

Learning Synergistic Attention for Light Field Salient Object Detection

Yi Zhang¹

yi.zhang1@insa-rennes.fr

Geng Chen*²

geng.chen.cs@gmail.com

Qian Chen³

poly@mail.ustc.edu.cn

YuJia Sun⁴

thograce@163.com

Yong Xia²

yxia@nwpu.edu.cn

Olivier Deforges¹

odeforge@insa-rennes.fr

Wassim Hamidouche¹

Wassim.Hamidouche@insa-rennes.fr

Lu Zhang¹

lu.ge@insa-rennes.fr

¹ Univ Rennes, INSA Rennes, CNRS, IETR (UMR 6164), France

² National Engineering Laboratory for Integrated Aero-Space-Ground-Ocean Big Data Application Technology, School of Computer Science and Engineering, Northwestern Polytechnical University, China

³ School of Information Science and Technology, University of Science and Technology of China, China

⁴ School of Computer Science, Inner Mongolia University, China

Abstract

In this work, we propose Synergistic Attention Network (*SA-Net*) to address the light field salient object detection by establishing a synergistic effect between multi-modal features with advanced attention mechanisms. Our *SA-Net* exploits the rich information of focal stacks via 3D convolutional neural networks, decodes the high-level features of multi-modal light field data with two cascaded synergistic attention modules, and predicts the saliency map using an effective feature fusion module in a progressive manner. Extensive experiments on three widely-used benchmark datasets show that our *SA-Net* outperforms 28 state-of-the-art models, sufficiently demonstrating its effectiveness and superiority. Our code is available at <https://github.com/PanoAsh/SA-Net>.

1 Introduction

Salient object detection (SOD) is a task aiming to segment the objects that grasp most of the human attention. It plays a key role in learning human visual mechanism and in various

* Corresponding Author: Geng Chen (geng.chen.cs@gmail.com)

© 2021. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

computer vision applications, such as instance segmentation and video object segmentation. According to the input modalities, the SOD task can be classified into three categories: 2/3/4D SOD [46]

Recently, light field SOD [19] (or 4D SOD) has attracted increasing attention owing to the introduction of various light field benchmark datasets, such as DUT-LF [57], LFSO [17], HFUT [44], DUT-MV [27], and Lytro Illum [45]. In addition to all-in-focus (AiF) images, light field datasets [17, 44, 45] also provide focal stacks (FSs), multi-view images, and depth maps, where the FS is usually known as a series of focal slices focusing at different depths of a given scene while the depth map contains holistic depth information. Unlike 3D (RGB-D) SOD models, which utilize only two modalities, i.e., RGB images and depth maps, the light field SOD models also use multi-view images (e.g., [27, 50]), or FSs [29, 57, 46, 47] as auxiliary inputs to further improve the performance. It is worth noting that, most recent FS-based deep learning light field SOD models (e.g., ERNet [29]) have achieved state-of-the-art performance on three widely-used light field benchmark datasets [17, 57, 44].

Despite their advantages, existing works suffer from two major limitations. First, they explore little about the complementarities between AiF images and the FSs. Existing FS-based methods [29, 57, 46, 47] apply only channel attention mechanisms to weight the key feature channels at the decoding stage, to aid the feature fusion between the modalities of AiF and FS. Considering the fact that salient objects usually appear at specific depths of a given scene, AiF image may include redundant texture details compared to FS, in which a focal slice focuses on a local region at specific depth and blurs the others. New cross-modal fusion strategy, which applies more sophisticated attention mechanisms learning robust cross-modal complementarities, may help solve the issue. Second, the methods [29, 57, 46, 47] all pay little attention to the inter-slice modeling during the encoding stage of FSs. In practice, the AiF images are generated from FSs with a photo-montage technique [11], implying that the former simultaneously depict the spatial details of each local region, while the latter asynchronously focus on different local details along the sequential dimension (Figure 1). The relationship between focal slices reflects the context of given scenes as the changes of depth, thus appropriate to be encoded in a progressive manner.

To this end, we propose Synergistic Attention Network (*SA-Net*) to conduct light field SOD with rich information from AiF images and FSs. Specifically, we first employ 3D convolutional neural networks (CNNs) to progressively extract the sequential features from FSs. At the decoding stage, we propose a synergistic attention (*SA*) module, where the features from AiF images and FSs are selectively fused and optimized to achieve a synergistic effect for SOD. Finally, the multi-modal features are fed to our progressive fusion (*PF*) module, which fuses multi-modal features and predicts the saliency map in a progressive manner. In a nutshell, we provide four main contributions as follows:

- We propose the *SA* module to decode the high-level features from both AiF images and FSs with a synergistic attention mechanism. Our *SA* module exploits the most meaningful information from the multi-modal multi-level features, allowing accurate

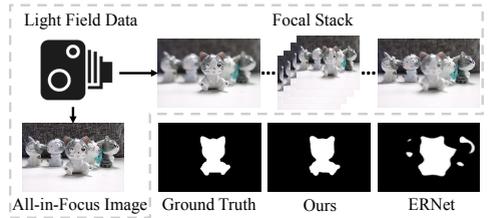


Figure 1: An example of light field SOD using our *SA-Net* and the state-of-the-art model, ERNet [29].

SOD by taking advantage of light field data.

- We introduce a dual-branch backbone to encode the AiF and FS information, simultaneously. To the best of our knowledge, our work is the first attempt to utilize 3D CNNs for the feature extraction of FSs in the field of light field SOD.
- We design the *PF* module to gradually fuse the selective high-level features for the final saliency prediction.
- Extensive experiments demonstrate that our *SA-Net* outperforms 28 state-of-the-art SOD models upon three widely-used light field datasets.

2 Related Works

In this section, we discuss recent works from the aspects of light field SOD, attention mechanisms, and 3D CNNs.

