BBS-Net: RGB-D Salient Object Detection with a Bifurcated Backbone Strategy Network

Deng-Ping Fan^{1,*[0000-0002-5245-7518]}, Yingjie Zhai^{2,*[0000-0001-7593-6055]}, Ali Borji^{3[0000-0001-8198-0335]}, Jufeng Yang^{2,#[0000-0003-0219-3443]}, and Ling Shao^{1,4[0000-0002-8264-6117]}

¹ Inception Institute of Artificial Intelligence, Abu Dhabi, UAE.
 ² Nankai University, Tianjin, China
 ³ HCL America, Manhattan, NY
 ⁴ Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE

dengpfan@gmail.com, zhajyingjie@mail.nankai.edu.cn, aliborji@gmail.com,

yangjufeng@nankai.edu.cn, ling.shao@ieee.org
 https://github.com/zyjwuyan/BBS-Net

Abstract. Multi-level feature fusion is a fundamental topic in computer vision for detecting, segmenting and classifying objects at various scales. When multi-level features meet multi-modal cues, the optimal fusion problem becomes a hot potato. In this paper, we make the first attempt to leverage the inherent multi-modal and multi-level nature of RGB-D salient object detection to develop a novel cascaded refinement network. In particular, we 1) propose a bifurcated backbone strategy (BBS) to split the multi-level features into teacher and student features, and 2) utilize a depth-enhanced module (DEM) to excavate informative parts of depth cues from the channel and spatial views. This fuses RGB and depth modalities in a complementary way. Our simple yet efficient architecture, dubbed **B**ifurcated **B**ackbone **S**trategy **Net**work (**BBS-Net**), is backbone independent and outperforms 18 SOTAs on seven challenging datasets using four metrics.

Keywords: RGB-D saliency detection · bifurcated backbone strategy

1 Introduction

Multi-modal and multi-level feature fusion [37] is essential for many computer vision tasks, such as object detection [8,21,26,42,70], semantic segmentation [29, 30, 32, 67], co-attention tasks [19, 72] and classification [38, 40, 53]. Here, we attempt to utilize this idea for RGB-D salient object detection (SOD) [4,74], which aims at finding and segmenting the most visually prominent object(s) [2,75] in a scene according to the RGB and depth cues.

To efficiently integrate the RGB and depth cues for SOD, researchers have explored several multi-modal strategies [3,5], and have achieved encouraging results. Existing RGB-D SOD methods, however, still face the following challenges:

^{*} Equal contributions.



Fig. 1: Saliency maps of state-of-the-art (SOTA) CNN-based methods (*i.e.*, DM-RA [52], CPFP [74], TANet [4], and our *BBS-Net*) and methods based on hand crafted features (*i.e.*, SE [27] and LBE [24]). Our method generates higher quality saliency maps and suppresses background distractors for challenging scenarios (first row: complex background; second row: depth with noise) more effectively.

(1) Effectively aggregating multi-level features. As discussed in [44,63], teacher features⁵ provide discriminative semantic information that serves as strong guidance for locating salient objects, while student features carry affluent details that are beneficial for refining edges. Therefore, previous RGB-D SOD algorithms focus on leveraging multi-level features, either via a progressive merging process [47, 76] or by using a dedicated aggregation strategy [52, 74]. However, these operations directly fuse multi-level features without considering level-specific characteristics, and thus suffer from the inherent noise often introduced by low-level features [4, 65]. Thus, some methods tend to get distracted by the background (e.g., first row in Fig. 1).

(2) Excavating informative cues from the *depth modality*. Previous methods combine RGB and depth cues by regarding the depth map as a fourthchannel input [13, 51] or fusing RGB and depth modalities by simple summation [22, 23] or multiplication [9, 78]. These algorithms treat depth and RGB information the same and ignore the fact that depth maps mainly focus on the spatial relations among objects, whereas RGB information captures color and texture. Thus, such simple combinations are not efficient due to the modality difference. Besides, depth maps are sometimes low-quality, which may introduce feature noise and redundancy into the network. As an example, the depth map shown in the second row of Fig. 1 is blurry and noisy, and that is why many methods, including the top-ranked model (DMRA-iccv19 [52]), fail to detect the complete salient object.

To address these issues, we propose a novel Bifurcated Backbone Strategy **Net**work (*BBS-Net*) for RGB-D salient object detection. As shown in Fig. 2 (b), *BBS-Net* consists of two cascaded decoder stages. In the first stage, teacher features are aggregated by a standard cascaded decoder \mathbf{F}_{CD1} to generate an initial saliency map S_1 . In the second stage, student features are refined by an element-wise multiplication with the initial saliency map S_1 and are then integrated by another cascaded decoder \mathbf{F}_{CD2} to predict the final map S_2 .

 $^{^{\}rm b}$ Note that we use the terms 'high-level features & low-level features' and 'teacher features & student features' interchangeably.



Fig. 2: (a) Existing multi-level feature aggregation methods for RGB-D SOD [3,4,47,52, 62,74,76]. (b) In this paper, we propose to adopt a bifurcated backbone strategy (BBS) to split the multi-level features into student and teacher features. The initial saliency map S_1 is utilized to refine the student features to effectively suppress distractors. Then, the refined features are passed to another cascaded decoder to generate the final saliency map S_2 .

To the best of our knowledge, *BBS-Net* is the first work to explore the cascaded refinement mechanism for the RGB-D SOD task. Our main **contributions** are as follows:

- 1 We exploit multi-level features in a bifurcated backbone strategy (BBS) to suppress distractors in the lower layers. This strategy is based on the observation that high-level features provide discriminative semantic information without redundant details [44,65], which may contribute significantly to eliminating distractors in lower layers.
- 2 To fully capture the informative cues in the depth map and improve the compatibility of RGB and depth features, we introduce a depth-enhanced module (DEM), which contains two sequential attention mechanisms: *channel attention* and *spatial attention*. The channel attention utilizes the interchannel relations of the depth features, while the spatial attention aims to determine where informative depth cues are carried.
- 3 We demonstrate that the proposed *BBS-Net* exceeds 18 SOTAs on seven public datasets, by a large margin. Our experiments show that our framework has strong scalability in terms of various backbones. This means that the bifurcated backbone strategy via a cascaded refinement mechanism is promising for multi-level and multi-modal learning tasks.

2 Related Works

Although RGB-based SOD has been thoroughly studied in recent years [7, 39, 60, 69, 71], most of algorithms fail under complicated scenarios (*e.g.*, cluttered backgrounds [16], low-intensity environments, or varying illuminations) [4, 52]. As a complementary modality to RGB information, depth cues contain rich spatial distance information [52] and contribute significantly to understanding challenging scenes. Therefore, researchers have started to solve the SOD problem by combining RGB images with depth information [15].

Traditional Models. Previous RGB-D SOD algorithms mainly focused on hand crafted features [9,78]. Some of these methods largely relied on contrastbased cues by calculating color, edge, texture and region contrast to measure the saliency in a local region. For example, [15] adopted the region based contrast to compute contrast strengths for the segmented regions. In [10], the saliency value of each pixel depended on the color contrast and surface normals. However, the local contrast methods focued on the boundaries of salient objects and were easily affected by high-frequency content [54]. Therefore, some algorithms, such as global contrast [11], spatial prior [9], and background prior [56], proposed to calculate the saliency by combining local and global information. To effectively combine saliency cues from the RGB and depth modalities, researchers have explored various fusion strategies. Some methods [13, 51] regarded depth images as the fourth-channel inputs and processed the RGB and depth channes together (early fusion). This operation seems simple but disregards the differences between the RGB and depth modalities and thus cannot achieve reliable results. Therefore, to effectively extract the saliency information from the two modalities separately, some algorithms [22, 78] first leveraged two backbones to predict saliency maps and then fused the saliency results (late fusion). Besides, considering that the RGB and depth modalities may positively influence each other, yet other methods [24,34] fused RGB and depth features in a middle stage and then predicted the saliency maps from the fused features (middle fusion). In fact, these three fusion strategies are also explored in the current deep models, and our model can be considered as a middle fusion.

