

Specificity-preserving RGB-D Saliency Detection

Tao Zhou^{1,2}, Deng-Ping Fan³ (✉), Geng Chen⁴, Yi Zhou⁵, and Huazhu Fu⁶

© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract Salient object detection (SOD) in RGB and depth images has attracted increasing research interest. Existing RGB-D SOD models usually adopt fusion strategies to learn a shared representation from RGB and depth modalities, while few methods explicitly consider how to preserve modality-specific characteristics. In this study, we propose a novel framework, the Specificity-preserving Network (SPNet), which improves SOD performance by exploring both the shared information and modality-specific properties. Specifically, we use two modality-specific networks and a shared learning network to generate individual and shared saliency prediction maps. To effectively fuse cross-modal features in the shared learning network, we propose a cross-enhanced integration module (CIM) and propagate the fused feature to the next layer to integrate cross-level information. Moreover, to capture rich complementary multi-modal information to boost SOD performance, we use a multi-modal feature aggregation (MFA) module to integrate the modality-specific features from each individual decoder into the shared decoder. By using skip connections between encoder and decoder layers, hierarchical features can be fully combined. Extensive experiments demonstrate that our SPNet outperforms cutting-edge approaches on six popular RGB-D SOD and three camouflaged object

detection benchmarks. The project is publicly available at: <https://github.com/taozh2017/SPNet>.

Keywords Salient Object Detection, RGB-D, Cross-enhanced Integration Module, Multi-modal Feature Aggregation.

1 Introduction

Salient object detection (SOD, also called saliency detection) aims to emulate the mechanisms of human visual attention and locate the most visually distinctive object(s) in a given scene [61]. SOD has been widely applied in various vision-related tasks, such as image understanding [108], action recognition [67, 71], video/semantic segmentation [71, 77], and person re-identification [95]. Although significant progress has been made, it is still challenging to accurately locate salient objects in many challenging scenarios, such as images with cluttered backgrounds, low-contrast lighting conditions, and salient object(s) having a similar appearance to the background. Recently, with the ready availability of depth sensors in smart devices, depth maps have been introduced to provide geometric and spatial information to improve SOD performance. Consequently, fusing RGB and depth images has gained increasing interest in the SOD community [5, 21, 24, 42, 52, 86, 87, 99, 101], and it is a challenging task to adaptively fuse RGB and depth modalities.

Over past years, various RGB-D SOD methods have been proposed; they often focus on how to effectively fuse RGB and depth images. Existing fusion strategies can be divided into categories using early fusion, late fusion, and intermediate fusion. The *early fusion* strategy often adopts a simple concatenation to integrate the two modalities. For example, methods in [53, 61, 68, 72, 79] directly integrate RGB and depth images to form four-channel input. However, this type of fusion does not consider the distribution gap

1 School of Computer Science and Engineering, Nanjing University of Science and Technology, China.

2 Key Laboratory of System Control and Information Processing, Ministry of Education, Shanghai, China.

3 CVL, ETH Zurich, Switzerland. (dengpfan@gmail.com)

4 School of Computer Science and Engineering, Northwestern Polytechnical University, Xi'an, China.

5 School of Computer Science and Engineering, Southeast University, Nanjing, China.

6 Inception Institute of Artificial Intelligence, UAE.

Manuscript received: 2021-10-19; accepted: 2022-01-04.

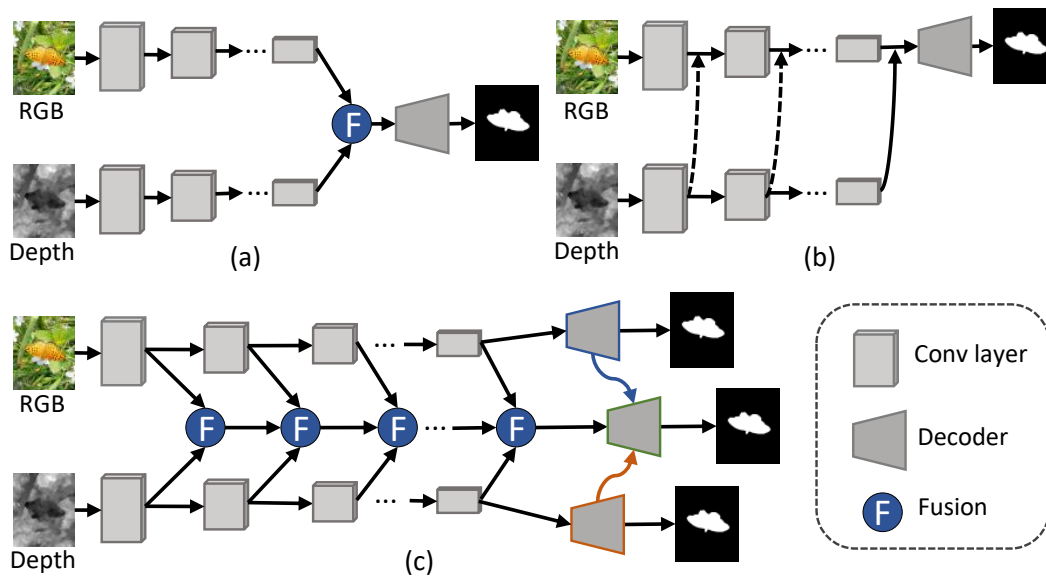


Fig. 1 Comparison of two existing RGB-D salient object detection frameworks and our proposed model. (a) RGB and depth images are fed into two independent network streams, and then fused high-level features are fed into a decoder to predict saliency maps (e.g., [3, 4, 28, 48]). (b) Depth features are integrated into the RGB network using an auxiliary subnetwork (e.g., [6, 22, 84, 94, 105]). (c) Our method adopts two modality-specific networks and a shared learning network to explicitly explore modality-specific characteristics and shared information. Features learned from the modality-specific decoders are integrated into the shared decoder to boost SOD performance.

between the two modalities, which could result in an inaccurate feature fusion. The *late fusion* strategy uses two parallel network streams to generate independent saliency maps for RGB and depth data, which are fused to obtain a final prediction map [16, 27, 75]. However, it is still challenging to capture the complex interactions between the two modalities.

Recent research mainly focuses on *intermediate fusion*, which utilizes two independent networks to learn intermediate features of the two modalities separately, and then the fused features are fed into a subsequent network or decoder (see Fig. 1(a)). Other methods carry out cross-modal fusion at multiple scales [3, 4, 7, 28, 32, 33, 48]. As a result, complex correlations can be effectively exploited from the two modalities. Further methods utilize depth information to enhance RGB features via an auxiliary subnetwork [6, 94, 105] (see Fig. 1(b)). For example, Zhao *et al.* [94] introduced a contrast prior into a CNN-based architecture to enhance the depth information, which was then integrated with RGB features using a fluid pyramid integration module. Zhu *et al.* [105] utilized an independent subnetwork to extract depth-based features, which were then incorporated into the RGB network. The above methods focus on learning shared representations by fusing them and then use a decoder to generate the final saliency

map. Furthermore, there is no supervised decoder to guide the depth-based feature learning [94, 105], which may prevent optimal depth features from being obtained. From a multi-modal learning perspective, several works [31, 54, 102, 104] have shown that exploring both the shared information and modality-specific characteristics can improve model performance. However, in the RGB-D SOD community, few methods explicitly exploit modality-specific characteristics.

Thus, in this paper, we propose a novel RGB-D SOD framework, the *Specificity-preserving Network* (SPNet), which can effectively explore the shared information as well as capture modality-specific characteristics to improve the SOD performance. Two encoder subnetworks are used to extract multi-scale features for the two modalities (*i.e.* RGB and depth), and a cross-enhanced integration module (CIM) is proposed to integrate cross-modal features in different feature layers. Then, we use a simple U-Net [69] structure to construct a modality-specific decoder, in which skip connections between the encoder and decoder layers are used to combine hierarchical features. In this way, we can learn powerful modality-specific features in each independent decoder, which also captures modality-specific characteristics to provide cross-modal complementarity. Further, we construct a shared decoder to combine hierarchical

features from outputs of the previous CIM via a skip connection. To make full use of the modality-specific features, a multi-modal feature aggregation module (MFA) is proposed to integrate them into the shared decoder. Finally, we formulate a unified and end-to-end trainable framework where shared and modality-specific information are simultaneously exploited to boost SOD performance.

The main contributions of our paper in summary are:

- A novel RGB-D salient object detection framework, the Specificity-preserving Network (SPNet), which explores shared information from RGB and depth images as well as preserving modality-specific characteristics.
- A cross-enhanced integration module (CIM) to integrate cross-modal features and learn shared representations for the two modalities. The output of each CIM is propagated to the next layer to explore rich cross-level information.
- An effective multi-modal feature aggregation (MFA) module to integrate learned modality-specific features. It allows our model to make full use of the features learned in the modality-specific decoder to improve salient object detection.
- Extensive experiments on six public RGB-D SOD and three camouflaged object detection (COD) datasets demonstrate the superiority of our model over other cutting-edge methods. Moreover, we carry out an attribute-based evaluation on various state-of-the-art RGB-D SOD methods under varying conditions (*e.g.* number of salient objects, indoors or outdoors, lighting, and object scale), which has not been done previously.

This paper significantly extends our previous work in [103], as follows.

- We discuss differences between (i) our proposed CIM and existing fusion strategies, and (ii) the proposed CIM and MFA.
- We provide further details, including (i) a review of existing RGB SOD methods. (ii) a discussion of the importance of integrating multi-level/scale features, (iii) better characterisation of our evaluation metrics.
- We provide an additional ablation study and attribute-based evaluation, to validate the effectiveness of the shared decoder, and to examine the effects of different numbers of CIMs. We also show that our model can effectively handle variations in object scale.
- We apply SPNet to a new RGB-D task: COD, and demonstrate its superiority over existing methods.

2 Related Work

In this section, we review three types of works most related to the proposed model, *i.e.* RGB salient object detection, RGB-D salient object detection, and multi-modal learning.

2.1 RGB Salient Object Detection

Early salient object detection methods were based on hand-crafted features and various saliency priors, such as a background prior [109], color contrast [1], a compactness prior [100], and a center prior [36]. However, the generalizability and effectiveness of these traditional methods are limited. With the breakthrough of deep learning in the field of computer vision, various deep learning-based salient object detection methods have been developed with promising results. For example, Hou *et al.* [30] proposed a novel salient object detection method by introducing short connections to the skip-layer structures within the holistically-nested edge detector architecture. Wang *et al.* [74] proposed a recurrent fully convolutional network framework for salient object detection with promising results. Liu *et al.* [50] proposed to hierarchically embed global and local context modules into the top-down pathway, which can generate attention over context regions for each pixel. Deng *et al.* [13] proposed a recurrent residual refinement network with residual refinement blocks to accurately detect salient objects. Further methods can be found in a survey [76]. Scale variation is a key challenge in the SOD task, so several methods have been proposed to integrate multi-level or scale features [60, 78, 89, 91] to improve SOD results. In our method, we consider how to effectively combine cross-modal (RGB and depth) features, and how multi-level information can be exploited via a cross-enhanced integration module.

