \in \mathbb{N}$, with $c (> 1) \in \mathbb{N}$, where $\lfloor \cdot \rfloor$ is floor function. Then, we perform normalized spectral clustering on $\{\mathbf{x}_i\}_{i=1}^n$. Finally, we generate new masks $\{\tilde{M}_i\}_{i=1}^k$ by combining the masks associated with the same cluster.

We call the entire procedure of combining masks using spectral clustering, including TMS, SCESAME: Spectral Clustering-based Ensemble for Segment Anything Model Estimation. SCESAME is designed to generate zero-shot segmentation masks such as AMG. While edge detection can be done using the SCESAME masks $\{\tilde{M}_i\}_{i=1}^k$, there

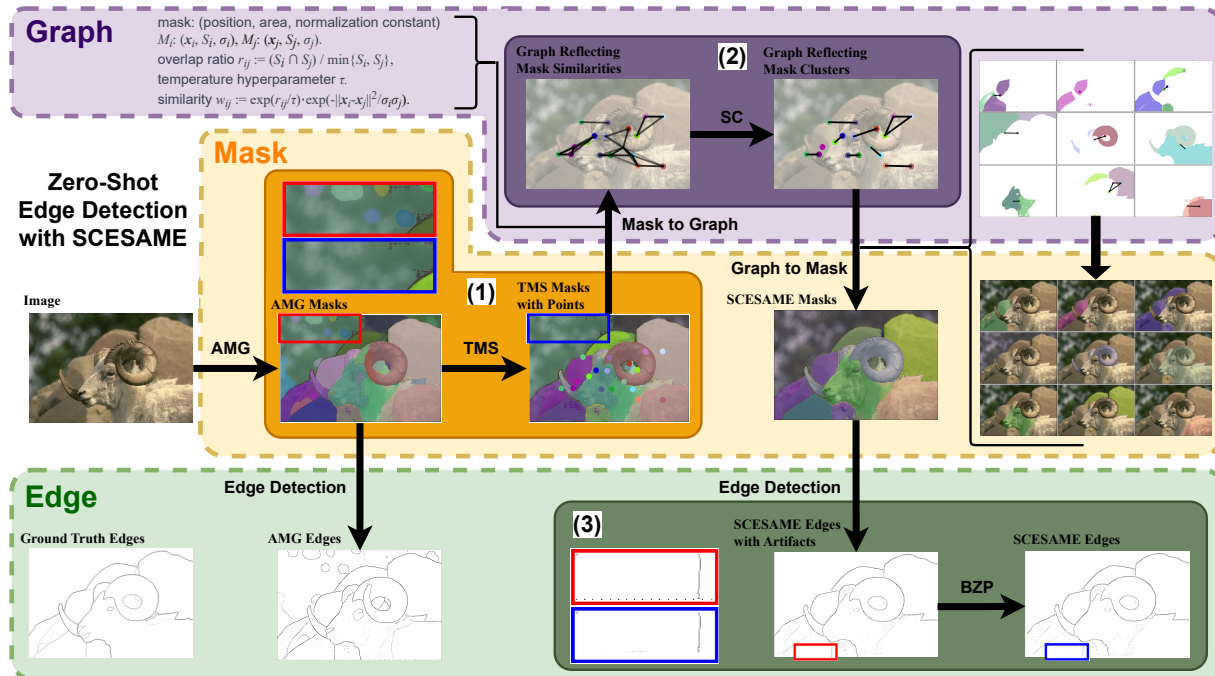


Figure 5. Representation of zero-shot edge detection with SCESAME for $t = 3, c = 2, p = 5$. For the AMG masks, (1) small masks are removed using TMS; (2) center points are assigned to each mask to define a graph, and SC is used to divide points into clusters, with masks corresponding to points in the same cluster combined to generate the SCESAME mask; (3) since the edges from the SCESAME masks contain boundary artifacts, BZP is used to remove them. In addition, in TMS and BZP, the changes before and after the process are magnified and shown in the red and blue regions, and in SC, all the points to be combined and their corresponding masks are shown.

