

SAUGE: Taming SAM for Uncertainty-Aligned Multi-Granularity Edge Detection

Xing Liufu¹, Chaolei Tan¹, Xiaotong Lin¹, Yonggang Qi², Jinxuan Li¹, Jian-Fang Hu^{1,3,4*}

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

²Beijing University of Posts and Telecommunications, Beijing, China

³Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, Guangzhou, China

⁴Guangdong Province Key Laboratory of Information Security Technology, China

{liufux,tanchlei,linxt29,lijx267}@mail2.sysu.edu.cn, qi yg@bupt.edu.cn, hujf5@mail.sysu.edu.cn

Abstract

Edge labels are typically at various granularity levels owing to the varying preferences of annotators, thus handling the subjectivity of per-pixel labels has been a focal point for edge detection. Previous methods often employ a simple voting strategy to diminish such label uncertainty or impose a strong assumption of labels with a pre-defined distribution, e.g., Gaussian. In this work, we unveil that the segment anything model (SAM) provides strong prior knowledge to model the uncertainty in edge labels. Our key insight is that the intermediate SAM features inherently correspond to object edges at various granularities, which reflects different edge options due to uncertainty. Therefore, we attempt to align uncertainty with granularity by regressing intermediate SAM features from different layers to object edges at multi-granularity levels. In doing so, the model can fully and explicitly explore diverse “uncertainties” in a data-driven fashion. Specifically, we inject a lightweight module ($\sim 1.5\%$ additional parameters) into the frozen SAM to progressively fuse and adapt its intermediate features to estimate edges from coarse to fine. It is crucial to normalize the granularity level of human edge labels to match their innate uncertainty. For this, we simply perform linear blending to the real edge labels at hand to create pseudo labels with varying granularities. Consequently, our uncertainty-aligned edge detector can flexibly produce edges at any desired granularity (including an optimal one). Thanks to SAM, our model uniquely demonstrates strong generalizability for cross-dataset edge detection. Extensive experimental results on BSDS500, Muticue and NYUDv2 validate our model’s superiority.

Introduction

Edge detection is a fundamental AI task in low-level vision that plays a crucial role in image understanding. It is of great value in supporting various high-level vision tasks, including semantic segmentation (Yu et al. 2021), image enhancement (Nazeri et al. 2019; Xu, Wang, and Lu 2023), object detection (Qin et al. 2019; Liu, Hou, and Cheng 2020; Yao and Wang 2023), image translation (Jiang et al. 2023), etc.

However, due to the observers’ diverse levels of visual perception (Zhou et al. 2023), there are often multiple

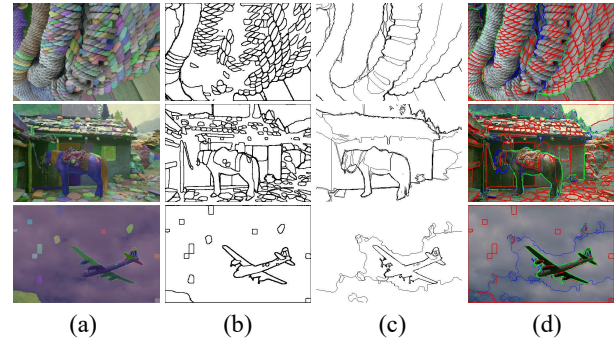


Figure 1: Comparison of edges obtained by SAM and manually annotated ground truth: (a) shows images from BSDS set; (b) illustrates edge maps generated by SAM; (c) presents edge annotations; (d) compares (b) and (c), where green indicates the shared edges, red represents edges in (b) but not in (c), and blue represents edges in (c) but not in (b).

ground truth edge maps for a given image. The inconsistency between different edge maps can cause uncertainty in decision-making, complicate the model training, and reduce the overall performance. The scarcity of human annotations may further exacerbate this issue.

Most previous works (Liu et al. 2017; He et al. 2019; Pu et al. 2022) tend to overlook the uncertainty in annotations by treating each edge map as equally valid, or fuse variant edge maps into a unified version by a simple voting strategy. As a result, these methods typically fail to capture the inherent subjectivity in the annotations. They can only produce a single edge map for a given image, lacking the ability to control granularity. This limitation severely declines their applicability and scalability in real-world scenarios.

To cope with the uncertainty issue of edge labels is non-trivial and remains relatively under-explored. UAED (Zhou et al. 2023) handles edge labels from a probabilistic perspective. It is assumed that the uncertainty of edge labels could be modeled using a Gaussian distribution, allowing the generation of multiple edges by sampling from the distribution. However, the distribution assumption imposes strict constraints, yielding limited diversity and a lack of control over granularity. RankED (Cetinkaya, Kalkan, and Akbas 2024) approached to calculate the label certainty of each

*Corresponding author.

pixel among different annotators, then favor the pixels with higher confidence during training. Consequently, pixel annotations with high uncertainty are largely overlooked. More recently, MuGE (Zhou et al. 2024) devised a granularity-controllable edge detector, which remedies label uncertainty by producing edge maps at various granularity levels. They naively assign a binary granularity score to its most simple (0) and complex (1) edge maps for an image to train a binary classifier, which is then applied to assign granularity scores (0-1) to rest edge maps. Then the granularity score is explicitly embedded in the edge detector, making it measurable and controllable. Unfortunately, this may easily introduce bias as a complex edge map might be incorrectly marked as simple if all the annotations are complex, and vice versa.

In this work, following the core idea of MuGE which models uncertainty by generating multi-scale edge maps, we seek to develop an edge detector that can produce granularity-controllable outputs without requiring granularity labels for edge maps. This can be achieved by progressively fusing and projecting intermediate features of the segmentation foundation model SAM (Kirillov et al. 2023) to edge map variants with increasing granularities. We demonstrate that SAM is highly effective in edge detection in the multi-scale setting. Intuitively, SAM excels in locating object boundaries, thus potentially providing strong prior knowledge of edge maps. Besides, we find out that the features from intermediate layers of SAM naturally encode rich information about various granularities for object edges.