2.1 Light Field SOD

Datasets. To the best of our knowledge, LFSD [17] is the earliest public light field SOD dataset, which provides AiF images, FSs, multi-view (MV) images, depth maps, and micro-lens (ML) images for 100 different scenes. Later, [44] constructed HFUT, which contains 255 AiF images and other light field data types, including FS, MV, ML, and depth. More recent datasets such as DUT-LF [37] and DUT-MV [27] contain 1,462 and 1,580 AiF images, respectively. Though being much larger than the early datasets, DUT-LF [37] does not provide ML and MV images, while DUT-MV [27] only provides MV images. The most recently proposed dataset, Lytro Illum [45], provides 640 AiF images as well as the four extra forms of light field data. Note that all the five datasets mentioned above are released with SOD ground truth (pixel-wise binary mask) for each AiF image.

Methods. As the light field SOD is an emerging field, to the best of our knowledge, there are only 18 (11/7 traditional/deep learning-based, respectively) published methods. For traditional ones, the early method [17] conducted light field SOD by considering background and location related prior knowledge. In addition, [18] proposed a unified architecture based on weighted sparse coding. Later methods [19, 31, 33, 34, 35, 43, 44] explored and further combined multiple visual cues (e.g., depth, color contrast, light field flows and boundary prior) to detect saliency. Most recent methods [28, 39] shifted more attention to depth information and employed cellular automata for the saliency detection in light field. With the development of public light field datasets, deep learning-based methods were proposed to conduct SOD task. Specifically, [27] developed a view synthesis network to detect salient objects by involving MVs. With MVs as inputs, [50] further established a unified structure to synchronously conduct salient object and edge detection. Besides, [45] applied DeepLabv2 for SOD with MLs. As a mainstream, [29, 37, 46, 47] all employed ConvLSTM and channel attention mechanisms at the decoding stage to detect salient objects in AiF images and FSs. [37] and [47] both modeled the AiF and FS with separate encoder-decoder architectures. Specifically, [37] added the outputs of the decoders of both the FS branch and AiF branch. [47] concatenated the outputs of AiF branch and FS branch and used the ConvLSTM [32] to refine the concatenated features. In [46], FS and AiF share the same encoder. The extracted multi-modal features were further fused with memory-oriented attention modules.

The most recent [29] designed teacher network and student network to encode the FS and AiF respectively, and used ConvLSTM based attention module to facilitate the distillation process between the teacher network and student network.

2.2 Attention Mechanism

Generally, there are three main categories of attention mechanisms, including: channel attention [15], spatial attention [14], and self-attention [58], which is also a widely known concept in the field of natural language processing. These attention mechanisms can be easily embedded into different CNN-based architectures. Besides, mutual(co)-attention, as a specific type of attention mechanism, has been used in the fields of video object segmentation (e.g., [23]), RGB-D SOD (e.g., [22]), etc. However, mutual attention mechanism has been seldom studied in the field of light field SOD. As for the SOD task, attention mechanisms are either used for multi-modal fusion (e.g., [22], [10]) or multi-level feature fusion (e.g., [5]). Specifically, [5] proposed a progressive attention guided recurrent network to decode the multi-level features with channel-wise and spatial attention sequentially.

2.3 3D CNNs

3D CNNs have proved great competence in modeling spatial-temporal information of video data, thus dominating the video-based detection fields, such as action recognition [8] and video object segmentation [25].

Recently, RD3D [7] was proposed to address the task of 3D (RGB-D) SOD by using a 3D CNN-based encoder-decoder structure, and achieved promising performance on widely-used RGB-D SOD benchmarks. As for light field SOD, MTCNet [6] applied 3D CNN-based encoder to extract the depth features from multi-view images. The rich high-level features gained from 3D CNNs were then used to infer depth maps and facilitate the SOD task, synchronously. Since FSs are sequences of focal slices focusing at different depths, learning FSs' features via 3D CNNs possesses great potential to boost the model performance for light field SOD, but so far lacks investigation.

3 Synergistic Attention Network

In *SA-Net*, we exploit rich cross-modal complementary information with channel attention and co-attention mechanisms to achieve a synergistic effect between multi-level AiF and FS features. In addition, to capture the inter-slice information of FS, we employ 3D CNNs to extract rich features from FSs. Figure 2 shows an overview of our *SA-Net*, which consists of three major components, including a multi-modal encoder consisting of 2D and 3D CNNs (§ 3.1), two cascaded *SA* modules (§ 3.2), and a *PF* module (§ 3.3).

3.1 Multi-Modal Encoder

As shown in Figure 2, the encoder of our network is a dual-branch architecture for synchronous feeding of the two modalities, i.e., AiF images and FSs. For the 2D branch, we encode an input AiF image with a group of convolutional blocks. On the other hand, FS is represented as a 4D tensor with the last dimension T denoting the number of focal slices. We encode the FS with a stack of 3D convolutional blocks, which are able to jointly capture the