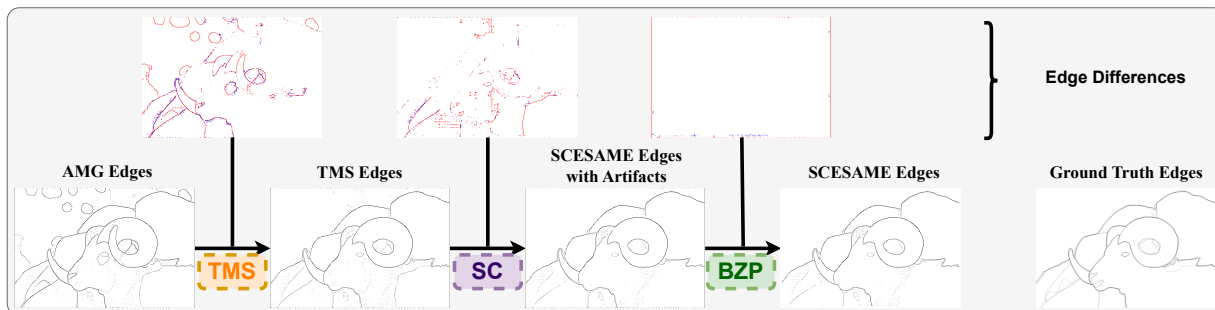


Figure 6. Differences in edges before and after processing in TMS, SC, BZP. Edges before processing are shown in red, and those after processing are shown in blue, with only differences greater than 0.05 shown for visibility.

may be artifacts when extracting edges from these masks. The next section describes a method for dealing with these artifacts.

5.3. Removal of Boundary Artifacts

Methods like AMG and SCESAME segment an entire image, resulting in the appearance of mask boundaries at the image boundaries. Consequently, when detecting edges from the mask, the mask boundaries tend to be detected as artifacts at the image border. We refer to these unintended artifacts as boundary artifacts. Figure 4 highlights in red the

SCESAME edges within 5 pixels of the image boundary and compares them with the ground truth edges. We can see that boundary artifacts appear in the SCESAME edges, although they are not in the ground truth edges.

For this reason, when detecting edges from AMG or SCESAME masks, we introduce a post-processing step termed Boundary Zero Padding (BZP). In this process, we fill all pixels within p pixels of the image boundary with zeros, where $p \in \mathbb{N}$. BZP step is applied after calculating the maximum probability for each pixel in the edge detection process. While there may be concerns about zero-padding