Deep Models. Early deep algorithms [54,56] extracted hand crafted features first, and then fed them to CNNs to compute saliency confidence scores. However, these methods need to design low-level features first and cannot be trained in an end-to-end manner. More recently, researchers have exploited CNNs to extract RGB and depth features in a bottom-up way [28]. Compared with hand crafted features, deep features contain more semantic and contextual information that can better capture representations of the RGB and depth modalities and achieve encouraging performance. The success of these deep models [5, 52] stems from two aspects of feature fusion. The first is the extracting of multi-scale features from different layers and then the effective fusion of these features. The second is the mechanism of fusing features from two different modalities.

To effectively aggregate multi-scale features, researchers have designed various network architectures. For example, [47] fed a four-channel RGB-D image into a single backbone and then obtained saliency map outputs from each sideout features (single stream). Chen *et al.* [3] leveraged two networks to extract RGB features and depth features respectively, and then fused them in a progressive complementary way (double stream). Further, to exploit cross-modal complements in the bottom-up feature extraction process, Chen *et al.* [4] proposed a three-stream network that contains two modality-specific streams and a parallel cross-modal distillation stream to learn supplementary features (three streams). However, depth maps are often of low quality and thus may contain a lot of noise and misleading information. This greatly decreases the performance



Fig. 3: The architecture of *BBS-Net*. Feature Extraction: '*Conv1*'~'*Conv5*' denote different layers from ResNet-50 [31]. Multi-level features $(f_1^d \sim f_5^d)$ from the depth branch are enhanced by the (a) DEM and then fused with features (*i.e.*, $f_1^{rgb} \sim f_5^{rgb}$) from the RGB branch. Stage 1: cross-modal teacher features $(f_3^{cm} \sim f_5^{cm})$ are first aggregated by the (b) cascaded decoder to produce the initial saliency map S_1 . Stage 2: Then, student features $(f_1^{cm} \sim f_3^{cm})$ are refined by the initial saliency map S_1 and are integrated by another cascaded decoder to predict the final saliency map S_2 .

of SOD models. To address this problem, Zhao *et al.* [74] designed a contrastenhanced network to improve the quality of depth maps by the contrast prior. Fan *et al.* [20] proposed a depth depurator unit that can evaluate the quality of the depth images and then filter out the low-quality maps automatically. Two recent famous works have also explored uncertainty [68], bilateral attention [73], graph neural network [48] and a joint learning strategy [25] and achieve good performance.

3 Proposed Method

3.1 Overview

Existing popular RGB-D SOD models directly aggregate multi-level features (Fig. 2(a)). As shown in Fig. 3, the network flow of our *BBS-Net* is different from the above mentioned models. We first introduce the bifurcated backbone strategy with the cascaded refinement mechanism in § 3.2. To fully use informative cues in the depth map, we introduce a new depth-enhanced module (§ 3.3).

3.2 Bifurcated Backbone Strategy (BBS)

We propose to excavate the rich semantic information in high-level cross-modal features to suppress background distractors in a cascaded refinement way. We adopt a bifurcated backbone strategy (BBS) to divide the multi-level cross-modal features into two groups, *i.e.*, $\mathbf{Q}_1 = \{Conv1, Conv2, Conv3\}$ and $\mathbf{Q}_2 = \{Conv3, Conv3, Conv3, Conv3\}$

Conv4, *Conv5*}, with the *Conv3* as the split point. Each group still preserves the original multi-scale information.

Cascaded Refinement Mechanism. To effectively leverage the features of the two groups, the whole network is trained with a cascaded refinement mechanism. This mechanism first produces an initial saliency map with three cross-modal teacher features (*i.e.*, \mathbf{Q}_2) and then improves the details of the initial saliency map S_1 with three cross-modal student features (*i.e.*, \mathbf{Q}_1) refined by the initial saliency map itself. Using this mechanism, our model can iteratively refine the details in the low-level features. This is based on the observation that high-level features contain rich global contextual information which is beneficial for locating salient objects, while low-level features carry much detailed micro-level information that can contribute significantly to refining the boundaries. In other words, this strategy efficiently eliminates noise in low-level cross-modal features, by exploring the specialties of multi-level features, and predicts the final saliency map in a progressive refinement manner.

Specifically, we first compute cross-modal features $\{f_i^{cm}; i = 1, 2, ..., 5\}$ by merging RGB and depth features processed by the DEM (Fig. 3(a)). In stage one, the three cross-modality teacher features $(i.e., f_3^{cm}, f_4^{cm}, f_5^{cm})$ are aggregated by the first cascaded decoder, which is formulated:

$$S_1 = \mathbf{T}_1 \big(\mathbf{F}_{CD1}(f_3^{cm}, f_4^{cm}, f_5^{cm}) \big), \tag{1}$$

where S_1 is the initial saliency map, \mathbf{F}_{CD1} is the first cascaded decoder and \mathbf{T}_1 represents two simple convolutional layers that change the channel number from 32 to 1. In stage two, the initial saliency map S_1 is leveraged to refine the three cross-modal student features, which is defined as:

$$f_i^{cm'} = f_i^{cm} \odot S_1, \tag{2}$$

where $f_i^{cm'}$ $(i \in \{1, 2, 3\})$ denotes the refined features and \odot represents the element-wise multiplication. Then, the three refined student features are integrated by another decoder followed by a progressively transposed module (PT-M), which is defined as,

$$S_2 = \mathbf{T}_2 \Big(\mathbf{F}_{CD2}(f_1^{cm'}, f_2^{cm'}, f_3^{cm'}) \Big),$$
(3)

where S_2 is the final saliency map. \mathbf{T}_2 represents the PTM module and \mathbf{F}_{CD2} denotes the second cascaded decoder. Finally, we jointly optimize the two stages by defining the total loss:

$$\mathcal{L} = \alpha \ell_{ce}(S_1, G) + (1 - \alpha) \ell_{ce}(S_2, G), \tag{4}$$

in which ℓ_{ce} represents the widely used binary cross entropy loss and $\alpha \in [0, 1]$ controls the trade-off between the two parts of the losses. The ℓ_{ce} is computed as:

$$\ell_{ce}(S,G) = G\log S + (1-G)\log(1-S),$$
(5)

in which S is the predicted saliency map and G denotes the ground-truth binary saliency map.

