

# Disentangled High Quality Salient Object Detection

Lv Tang Bo Li<sup>1\*</sup> Yijie Zhong Shouhong Ding<sup>1</sup> Mofei Song<sup>2,3</sup>

<sup>1</sup>Youtu Lab, Tencent, Shanghai, China

<sup>2</sup>The School of Computer Science and Engineering,

<sup>3</sup>The Key Lab of Computer Network and Information Integration (Ministry of Education),  
Southeast University, Nanjing, China

luckybird1994@gmail.com, libraboli@tencent.com, dun.hasaki@gmail.com,

ericshding@tencent.com, songmf@seu.edu.cn

## Abstract

*Aiming at discovering and locating most distinctive objects from visual scenes, salient object detection (SOD) plays an essential role in various computer vision systems. Coming to the era of high resolution, SOD methods are facing new challenges. The major limitation of previous methods is that they try to identify the salient regions and estimate the accurate objects boundaries simultaneously with a single regression task at low-resolution. This practice ignores the inherent difference between the two difficult problems, resulting in poor detection quality. In this paper, we propose a novel deep learning framework for high-resolution SOD task, which disentangles the task into a low-resolution saliency classification network (LRSCN) and a high-resolution refinement network (HRRN). As a pixel-wise classification task, LRSCN is designed to capture sufficient semantics at low-resolution to identify the definite salient, background and uncertain image regions. HRRN is a regression task, which aims at accurately refining the saliency value of pixels in the uncertain region to preserve a clear object boundary at high-resolution with limited GPU memory. It is worth noting that by introducing uncertainty into the training process, our HRRN can well address the high-resolution refinement task without using any high-resolution training data. Extensive experiments on high-resolution saliency datasets as well as some widely used saliency benchmarks show that the proposed method achieves superior performance compared to the state-of-the-art methods.*

\*Corresponding author and equal contribution to first author. This work was supported by National Natural Science Foundation of China 61906036 and the Fundamental Research Funds for the Central Universities (2242021k30056).

## 1. Introduction

Salient object detection (SOD) is derived with the goal of accurately detecting and segmenting the most distinctive objects from visual scenes. As a preliminary step, it plays an essential role in various visual systems, such as video object segmentation [43], light field image segmentation [39], image-sentence matching [18], person re-identification [23] and instance segmentation [64].

Recently, the rapid development of the commodity imaging and display device, has resulted in higher requirements for the producing and editing of high-resolution (e.g., 720p, 1080p and 4K) images. Salient object detection as well as many state-of-the-art computer vision tasks are facing various challenges when encountering high-resolution scenarios. A good high-resolution salient object detection method should not only accurately detect the whole salient object but also predict the precise boundaries of salient objects. Despite the conventional Deep Neural Networks (DNNS) based SOD models have achieved remarkable performance at low-resolution (e.g., typical size  $224 \times 224$ ,  $384 \times 384$ ), they often fail to generate high quality detection results for high-resolution images. The major reason for this drawback is that the most previous methods try to identify the salient regions and estimate the accurate objects boundaries simultaneously in one step, which are two difficult and inherently different problems for high-resolution salient object detection. To address the first problem, a network is required to capture sufficient semantics by maintaining a larger receptive field. However, since the memory usage increases dramatically along with the image resolution, it is impractical for these models to directly learn sufficient semantics for high-resolution images. One plausible way is introducing downsample operations, but the structure details are inevitably lost during the downsampling, which however is precisely the key to solving the second problem.

Unfortunately, most of the existing low-resolution SOD

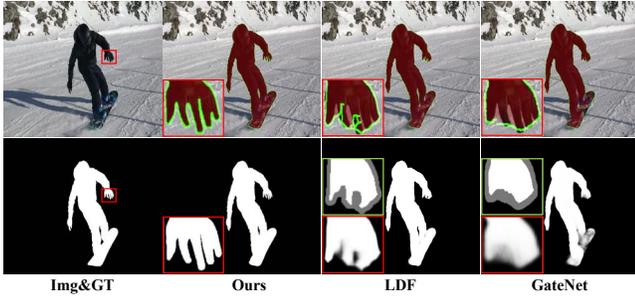


Figure 1. Comparison with state-of-the-art method in high-resolution SOD. Best viewed by zooming in.

methods [24, 57, 4, 47] try to address the aforementioned two problems with a single regression framework, which ignore the inherent difference between the two problems and result in blurry boundaries. As shown in Fig.1, if we take a deeper look at the saliency map generated by the representative existing methods LDF [47] and GateNet [62], we can observe that pixels can be divided into three different sets: (1) most pixels inside the salient object have the highest saliency value, and we call these pixels as *definite salient pixels*; (2) most pixels in the background regions have the lowest saliency value, which belong to *definite background pixels*; (3) saliency values of the pixels at blurry object boundaries fluctuate between 0 and 1, so we call these pixels as *uncertain pixels*. An ideal SOD method should effectively identify the definite salient and background regions in the image and accurately calculate the saliency value of pixels in the uncertain region to preserve a clear object boundary. From this perspective, there are essentially two tasks in SOD which demand quite different abilities to address the aforementioned two problems. The former task can be viewed as a classic classification task, while the later one is a typical regression task.