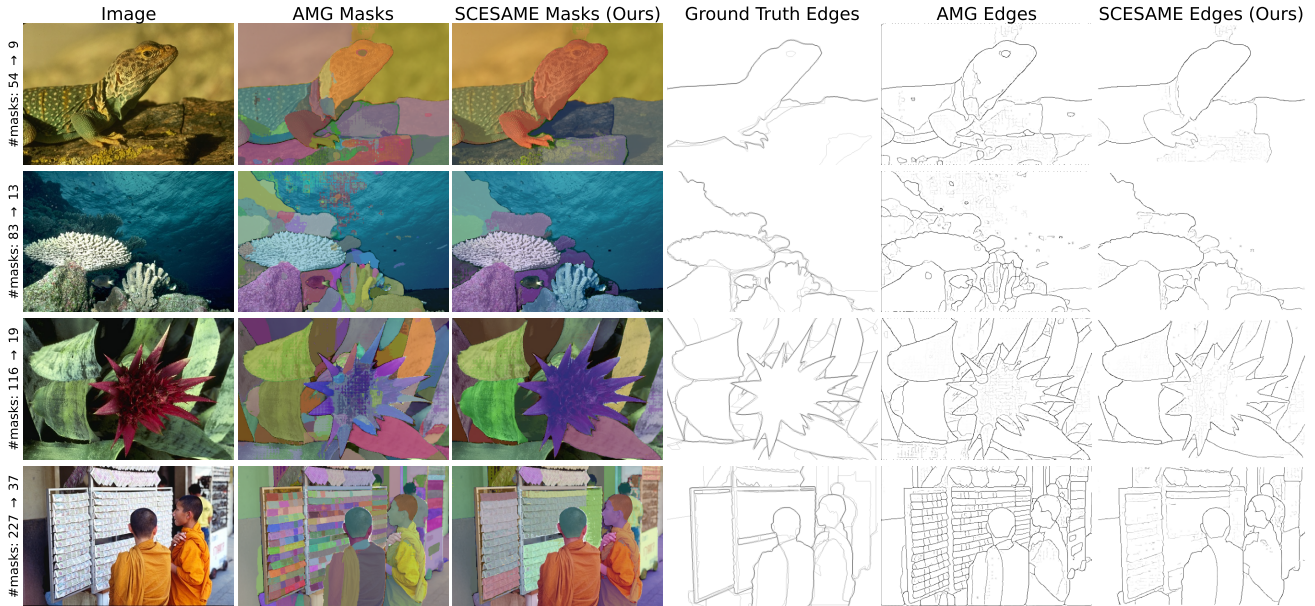


Figure 7. For some BSDS500 [1] images, this figure shows the images, the masks generated by AMG and SCESAME, the ground truth edges, and the edges generated by AMG and SCESAME. It also shows the change in the number of masks between AMG and SCESAME.

potentially obscuring true positive edges and thereby degrading performance, our experiments demonstrate the high effectiveness of BZP. Further details can be found in § 6.5.

5.4. Zero-Shot Edge Detection with SCESAME

In this paper, unless otherwise noted, zero-shot edge detection with SCESAME includes BZP processing. Figure 5 illustrates how zero-shot edge detection with SCESAME is constructed from the procedures described in § 5.1, § 5.2, and § 5.3. We can see that small masks are removed by TMS, remaining masks are adaptively combined by SC, and boundary artifacts are removed by BZP.

Figure 6 shows the edges generated by TMS, SC, and BZP, along with the differences between them. First, most of the edges from small masks are removed by TMS, followed by the removal of detail edges such as shadows by SC. Finally, boundary artifacts are removed by BZP.

Figure 7 shows examples of edge detection with AMG and SCESAME for some BSDS500 [1] images.

6. Experiments

6.1. Datasets

BSDS500 [1] consists of 500 RGB natural images, divided into 100 for training, 200 for validation, and 200 for testing. Each image was manually annotated by 4-9 annotators, with an average of 5 annotations per image.

NYUDv2 [43] contains 1449 indoor scenes consisting of RGB and HHA image pairs, divided into 381 for training, 414 for validation, and 654 for testing.

Since edge detection with SCESAME is a zero-shot technique designed for RGB images, we use only the test RGB images from both datasets.

6.2. Implementation Details

Since the original implementation of edge detection with AMG used in the SAM paper [23] is not publicly available, we reimplemented it based on the description in the paper. Specifically, we set the NMS threshold to 0.7, determined the mask boundary using the Sobel filter, applied a Gaussian blur with kernel size 3 before edge NMS, and used OpenCV [5]’s Structured Forests [12] model¹ for edge NMS².

For the value of τ in (4), as seen in § 5.2, we set $\tau = 0.5 < 1$ to emphasize the ratio of overlapping area between masks rather than their distance. For BZP, we set $p = 5$, and the values of t and c used in SCESAME are discussed in § 6.4. We use scikit-learn [36] to perform normalized spectral clustering and Python implementation for prediction evaluation³.

6.3. Evaluation Metric

We use Optimal Dataset Scale (ODS), Optimal Image Scale (OIS), and Average Precision (AP) [23, 37] as evalu-

¹https://github.com/opencv/opencv_extra/blob/master/testdata/cv/ximgproc/model.yml.gz

²In the original paper, Canny edge NMS [7] was used for edge NMS. However, in our environment, it did not produce the edges reported in the paper. This part needs further investigation and improvement.