However, unlocking the full potential of SAM for multi-granularity edge detection presents significant challenges. As aforementioned, SAM inclines to provide redundant details and coarse object boundaries rather than detailed internal edges, as shown in Figure 1 (d). To fill the gap and capture uncertainty through granularity modeling, we incorporate a lightweight feature transfer network into SAM that intermediate features are gradually fused and projected into side outputs of edges at increasingly complex levels. To supervise the transfer network, linear blending is performed to the real edge labels to synthesize pseudo edges at diverse granularities. This step crucially normalizes the granularity levels of all edge labels, easing the convergence of network training. Moreover, we develop a novel diversity loss to encourage the obtained side edge maps to be sufficiently diverse, further enhancing uncertainty modeling.

Our main contributions can be summarized as follows:

- We propose a novel edge detector, named SAUGE, which sidesteps the uncertainty difficulty by modeling edge detection in a multi-granularity setting based on SAM. The obtained edge detector enables edge detection at any desired granularity level.
- A lightweight module is proposed to be injected into the frozen SAM to progressively fuse and transfer its intermediate features to generate edges from coarse to fine. Linear blending is performed to real edge labels to synthesize normalized supervision cues.
- We conduct extensive experiments on the BSDS500, Multicue and NYUDv2 datasets, and the results demonstrate that our model achieves the new state-of-the-art

and suggest a strong generalizability on unseen datasets.

Related Work

Edge detection. Edge detection has been a significant research area for a long time. Traditional techniques such as (Kittler 1983; Canny 1986) rely on gradient computation on low-level features which are easy to be affected by noise. Over the past decade, deep learning-based approaches, such as (Shen et al. 2015; Liu and Lew 2016; He et al. 2019; Su et al. 2021), have risen to prominence, focusing on designing network architectures that surpass human-level performance. The majority of these methods, including RCF (Liu et al. 2017), leverage pre-trained VGG16 (Simonyan and Zisserman 2014) as their backbone. More recently, advanced methods like EDTER (Pu et al. 2022) incorporated Transformers to enhance edge detection. Research based on uncertainty explores the uncertainty caused by multiple labels. UAED (Zhou et al. 2023) models edge maps as multivariate Gaussian distribution and use predicted variance to measure the uncertainty, RankED (Cetinkaya, Kalkan, and Akbas 2024) sorts pixels to balance edge and non edge pixels and promote higher label certainty for high confidence edges. MuGE (Zhou et al. 2024) integrates encoded edge granularity into feature maps to produce edge maps at multiple granularities for alleviating uncertainty. Recent image generation methods such as DiffusionEdge (Ye et al. 2024) applies diffusion models to generate crisp edge maps.

Most of these methods do not utilize prior knowledge from advanced tasks. In contrast, we explore the priors of semantic-aware features in SAM and construct multiple granularity edge maps to address the uncertainty.

Exploring SAM for downstream tasks. The Segment Anything Model (SAM) (Kirillov et al. 2023) accepts intuitive prompts (points or bounding boxes), and has established a new benchmark in natural image segmentation. It has demonstrated impressive performance across various downstream tasks, including medical imaging (Ma et al. 2024; Gu et al. 2024), object segmentation in challenging conditions (Chen et al. 2024; Zhang et al. 2024; Ke et al. 2024) and image inpainting (Yu et al. 2023).

The work most closely related to ours is EdgeSAM (Yang et al. 2024), which introduces an adapter module to fine-tune SAM for edge detection. However, their approach uses SAM in a simplistic manner, without fully exploring its potential and modeling the uncertainty. In contrast, we leverage the granularity-aware prior knowledge embedded in SAM’s features and attempt to align uncertainty with granularity.

Method

The overall framework of the proposed uncertainty-aligned edge detector (SAUGE) is presented in Figure 2. As shown, our SAUGE is built upon the pre-trained SAM, utilizing the proposed lightweight Side Transfer Network (STN) to explicitly explore the knowledge embedded in SAM for multiple granularity edge detection. STN gradually uses Feature Fuse Block (FFB) to fuse intermediate features extracted from different stages of SAM to supplement edge-aware details, thereby gradually constructing side outputs represent-