2.2 RGB-D Salient Object Detection

Early RGB-D based SOD methods often extracted hand-crafted features from the input RGB-D data. For example, Lang *et al.* [38] in the first RGB-D SOD work utilized Gaussian mixture models to model the distribution of depth-induced saliency. Subsequently, several methods explored different principles, such as center-surround difference [27, 37], contrast [14, 61, 68], a center/boundary prior [46, 107], and background enclosure [23]. However, these methods typically provide poor results due to the limited expressivity of handcrafted features. Benefiting from the rapid development of deep convolutional neural networks (CNNs), several deep learning-based works [21, 63, 66,

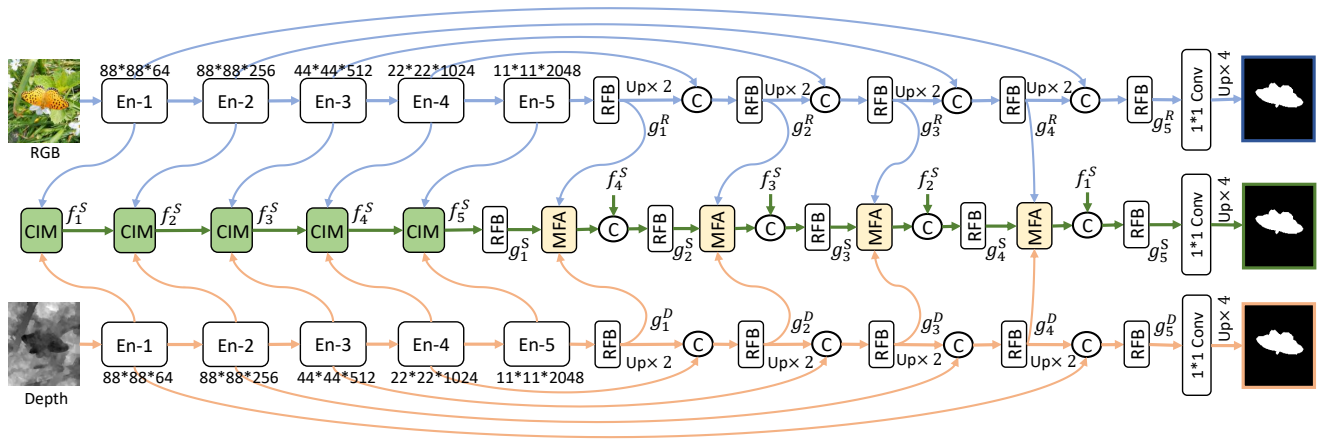


Fig. 2 Architecture of SPNet, consisting of two modality-specific learning networks and a shared learning network. The former preserve individual properties for RGB or depth, while the shared network fuses cross-modal features and explores complementary information. Skip connections combine hierarchical features between encoder and decoder layers. Learned features from the modality-specific decoder are integrated into the shared decoder to provide rich multi-modal complementary information, boosting saliency detection. C denotes feature concatenation.

[86, 94] have recently obtained promising results. For example, Qu *et al.* [66] used a CNN model to fuse saliency cues from different low levels into hierarchical features to boost SOD abilities. Chen *et al.* [3] proposed a complementarity-aware fusion module to effectively integrate cross-modal and cross-level features for RGB and depth modalities. Piao *et al.* [63] proposed a depth-induced multi-scale recurrent attention network to enhance cross-modality feature fusion. Fan *et al.* [21] designed a depth purification unit to remove some low-quality depth maps. Most other models [4, 7, 28, 41, 44, 48] employ cross-modal fusion at multiple scales using different integration strategies.

2.3 Multi-modal Learning

Recently, multi-modal (or multi-view) learning has attracted increasing attention: much data can be collected from multiple sources or represented using different types of features. One traditional strategy directly concatenates feature vectors from such multi-modal data into a feature vector. However, this may fail to exploit the complex correlations within multi-modal data. Thus, several multi-modal learning methods have been developed to explicitly fuse the complementary information from different modalities to improve model results. These popular methods can be divided into three types (i) co-training [2, 15] tries to minimize the disagreement between different modalities, (ii) multiple kernel learning [26] utilizes a predefined set of kernels for multiple modalities and integrates these modalities using the learned kernel weights, and (iii) subspace learning [81, 85] assumes that a latent subspace exists shared by different modalities, with one underlying

latent representation. To effectively fuse multi-modal data, several deep learning-based models have also been explored. For example, Ngiam *et al.* [56] proposed to learn a shared representation from audio and video inputs. Eitel *et al.* [17] adopted two separate CNN streams for RGB and depth, combining them using a late fusion network for RGB-D object recognition. Hu *et al.* [31] presented a shared and individual multi-view learning algorithm to explore further properties of multi-modal data. Lu *et al.* [54] presented a shared-specific feature transfer framework to perform a cross-modal person ReID task.

3 Methodology

In this section, we first present the overall SPNet. Then we describe the two key components in our model the modality-specific learning network and shared learning network, and finally provide the overall loss function.

3.1 Overview

Fig. 2 shows the framework of our proposed specificity-preserving network for RGB-D SOD. First, RGB and depth images are fed into two stream modality-specific learning networks to obtain their multi-level feature representations, and a CIM learns their shared feature representation. Secondly, the individual and shared decoder subnetworks are each utilized to generate saliency prediction maps. The original features from the encoder networks are integrated into the decoder via skip connections. Finally, to make full use of the features learned by using the modality-specific decoder, an MFA module




Fig. 3 Cross-enhanced integration module (CIM). C, +, \times , and M denote feature concatenation, element-wise addition, multiplication, and maximization, respectively.

integrates these features into the shared decoder. We detail each key part below.

3.2 Modality-specific Learning Network

As Fig. 2 shows, the modality-specific subnetwork is built upon Res2Net-50 [25], pretrained on the ImageNet [70] dataset. Thus, there are five multi-level features, *i.e.* $F^R = [f_m^R]$ and $F^D = [f_m^D]$, $m = 1, \dots, 5$, in the modality-specific encoder subnetworks for RGB and depth respectively. The input resolution of the modality-specific encoder subnetwork is $W \times H$. Thus, we have a feature resolution of $(H/8) \times (W/8)$ for the first layer, and a general resolution of $(H/2^m) \times (W/2^m)$ (for $m > 1$). The number of channel features in the m -th layer is denoted C_m , where $C_m = [64, 256, 512, 1024, 2048]$.

After obtaining the high-level features f_5^R and f_5^D , they are then fed into the modality-specific decoder subnetworks to generate individual saliency maps. We further utilize a U-Net [69] structure to construct the modality-specific decoder, where the skip connections between encoder and decoder layers are used to combine hierarchical features. Moreover, the concatenated features (only f_5^R and f_5^D in the first stage of the decoder subnetwork) are fed to the receptive field block (RFB) [82] to capture global context information. This modality-specific learning network enables us to learn effective and powerful individual features for each modality by retaining its specific properties. These features are then integrated into the shared decoder subnetwork to improve saliency detection.

3.3 Shared Learning Network

3.3.1 Structure

As Fig. 2 shows, in the shared learning network, we fuse the cross-modal features from the RGB and depth modalities to learn their shared representation, which is fed into the shared decoder to generate the




Fig. 4 Multi-modal feature aggregation module (MFA). C, +, and \times denote feature concatenation, element-wise addition, and element-wise multiplication respectively.

final saliency map. We again adopt skip connections between the encoder and decoder layers to combine hierarchical features. Moreover, to make full use of the features learned by the modality-specific decoder, we integrate them into the shared decoder to improve saliency detection.

3.3.2 Cross-enhanced Integration Module

Our CIM is used to effectively fuse cross-modal features. Let the width, height, and number of channels for the m -th layer be denoted W_m , H_m , and C_m . Taking $f_m^R \in \mathbb{R}^{W_m \times H_m \times C_m}$ and $f_m^D \in \mathbb{R}^{W_m \times H_m \times C_m}$ as an example, we use a 1×1 convolutional layer to reduce the number of channels channel to $C_m/2$ for speed. The CIM has two parts, for cross-modal feature enhancement and adaptive feature fusion. First, we use a cross-enhanced strategy to exploit correlations between the two modalities by learning their enhanced features. Specifically, as shown in Fig. 3, the two features are fed into a 3×3 convolutional layer with a sigmoid activation function to obtain the normalized feature maps, $w_m^R = \sigma(\text{Conv}_3(f_m^R)) \in [0, 1]$ and $w_m^D = \sigma(\text{Conv}_3(f_m^D)) \in [0, 1]$, where σ is the logistic sigmoid activation function. To exploit correlations between the two modalities, the normalized feature maps can be regarded as feature-level attention maps to adaptively enhance the feature representation. In this way, the feature map from one modality can be used to enhance the other modality. To preserve the original information of each modality, a residual connection is used to combine the enhanced features with the original features. Thus, the cross-enhanced feature representations for the two modalities are as follows:

$$\begin{aligned} f_m^{R'} &= f_m^R + f_m^R \otimes w_m^D, \\ f_m^{D'} &= f_m^D + f_m^D \otimes w_m^R, \end{aligned} \quad (1)$$

where \otimes denotes element-wise multiplication.

Having obtained the cross-enhanced feature representations $f_m^{R'}$ and $f_m^{D'}$, the critical task is to effectively fuse them. Various strategies can be used to fuse features from different modalities, including

element-wise multiplication and maximization. However, it is unclear which is best for specific tasks. In order to benefit from the advantages of different strategies, we apply element-wise multiplication and maximization, and concatenate the results. Specifically, the two features $f_m^{R'}$ and $f_m^{D'}$ are first fed into a 3×3 convolutional layer to obtain smooth representations, and then we carry out element-wise multiplication and maximization, giving:

$$\begin{aligned} p_{\text{mul}} &= \text{BConv}_3(f_m^{R'}) \otimes \text{BConv}_3(f_m^{D'}), \\ p_{\text{max}} &= \max(\text{BConv}_3(f_m^{R'}), \text{BConv}_3(f_m^{D'})), \end{aligned} \quad (2)$$

where $\text{BConv}(\cdot)$ is a sequential operation that applies 3×3 convolution followed by batch normalization, then a ReLU function. Then, we concatenate the results as $p_{\text{cat}} = [p_{\text{mul}}, p_{\text{max}}] \in \mathbb{R}^{W_m \times H_m \times C_m}$, and obtain $p_{\text{cat}}^1 = \text{BConv}_3(p_{\text{cat}})$ through a BConv_3 operation to adaptively weigh the two parts. Further, the output p_{cat}^1 is concatenated with the previous output f_{m-1}^S of the $(m-1)$ -th CIM, and fed into the second BConv_3 operation. Finally, we obtain the output f_m^S of the m -th CIM. Note that, when $m = 1$, we do not need to use a 1×1 convolutional layer to reduce the number of channels. Furthermore, there is no previous output f_{m-1}^S when $m = 1$, so we only feed the concatenated features into a BConv_3 operation.