³<https://github.com/Britefury/py-bsds500/>

Method		Pub.'Year	ODS	OIS	AP
Human [25]		ICLR'16	0.803	-	-
Traditional	Canny [7]	PAMI'86	0.600	0.640	0.580
	Felz-Hutt [15]	IJCV'04	0.610	0.640	0.560
	gPb-owt-ucm [1]	PAMI'10	0.726	0.757	0.696
	SCG [41]	NeurIPS'12	0.739	0.758	0.773
	Sketch Tokens [27]	CVPR'13	0.727	0.746	0.780
	PMI [21]	ECCV'14	0.741	0.769	0.799
	SE [12]	PAMI'14	0.746	0.767	0.803
	OEF [18]	CVPR'15	0.746	0.770	0.820
MES [44]	ICCV'15	0.756	0.776	0.756	
7 to 8-Year-Old CNN	DeepEdge [2]	CVPR'15	0.753	0.772	0.807
	CSCNN [20]	ArXiv'15	0.756	0.775	0.798
	MSC [45]	PAMI'15	0.756	0.776	0.787
	DeepContour [42]	CVPR'15	0.757	0.776	0.800
	HFL [3]	ICCV'15	0.767	0.788	0.795
	HED [50]	ICCV'15	0.788	0.808	0.840
	Deep Boundary [25]	ICLR'16	0.813	0.831	0.866
	CEDN [53]	CVPR'16	0.788	0.804	-
	RDS [31]	CVPR'16	0.792	0.810	0.818
	COB [32]	ECCV'16	0.793	0.820	0.859
SAM	SAM [23]	ICCV'23	0.768	0.786	0.794
	SAM [23] (Recalc.)	ICCV'23	0.730	0.754	0.729
	SAM-p5 (Our Baseline)	-	0.754	0.779	0.763
Ours	SCESAME-t2c2p5	-	0.796	0.812	0.780
	SCESAME-t2c3p5	-	0.797	0.811	0.768
	SCESAME-t3c2p5	-	0.800	0.814	0.773
	SCESAME-t3c3p5	-	0.796	0.809	0.753
SOTA	EDTER-MS [37]	CVPR'22	0.840	0.858	0.896
	EDTER-MS-VOC [37]	CVPR'22	0.848	0.865	0.903
	UAED-MS [58]	CVPR'23	0.837	0.855	0.897
	UAED-MS-VOC [58]	CVPR'23	0.844	0.864	0.905

Table 1. Results on BSDS500 [1] testing set. The notation t3c2p5 represents $t = 3, c = 2, p = 5$ and so on. The best three results, excluding SOTA methods, are highlighted in **red**, **blue**, and **purple**. SOTA methods are highlighted in **bold**. MS indicates multi-scale testing [37, 58], and VOC indicates training with additional PASCAL VOC data [14].

ation metrics. ODS is the F-score when selecting the optimal threshold for the entire dataset, and OIS is the F-score when selecting the optimal threshold for each image, with thresholds ranging from 0.01 to 0.99. AP is the integrated value of the precision-recall curve. Following previous works [30, 37, 50], the localization tolerance is set to 0.0075 for BSDS500 and 0.011 for NYUDv2. This value determines the maximum distance allowed between the predicted edge results and the ground truth in matching.

6.4. Results

On BSDS500. We compare zero-shot edge detection with SCESAME to the following models: human performance [25], traditional methods such as Canny [7], Felz-Hutt [15], gPb-owt-ucm [1], SCG [41], Sketch Tokens [27], PMI [21], SE [12], OEF [18], MES [44], as well as CNN-based models from 7-8 years ago including DeepEdge [2], CSCNN [20], MSC [45], DeepContour [42], HFL [3], HED [50], Deep Boundary [25], CEDN [53], RDS [31], COB [32], and state-of-the-art methods such as




Figure 8. The precision-recall curves on BSDS500.

EDTER [37], UAED [58]. The results of these experiments are taken from previous works [25, 37, 58]. We also compare the original results of SAM [23], our reimplementations of AMG, and the results of AMG using BZP. For the hyperparameters t and c of TMS and SC, we considered the values $(t, c) = (2, 2), (2, 3), (3, 2), (3, 3)$. The results are presented in Table 1. For BSDS500, the best ODS and OIS are obtained at $(t, c) = (3, 2)$. Edge detection with SCESAME surpasses traditional methods for both ODS and OIS. It outperforms most CNN-based methods from 7-8 years ago for ODS except Deep Boundary, and for OIS except Deep Boundary and COB. It also comes close to human performance. Compared to AMG, we observe improvements for both OIS and ODS. However, there is still a considerable gap in the results compared to the state-of-the-art methods.

For methods where results are available at different thresholds, the precision-recall curves are shown in Figure 8. Edge detection with SCESAME achieves a high F-score (ODS) and is close to human performance. In Table 1, edge detection with SCESAME does not show such a high AP compared to ODS and OIS, and its cause, as seen in Figure 8, is the lack of high recall values. We will discuss this further in § 7.