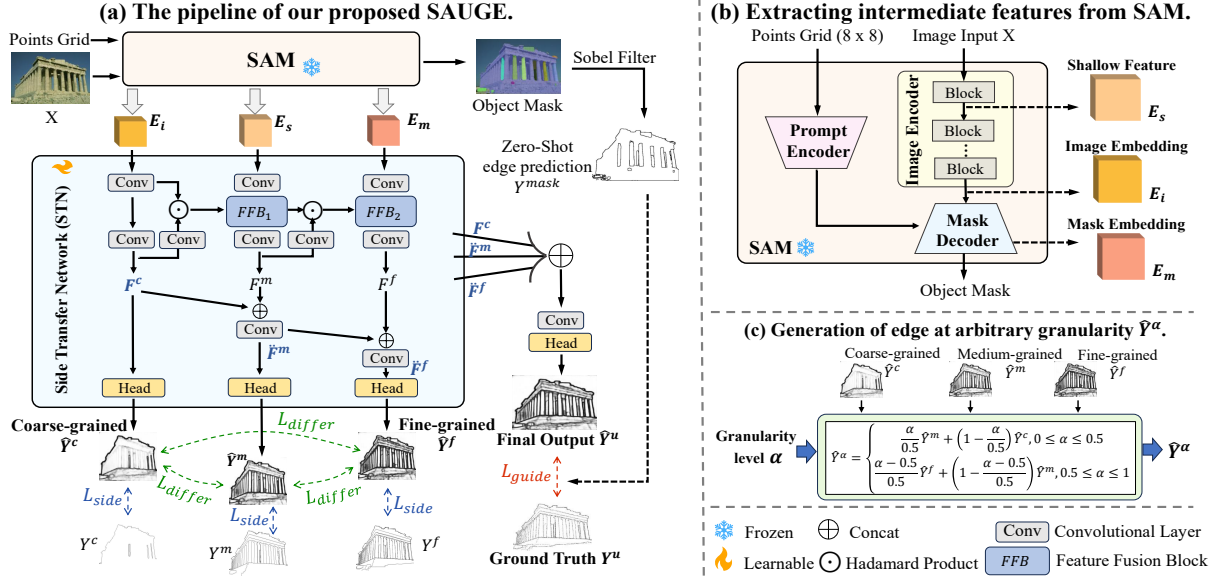


Figure 2: The overall framework of SAUGE. (a) illustrates the pipeline of SAUGE. We extract the intermediate SAM features and feed them into STN, which constructs edges at multiple granularity levels to align uncertainty with granularity. The final output \hat{Y}^u is obtained by merging the features of side outputs. We devise losses $\{L_{side}, L_{differ}, L_{guide}\}$ to supervise the side outputs, promote the pairwise diversity among them, and guide edge learning using SAM masks, respectively. (b) shows the features extracted from SAM. (c) demonstrates the generation of \hat{Y}^α at any granularity level α in a controllable manner.

ing edge at varied granularity levels. Finally, STN fuses all the features of side outputs to generate final output. To ensure the quality of side outputs, we construct pseudo labels corresponding to varied granularity levels to supervise and emphasize diversity of side outputs. We also propose to use the masks outputted by SAM to improve our edge learning.

Aligning uncertainty through granularity modeling

The proposed Side Transfer Network (STN) aims to fill the gap and capture uncertainty through granularity modeling. Intermediate features of SAM are gradually fused and projected into side outputs of edges at increasingly complex levels, allowing the model to fully and explicitly explore diverse uncertainties with granularity.

Specifically, we feed the image $X \in \mathbb{R}^{H \times W \times 3}$ into the frozen SAM with the prompt of 8×8 points grid, obtaining shallow feature maps E_s (i.e. the output of the first block in the encoder), image embeddings E_i from the SAM encoder, and mask embeddings E_m from the decoder.

For features $\{E_s, E_i\}$, we transform them into edge-aware features $\{E_s^e, E_i^e\} \in \mathbb{R}^{D \times D \times C_1}$ using learnable convolution layers, where $D \times D$ represents the spatial resolution and C_1 is feature channels. For the mask embedding feature E_m , we first rescale the feature channel to C_2 using learnable convolutional layers, and then reshape the features into edge-aware features $E_m^e \in \mathbb{R}^{D \times D \times (C_2 \times (8 \times 8))}$, ensuring all edge-aware details from the prompt is integrated. The detailed process is formulated as follows:

$$E_i^e = W_i^t(E_i), E_s^e = W_s^t(E_s), E_m^e = R(W_m^t(E_m)) \quad (1)$$

where $W_i^t(\cdot), W_s^t(\cdot), W_m^t(\cdot)$ are convolution layers, $R(\cdot)$ is

the reshape operator.

The Feature Fusion Block (FFB) is developed inside the STN to gradually aggregate edge-aware features $\{E_s^e, E_i^e, E_m^e\}$. Inspired by (Wu et al. 2023), we construct the FFB as a stack of cross-attention operations and Gated-Dconv Feed-Forward Network proposed by (Zamir et al. 2022). The process of constructing edge features at different granularity levels can be summarized as follows:

$$\begin{aligned} F^c &= W_c^h(E_i^e), E^m = \text{FFB}_1(E_i^e \cdot W_c^g(F^c), E_s^e), \\ F^m &= W_m^h(E^m), E^f = E^m \cdot W_m^g(F^m), \\ F^f &= W_f^h(E^f), E^{f'} = \text{FFB}_2(E^f, E_m^e), \end{aligned} \quad (2)$$

where $W_c^h(\cdot), W_m^h(\cdot), W_f^h(\cdot), W_c^g(\cdot), W_m^g(\cdot)$ are convolution layers, $\text{FFB}(\cdot)$ is the Feature Fuse Block. Specifically, we directly generate the coarse-grained edge feature F^c using the E_i^e . Then, we use the first FFB to fuse E_i^e and E_s^e for including more details, which forms our medium-grained edge feature F^m . The features are further fused with E_m^e by another FFB, resulting our fine-grained edge feature F^f .

With the edge features $\{F^c, F^m, F^f\}$ prepared, we then employ a shared classification head H to generate the side outputs $\{\hat{Y}^c, \hat{Y}^m, \hat{Y}^f\}$, which represents edges at different granularity levels. The features of these side outputs are then concatenated to form the final output \hat{Y}^u . The detailed process is elaborated as follows:

$$\begin{aligned} \hat{Y}^c &= H(F^c), \hat{F}^m = W_m^a([F^m, F^c]), \\ \hat{Y}^m &= H(\hat{F}^m), \hat{F}^f = W_f^a([F^f, \hat{F}^m]), \\ \hat{Y}^f &= H(\hat{F}^f), \hat{Y}^u = H(W_u^a([F^c, \hat{F}^m, \hat{F}^f])), \end{aligned} \quad (3)$$




Figure 3: Qualitative comparison results on BSDS test set. * indicates the Zero Shot method.




Figure 4: Precision-recall curves for the BSDS500 test set.

where $W_m^a(\cdot)$, $W_f^a(\cdot)$, $W_u^a(\cdot)$ are convolutions, H is output head shared across varied side outputs, $[\cdot]$ is concatenation.

