

STENet: Superpixel Token Enhancing Network for RGB-D Salient Object Detection

Jianlin Chen, Gongyang Li, *Member, IEEE*, Zhiqiang Zhang, Liang Chang, and Dan Zeng, *Senior Member, IEEE*

Abstract—Transformer-based methods for RGB-D Salient Object Detection (SOD) have gained significant interest, owing to the transformer’s exceptional capacity to capture long-range pixel dependencies. Nevertheless, current RGB-D SOD methods face challenges, such as the quadratic complexity of the attention mechanism and the limited local detail extraction. To overcome these limitations, we propose a novel Superpixel Token Enhancing Network (STENet), which introduces superpixels into cross-modal interaction. STENet follows the two-stream encoder-decoder structure. Its cores are two tailored superpixel-driven cross-modal interaction modules, responsible for global and local feature enhancement. Specifically, we update the superpixel generation method by expanding the neighborhood range of each superpixel, allowing for flexible transformation between pixels and superpixels. With the updated superpixel generation method, we first propose the Superpixel Attention Global Enhancing Module to model the global pixel-to-superpixel relationship rather than the traditional global pixel-to-pixel relationship, which can capture region-level information and reduce computational complexity. We also propose the Superpixel Attention Local Refining Module, which leverages pixel similarity within superpixels to filter out a subset of pixels (*i.e.*, local pixels) and then performs feature enhancement on these local pixels, thereby capturing concerned local details. Furthermore, we fuse the globally and locally enhanced features along with the cross-scale features to achieve comprehensive feature representation. Experiments on seven RGB-D SOD datasets reveal that our STENet achieves competitive performance compared to state-of-the-art methods.

Index Terms—RGB-D salient object detection, superpixel generation, pixel-to-superpixel relationship, global and local feature enhancement.

I. INTRODUCTION

SALIENT Object Detection (SOD) stands as a fundamental and crucial task in computer vision. Its primary objective is to identify the most visually prominent or important targets within images, providing key prerequisite information for numerous advanced visual processing tasks. SOD plays a vital

Jianlin Chen and Zhiqiang Zhang are with School of Communication and Information Engineering, Shanghai University, Shanghai 200444, China (email: chen1026@shu.edu.cn; zjzhang@staff.shu.edu.cn).

Dan Zeng is with School of Communication and Information Engineering, Shanghai University, Shanghai 200444, China, and the Institute for Urban Governance, Shanghai University, Shanghai 200444, China (email: dzeng@shu.edu.cn).

Gongyang Li is with the School of Communication and Information Engineering, Shanghai University, Shanghai 200444, China, and Yunnan Key Laboratory of Service Computing, Yunnan University of Finance and Economics, Kunming 650000, China (email: ligongyang@shu.edu.cn).

Liang Chang is with Innovation Academy for Microsatellites of Chinese Academy of Science, Shanghai 201203, China (email: changl@microsate.com).

Corresponding authors: Gongyang Li and Zhiqiang Zhang.

role in industrial defect detection [1], target tracking [2], remote sensing scene understanding [3], [4], and medical image analysis [5]. However, traditional RGB SOD methods [6]–[8] often encounter challenges in complex or low-contrast scenes. To well handle these scenes, researchers have explored the integration of additional information, particularly depth data, and focused on RGB-D SOD. The combination of depth information has exhibited impressive potential in bolstering the stability and accuracy of RGB-D SOD models in various scenes [9]–[12].

Contemporary RGB-D SOD approaches predominantly employ Convolutional Neural Networks (CNNs) [13], with research efforts concentrating on devising multi-modal multi-scale fusion mechanisms [14]–[16], that effectively harness complementary relationships between heterogeneous visual features. Zhang *et al.* [17] designed a bilateral attention network to learn features from both the foreground and background separately, in order to refine the prediction results. However, CNNs exhibit inherent limitations due to their local connectivity, which restricts them to extracting only limited local detail information and hinders their ability to capture global contextual information. Their reliance on fixed parameters impedes adaptability across diverse scenarios.

The recent emergence of transformers [18] has revolutionized long-range dependency modeling through synergistic integration of self-attention [19] and cross-attention [20] operations, effectively addressing the inherent locality constraints of conventional CNN-based frameworks. In RGB-D SOD, transformers are typically deployed in dual capacities: serving as backbone networks for hierarchical feature extraction, and functioning as cross-modal integration modules for heterogeneous feature fusion. Wu *et al.* [21] effectively captured long-range cross-modal dependencies through parallelized attention mechanisms in the cross-modal interactive parallel transformer module, generating more discriminative fused feature representations. However, this fusion method is typically implemented only within the channel dimension due to the quadratic computational complexity issue inherent in transformer attention architectures. In order to address this issue, Pang *et al.* [22] devised a parameter-free operation to implement spatial attention within the transformer, significantly reducing computation and memory requirements. However, this fusion method only focuses on the global perspective and pays insufficient attention to local details, making it difficult to accurately capture changes in local details in certain specific scenarios.

Inspired by the above observations, we propose the Superpixel Token Enhancing Network (STENet) based on the

two-stream encoder-decoder architecture, which introduces the efficient superpixels to RGB-D SOD. Superpixels [23], [24] aggregate pixels into geometrically coherent units to reduce computational redundancy while preserving object-level structural integrity, which is widely adopted in computer vision. Here, based on superpixels, we explore global enhancement and local refinement for RGB-D SOD by leveraging superpixels as coherent structural units. In particular, our STENet contains two specialized superpixel-driven modules for global enhancement and local refinement, named the Superpixel Attention Global Enhancing Module (SAGEM) and the Superpixel Attention Local Refining Module (SALRM), respectively. SAGEM is responsible for modeling global pixel-to-superpixel relationships instead of traditional pixel-to-pixel interactions, capturing region-level relationships while reducing computational complexity. SALRM is responsible for identifying and enhancing similar pixels within superpixels, with a focus on fine-grained local details. The globally and locally enhanced features generated by SAGEM and SALRM are further fused with the cross-scale features to achieve powerful feature representation. In this way, the proposed STENet can generate high-quality saliency maps while having efficient computation.

Our main contributions are summarized as follows:

- We propose a novel Superpixel Token Enhancing Network, which integrates the classical pixel clustering technique, *i.e.*, superpixels, into RGB-D SOD to preserve locally coherent structural information while maintaining efficient parallel computing capabilities.
- We update the superpixel generation method, which expands the neighborhood range of each superpixel, allowing for flexible transformation between pixels and superpixels.
- We propose the Superpixel Attention Global Enhancing Module to employ a cross-modal superpixel feature for cross-attention mechanisms, resulting in enhanced global feature representation.
- We propose the Superpixel Attention Local Refining Module to utilize internal statistical information from cross-modal superpixels to implement a self-attention mechanism for selected pixels within the pixel features, achieving enhanced local feature representation.

The rest of this paper is organized as follows. In Sec. II, we review the related work of RGB-D SOD and superpixel representation. In Sec. III, we elaborate the proposed STENet. In Sec. IV, we present comprehensive experiments and analyses. Finally, in Sec. V, we conclude this work.

II. RELATED WORK

A. RGB-D Salient Object Detection

RGB SOD continues to face challenges such as complex backgrounds, low contrast, and varying lighting conditions, which restrict detection accuracy. To mitigate these issues, researchers have incorporated depth information into SOD and focused on RGB-D SOD. Depth information effectively enhances the ability to distinguish between foreground and background, demonstrating significant application value in

complex environments. In the advancing field of RGB-D SOD, CNN-based methods have achieved remarkable performance due to deep learning advances. The existing methods can be categorized into three fusion strategies: early fusion [25]–[27], intermediate fusion [10], [17], [28]–[33], and late fusion [34]–[37].

Intermediate fusion methods are currently the mainstream scheme, and focus on cross-modal and cross-level interaction and fusion in RGB-D SOD. For example, Li *et al.* [10] proposed an alternate interaction approach to gradually extract and fuse cross-modal features. Zhang *et al.* [17] designed a bilateral attention network to learn features from both the foreground and background separately, in order to refine the prediction results. Zhang *et al.* [38] employed a decoupled dynamic filtering convolutional architecture to achieve cross-modal dynamic feature fusion. Sun *et al.* [29] presented a dual-branch multi-modal fusion network, aimed at addressing the challenges of multi-modal fusion and multi-level aggregation in RGB-D SOD. Wu *et al.* [32] enhanced the discrimination ability of RGB and depth features through a granularity-based attention mechanism and introduced a unified cross dual-attention module to achieve multi-modal and multi-level fusion. However, due to the limitations of convolution, these methods have limited capabilities in modeling global correlations.

Concurrently, the transformer [18], [39], renowned for its robust capability in modeling long-range dependencies, has garnered increasing attention within the computer vision fields. In RGB-D SOD, the transformer is frequently employed as backbone network for feature extraction because of its ability to capture rich global context information. Liu *et al.* [40] utilized the swin transformer [39] to extract features and integrated them with edge information for cross-modal fusion. Similarly, Liu *et al.* [41] enhanced multi-level features using three transformer encoders with shared weights. Hu *et al.* [41] employed a two-stream swin transformer encoder to extract multi-level and multi-scale features from RGB images and depth images, in order to model global information.