Although there are discrepancies between the original SAM results and our reimplementations, the trends observed using AMG with BZP are consistent with the original. Therefore, we use it as the baseline for our experiments.

On NYUDv2. We also evaluated the performance on RGB images using the NYUDv2 dataset. Our comparison involved SCESAME-t3c2p5 and SAM-p5 against var-

	Method	Pub.'Year	ODS	OIS	AP
Traditional	gPb-ucm [1]	PAMI'11	0.632	0.661	0.562
	Silberman <i>et al.</i> [43]	ECCV'12	0.658	0.661	-
	gPb+NG [16]	CVPR'13	0.687	0.716	0.629
	SE [12]	PAMI'14	0.695	0.708	0.679
	SE+NG+ [17]	ECCV'14	0.706	0.734	0.738
	OEF [18]	CVPR'15	0.651	0.667	-
	SemiContour [57]	CVPR'16	0.680	0.700	0.690
CNN-based	HED [50]	ICCV'15	0.720	0.734	0.734
	RCF [30]	CVPR'17	0.729	0.742	-
	AMH-Net [51]	NeurIPS'17	0.744	0.758	0.765
	LPCB [10]	ECCV'18	0.739	0.754	-
	BDCN [19]	CVPR'19	0.748	0.763	0.770
	PiDiNet [46]	ICCV'21	0.733	0.747	-
	SAM-p5 (Our Baseline)	-	0.699	0.719	0.707
SCESAME-t3c2p5 (Ours)	-	0.742	0.754	0.707	
EDTER [37] (SOTA)		CVPR'22	0.774	0.789	0.797

Table 2. Results on NYUDv2 [43] testing set. The notation t3c2p5 represents $t = 3, c = 2, p = 5$ and so on. The best three results, excluding the SOTA method, are highlighted in **red**, **blue**, and **purple**. The SOTA method is highlighted in **bold**.

ious models, including traditional methods such as gPb-ucm [1], Silberman *et al.* [43], gPb+NG [16], SE [12], SE+NG+ [17], OEF [18], SemiContour [57], and CNN-based models such as HED [50], RCF [30], AMH-Net [51], LPCB [10], BDCN [19], PiDiNet [46], and the state-of-the-art method EDTER [37]. The experimental results are taken from previous work [37]. These results are presented in Table 2. Edge detection with SCESAME outperforms traditional methods and performs almost as well as CNN-based methods for ODS and OIS. Similar to the BSDS500 results, we observe an improvement over AMG for OIS and ODS. However, there remains a noticeable performance gap when compared with the state-of-the-art method.

Note that there are fewer methods tested on NYUDv2 compared to BSDS500, and CNN-based methods can further improve their performance by using both RGB and HHA images during training.

6.5. Ablation Study

As seen in § 5, TMS, SC, and BZP are independent processes. Therefore, an ablation study is performed on BSDS500 using the parameters that gave the best ODS and OIS performance: $t = 3, c = 2$, and $p = 5$. The experimental results are presented in Table 3. It is evident that the performance improves when BZP is used in AMG, TMS, SC, and SCESAME. While TMS-t3p5 outperforms SC-c2p5, SCESAME-t3c2p5 shows superior ODS and OIS compared to TMS-t3p5. This suggests the importance of combining TMS and SC. Note that for $t = 3$ and $c = 2$, the proportion of mask removal and mask combination is different and that TMS-t3p5 had the highest AP among them.

Method	TMS	SC	BZP	ODS	OIS	AP
SAM (Recalc.)				0.730	0.754	0.729
SAM-p5			✓	0.754	0.779	0.763
TMS-t3	✓			0.757	0.769	0.718
TMS-t3p5	✓		✓	0.797	0.812	0.792
SC-c2		✓		0.743	0.762	0.731
SC-c2p5		✓	✓	0.771	0.792	0.773
SCESAME-t3c2	✓	✓		0.753	0.764	0.693
SCESAME-t3c2p5	✓	✓	✓	0.800	0.814	0.773

Table 3. Ablation study on BSDS500 [1] testing set. The notation t3c2p5 represents $t = 3, c = 2, p = 5$ and so on.