Despite the demand for effective high-resolution SOD methods, this line of work is rarely studied. In this paper, motivated by the new observation that SOD should be disentangled into two tasks, we propose a novel deep learning framework for high-resolution salient object detection. Specifically, we decouple the high-resolution salient object detection into a low-resolution saliency classification network (LRSCN) and a high-resolution refinement network (HRRN). LRSCN is designed to capture sufficient semantics at low-resolution and classify the pixels into three different sets for later process. HRRN aims at accurately refining the saliency value of pixels in the uncertain region to preserve a clear object boundary at high-resolution with limited GPU memory. As discussed above, HRRN requires structure details in high-resolution image. However, widely used low-resolution saliency datasets generally have some problems in annotation quality [52], making it almost impossible to directly obtain enough object boundary details from these defective datasets to train the high-resolution

network. In the very recent work, Zeng et al. [52] proposed to train their SOD network by using high-resolution images with accurate annotation. However, such high-quality image annotation requires heavy labor costs. In our paper, we argue that it is unnecessary to use such accurately annotated high-resolution images in network training. By introducing uncertainty [19] in the training process, our HRRN can well address the high-resolution refinement task only using the low-resolution training datasets with poor annotation.

Our major contributions can be summarized as:

- We provide a new perspective that high-resolution salient object detection should be disentangled into two tasks, and demonstrate that the disentanglement of the two tasks is essential for improving the performance of DNN based SOD models.
- Motivated by the principle of disentanglement, we propose a novel framework for high-resolution salient object detection, which uses LRSCN to capture sufficient semantics at low-resolution and HRRN for accurate boundary refinement at high-resolution.
- We make the earliest efforts to introduce the uncertainty into SOD network training, which empowers HRRN to well address the high-resolution refinement task without any high-resolution training datasets.
- We perform extensive experiments to demonstrate the proposed method refreshes the SOTA performance on high-resolution saliency datasets as well as some widely used saliency benchmarks by a large margin.

## 2. Related Work

Over the past decades, a large amount of SOD algorithms have been developed. Traditional models [17, 6, 36, 40, 20] detect salient objects by utilizing various heuristic saliency priors with hand-crafted features. More details about the traditional methods can be found in the survey [1]. Recently, with the development of deep learning, the performance of saliency detection has archived great improvement [24, 9, 26, 34, 65, 38, 59]. Here we mainly focus on deep learning based saliency detection models.

Recently, some DNN-based models use various feature enhancement strategies to improve the ability of localization and awareness of salient objects [16, 48, 35, 5, 46, 28, 62, 12], or take advantage of edge features to restore the structural details of salient objects [44, 60, 49, 63]. For example, Pang et al. [28] applied the transformation-interaction-fusion strategy on multi-level and multi-scale features to learn discriminant feature representation. Zhao et al. [62] designed a gated dual branch structure to build the cooperation among different levels of features and improve the discriminability of the whole network. In [60],

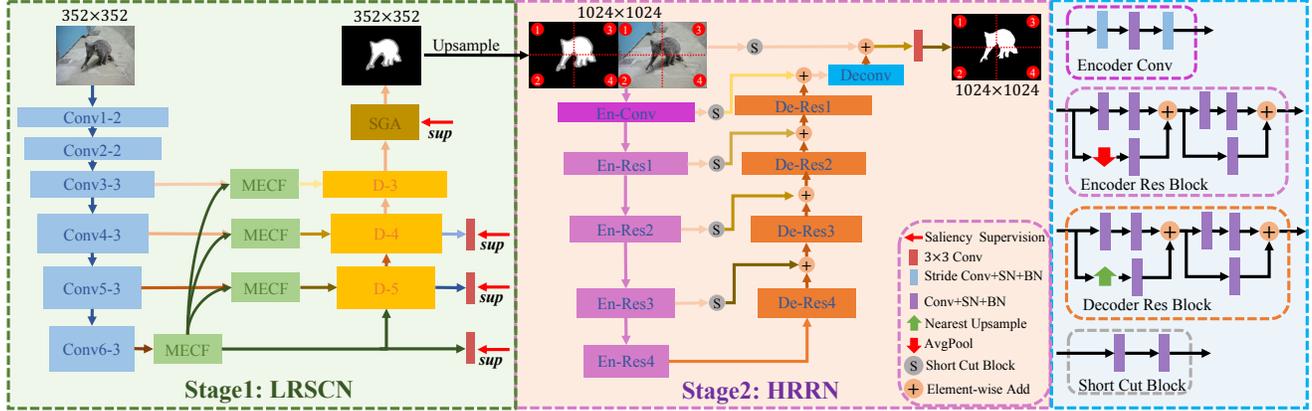


Figure 2. The framework of the proposed disentangled high quality salient object detection method.

edge features from edge detection branch was fused with salient features as complementary information to enhance the structural details for accurate saliency detection. Zhou et al. [63] used two individual branches for representing saliency and contour stream respectively, and a novel feature fusion module for their correlation combination.