Furthermore, the transformer has been employed for feature interaction and enhancement. Tang *et al.* [42] used inter-feature and intra-feature interactive transformers to aggregate output features of all different resolutions. Cong *et al.* [43] proposed a cross-modal point-wise interaction module based on transformer attention triggering, exploring feature interaction between different modalities under positional constraints. Pang *et al.* [22] proposed a cross-modal perspective-mixed Transformer approach, achieving a balanced trade-off between computational efficiency and accuracy. However, the transformer-based model also suffers from weaknesses in capturing local information and high computational costs, restricting its application.

B. Superpixel Representation

In the field of computer vision, superpixel representation has long been a focal point of research due to its efficiency as an image representation method. Traditional superpixel representation techniques mainly rely on unsupervised clustering

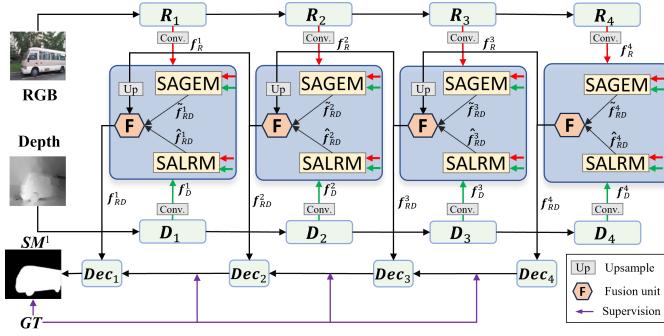


Fig. 1. Overview of the proposed STENet.

algorithms, such as graph-based approaches [23] [44], mean shift [45], or K-means clustering [46] [47]. These methods generate locally consistent pixel blocks (superpixels), which effectively preserve critical structural information within images and facilitate subsequent processing.

With the emergence of deep learning, the integration of superpixel segmentation with deep learning framework has emerged as a new trend in computer vision. Varun *et al.* [24] developed a new differentiable model named SSN for superpixel sampling, utilizing deep networks to learn superpixel segmentation. Zhang *et al.* [48] combined SLIC with a vision transformer to convert images into superpixel tokens for region-level semantic learning. Huang *et al.* [49] employed sparse connections and local sampling strategies to generate super tokens, effectively reducing computational costs. Zhu *et al.* [50] decomposed the pixel space into a low-dimensional superpixel space through local cross-attention for category prediction. Mei *et al.* [51] utilized superpixels to perform irregular yet semantically coherent partitioning of images, effectively capturing intricate image details. Zhang *et al.* [52] further integrated superpixels with attention mechanisms in super-resolution, effectively blending global and local attention mechanisms.

To address computational redundancy and cross-modal interaction challenges in RGB-D SOD, we introduce a superpixel representation method that compresses RGB-D features into hierarchical superpixel structural units. Our method achieves efficient cross-modal fusion through dual-phase attention interactions. Specifically, cross-superpixel attention mechanisms model global contextual relationships, while intensity-geometry guided refinement operations enhance local details within adaptively partitioned superpixel regions.

III. PROPOSED METHOD

In this section, we present the proposed STENet in detail. In Sec. III-A, we introduce the network overview. In Sec. III-B, Sec. III-C, and Sec. III-D, we elaborate the superpixel generation, SAGEM, and SALRM, respectively. In Sec. III-E, we present the loss function.

A. Network Overview

As depicted in Fig. 1, our STENet embraces the classic encoder-decoder architecture. The encoder encompasses RGB

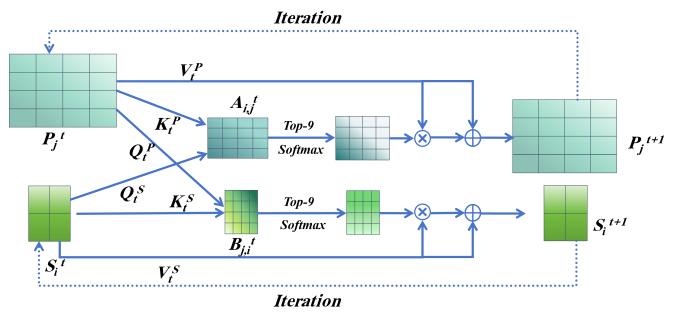


Fig. 2. The process of the superpixel generation.

and depth branches. Both branches adopt the Swin-B Transformer [39] for feature extraction with the input sizes of $3 \times 384 \times 384$ and $1 \times 384 \times 384$, respectively. Each branch consists of four blocks denoted as R_i/D_i ($i = 1, 2, 3, 4$), yielding output cross-modal features f_R^i/f_D^i . Subsequently, we simultaneously use SAGEM and SALRM to perform global enhancement and local refinement on cross-modal features, generating \tilde{f}_{RD}^i and \hat{f}_{RD}^i , respectively. Then, we adopt the fusion unit to integrate the global feature \tilde{f}_{RD}^i , the local feature \hat{f}_{RD}^i , and cross-scale features. In the fusion unit, we first concatenate \tilde{f}_{RD}^i and \hat{f}_{RD}^i , then add them to cross-scale features (except for the highest-level one), and finally perform the self-attention operation to capture long-range pixel dependencies. Finally, our STENet produces the final saliency map SM^1 in the decoder. The decoder consists of four blocks, named Dec_i . Each block includes a 5×5 depthwise convolution layer, a layer normalization layer, and two 1×1 pointwise convolution layers. A residual connection adds the output features to the input features, and an upsampling operation adjusts the final sizes. For comprehensive and efficient supervision, we implement a multi-scale supervision strategy at each block and the final saliency map SM^1 .

B. Superpixel Generation

Distinct from previous methods, we introduce superpixel into RGB-D SOD. By clustering similar pixels, superpixels preserve crucial local edge information while reducing computational complexity, striking a balance between accuracy and performance. To enable more flexible and dynamic adjustments of relationships between pixels and superpixels, we adopt the local cross-attention [50] to generate superpixel tokens. Notably, we expand the neighborhood range of each superpixel. As shown in Fig. 2, we first initialize the superpixel features $S \in \mathbb{R}^{M \times d}$ as the average of pixel features $P \in \mathbb{R}^{N \times d}$, where M and N represents the number of superpixels and pixels respectively and d is the feature dimension. Then, we obtain the corresponding Query, Key, and Value from both superpixel and pixel features, respectively. And we compute the local cross-attention weights $A_{i,j}^t$ and $B_{j,i}^t$ between each superpixel and its neighboring pixels, update the superpixel features based on these weights, and repeat this process for T iteration.

Under normal circumstances, to maintain the locality of superpixels, a masking operation is performed on the acquired

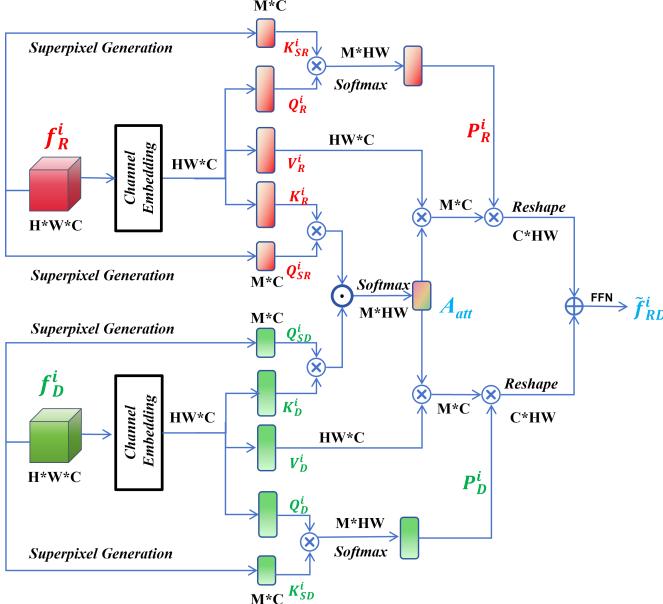


Fig. 3. Illustration of the Superpixel Attention Global Enhancing Module.

weights, ensuring that each pixel is exclusively correlated with the 9 adjacent superpixels. However, in shallow layers, high-resolution features lead to dense superpixel grids (*e.g.*, 96×96 pixels), where 3×3 neighborhoods may miss semantically relevant superpixels due to strict spatial constraints. Expanding to 5×5 allows pixels to search for similar superpixels over a wider spatial range, improving the quality of selected features.

For each pixel i , we choose the top-9 most similar superpixels from its correlation map $A_{i,j}^t$ and set the corresponding elements in $\Omega_{9,i,j}^t$ to 1, with the rest set to 0. Similarly, for each superpixel j , we select the top $9 \times (M/N)$ most similar pixels from its correlation map $B_{j,i}^t$ and set $\Omega_{9,j,i}^t$ to 1, with the others set to 0. By maintaining the top-9 selection (same as 3×3 neighborhoods), the computational complexity remains constant. The mask operation still filters only the most relevant superpixels, avoiding redundant calculations while leveraging the expanded spatial context. Using these masks, we update the features as follows:

$$\begin{aligned} S_i^{t+1} &= \text{softmax} \left(\sum_j \Omega_{9,i,j}^t A_{i,j}^t \right) \odot V_{t,i}^S + S_i^t, \\ P_j^{t+1} &= \text{softmax} \left(\sum_i \Omega_{9,j,i}^t B_{j,i}^t \right) \odot V_{t,j}^P + P_j^t, \end{aligned} \quad (1)$$

where \odot denotes the element-wise multiplication, and S_i^t and P_j^t represent the superpixel and pixel features from the previous iteration, respectively. This approach combines the retention of essential information with residual learning. Through T iterations of this process, we ultimately get the refined superpixel features S_i^T . In practical implementation, T is set to 2. This is because excessive iterations offer no benefit to performance but significantly increase computational costs.