7. Discussion

In this section, we first discuss the limitations of edge detection with SCESAME based on the results in § 6.4. Then we give some suggestions for future work and conclude.

Limitation. Edge detection with SCESAME shows a gap compared to state-of-the-art methods. The lower AP compared to ODS and OIS can be attributed to the suppression of edges during mask removal and combination. While BZP is effective, it can also fill true positive edges with zeros, contributing to low recall. AMG may be more practical than SCESAME for detecting finer edges. In addition, the datasets we used are based on a few annotations, so the edges removed by SCESAME (or overdetected by AMG) are not necessarily redundant.

Future Work. Instead of TMS, selection based on mask importance, random selection, or selection from mask features can preserve even small masks that are critical for edge detection. Affinity in (4), based on mask position and overlap ratio, is a simple model with room for improvement. The parameters t, c, p in TMS, SC, and BZP are currently fixed, so choosing optimal values for each image may be beneficial. Consideration of a few-shot fine-tuning model could further improve its performance. In this study, SCESAME was used for edge detection, but if AMG is used in downstream tasks, SCESAME can also be used.

Conclusion. This paper proposes a novel zero-shot edge detection with SCESAME based on AMG. This method, which consists of three steps, overcomes the over-detection problem of AMG edges. Experimental results on the BSDS500 and NYUDv2 show that despite being a simple zero-shot method, our approach exhibits performance comparable to human performance and recent CNN-based methods. These results indicate that our method effectively enhances the utility of SAM and can be a new direction in zero-shot edge detection methods.

References

- [1] Pablo Arbelaez, Michael Maire, Charless C. Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):898–916, 2011. [2](#), [6](#), [7](#), [8](#)
- [2] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Deepedge: A multi-scale bifurcated deep network for top-down contour detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4380–4389. IEEE Computer Society, 2015. [7](#)
- [3] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. High-for-low and low-for-high: Efficient boundary detection from deep object features and its applications to high-level vision. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 504–512. IEEE Computer Society, 2015. [7](#)
- [4] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ B. Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen Creel, Jared Quincy Davis, Dorottya Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kavin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kudithipudi, and et al. On the opportunities and risks of foundation models. *CoRR*, abs/2108.07258, 2021. [1](#)
- [5] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000. [6](#)
- [6] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5320–5329. IEEE Computer Society, 2017. [1](#)
- [7] John F. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(6):679–698, 1986. [2](#), [6](#), [7](#)
- [8] Yunkang Cao, Xiaohao Xu, Chen Sun, Yuqi Cheng, Zongwei Du, Liang Gao, and Weiming Shen. Segment any anomaly without training via hybrid prompt regularization. *CoRR*, abs/2305.10724, 2023. [2](#)
- [9] Tianheng Cheng, Xinggang Wang, Lichao Huang, and Wenyu Liu. Boundary-preserving mask R-CNN. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XIV*, volume 12359 of *Lecture Notes in Computer Science*, pages 660–676. Springer, 2020. [1](#)
- [10] Ruoxi Deng, Chunhua Shen, Shengjun Liu, Huibing Wang, and Xinru Liu. Learning to predict crisp boundaries. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part VI*, volume 11210 of *Lecture Notes in Computer Science*, pages 570–586. Springer, 2018. [8](#)
- [11] Zhijie Deng and Yucen Luo. Learning neural eigenfunctions for unsupervised semantic segmentation. *CoRR*, abs/2304.02841, 2023. [2](#)
- [12] Piotr Dollár and C. Lawrence Zitnick. Fast edge detection using structured forests. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(8):1558–1570, 2015. [2](#), [6](#), [7](#), [8](#)
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [1](#)
- [14] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010. [7](#)
- [15] Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Efficient graph-based image segmentation. *Int. J. Comput. Vis.*, 59(2):167–181, 2004. [2](#), [7](#)
- [16] Saurabh Gupta, Pablo Arbelaez, and Jitendra Malik. Perceptual organization and recognition of indoor scenes from RGB-D images. In *2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, June 23-28, 2013*, pages 564–571. IEEE Computer Society, 2013. [8](#)
- [17] Saurabh Gupta, Ross B. Girshick, Pablo Arbelaez, and Jitendra Malik. Learning rich features from RGB-D images for object detection and segmentation. *CoRR*, abs/1407.5736, 2014. [8](#)
- [18] Sam Hallman and Charless C. Fowlkes. Oriented edge forests for boundary detection. *CoRR*, abs/1412.4181, 2014. [7](#), [8](#)
- [19] Jianzhong He, Shiliang Zhang, Ming Yang, Yanhu Shan, and Tiejun Huang. Bi-directional cascade network for perceptual edge detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 3828–3837. Computer Vision Foundation / IEEE, 2019. [8](#)
- [20] Jyh-Jing Hwang and Tyng-Luh Liu. Pixel-wise deep learning for contour detection, 2015. [7](#)
- [21] Phillip Isola, Daniel Zoran, Dilip Krishnan, and Edward H. Adelson. Crisp boundary detection using pointwise mutual information. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part III*, volume 8691 of *Lecture Notes in Computer Science*, pages 799–814. Springer, 2014. [7](#)