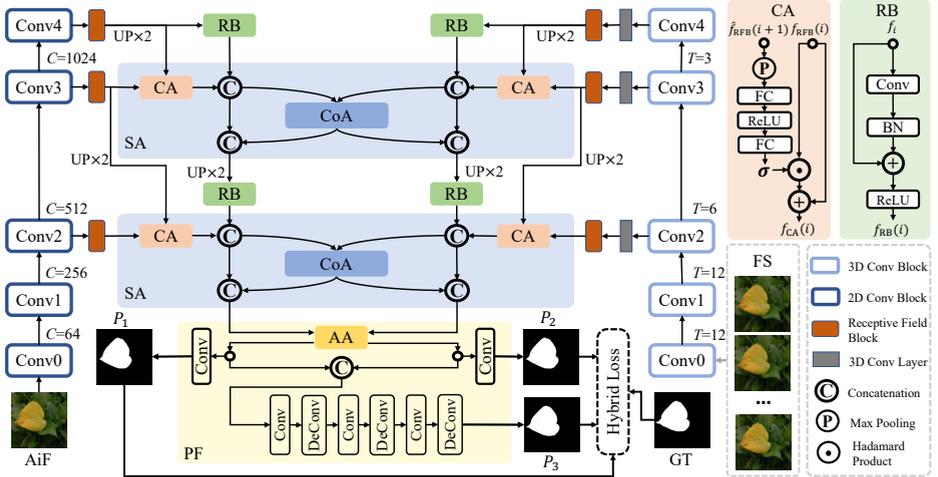


Figure 2: An overview of our *SA-Net*. Multi-modal multi-level features extracted from our multi-modal encoder are fed to two cascaded synergistic attention (SA) modules followed by a progressive fusion (PF) module. The short names in the figure are detailed as follows: CoA = co-attention component. CA = channel attention component. AA = AiF-induced attention component. RB = residual block. P_n = the n th saliency prediction. (De)Conv = (de-)convolutional layer. BN = batch normalization layer. FC = fully connected layer.

rich intra- and inter-slice information for accurate SOD. Note that the same setting ($T = 12$) as in [24] is adopted in our 3D branch, and a zero-padding strategy is applied to the FS with less than 12 focal slices.

3.2 Synergistic Attention Module

As high-level features tend to reserve the essential cues (e.g., location, shape) of salient objects while the low-level ones contain relative trivial information (e.g., edge) [24], our decoder only integrates high-level features to avoid redundant computational complexity. Specifically, we use $\{f_i^{2D}\}_{i=2}^4$ and $\{f_i^{3D}\}_{i=2}^4$ to denote the high-level AiF and FS features extracted from the 2D and 3D CNNs of our dual-branch backbone network (§ 3.1).

Multi-Level Attention. As shown in Figure 2, a receptive field block (RFB) [24] is first employed to enrich the global context information for each convolution block. Taking the AiF branch as an example, the adjacent high-level features from the encoder are then combined with a channel attention (CA) mechanism from [25], i.e.,

$$f_{CA}^{2D}(i) = \sigma(FC(ReLU(FC(P(\hat{f}_{RFB}^{2D}(i+1)))))) \odot f_{RFB}^{2D}(i) + f_{RFB}^{2D}(i), \quad (1)$$

where $f_{RFB}^{2D}(i)$ represents the i th level features provided by RFB; $\hat{f}_{RFB}^{2D}(i+1)$ is the up-sampled version of $f_{RFB}^{2D}(i+1)$; $\sigma(\cdot)$, $FC(\cdot)$, $P(\cdot)$, and \odot denotes the Sigmoid function, fully connected layer, max pooling, and Hadamard product, respectively. The resulting feature $f_{CA}^{2D}(i)$ is further concatenated with the upper level feature $f_{RB}^{2D}(i+1)$ provided by a residual block (RB) for the feature $f_{Cat}^{2D}(i)$, which is one of the pair-wise inputs ($\{f_{Cat}^{2D}(i), f_{Cat}^{3D}(i)\}$) for the second stage of our SA module. Note that the FS branch follows the consistent procedure as in AiF's since two branches are symmetric.

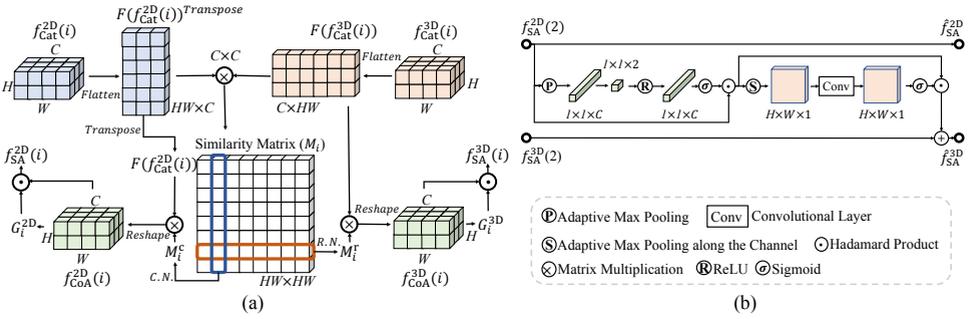


Figure 3: (a) The architecture of the co-attention (CoA) module. C.N. = column-wise normalization. R.N. = row-wise normalization. (b) The architecture of AiF-induced attention component.

Multi-Modal Attention. Inspired by a mutual attention mechanism [23] used for cross-frame feature fusion in the field of video object segmentation, the high-level feature interaction between the two modalities is conducted with two cascaded co-attention (CoA) modules (Figure 2). As shown in Figure 3 (a), given the pair-wise features $\{f_{\text{Cat}}^{2D}(i), f_{\text{Cat}}^{3D}(i)\}$ at i th layer as inputs, a similarity matrix M_i can be computed as:

$$M_i = F(f_{\text{Cat}}^{2D}(i))^T \otimes F(f_{\text{Cat}}^{3D}(i)), \quad (2)$$

where $F(\cdot)$ represents a flatten operation reshaping the 3D feature matrix $f_{\text{Cat}}^{2D}(i) \in \mathbb{R}^{H \times W \times C}$ to a 2D one with a dimension of $HW \times C$, \otimes denotes matrix multiplication. Note that we do not apply extra weight matrix as in [23] to compute M_i , since the CoA module aims at fusing the cross-modal features with equally assigned attention. The M_i is then column-/row-wisely normalized via:

$$\begin{aligned} M_i^c &= \text{Softmax}(M_i) \in [0, 1]^{HW \times HW}, \\ M_i^r &= \text{Softmax}(M_i^T) \in [0, 1]^{HW \times HW}, \end{aligned} \quad (3)$$

where $\text{Softmax}(\cdot)$ normalizes each column of the similarity matrix. Therefore, the co-attention-based pair-wise features $(\{f_{\text{CoA}}^{2D}(i), f_{\text{CoA}}^{3D}(i)\})$ at i th layer are further defined as:

$$\begin{aligned} f_{\text{CoA}}^{2D}(i) &= R(f_{\text{Cat}}^{2D}(i) \otimes M_i^c) \in [0, 1]^{H \times W \times C}, \\ f_{\text{CoA}}^{3D}(i) &= R(f_{\text{Cat}}^{3D}(i) \otimes M_i^r) \in [0, 1]^{H \times W \times C}, \end{aligned} \quad (4)$$

where $R(\cdot)$ reshapes the given matrix from a dimension of $C \times HW$ to $H \times W \times C$. A self-gate mechanism [23] is further employed to automatically learn the co-attention confidences (G_i^{2D}, G_i^{3D}) for $f_{\text{CoA}}^{2D}(i)$ and $f_{\text{CoA}}^{3D}(i)$. Therefore the final outputs $\{f_{\text{SA}}^{2D}(i), f_{\text{SA}}^{3D}(i)\}$ of our SA module at i th layer are computed as:

$$f_{\text{SA}}^{2D}(i) = G_i^{2D} \odot f_{\text{CoA}}^{2D}(i) \text{ and } f_{\text{SA}}^{3D}(i) = G_i^{3D} \odot f_{\text{CoA}}^{3D}(i), \quad (5)$$

where the co-attention confidence $G_i^{2D} = \sigma(\text{Conv}(f_{\text{CoA}}^{2D}(i)))$ with $\text{Conv}(\cdot)$ denoting a convolutional layer.

By combining the CA and CoA module, our SA module is particularly effective in exploiting the multi-level and multi-modal complementary information, which, therefore, provides significantly improved performance, as demonstrated by our ablation studies in § 4.3.

3.3 Progressive Fusion Module

To obtain the final prediction, we further add a *PF* module to gradually up-sample the selective high-level features provided by our *SA* module (Figure 2). Specifically, we first balance the FS and AiF features with an AiF-induced attention (AA) component (Figure 3 (b)) before the final fusion of the two modalities. The AA component follows the same procedure applied in RGB-D fusion [12], i.e., unifying the channel and spatial attention by computing:

$$\hat{f}_{SA}^{2D} = f_{SA}^{2D}(2) \text{ and } \hat{f}_{SA}^{3D} = SA(CA(f_{SA}^{2D}(2))) + f_{SA}^{3D}(2), \quad (6)$$

where $CA(\cdot)$ and $SA(\cdot)$ denote spatial and channel attention components, respectively. We then concatenate the balanced cross-modal features and feed them to a deconvolutional block for the final prediction P_3 , i.e.,

$$P_3 = DB(Cat(\hat{f}_{SA}^{2D}, \hat{f}_{SA}^{3D})), \quad (7)$$

where $Cat(\cdot)$ denotes the concatenation operation, and $DB(\cdot)$ represents a deconvolutional block consisting of three deconvolutional layers [12] and convolutional layers that are organized in a cascaded manner (Figure 2).

3.4 Loss Function

As shown in Figure 2, our model predicts three saliency maps: $\{P_n\}_{n=1}^3 \in [0, 1]$. Let $G \in \{0, 1\}$ denotes the ground-truth saliency map, we jointly optimize the three-way predictions by defining a hybrid loss ℓ :

$$\ell = \sum_{n=1}^N \ell_{BCE}(P_n, G) + \ell_{IoU}(P_n, G) + \ell_{EM}(P_n, G), \quad (8)$$

where ℓ_{BCE} and ℓ_{IoU} denote Binary Cross Entropy (BCE) and Intersection over Union (IoU) loss, respectively; the loss $\ell_{EM} = 1 - E_\phi$ with E_ϕ denoting E-Measure [13].

3.5 Implementation Details

Our *SA-Net* is implemented in PyTorch and optimized with Adam algorithm [14]. The backbone of *SA-Net* is based on a 2D standard ResNet50 for AiF images and an inflated 3D ResNet50 [9] for FSs. The 2D convolution layers in our backbone are initialized with ImageNet-pretrained ResNet50, while the 3D convolutional layers are initialized with a 2D weight transfer strategy [9]. During the training stage, the batch size is set to 2, the learning rate is initialized as 1e-5 and decreased by 10% when training loss reaches a flat. It takes about 14 hours to train the proposed model based on a platform consists of Intel[®] i9-7900X CPU@3.30GHz and one TITAN XP.

4 Experiments

4.1 Datasets and Evaluation Metrics

Datasets. We evaluate our *SA-Net* and 28 state-of-the-art SOD methods based on three widely-used light field datasets: DUT-LF, HFUT and LFSD, which all provide FS and semantic ground truth corresponding to each of the AiF images (see details in § 2.1). For fair