Different from the above methods, some methods consider leveraging predict-refine architecture to generate fine salient objects. For example, Wang et al. [41] proposed to localize salient objects globally and then refine them by a local boundary refinement module. Qin et al. [31] was composed of an Encoder-Decoder network and a residual refinement module, which were respectively in charge of saliency prediction and saliency map refinement.

However, all these methods cannot handle high-resolution salient object detection problem well since such simple regression framework cannot identify the salient regions and estimate the accurate objects boundaries simultaneously and their architectures are not optimized for high-resolution SOD. Zeng et al. [52] tried to alleviate this problem by leveraging both global semantic information and local high-resolution details to accurately detect salient objects in high-resolution images. However, Zeng et al. [52] relies on high-resolution training images with accurate annotation, which requires heavy labor costs. Different from the above methods, we disentangle high-resolution SOD into two tasks at different resolutions: identifying the salient regions at low-resolution and estimating the accurate objects boundaries at high-resolution. Moreover, unlike Zeng et al. [52], we introduce novel uncertainty loss, which empowers our HRRN to well address the high-resolution refinement task without using any high-resolution training datasets. Recently, Wei et al. [47] and Zhang et al. [55] also leverage disentanglement in their SOD methods. However, they still try to address the SOD task under a single regression framework but with decoupled supervisions. Unlike our proposed methods, their disentanglement frameworks barely touch the very nature of the SOD, which essentially

contains two different tasks. For more information about the DNN-based methods, please refer to survey [42, 14].

### 3. Proposed Method

In this section, we first describe the overall architecture of the proposed disentangled high quality salient object detection network, then elaborate our main contributions, which are corresponding to LRSCN and HRRN.

#### 3.1. Network Overview

The architecture of the proposed approach is illustrated in Fig.2. As can be seen, the disentanglement includes two decoupled tasks at two different resolutions. LRSCN aims at capturing sufficient semantics at low-resolution and classifying the pixels into three different sets, which also can save the memory usage. While estimating the accurate objects boundaries needs more local details at high-resolution. So, we design HRRN to regress the saliency value of pixels and preserve a clear object boundary at high-resolution.

LRSCN has a simple U-Net like Encoder-Decoder architecture [32]. VGG-16 [33] is used as backbone. Following [16, 60], we connect another side path to the last pooling layer in VGG-16. Hence, we obtain six side features Conv1-2, Conv2-2, Conv3-3, Conv4-3, Conv5-3 and Conv6-3 from backbone network. Because Conv1-2 and Conv2-2 are too close to the input and their receptive fields are too small, following [60, 48], we only use the last four levels features for the following process. Conv6-3 is denoted as  $\{F_h|h = 6\}$ , the other three levels features are denoted as  $\{F_l|l = 3, 4, 5\}$ . Multi-scale feature extraction and Cross-level feature fusion (MECF) module is added between encoder and decoder to help improve the discriminability of feature representations. Decoder fuses the output features from MECF and the upsampled features from the previous stage in a bottom-up manner. The output of each decoder is defined as  $\{D_i|i = 3, 4, 5, 6\}$ . Finally, SGA module is built upon  $D_3$  for accurate trimap  $T$  generation.

As described, LRSCN is classification task and aims at

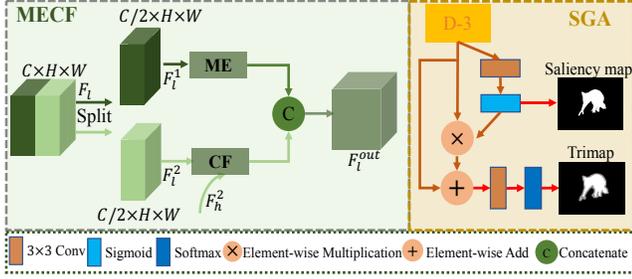


Figure 3. Architecture of MECF and SGA Modules.

capturing sufficient semantics at low-resolution. To regress a clear object boundary value, the input of HRRN is a high-resolution image under guidance of the trimap provided by LRSCN. HRRN has a basic Encoder-Decoder architecture and with the help of uncertainty loss, the network can be more robust to noisy data and predict a high-resolution saliency map with clear boundary.