- [22] Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan Liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Segment anything in high quality, 2023. [2](#)
- [23] 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, October 2023. [1](#), [6](#), [7](#)
- [24] Josef Kittler. On the accuracy of the sobel edge detector. *Image Vis. Comput.*, 1:37–42, 1983. [2](#)
- [25] Iasonas Kokkinos. Surpassing humans in boundary detection using deep learning. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. [2](#), [7](#)
- [26] Zhengqin Li and Jiansheng Chen. Superpixel segmentation using linear spectral clustering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 1356–1363. IEEE Computer Society, 2015. [2](#)
- [27] Joseph J. Lim, C. Lawrence Zitnick, and Piotr Dollár. Sketch tokens: A learned mid-level representation for contour and object detection. In *2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, June 23-28, 2013*, pages 3158–3165. IEEE Computer Society, 2013. [7](#)
- [28] Jiang-Jiang Liu, Qibin Hou, and Ming-Ming Cheng. Dynamic feature integration for simultaneous detection of salient object, edge, and skeleton. *IEEE Trans. Image Process.*, 29:8652–8667, 2020. [1](#)
- [29] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. Grounding DINO: marrying DINO with grounded pre-training for open-set object detection. *CoRR*, abs/2303.05499, 2023. [2](#)
- [30] Yun Liu, Ming-Ming Cheng, Xiaowei Hu, Kai Wang, and Xiang Bai. Richer convolutional features for edge detection. *CoRR*, abs/1612.02103, 2016. [7](#), [8](#)
- [31] Yu Liu and Michael S. Lew. Learning relaxed deep supervision for better edge detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 231–240. IEEE Computer Society, 2016. [7](#)
- [32] Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Andrés Arbeláez, and Luc Van Gool. Convolutional oriented boundaries. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I*, volume 9905 of *Lecture Notes in Computer Science*, pages 580–596. Springer, 2016. [7](#)
- [33] Angulakshmi Maruthamuthu and Lakshmi Priya Gnanapandithan G. Brain tumour segmentation from MRI using superpixels based spectral clustering. *J. King Saud Univ. Comput. Inf. Sci.*, 32(10):1182–1193, 2020. [2](#)
- [34] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In Thomas G. Dietterich, Suzanna Becker, and Zoubin Ghahramani, editors, *Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001, December 3-8, 2001, Vancouver, British Columbia, Canada]*, pages 849–856. MIT Press, 2001. [3](#)
- [35] OpenAI. Gpt-4v(ision) system card. 2023. [1](#)
- [36] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. [6](#)
- [37] Mengyang Pu, Yaping Huang, Yuming Liu, Qingji Guan, and Haibin Ling. EDTER: edge detection with transformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 1392–1402. IEEE, 2022. [1](#), [2](#), [6](#), [7](#), [8](#)
- [38] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jägersand. Basnet: Boundary-aware salient object detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 7479–7489. Computer Vision Foundation / IEEE, 2019. [1](#)
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021. [1](#)
- [40] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *CoRR*, abs/2204.06125, 2022. [1](#)
- [41] Xiaofeng Ren and Liefeng Bo. Discriminatively trained sparse code gradients for contour detection. In Peter L. Bartlett, Fernando C. N. Pereira, Christopher J. C. Burges, Léon Bottou, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States*, pages 593–601, 2012. [7](#)
- [42] Wei Shen, Xinggang Wang, Yan Wang, Xiang Bai, and Zhijiang Zhang. Deepcontour: A deep convolutional feature learned by positive-sharing loss for contour detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 3982–3991. IEEE Computer Society, 2015. [7](#)
- [43] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from RGBD images. In Andrew W. Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid, editors, *Computer Vision - ECCV 2012 - 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part V*, volume 7576 of *Lecture Notes in Computer Science*, pages 746–760. Springer, 2012. [2](#), [6](#), [8](#)