C. Superpixel Attention Global Enhancing Module

The use of transformers for cross-modal feature interaction and enhancement of global information has become widespread. However, in most cases, only the channel dimension is considered due to the heavy computational workload in the spatial dimension. Therefore, we propose a Superpixel Attention Global Enhancement Module (SAGEM), which establishes a relationship between superpixels and pixels in the spatial dimension, thereby facilitating cross-modal feature interaction and enhancing the dissemination of global information.

As depicted in Fig. 3, the inputs of SAGEM are f_R^i and f_D^i . We flatten their sizes from $H \times W \times C$ to $HW \times C$ and perform linear embedding to acquire the respective query (Q_R^i/Q_D^i), key (K_R^i/K_D^i), and value (V_R^i/V_D^i). We begin by applying the superpixel generation method, as detailed in Sec. III-B, to obtain superpixels Q_{SR}^i and Q_{SD}^i for each modality. These superpixels are subsequently flattened to a dimension of $M \times C$, where M represents the quantity of superpixels and C denotes the number of channels.

Then, we establish dependencies between pixels and superpixels via the local cross-attention mechanism, and then utilize superpixels for global information propagation. We compute the global attention maps A_R^i and A_D^i through matrix multiplication between the queries of superpixels and the keys of pixels, yielding attention maps with sizes of $M \times HW$ as follows:

$$\begin{aligned} A_R^i &= \text{softmax} \left(\frac{Q_{SR}^i \otimes K_R^i{}^\top}{\sqrt{D_k}} \right), \\ A_D^i &= \text{softmax} \left(\frac{Q_{SD}^i \otimes K_D^i{}^\top}{\sqrt{D_k}} \right), \end{aligned} \quad (2)$$

where \otimes represents the matrix multiplication. This process facilitates the derivation of attention maps with sizes of $M \times HW$, where M represents the number of superpixels and HW corresponds to the flattened pixel dimensions. The softmax function serves to normalize the values, while D_k operates as a scaling factor to mitigate gradient vanishing. To facilitate information interaction among cross-modal features, we perform an element-wise multiplication between the global attention maps of the two modalities, yielding a shared attention map A_{att}^i :

$$A_{att}^i = A_R^i \odot A_D^i. \quad (3)$$

We then multiply the shared attention map A_{att}^i with the corresponding value matrices, generating outputs of dimensions $M \times C$. Global features of the two modalities interact through superpixels, and the co-attentioned features are enhanced. Next, we employ another self-attention operation for the two modalities. By utilizing superpixel keys and pixel queries, we distribute the global information from the superpixels into the pixel space, obtaining attention maps P_R^i

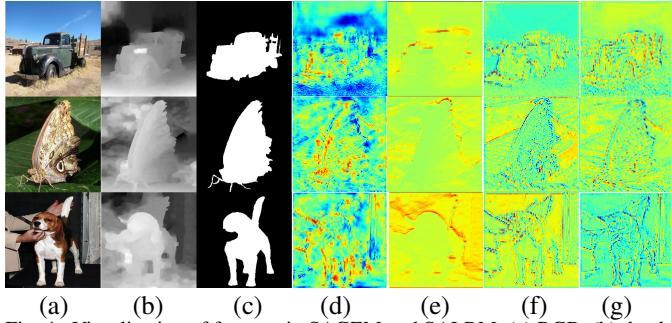


Fig. 4. Visualization of features in SAGEM and SALRM. (a) RGB, (b) depth, (c) GT, (d) RGB features, (e) depth features, (f) SAGEM features, and (g) SALRM features.

and P_D^i as follows:

$$\begin{aligned} P_R^i &= \text{Softmax} \left(\frac{K_{SR}^i \otimes Q_{SR}^{i \top}}{\sqrt{D_k}} \right), \\ P_D^i &= \text{Softmax} \left(\frac{K_{SD}^i \otimes Q_{SD}^{i \top}}{\sqrt{D_k}} \right). \end{aligned} \quad (4)$$

These mapping matrices are utilized to propagate the updated superpixel feature information back into the original pixel space:

$$\tilde{f}_{RD}^i = P_R^i \otimes (A_{att}^i \otimes V_R^i) + P_D^i \otimes (A_{att}^i \otimes V_D^i), \quad (5)$$

where \tilde{f}_{RD}^i signifies the enhanced features remapped to the original pixel space.

Finally, the globally enhanced features \tilde{f}_{RD}^i are further transformed and output through a Feedforward Network (FFN). We visualize features in our SAGEM in Fig. 4 (f). The SAGEM extracts features that prioritize the integration of cross-modal global contextual information. By capturing long-range dependencies, the module enhances the overall contours and semantic consistency of objects. The attention mechanism design enables our model to focus on cross-regional semantic associations. For instance, it can connect scattered object components into a coherent structure via contextual cues, thereby improving the overall performance of objects in complex scenes.

D. Superpixel Attention Local Refining Module

Although SAGEM can enhance global features, it inevitably loses some local information. Additionally, superpixel token generation may not be precise due to noise and other factors, resulting in semantic confusion of different regions. To address the above issues, we propose a Superpixel Attention Local Refining Module (SALRM) to utilize pixel similarity within superpixels to significantly enhance local features, boosting the model's ability to capture and express details.

As shown in Fig. 5, we first perform superpixel generation on the input RGB and depth features of SALRM, generating association matrices, denoted as S_R^i and S_D^i . To address the potential semantic confusion caused by superpixel generation, we need to filter the co-associated pixels of superpixels to highlight local features. The association matrix represents the weight relationship between each pixel and superpixels.

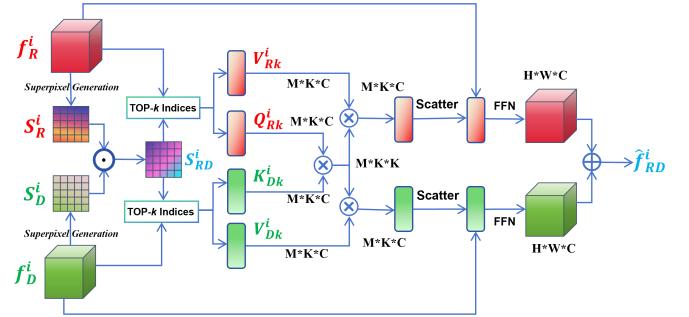


Fig. 5. Illustration of the Superpixel Attention Local Refining Module.

By performing element-wise multiplication on the association matrices of the two modalities, we generate a combined association matrix S_{RD}^i to determine the pixels that are semantically best-matched to the superpixel tokens:

$$S_{RD}^i = S_R^i \odot S_D^i. \quad (6)$$

For the combined association matrix, the weights of pixels that are similar within superpixels of both modalities are enhanced. Next, we need to determine which pixel features require individual local enhancement by sorting the weights. Given a superpixel denoted as $f = \{x(i)\}_N \in \mathbb{R}^{N \times C}$, encompassing numerous pixels where $x(i)$ represents the pixel, N signifies the count of selected pixels and C represents the channel count. It is worth noting that the pixel count within each superpixel can fluctuate. To streamline computations, we exclusively amplify the k most similar pixels in each superpixel. Through the application of the top- k function, we efficiently pinpoint the indices of these remarkably similar pixels.

Subsequently, we perform linear embedding on the input features to generate corresponding queries, keys, and values. Utilizing the previously acquired index data, we employ the gather function to extract values from Q , K , and V that are aligned with each superpixel position, and then appropriately rearrange their dimensions. The resultant local feature set possesses a dimension of $M \times N \times C$, where M indicates the superpixel count and N signifies the number of extracted components.