We note that our CIM can effectively exploit correlations between the two modalities via cross-enhanced feature learning, and fuse them by adaptively weighting the different feature representations. The fused feature representation f_m^S is propagated to the next layer to capture and integrate cross-level information. Some works [53, 61, 72] directly integrate RGB images and depth maps to form four-channel input (cascaded operation), and other methods carry out cross-modal fusion strategies, *e.g.* using attention-based fusion modules [4, 7], fusion-refinement modules (*e.g.* using summation) [48], etc. Unlike these methods, our proposed CIM mainly exploits the correlation between RGB and depth images, and then adaptively integrates enhanced cross-modal features to obtain a fused feature representation.

3.3.3 Multi-modal Feature Aggregation

To make full use of the features learned in the modality-specific decoder, we propose a simple but effective MFA module to integrate them into the shared decoder. Specifically, in the m -th layer of the shared decoder, we have the shared representation g_m^S , and the learned features g_m^R and g_m^D in the modality-specific decoder. As Fig. 4 shows, two features g_m^R and g_m^D are multiplied by the shared features of the current

layer: $g_m^{RS} = g_m^S \otimes g_m^R$ and $g_m^{DS} = g_m^S \otimes g_m^D$. The two features are further concatenated ($[g_m^{DR}, g_m^{DS}]$) and then fed into a $\text{BConv}(\cdot)$ operation to obtain g_m^{Sc} . Finally, we obtain the output of the MFA module to combine the convolutional feature g_m^{Sc} with the original feature g_m^S via an addition operation.

In the MFA, the learned modality-specific features are used to enhance the shared features and provide rich and complementary cross-modal information. Specifically, we use the two modality-specific features g_m^R and g_m^D to enhance g_m^S . More importantly, the modality-specific decoder is given a supervision signal to guide feature learning for modality-specific property preservation, which benefits the final prediction results when integrating them in the shared decoder. We also note the differences between the CIM and the MFA: the CIM is used to learn the fused multi-modal (RGB and depth) feature representation, while the MFA utilizes the learned modality-specific features to form an aggregate feature representation in the shared decoder.

3.4 Loss Function

We may now formulate a unified, end-to-end trainable framework. The overall loss function has two parts, \mathcal{L}_{sp} and \mathcal{L}_{sh} , for the modality-specific and decoders, respectively. For convenience, S_R and S_D denote the prediction maps for RGB and depth images, respectively, S_{sh} denotes the prediction map using their shared representation, and G denotes the ground truth map. Therefore, the overall loss function can be formulated as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{sh}(S_{sh}, G) + \mathcal{L}_{sp}(S_R, G) + \mathcal{L}_{sp}(S_D, G). \quad (3)$$

Here, we utilize the pixel position-aware loss [80] for \mathcal{L}_{sp} and \mathcal{L}_{sh} , which can pay different attention to hard and easy pixels to improve results.

4 Experimental Results and Analysis

In this section, we first give the experimental setup, including datasets, evaluation metrics, and implementation details. Then we carry out a quantitatively and qualitatively evaluation, as well as conducting ablation studies to validate the effectiveness of each key component. Finally, we conduct an attribute-based evaluation to show the effectiveness of our model in dealing with different challenges.

4.1 Experimental Setup

4.1.1 Datasets

To validate the effectiveness of the proposed model, we have evaluated it on six public RGB-D SOD

Tab. 1 Benchmarking results using 8 representative traditional models and 23 deep models on six public RGB-D saliency detection datasets using four widely used evaluation metrics: S_α [8], max E_ϕ [18], max F_β [1], and \mathcal{M} [62]). \uparrow, \downarrow indicate that larger or smaller is better. The subscript for each model denotes the publication year. Best results are highlighted in bold.

Model	NJU2K [37]				STERE [57]				DES [10]				NLPR [61]				SSD [106]				SIP [21]			
	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$
LHM ₁₄ [61]	.514	.632	.724	.205	.562	.683	.771	.172	.562	.511	.653	.114	.630	.622	.766	.108	.566	.568	.717	.195	.511	.574	.716	.184
ACSD ₁₄ [37]	.699	.711	.803	.202	.692	.669	.806	.200	.728	.756	.850	.169	.673	.607	.780	.179	.675	.682	.785	.203	.732	.763	.838	.172
LBE ₁₆ [23]	.695	.748	.803	.153	.660	.633	.787	.250	.703	.788	.890	.208	.762	.745	.855	.081	.621	.619	.736	.278	.727	.751	.853	.200
DCMC ₁₆ [12]	.686	.715	.799	.172	.731	.740	.819	.148	.707	.666	.773	.111	.724	.648	.793	.117	.704	.711	.786	.169	.683	.618	.743	.186
SE ₁₆ [27]	.664	.748	.813	.169	.708	.755	.846	.143	.741	.741	.856	.090	.756	.713	.847	.091	.675	.710	.800	.165	.628	.661	.771	.164
MDSF ₁₇ [72]	.748	.775	.838	.157	.728	.719	.809	.176	.741	.746	.851	.122	.805	.793	.885	.095	.673	.703	.779	.192	.717	.698	.798	.167
CDCP ₁₇ [107]	.669	.621	.741	.180	.713	.664	.786	.149	.709	.631	.811	.115	.669	.621	.741	.180	.603	.535	.700	.214	.595	.505	.721	.224
DTM ₂₀ [11]	.766	.716	.799	.190	.747	.743	.837	.168	.752	.697	.858	.123	.733	.677	.833	.145	.677	.651	.773	.199	.690	.659	.778	.203
DF ₁₇ [66]	.763	.804	.864	.141	.757	.757	.847	.141	.752	.766	.870	.093	.802	.778	.880	.085	.747	.735	.828	.142	.653	.657	.759	.185
CTMF ₁₈ [28]	.849	.845	.913	.085	.848	.831	.912	.086	.863	.844	.932	.055	.860	.825	.929	.056	.776	.729	.865	.099	.716	.694	.829	.139
PCF ₁₈ [3]	.877	.872	.924	.059	.875	.860	.925	.064	.842	.804	.893	.049	.874	.841	.925	.044	.841	.807	.894	.062	.842	.838	.901	.071
AFNet ₁₉ [75]	.772	.775	.853	.100	.825	.823	.887	.075	.770	.729	.881	.068	.799	.771	.879	.058	.714	.687	.807	.118	.720	.712	.819	.118
CPFP ₁₉ [94]	.878	.877	.923	.053	.879	.874	.925	.051	.872	.846	.923	.038	.888	.867	.932	.036	.807	.766	.852	.082	.850	.851	.903	.064
MMCI ₁₉ [6]	.859	.853	.915	.079	.873	.863	.927	.068	.848	.822	.928	.065	.856	.815	.913	.059	.813	.781	.882	.082	.833	.818	.897	.086
TANet ₁₉ [4]	.878	.874	.925	.060	.871	.861	.923	.060	.858	.827	.910	.046	.886	.863	.941	.041	.839	.810	.897	.063	.835	.830	.895	.075
DMRA ₁₉ [63]	.886	.886	.927	.051	.886	.886	.938	.047	.900	.888	.943	.030	.899	.879	.947	.031	.857	.844	.906	.058	.806	.821	.875	.085
cmSalGAN ₂₀ [35]	.903	.896	.940	.046	.900	.894	.936	.050	.913	.899	.943	.028	.922	.907	.957	.027	.791	.735	.867	.086	.865	.864	.906	.064
ASIFNet ₂₀ [40]	.889	.888	.927	.047	.878	.878	.927	.049	.934	.935	.974	.019	.906	.888	.944	.030	.857	.834	.884	.056	.857	.859	.896	.061
ICNet ₂₀ [43]	.894	.891	.926	.052	.903	.898	.942	.045	.920	.913	.960	.027	.923	.908	.952	.028	.848	.841	.902	.064	.854	.857	.903	.069
A2dele ₂₀ [64]	.871	.874	.916	.051	.878	.879	.928	.044	.886	.872	.920	.029	.898	.882	.944	.029	.802	.776	.861	.070	.828	.833	.889	.070
JL-DCF ₂₀ [24]	.903	.903	.944	.043	.905	.901	.946	.042	.929	.919	.968	.022	.925	.916	.962	.022	.830	.795	.885	.068	.879	.885	.923	.051
S ² MA ₂₀ [51]	.894	.889	.930	.053	.890	.882	.932	.051	.941	.935	.973	.021	.915	.902	.953	.030	.868	.848	.909	.052	.872	.877	.919	.057
UCNet ₂₀ [86]	.897	.895	.936	.043	.903	.899	.944	.039	.933	.930	.976	.018	.920	.903	.956	.025	.865	.854	.907	.049	.875	.879	.919	.051
SSF ₂₀ [90]	.899	.896	.935	.043	.893	.890	.936	.044	.904	.884	.941	.026	.914	.896	.953	.026	.845	.824	.897	.058	.876	.882	.922	.052
HDFNet ₂₀ [58]	.908	.911	.944	.038	.900	.900	.943	.041	.926	.921	.970	.021	.923	.917	.963	.023	.879	.870	.925	.045	.886	.894	.930	.047
Cas-GNN ₂₀ [55]	.911	.903	.933	.035	.899	.901	.930	.039	.905	.906	.947	.028	.919	.904	.947	.028	.872	.862	.915	.047	.875	.879	.919	.051
CMMS ₂₀ [41]	.900	.897	.936	.044	.895	.893	.939	.043	.937	.930	.976	.018	.915	.896	.949	.027	.874	.864	.922	.046	.872	.877	.911	.058
CoNet ₂₀ [34]	.895	.893	.937	.046	.908	.905	.949	.040	.909	.896	.945	.028	.908	.887	.945	.031	.853	.840	.915	.059	.858	.867	.913	.063
DANet ₂₀ [97]	.899	.910	.935	.045	.901	.892	.937	.043	.924	.928	.968	.023	.915	.916	.953	.028	.864	.866	.914	.050	.875	.892	.918	.054
PGAR ₂₀ [9]	.909	.907	.940	.042	.907	.898	.939	.041	.913	.902	.945	.026	.930	.916	.961	.024	.865	.838	.898	.057	.876	.876	.915	.055
D ³ Net ₂₁ [21]	.900	.900	.950	.041	.899	.891	.938	.046	.898	.885	.946	.031	.912	.897	.953	.030	.857	.834	.910	.058	.860	.861	.909	.063
SPNet (Ours)	.925	.935	.954	.028	.907	.915	.944	.037	.945	.950	.980	.014	.927	.925	.959	.021	.871	.883	.915	.044	.894	.916	.930	.043

datasets: NJU2K [37], NLPR [61], DES [10], SSD [106], STERE [57] and SIP [21]. Details of each dataset can be found at <https://github.com/taozh2017/RGBD-SODsurvey>.