Types	Models	DUT-LF [49]				HFUT [49]				LFSD [49]			
		$F_\beta \uparrow$	$S_\alpha \uparrow$	$E_\phi \uparrow$	$M \downarrow$	$F_\beta \uparrow$	$S_\alpha \uparrow$	$E_\phi \uparrow$	$M \downarrow$	$F_\beta \uparrow$	$S_\alpha \uparrow$	$E_\phi \uparrow$	$M \downarrow$
4D	SA-Net	0.920	0.918	0.954	0.032	0.736	0.784	0.849	0.078	0.844	0.841	0.889	0.074
	ERNetT [49]	0.889	0.899	0.943	0.040	0.705	0.777	0.831	0.082	0.842	0.838	0.889	0.080
	ERNetS [49]	0.838	0.848	0.916	0.061	0.651	0.736	0.824	0.085	0.721	0.726	0.820	0.137
	DLFS [49]	0.801	0.841	0.891	0.076	0.615	0.741	0.783	0.098	0.715	0.737	0.806	0.147
	LFS* [49]	0.484	0.563	0.728	0.240	0.430	0.579	0.686	0.205	0.740	0.680	0.771	0.208
	MCA* [49]	-	-	-	-	-	-	-	-	0.815	0.749	0.841	0.150
	WSC* [49]	-	-	-	-	-	-	-	-	0.706	0.706	0.794	0.156
	DILF* [49]	0.641	0.705	0.805	0.168	0.555	0.695	0.736	0.131	0.728	0.755	0.810	0.168
3D	S2MA [49]	0.754	0.787	0.841	0.103	0.647	0.761	0.787	0.100	0.819	0.837	0.863	0.095
	D3Net [49]	0.790	0.822	0.869	0.084	0.692	0.778	0.827	0.080	0.804	0.825	0.853	0.095
	CPFF [49]	0.730	0.741	0.808	0.101	0.594	0.701	0.768	0.096	0.524	0.599	0.669	0.186
	TANet [8]	0.771	0.803	0.861	0.096	0.638	0.744	0.789	0.096	0.804	0.803	0.849	0.112
	MMCI [8]	0.750	0.785	0.853	0.116	0.645	0.741	0.787	0.104	0.796	0.799	0.848	0.128
	PDNet [49]	0.763	0.803	0.864	0.111	0.629	0.770	0.786	0.105	0.780	0.786	0.849	0.116
	PCA [8]	0.762	0.800	0.857	0.100	0.644	0.748	0.782	0.095	0.801	0.807	0.846	0.112
	CTMF [49]	0.790	0.823	0.881	0.100	0.620	0.752	0.784	0.103	0.791	0.801	0.856	0.119
DF [8a]	0.733	0.716	0.838	0.151	0.562	0.670	0.742	0.138	0.756	0.751	0.816	0.162	
2D	F3Net [49]	0.882	0.888	0.900	0.057	0.718	0.777	0.815	0.095	0.797	0.806	0.824	0.106
	GCPANet [49]	0.867	0.885	0.898	0.064	0.691	0.777	0.799	0.105	0.805	0.822	0.809	0.097
	EGNet [49]	0.870	0.886	0.914	0.053	0.672	0.772	0.794	0.094	0.762	0.784	0.776	0.118
	PoolNet [49]	0.868	0.889	0.919	0.051	0.683	0.776	0.802	0.092	0.769	0.800	0.786	0.118
	PAGRN [49]	0.828	0.822	0.878	0.084	0.635	0.717	0.773	0.114	0.725	0.727	0.805	0.147
	C2S [49]	0.791	0.844	0.874	0.084	0.650	0.763	0.786	0.111	0.749	0.806	0.820	0.113
	R ³ Net [8]	0.783	0.819	0.833	0.113	0.625	0.727	0.728	0.151	0.781	0.789	0.838	0.128
	Amulet [49]	0.805	0.847	0.882	0.083	0.636	0.767	0.760	0.110	0.757	0.773	0.821	0.135
	UCF [49]	0.769	0.837	0.850	0.107	0.623	0.754	0.764	0.130	0.710	0.762	0.776	0.169
	SRM [49]	0.832	0.848	0.899	0.072	0.672	0.762	0.801	0.096	0.827	0.826	0.863	0.099
NLDF [49]	0.778	0.786	0.862	0.103	0.636	0.729	0.807	0.091	0.748	0.745	0.810	0.138	
DSS [49]	0.728	0.764	0.827	0.128	0.626	0.715	0.778	0.133	0.644	0.677	0.749	0.190	

Table 1: Quantitative results for different models on three benchmark datasets. The best scores are in **boldface**. We train and test our *SA-Net* with the settings that are consistent with [49], which is the state-of-the-art model at present. \star indicates tradition methods. - denotes no available result. \uparrow indicates the higher the score the better, and vice versa for \downarrow .

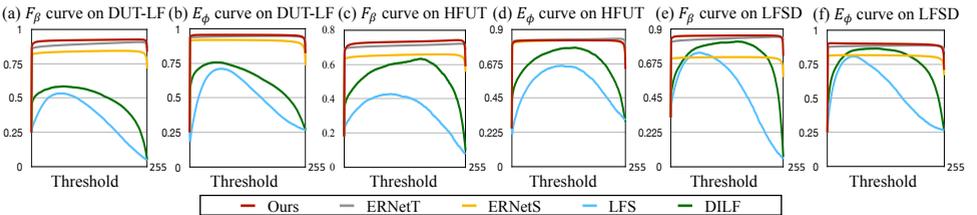


Figure 4: F-Measure (F_β) and E-Measure (E_ϕ) curves of cutting edge light field SOD models and our *SA-Net* upon three datasets.

comparison, we simply follow the settings of a top-ranking method, i.e., ERNet [49]. To be specific, 1000/100 AiF images of DUT-LF/HFUT are randomly selected as the training set, respectively, while the remains (462+155) and the whole LFSD are used for testing. Notably, as for competing methods, we report the results directly provided by authors or generated by officially released codes.

Metrics. We adopt the newly proposed S-Measure (S_α) [49] and E-Measure (E_ϕ) [49], also the generally agreed Mean Absolute Error (M) [49] and F-Measure (F_β) [49] as evaluation metrics for the quantitative comparison between benchmark models and *SA-Net*. Following the benchmark in [49], we report the adaptive F/E-Measure scores.