For these pixels with similar semantic information, we utilize the cross-modal local features Q_{Rk}^i and K_{Dk}^i to construct a local attention map through the cross-attention mechanism, thereby facilitating cross-modal information interaction of local features. The precise computation procedure is outlined below:

$$F_{att} = \text{Softmax} \left(\frac{Q_{Rk}^i \otimes K_{Dk}^{i \top}}{\sqrt{D_k}} \right). \quad (7)$$

Following this, we execute the matrix multiplication between the generated local attention map and the respective local values V_R^i and V_D^i , thereby obtaining the refined local features:

$$\begin{aligned} \hat{f}_{Rk}^i &= F_{att} \otimes V_{Rk}^i, \\ \hat{f}_{Dk}^i &= F_{att} \otimes V_{Dk}^i. \end{aligned} \quad (8)$$

Utilizing the previously secured index value information, we scatter the enhanced local features back to their original

positions and integrate them with the original features. We then employ the FFN to further refine the feature expression. Ultimately, we perform the element-wise summation (*i.e.*, $+$) to merge the enhanced features from both modalities, deriving the ultimate feature representation \hat{f}_{RD}^i :

$$\hat{f}_{RD}^i = \text{FFN}(\hat{f}_{Rk}^i + \hat{f}_R^i) + \text{FFN}(\hat{f}_{Dk}^i + \hat{f}_D^i). \quad (9)$$

By maximizing the utilization of the information embedded within the association matrices generated by superpixels, we can precisely pinpoint the most similar pixels within a superpixel and harness a cross-attention mechanism for local feature refinement. We visualize features in our SALRM in Fig. 4 (g). The SALRM extracts features that prioritize enhancing local structural details (*e.g.*, textures and edges). The attention mechanism suppresses noise by strengthening intra-superpixel feature similarity, refines object boundaries, and enhances the structural representation accuracy of salient regions.

E. Loss Function

We utilize a hybrid loss function that combines Binary Cross-Entropy (BCE) loss and Intersection over Union (IoU) loss for comprehensive supervision. This hybrid approach ensures meticulous supervision at both pixel level and patch level. Furthermore, we incorporate a deep supervision strategy to bolster the network's supervision power and expedite convergence.

The overall loss function L_{total} is computed as:

$$L_{\text{total}} = -\frac{1}{N} \sum_{i=1}^N \underbrace{[g_i \log(s_i) + (1 - g_i) \log(1 - s_i)]}_{\text{BCE Loss}} + 1 - \underbrace{\frac{\sum_{i=1}^N s_i \times g_i}{\sum_{i=1}^N (s_i + g_i - s_i \times g_i)}}, \quad (10)$$

where N denotes the total number of pixels, s_i signifies the saliency value for the i -th pixel in the predicted saliency map, and g_i corresponds to the ground truth value for the same i -th pixel in GT .

IV. EXPERIMENTS

A. Experimental Protocol

1) *Datasets and Evaluation Metrics.* To validate STENet's effectiveness, we conduct experiments on seven datasets: NJUD (1985 image pairs, large-scale analysis) [53]; NLPR (1000 RGB/depth images, indoor/outdoor) [54]; LFSD (100 light field RGB-D pairs) [65]; SIP (929 high-quality images, multi-salient objects) [66]; SSD (80 stereo movie samples) [68]; STEREO (1000 binocular pairs, stereo SOD benchmark) [67]; DUT-RGBD (1200 pairs, diverse scenes) [55]. Training data combines 700 NLPR, 1485 NJUD, and 800 DUT-RGBD images per recent works [22], [63], [64].

We employ five general metrics: MAE (M) [69], E-measure (E_ξ) [70], S-measure (S_m) [71], F-measure (F_β) [72], and Weighted F-measure (F_β^ω) [73].

2) *Implementation Details.* We implement STENet using PyTorch [74] on a NVIDIA GTX 3090 GPU (24GB RAM). The pre-trained Swin-B Transformer [39] serves as the backbone, with RGB/depth images resized to 384×384 (depth replicated to three channels). We use Adam optimizer [75] with cosine scheduling ($lr=5e^{-5}$, batch=5, 100 epochs) and augmentations like affine transforms and color jittering.

For ResNet-50 [76] backbone experiments, inputs are 256×256, optimized with Adam [75] optimizer ($lr=5e^{-5}$, batch=24, 100 epochs). Test predictions are resized to original dimensions for evaluation.

B. Comparison with State-of-the-arts

We conduct a comparison of our model against 14 state-of-the-art RGB-D SOD methods that use CNN and transformer as the backbone. We evaluate them uniformly by running the provided code and pre-trained models or using the saliency maps provided by the authors. The compared methods with ResNet-50 as the backbone include DSNet [56], DCFNet [30], CIRNet [57], C2DFNet [38], DIGRNet [58], HINet [60], CFIDNet [59], CAVER [22], RD3D+ [61], and LAFB [62]. The compared methods with Swin-B as the backbone include SwinNet [40], CATNet [63], PICRNet [43], and CPNet [64].

1) *Quantitative Evaluation:* As shown in Tab. I and II, STENet achieves superior performance across five metrics with both backbones, demonstrating strong generalization and robustness. Using Swin-B as the backbone, STENet outperforms most state-of-the-art methods, while ResNet-50 based STENet still delivers competitive results. Although DIGRNet [58] slightly outperforms our model in E_m and F_m on NJUD, it lags far behind in complex scenarios like SIP. On SIP, STENet with ResNet-50 surpasses the second-best CAVER [22] by 0.06 in M , 0.09 in F_β , 0.012 in E_m , 0.016 in S_m , and 0.022 in F_β^ω . It also shows strong competitiveness on complex datasets (STEREO and SSD). Compared to transformer-based models, STENet ranks top in most metrics, excelling in M while slightly lower in S_m and F_β^ω than CATNet [63] and CPNet [64]. This is attributed to its enhanced local information processing for complex scene details.

2) *Qualitative Evaluation:* To visually demonstrate the superiority of our proposed STENet, we show challenging scenes from the test set and compare the saliency maps of seven methods that use ResNet-50 as the backbone in Fig. 6. Rows 1 and 2 represent the scene of small salient object, while Row 3 represents the scene of large salient object. It can be seen that our method generates excellent saliency maps in handling the details of salient objects. Rows 4 and 5 represent multi-object scenes, row 6 represents complex background scenes, and rows 7 and 8 represent scenes where the background and foreground are similar. Rows 9-10 represent cases where the depth map is unreliable, and row 11 represents weak texture scenes. Rows 12-13 represent low-light scenes. It is difficult to see salient objects from RGB images. Our method stably learns features from depth maps, highlighting salient objects.

When faced with different challenging scenes, our method can still generate better saliency maps compared to recent methods, demonstrating its good robustness.

TABLE I

QUANTITATIVE COMPARISONS (%) ON THE TEST SET OF NJUD [53], NLPR [54], AND DUT-RGBD [55] DATASET. THE TOP TWO RESULTS IN EACH COLUMN ARE HIGHLIGHTED IN **RED** AND **BLUE**.

Model	Publisher	Year	Backbone	#Param (M)↓	FLOPs (G)↓	NJUD [53]				NLPR [54]				DUT-RGBD [55]						
						$M \downarrow$	$F_\beta \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_\beta^\omega \uparrow$	$M \downarrow$	$F_\beta \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_\beta^\omega \uparrow$	$M \downarrow$	$F_\beta \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_\beta^\omega \uparrow$
DSNet [56]	TIP	2021	Res50	-	-	.034	.922	.947	.921	.898	.024	.917	.956	.926	.886	.079	.821	.886	.841	.774
DCFNet [30]	CVPR	2021	Res50	108.5	107.8	.035	.915	.944	.911	.893	.021	.912	.959	.923	.892	.071	.835	.891	.836	.766
CIRNet [57]	TIP	2022	Res50	103.1	22.5	.035	.927	.943	.925	.895	.023	.924	.955	.933	.889	.031	.940	.958	.932	.904
C2DFNet [38]	TMM	2022	Res50	47.5	22.1	.038	.909	.941	.907	.885	.021	.916	.961	.927	.897	.026	.943	.963	.933	.918
DIGRNet [58]	TMM	2022	Res50	201.8	68.2	.028	.936	.954	.933	.909	.023	.928	.956	.935	.895	.033	.938	.951	.926	.898
CFIDNet [59]	NCA	2022	Res50	53.9	42.9	.038	.914	.929	.914	.886	.026	.906	.951	.922	.881	.039	.920	.944	.916	.887
HINet [60]	PR	2023	Res50	98.9	389.7	.039	.913	.938	.915	.881	.026	.906	.950	.922	.876	.054	.873	.921	.884	.826
CAVER [22]	TIP	2023	Res50	55.5	44.3	.031	.924	.953	.920	.906	.020	.921	.963	.928	.901	.042	.904	.937	.903	.874
RD3D+ [61]	TNNLS	2024	Res50	28.9	43.3	.033	.928	.948	.927	.899	.022	.920	.958	.933	.889	.031	.944	.960	.936	.908
LAFB [62]	TCSV	2024	Res50	451.8	137.6	.033	.918	.947	.916	.897	.027	.909	.946	.913	.876	.030	.935	.959	.927	.909
Ours		2025	Res50	56.4	25.5	.029	.930	.953	.923	.912	.020	.929	.964	.933	.909	.026	.944	.966	.935	.922
SwinNet [40]	TCSV	2022	SwinB	198.7	124.3	.027	.938	.954	.934	.917	.018	.936	.969	.941	.913	.021	.958	.973	.948	.933
CATNet [63]	TMM	2023	SwinB	262.6	341.8	.026	.937	.957	.932	.922	.018	.934	.970	.940	.916	.020	.960	.974	.953	.942
PICRNet [43]	ACM MM	2023	SwinB	106.8	121.3	.029	.930	.950	.927	.912	.019	.931	.967	.935	.911	.021	.953	.970	.943	.933
CPNet [64]	IJCV	2024	SwinB	216.5	129.3	.025	.941	.960	.935	.927	.016	.936	.972	.940	.922	.019	.961	.976	.951	.948
Ours	TMM	2025	SwinB	180.5	118.3	.023	.941	.967	.932	.931	.015	.938	.975	.939	.929	.018	.961	.979	.950	.952

TABLE II

QUANTITATIVE COMPARISONS (%) ON THE TEST SET OF LFSD [65], SIP [66], STEREO [67], AND SSD [68] DATASET. THE TOP TWO RESULTS IN EACH COLUMN ARE HIGHLIGHTED IN **RED** AND **BLUE**.