For a fair comparison, we utilized the same protocol to form the training and test sets as introduced in [21, 63]. The training set includes 2,195 samples in total, with 1,485 samples from NJU2K [37] and 700 samples from NLPR [61]. The remaining samples from NJU2K (500) and NLPR (300), and the entire DES (135), SSD (80), STERE (1,000), and SIP (929) datasets were used for testing.

4.1.2 Evaluation Metrics

We adopt four widely used metrics to evaluate the effectiveness of the proposed model. Their definitions are as follows.

- **Structure Measure.** The S-measure S_α [8] assesses the structural similarity between regional perception (S_r) and object perception (S_o), and is defined as:

$$S_\alpha = \alpha S_o + (1 - \alpha) S_r, \quad (4)$$

where $\alpha \in [0, 1]$ is a trade-off parameter, set to 0.5 by default [8].

- **F-measure.** Given a saliency map S , we convert it to a binary map M , and then compute the Precision and Recall [1] using

$$\text{Precision} = \frac{|M \cap G|}{|M|}, \quad \text{Recall} = \frac{|M \cap G|}{|G|}, \quad (5)$$

where G denotes the ground-truth. A popular strategy is to partition S using a set of thresholds varying from 0 to 255. For each threshold, we calculate a pair of recall and precision scores, and then combine all scores to obtain a PR curve.

The F-measure F_β [1] combines both precision and recall, via a weighted harmonic mean:

$$F_\beta = (1 + \beta^2) \frac{\text{Precision} \times \text{Recall}}{\beta^2 \text{Precision} + \text{Recall}}, \quad (6)$$

where β^2 is set to 0.3 to emphasize precision [1]. We use different fixed $[0, 255]$ thresholds to compute the F -measure. This yields a set of F -measure values; we report the maximum F_β values from our experiments.

- **Enhanced-alignment Measure.** E_ϕ [18] is used to capture image-level statistics and local pixel matching information. It is defined as

$$E_\phi = \frac{1}{WH} \sum_{i=1}^W \sum_{j=1}^H \phi_{FM}(i, j), \quad (7)$$

where ϕ_{FM} denotes the enhanced-alignment matrix [18].

- **Mean Absolute Error (\mathcal{M}).** It is adopted to evaluate the average pixel-level relative error between the ground truth (*i.e.*, G) and normalized prediction (*i.e.*, S), which is defined by

$$\mathcal{M} = \frac{1}{W * H} \sum_{i=1}^W \sum_{j=1}^H |S(i, j) - G(i, j)|, \quad (8)$$

where W and H denote the width and height of the map, respectively. \mathcal{M} estimates the similarity between the saliency map and the ground-truth map, and normalizes it to $[0, 1]$.

4.1.3 Implementation Details

Our proposed model was implemented with the PyTorch library, and trained on an nVidia Tesla V100 GPU with 32 GB memory. Res2Net-50 [25], pre-trained on ImageNet [70], was used as the backbone network. Since RGB and depth images have different numbers of channels, the input channel for the depth encoder was modified to 1. We utilized the Adam algorithm to optimize the proposed model. The initial learning rate was set to 10^{-4} and divided by 10 every 60 epochs. The input RGB and depth images were resized to 352×352 . To enhance the generalizability of the proposed learning algorithm, we adopted multiple data augmented strategies: random flipping, rotation, and border clipping. The batch size was set to 20 and the model was trained over 200 epochs.

For testing, the RGB and depth images were first resized to 352×352 and then fed into the model to obtain the predicted saliency map. The predicted saliency map was then resized back to the original size of the input images. The output of the shared decoder is regarded as the final prediction of our model.

4.2 Comparison

4.2.1 Models Compared

We compared our proposed SPNet with 30 RGB-D saliency detection methods, including 8 handcrafted traditional models: LHM [61], ACSF [37], LBE [23], DCMC [12], SE [27], MDSF [72], CDCP [107], and DTM [11], and 23 deep models: DF [66], CTMF [28], PCF [3], AFNet [75], CFPF [94], MMCI [6], TANet [4], DMRA [63], cmSalGAN [35], ASIFNet [40], ICNet [43],

A2dele [64], JL-DCF [24], S²MA [51], UCNNet [86], SSF [90], HDFNet [58], Cas-GNN [55], CMMS [41], D³Net [21], CoNet [34], DANet [97], and PGAR [9]. See also the survey in [101].

4.2.2 Quantitative Evaluation

As Table 1 shows, our method is superior to the eight traditional methods LHM [61], ACSF [37], LBE [23], DCMC [12], SE [27], MDSF [72], and CDCP [107], by a large margin, on all six datasets. Our method furthermore outperforms all compared state-of-the-art methods and obtains the best performance in terms of the four evaluation metrics on NJU2K, DES, and SIP datasets. It is worth noting that our model obtains better results on STERE and NLPR than most compared RGB-D saliency detection methods. Our model is also comparable with CoNet on the STERE dataset, and JL-DCF and PGAR on the NLPR dataset. Overall, our proposed SPNet obtains promising results in locating salient object(s) in a given scene. We further show PR curves in Fig. 5 and F-measure curves in Fig. 6, giving results for 29 RGB-D saliency detection methods, including 28 state-of-the-art models with complete saliency maps. The superiority of our model is clearly visible on these datasets.

In addition, we compared our SPNet to 13 recent state-of-the-art models on the ReDWeb-S dataset. Results for the other methods are from <https://github.com/nmizhang/SMAC>, while results for our method were obtained by testing the model (trained using NJU2K [37] and NLPR [61]) on the ReDWeb-S dataset. The comparison is shown in Table 2. Our method works better than most compared methods, and is comparable to UCNNet and JL-DCF on the ReDWeb-S dataset.

We further compared using different backbone networks in the proposed model, with the results shown in Table 3. The proposed model works better when using Res2Net-50 as the backbone, yet the model using ResNet-50 as backbone still performs better than other methods (see Table 1).

4.2.3 Qualitative Evaluation

Fig. 7 shows several representative samples of results comparing our model with those from eight state-of-the-art methods. The first row shows a scene with a small object. Our method, A2dele, PGAR, and D3Net accurately detect the salient object, while JL-DCF, S2MA, SSF, and UCNNet predict some non-object regions. Rows 2,3 show two examples of scenes with complex backgrounds. Our method and S2MA produce reliable results, while other RGB-D saliency

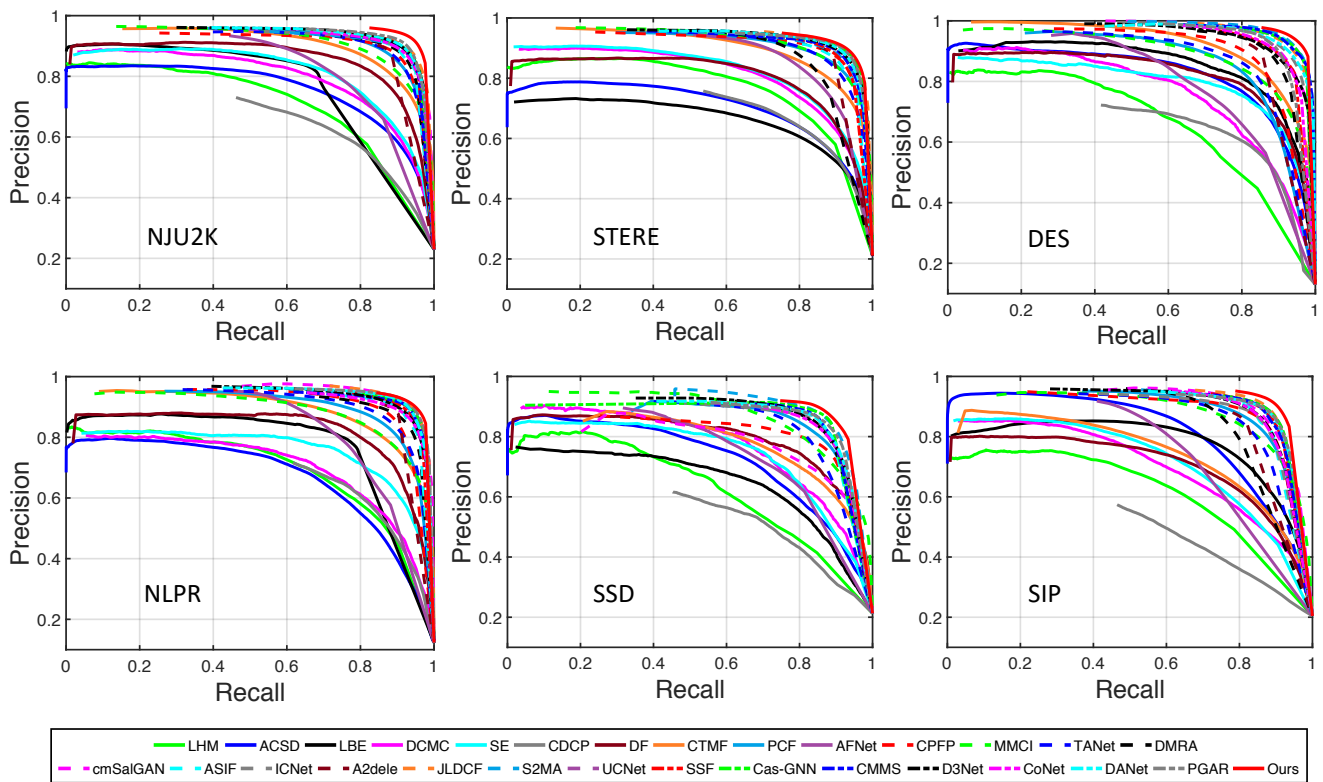


Fig. 5 PR curves for six datasets: NJU2K [37], STERE [57], DES [10], NLPR [61], SSD [106], and SIP [21].