Generation of edge at arbitrary granularity

Given side outputs at three granularity levels (coarse-grained \hat{Y}^c , medium-grained \hat{Y}^m , fine-grained \hat{Y}^f), we propose to generate edge maps of arbitrary granularity \hat{Y}^α by a simple linear weighting strategy. Specifically, following the work of MuGE (Zhou et al. 2024), we measure the granularity level using $\alpha \in [0, 1]$, where 0 represents the coarsest and 1 represents the finest. Then, the process of generating arbitrary granularity edge graph \hat{Y}^α can be formulated as follows:

$$\hat{Y}^\alpha = \begin{cases} \frac{\alpha}{0.5} \hat{Y}^m + (1 - \frac{\alpha}{0.5}) \hat{Y}^c, & 0 \leq \alpha \leq 0.5 \\ \frac{\alpha-0.5}{0.5} \hat{Y}^f + (1 - \frac{\alpha-0.5}{0.5}) \hat{Y}^m, & 0.5 < \alpha \leq 1 \end{cases} \quad (4)$$

Loss functions for model training

Three loss functions are formulated to train our framework with three side outputs and one final output for detecting edges at multiple granularity: the granularity-aligned loss function for supervising side outputs with pseudo labels, the diversity loss function for enhancing the difference between different side outputs, and the guide loss constraints the produced edges to be compatible with SAM masks.

Pseudo labels at varied granularities. Given the input image X and its manual annotations $\{Y_n\}_{n=1}^N$, where N is

the number of annotations, and $Y_n \in \{0, 1\}^{H \times W}$ is the annotation given by the n -th annotator. We sort all annotations in ascending order based on the number of pixels labeled as edges. Denote the sorted annotations as $\{Y_n^S\}_{n=1}^N$, then the coarse-grained annotation Y^c , medium-grained annotation Y^m , and fine-grained annotation Y^f can be constructed as:

$$Y^c = Y_1^S, Y^m = Y^c \vee Y_{\lfloor \frac{N}{2} \rfloor}^S, Y^f = Y^m \vee Y_N^S, \quad (5)$$

where \vee is element-wise OR operation. The side outputs at different granularity levels $\{\hat{Y}^c, \hat{Y}^m, \hat{Y}^f\}$ are supervised by the corresponding $\{Y^c, Y^m, Y^f\}$.

In order to comprehensively consider all annotations and catch the uncertainty caused by multi-label, inspired by (Zhou et al. 2023), the distribution of final label \hat{Y}^u is obtained by randomly sampling from a multivariate Gaussian distribution $N(\mu_Y, \sigma_Y)$ and then binarizing \hat{Y}^u to the final label Y^u using the threshold ζ , where the mean μ_Y and variance σ_Y are pixel-wise calculated by the label set $\{Y_n\}_{n=1}^N$.

Granularity-aligned loss for side outputs. Edge detection is a binary classification task where each pixel needs to be predicted as an edge (positive) or not (negative), hence binary cross-entropy(BCE) loss is widely used in this task. However, the edge pixels in an image are usually only a tiny fraction of all pixels. Following (Xie and Tu 2015), we adaptively weight each pixel based on the ratio of positive and negative samples in the ground truth. For the side outputs $\{\hat{Y}^c, \hat{Y}^m, \hat{Y}^f\}$, we minimize the following BCE loss:

$$L_{detect}^g = - \sum_{j=1}^{HW} (Y_j^g \xi \log(\hat{Y}_j^g)) + (1 - Y_j^g)(1 - \xi) \log(1 - \hat{Y}_j^g), \quad (6)$$

where ξ is used to balance the contribution of varied pixels, $g \in \{c, m, f, u\}$ represents the granularity level, and j represents the j -th pixel in the prediction and ground truth. In addition, $\xi = |Y_-^g| / (|Y_-^g| + |Y_+^g|)$, $|\cdot|$ represents the total number of pixels, and Y_+^g and Y_-^g represents positive and negative samples in ground truth Y^g .

Given the pseudo labels at varied granularities $\{Y^c, Y^m, Y^f\}$, the loss for supervising the side outputs to align granularity is defined as:

$$L_{side} = \sum_{g \in \{c, m, f\}} L_{detect}^g, \quad (7)$$

	Method	Param. (M)	ODS	OIS	AP
Single-Scale (SS)	Canny (PAMI'86)	-	.611	.676	.520
	OEF (CVPR'15)	-	.746	.770	.815
	DeepContour (CVPR'15)	27.5	.757	.776	.790
	DeepBoundary (ICLR'15)	-	.789	.811	.789
	HED (ICCV'15)	-	.788	.808	.840
	RDS (CVPR'16)	-	.792	.810	.818
	RCF (CVPR'17)	14.8	.798	.815	-
	CED (CVPR'17)	21.8	.803	.820	.871
	BDCN (CVPR'19)	16.3	.806	.826	.847
	DSCD (MM'20)	-	.802	.817	-
	LDC (MM'21)	-	.799	.816	.837
	EDTER (CVPR'22)	468.8	.824	.841	.880
	UAED (CVPR'23)	69.2	.829	.847	.892
	RankED (CVPR'24)	114.4	.824	.840	<u>.895</u>
	DiffusionEdge (AAAI'24)	224.9	.834	.848	.815
	EdgeSAM (TII'24)	7.4	.838	.852	.893
	SAUGE (Ours)	1.3	<u>.839</u>	<u>.860</u>	.893
	SAUGE-L (Ours)	<u>2.1</u>	.847	.868	.898
Single-Scale-VOC (SS-VOC)	DeepBoundary (ICLR'15)	-	.809	.827	.861
	RCF (CVPR'17)	14.8	.806	.823	-
	CED (CVPR'17)	21.8	.812	.833	.889
	BDCN (CVPR'19)	16.3	.820	.838	.888
	DSCD (MM'20)	-	.813	.836	-
	LDC (MM'21)	-	.812	.826	.857
	EDTER (CVPR'22)	468.8	.832	.847	.886
	UAED (CVPR'23)	69.2	.838	.855	.902
	RANKED (CVPR'24)	114.4	.833	.848	<u>.901</u>
	SAUGE (Ours)	1.3	<u>.842</u>	<u>.862</u>	.896
	SAUGE-L (Ours)	<u>2.1</u>	.849	.869	.899

Table 1: Comparisons of approaches with single edge output on BSDS500 test set. The best and second-best results are shown with bold and underlined texts.