Model	Backbone	LFSD [65]				SIP [66]				STEREO [67]				SSD [68]							
		$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$
DSNet [56]	Res50	.069	.858	.905	.868	.826	.052	.881	.920	.876	.840	.036	.911	.947	.915	.882	.045	.878	.923	.885	.838
DCFNet [30]	Res50	.075	.841	.883	.841	.805	.051	.884	.921	.875	.848	.039	.901	.941	.901	.875	.049	.851	.905	.864	.814
CIRNet [57]	Res50	.068	.883	.904	.875	.838	.052	.896	.918	.888	.848	.039	.913	.939	.915	.872	.049	.864	.911	.878	.816
C2DFNet [38]	Res50	.065	.867	.902	.863	.835	.052	.877	.918	.871	.841	.038	.897	.942	.902	.871	.047	.860	.920	.872	.827
DIGRNet [58]	Res50	.067	.865	.906	.873	.828	.053	.897	.918	.885	.849	.038	.914	.943	.916	.877	.053	.846	.898	.866	.804
CFIDNet [59]	Res50	.071	.865	.901	.870	.828	.060	.870	.906	.864	.825	.043	.897	.932	.901	.867	.050	.871	.913	.879	.829
HINet [60]	Res50	.076	.847	.898	.852	.802	.066	.855	.899	.856	.805	.049	.883	.927	.892	.839	.049	.852	.916	.865	.808
CAVER [22]	Res50	.063	.876	.914	.873	.844	.042	.901	.933	.892	.872	.033	.912	.949	.913	.889	.041	.859	.921	.878	.834
RD3D+ [61]	Res50	.076	.849	.896	.861	.807	.047	.899	.928	.891	.857	.039	.905	.940	.914	.867	.044	.859	.920	.882	.820
LAFB [62]	Res50	.065	.870	.904	.866	.833	.043	.903	.932	.894	.872	.037	.903	.943	.904	.875	.042	.864	.923	.878	.835
Ours	Res50	.057	.880	.914	.877	.848	.036	.920	.945	.908	.894	.034	.910	.949	.911	.888	.040	.871	.927	.882	.844
SwinNet [40]	SwinB	.059	.889	.917	.886	.854	.035	.927	.947	.911	.896	.033	.918	.949	.919	.889	.040	.879	.925	.892	.851
CATNet [63]	SwinB	.051	.891	.926	.894	.863	.035	.928	.948	.911	.897	.030	.922	.954	.921	.900	-	-	-	-	-
PICRNet [43]	SwinB	.053	.894	.922	.888	.864	.040	.914	.937	.898	.883	.031	.920	.952	.920	.898	.047	.862	.922	.874	.832
CPNet [64]	SwinB	.050	.892	.926	.893	.869	.035	.927	.944	.907	.900	.029	.922	.955	.920	.901	.035	.892	.932	.894	.864
Ours	SwinB	.048	.891	.928	.890	.873	.032	.931	.952	.910	.911	.028	.925	.957	.926	.907	.035	.895	.935	.896	.870

TABLE III

QUANTITATIVE RESULTS OF ASSESSING THE INDIVIDUAL AND JOINT CONTRIBUTIONS OF THE TWO MODULES IN STENET. THE BEST RESULT IN EACH COLUMN IS HIGHLIGHTED IN **BOLD**.

Model	NJUD [53]				LFSD [65]				SIP [66]				Average							
	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$
Baseline	.029	.930	.954	.926	.913	.057	.880	.914	.877	.848	.036	.920	.945	.908	.894	.041	.910	.938	.904	.885
+SAGEM	.026	.936	.963	.927	.925	.050	.884	.917	.880	.861	.037	.924	.941	.901	.899	.038	.915	.940	.903	.895
+SALRM	.025	.935	.961	.927	.925	.051	.879	.917	.879	.857	.036	.924	.944	.900	.903	.037	.913	.941	.902	.894
+SAGEM+SALRM (Ours)	.023	.941	.967	.932	.931	.048	.891	.928	.890	.873	.032	.931	.952	.910	.911	.034	.921	.949	.911	.905
+CAF+SALRM	.023	.936	.963	.929	.926	.047	.891	.929	.891	.873	.036	.920	.943	.902	.897	.035	.916	.945	.907	.899
+VAM+SALRM	.024	.936	.962	.929	.926	.049	.889	.925	.886	.869	.033	.924	.948	.907	.904	.035	.916	.945	.907	.900
+SAGEM+CA	.025	.935	.962	.929	.926	.053	.876	.917	.877	.855	.037	.921	.940	.900	.897	.038	.911	.940	.902	.893
+SAGEM+SA	.028	.930	.957	.923	.917	.051	.885	.919	.882	.861	.035	.923	.945	.902	.900	.038	.913	.940</		



Fig. 6. Qualitative comparisons of our method and state-of-the-art methods in challenging RGB-D SOD scenes.

C. Ablation Studies

To thoroughly evaluate the effectiveness of our framework, we conduct ablation experiments using Swin-B as the backbone. These experiments focus on analyzing the individual and combined contributions of key components, including the roles of SAGEM and SALRM, and the configuration analysis and the mask size selection in superpixel generation. Through these studies, we gain valuable insights into the design choices and validate the robustness of our method.

1) *The Individual and Joint Contributions of SAGEM and SALRM.* We evaluated the individual and joint impacts of SAGEM and SALRM on the baseline model, focusing on key performance metrics. As shown in Tab. III, integrating SAGEM alone leads to notable improvements: the baseline's average metrics demonstrate an increase of 0.005 in F_m , a 0.002 rise in E_m , a 0.010 improvement in F_m^ω , and a 0.003 reduction in M . Similarly, SALRM alone exhibits consistent gains: F_m increases by 0.003, E_m by 0.003, F_m^ω by 0.009, and M decreases by 0.004. The combined application of SAGEM and SALRM achieves the most significant advancements: F_m

improves by 0.011, E_m by 0.011, S_m by 0.007, and F_m^ω by 0.020, with M reducing by 0.007. Notably, on the complex SIP dataset, the modules demonstrate pronounced synergistic effects, balancing and enhancing performance across all evaluated metrics.

2) *Effectiveness of Superpixels in SAGEM.* The proposed SAGEM leverages superpixels to implement spatial-wise cross-attention, enabling cross-modal feature enhancement through global information interaction. To analyze the role of superpixels in SAGEM, we explore the effect of removing superpixels on global cross-modal feature fusion. Without superpixels, the mechanism would revert to standard spatial-wise self-attention, leading to a memory explosion. Therefore, we employ two existing alternative methods for comparison, demonstrating the critical contribution of superpixels to global feature fusion:

- (a) Channel Attention Fusion (CAF) from CMX [77], which applies channel-wise cross-attention to modality-specific features;
- (b) View-Mixed Fusion (VMF) in CAVER [22], which in-

TABLE IV
CONFIGURATION ANALYSIS IN SUPERPIXEL GENERATION. THE BEST RESULT IN EACH COLUMN IS HIGHLIGHTED IN **BOLD**.

Size	FLOPs(G)↓	NJUD [53]					LFSD [65]					SIP [66]					Average				
		$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$
(1, 1, 1, 1)	118.4	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
(4, 4, 4, 4)	118.3	.025	.939	.957	.935	.924	.059	.880	.912	.883	.855	.037	.920	.941	.903	.893	.040	.913	.937	.907	.891
(12, 12, 6, 6)	118.3	.023	.941	.967	.932	.931	.048	.891	.928	.890	.873	.032	.931	.952	.910	.911	.034	.921	.949	.911	.905
(12, 12, 12, 12)	118.3	.024	.938	.964	.930	.927	.046	.891	.926	.889	.872	.034	.925	.947	.905	.902	.035	.918	.946	.908	.900
(24, 24, 12, 12)	118.3	.023	.943	.958	.936	.926	.048	.891	.925	.886	.870	.031	.930	.954	.911	.904	.034	.921	.946	.911	.900
(48, 24, 12, 12)	118.3	.024	.936	.963	.929	.926	.047	.889	.928	.888	.870	.033	.925	.949	.907	.904	.035	.917	.947	.908	.900
(48, 48, 24, 12)	118.3	.026	.932	.960	.925	.921	.056	.878	.913	.877	.857	.033	.925	.949	.906	.905	.038	.912	.941	.903	.894
(96, 48, 24, 12)	118.3	.027	.935	.960	.926	.924	.058	.879	.911	.873	.855	.038	.919	.940	.898	.894	.041	.911	.937	.899	.891

TABLE V
QUANTITATIVE RESULTS OF DIFFERENT METHODS IN SUPERPIXEL GENERATION. THE BEST RESULT IN EACH COLUMN IS HIGHLIGHTED IN **BOLD**.