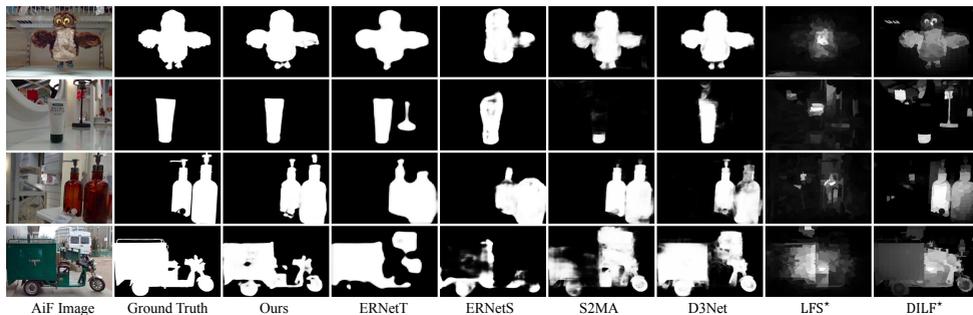


Figure 5: Qualitative comparison between our *SA-Net* and state-of-the-art light field SOD models. \star denotes traditional methods. Our *SA-Net* provides predictions closest to the ground truth on various aspects.

4.2 Comparison with State-of-the-Arts

We quantitatively compare our *SA-Net* with 12/9/7 state-of-the-art RGB/RGB-D/light field SOD methods, respectively. As shown in Table 1, our *SA-Net* outperforms all the state-of-the-art SOD models by a large margin in terms of all four evaluation metrics. We also perform a detailed comparison between our *SA-Net* and the top-ranking 4D SOD methods by using F/E-Measure curves. The results, shown in Figure 4, indicate the F/E-Measure curves of our *SA-Net* are higher than those of competing models. Furthermore, we show some of the predicted saliency maps in Figure 5. As can be observed, our *SA-Net* provides the saliency maps closest to the ground truth on various aspects, e.g., correct localization, intact object structure and clear details. It is worth nothing that our *SA-Net* trained on DUT-LF [57] and HFUT [44] also achieves promising performance on the unseen dataset, i.e., LFSD [47], indicating its superior generalization ability and robustness. Theoretically, the robustness of *SA-Net* owes to the synergistic attention mechanism. In practice, attention mechanisms can improve network robustness [24, 41, 51] since they emphasize the most informative features and reduce the disturbance of noisy features. Our *SA* module employs both channel attention and co-attention for better feature representation, which can also improve the robustness of our model.

4.3 Ablation Studies

To verify the effectiveness of each proposed module of our *SA-Net*, we conduct thorough ablation studies by gradually adding key components. We first construct a baseline “B”, which extracts AiF and FS features with two 2D ResNet50 backbones, simply concatenates, and up-samples the pair-wise high-level features for SOD.

Effectiveness of multi-modal encoder. To investigate the effectiveness of our multi-modal encoder, we construct the second ablated version “ME”, which is similar to “B”, but using a 3D backbone to extract FS features, consistent with our multi-modal encoder (§ 3.1). The results, shown in Table 2, indicate that “ME” outperforms “B” in terms of all evaluations, demonstrating the effectiveness of our 3D CNN-based encoder. Besides, to confirm the effectiveness of RFB for multi-level feature refinement, we also construct “ME0” without using the RFB, when compared to “ME”. The result (Table 2) shows that RFB benefits significantly to the task.

Metric	B	ME0	ME	SA1	SA2	PF1	PF2	F-SA	SA-Net	
DUT-LF	F_{β} \uparrow	0.871	0.874	0.881	0.890	0.899	0.912	0.919	0.913	0.920
	M \downarrow	0.051	0.051	0.048	0.041	0.037	0.035	0.034	0.037	0.032
LFSD	F_{β} \uparrow	0.811	0.825	0.835	0.836	0.835	0.838	0.839	0.844	0.845
	M \downarrow	0.095	0.089	0.080	0.079	0.077	0.075	0.075	0.078	0.074

Table 2: Quantitative results for the ablation studies of SA-Net on DUT-LF [57] and LFSD [4]. The best scores are in **boldface**. \uparrow indicates the higher the score the better, and vice versa for \downarrow .

importance of synergistic attention for learning the complementarities of multi-modal features. Besides, we compare SA-Net with F-SA (Figure 6), which consists of full four SA modules that fuse both the high and low level features for light field SOD tasks. An interesting finding is that an increase of parameters (about 1.3 million of increment) focusing on low level features do not contribute to performance improvement (Table 2), which is also consistent with the conclusion in [47].

Effectiveness of PF module. Compared with “SA2”, “PF1” uses the deconvolutional block (Figure 2) to gradually up-sample the features for predicting the saliency map. Besides, a three-way supervision (“PF2”) is further employed to provide a deep supervision for the training. Finally, with the AA component (§ 3.3), our SA-Net achieves the best performance (Table 2) and provides the saliency maps closest to ground truth (Figure 7).

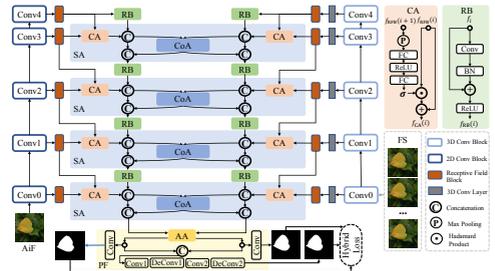


Figure 6: An overview of F-SA with four SA modules. CoA = co-attention component. CA = channel attention component. AA = AiF-induced attention component. RB = residual block.