Emphasize the difference between side outputs. Different side outputs represent edge maps at different granularity levels, we use mean absolute error (MAE) to emphasize the diversity of side outputs:

$$L_{diff}^{(i,k)} = - \sum_{j=1}^{HW} (|\hat{Y}_j^i - \hat{Y}_j^k|) \cdot (Y_j^i \oplus Y_j^k), \quad (8)$$

where (i, k) is about the combination of $\{c, m, f\}$ with a size of 2, $|\cdot|$ represents taking the absolute value, \oplus represents element-wise XOR operation. In this way, different pixels with edge definitions can be located, and two different side outputs can have differences on these pixels. L_{diff} is obtained by enumerating combinations (i, k) and summing $L_{diff}^{(i,k)}$ as:

$$L_{diff} = L_{diff}^{(c,m)} + L_{diff}^{(c,f)} + L_{diff}^{(m,f)}, \quad (9)$$

The use of L_{diff} promotes the generation of rougher \hat{Y}^c (Figure 6), ensuring the difference between the side outputs.

Guide loss for final output. Here, we further use the object masks outputted by SAM to guide our edge learning. Given the object mask M outputted by SAM, we use the Sobel (Kittler 1983) operator to extract the edges of each mask and denote it as Y^{mask} . Meanwhile, we calculate maps $\bar{Y}^{mask} \in \mathbb{R}^{H \times W}$ which indicate the frequency of the corresponding pixel is recognized as edge on each mask.

	Method	Param. (M)	ODS	OIS	AP
SS	MuGE (M=3) (CVPR'24)	93.9	.845	.854	.876
	SAUGE (M=3) (Ours)	1.3	<u>.854</u>	.865	<u>.894</u>
	MuGE (M=11) (CVPR'24)	93.9	.850	.856	.882
	SAUGE (M=11) (Ours)	1.3	.857	.865	.896
SS-VOC	MuGE (M=3) (CVPR'24)	93.9	.852	.859	.889
	SAUGE (M=3) (Ours)	1.3	<u>.857</u>	<u>.867</u>	<u>.900</u>
	MuGE (M=11) (CVPR'24)	93.9	.855	.860	.894
	SAUGE (M=11) (Ours)	1.3	.859	.868	.901

Table 2: Comparison of multi-granularity approaches on the BSDS500 test set. The best and second-best results are shown with bold and underlined texts, respectively.

The loss for supervising the final output \hat{Y}^u is defined as:

$$L_{guide} = \sum_{j=1}^{HW} e^{(\psi_j + \omega_j)} L_{detect}^u \quad (10)$$

where $\psi_j = -(Y_j^{mask}) \cdot (Y_j^{mask} \oplus Y_j) \cdot \bar{Y}_j^{mask}$, $\omega_j = \cos(\tilde{Y}_j^u)$. The weight map ψ is defined in the way such that the easily confused pixels indicated by the object mask contributes less in the edge learning. ω is constructed in order to assign higher weights to edge pixels with lower confidence.

Overall loss function. The final optimization objective is defined as the weighted sum of L_{side} , L_{diff} , and L_{guide} :

$$L_{total} = L_{guide} + \lambda L_{diff} + \beta L_{side} \quad (11)$$

where λ and β are the coefficient to control the balance of different losses. In this work, we fix $\lambda = 0.1, \beta = 0.5$.

Experiments

Datasets

We conduct experiments on three widely-used edge detection datasets: BSDS500 (Arbelaez et al. 2010), Multicue (Mély et al. 2016) and NYUDv2 (Silberman et al. 2012). For data augmentation, we adopt the same strategy as UAED (Zhou et al. 2023) across both datasets.

BSDS500 consists of 500 natural images, with 200 for training, 100 for validation, and the remaining for test. Each image has 4 to 9 manual annotations. Additionally, the PAS-CAL VOC set (Everingham et al. 2010) with 10,103 images is used as supplementary training data, with edge annotations derived from semantic masks using Laplacian detector.

Multicue includes 100 images from complex natural scenes, each of which is annotated by multiple individuals. We randomly split these images into training and evaluation sets, with 80 images for training and 20 for testing. This process is repeated three times and average scores are reported.

NYUDv2 is a dataset for indoor scene parsing and edge detection, containing 1,449 paired RGB-D images. Each image has a single ground-truth edge map, with the dataset split into 381 training, 414 validation, and 654 testing images.