Superpixel method	NJUD [53]					LFSD [65]					SIP [66]					Average				
	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$
K-means clustering [52]	.025	.938	.962	.928	.925	.050	.891	.926	.883	.868	.035	.927	.947	.903	.903	.037	.919	.945	.905	.899
Self-attention based method [49]	.028	.927	.957	.922	.916	.054	.872	.913	.877	.854	.037	.920	.941	.898	.895	.040	.906	.937	.899	.888
Cross-attention based method [51]	.024	.941	.957	.932	.925	.053	.877	.916	.882	.858	.033	.921	.950	.902	.911	.036	.916	.942	.908	.898
Ours	.023	.941	.967	.932	.931	.048	.891	.928	.890	.873	.032	.931	.952	.910	.911	.034	.921	.949	.911	.905

TABLE VI

QUANTITATIVE RESULTS OF MASK SIZE SELECTION IN SUPERPIXEL GENERATION. THE BEST RESULT IN EACH COLUMN IS HIGHLIGHTED IN **BOLD**.

Mask size	NJUD [53]					LFSD [65]					SIP [66]					Average				
	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$	$M \downarrow$	$F_m \uparrow$	$E_m \uparrow$	$S_m \uparrow$	$F_m^\omega \uparrow$
3×3	.024	.940	.965	.931	.926	.048	.889	.923	.888	.872	.034	.926	.948	.903	.901	.035	.918	.945	.907	.900
5×5	.023	.941	.967	.932	.931	.048	.891	.928	.890	.873	.032	.931	.952	.910	.911	.034	.921	.949	.911	.905
7×7	.029	.923	.955	.919	.911	.053	.883	.921	.879	.862	.037	.916	.942	.900	.893	.040	.907	.939	.899	.889

tegrates features through a parameter-free spatial cross-attention mechanism combined with patch-wise token re-embedding.

Tab. III shows SAGEM+SALRM outperforms CAF+SALRM and VMF+SALRM. Compared to VMF+SALRM, it achieves a 0.54% higher F_β and 2.86% lower M . SAGEM excels in global feature enhancement versus pixel-based cross-attention, especially on the SIP dataset, indicating better suitability for complex scenes. Notably, CAF+SALRM performs slightly better than our method on the LFSD dataset. The main reason is that the scenes and objects in the LFSD dataset are relatively simple, and the quality of depth maps is high. This means the advantages of SAGEM's global modeling are not fully demonstrated. On the contrary, our method may fail to achieve optimal performance due to potential overfitting.

3) *Effectiveness of Superpixels in SALRM*. The SALRM leverages statistical information derived from superpixels to perform local pixel-level operations, enabling adaptive feature enhancement for detailed refinement. We employ two alternative methods, specifically commonly used convolution-based attention operations, to refine local details, thereby validating the role of superpixels in SALRM:

- Channel Attention (CA), which strengthens local regions by adaptively recalibrating feature channels, enhancing discriminative information within pixels while suppressing irrelevant features;
- Spatial Attention (SA), which refines local regions by

dynamically weighting spatial locations, emphasizing structural details and improving the representation of fine-grained patterns within pixels.

As shown in Tab. III, SAGEM+SALRM outperforms both SAGEM+CA and SAGEM+SA across all metrics. Notably, compared to SAGEM+CA, it achieves a 1.1% increase in F_β and a 10.5% reduction in M . Against SAGEM+SA, it yields a 1.3% improvement in F_m^ω . The gains in E_m and S_m validate that local feature enhancement improves boundary detail and structural consistency. Experiments also show SALRM overcomes traditional CNN neighborhood window limitations for more precise local region enhancement.

5) *Configuration Analysis in Superpixel Generation*. We explore the impact of superpixel token sizes on model performance, given the encoder's hierarchical features (resolutions from 96×96 to 12×12). Superpixel grid sizes $[p \times p]$ determine the number of superpixels M per layer, affecting both efficiency and performance. As shown in Tab. IV, a moderate configuration (12, 12, 6, 6) achieves optimal average performance by balancing local-global feature extraction. A (1, 1, 1, 1) setup mimics standard spatial attention, causing memory explosion and training failure, while (96, 48, 24, 12)—with one superpixel per layer—shows competitive performance despite reduced granularity. All configurations maintain consistent FLOPs, confirming superpixel size has minimal impact on computational efficiency.

6) *Different Superpixel Sampling Methods*. To evaluate the effectiveness of our superpixel generation, we replace our

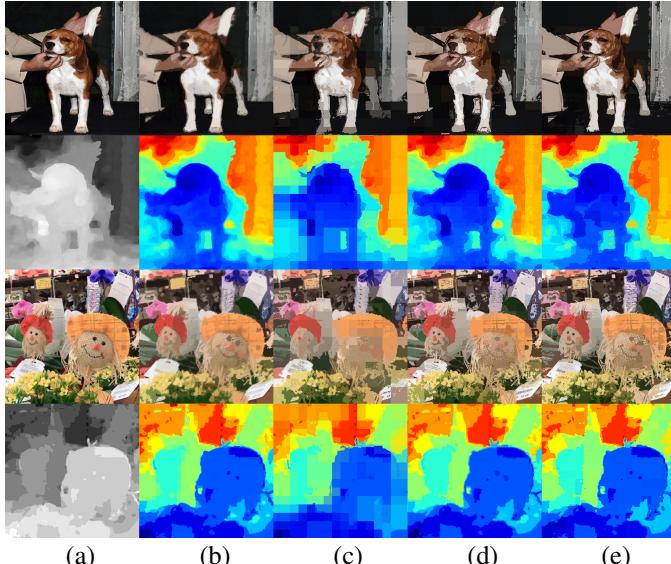


Fig. 7. Visualization of different superpixel representations. All images in the figure are 256×256 in size with a superpixel grid size of $[16 \times 16]$. (a) RGB and depth, (b) K-means clustering [52], (c) self-attention based method [49], (d) cross-attention based method [51], and (e) our method (with neighborhood expansion).

superpixel generation with three types of superpixel generation methods, including K-means clustering [52], self-attention based method [49], and cross-attention based method [51]. We present the quantitative results in Tab. V.

The results show that our method performs better in most metrics across multiple datasets, and its average performance across all metrics is also the best. We visualize the superpixel representations based on the connection matrix. The connection matrix represents the pixels contained in each superpixel. We perform the argmax operation on the connection matrix to obtain the superpixel with the highest probability that each pixel belongs to. The pixel positions belonging to the same superpixel are replaced with the average pixel value. In Fig. 7, we visualize the learned soft associations by selecting the argmax over the superpixels.

Visualization results demonstrate that our superpixel generation method has the strongest semantic discriminative ability and the best detail preservation effect. The K-means clustering method relies on Euclidean space distance, resulting in relatively uniform superpixel distribution but details that tend to be blurred. The self-attention method has poor clustering performance and unstable effectiveness. The cross-attention method is limited by neighborhood constraints, and its clustering performance degrades in complex environments. In contrast, our method performs more prominently in clustering, especially in detail preservation. This is because the superpixel neighborhood expansion in our method enhances semantic discriminative ability, overcoming distance constraints.

7) *Mask Size Selection in Superpixel Generation.* We evaluate mask sizes (3×3 , 5×5 , and 7×7) for superpixel-pixel interaction. To improve efficiency, we limit each superpixel to 9 neighboring superpixels via masking. Tab. VI shows the 5×5 mask achieves optimal performance, with the highest F_m and E_m on SIP and superior average metrics with low M . The 3×3 mask performs comparably, while 7×7 degrades performance

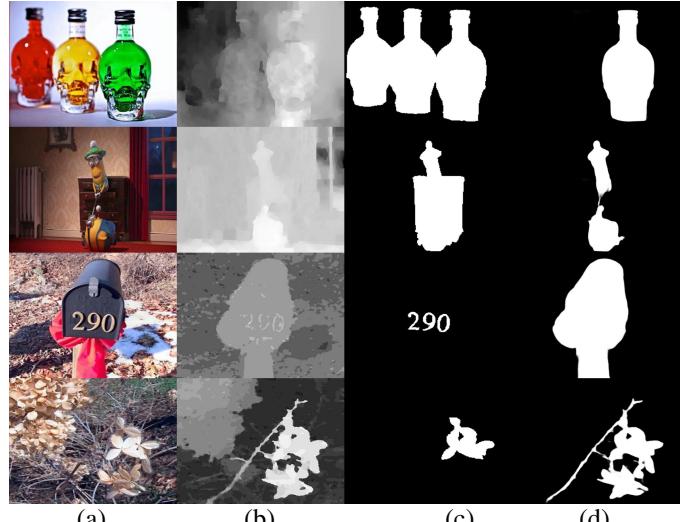


Fig. 8. Visual examples of failure cases. (a) RGB image. (b) Depth map. (c) GT. (d) Ours.

(e.g., S_m drops 0.010 on SIP) due to irrelevant pixel inclusion. The 5×5 mask balances global context and local precision, chosen as the default for performance-efficiency trade-off.