5 Conclusion

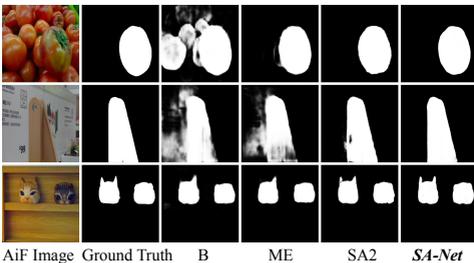


Figure 7: Representative visual results of ablation studies.

Extensive qualitative and quantitative experimental results on three light field datasets demonstrate the better performance of our SA-Net compared to 28 competing models.

Effectiveness of SA module. To investigate the effectiveness of our SA module, we further construct “SA1” and “SA2”, which incorporate the SA into “ME” without and with CoA, respectively. As shown in Table 2, both “SA1” and “SA2” improve the performance in comparison with “ME”. In particular, the full version of SA (“SA2”) provides a significant improvement compared to “ME”, indicating the

In this paper, we propose SA-Net, which addresses the light field SOD by learning the synergistic attention for AiF and FS features. The innovative attributes of our SA-Net are three-fold: (i) it exploits the cross-modal complementary information by establishing a synergistic effect between multi-modal features, (ii) it is the first attempt to learn both the spatial and interslice features of FSs with 3D CNNs, and (iii) it predicts the saliency map with an effective fusion model in a progressive manner.

6 Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grants 61771397 and 62171377.

References

- [1] Aseem Agarwala, Mira Dontcheva, Maneesh Agrawala, Steven Drucker, Alex Colburn, Brian Curless, David Salesin, and Michael Cohen. Interactive digital photomontage. In *SIGGRAPH*, 2004.
- [2] Ali Borji, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. Salient object detection: A benchmark. *TIP*, 2015.
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [4] Hao Chen and Youfu Li. Progressively complementarity-aware fusion network for rgb-d salient object detection. In *CVPR*, 2018.
- [5] Hao Chen and Youfu Li. Three-stream attention-aware network for RGB-D salient object detection. *TIP*, 2019.
- [6] Hao Chen, Youfu Li, and Dan Su. Multi-modal fusion network with multi-scale multi-path and cross-modal interactions for RGB-D salient object detection. *Pattern Recognition*, 2019.
- [7] Qian Chen, Ze Liu, Yi Zhang, Keren Fu, Qijun Zhao, and Hongwei Du. Rgb-d salient object detection via 3d convolutional neural. *AAAI*, 2021.
- [8] Zijun Deng, Xiaowei Hu, Lei Zhu, Xuemiao Xu, Jing Qin, Guoqiang Han, and Pheng-Ann Heng. R3net: Recurrent residual refinement network for saliency detection. In *IJCAI*, 2018.
- [9] Deng-Ping Fan, Ming-Ming Cheng, Yun Liu, Tao Li, and Ali Borji. Structure-measure: A new way to evaluate foreground maps. In *ICCV*, 2017.
- [10] Deng-Ping Fan, Cheng Gong, Yang Cao, Bo Ren, Ming-Ming Cheng, and Ali Borji. Enhanced-alignment measure for binary foreground map evaluation. In *IJCAI*, 2018.
- [11] Deng-Ping Fan, Zheng Lin, Zhao Zhang, Menglong Zhu, and Ming-Ming Cheng. Re-thinking rgb-d salient object detection: Models, data sets, and large-scale benchmarks. *TNNLS*, 2020.
- [12] Deng-Ping Fan, Yingjie Zhai, Ali Borji, Jufeng Yang, and Ling Shao. Bbs-net: Rgb-d salient object detection with a bifurcated backbone strategy network. In *ECCV*, 2020.
- [13] Junwei Han, Hao Chen, Nian Liu, Chenggang Yan, and Xuelong Li. CNNs-based RGB-D saliency detection via cross-view transfer and multiview fusion. *IEEE Transactions on Cybernetics*, 2017.

- [14] Qibin Hou, Ming-Ming Cheng, Xiaowei Hu, Ali Borji, Zhuowen Tu, and Philip HS Torr. Deeply supervised salient object detection with short connections. In *CVPR*, 2017.
- [15] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [17] Nianyi Li, Jinwei Ye, Yu Ji, Haibin Ling, and Jingyi Yu. Saliency detection on light field. In *CVPR*, 2014.
- [18] Nianyi Li, Bilin Sun, and Jingyi Yu. A weighted sparse coding framework for saliency detection. In *CVPR*, 2015.
- [19] Nianyi Li, Jinwei Ye, Yu Ji, Haibin Ling, and Jingyi Yu. Saliency detection on light field. *TPAMI*, 2017.
- [20] Xin Li, Fan Yang, Hong Cheng, Wei Liu, and Dinggang Shen. Contour knowledge transfer for salient object detection. In *ECCV*, 2018.
- [21] Jiang-Jiang Liu, Qibin Hou, Ming-Ming Cheng, Jiashi Feng, and Jianmin Jiang. A simple pooling-based design for real-time salient object detection. In *ICCV*, 2019.
- [22] Nian Liu, Ni Zhang, and Junwei Han. Learning selective self-mutual attention for rgb-d saliency detection. In *CVPR*, 2020.
- [23] Xiankai Lu, Wenguan Wang, Jianbing Shen, David Crandall, and Jiebo Luo. Zero-shot video object segmentation with co-attention siamese networks. *TPAMI*, 2020.
- [24] Zhiming Luo, Akshaya Mishra, Andrew Achkar, Justin Eichel, Shaozi Li, and Pierre-Marc Jodoin. Non-local deep features for salient object detection. In *CVPR*, 2017.
- [25] Sabarinath Mahadevan, Ali Athar, Aljosa Osep, Sebastian Hennen, Laura Leal-Taixé, and Bastian Leibe. Making a case for 3d convolutions for object segmentation in videos. In *BMVC*, 2020.
- [26] Federico Perazzi, Philipp Krähenbühl, Yael Pritch, and Alexander Hornung. Saliency filters: Contrast based filtering for salient region detection. In *CVPR*, 2012.
- [27] Yongri Piao, Zhengkun Rong, Miao Zhang, Xin Li, and Huchuan Lu. Deep light-field-driven saliency detection from a single view. In *IJCAI*, 2019.
- [28] Yongri Piao, Xiao Li, Miao Zhang, Jingyi Yu, and Huchuan Lu. Saliency detection via depth-induced cellular automata on light field. *TIP*, 2020.
- [29] Yongri Piao, Zhengkun Rong, Miao Zhang, and Huchuan Lu. Exploit and replace: an asymmetrical two-stream architecture for versatile light field saliency detection. In *AAAI*, 2020.
- [30] Liangqiong Qu, Shengfeng He, Jiawei Zhang, Jiandong Tian, Yandong Tang, and Qingxiong Yang. Rgb-d salient object detection via deep fusion. *TIP*, 2017.