### 3.2. Architecture of LRSCN

To capture sufficient semantics at low-resolution, learning discriminant feature representations is essential. The network should not only consider scale and location variations of different salient objects, but also distinguish the appearance difference between the salient object and the non-salient regions. To achieve the first goal, we develop a multi-scale feature extraction module (ME) based on Global Convolutional Network (GCN) [29] to enlarge the feature receptive field and obtain multi-scale information. To achieve the second goal, we utilize cross-level feature fusion module (CF) to leverage the advantages of features at different levels. Moreover, in designing the network architecture, inspired by [50], we use split-transform-merge strategy to further enlarge feature receptive fields and hence results in more discriminative feature representations. Specifically, we uniformly split the input  $F$  into two portions  $\{F^1, F^2\}$  by channel dimension, then  $F^1$  is sent into multi-scale feature extraction pathway and  $F^2$  is sent into cross-level feature fusion pathway. The outputs of these two pathways are concatenated together as the final output. We call this bridge module as **MECF module**, which is shown in Fig.3. More details about MECF module can be found in section 6 of supplementary materials.

**SGA Module.** As illustrated in Fig.2, each decoder fuses features from MECF module and previous decoder stage, then uses  $3 \times 3$  convolutional layer for final prediction. To maintain consistency between trimap and saliency map and ensure the uncertain regions of the trimap can accurately cover the boundary of saliency map, we design a saliency guide attention module (SGA) on  $D_3$ . Specifically, we first use a  $3 \times 3$  convolution and sigmoid function to compute a saliency map. Then, the saliency map is treated as spatial weight map which can help refine feature and generate an

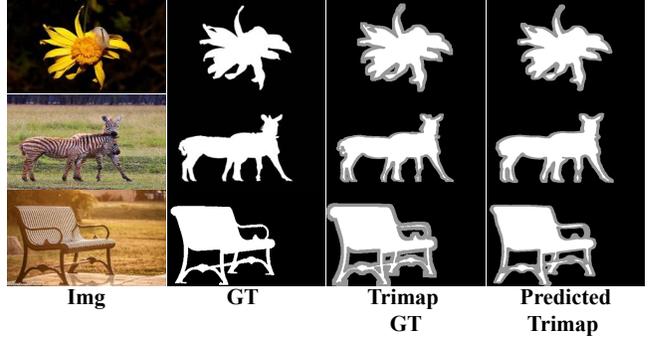


Figure 4. Examples of trimap. Column 3 shows the trimaps generated from GT. Column 4 shows the trimaps predicted by LRSCN.

accurate trimap. Finally, the output trimap  $T$  is 3-channel classification logits. The whole SGA module guarantees the alignment of trimap and saliency map.

### 3.3. Architecture of HRRN

Following the principle of disentanglement, HRRN aims at accurately refining the saliency value of pixels in the uncertain region to preserve a clear object boundary at high-resolution under the guidance of the trimap provided by LRSCN. The architecture of HRRN is shown in Fig.2. HRRN has a simple U-NET like architecture. For better prediction at high-resolution, we then do some non-trivial modifications. First, lower-level features contain rich spatial and detail information which play a crucial role in restoring a clear object boundary, so decoder combines encoder features before each upsampling block instead of after each upsampling block. Moreover, we use a two layers short cut block to align channels of encoder features for feature fusion. Second, to let network pay more attention to detail information, we directly feed the original input to the last convolutional layer through a short cut block to generate better results. Finally, learning from image generation tasks [3, 53], we use the spectral normalization [27] to each convolutional layer to add a constraint on Lipschitz constant of the network and stable the training.

### 3.4. Loss Function of LRSCN

To supervise LRSCN, we should generate trimap groundtruth  $T^{gt}$ , which can represent the definite salient, definite background and uncertain regions. As described, uncertain regions exist mainly at the boundaries of the objects. So we erase and dilate binary groundtruth maps at the object boundaries with a random pixel number (5,7,9,11,13) to generate the GT uncertain regions. The remaining foreground and background regions represent definite salient and background regions.  $T^{gt}$  is defined as:

$$T^{gt}(x, y) = \begin{cases} 2, & T^{gt}(x, y) \in \text{definite salient} \\ 0, & T^{gt}(x, y) \in \text{definite background} \\ 1, & T^{gt}(x, y) \in \text{uncertain region} \end{cases} \quad (1)$$

where  $(x, y)$  stands for each pixel location on the image. Some examples can be seen in Fig.4.

For trimap supervision, we use Softmax cross-entropy loss, which is defined as:

$$L_{trimap} = \frac{1}{N} \sum_i -\log\left(\frac{e^{T_i}}{\sum_j e^{T_j}}\right). \quad (2)$$

To guarantee the accuracy of trimap, we add extra saliency supervision  $L_{saliency}$  as the supplement of trimap supervision. Similar to BASNet [31], we use pixel-level, region-level and object-level supervision strategy on multi-levels to better keep the uniformity and wholeness of the salient objects. Specifically, binary cross-entropy (BCE) [7], SSIM [45] and F-measure loss [61] are denoted as pixel-level, region-level and object-level loss. Note that all parts of LRSCN are trained jointly, so the overall loss function is given as:

$$L_{LRSCN} = L_{saliency} + L_{trimap}. \quad (3)$$

We do not use uncertainty loss because the main goal of LRSCN is to capture sufficient semantics, not accurate boundary. More details about  $L_{saliency}$  can be found in section 5 of supplementary materials.