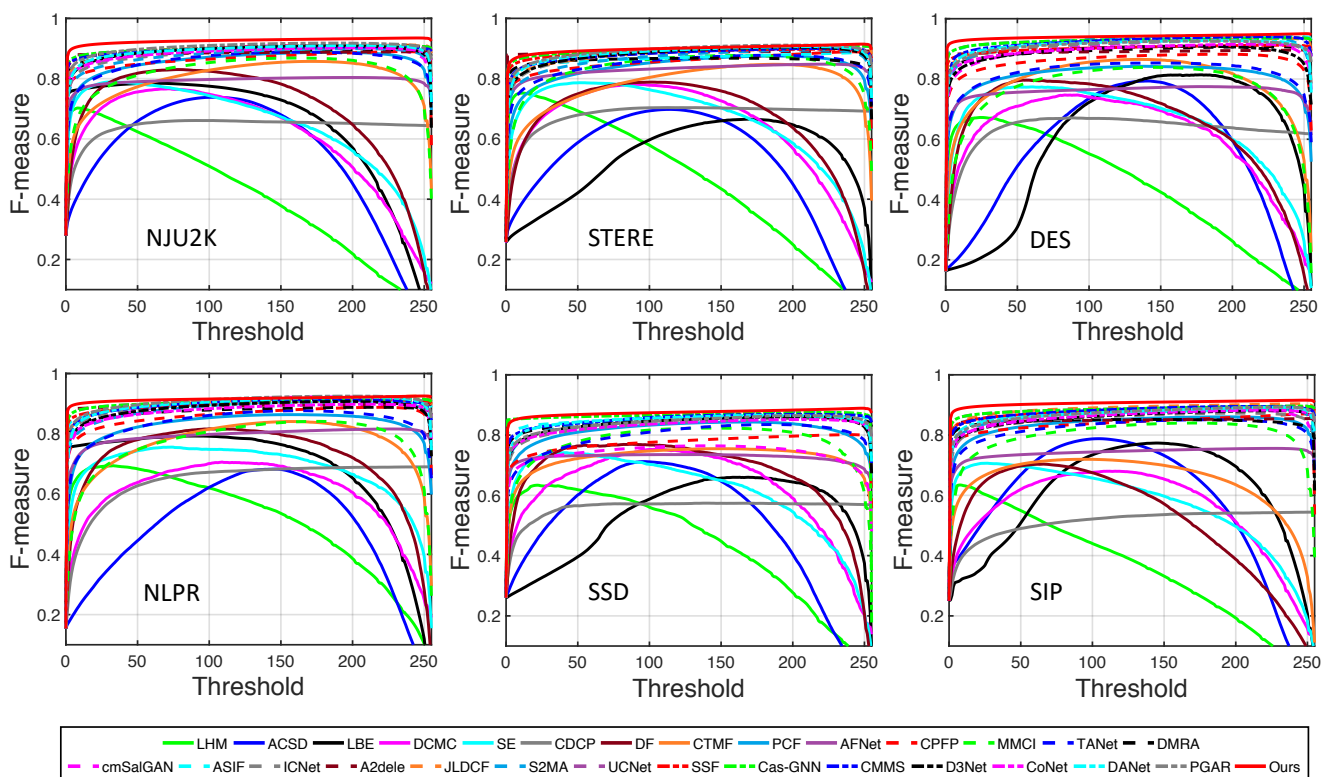


Fig. 6 F-measure curves for different thresholds, for NJU2K [37], STERE [57], DES [10], NLPR [61], SSD [106], and SIP [21].




Fig. 7 Visual comparison of results from our method and eight state-of-the-art methods: A2dele [64], JL-DCF [24], S2MA [51], UCNet [86], SSF [90], D3Net [21], DANet [97], and PGAR [9].

Tab. 2 Results from our model and 13 state-of-the-art methods: CTMFF [28], PCF [3], AFNet [75], MCCI [6], CFPF [94], DMRA [63], TANet [4], A2dele [64], UCNet [86], JL-DCF [24], S²MA [51], SSF [90], and D³Net [21] on the ReDWeb-S dataset.

Model	CTMF	PCF	AFNet	MCCI	CPFP	DMRA	TANet	A2dele	UCNet	JL-DCF	S ² MA	SSF	D ³ Net	Ours
$S_\alpha \uparrow$	0.641	0.655	0.546	0.660	0.685	0.592	0.656	0.641	0.713	0.734	0.711	0.595	0.689	0.710
$F_\beta \uparrow$	0.607	0.627	0.549	0.641	0.645	0.579	0.623	0.603	0.710	0.727	0.696	0.558	0.673	0.715
$E_\phi \uparrow$	0.739	0.743	0.693	0.754	0.744	0.721	0.741	0.672	0.794	0.805	0.781	0.710	0.768	0.800
$\mathcal{M} \downarrow$	0.204	0.166	0.213	0.176	0.142	0.188	0.165	0.160	0.130	0.128	0.139	0.189	0.149	0.129

Tab. 3 Results from our model using different backbone networks.

Backbone	NJU2K [37]				STERE [57]				DES [10]				NLPR [61]				SSD [106]				SIP [21]			
	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$
ResNet-50	.922	.934	.952	.030	.904	.914	.942	.037	.936	.944	.974	.016	.930	.931	.965	.020	.869	.876	.906	.044	.896	.916	.934	.041
Res2Net-50	.925	.935	.954	.028	.907	.915	.944	.037	.945	.950	.980	.014	.927	.925	.959	.021	.871	.883	.915	.044	.894	.916	.930	.043

detection models fail to locate the object or confuse the background with a salient object. In row 4, the compared methods other than D3Net locate a non-salient and small object. In row 5, we show an example with multiple salient objects, where it is challenging to accurately locate them all. Our method locates all salient objects and segments them more accurately, generating sharper edges than other approaches. We show an example under low-light conditions in the last row. While some approaches fail to detect the entire extent of the salient object, our model suppresses background distractors and gives good saliency detection results.

Tab. 4 Comparisons of inference time (ms) and model size (MB) for different methods.

Method	Ours	JL-DCF	S2MA
Model Size	175.3	124.5	82.7
Inference Time	91.7	21.8	22.1
Method	UCNet	SSF	HDFNet
Model Size	31.3	32.9	153.2
Inference Time	31.8	45.7	57.1

4.2.4 Inference Time and Model Size

We tested the inference time for different methods on an NVIDIA TESLA P40 GPU with 24G memory. The inference times and model sizes of different methods,

Tab. 5 Quantitative evaluation for ablation studies.

	NJU2K [37]		STERE [57]		DES [10]		NLPR [61]		SSD [106]		SIP [21]	
	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$
Ours	.925	.028	.907	.037	.945	.014	.927	.021	.871	.044	.894	.043
A1	.916	.034	.898	.042	.939	.016	.926	.022	.869	.047	.892	.044
A2	.921	.031	.895	.042	.938	.016	.925	.022	.865	.051	.896	.042
A3	.919	.032	.895	.043	.938	.016	.929	.020	.864	.049	.887	.048
A4	.924	.029	.903	.038	.930	.019	.927	.023	.867	.049	.888	.046
B1	.918	.034	.901	.041	.939	.017	.922	.024	.858	.050	.885	.048
B2	.924	.029	.900	.041	.941	.015	.926	.022	.864	.049	.893	.044
B3	.921	.031	.903	.039	.938	.016	.925	.022	.863	.050	.891	.045
C1	.913	.037	.900	.047	.935	.019	.922	.025	.861	.055	.880	.051
C2	.916	.034	.906	.040	.923	.021	.924	.022	.866	.049	.882	.051




Fig. 8 Comparison of MFA module with other fusion strategies.

including our SPNet, JL-DCF [24], S2MA [51], UCNNet [86], SSF [90], and HDFNet [58], are shown in Table 4. Because our model adopts two modality-specific networks and a shared learning network to generate individual and shared saliency prediction maps, it has a relatively large model size and takes more inference time for saliency prediction than other methods. We thus hope to design lightweight networks to improve the efficiency of SPNet in future work.

4.3 Ablation Studies

To verify the relative importance of different key components of our model, we conducted ablation studies by removing or replacing them.

4.3.1 Effectiveness of CIM

Since the proposed CIM is used to fuse cross-modal features and learn their shared representation, compared it to an alternative of a direct concatenation strategy. Specifically, the two features f_m^R and f_m^D (see Fig. 3) are directly concatenated and then fed into a 3×3 convolutional layer to obtain the fused representation in each layer. We denote this approach as A1 in Table 5, which shows that our model performs better when using the proposed CIM than using a simple feature concatenation strategy. This also indicates

the contribution of the CIM in improving the saliency detection results. Going further, there are two parts to CIM: cross-modal feature enhancement and adaptive feature fusion. Thus, to evaluate the contribution of each part, we modified CIM to have only cross-modal feature enhancement or adaptive feature fusion, with results denoted A2 and A3, respectively. When comparing them to the full version of CIM, we can see the effectiveness of the proposed CIM. Moreover, in CIM, the features of the last layer are propagated to the next layer to capture cross-level correlations. To validate the effectiveness of the propagation strategy, we removed this propagation in the CIM, with results denoted A4, showing that the propagation strategy does improve saliency detection results.

4.3.2 Effectiveness of MFA

In the proposed framework, the MFA is proposed to make full use of the features learned in the modality-specific decoder, which are then integrated into the shared decoder to provide more multi-modal complementary information. To validate its effectiveness, we deleted this module in an approach denoted B1. We also considered two other feature fusion strategies: see Fig. 8. One provides cross-modal feature enhancement fusion; the other is a simple concatenation strategy. Results for the two strategies are denoted B2 and B3. Table 5 demonstrates, by comparing results of B1 and our full model, the effectiveness of integrating the features learned into the shared decoder. Comparing results of B2 and B3 with our full model, we can see that the MFA module outperforms both other fusion strategies.




Fig. 9 Attribute-based evaluation with respect to (a) number of salient objects (one, or multiple), (b) indoor versus outdoor environments, and (c) lighting conditions (low-light versus sunny).

4.3.3 Effectiveness of Modality-specific Decoders

We deleted the two modality-specific decoders, with results shown in C1 in Table 5. Performance degrades when not using the two parts. This indicates the effectiveness of the modality-specific decoders, which provide supervision signals to ensure that modality-specific properties can be learned.

To further evaluate the effectiveness of the combination of the two modality-specific decoders, we added an experiment to compare the SOD results when using the output from the shared decoder and the combination of the two modality-specific decoders. Results are shown in C2 of Table 5. We can see that the shared decoder outperforms the combination of the two modality-specific decoders, indicating that the shared decoder can combine multi-modal shared information and modality-specific characteristics to improve SOD results.