Cascaded Decoder. Given the two groups of multi-level, cross-modal features $(\{f_i^{cm}, f_{i+1}^{cm}, f_{i+2}^{cm}\}, i \in \{1, 3\})$ fused by the RGB and depth features from different layers, we need to efficiently utilize the multi-level, multi-scale information in each group to carry out our cascaded refinement strategy. Thus, we introduce a light-weight cascaded decoder [65] to aggregate the two groups of multi-level, cross-modal features. As shown in Fig. 3(b), the cascaded decoder contains three global context modules (GCM) and a simple feature aggregation strategy. The GCM is refined from the RFB module [46] with an additional branch to enlarge the receptive field and a residual connection [31] to preserve the original information. Specifically, as illustrated in Fig. 3(c), the GCM module contains four parallel branches. For all of these branches, a 1×1 convolution is first applied to reduce the channel size to 32. For the k^{th} branch $(k \in \{2, 3, 4\})$, a convolution operation with a kernel size of 2k - 1 and dilation rate of 1 is applied. This is followed by another 3×3 convolution operation with a dilation rate of 2k-1. The goal here is to extract the global contextual information from the cross-modal features. Next, the outputs of the four branches are concatenated together and their channel number is reduced to 32 with a 1×1 convolution operation. Finally, the concatenated features form a residual connection with the input feature. The outputs of the GCM modules in the two cascaded decoders are defined by:

$$f_i^{gcm} = \mathbf{F}_{GCM}(f_i),\tag{6}$$

To further improve the representational ability of cross-modal features, we leverage a pyramid multiplication and concatenation feature aggregation strategy to integrate the cross-modal features ($\{f_i^{gem}, f_{i+1}^{gem}, f_{i+2}^{gem}\}, i \in \{1, 3\}$). As shown in Fig. 3(b), first, each refined feature f_i^{gem} is updated by multiplying it with all higher-level features:

$$f_i^{gcm'} = f_i^{gcm} \odot \Pi_{k=i+1}^{k_{max}} Conv \Big(\mathbf{F}_{UP} \left(f_k^{gcm} \right) \Big), \tag{7}$$

where $i \in \{1, 2, 3\}$, $k_{max} = 3$ or $i \in \{3, 4, 5\}$, $k_{max} = 5$. $Conv(\cdot)$ represents the standard 3×3 convolution operation, and \mathbf{F}_{UP} denotes the upsampling operation if these features are not in the same scale. \odot represents the element-wise multiplication. Second, the updated features are aggregated by a progressive concatenation strategy to generate the output:

$$S = \mathbf{T}\left(\left[f_k^{gcm'}; Conv\left(\mathbf{F}_{UP}\left[f_{k+1}^{gcm'}; Conv\left(\mathbf{F}_{UP}(f_{k+2}^{gcm'})\right)\right]\right)\right]\right),$$
(8)

where S is the generated salient map, $k \in \{1,3\}$, and [x;y] denotes the concatenation operation of x and y. In the first stage, **T** represents two sequential convolutional layers (**T**₁), while it denotes the PTM module (**T**₂) for the second stage. The output (88×88) of the second decoder is 1/4 of the ground-truth resolution (352×352), so directly up-sampling the output to the ground-truth size will result in a loss of some details. To address this problem, we design a simple yet effective progressively transposed module (PTM, Fig. 3(d)) to predict the final saliency map (S₂) in a progressive upsampling way. It is composed of two

sequential residual-based transposed blocks [33] and three sequential 1×1 convolutions. Each residual-based transposed block consists of a 3×3 convolution and a residual-based transposed convolution.

Note that our cascaded refinement mechanism is different from the recent refinement mechanisms R3Net [14], CRN [6], and RFCN [61] in its usage of multilevel features and the initial map. The obvious difference and superiority of our design is that we only need one round of saliency refinement to obtain a good performance, while R3Net, CRN, and RFCN all need more iterations, which will increase the training time and computational resources. In addition, our cascaded strategy is different from CPD [65] in that it exploits the details in low-level features and semantic information in high-level features, while suppressing the noise in low-level features, simultaneously.

3.3 Depth-Enhanced Module (DEM)

There are two main problems when trying to fuse RGB and depth features. One is the compatibility of the two due to the intrinsic modality difference, and the other is the redundancy and noise in low-quality depth features. Inspired by [64], we introduce a depth-enhanced module (DEM) to improve the compatibility of multi-modal features and to excavate the informative cues from the depth features.

Specifically, let f_i^{rgb} , f_i^d denote the feature maps of the i^{th} $(i \in 1, 2, ..., 5)$ side-out layer from the RGB and depth branches, respectively. As shown in Fig. 3, each DEM is added before each side-out feature map from the depth branch to improve the compatibility of the depth features. Such a side-out process enhances the saliency representation of depth features and preserves the multi-level, multi-scale information. The fusion process of the two modalities is formulated as:

$$f_i^{cm} = f_i^{rgb} + \mathbf{F}_{DEM}(f_i^d), \tag{9}$$

where f_i^{cm} represents the cross-modal features of the i^{th} layer. As illustrated in Fig. 3(a), the DEM includes a sequential channel attention operation and a spatial attention operation. The operation of the DEM is defined as:

$$\mathbf{F}_{DEM}(f_i^d) = \mathbf{S}_{att} \Big(\mathbf{C}_{att}(f_i^d) \Big), \tag{10}$$

where $\mathbf{C}_{att}(\cdot)$ and $\mathbf{S}_{att}(\cdot)$ denote the spatial and channel attention, respectively. More specifically,

$$\mathbf{C}_{att}(f) = \mathbf{M} \Big(\mathbf{P}_{max}(f) \Big) \otimes f, \tag{11}$$

where $\mathbf{P}_{max}(\cdot)$ represents the global max pooling operation for each feature map, f denotes the input feature map, $\mathbf{M}(\cdot)$ is a multi-layer (two-layer) perceptron, and \otimes denotes the multiplication by the dimension broadcast. The spatial attention is implemented as:

$$\mathbf{S}_{att}(f) = Conv \Big(\mathbf{R}_{max}(f) \Big) \odot f, \tag{12}$$

where $\mathbf{R}_{max}(\cdot)$ represents the global max pooling operation for each point in the feature map along the channel axis. Our depth enhanced module is different from previous RGB-D models. Previous models fuse the corresponding multilevel features from the RGB and depth branches by direct concatenation [3,4,76], enhance the depth map by contrast prior [74] or process the multi-level depth features by a simple convolutional layer [52]. To the best of our knowledge, we are the first to introduce the attention mechanism to excavate informative cues from depth features in multiple side-out layers. Our experiments (see Tab. 4 and Fig. 5) show the effectiveness of this approach in improving the compatibility of multi-modal features.

Moreover, the spatial and channel attention mechanisms are different from the operation proposed in [64]. We only leverage a single global max pooling [50] operation to excavate the most critical cues in the depth features and reduce the complexity of the module simultaneously, which is based on the intuition that SOD aims at finding the most important area in an image.

4 Experiments

4.1 Experimental Settings

Datasets. We tested seven challenging RGB-D SOD datasets, *i.e.*, NJU2K [34], NLPR [51], STERE [49], SIP [20], DES [9], LFSD [41], and SSD [77].

Training/Testing. Following the same training settings as in [52, 74], we use 1, 485 samples from the NJU2K dataset and 700 samples from the NLPR dataset as our training set. The remaining images in the NJU2K and NLPR datasets and the whole datasets of STERE, DES, LFSD, SSD, and SIP are used for testing. **Evaluation Metrics.** We adopt four widely used metrics, including S-measure (S_{α}) [17], maximum E-measure (E_{ξ}) [18], maximum F-measure (F_{β}) [1], mean absolute error (MAE). Evaluation code: http://dpfan.net/d3netbenchmark/. **Contenders.** We compare the proposed *BBS-Net* with ten models that use hand crafted features [9, 12, 24, 27, 34, 43, 51, 55, 58, 78] and eight models [3–5, 28, 52, 54, 62, 74] based on deep learning. We train and test the above models using their default settings, as proposed in the original papers. For those models without released source codes, we used their published results for comparisons. **Inference Time.** In terms of speed, *BBS-Net* achieves 14fps and 48fps on a single GTX 1080Ti with a batch size of 1 and 10, respectively.