### 3.5. Loss Function of HRRN

We perform a  $L_1$  loss and novel uncertainty loss to restore the fine structures and boundaries of salient objects. For an input high-resolution image  $I$ , let  $G^H$  denote its groundtruth, and predicted saliency map is  $S^H$ .

We leverage the  $L_1$  loss to compare an absolute difference between predicted saliency map and groundtruth over the definite salient and background regions:

$$L_1 = \frac{1}{E} \sum_{i \in E} |S_i^H - G_i^H|, \quad (4)$$

where  $E$  indicates the number of pixels which are labeled as definite salient or background in the trimap,  $S_i^H$  and  $G_i^H$  denote the predicted and groundtruth value at position  $i$ .

We cannot directly compute  $L_1$  loss between predicted saliency map and groundtruth over the uncertain regions because widely used saliency training datasets have some problems in annotation quality [52]. We show these low quality annotations in section 4 of supplementary materials. It is almost impossible to directly obtain enough object boundary details from these defective datasets to train the high-resolution network. To address this problem, we design uncertainty loss, which empowers our HRRN to well address the high-resolution refinement task only using these defective low-resolution training datasets. It is worth noting that there are some previous works [54, 58] involving ‘‘uncertainty’’ in their titles, which seem relevant to our method.

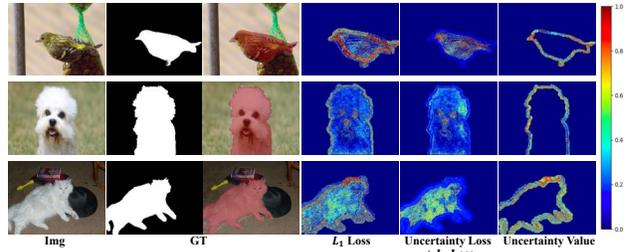


Figure 5. The impact of the losses. Best viewed by zooming in.

However, in [54], ‘‘uncertainty’’ means the human perceptual uncertainty modeled by CVAE. While in [58], ‘‘uncertainty’’ indicates the saliency prediction system uncertainty modeled by R-dropout. Obviously, their usages of uncertainty are different from ours.

Inspired by [19], Gaussian likelihood is used to model the uncertainty. Let  $x$  and  $f(x)$  be the input and output of HRRN, and Gaussian likelihood is defined as:

$$p(y|f(x)) = \mathcal{N}(f(x), \sigma^2), \quad (5)$$

where  $\sigma$  measures uncertainty of the estimation,  $y$  is the label of output. In maximum likelihood inference, we maximize the log likelihood of the model, which is written as:

$$\log p(y|f(x)) \propto -\frac{\|y - f(x)\|^2}{2\sigma^2} - \frac{1}{2} \log \sigma^2, \quad (6)$$

so the proposed uncertainty loss is defined as:

$$L_{uncertainty} = \frac{\|y - f(x)\|^2}{2\sigma^2} + \frac{1}{2} \log \sigma^2. \quad (7)$$

We only care about the pixels in uncertain regions, so  $L_{uncertainty}$  is written as:

$$L_{uncertainty} = \frac{1}{U} \sum_{i \in U} \frac{\|S_i^H - G_i^H\|^2}{2\sigma_i^2} + \frac{1}{2} \log \sigma_i^2, \quad (8)$$

where  $U$  is the total number of pixels in uncertain region,  $\sigma_i$  is the uncertainty of each pixel and is generated from HRRN. Different from directly learning from noisy data, uncertainty loss can allow the network to learn how to attenuate the effect from erroneous labels. Specifically, pixels for which the network learned to predict high uncertainty will have a smaller value of the first term of Eq.8, so have little effect on the loss. Meanwhile, large uncertainty increases the contribution of the second term of Eq.8, and in turn penalizes the model and lets the model make a better prediction that has low uncertainty. Note that all parts of HRRN are trained jointly, so the overall loss function is given as:

$$L_{HRRN} = L_{uncertainty} + L_1. \quad (9)$$

To show how proposed uncertainty loss makes the network attenuate the effect from erroneous labels during training, we visualize the impact of  $L_1$  loss and uncertainty loss

at the same training iteration in Fig.5. We show the impact of both these losses in the same image. The images shown in Fig.5 are these which have problems in annotation quality. Compared with column.4 and column.5, if we only use  $L_1$  loss, the weight of the loss in the uncertainty region will be large, which leads the network hard to converge. While uncertainty loss will make the weight of the loss in the uncertainty region be small and let the network ignore effects from noisy data as much as possible. Column.6 shows the uncertainty value of pixels in uncertain regions. It can be seen that pixels in uncertain regions usually have higher uncertainty value. In general, compared to  $L_1$  loss, uncertainty loss reduces the weight of the loss in the uncertainty region, thus mitigating the impact of noisy data on the network. But due to the uncertainty value, it will allow the network to learn how to predict a better prediction with a low certainty value, instead of ignoring the learning of the uncertainty region completely. These visual comparisons show how uncertainty loss makes network more robust to noisy data.