4.3.4 Effects of Varying Numbers of CIMs

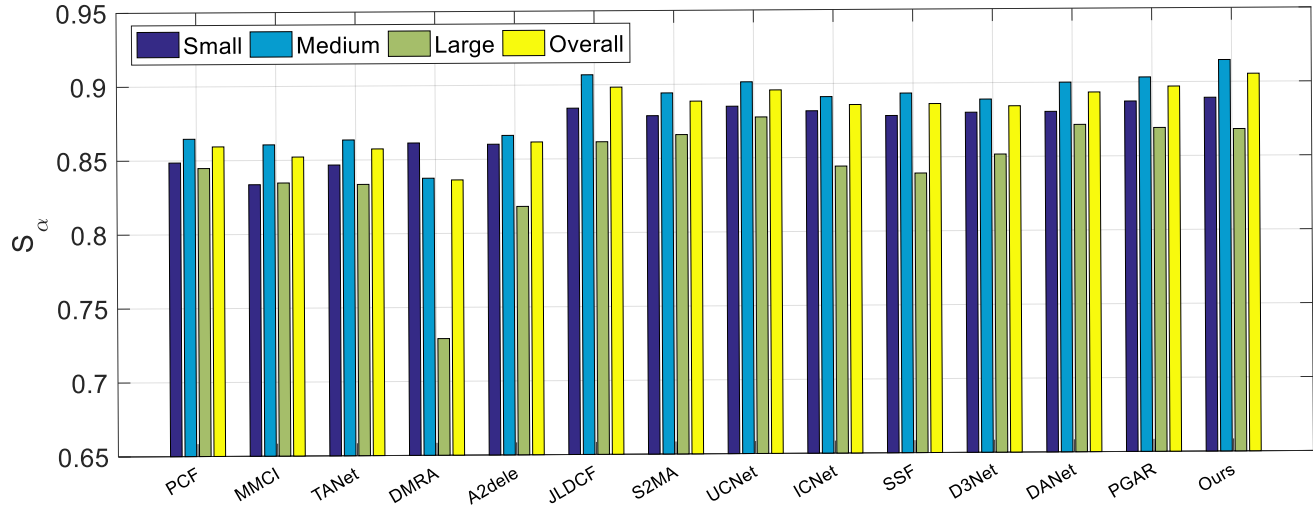
To investigate the effects of changing the numbers of CIMs, we compare our full model using five CIMs with two degraded versions, CIM_1 , which only applies a CIM to the features from the last layer in the encoder network, and CIM_3 , using CIMs on the features from each of the last three layers in the encoder network. Table 6 shows the results; our model with five CIMs works better for most datasets.

4.4 Attribute-based Evaluation

There are several challenging factors that may affect results from RGB-D saliency detection models, such as the number of salient objects, indoor versus outdoor environments, lighting conditions, and so on. Thus, we evaluated saliency detection results under different conditions, to show the strengths and weaknesses of state-of-the-art models in handling these challenges.

Tab. 6 Results for different numbers of CIMs.

	NJU2K		STERE		DES		NLPR		SSD		SIP	
	$S_\alpha \uparrow$	$M \downarrow$	$S_\alpha \uparrow$	$M \downarrow$	$S_\alpha \uparrow$	$M \downarrow$	$S_\alpha \uparrow$	$M \downarrow$	$S_\alpha \uparrow$	$M \downarrow$	$S_\alpha \uparrow$	$M \downarrow$
CIM ₁	.918	.034	.908	.039	.929	.019	.928	.022	.865	.047	.889	.046
CIM ₃	.920	.032	.900	.041	.935	.017	.928	.021	.857	.049	.891	.045
Ours	.925	.028	.907	.037	.945	.014	.927	.021	.871	.044	.894	.043

**Fig. 10** Attribute-based evaluation with respect to scale of the salient object.

4.4.1 Single vs. Multiple Objects


In this evaluation, we constructed a hybrid dataset with 1,229 images from the NLPR [61] and SIP [21] datasets. Results using S_α are shown in Fig. 9(a). As can be observed, it is easier to detect a single salient object than several. Our model outperforms other state-of-the-art methods in locating single and multiple objects.

4.4.2 Indoor vs. Outdoor

We evaluated the results of different RGB-D SOD models on indoor and outdoor scenes. As DES [10] and NLPR [61] include indoor and outdoor scenes, we constructed a hybrid dataset collected from the two datasets. Results are shown in Fig. 9(b). As can be observed, many models find it harder to detect salient objects in indoor scenes than outdoor scenes, while JLDCF, S2MA, UCNet, ICNet, SSF, DANet, and our model work a little better on outdoor scenes.

4.4.3 Lighting Conditions

We carried out this evaluation on the SIP dataset [21], with examples grouped into two categories, sunny and low-light. Results are shown in Fig. 9(c). All models struggle to detect salient objects in low-light conditions, confirming that low-light negatively impacts SOD performance.

**Fig. 11** Cases in which our model fails.

4.4.4 Object Scale

To characterize the scale of a salient object, we compute the ratio r of the size of the salient region to the whole image, and define three object scales: small, when the $r < 0.1$, large, when $r > 0.4$, and medium otherwise. To evaluate how different methods handle scale variation, we constructed a hybrid dataset with 2,444 images from STERE [57], NLPR [61], SSD [106], DES [10], and SIP [21]. Fig. 10 shows results of

Tab. 7 Results for camouflaged object detection models on benchmark datasets using evaluation metrics S_α [8] and \mathcal{M} [62]. \uparrow, \downarrow indicate that larger or smaller is better.

Model	CHAMELEON		CAMO		COD10K	
	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$\mathcal{M} \downarrow$
FPN [47]	0.794	0.075	0.684	0.131	0.697	0.075
MaskRCNN [29]	0.643	0.099	0.574	0.151	0.613	0.080
PSPNet [92]	0.773	0.085	0.663	0.139	0.678	0.080
PiCANet [49]	0.769	0.085	0.609	0.156	0.649	0.090
BASNet [65]	0.687	0.118	0.618	0.159	0.634	0.105
PFANet [96]	0.679	0.144	0.659	0.172	0.636	0.128
CPD [83]	0.853	0.052	0.726	0.115	0.747	0.059
EGNet [93]	0.848	0.050	0.732	0.104	0.737	0.056
SINet [19]	0.869	0.044	0.751	0.100	0.771	0.051
DANet [98]	0.874	0.043	0.752	0.100	0.765	0.051
HDFNet [59]	0.875	0.032	0.778	0.085	0.779	0.045
SPNet (Ours)	0.895	0.027	0.795	0.082	0.797	0.042

this attribute-based evaluation with respect to the scales of the salient objects. All methods work better at detecting small salient objects and relatively at detecting large salient objects. The most recent models, JL-DCF, DANet, PGAR, and our model, obtain the promising results.

4.5 Failures and Discussion

Our proposed SPNet shows good RGB-D saliency detection in most cases. However, it fails to detect salient objects in some challenging scenes such as those with complex backgrounds and low-quality depth data. Some failures of our model are shown in Fig. 11. In the first row, the depth data quality is very poor, so our model can only roughly locate the boat without fine details. This suggests that it is helpful to enhance or filter depth maps to improve saliency detection results. In the second row, the annotated salient object has a similar appearance to other objects in the scene, so it is challenging to accurately detect the salient object. In the third row, the object has fine details, but our model only locates the main regions without the fine details. There is still considerable room to improve our model to handle such scenes with fine structures.

4.6 Application to RGB-D Camouflaged Object Detection

SPNet was originally designed for the RGB-D SOD task, which can be easily extended to other related RGB-D tasks, *e.g.*, RGB-D based camouflaged object

detection (COD). The aim of COD is to identify objects that are ‘seamlessly’ embedded in their background surroundings. This is a very challenging task due to the high intrinsic similarities between the target object and the background [19, 45, 73]. Recent research [88] suggests that depth can provide useful spatial information to improve COD results. Thus, we extended SPNet to the RGB-D COD task.

We conducted this experiment on three public benchmark datasets for camouflaged object detection: (i) CHAMELEON [19] consisting of 76 camouflaged images, (ii) CAMO [39], with 1,250 images (1,000 for training, 250 for testing) in 8 categories, and (iii) COD10K [19], with 5,066 camouflaged images (3,040 for training, 2,026 for testing) in 5 super-classes and 69 sub-classes. Following the same setting in [20], we divided the training and testing sets and then trained our model on the training set.

We compare our method to other existing COD models, including FPN [47], MaskRCNN [29], PSPNet [92], PiCANet [49], BASNet [65], PFANet [96], CPD [83], EGNet [93], and SINet [20] (results are from [20]). Since there are few works for RGB-D camouflaged object detection, we also compared two recent RGB-D salient object detection methods, DANet [98], and HDFNet [59], in this experiment. We re-trained the two RGB-D SOD models and our model using RGB and depth images.

Table 7 shows quantitative results for three public datasets. Our model performs better than the other




Fig. 12 COD results of our SPNet and three state-of-the-art COD methods: SINet [20], DANet [98], and HDFNet [59].

COD methods. Our model and the two RGB-D COD methods use depth cues, and work better than other methods which do not, indicating that depth cues can provide spatial information to improve COD results. Fig. 12 shows qualitative results for different COD methods. Compared to other COD models, our SPNet achieves better results by detecting more accurate boundaries of camouflaged objects.

5 Conclusions

In this paper, we have presented a novel RGB-D salient object detection framework, SPNet. Unlike most existing RGB-D SOD methods, which focus on learning shared representations, SPNet not only explores shared cross-modal information but also uses modality-specific characteristics to improve SOD results. To learn the shared representations for the two modalities, we introduce a cross-enhanced integration module (CIM) to fuse the cross-modal features, and the output of each CIM is propagated to the next layer to explore rich cross-level information. We further adopt a multi-modal feature aggregation (MFA) module to integrate the learned modality-specific features to enhance the complementary multi-modal information. Extensive results on benchmark datasets show the effectiveness of our model in comparison to other state-of-the-art RGB-D SOD methods. Moreover,

we have thoroughly validated the effectiveness of key components of our framework, and an attribute-based evaluation was conducted to study the ability of many cutting-edge RGB-D SOD approaches to meet different challenges. Finally, we extended SPNet to the recently proposed RGB-D camouflaged object detection task, and its effectiveness was verified.

Acknowledgements

This work was supported in part by the National Science Fund of China under Grant 62172228, in part by an Open Project of the Key Laboratory of System Control and Information Processing, Ministry of Education (Shanghai Jiao Tong University, ID: Scip202102).