Implementation Details. We perform our experiments using the PyTorch [59] framework on a single 1080Ti GPU. Parameters of the backbone network (ResNet-50 [31]) are initialized from the model pre-trained on ImageNet [36]. We discard the last pooling and fully connected layers of ResNet-50 and leverage each middle output of the five convolutional blocks as the side-out feature maps. The two branches do not share weights and the only difference between them is that the depth branch has the input channel number set to 1. The other parameters are initialized as the PyTorch default settings. We use the Adam algorithm [35] to optimize the proposed model. The initial learning rate is set to 1e-4 and is

Table 1: Quantitative comparison of models using S-measure (S_{α}) , max F-measure (F_{β}) , max E-measure (E_{ξ}) and MAE (M) scores on seven datasets. $\uparrow (\downarrow)$ denotes that the higher (lower) the better. The best score in each row is highlighted in **bold**.

					Hand	l-craft	ed-Feat	ures-B	ased	Models					Cl	Ns-Ba	sed M	odels			
Da	taset	Metric	LHM	CDB	DESM	GP	CDCP	ACSD	LBE	DCMC	MDSF	SE	DF	AFNet	CTMF	MMCI	PCF	TANet	CPFP	DMRA	BBS-Net
			[51]	[43]	[9]	[55]	[78]	[34]	[24]	[12]	[58]	[27]	[54]	[62]	[28]	[5]	[3]	[4]	[74]	[52]	Ours
м		$S_{\alpha} \uparrow$.514	.624	.665	.527	.669	.699	.695	.686	.748	.664	.763	.772	.849	.858	.877	.878	.879	.886	.921
<u>1</u> 2	[24]	$F_{\beta} \uparrow$.632	.648	.717	.647	.621	.711	.748	.715	.775	.748	.804	.775	.845	.852	.872	.874	.877	.886	.920
Ē.	[<mark>04</mark>]	$E_{\xi} \uparrow$.724	.742	.791	.703	.741	.803	.803	.799	.838	.813	.864	.853	.913	.915	.924	.925	.926	.927	.949
~		$M\downarrow$.205	.203	.283	.211	.180	.202	.153	.172	.157	.169	.141	.100	.085	.079	.059	.060	.053	.051	.035
~		$S_{\alpha} \uparrow$.630	.629	.572	.654	.727	.673	.762	.724	.805	.756	.802	.799	.860	.856	.874	.886	.888	.899	.930
NLPF	[51]	$F_{\beta} \uparrow$.622	.618	.640	.611	.645	.607	.745	.648	.793	.713	.778	.771	.825	.815	.841	.863	.867	.879	.918
	[01]	$E_{\xi} \uparrow$.766	.791	.805	.723	.820	.780	.855	.793	.885	.847	.880	.879	.929	.913	.925	.941	.932	.947	.961
		$M\downarrow$.108	.114	.312	.146	.112	.179	.081	.117	.095	.091	.085	.058	.056	.059	.044	.041	.036	.031	.023
TERE		$S_{\alpha} \uparrow$.562	.615	.642	.588	.713	.692	.660	.731	.728	.708	.757	.825	.848	.873	.875	.871	.879	.835	.908
	[40]	$F_{\beta} \uparrow$.683	.717	.700	.671	.664	.669	.633	.740	.719	.755	.757	.823	.831	.863	.860	.861	.874	.847	.903
	[10]	$E_{\xi} \uparrow$.771	.823	.811	.743	.786	.806	.787	.819	.809	.846	.847	.887	.912	.927	.925	.923	.925	.911	.942
0,1		$M\downarrow$.172	.166	.295	.182	.149	.200	.250	.148	.176	.143	.141	.075	.086	.068	.064	.060	.051	.066	.041
		$S_{\alpha} \uparrow$.578	.645	.622	.636	.709	.728	.703	.707	.741	.741	.752	.770	.863	.848	.842	.858	.872	.900	.933
ES	[0]	$F_{\beta} \uparrow$.511	.723	.765	.597	.631	.756	.788	.666	.746	.741	.766	.728	.844	.822	.804	.827	.846	.888	.927
Ω	[•]	$E_{\xi} \uparrow$.653	.830	.868	.670	.811	.850	.890	.773	.851	.856	.870	.881	.932	.928	.893	.910	.923	.943	.966
_		$M\downarrow$.114	.100	.299	.168	.115	.169	.208	.111	.122	.090	.093	.068	.055	.065	.049	.046	.038	.030	.021
~		$S_{\alpha} \uparrow$.553	.515	.716	.635	.712	.727	.729	.753	.694	.692	.783	.738	.788	.787	.786	.801	.828	.839	.864
ŝ	[41]	$F_{\beta} \uparrow$.708	.677	.762	.783	.702	.763	.722	.817	.779	.786	.813	.744	.787	.771	.775	.796	.826	.852	.858
Ë	[11]	$E_{\xi} \uparrow$.763	.766	.811	.824	.780	.829	.797	.856	.819	.832	.857	.815	.857	.839	.827	.847	.863	.893	.901
		$M\downarrow$.218	.225	.253	.190	.172	.195	.214	.155	.197	.174	.145	.133	.127	.132	.119	.111	.088	.083	.072
		$S_{\alpha} \uparrow$.566	.562	.602	.615	.603	.675	.621	.704	.673	.675	.747	.714	.776	.813	.841	.839	.807	.857	.882
8	[77]	$F_{\beta} \uparrow$.568	.592	.680	.740	.535	.682	.619	.711	.703	.710	.735	.687	.729	.781	.807	.810	.766	.844	.859
ŝ	[]	$E_{\xi} \uparrow$.717	.698	.769	.782	.700	.785	.736	.786	.779	.800	.828	.807	.865	.882	.894	.897	.852	.906	.919
		$M\downarrow$.195	.196	.308	.180	.214	.203	.278	.169	.192	.165	.142	.118	.099	.082	.062	.063	.082	.058	.044
		$S_{\alpha} \uparrow$.511	.557	.616	.588	.595	.732	.727	.683	.717	.628	.653	.720	.716	.833	.842	.835	.850	.806	.879
£	[20]	$F_{\beta} \uparrow$.574	.620	.669	.687	.505	.763	.751	.618	.698	.661	.657	.712	.694	.818	.838	.830	.851	.821	.883
SI	[-0]	$E_{\xi} \uparrow$.716	.737	.770	.768	.721	.838	.853	.743	.798	.771	.759	.819	.829	.897	.901	.895	.903	.875	.922
		$M\downarrow$.184	.192	.298	.173	.224	.172	.200	.186	.167	.164	.185	.118	.139	.086	.071	.075	.064	.085	.055

divided by 10 every 60 epochs. We resize the input RGB and depth images to 352×352 for both the training and test phases. All the training images are augmented using multiple strategies (*i.e.*, random flipping, rotating and border clipping). It takes about 10 hours to train our model with a mini-batch size of 10 for 200 epochs.