## 4. Experiment

### 4.1. Experimental Settings

**Implementation Details.** Following the works [31, 60, 28, 47], we train our proposed network on DUTS-TR. We use Pytorch<sup>1</sup> to implement our model. A GTX 1080Ti GPU is used for acceleration. VGG-16 [33] is used as the backbone network of LRSCN, and the whole network is trained end-to-end by stochastic gradient descent (SGD). For a more comprehensive demonstration, we also trained our network with ResNet-50 [15] backbone. Maximum learning rate is set to 0.001 for backbone and 0.01 for other parts. Warm-up and linear decay strategies are used to adjust the learning rate. Momentum and weight decay are set to 0.9 and 0.0005 respectively. Batchsize is set to 32 and maximum epoch is set to 100. Horizontal flip and multi-scale input images are utilized for data augmentation as done in [31, 60, 46]. During testing, the input of LRSCN is about  $352 \times 352$  resolution.

The learning rate of HRRN is initialized to 0.0005. Warmup and cosine decay are applied to the learning rate. The network HRRN is trained for 10000 iterations with a batch size of 20. During training, the resolution of input image and trimap is  $512 \times 512$ . During testing, we first resize the image and trimap to  $1024 \times 1024$ , then we split the image and trimap into four sub-images and sub-trimaps with  $512 \times 512$  resolution, as shown in Fig.2. Finally, we send each sub-image and sub-trimap together to HRRN to generate sub-prediction result, and use 4 sub-predictions stitched together to make one high-resolution saliency result.

<sup>1</sup><https://pytorch.org/>

**Evaluation Datasets.** Following work [52], we evaluate our method on two high-resolution saliency detection datasets, including HRSOD-TE and DAVIS-S, which contain 400 and 92 images. DAVIS-S dataset is collected from DAVIS [30]. Images in these two datasets are precisely annotated and have very high resolutions (i.e.,  $1920 \times 1080$ ). We also evaluate our method on three low-resolution datasets, including DUT-OMRON [51], DUTS-TE [37] and HKU-IS [22], which contain 5168, 5019 and 4447 images. Our results are available at <https://github.com/luckybird1994/HQSOD>.

**Evaluation Metrics.** Six metrics are used to evaluate the performance of our method. The first is Mean Absolute Error (MAE), which characterize the average 1-norm distance between ground truth maps and predictions. The second is F-measure ( $F_\beta$  and  $F_\beta^{max}$ ), a weighted mean of average precision and average recall, calculated by  $F_\beta = \frac{(1+\beta^2) \times Precision \times Recall}{\beta^2 \times Precision + Recall}$ . We set  $\beta^2$  to be 0.3 as suggested in [2]. The third is Structure Measure ( $S_m$ ), a metric to evaluate the spatial structure similarities of saliency maps based on both region-aware structural similarity  $S_r$  and object-aware structural similarity  $S_o$ , defined as  $S_\alpha = \alpha * S_r + (1 - \alpha) * S_o$ , where  $\alpha = 0.5$  [10]. In addition, precision-recall (PR) curve is used to show the whole performance. To further evaluate the boundary quality, Following [52] and [56], we use Boundary Displacement Error (BDE) [11] and  $B_\mu$  metrics. More details about BDE and  $B_\mu$  can be found in section 7 of supplementary materials. The last two metrics are only used in two high-resolution datasets, because their boundaries annotation is accurate and evaluating results are reliable.

### 4.2. Comparisons with the State-of-the-Arts

We compare our approach with 16 SOTA methods, including Amulet [57], R3Net [8], DGRL [41], DSS [16], BASNet [31], CPD [48], EGNet [60], PFPN [35], GCPA [5], F3N [46], MINet [28], ITSD [63], LDF [47], GateNet [62], CSF [12] and HRNet [52]. For a fair comparison, we use either the implementations with recommended parameter settings or the saliency maps provided by the authors. The evaluation toolbox used in this paper is same as F3N [46].

**Quantitative Evaluation.** From Table.1, when we train our network only using DUTS (Ours), our method can already improve the  $F_\beta^{max}$ ,  $F_\beta$ ,  $S_m$  and MAE achieved by the best-performing existing algorithms, especially two high-resolution test datasets. It is worth noting that for boundary accuracy, our method is far better than other methods on two high-resolution. These results demonstrate the efficiency of the proposed disentangled SOD framework in both identifying the salient regions and estimating the accurate objects boundaries. Other than numerical results, we also show the PR curves on two high-resolution datasets

Table 1. Quantitative comparison with SOTA on two high-resolution and three low-resolution datasets. The best three results are in red , green and blue fonts. "†" means the results are post-processed by dense conditional random field(CRF) [21]. "\*" means using ResNeXt-101 [50] backbone. "★" means using ResNet-101 backbone. "‡" means using Res2Net50 [13] backbone. MK: MSRA10K [6], DUTS: DUTS-TR [37], MB: MSRA-B [25], HR: HRSOD-Training [52], HR-L: HRSOD-Training resized in low-resolution. Smaller MAE, BDE and  $B_{\mu}$ , larger  $F_{\beta}^{max}$ ,  $F_{\beta}$  and  $S_m$  correspond to better performance.