D. Failure Cases

Although our method exhibits several advantages over existing methods, it also has some limitations. We show some failure cases of our method in Fig. 8.

First, noisy depth maps make our method difficult to distinguish foreground from background. The saliency maps of our method may mistakenly include low-contrast regions in the background. For example, in the first row of Fig. 8, only one of the three bottles in the GT is segmented, and the cabinet in the second row is not identified.

Second, in weak-texture scenes, local feature similarity may cause superpixels to merge cross-semantic regions. For instance, the salient digits in the third row are merged with surrounding objects semantically, and color-similar leaves in the fourth row are misidentified as salient objects.

In future work, we plan to introduce cross-modal consistency constraints or depth map enhancement preprocessing to alleviate the above limitations.

V. CONCLUSION

In this paper, we present a novel solution, named STENet, to address the limitations of existing transformer-based RGB-D SOD methods, particularly their quadratic computational complexity and insufficient local detail extraction. By integrating superpixels into multi-head self-attention mechanisms, our STENet introduces two innovative modules, *i.e.*, SAGEM and SALRM. SAGEM reduces computational overhead by modeling global pixel-to-superpixel relationships instead of traditional pixel-to-pixel interactions, effectively capturing region-level dependencies while maintaining efficiency. SALRM enhances local details by leveraging pixel similarity within superpixels to filter and refine critical local pixels, thereby addressing the challenge of fine-grained feature extraction.

Last, the globally and locally enhanced features generated by SAGEM and SALRM are effectively fused with cross-scale features, resulting in a robust and comprehensive feature representation that significantly enhances the overall performance of the model. Extensive experiments on seven RGB-D SOD datasets demonstrate that our STENet achieves competitive performance against state-of-the-art methods, validating its effectiveness in balancing computational efficiency and detection accuracy.

REFERENCES

- [1] G. Li, C. Han, and Z. Liu, "No-service rail surface defect segmentation via normalized attention and dual-scale interaction," *IEEE Trans. Instrum. Meas.*, pp. 1–10, 2023.
- [2] X. Jia, H. Lu, and M.-H. Yang, "Visual tracking via coarse and fine structural local sparse appearance models," *IEEE Trans. Image Process.*, vol. 25, no. 10, pp. 4555–4564, Oct. 2016.
- [3] G. Li, Z. Bai, and Z. Liu, "Texture-semantic collaboration network for orsi salient object detection," *IEEE Trans. Circuits Syst. II-Express Briefs*, vol. 71, no. 4, pp. 2464–2468, 2024.
- [4] G. Li, Z. Bai, Z. Liu, X. Zhang, and H. Ling, "Salient object detection in optical remote sensing images driven by transformer," *IEEE Trans. Image Process.*, vol. 32, pp. 5257–5269, 2023.
- [5] G. Li, Z. Liu, R. Shi, Z. Hu, W. Wei, Y. Wu, M. Huang, and H. Ling, "Personal fixations-based object segmentation with object localization and boundary preservation," *IEEE Trans. Image Process.*, vol. 30, pp. 1461–1475, Jan. 2021.
- [6] X. Zhou, K. Shen, and Z. Liu, "Admnet: Attention-guided densely multi-scale network for lightweight salient object detection," *IEEE Trans. Multimedia*, vol. 26, pp. 10828–10841, 2024.
- [7] G. Zhu, L. Wang, and J. Tang, "Learning discriminative context for salient object detection," *Eng. Appl. Artif. Intell.*, vol. 131, p. 107820, 2024.
- [8] Y. K. Yun and W. Lin, "Towards a complete and detail-preserved salient object detection," *IEEE Trans. Multimedia*, vol. 26, pp. 4667–4680, 2024.
- [9] G. Li, Z. Liu, and H. Ling, "Icnet: Information conversion network for rgb-d based salient object detection," *IEEE Trans. Image Process.*, vol. 29, pp. 4873–4884, 2020.
- [10] G. Li, Z. Liu, M. Chen, Z. Bai, W. Lin, and H. Ling, "Hierarchical alternate interaction network for rgb-d salient object detection," *IEEE Trans. Image Process.*, vol. 30, pp. 3528–3542, 2021.
- [11] A. Mou, Y. Lu, J. He, D. Min, K. Fu, and Q. Zhao, "Salient object detection in rgb-d videos," *IEEE Transactions on Image Processing*, vol. 33, pp. 6660–6675, 2024.
- [12] J. He, K. Fu, X. Liu, and Q. Zhao, "Samba: A unified mamba-based framework for general salient object detection," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 25 314–25 324.
- [13] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," *Neural Comput.*, vol. 1, no. 4, pp. 541–551, Dec. 1989.
- [14] N. Wang and X. Gong, "Adaptive fusion for RGB-D salient object detection," *IEEE Access*, vol. 7, pp. 55 277–55 284, May 2019.
- [15] J. Han, H. Chen, N. Liu, C. Yan, and X. Li, "CNNs-based RGB-D saliency detection via cross-view transfer and multiview fusion," *IEEE Trans. Cybern.*, vol. 48, no. 11, pp. 3171–3183, Nov. 2018.
- [16] H. Chen, Y. Li, and D. Su, "Multi-modal fusion network with multi-scale multi-path and cross-modal interactions for RGB-D salient object detection," *Pattern Recognit.*, vol. 86, pp. 376–385, Feb. 2019.
- [17] Z. Zhang, Z. Lin, J. Xu, W.-D. Jin, S.-P. Lu, and D.-P. Fan, "Bilateral attention network for rgb-d salient object detection," *IEEE Trans. Image Process.*, vol. 30, pp. 1949–1961, 2021.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NeurIPS*, Dec. 2017, pp. 6000–6010.
- [19] C. Yang, Y. Wang, J. Zhang, H. Zhang, Z. Wei, Z. Lin, and A. Yuille, "Lite vision transformer with enhanced self-attention," in *Proc. IEEE CVPR*, 2022, pp. 11 998–12 008.
- [20] H. Lin, X. Cheng, X. Wu, and D. Shen, "Cat: Cross attention in vision transformer," in *Proc. IEEE ICME*, 2022, pp. 1–6.
- [21] J. Wu, F. Hao, W. Liang, and J. Xu, "Transformer fusion and pixel-level contrastive learning for rgb-d salient object detection," *IEEE Trans. Multimedia*, vol. 26, pp. 1011–1026, 2024.
- [22] Y. Pang, X. Zhao, L. Zhang, and H. Lu, "CAVER: Cross-modal view-mixed transformer for bi-modal salient object detection," *IEEE Trans. Image Process.*, vol. 32, pp. 892–904, 2023.
- [23] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 888–905, 2000.
- [24] V. Jampani, D. Sun, M.-Y. Liu, M.-H. Yang, and J. Kautz, "Superpixel sampling networks," in *Proc. ECCV*, 2018, pp. 352–368.
- [25] Z. Liu, S. Shi, Q. Duan, W. Zhang, and P. Zhao, "Salient object detection for rgb-d image by single stream recurrent convolution neural network," *Neurocomputing*, vol. 363, pp. 46–57, 2019.
- [26] X. Zhao, L. Zhang, Y. Pang, H. Lu, and L. Zhang, "A single stream network for robust and real-time rgb-d salient object detection," in *Proc. ECCV*, 2020, pp. 646–662.
- [27] Q. Chen, Z. Liu, Y. Zhang, K. Fu, Q. Zhao, and H. Du, "Rgb-d salient object detection via 3d convolutional neural networks," in *Proc. AAAI*, vol. 35, no. 2, 2021, pp. 1063–1071.
- [28] K. Fu, D.-P. Fan, G.-P. Ji, Q. Zhao, J. Shen, and C. Zhu, "Siamese network for rgb-d salient object detection and beyond," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 9, pp. 5541–5559, 2022.
- [29] C. Sun, Q. Zhang, C. Zhuang, and M. Zhang, "Bmfnet: Bifurcated multi-modal fusion network for rgb-d salient object detection," *Image Vis. Comput.*, vol. 147, p. 105048, 2024.
- [30] W. Ji, J. Li, S. Yu, M. Zhang, Y. Piao, S. Yao, Q. Bi, K. Ma, Y. Zheng, H. Lu *et al.*, "Calibrated rgb-d salient object detection," in *Proc. IEEE CVPR*, 2021, pp. 9471–9481.
- [31] P. Sun, W. Zhang, H. Wang, S. Li, and X. Li, "Deep rgb-d saliency detection with depth-sensitive attention and automatic multi-modal fusion," in *Proc. IEEE CVPR*, 2021, pp. 1407–1417.
- [32] Z. Wu, G. Allibert, F. Meriaudeau, C. Ma, and C. Demonceaux, "Hidancet: Rgb-d salient object detection via hierarchical depth awareness," *IEEE Trans. Image Process.*, vol. 32, pp. 2160–2173, 2023.
- [33] W. Zhang, G.-P. Ji, Z. Wang, K. Fu, and Q. Zhao, "Depth quality-inspired feature manipulation for efficient rgb-d salient object detection," in *Proceedings of the 29th ACM international conference on multimedia*, 2021, pp. 731–740.
- [34] J. Han, H. Chen, N. Liu, C. Yan, and X. Li, "Cnns-based rgb-d saliency detection via cross-view transfer and multiview fusion," *IEEE Trans. Cybern.*, vol. 48, no. 11, pp. 3171–3183, Nov 2017.
- [35] N. Wang and X. Gong, "Adaptive fusion for rgb-d salient object detection," *IEEE Access*, vol. 7, pp. 55 277–55 284, 2019.
- [36] H. Chen, Y. Deng, Y. Li, T.-Y. Hung, and G. Lin, "Rgbd salient object detection via disentangled cross-modal fusion," *IEEE Trans. Image Process.*, vol. 29, pp. 8407–8416, 2020.
- [37] Y. Pang, L. Zhang, X. Zhao, and H. Lu, "Hierarchical dynamic filtering network for rgb-d salient object detection," *Proc. ECCV*, pp. 235–252, 2020.
- [38] M. Zhang, S. Yao, B. Hu, Y. Piao, and W. Ji, "C²DFNet: Criss-cross dynamic filter network for RGB-D salient object detection," *IEEE Trans. Multimedia*, vol. 25, pp. 5142–5154, 2022.
- [39] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE CVPR*, 2021, pp. 10 012–10 022.
- [40] Z. Liu, Y. Tan, Q. He, and Y. Xiao, "Swinnet: Swin transformer drives edge-aware rgbd and rgbt salient object detection," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 7, pp. 4486–4497, 2022.
- [41] Z. Liu, Y. Wang, Z. Tu, Y. Xiao, and B. Tang, "Tritransnet: Rgb-d salient object detection with a triplet transformer embedding network," in *Proc. ACM MM*, 2021, pp. 4481–4490.
- [42] B. Tang, Z. Liu, Y. Tan, and Q. He, "Hrtransnet: Hrformer-driven two-modality salient object detection," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 33, no. 2, pp. 728–742, 2022.
- [43] R. Cong, H. Liu, C. Zhang, W. Zhang, F. Zheng, R. Song, and S. Kwong, "Point-aware interaction and cnn-induced refinement network for rgbd salient object detection," in *Proc. ACM MM*, 2023, pp. 406–416.
- [44] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *Int. J. Comput. Vis.*, vol. 59, pp. 167–181, 2004.
- [45] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 5, pp. 603–619, 2002.
- [46] S. Na, L. Xumin, and G. Yong, "Research on k-means clustering algorithm: An improved k-means clustering algorithm," in *Proc. IITSI*, 2010, pp. 63–67.