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

Conflict of Interest Statement: We declare that we have no conflicts of interest.

References

- [1] R. Achanta, S. Hemami, F. Estrada, and S. Susstrunk. Frequency-tuned salient region detection. In *CVPR*, pages 1597–1604, 2009.
- [2] K. Chaudhuri, S. M. Kakade, K. Livescu, and K. Sridharan. Multi-view clustering via canonical correlation analysis. In *ICML*, pages 129–136, 2009.
- [3] H. Chen and Y. Li. Progressively complementarity-aware fusion network for RGB-D salient object detection. In *CVPR*, pages 3051–3060, 2018.
- [4] H. Chen and Y. Li. Three-stream attention-aware network for RGB-D salient object detection. *IEEE TIP*, 28(6):2825–2835, 2019.
- [5] H. Chen, Y. Li, Y. Deng, and G. Lin. CNN-based RGB-D salient object detection: Learn, select, and fuse. *IJCV*, pages 1–21, 2021.
- [6] 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 Recognition*, 86:376–385, 2019.
- [7] H. Chen, Y.-F. Li, and D. Su. Attention-aware cross-modal cross-level fusion network for RGB-D salient object detection. In *IROS*, pages 6821–6826, 2018.
- [8] M.-M. Chen and D.-P. Fan. Structure-measure: A new way to evaluate foreground maps. *IJCV*, 129:2622–2638, 2021.
- [9] S. Chen and Y. Fu. Progressively guided alternate refinement network for RGB-D salient object detection. In *ECCV*, 2020.
- [10] Y. Cheng, H. Fu, X. Wei, J. Xiao, and X. Cao. Depth enhanced saliency detection method. In *ICIMCS*, pages 23–27, 2014.
- [11] R. Cong, J. Lei, H. Fu, J. Hou, Q. Huang, and S. Kwong. Going from RGB to RGBD saliency: A depth-guided transformation model. *IEEE TCYB*, 2019.
- [12] R. Cong, J. Lei, C. Zhang, Q. Huang, X. Cao, and C. Hou. Saliency detection for stereoscopic images based on depth confidence analysis and multiple cues fusion. *SPL*, 23(6):819–823, 2016.
- [13] Z. Deng, X. Hu, L. Zhu, X. Xu, J. Qin, G. Han, and P.-A. Heng. R3net: Recurrent residual refinement network for saliency detection. In *IJCAI*, pages 684–690, 2018.
- [14] K. Desingh, K. M. Krishna, D. Rajan, and C. Jawahar. Depth really matters: Improving visual salient region detection with depth. In *BMVC*, 2013.
- [15] C. Ding and D. Tao. Robust face recognition via multimodal deep face representation. *IEEE TMM*, 17(11):2049–2058, 2015.
- [16] Y. Ding, Z. Liu, M. Huang, R. Shi, and X. Wang. Depth-aware saliency detection using convolutional neural networks. *JVCIR*, 61:1–9, 2019.
- [17] A. Eitel, J. T. Springenberg, L. Spinello, M. Riedmiller, and W. Burgard. Multimodal deep learning for robust rgb-d object recognition. In *IROS*, pages 681–687, 2015.
- [18] D.-P. Fan, C. Gong, Y. Cao, B. Ren, M.-M. Cheng, and A. Borji. Enhanced-alignment measure for binary foreground map evaluation. In *IJCAI*, pages 698–704, 2018.
- [19] D.-P. Fan, G.-P. Ji, M.-M. Cheng, and L. Shao. Concealed object detection. *IEEE TPAMI*, 2021.
- [20] D.-P. Fan, G.-P. Ji, G. Sun, M.-M. Cheng, J. Shen, and L. Shao. Camouflaged object detection. In *CVPR*, pages 2777–2787, 2020.
- [21] 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 TNNLS*, 32(5):2075–2089, 2021.
- [22] D.-P. Fan, Y. Zhai, A. Borji, J. Yang, and L. Shao. Bbs-net: RGB-D salient object detection with a bifurcated backbone strategy network. In *ECCV*, pages 275–292, 2020.
- [23] D. Feng, N. Barnes, S. You, and C. McCarthy. Local background enclosure for RGB-D salient object detection. In *CVPR*, pages 2343–2350, 2016.
- [24] 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 TPAMI*, 2021.
- [25] S.-H. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. Torr. Res2net: A new multi-scale backbone architecture. *IEEE TPAMI*, 2020.
- [26] M. Gönen and E. Alpaydm. Multiple kernel learning algorithms. *JMLR*, 12:2211–2268,


- 2011.
- [27] J. Guo, T. Ren, and J. Bei. Salient object detection for RGB-D image via saliency evolution. In *ICME*, pages 1–6, 2016.
- [28] 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 TCYB*, 48(11):3171–3183, 2017.
- [29] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *ICCV*, 2017.
- [30] Q. Hou, M.-M. Cheng, X. Hu, A. Borji, Z. Tu, and P. H. Torr. Deeply supervised salient object detection with short connections. In *CVPR*, pages 3203–3212, 2017.
- [31] J. Hu, J. Lu, and Y.-P. Tan. Sharable and individual multi-view metric learning. *IEEE TPAMI*, 40(9):2281–2288, 2017.
- [32] Z. Huang, H.-X. Chen, T. Zhou, Y.-Z. Yang, and B.-Y. Liu. Multi-level cross-modal interaction network for rgb-d salient object detection. *Neurocomputing*, 452:200–211, 2021.
- [33] 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 *CVPR*, pages 9471–9481, 2021.
- [34] W. Ji, J. Li, M. Zhang, Y. Piao, and H. Lu. Accurate RGB-D salient object detection via collaborative learning. In *ECCV*, 2020.
- [35] B. Jiang, Z. Zhou, X. Wang, J. Tang, and B. Luo. cmsalgan: RGB-D salient object detection with cross-view generative adversarial networks. *IEEE TMM*, 2020.
- [36] Z. Jiang and L. S. Davis. Submodular salient region detection. In *CVPR*, pages 2043–2050, 2013.
- [37] R. Ju, L. Ge, W. Geng, T. Ren, and G. Wu. Depth saliency based on anisotropic center-surround difference. In *ICIP*, pages 1115–1119, 2014.
- [38] C. Lang, T. V. Nguyen, H. Katti, K. Yadati, M. Kankanhalli, and S. Yan. Depth matters: Influence of depth cues on visual saliency. In *ECCV*, pages 101–115, 2012.
- [39] T. Le, T. Nguyen, Z. Nie, M. Tran, and A. Sugimoto. Anabranh network for camouflaged object segmentation. *CVIU*, 2019.
- [40] C. Li, R. Cong, S. Kwong, J. Hou, H. Fu, G. Zhu, D. Zhang, and Q. Huang. ASIF-Net: Attention steered interweave fusion network for RGB-D salient object detection. *IEEE TCYB*, 2020.
- [41] C. Li, R. Cong, Y. Piao, Q. Xu, and C. C. Loy. RGB-D salient object detection with cross-modality modulation and selection. In *ECCV*, 2020.
- [42] G. Li, Z. Liu, M. Chen, Z. Bai, W. Lin, and H. Ling. Hierarchical alternate interaction network for rgb-d salient object detection. *IEEE Transactions on Image Processing*, 30:3528–3542, 2021.
- [43] G. Li, Z. Liu, and H. Ling. Icnnet: Information conversion network for RGB-D based salient object detection. *IEEE TIP*, 29:4873–4884, 2020.
- [44] G. Li, Z. Liu, L. Ye, Y. Wang, and H. Ling. Cross-modal weighting network for RGB-D salient object detection. In *ECCV*, 2020.
- [45] L. Li, B. Dong, E. Rigall, T. Zhou, J. Dong, and G. Chen. Marine animal segmentation. *IEEE TCSVT*, 2021.
- [46] F. Liang, L. Duan, W. Ma, Y. Qiao, Z. Cai, and L. Qing. Stereoscopic saliency model using contrast and depth-guided-background prior. *Neurocomputing*, 275:2227–2238, 2018.
- [47] T. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [48] D. Liu, Y. Hu, K. Zhang, and Z. Chen. Two-stream refinement network for RGB-D saliency detection. In *ICIP*, pages 3925–3929, 2019.
- [49] N. Liu, J. Han, and M. Yang. Picanet: Learning pixel-wise contextual attention for saliency detection. In *CVPR*, 2018.
- [50] N. Liu, J. Han, and M.-H. Yang. Picanet: Learning pixel-wise contextual attention for saliency detection. In *CVPR*, pages 3089–3098, 2018.
- [51] N. Liu, N. Zhang, and J. Han. Learning selective self-mutual attention for RGB-D saliency detection. In *CVPR*, 2020.
- [52] N. Liu, N. Zhang, K. Wan, L. Shao, and J. Han. Visual saliency transformer. In *ICCV*, 2021.
- [53] 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*, 363:46–57, 2019.