Models	Training datasets	HRSOD-TE						DAVIS-S						DUT-OMRON				DUTS-TE				HKU-IS			
		$F_{\beta}^{max}$	$F_{\beta}$	$S_m$	MAE	BDE	$B_{\mu}$	$F_{\beta}^{max}$	$F_{\beta}$	$S_m$	MAE	BDE	$B_{\mu}$	$F_{\beta}^{max}$	$F_{\beta}$	$S_m$	MAE	$F_{\beta}^{max}$	$F_{\beta}$	$S_m$	MAE	$F_{\beta}^{max}$	$F_{\beta}$	$S_m$	MAE
VGG-16 backbone																									
Amulet(ICCV2017)	MK	0.799	0.717	0.829	0.075	139.889	0.947	0.802	0.755	0.848	0.042	64.827	0.856	0.743	0.647	0.781	0.098	0.778	0.678	0.804	0.085	0.897	0.841	0.886	0.051
DGRL(CVPR2018)	DUTS	0.821	0.789	0.847	0.055	95.034	0.889	0.803	0.772	0.859	0.038	50.323	0.826	0.774	0.709	0.810	0.063	0.828	0.794	0.842	0.050	0.910	0.881	0.896	0.037
DSS(TPAMI2019)	MB	0.826	0.756	0.840	0.060	145.403	0.952	0.830	0.728	0.865	0.041	94.069	0.890	0.781	0.740	0.790	0.062	0.825	0.808	0.820	0.057	0.916	0.902	0.878	0.040
CPD(CVPR2019)	DUTS	0.876	0.829	0.887	0.039	72.686	0.824	0.878	0.822	0.903	0.025	36.649	0.703	0.794	0.745	0.818	0.057	0.864	0.813	0.867	0.043	0.924	0.896	0.904	0.033
EGNet(CVPR2019)	DUTS	0.883	0.814	0.888	0.044	73.500	0.896	0.886	0.794	0.897	0.030	37.369	0.799	0.803	0.744	0.813	0.057	0.877	0.800	0.866	0.044	0.927	0.893	0.910	0.035
MINet(CVPR2020)	DUTS	0.902	0.851	0.903	0.032	76.291	0.849	0.915	0.864	0.926	0.019	32.304	0.742	0.794	0.741	0.822	0.057	0.877	0.823	0.875	0.039	0.930	0.904	0.912	0.031
ITSD(CVPR2020)	DUTS	0.824	0.715	0.834	0.071	139.943	0.924	0.806	0.687	0.843	0.055	92.864	0.861	0.802	0.745	0.828	0.063	0.876	0.798	0.877	0.042	0.927	0.890	0.906	0.035
GateNet(CVPR2020)	DUTS	0.905	0.825	0.906	0.035	79.468	0.886	0.914	0.825	0.923	0.023	44.827	0.778	0.794	0.723	0.821	0.061	0.870	0.783	0.870	0.045	0.929	0.889	0.910	0.036
HRNet(ECCV2019)	DUTS+HR	0.918	0.888	0.897	0.030	88.017	0.888	0.899	0.888	0.926	0.026	44.359	0.801	0.743	0.690	0.762	0.065	0.835	0.788	0.824	0.050	0.910	0.886	0.877	0.042
Ours	DUTS	0.918	0.902	0.912	0.027	48.468	0.711	0.933	0.919	0.933	0.015	15.676	0.536	0.804	0.769	0.829	0.053	0.882	0.855	0.879	0.036	0.935	0.918	0.913	0.029
Ours-DH	DUTS+HR-L	0.921	0.907	0.917	0.024	45.462	0.706	0.938	0.926	0.936	0.014	14.412	0.531	0.795	0.764	0.820	0.052	0.894	0.865	0.879	0.035	0.933	0.914	0.904	0.031
ResNeXt-101/ResNeXt-101/Res2Net50 backbone																									
R3Net*(ICAAI2018)	MK	0.798	0.744	0.812	0.081	108.910	0.931	0.806	0.753	0.835	0.041	47.373	0.868	0.690	0.819	0.073	0.778	0.716	0.837	0.067	0.915	0.853	0.894	0.047	
BASNet(CVPR2019)	DUTS	0.878	0.831	0.890	0.038	67.643	0.823	0.857	0.806	0.881	0.039	46.283	0.705	0.805	0.766	0.838	0.056	0.859	0.791	0.866	0.048	0.928	0.895	0.909	0.032
PPFN(AAAI2020)	DUTS	0.889	0.825	0.897	0.042	65.048	0.896	0.886	0.822	0.912	0.025	30.488	0.748	0.818	0.748	0.841	0.057	0.885	0.805	0.887	0.041	0.937	0.896	0.919	0.033
GCFA(AAAI2020)	DUTS	0.889	0.827	0.894	0.039	70.320	0.873	0.912	0.833	0.924	0.021	24.132	0.759	0.812	0.748	0.838	0.056	0.888	0.817	0.891	0.038	0.938	0.898	0.920	0.031
F3N(AAAI2020)	DUTS	0.900	0.853	0.897	0.035	65.901	0.817	0.915	0.845	0.913	0.020	45.106	0.719	0.813	0.766	0.838	0.053	0.891	0.840	0.888	0.035	0.937	0.910	0.917	0.028
LDf(CVPR2020)	DUTS	0.905	0.866	0.905	0.032	58.655	0.812	0.911	0.864	0.922	0.019	35.496	0.713	0.817	0.773	0.839	0.052	0.894	0.855	0.890	0.034	0.939	0.914	0.919	0.028
ITSD(CVPR2020)	DUTS	0.894	0.832	0.900	0.038	71.293	0.922	0.899	0.822	0.912	0.025	30.488	0.848	0.815	0.750	0.838	0.055	0.894	0.823	0.890	0.038	0.935	0.902	0.921	0.030
CSF(ECCV2020)	DUTS	0.915	0.902	0.919	0.024	47.804	0.750	0.935	0.923	0.937	0.013	14.396	0.576	0.818	0.785	0.842	0.051	0.895	0.870	0.892	0.033	0.943	0.928	0.923	0.025
Ours	DUTS	0.922	0.909	0.922	0.022	46.495	0.746	0.938	0.926	0.939	0.012	14.266	0.571	0.820	0.791	0.843	0.048	0.900	0.876	0.892	0.031	0.944	0.929	0.922	0.026