4.2 Comparison with SOTAs

Quantitative Results. As shown in Tab. 1, our method performs favorably against all algorithms based on hand crafted features as well as SOTA CNN-based methods by a large margin, in terms of all four evaluation metrics. Performance gains over the best compared algorithms (ICCV'19 DMRA [52] and CVPR'19 CPFP [74]) are $(2.5\% \sim 3.5\%, 0.7\% \sim 3.9\%, 0.8\% \sim 2.3\%, 0.009 \sim 0.016)$ for the metrics $(S_{\alpha}, maxF_{\beta}, maxE_{\xi}, M)$ on seven challenging datasets.

Visual Comparison. Fig. 4 provides sample saliency maps predicted by the proposed method and several SOTA algorithms. Visualizations cover simple scenes (a) and various challenging scenarios, including small objects (b), multiple objects (c), complex backgrounds (d), and low-contrast scenes (e). First, (a) is an easy example. The flower in the foreground is evident in the original RGB image, but the depth image is low-quality and contains some misleading information. The top two models, *i.e.*, DMRA and CPFP, fail to predict the whole extent of the salient object due to the interference from the depth map. Our method can eliminate the side-effects of the depth map by utilizing the



Fig. 4: Qualitative visual comparison of the proposed model versus 8 SOTAs.

Table 2: Performance comparison using different backbones.

							0							
Models	NJU2 $ S_{\alpha}\uparrow$		$\stackrel{\text{NLPI}}{S_{\alpha}}\uparrow$	$\begin{array}{c} \mathbf{R} \ [51] \\ M \downarrow \end{array}$	${\scriptstyle {\rm STEF}\ S_{lpha}\uparrow}$	$\stackrel{\text{RE }[49]}{M \downarrow}$	$ DES S_{\alpha} \uparrow$	$\begin{bmatrix} 9 \\ M \end{bmatrix}$	$ $ LFSI $ S_{\alpha} \uparrow$	$\begin{bmatrix} 0 & [41] \\ M \downarrow \end{bmatrix}$	$SSD \\ S_{\alpha} \uparrow$	$\begin{bmatrix} 77\\ M \downarrow \end{bmatrix}$	$SIP \\ S_{\alpha} \uparrow$	$\begin{bmatrix} 20 \\ M \downarrow \end{bmatrix}$
CPFP [74]	.879	.053	.888	.036	.879	.051	.872	.038	.828	.088	.807	.082	.850	.064
DMRA ^[52]	.886	.051	.899	.031	.835	.066	.900	.030	.839	.083	.857	.058	.806	.085
BBS-Net(VGG-16)	.916	.039	.923	.026	.896	.046	.908	.028	.845	.080	.858	.055	.874	.056
BBS-Net(VGG-19)	.918	.037	.925	.025	.901	.043	.915	.026	.852	.074	.855	.056	.878	.054
BBS-Net(ResNet-50)	.921	.035	.930	.023	.908	.041	.933	.021	.864	.072	.882	.044	.879	.055

complementary depth information more effectively. Second, two examples of small objects are shown in (b). Despite the handle of the teapot in the first row being tiny, our method can accurately detect it. Third, we show two examples with multiple objects in the image in (c). Our method locates all salient objects in the image. It segments the objects better and generates sharper edges compared to other algorithms. Even though the depth map in the first row of (c) lacks clear information, our algorithm predicts the salient objects correctly. Fourth, (d) shows two examples with complex backgrounds. Here, our method produces reliable results, while other algorithms confuse the background as a salient object. Finally, (e) presents two examples in which the contrast between the object and background is low. Many algorithms fail to detect and segment the entire extent of the salient object. Our method produces satisfactory results by suppressing background distractors and exploring the informative cues from the depth map.

	rabie of com	parise				avaro	~88*	8		10008	,100 0		on a	arease	
#	Settings	$ NJU2 S_{\alpha}\uparrow$	$ \begin{array}{c} \mathrm{K} \ [34] \\ M \downarrow \end{array} $	$ $ NLPI $S_{\alpha} \uparrow$	$ \begin{array}{c} \mathbf{R} \ [51] \\ M \downarrow \end{array} $	$STEF S_{\alpha} \uparrow$	$ \begin{array}{c} \operatorname{RE} \ [49] \\ M \downarrow \end{array} $	$\begin{vmatrix} \text{DES} \\ S_{\alpha} \uparrow \end{vmatrix}$	$\begin{bmatrix} 9 \\ M \downarrow \end{bmatrix}$	$ $ LFSI $ S_{\alpha}\uparrow$	$\begin{bmatrix} 0 & [41] \\ M \downarrow \end{bmatrix}$	$\left \begin{array}{c} \mathrm{SSD} \\ S_{\alpha} \end{array} \right. \uparrow$	$\begin{bmatrix} 77 \\ M \downarrow \end{bmatrix}$	$SIP \\ S_{\alpha} \uparrow$	$\begin{bmatrix} 20 \\ M \downarrow \end{bmatrix}$
1	Low3	.881	.051	.882	.038	.832	.070	.853	.044	.779	.110	.805	.080	.760	.108
2	High3	.902	.042	.911	.029	.886	.048	.912	.026	.845	.080	.850	.058	.833	.073
3	All5	.905	.042	.915	.027	.891	.045	.901	.028	.845	.082	.848	.060	.839	.071
4	BBS-NoRF	.893	.050	.904	.035	.843	.072	.886	.039	.804.	.105	.839	.069	.843	.076
5	BBS-RH	.913	.040	.922	.028	.881	.054	.919	.027	.833	.085	.872	.053	.866	.063
6	BBS-RL (ours)	.921	.035	.930	.023	.908	.041	.933	.021	.864	.072	.882	.044	.879	.055

Table 3: Comparison of different feature aggregation strategies on seven datasets.

5 Discussion

Scalability. There are three popular backbone architectures (*i.e.*, VGG-16 [57], VGG-19 [57] and ResNet-50 [31]) that are used in deep RGB-D models. To further validate the scalability of the proposed method, we provide performance comparisons using different backbones in Tab. 2. We find that our *BBS-Net* exceeds the SOTA methods (*e.g.*, CPFP [74], and DMRA [52]) with all of these popular backbones, showing the strong scalability of our framework.

Aggregation Strategies. We conduct several experiments to validate the effectiveness of our cascaded refinement mechanism. Results are shown in Tab. 3 and Fig. 5(a). 'Low3' means that we only integrate the student features ($Conv1 \sim 3$) using the decoder without the refinement from the initial map for training and testing. Student features contain abundant details that are beneficial for refining the object edges, but at the same time introduce a lot of background distraction. Integrating only low-level features produces unsatisfactory results and generates many distractors (e.g., $1^{st}-2^{nd}$ row in Fig. 5(a)) or fails to locate the salient objects (e.g., the 3^{rd} row in Fig. 5(a)). 'High3' only integrates the teacher features ($Conv3 \sim 5$) using the decoder to predict the saliency map. Compared with student features, teacher features are 'sophisticated' and thus contain more semantic information. As a result, they help locate the salient objects and preserve edge information. Thus, integrating teacher features leads to better results. 'All5' aggregates features from all five levels ($Conv1 \sim 5$) directly using a single decoder for training and testing. It achieves comparable results with the 'High3' but may generate background noise introduced by the student features. 'BBS-NoRF' indicates that we directly remove the refinement flow of our model. This leads to poor performance. 'BBS-RH' can be seen as a reverse refinement strategy to our cascaded refinement mechanism, where teacher features ($Conv3 \sim 5$) are first refined by the initial map aggregated by student features ($Conv1 \sim 3$) and are then integrated to generate the final saliency map. It performs worse than our final mechanism (BBS-RL), because, with this reverse refinement strategy, noise in student features cannot be effectively suppressed. Besides, compared to 'All5', our method fully utilizes the features at different levels, and thus achieves significant performance improvement with fewer background distractors and sharper edges (*i.e.*, 'BBS-RL' in Fig. 5(a)).