[47] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, “Slic superpixels compared to state-of-the-art superpixel methods,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2274–2282, 2012.

[48] Y. Zhang, B. Pang, and C. Lu, “Semantic segmentation by early region proxy,” in *Proc. IEEE CVPR*, 2022, pp. 1258–1268.

[49] H. Huang, X. Zhou, J. Cao, R. He, and T. Tan, “Vision transformer with super token sampling,” in *Proc. IEEE CVPR*, 2023, pp. 22690–22699.

[50] A. Z. Zhu, J. Mei, S. Qiao, H. Yan, Y. Zhu, L.-C. Chen, and H. Kretzschmar, “Superpixel transformers for efficient semantic segmentation,” in *Proc. IEEE/RSJ IROS*, 2023, pp. 7651–7658.

[51] J. Mei, L.-C. Chen, A. Yuille, and C. Xie, “Spformer: Enhancing vision transformer with superpixel representation,” *arXiv preprint arXiv:2401.02931*, 2024.

[52] A. Zhang, W. Ren, Y. Liu, and X. Cao, “Lightweight image super-resolution with superpixel token interaction,” in *Proc. IEEE ICCV*, 2023, pp. 12728–12737.

[53] R. Ju, L. Ge, W. Geng, T. Ren, and G. Wu, “Depth saliency based on anisotropic center-surround difference,” in *Proc. IEEE ICIP*, 2014, pp. 1115–1119.

[54] H. Peng, B. Li, W. Xiong, W. Hu, and R. Ji, “Rgbd salient object detection: A benchmark and algorithms,” in *Proc. ECCV*, 2014, pp. 92–109.

[55] Y. Piao, W. Ji, J. Li, M. Zhang, and H. Lu, “Depth-induced multi-scale recurrent attention network for saliency detection,” in *Proc. IEEE ICCV*, 2019, pp. 7254–7263.

[56] H. Wen, C. Yan, X. Zhou, R. Cong, Y. Sun, B. Zheng, J. Zhang, Y. Bao, and G. Ding, “Dynamic selective network for rgb-d salient object detection,” *IEEE Trans. Image Process.*, vol. 30, pp. 9179–9192, 2021.

[57] R. Cong, Q. Lin, C. Zhang, C. Li, X. Cao, Q. Huang, and Y. Zhao, “Cirnet: Cross-modality interaction and refinement for rgb-d salient object detection,” *IEEE Trans. Image Process.*, vol. 31, pp. 6800–6815, 2022.

[58] X. Cheng, X. Zheng, J. Pei, H. Tang, Z. Lyu, and C. Chen, “Depth-induced gap-reducing network for rgb-d salient object detection: An interaction, guidance and refinement approach,” *IEEE Trans. Multimedia*, vol. 25, pp. 4253–4266, 2023.

[59] T. Chen, X. Hu, J. Xiao, G. Zhang, and S. Wang, “Cfidnet: Cascaded feature interaction decoder for rgb-d salient object detection,” *Neural. Comput. Appl.*, vol. 34, no. 10, pp. 7547–7563, 2022.

[60] H. Bi, R. Wu, Z. Liu, H. Zhu, C. Zhang, and T.-Z. Xiang, “Cross-modal hierarchical interaction network for rgb-d salient object detection,” *Pattern Recognit.*, vol. 136, p. 109194, 2023.

[61] Q. Chen, Z. Zhang, Y. Lu, K. Fu, and Q. Zhao, “3-d convolutional neural networks for rgb-d salient object detection and beyond,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 3, pp. 4309–4323, 2024.

[62] K. Wang, Z. Tu, C. Li, C. Zhang, and B. Luo, “Learning adaptive fusion bank for multi-modal salient object detection,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 34, no. 8, pp. 7344–7358, 2024.

[63] F. Sun, P. Ren, B. Yin, F. Wang, and H. Li, “Catnet: A cascaded and aggregated transformer network for rgb-d salient object detection,” *IEEE Trans. Multimedia*, vol. 26, pp. 2249–2262, 2024.

[64] X. Hu, F. Sun, J. Sun, F. Wang, and H. Li, “Cross-modal fusion and progressive decoding network for rgb-d salient object detection,” *Int. J. Comput. Vis.*, vol. 132, pp. 3067–3085, 2024.

[65] N. Li, J. Ye, Y. Ji, H. Ling, and J. Yu, “Saliency detection on light field,” in *Proc. IEEE CVPR*, 2014, pp. 2806–2813.

[66] D.-P. Fan, Z. Lin, Z. Zhang, M. Zhu, and M.-M. Cheng, “Rethinking rgb-d salient object detection: Models, data sets, and large-scale benchmarks,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 5, pp. 2075–2089, 2020.

[67] Y. Niu, Y. Geng, X. Li, and F. Liu, “Leveraging stereopsis for saliency analysis,” in *Proc. IEEE CVPR*, 2012, pp. 454–461.

[68] C. Zhu and G. Li, “A three-pathway psychobiological framework of salient object detection using stereoscopic technology,” in *Proc. IEEE ICCVW*, 2017, pp. 3008–3014.

[69] F. Perazzi, P. Krähenbühl, Y. Pritch, and A. Hornung, “Saliency filters: Contrast based filtering for salient region detection,” in *Proc. IEEE CVPR*, 2012, pp. 733–740.

[70] D.-P. Fan, C. Gong, Y. Cao, B. Ren, M.-M. Cheng, and A. Borji, “Enhanced-alignment measure for binary foreground map evaluation,” in *Proc. IJCAI*, Jul. 2018, pp. 698–704.

[71] D.-P. Fan, M.-M. Cheng, Y. Liu, T. Li, and A. Borji, “Structure-measure: A new way to evaluate foreground maps,” in *Proc. IEEE ICCV*, 2017, pp. 4548–4557.

[72] R. Achanta, S. Hemami, F. Estrada, and S. Süsstrunk, “Frequency-tuned salient region detection,” in *Proc. IEEE CVPR*, 2009, pp. 1597–1604.

[73] R. Margolin, L. Zelnik-Manor, and A. Tal, “How to evaluate foreground maps?” in *Proc. IEEE CVPR*, 2014, pp. 248–255.

[74] A. Paszke *et al.*, “PyTorch: An imperative style, high-performance deep learning library,” in *Proc. NeurIPS*, Dec. 2019, pp. 8024–8035.

[75] D. P. Kingma and J. L. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, May 2015, pp. 1–15.

[76] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE CVPR*, Jun. 2016, pp. 770–778.

[77] J. Zhang, H. Liu, K. Yang, X. Hu, R. Liu, and R. Stiefelhagen, “CMX: Cross-modal fusion for rgb-x semantic segmentation with transformers,” *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 12, pp. 14679–14694, 2023.