Implementation Details

We implement SAUGE based on PyTorch (Paszke et al. 2019), and use SAM pre-trained on the SA-1B dataset (Kir-

Methods	ODS	OIS	AP
Human (VR'16)	0.750	-	-
Multicue (VR'16)	0.830	-	-
HED (ICCV'15)	0.851	0.864	-
RCF (CVPR'17)	0.857	0.862	-
BDCN (CVPR'19)	0.891	0.898	0.835
DSCD (MM'20)	0.871	0.876	-
LDC (MM'21)	0.881	0.893	-
EDTER (CVPR'22)	0.894	0.900	0.944
UAED (CVPR'23)	0.895	0.902	<u>0.949</u>
MuGE (CVPR'24)	0.898	0.900	0.950
DifussionEdge (AAAI'24)	<u>0.904</u>	0.909	-
SAUGE (Ours)	0.905	<u>0.907</u>	0.939

Table 3: Comparisons on Multicue. Best and second-best results are shown with bold and underlined texts, respectively.

illov et al. 2023) as our backbone. The Adam optimizer (Kinga, Adam et al. 2015) is used to update all parameters. The learning rate is initialized as $1e-4$ with step scheduling and weight decay is set to $5e-4$. For BSDS, we set the ζ for the thresholding label to 0.2, and the model was trained for 6 epochs with a batch size of 3. For Multicue, we randomly crop the images into 512×512 and set ζ to 0.3. The model is trained for 20 epochs using a batch size of 3. All experiments were conducted on RTX 3090, where training the model on BSDS500 requires approximately 20 GPU hours and 16GB GPU memory.

Evaluation Metric

We evaluate the performance using widely adopted metrics, including F-scores for both Optimal Dataset Scale (ODS) and Optimal Image Scale (OIS), as well as Average Precision (AP). ODS applies a fixed threshold to binarize the predicted edge map across the entire dataset, while OIS selects an optimal threshold for each image. We also report the learnable parameter. Following previous works (Liu et al. 2017; Pu et al. 2022; Zhou et al. 2023), we apply non-maximum suppression (NMS) to the predictions before evaluation. The localization tolerance is set to 0.011 for NYUDv2 and 0.0075 for other datasets, defining the maximum allowed distance for matching predictions with GT. For AP evaluation, we select the candidate with the highest F-value across M candidates at each threshold, providing a more comprehensive assessment of the overall performance across all candidates. To assess the quality of multiple granularity edge maps ($M=3, 11$, etc), we follow the MuGE (Zhou et al. 2024) that uses the best-matching strategy, selecting the candidate closest to the GT from multiple candidates.

Comparison with State-of-the-arts

Existing methods can be divided into two categories: single definite output and multi-granularity output. For single definite output, traditional methods include Canny (Canny 1986) and OEF (Hallman and Fowlkes 2015). CNN-based methods, such as DeepContour (Shen et al. 2015), DeepBoundary (Kokkinos 2015), HED (Xie and Tu 2015), RDS (Liu and Lew 2016), CED (Wang, Zhao, and Huang 2017), BDCN (He et al. 2019), DSCD (Deng and Liu 2020), and

Methods	ODS	OIS	AP
OEF (CVPR'15)	0.651	0.667	-
HED (ICCV'16)	0.720	0.734	0.734
RCF (CVPR'17)	0.729	0.742	-
BDCN (CVPR'19)	0.748	0.763	0.770
EDTER (CVPR'22)	0.774	0.789	0.797
SAM* (ICCV'23)	0.699	0.719	0.707
DiffusionEdge (AAAI'24)	0.761	0.766	-
EdgeSAM (TII'24)	<u>0.783</u>	<u>0.797</u>	0.805
RankED (CVPR'24)	0.780	0.793	0.826
SAUGE-L* (Ours)	0.794	0.803	<u>0.813</u>

Table 4: Comparisons on the NYUDv2. The best and second-best results are shown with bold and underlined texts respectively. * indicates the Zero Shot method.

LDC (Deng et al. 2021), are widely adopted. Transformer-based methods involve EDTER (Pu et al. 2022), RankED (Cetinkaya, Kalkan, and Akbas 2024) and EdgeSAM (Yang et al. 2024). The diffusion-based method involves DiffusionEdge (Ye et al. 2024). Currently, only MuGE (Zhou et al. 2024) is developed for multi-granularity output.

Results on BSDS500. Since SAUGE can generate both definite output \hat{Y}^u and multi-granularity outputs \hat{Y}^α , we compare it with excellent representatives of these two types of methods respectively. The quantitative results are presented in Table 1 and Table 2, where $M=3$ represents $\alpha=\{0, 0.5, 1\}$, and $M=11$ represents the 11 outputs generated from 0 to 1 at intervals of 0.1. Experiments were conducted under two settings: Single-Scale input (SS) and Single-Scale input with additional PASCAL VOC data for training (SS-VOC).

In Table 1, we report the performance of existing state-of-the-art methods for definite output. As shown, our SAUGE achieves the best performances in both ODS and OIS metrics under all experimental settings (SS and SS-VOC), which significantly outperforms EdgeSAM with the same SAM backbone. It is worth noting that our model only needs to tune a very small number of parameters—82% fewer compared to EdgeSAM. When using a larger SAM (ViT-L SAM) as the backbone (denoted as SAUGE-L), the performance improves by 0.009 in ODS and 0.016 in OIS under the SS setting, and by 0.011 in ODS and 0.014 in OIS under the SS-VOC setting. Qualitatively, Figure 3 compares our SAUGE with other methods, demonstrating that SAUGE can produce more detailed predictions. We note that the background edges missed by SAM can be successfully obtained by our model. The Precision-Recall curve is shown in Figure 4.

For multi-granularity output, Table 2 shows that SAUGE significantly outperforms MuGE in all settings. It can be found that increasing the number of candidates (M) leads to higher performance, demonstrating the advantages of multi-granularity edge modeling in addressing uncertainty.

Results on Multicue. We also conduct experiments on Multicue edges, and the results are illustrated in Table ???. As can be seen, the proposed SAUGE outperforms the previous state-of-the-art in the ODS metric and achieves second-place performance in the OIS metric, with scores of 90.5% in ODS and 90.7% in the OIS metric. The relatively lower performance in the AP metric may be attributed to overfit-




Figure 5: Qualitative comparison with MuGE for different edge granularities α under the SS-VOC setting.