Impact of Different Modules. As shown in Tab. 4 and Fig. 5(b), we conduct an ablation study to test the effectiveness of different modules in our *BBS-Net*. The base model (BM) is our *BBS-Net* without additional modules (*i.e.*, CA, SA,

Table 4: Ablation study of our *BBS-Net.* 'BM' = base model. 'CA' = channel attention. 'SA' = spatial attention. 'PTM' = progressively transposed module.

		1						1	0			1						
-#		Set	ting	s	NJU2	K [34]	NLP	R [51]	STEF	RE [49]	DES	5 [<mark>9</mark>]	LFSI) [41]	SSD	[77]	SIP	[20]
#	BM	CA	SA	PTM	$ S_{\alpha}\uparrow$	$M\downarrow$	$ S_{\alpha}\uparrow$	$M\downarrow$	$S_{\alpha} \uparrow$	$M\downarrow$	$S_{\alpha} \uparrow$	$M\downarrow$	$ S_{\alpha}\uparrow$	$M\downarrow$	$S_{\alpha} \uparrow$	$M\downarrow$	$ S_{\alpha}\uparrow$	$M\downarrow$
1	√				908	.045	.918	.029	.882	.055	.917	.027	.842	.083	.862	.057	.864	.066
2		\checkmark			.913	.042	.922	.027	.896	.048	.923	.025	.840	.086	.855	.057	.868	.063
3	\checkmark		\checkmark		.912	.045	.918	.029	.891	.054	.914	.029	.855	.083	.872	.054	.869	.063
4	\checkmark	\checkmark	\checkmark		.919	.037	.928	.026	.900	.045	.924	.024	.861	.074	.873	.052	.869	.061
5	\checkmark	\checkmark	\checkmark	\checkmark	.921	.035	.930	.023	.908	.041	.933	.021	.864	.072	.882	.044	.879	.055

Table 5: Effectiveness analysis of the cascaded decoder.

Strategies	$\begin{vmatrix} NJU2 \\ S_{\alpha} \uparrow \end{vmatrix}$	$\begin{array}{c} \mathrm{K} \ [34] \\ M \downarrow \end{array}$	$\begin{array}{c} \text{NLPI} \\ S_{\alpha} \uparrow \end{array}$	$\begin{array}{c} \mathbf{R} \ [51] \\ M \downarrow \end{array}$	$\begin{array}{c} \text{STER} \\ S_{\alpha} \uparrow \end{array}$	$ \begin{array}{c} \text{te} \ [49] \\ M \downarrow \end{array} $	$DES \\ S_{\alpha} \uparrow$	$\begin{bmatrix} 9 \\ M \end{bmatrix}$	LFSI $S_{\alpha} \uparrow$	$M \downarrow M$	$\left \begin{array}{c} \mathrm{SSD} \\ S_{\alpha} \end{array} \right. \uparrow$	$\begin{matrix} [77] \\ M \downarrow \end{matrix}$	$\begin{array}{c} \text{SIP} \\ S_{\alpha} \uparrow \end{array}$	$\begin{matrix} [20] \\ M \downarrow \end{matrix}$
Element-wise sum	.915	.037	.925	.025	.897	.045	.925	.022	.856	.073	.868	.050	.880	.052
Cascaded decoder	.921	.035	.930	.023	.908	.041	.933	.021	.864	.072	.882	.044	.879	.055



Fig. 5: (a): Visual comparison of different aggregation strategies, (b): Visual effectiveness of gradually adding modules. '#' denotes the corresponding row of Tab. 4.

and PTM). Note that just the BM performs better than the SOTA methods over almost all datasets, as shown in Tab. 1 and Tab. 4. Adding the channel attention (CA) and spatial attention (SA) modules enhances performance on most of the datasets. See the results shown in the second and third rows of Tab. 4. When we combine the two modules (fourth row in Tab. 4), the performance is greatly improved on all datasets, compared with the BM. We can easily conclude from the '#3' and '#4' columns in Fig. 5(b) that the spatial attention and channel attention mechanisms in DEM allow the model to focus on the informative parts of the depth features, which results in better suppression of background distraction. Finally, we add a progressively transposed block before the second decoder to gradually upsample the feature map to the same resolution as the ground truth. The results in the fifth row of Tab. 4 and the '#5' column of Fig. 5(b) show that the 'PTM' achieves impressive performance gains on all datasets and generates sharper edges with fine details.

To further analyze the effectiveness of the cascaded decoder, we experiment by changing the decoder to an element-wise summation mechanism. That is to say, we first change the features from different layers using 1×1 convolution and upsampling operation to the same dimension and then fuse them by element-

Table 6: S_{α} comparison with SOTA RGB SOD methods on three datasets. 'w/o D' and 'w/ D' represent training and testing the proposed method without/with the depth.

Methods	$\begin{array}{c} \text{CPD [65] PoolN} \\ S_{\alpha} \uparrow M \downarrow S_{\alpha} \uparrow \end{array}$	$\begin{bmatrix} 44 \\ M \downarrow \end{bmatrix} \begin{bmatrix} \text{PiC.} \\ S_{\alpha} \end{bmatrix}$	ANet [45] $M \downarrow$	$ PAGR S_{\alpha} \uparrow$	$M \downarrow M$	$\begin{array}{c} \text{R3Net } [14] \\ S_{\alpha} \uparrow M \downarrow \end{array}$	$\underset{S_{\alpha}}{\operatorname{Ours}}$	(w/o D) $M\downarrow$	$\begin{array}{c} \text{Ours} \\ S_{\alpha} \uparrow \end{array}$	$\begin{pmatrix} w/D \\ M \downarrow \end{pmatrix}$
NJU2K [34	.894 .046 .887	.045 .847	.071	.829	.081	.837 .092	.914	.038	.921	.035
NLPR [51]	.915 .025 .900	.029 .834	.053	.844	.051	.798 .101	.925	.026	.930	.023
DES $[9]$.897 .028 .873	.034 .854	.042	.858	.044	.847 .066	.912	.025	.933	.021

wise summation. Experimental results in Tab. 5 demonstrate the effectiveness of the cascaded decoder.

Benefits of the Depth Map. To explore whether or not the depth information can really contribute to the performance of SOD, we conduct two experiments in Tab. 6: (i) We compare the proposed method with five SOTA RGB SOD methods (*i.e.*, CPD [66], PoolNet [44], PiCANet [45], PAGRN [71] and R3Net [14]) when neglecting the depth information. We train and test these methods using the same training and testing sets as our *BBS-Net*. It is shown that the proposed methods (*i.e.*, Ours (w/ D)) can significantly exceed SOTA RGB SOD methods due to the usage of depth information. (ii) We train and test the proposed method without using the depth information by setting the inputs of the depth branch to zero (*i.e.*, Ours (w/o D)). Comparing the results of Ours (w/ D) with Ours (w/o D), we find that the depth information can effectively improve the performance of the proposed model. The above two experiments together demonstrate the benefits of the depth information for SOD, since depth maps can be seen as prior knowledge that provides spatial-distance information and contour guidance to detect salient objects.