- [44] Amos Sironi, Vincent Lepetit, and Pascal Fua. Projection onto the manifold of elongated structures for accurate extraction. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 316–324. IEEE Computer Society, 2015. [7](#)
- [45] Amos Sironi, Engin Türetken, Vincent Lepetit, and Pascal Fua. Multiscale centerline detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(7):1327–1341, 2016. [7](#)
- [46] Zhuo Su, Wenzhe Liu, Zitong Yu, Dewen Hu, Qing Liao, Qi Tian, Matti Pietikäinen, and Li Liu. Pixel difference networks for efficient edge detection. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 5097–5107. IEEE, 2021. [8](#)
- [47] Frederick Tung, Alexander Wong, and David A. Clausi. Enabling scalable spectral clustering for image segmentation. *Pattern Recognit.*, 43(12):4069–4076, 2010. [2](#)
- [48] Ulrike von Luxburg. A tutorial on spectral clustering. *Stat. Comput.*, 17(4):395–416, 2007. [3, 4](#)
- [49] Kaijian Xia, Xiaoqing Gu, and Yudong Zhang. Oriented grouping-constrained spectral clustering for medical imaging segmentation. *Multim. Syst.*, 26(1):27–36, 2020. [2](#)
- [50] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 1395–1403. IEEE Computer Society, 2015. [2, 7, 8](#)
- [51] Dan Xu, Wanli Ouyang, Xavier Alameda-Pineda, Elisa Ricci, Xiaogang Wang, and Nicu Sebe. Learning deep structured multi-scale features using attention-gated crfs for contour prediction. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 3961–3970, 2017. [8](#)
- [52] Jinyu Yang, Mingqi Gao, Zhe Li, Shang Gao, Fangjing Wang, and Feng Zheng. Track anything: Segment anything meets videos. *CoRR*, abs/2304.11968, 2023. [2](#)
- [53] Jimei Yang, Brian L. Price, Scott Cohen, Honglak Lee, and Ming-Hsuan Yang. Object contour detection with a fully convolutional encoder-decoder network. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 193–202. IEEE Computer Society, 2016. [7](#)
- [54] Zhiding Yu, Rui Huang, Wonmin Byeon, Sifei Liu, Guilin Liu, Thomas M. Breuel, Anima Anandkumar, and Jan Kautz. Coupled segmentation and edge learning via dynamic graph propagation. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 4919–4932, 2021. [1](#)
- [55] Lihi Zelnik-Manor and Pietro Perona. Self-tuning spectral clustering. In *Advances in Neural Information Processing Systems 17 [Neural Information Processing Systems, NIPS 2004, December 13-18, 2004, Vancouver, British Columbia, Canada]*, pages 1601–1608, 2004. [3, 4](#)
- [56] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *CoRR*, abs/2305.03048, 2023. [2](#)
- [57] Zizhao Zhang, Fuyong Xing, Xiaoshuang Shi, and Lin Yang. Semicontour: A semi-supervised learning approach for contour detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 251–259. IEEE Computer Society, 2016. [8](#)
- [58] Caixia Zhou, Yaping Huang, Mengyang Pu, Qingji Guan, Li Huang, and Haibin Ling. The treasure beneath multiple annotations: An uncertainty-aware edge detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15507–15517, June 2023. [2, 7](#)
- [59] Hongwei Zhu, Peng Li, Haoran Xie, Xuefeng Yan, Dong Liang, Dapeng Chen, Mingqiang Wei, and Jing Qin. I can find you! boundary-guided separated attention network for camouflaged object detection. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 3608–3616. AAAI Press, 2022. [1](#)
- [60] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. In *NeurIPS 2023*, July 2023. [1](#)