- [31] Hao Sheng, Shuo Zhang, Xiaoyu Liu, and Zhang Xiong. Relative location for light field saliency detection. In *ICASSP*, 2016.
- [32] Xingjian Shi, Zhouong Chen, Hao Wang, Dityan Yeung, Waikin Wong, and Wangchun Woo. Convolutional lstm network: a machine learning approach for precipitation now-casting. In *NeurIPS*, 2015.
- [33] Anzhi Wang, Minghui Wang, Xiaoyan Li, Zetian Mi, and Huan Zhou. A two-stage bayesian integration framework for salient object detection on light field. *Neural Processing Letters*, 2017.
- [34] Haoqian Wang, Bing Yan, Xingzheng Wang, Yongbing Zhang, and Yi Yang. Accurate saliency detection based on depth feature of 3d images. *Multimedia Tools and Applications*, 2018.
- [35] Shizheng Wang, Wenjuan Liao, Phil Surman, Zhigang Tu, Yuanjin Zheng, and Junsong Yuan. Saliency guided depth calibration for perceptually optimized compressive light field 3d display. In *CVPR*, 2018.
- [36] Tiantian Wang, Ali Borji, Lihe Zhang, Pingping Zhang, and Huchuan Lu. A stagewise refinement model for detecting salient objects in images. In *ICCV*, 2017.
- [37] Tiantian Wang, Yongri Piao, Huchuan Lu, Xue chun Li, and Li Zhang. Deep learning for light field saliency detection. In *ICCV*, 2019.
- [38] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [39] Xue Wang, Yingying Dong, Qicheng Zhang, and Qiaoyu Wang. Region-based depth feature descriptor for saliency detection on light field. *Multimedia Tools and Applications*, 2020.
- [40] Jun Wei, Shuhui Wang, and Qingming Huang. F³net: Fusion, feedback and focus for salient object detection. In *AAAI*, 2020.
- [41] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *ECCV*, 2018.
- [42] Zhe Wu, Li Su, and Qingming Huang. Cascaded partial decoder for fast and accurate salient object detection. In *CVPR*, 2019.
- [43] Jun Zhang, Meng Wang, Jun Gao, Yi Wang, Xudong Zhang, and Xindong Wu. Saliency detection with a deeper investigation of light field. In *IJCAI*, 2015.
- [44] Jun Zhang, Meng Wang, Liang Lin, Xun Yang, Jun Gao, and Yong Rui. Saliency detection on light field: A multi-cue approach. *TOMM*, 2017.
- [45] Jun Zhang, Yamei Liu, Shengping Zhang, Ronald Poppe, and Meng Wang. Light field saliency detection with deep convolutional networks. *TIP*, 2020.
- [46] Miao Zhang, Jingjing Li, Ji Wei, Yongri Piao, and Huchuan Lu. Memory-oriented decoder for light field salient object detection. In *NeurIPS*, 2019.

- [47] Miao Zhang, Wei Ji, Yongri Piao, Jishun Li, Yu Zhang, Shuang Xu, and Huchuan Lu. Lfnet: Light field fusion network for salient object detection. *TIP*, 2020.
- [48] Pingping Zhang, Dong Wang, Huchuan Lu, Hongyu Wang, and Xiang Ruan. Amulet: Aggregating multi-level convolutional features for salient object detection. In *ICCV*, 2017.
- [49] Pingping Zhang, Dong Wang, Huchuan Lu, Hongyu Wang, and Baocai Yin. Learning uncertain convolutional features for accurate saliency detection. In *ICCV*, 2017.
- [50] Qiudan Zhang, Shiqi Wang, Xu Wang, Zhenhao Sun, Sam Kwong, and Jianmin Jiang. A multi-task collaborative network for light field salient object detection. *TCSVT*, 2020.
- [51] Xiaoning Zhang, Tiantian Wang, Jinqing Qi, Huchuan Lu, and Gang Wang. Progressive attention guided recurrent network for salient object detection. In *CVPR*, 2018.
- [52] Jia-Xing Zhao, Yang Cao, Deng-Ping Fan, Ming-Ming Cheng, Xuan-Yi Li, and Le Zhang. Contrast prior and fluid pyramid integration for rgb-d salient object detection. In *CVPR*, 2019.
- [53] Jia-Xing Zhao, Jiang-Jiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. Egnnet: Edge guidance network for salient object detection. In *ICCV*, 2019.
- [54] Chunbiao Zhu, Xing Cai, Kan Huang, Thomas H Li, and Ge Li. PDNet: Prior-model guided depth-enhanced network for salient object detection. In *ICME*, 2019.
- [55] Chen Zuyao, Xu Qianqian, Cong Runmin, and Huang Qingming. Global context-aware progressive aggregation network for salient object detection. In *AAAI*, 2020.