Figure 6: Comparison of coarse-grained and fine-grained side output $\{\hat{Y}^c, \hat{Y}^f\}$ with/without L_{diff} . The ODS is obtained on BSDS500 test set under Single-Scale (SS) setting.

ting, likely due to the limited size of the Multicue dataset.

Results on NYUDv2. To verify the generalization of our SAUGE, we directly use the SAUGE-L model trained on BSDS500 and PASCAL VOC for zero-shot edge detection on this set. The results in Table 4 indicate that our method outperforms previous supervised methods on both ODS and OIS without parameter tuning, which demonstrates the effectiveness of our approach.

Ablation Study and Further Analysis

Effect of key components. The key designs of our SAUGE lie in the lightweight Side Transfer Network (STN), the constraints on the multi-granularity side outputs, and the mask-guided loss L_{guide} . To verify the effectiveness of these designs, we begin with SAM as the baseline, and then sequentially add the Side Transfer Network, the side output constraints (denoted as SOC), and the mask-guided loss. The results on the BSDS500 dataset are presented in Table 5. As shown, for both pre-trained ViT-B SAM and ViT-L SAM baseline, the inclusion of our lightweight STN module leads to a significant improvement across all metrics, by more than 9% in ODS, 8.7% in OIS, 10.8% in AP metrics. Also, with the constraints on side output (SOC), our framework can generate edges at arbitrary granularity and achieve a notable enhancement in the AP metric. We attribute this to the fact that SOC explicitly facilitates the model in capturing different granularity information, thus enabling better alignment with the uncertainty in edge decisions. Additionally, mask-guided loss L_{guide} also contributes positively to all metrics. To study the impact of L_{diff} , Figure 6 visualizes model outputs with and without it. As shown, the L_{diff} amplifies the differences between outputs at various granularities, further enhancing alignment with edge decision uncertainty.

Method	SOC	\mathcal{L}_{guide}	ODS	OIS	AP	MG
SAM*			.768	.786	.794	×
Ours	✓		.837	.855	.880	×
	✓	✓	.836	.857	.891	✓
Ours-L			.839	.860	.893	✓
	✓		.844	.865	.822	×
	✓		.846	.866	.895	✓
		✓	.847	.868	.898	✓

Table 5: Ablation study on key components of our model on the BSDS500 set. MG refers to whether the method can output multi-granularity edges. * indicates Zero Shot method.

Performance	$\alpha = 0$	$\alpha = 0.2$	$\alpha = 0.4$	$\alpha = 0.5$	$\alpha = 1$
ODS	0.832	0.836	0.839	0.839	0.836
OIS	0.851	0.857	0.859	0.859	0.858
AP	0.771	0.814	0.882	0.890	0.891

Table 6: Comparisons on BSDS500 for varied granularities.

Performance of specific granularity outputs. Given a granularity level α , our method can generate edge maps at arbitrary granularity in a controlled manner. In Figure 5, we visualize the edge maps at different granularities generated by MuGE and our SAUGE. For all α , SAUGE consistently outperforms MuGE and produces edge maps with high quality. Furthermore, we report the quantitative results of various granularity $\alpha = \{0, 0.2, 0.4, 0.5, 1\}$ in Table 6. As shown, the edge maps at varied granularities exhibit both good performance and diversity in ODS, OIS and AP metrics.

Computational efficiency. We evaluate the inference efficiency of our proposed SAUGE on a single RTX 3090 GPU. SAUGE consumes 6.5GB of GPU memory and achieves a speed of 5.3 FPS, exhibiting similar efficiency to the transformer-based EDTER (6.0GB, 5.3 FPS) while outperforming the zero-shot-based SAM (8.1GB, 0.9 FPS) and the diffusion-based DiffusionEdge (3.4GB, 0.4 FPS) in speed.

Conclusion and Limitation

In this paper, we designed a novel uncertainty-aligned edge detector called SAUGE. In our SAUGE, we developed a lightweight Side Transfer Network (STN) to explicitly explore the knowledge embedded in SAM for multiple granularity edge detection. Extensive experiments on three edge sets are reported to demonstrate the superiority of SAUGE.

Limitation. Our SAUGE is constructed based on the SAM backbone model. In the future, we would like to extend our framework for a more efficient backbone model.

Acknowledgments

This work was supported partially by the NSFC (U21A20-471, U22A2095, 62076260, 61772570), Guangdong Natural Science Funds Project (2023B1515040025), Guangdong NSF for Distinguished Young Scholar (2022B15-15020009), Guangdong Provincial Key Laboratory of Information Security Technology (2023B1212060026), open research fund of Key Laboratory of Machine Intelligence and System Control, Ministry of Education (No. MISC-202407).