6 Conclusion

We presented a new multi-level multi-modality learning framework that demonstrates state-of-the-art performance on seven challenging RGB-D salient object detection datasets using several evaluation measures. Our *BBS-Net* is based on a novel bifurcated backbone strategy (BBS) with a cascaded refinement mechanism. Importantly, our simple architecture is backbone independent, making it promising for further research on other related topics, including semantic segmentation, object detection and classification.

Acknowledgments

This work was supported by the Major Project for New Generation of AI Grant (NO. 2018AAA0100403), NSFC (NO.61876094, U1933114), Natural Science Foundation of Tianjin, China (NO.18JCYBJC15400, 18ZXZNGX00110), the Open Project Program of the National Laboratory of Pattern Recognition (NLPR), and the Fundamental Research Funds for the Central Universities.

References

- Achanta, R., Hemami, S., Estrada, F., Susstrunk, S.: Frequency-tuned salient region detection. In: CVPR. pp. 1597–1604 (2009)
- Borji, A., Cheng, M.M., Jiang, H., Li, J.: Salient object detection: A benchmark. IEEE TIP 24(12), 5706–5722 (2015)
- Chen, H., Li, Y.: Progressively complementarity-aware fusion network for RGB-D salient object detection. In: CVPR. pp. 3051–3060 (2018)
- Chen, H., Li, Y.: Three-stream attention-aware network for RGB-D salient object detection. IEEE TIP 28(6), 2825–2835 (2019)
- Chen, H., Li, Y., Su, D.: Multi-modal fusion network with multi-scale multi-path and cross-modal interactions for RGB-D salient object detection. IEEE TOC 86, 376–385 (2019)
- Chen, Q., Koltun, V.: Photographic image synthesis with cascaded refinement networks. In: CVPR. pp. 1511–1520 (2017)
- Chen, S., Tan, X., Wang, B., Lu, H., Hu, X., Fu, Y.: Reverse attention-based residual network for salient object detection. IEEE TIP 29, 3763–3776 (2020)
- Cheng, G., Han, J., Zhou, P., Xu, D.: Learning rotation-invariant and fisher discriminative convolutional neural networks for object detection. IEEE TIP 28(1) (2018)
- Cheng, Y., Fu, H., Wei, X., Xiao, J., Cao, X.: Depth enhanced saliency detection method. In: ICIMCS. pp. 23–27 (2014)
- Ciptadi, A., Hermans, T., M. Rehg, J.: An in depth view of saliency. In: BMVC (2013)
- Cong, R., Lei, J., Fu, H., Hou, J., Huang, Q., Kwong, S.: Going from RGB to RGBD saliency: A depth-guided transformation model. IEEE TOC pp. 1–13 (2019)
- Cong, R., Lei, J., Zhang, C., Huang, Q., Cao, X., Hou, C.: Saliency detection for stereoscopic images based on depth confidence analysis and multiple cues fusion. IEEE SPL 23(6), 819–823 (2016)
- Cong, R., Lei, J., Fu, H., Huang, Q., Cao, X., Ling, N.: HSCS: Hierarchical sparsity based co-saliency detection for RGBD images. IEEE TMM 21(7), 1660–1671 (2019)
- 14. Deng, Z., Hu, X., Zhu, L., Xu, X., Qin, J., Han, G., Heng, P.A.: R3net: Recurrent residual refinement network for saliency detection. In: IJCAI. pp. 684–690 (2018)
- Desingh, K., Krishna, K., Rajanand, D., Jawahar, C.: Depth really matters: Improving visual salient region detection with depth. In: BMVC. pp. 1–11 (2013)
- Fan, D.P., Cheng, M.M., Liu, J.J., Gao, S.H., Hou, Q., Borji, A.: Salient objects in clutter: Bringing salient object detection to the foreground. In: ECCV. pp. 186–202 (2018)
- 17. Fan, D.P., Cheng, M.M., Liu, Y., Li, T., Borji, A.: Structure-measure: A new way to evaluate foreground maps. In: ICCV. pp. 4548–4557 (2017)
- Fan, D.P., Gong, C., Cao, Y., Ren, B., Cheng, M.M., Borji, A.: Enhanced-alignment measure for binary foreground map evaluation. In: IJCAI. pp. 698–704 (2018)
- Fan, D.P., Lin, Z., Ji, G.P., Zhang, D., Fu, H., Cheng, M.M.: Taking a deeper look at co-salient object detection. In: CVPR. pp. 2919–2929 (2020)
- Fan, D.P., Lin, Z., Zhang, Z., Zhu, M., Cheng, M.M.: Rethinking RGB-D salient object detection: Models, datasets, and large-scale benchmarks. IEEE TNNLS (2020)
- Fan, D.P., Wang, W., Cheng, M.M., Shen, J.: Shifting more attention to video salient object detection. In: CVPR. pp. 8554–8564 (2019)

- 16 D.-P. Fan et al.
- Fan, X., Liu, Z., Sun, G.: Salient region detection for stereoscopic images. In: DSP. pp. 454–458 (2014)
- Fang, Y., Wang, J., Narwaria, M., Le Callet, P., Lin, W.: Saliency detection for stereoscopic images. IEEE TIP 23(6), 2625–2636 (2014)
- Feng, D., Barnes, N., You, S., McCarthy, C.: Local background enclosure for RGB-D salient object detection. In: CVPR. pp. 2343–2350 (2016)
- Fu, K., Fan, D.P., Ji, G.P., Zhao, Q.: JL-DCF: Joint learning and denselycooperative fusion framework for RGB-D salient object detection. In: CVPR. pp. 3052–3062 (2020)
- Gao, S.H., Tan, Y.Q., Cheng, M.M., Lu, C., Chen, Y., Yan, S.: Highly efficient salient object detection with 100k parameters. In: ECCV (2020)
- Guo, J., Ren, T., Bei, J.: Salient object detection for RGB-D image via saliency evolution. In: ICME. pp. 1–6 (2016)
- Han, J., Chen, H., Liu, N., Yan, C., Li, X.: CNNs-Based RGB-D saliency detection via cross-view transfer and multiview fusion. IEEE TOC 48(11), 3171–3183 (2018)
- Han, J., Yang, L., Zhang, D., Chang, X., Liang, X.: Reinforcement cutting-agent learning for video object segmentation. In: CVPR. pp. 9080–9089 (2018)
- Han, Q., Zhao, K., Xu, J., Cheng, M.M.: Deep hough transform for semantic line detection. In: ECCV) (2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
- He, X., Yang, S., Li, G., Li, H., Chang, H., Yu, Y.: Non-local context encoder: Robust biomedical image segmentation against adversarial attacks. In: AAAI019. pp. 8417–8424 (2019)
- Hu, X., Yang, K., Fei, L., Wang, K.: Acnet: Attention based network to exploit complementary features for RGBD semantic segmentation. In: ICIP. pp. 1440–1444 (2019)
- Ju, R., Ge, L., Geng, W., Ren, T., Wu, G.: Depth saliency based on anisotropic center-surround difference. In: ICIP. pp. 1115–1119 (2014)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS. pp. 1106–1114 (2012)
- Li, G., Yu, Y.: Visual saliency based on multiscale deep features. In: CVPR. pp. 5455–5463 (2015)
- Li, G., Zhu, X., Zeng, Y., Wang, Q., Lin, L.: Semantic relationships guided representation learning for facial action unit recognition. In: AAAI. pp. 8594–8601 (2019)
- Li, H., Chen, G., Li, G., Yu, Y.: Motion guided attention for video salient object detection. In: ICCV. pp. 7274–7283 (2019)
- Li, J., Song, Y., Zhu, J., Cheng, L., Su, Y., Ye, L., Yuan, P., Han, S.: Learning from large-scale noisy web data with ubiquitous reweighting for image classification. IEEE TPAMI (2019)
- Li, N., Ye, J., Ji, Y., Ling, H., Yu, J.: Saliency detection on light field. In: CVPR. pp. 2806–2813 (2014)
- Li, X., Yang, F., Cheng, H., Liu, W., Shen, D.: Contour knowledge transfer for salient object detection. In: ECCV. pp. 355–370 (2018)
- Liang, F., Duan, L., Ma, W., Qiao, Y., Cai, Z., Qing, L.: Stereoscopic saliency model using contrast and depth-guided-background prior. Neurocomputing 275, 2227–2238 (2018)