- [54] Y. Lu, Y. Wu, B. Liu, T. Zhang, B. Li, Q. Chu, and N. Yu. Cross-modality person re-identification with shared-specific feature transfer. In *CVPR*, pages 13379–13389, 2020.
- [55] A. Luo, X. Li, F. Yang, Z. Jiao, H. Cheng, and S. Lyu. Cascade graph neural networks for RGB-D salient object detection. In *ECCV*, 2020.
- [56] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng. Multimodal deep learning. In *ICML*, 2011.
- [57] Y. Niu, Y. Geng, X. Li, and F. Liu. Leveraging stereopsis for saliency analysis. In *CVPR*, pages 454–461, 2012.
- [58] Y. Pang, L. Zhang, X. Zhao, and H. Lu. Hierarchical dynamic filtering network for RGB-D salient object detection. In *ECCV*, 2020.
- [59] Y. Pang, L. Zhang, X. Zhao, and H. Lu. Hierarchical dynamic filtering network for RGB-D salient object detection. In *ECCV*, 2020.
- [60] Y. Pang, X. Zhao, L. Zhang, and H. Lu. Multi-scale interactive network for salient object detection. In *CVPR*, pages 9413–9422, 2020.
- [61] H. Peng, B. Li, W. Xiong, W. Hu, and R. Ji. RGBD salient object detection: a benchmark and algorithms. In *ECCV*, pages 92–109, 2014.
- [62] F. Perazzi, P. Krähenbühl, Y. Pritch, and A. Hornung. Saliency filters: Contrast based filtering for salient region detection. In *CVPR*, pages 733–740, 2012.
- [63] Y. Piao, W. Ji, J. Li, M. Zhang, and H. Lu. Depth-induced multi-scale recurrent attention network for saliency detection. In *ICCV*, pages 7254–7263, 2019.
- [64] Y. Piao, Z. Rong, M. Zhang, W. Ren, and H. Lu. A2dele: Adaptive and attentive depth distiller for efficient RGB-D salient object detection. In *CVPR*, 2020.
- [65] X. Qin, Z. Zhang, C. Huang, C. Gao, M. Dehghan, and M. Jagersand. Basnet: Boundary-aware salient object detection. In *CVPR*, 2019.
- [66] L. Qu, S. He, J. Zhang, J. Tian, Y. Tang, and Q. Yang. RGBD salient object detection via deep fusion. *IEEE TIP*, 26(5):2274–2285, 2017.
- [67] K. Rapantzikos, Y. Avrithis, and S. Kollias. Dense saliency-based spatiotemporal feature points for action recognition. In *CVPR*, pages 1454–1461, 2009.
- [68] J. Ren, X. Gong, L. Yu, W. Zhou, and M. Ying Yang. Exploiting global priors for RGB-D saliency detection. In *CVPRW*, pages 25–32, 2015.
- [69] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015.
- [70] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [71] W. Shimoda and K. Yanai. Distinct class-specific saliency maps for weakly supervised semantic segmentation. In *ECCV*, pages 218–234, 2016.
- [72] H. Song, Z. Liu, H. Du, G. Sun, O. Le Meur, and T. Ren. Depth-aware salient object detection and segmentation via multiscale discriminative saliency fusion and bootstrap learning. *IEEE TIP*, 26(9):4204–4216, 2017.
- [73] Y. Sun, G. Chen, T. Zhou, Y. Zhang, and N. Liu. Context-aware cross-level fusion network for camouflaged object detection. *IJCAI*, 2021.
- [74] L. Wang, L. Wang, H. Lu, P. Zhang, and X. Ruan. Salient object detection with recurrent fully convolutional networks. *IEEE TPAMI*, 41(7):1734–1746, 2018.
- [75] N. Wang and X. Gong. Adaptive fusion for RGB-D salient object detection. *IEEE Access*, 7:55277–55284, 2019.
- [76] W. Wang, Q. Lai, H. Fu, J. Shen, H. Ling, and R. Yang. Salient object detection in the deep learning era: An in-depth survey. *IEEE TPAMI*, 2021.
- [77] W. Wang, J. Shen, R. Yang, and F. Porikli. Saliency-aware video object segmentation. *IEEE TPAMI*, 40(1):20–33, 2017.
- [78] X. Wang, H. Ma, X. Chen, and S. You. Edge preserving and multi-scale contextual neural network for salient object detection. *IEEE TIP*, 27(1):121–134, 2017.
- [79] X. Wang, L. Zhu, S. Tang, H. Fu, P. Li, F. Wu, Y. Yang, and Y. Zhuang. Boosting rgb-d saliency detection by leveraging unlabeled

- rgb images. *IEEE Transactions on Image Processing*, 31:1107–1119, 2022.
- [80] J. Wei, S. Wang, and Q. Huang. F3Net: Fusion, feedback and focus for salient object detection. *AAAI*, 2019.
- [81] M. White, X. Zhang, D. Schuurmans, and Y.-I. Yu. Convex multi-view subspace learning. In *NeurIPS*, pages 1673–1681, 2012.
- [82] Z. Wu, L. Su, and Q. Huang. Cascaded partial decoder for fast and accurate salient object detection. In *CVPR*, pages 3907–3916, 2019.
- [83] Z. Wu, L. Su, and Q. Huang. Cascaded partial decoder for fast and accurate salient object detection. In *CVPR*, 2019.
- [84] Y. Zhai, D.-P. Fan, J. Yang, A. Borji, L. Shao, J. Han, and L. Wang. Bifurcated backbone strategy for rgb-d salient object detection. *IEEE TIP*, 2021.
- [85] C. Zhang, Q. Hu, H. Fu, P. Zhu, and X. Cao. Latent multi-view subspace clustering. In *CVPR*, pages 4279–4287, 2017.
- [86] J. Zhang, D.-P. Fan, Y. Dai, S. Anwar, F. Saleh, S. Aliakbarian, and N. Barnes. Uncertainty inspired RGB-D saliency detection. *IEEE TPAMI*, 2021.
- [87] J. Zhang, D.-P. Fan, Y. Dai, X. Yu, Y. Zhong, N. Barnes, and L. Shao. RGB-D saliency detection via cascaded mutual information minimization. In *ICCV*, 2021.
- [88] J. Zhang, Y. Lv, M. Xiang, A. Li, Y. Dai, and Y. Zhong. Depth confidence-aware camouflaged object detection. *arXiv preprint arXiv:2106.13217*, 2021.
- [89] L. Zhang, J. Dai, H. Lu, Y. He, and G. Wang. A bi-directional message passing model for salient object detection. In *CVPR*, pages 1741–1750, 2018.
- [90] M. Zhang, W. Ren, Y. Piao, Z. Rong, and H. Lu. Select, supplement and focus for RGB-D saliency detection. In *CVPR*, 2020.
- [91] P. Zhang, D. Wang, H. Lu, H. Wang, and X. Ruan. Amulet: Aggregating multi-level convolutional features for salient object detection. In *ICCV*, pages 202–211, 2017.
- [92] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *CVPR*, 2017.
- [93] J. Zhao, J. Liu, D. Fan, Y. Cao, J. Yang, and M. Cheng. Egnnet: Edge guidance network for salient object detection. In *CVPR*, 2019.
- [94] J.-X. Zhao, Y. Cao, D.-P. Fan, M.-M. Cheng, X.-Y. Li, and L. Zhang. Contrast prior and fluid pyramid integration for RGBD salient object detection. In *CVPR*, pages 3927–3936, 2019.
- [95] R. Zhao, W. Oyang, and X. Wang. Person re-identification by saliency learning. *IEEE TPAMI*, 39(2):356–370, 2016.
- [96] T. Zhao and X. Wu. Pyramid feature attention network for saliency detection. In *CVPR*, 2019.
- [97] 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 *ECCV*, 2020.
- [98] 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 *ECCV*, pages 646–662, 2020.
- [99] Y. Zhao, J. Zhao, J. Li, and X. Chen. Rgb-d salient object detection with ubiquitous target awareness. *IEEE Transactions on Image Processing*, 30:7717–7731, 2021.
- [100] L. Zhou, Z. Yang, Q. Yuan, Z. Zhou, and D. Hu. Salient region detection via integrating diffusion-based compactness and local contrast. *IEEE TIP*, 24(11):3308–3320, 2015.
- [101] T. Zhou, D.-P. Fan, M.-M. Cheng, J. Shen, and L. Shao. RGB-D salient object detection: A survey. *CVMJ*, pages 1–33, 2021.
- [102] T. Zhou, H. Fu, G. Chen, J. Shen, and L. Shao. Hi-net: hybrid-fusion network for multi-modal MR image synthesis. *IEEE TMI*, 39(9):2772–2781, 2020.
- [103] T. Zhou, H. Fu, G. Chen, Y. Zhou, D.-P. Fan, and L. Shao. Specificity-preserving RGB-D saliency detection. In *ICCV*, 2021.
- [104] T. Zhou, C. Zhang, X. Peng, H. Bhaskar, and J. Yang. Dual shared-specific multiview subspace clustering. *IEEE TCYB*, 50(8):3517–3530, 2019.
- [105] C. Zhu, X. Cai, K. Huang, T. H. Li, and G. Li. PDNet: Prior-model guided depth-enhanced network for salient object detection. In *ICME*, pages 199–204, 2019.
- [106] C. Zhu and G. Li. A three-pathway psychobiological framework of salient object detection using stereoscopic technology. In

ICCVW, pages 3008–3014, 2017.

- [107] C. Zhu, G. Li, W. Wang, and R. Wang. An innovative salient object detection using center-dark channel prior. In *ICCVW*, pages 1509–1515, 2017.
- [108] J.-Y. Zhu, J. Wu, Y. Xu, E. Chang, and Z. Tu. Unsupervised object class discovery via saliency-guided multiple class learning. *IEEE TPAMI*, 37(4):862–875, 2014.
- [109] W. Zhu, S. Liang, Y. Wei, and J. Sun. Saliency optimization from robust background detection. In *CVPR*, pages 2814–2821, 2014.



Tao Zhou received a Ph.D. degree in Pattern Recognition and Intelligent Systems from the Institute of Image Processing and Pattern Recognition, Shanghai Jiao Tong University, in 2016. From 2016 to 2018, he was a Postdoctoral Fellow in the BRIC and IDEA lab, University of North Carolina

at Chapel Hill. From 2018 to 2020, he was a Research Scientist at the Inception Institute of Artificial Intelligence (IIAI), Abu Dhabi. He is currently a Professor in the School of Computer Science and Engineering, Nanjing University of Science and Technology, China. His research interests include machine learning, computer vision, and medical image analysis.



Deng-Ping Fan is a postdoctoral researcher at ETH Zurich, Switzerland. He received his Ph.D. degree from Nankai University in 2019. He joined IIAI in 2019. He has published about 30 top journal and conference papers in outlets such as *IEEE TPAMI*, *CVPR*, and *ICCV*. His research interests

include computer vision, deep learning, saliency detection, especially co-salient object detection, RGB salient object

detection, RGB-D salient object detection, and video salient object detection.



Geng Chen is a Professor at Northwestern Polytechnical University, China, where he received his Ph.D. in 2016. He was a research scientist at IIAI from 2019 to 2021, and a postdoctoral research associate at the University of North Carolina at Chapel Hill, USA, from 2016 to 2019. He has published over 50 papers in peer-reviewed international conference proceedings and journals. His research interests lie in computer vision and medical image analysis.



Yi Zhou is currently an Associate Professor at Southeast University. He received an M.Sc. degree from the Department of Electronic and Electrical Engineering, University of Sheffield, U.K. in 2014, and a Ph.D. degree from the School of Computing Sciences, University of East Anglia, U.K. in 2018.

He was a Research Scientist with IIAI. His research interests include computer vision, pattern recognition, machine learning, and medical imaging.



Huazhu Fu is currently a Senior Scientist at Inception Institute of Artificial Intelligence, UAE. He received a Ph.D. degree from Tianjin University in 2013 and was a Research Fellow at NTU for two years. From 2015 to 2018, he was a Research Scientist with I2R, A*STAR, Singapore. His research

interests include computer vision, machine learning, and AI in healthcare. He serves as an Associate Editor for *IEEE TMI* and *IEEE JBHI*, and also served as the Area Chair for *MICCAI 2021* and Co-Chair for the *OMIA Workshop*.