References

- Arbelaez, P.; Maire, M.; Fowlkes, C.; and Malik, J. 2010. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5): 898–916.
- Canny, J. 1986. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6): 679–698.
- Cetinkaya, B.; Kalkan, S.; and Akbas, E. 2024. RankED: Addressing Imbalance and Uncertainty in Edge Detection Using Ranking-based Losses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3239–3249.
- Chen, W.-T.; Vong, Y.-J.; Kuo, S.-Y.; Ma, S.; and Wang, J. 2024. RobustSAM: Segment Anything Robustly on Degraded Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4081–4091.
- Deng, R.; and Liu, S. 2020. Deep structural contour detection. In *Proceedings of the 28th ACM international conference on multimedia*, 304–312.
- Deng, R.; Liu, S.; Wang, J.; Wang, H.; Zhao, H.; and Zhang, X. 2021. Learning to decode contextual information for efficient contour detection. In *Proceedings of the 29th ACM International Conference on Multimedia*, 4435–4443.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88: 303–338.
- Gu, P.; Zhao, Z.; Wang, H.; Peng, Y.; Zhang, Y.; Sapkota, N.; Wang, C.; and Chen, D. Z. 2024. Boosting Medical Image Classification with Segmentation Foundation Model. *arXiv preprint arXiv:2406.11026*.
- Hallman, S.; and Fowlkes, C. C. 2015. Oriented edge forests for boundary detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1732–1740.
- He, J.; Zhang, S.; Yang, M.; Shan, Y.; and Huang, T. 2019. Bi-directional cascade network for perceptual edge detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3828–3837.
- Jiang, C.; Gao, F.; Ma, B.; Lin, Y.; Wang, N.; and Xu, G. 2023. Masked and adaptive transformer for exemplar based image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22418–22427.
- Ke, L.; Ye, M.; Danelljan, M.; Tai, Y.-W.; Tang, C.-K.; Yu, F.; et al. 2024. Segment anything in high quality. *Advances in Neural Information Processing Systems*, 36.
- Kinga, D.; Adam, J. B.; et al. 2015. A method for stochastic optimization. In *International conference on learning representations (ICLR)*, volume 5, 6. San Diego, California;.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Kittler, J. 1983. On the accuracy of the Sobel edge detector. *Image and Vision Computing*, 1(1): 37–42.
- Kokkinos, I. 2015. Pushing the boundaries of boundary detection using deep learning. *arXiv preprint arXiv:1511.07386*.
- Liu, J.-J.; Hou, Q.; and Cheng, M.-M. 2020. Dynamic feature integration for simultaneous detection of salient object, edge, and skeleton. *IEEE Transactions on Image Processing*, 29: 8652–8667.
- Liu, Y.; Cheng, M.-M.; Hu, X.; Wang, K.; and Bai, X. 2017. Richer convolutional features for edge detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3000–3009.
- Liu, Y.; and Lew, M. S. 2016. Learning relaxed deep supervision for better edge detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 231–240.
- Ma, J.; He, Y.; Li, F.; Han, L.; You, C.; and Wang, B. 2024. Segment anything in medical images. *Nature Communications*, 15(1): 654.
- Mély, D. A.; Kim, J.; McGill, M.; Guo, Y.; and Serre, T. 2016. A systematic comparison between visual cues for boundary detection. *Vision research*, 120: 93–107.
- Nazeri, K.; Ng, E.; Joseph, T.; Qureshi, F.; and Ebrahimi, M. 2019. Edgeconnect: Structure guided image inpainting using edge prediction. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, 0–0.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Pu, M.; Huang, Y.; Liu, Y.; Guan, Q.; and Ling, H. 2022. Edter: Edge detection with transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1402–1412.
- Qin, X.; Zhang, Z.; Huang, C.; Gao, C.; Dehghan, M.; and Jagersand, M. 2019. Basnet: Boundary-aware salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7479–7489.
- Shen, W.; Wang, X.; Wang, Y.; Bai, X.; and Zhang, Z. 2015. Deepcontour: A deep convolutional feature learned by positive-sharing loss for contour detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3982–3991.

- Silberman, N.; Hoiem, D.; Kohli, P.; and Fergus, R. 2012. Indoor segmentation and support inference from rgbd images. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part V 12*, 746–760. Springer.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Su, Z.; Liu, W.; Yu, Z.; Hu, D.; Liao, Q.; Tian, Q.; Pietikäinen, M.; and Liu, L. 2021. Pixel difference networks for efficient edge detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5117–5127.
- Wang, Y.; Zhao, X.; and Huang, K. 2017. Deep crisp boundaries. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3892–3900.
- Wu, Y.; Pan, C.; Wang, G.; Yang, Y.; Wei, J.; Li, C.; and Shen, H. T. 2023. Learning semantic-aware knowledge guidance for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1662–1671.
- Xie, S.; and Tu, Z. 2015. Holistically-nested edge detection. In *Proceedings of the IEEE international conference on computer vision*, 1395–1403.
- Xu, X.; Wang, R.; and Lu, J. 2023. Low-light image enhancement via structure modeling and guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9893–9903.
- Yang, W.; Chen, X.-D.; Wu, W.; Qin, H.; Yan, K.; Mao, X.; and Song, H. 2024. Boosting Deep Unsupervised Edge Detection via Segment Anything Model. *IEEE Transactions on Industrial Informatics*.
- Yao, Z.; and Wang, L. 2023. Object localization and edge refinement network for salient object detection. *Expert Systems with Applications*, 213: 118973.
- Ye, Y.; Xu, K.; Huang, Y.; Yi, R.; and Cai, Z. 2024. DiffusionEdge: Diffusion Probabilistic Model for Crisp Edge Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6675–6683.
- Yu, T.; Feng, R.; Feng, R.; Liu, J.; Jin, X.; Zeng, W.; and Chen, Z. 2023. Inpaint anything: Segment anything meets image inpainting. *arXiv preprint arXiv:2304.06790*.
- Yu, Z.; Huang, R.; Byeon, W.; Liu, S.; Liu, G.; Breuel, T.; Anandkumar, A.; and Kautz, J. 2021. Coupled segmentation and edge learning via dynamic graph propagation. *Advances in Neural Information Processing Systems*, 34: 4919–4932.
- Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; and Yang, M.-H. 2022. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5728–5739.
- Zhang, M.; Wang, Y.; Guo, J.; Li, Y.; Gao, X.; and Zhang, J. 2024. IRSAM: Advancing Segment Anything Model for Infrared Small Target Detection. *arXiv preprint arXiv:2407.07520*.
- Zhou, C.; Huang, Y.; Pu, M.; Guan, Q.; Deng, R.; and Ling, H. 2024. MuGE: Multiple Granularity Edge Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25952–25962.
- Zhou, C.; Huang, Y.; Pu, M.; Guan, Q.; Huang, L.; and Ling, H. 2023. The treasure beneath multiple annotations: An uncertainty-aware edge detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15507–15517.