- 44. Liu, J.J., Hou, Q., Cheng, M.M., Feng, J., Jiang, J.: A simple pooling-based design for real-time salient object detection. In: CVPR. pp. 3917–3926 (2019)
- Liu, N., Han, J., Yang, M.H.: PiCANet: Learning pixel-wise contextual attention for saliency detection. In: CVPR. pp. 3089–3098 (2018)
- Liu, S., Huang, D., Wang, Y.: Receptive field block net for accurate and fast object detection. In: ECCV. pp. 404–419 (2018)
- 47. Liu, Z., Shi, S., Duan, Q., Zhang, W., Zhao, P.: Salient object detection for RGB-D image by single stream recurrent convolution neural network. Neurocomputing 363, 46–57 (2019)
- 48. Luo, A., Li, X., Yang, F., Jiao, Z., Cheng, H., Lyu, S.: Cascade graph neural networks for rgb-d salient object detection. In: ECCV (2020)
- Niu, Y., Geng, Y., Li, X., Liu, F.: Leveraging stereopsis for saliency analysis. In: CVPR. pp. 454–461 (2012)
- Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Is object localization for free? weaklysupervised learning with convolutional neural networks. In: CVPR. pp. 685–694 (2015)
- 51. Peng, H., Li, B., Xiong, W., Hu, W., Ji, R.: RGBD salient object detection: a benchmark and algorithms. In: ECCV. pp. 92–109. Springer (2014)
- 52. Piao, Y., Ji, W., Li, J., Zhang, M., Lu, H.: Depth-induced multi-scale recurrent attention network for saliency detection. In: ICCV. pp. 7254–7263 (2019)
- 53. Qiao, L., Shi, Y., Li, J., Wang, Y., Huang, T., Tian, Y.: Transductive episodic-wise adaptive metric for few-shot learning. In: ICCV. pp. 3603–3612 (2019)
- 54. Qu, L., He, S., Zhang, J., Tian, J., Tang, Y., Yang, Q.: RGBD salient object detection via deep fusion. IEEE TIP 26(5), 2274–2285 (2017)
- Ren, J., Gong, X., Yu, L., Zhou, W., Ying Yang, M.: Exploiting global priors for RGB-D saliency detection. In: CVPRW. pp. 25–32 (2015)
- Shigematsu, R., Feng, D., You, S., Barnes, N.: Learning RGB-D salient object detection using background enclosure, depth contrast, and top-down features. In: ICCVW. pp. 2749–2757 (2017)
- 57. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
- Song, H., Liu, Z., Du, H., Sun, G., Le Meur, O., Ren, T.: Depth-aware salient object detection and segmentation via multiscale discriminative saliency fusion and bootstrap learning. IEEE TIP 26(9), 4204–4216 (2017)
- 59. Steiner, B., DeVito, Z., Chintala, S., Gross, S., Paszke, A., Massa, F., Lerer, A., Chanan, G., Lin, Z., Yang, E., et al.: PyTorch: An imperative style, highperformance deep learning library. In: NIPS. pp. 8024–8035 (2019)
- Su, J., Li, J., Zhang, Y., Xia, C., Tian, Y.: Selectivity or invariance: Boundaryaware salient object detection. In: ICCV. pp. 3798–3807 (2019)
- Wang, L., Wang, L., Lu, H., Zhang, P., Ruan, X.: Salient object detection with recurrent fully convolutional networks. IEEE TPAMI 41(7), 1734–1746 (2018)
- Wang, N., Gong, X.: Adaptive fusion for RGB-D salient object detection. IEEE Access 7, 55277–55284 (2019)
- Wang, T., Zhang, L., Wang, S., Lu, H., Yang, G., Ruan, X., Borji, A.: Detect globally, refine locally: A novel approach to saliency detection. In: CVPR. pp. 3127–3135 (2018)
- Woo, S., Park, J., Lee, J.Y., So Kweon, I.: CBAM: Convolutional block attention module. In: ECCV. pp. 3–19 (2018)
- Wu, Z., Su, L., Huang, Q.: Cascaded partial decoder for fast and accurate salient object detection. In: CVPR. pp. 3907–3916 (2019)

- 18 D.-P. Fan et al.
- Wu, Z., Su, L., Huang, Q.: Stacked cross refinement network for edge-aware salient object detection. In: ICCV. pp. 7264–7273 (2019)
- Zeng, Y., Zhuge, Y., Lu, H., Zhang, L.: Joint learning of saliency detection and weakly supervised semantic segmentation. In: ICCV. pp. 7223–7233 (2019)
- Zhang, J., Fan, D.P., Dai, Y., Anwar, S., Saleh, F.S., Zhang, T., Barnes, N.: UC-Net: uncertainty inspired RGB-D saliency detection via conditional variational autoencoders. In: CVPR. pp. 8582–8591 (2020)
- Zhang, L., Wu, J., Wang, T., Borji, A., Wei, G., Lu, H.: A multistage refinement network for salient object detection. IEEE TIP 29, 3534–3545 (2020)
- Zhang, Q., Huang, N., Yao, L., Zhang, D., Shan, C., Han, J.: Rgb-t salient object detection via fusing multi-level cnn features. IEEE TIP 29, 3321–3335 (2020)
- Zhang, X., Wang, T., Qi, J., Lu, H., Wang, G.: Progressive attention guided recurrent network for salient object detection. In: CVPR. pp. 714–722 (2018)
- Zhang, Z., Jin, W., Xu, J., Cheng, M.M.: Gradient-induced co-saliency detection. In: ECCV (2020)
- 73. Zhang, Z., Lin, Z., Xu, J., Jin, W., Lu, S.P., Fan, D.P.: Bilateral attention network for RGB-D salient object detection. arXiv preprint arXiv:2004.14582 (2020)
- Zhao, J.X., Cao, Y., Fan, D.P., Cheng, M.M., Li, X.Y., Zhang, L.: Contrast prior and fluid pyramid integration for RGBD salient object detection. In: CVPR. pp. 3927–3936 (2019)
- Zhao, J.X., Liu, J.J., Fan, D.P., Cao, Y., Yang, J., Cheng, M.M.: EGNet: Edge guidance network for salient object detection. In: CVPR. pp. 8779–8788 (2019)
- Zhu, C., Cai, X., Huang, K., Li, T.H., Li, G.: Pdnet: Prior-model guided depthenhanced network for salient object detection. In: ICME. pp. 199–204 (2019)
- 77. Zhu, C., Li, G.: A three-pathway psychobiological framework of salient object detection using stereoscopic technology. In: ICCVW. pp. 3008–3014 (2017)
- Zhu, C., Li, G., Wang, W., Wang, R.: An innovative salient object detection using center-dark channel prior. In: ICCVW. pp. 1509–1515 (2017)