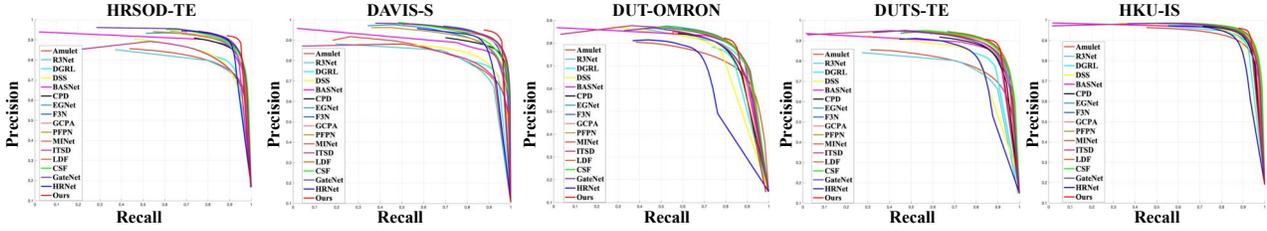


Figure 6. Comparison of PR curves across two high-resolution and three low-resolution datasets.

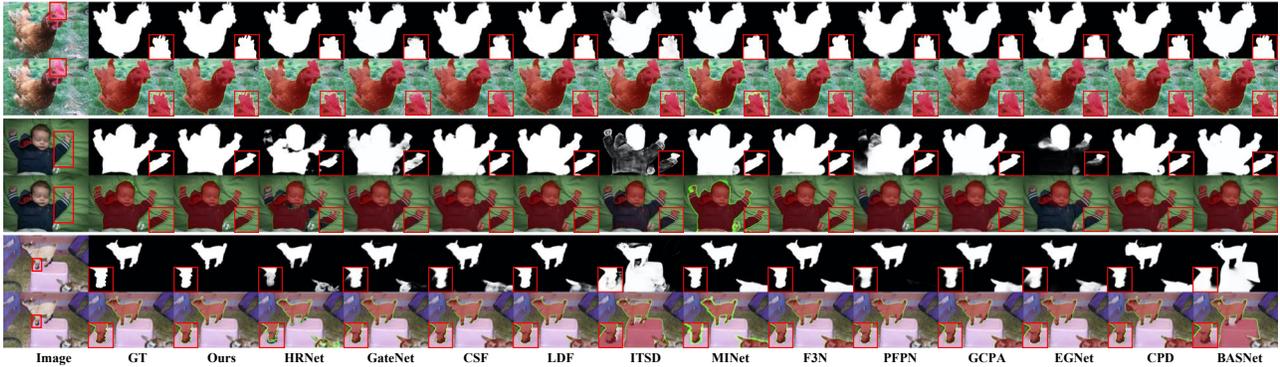


Figure 7. Visual comparison between our method and other SOTA methods. Each sample occupies two rows. Best viewed by zooming in. It can be clearly observed that our method achieves impressive performance in all these cases.

and three low-resolution datasets in Fig.6. As can be seen, the PR curves by our method (red ones) are especially outstanding compared to all other previous methods. Besides, shorter PR curves imply that our saliency maps are usually more assertive with sharper boundaries than the results of other methods. An interesting observation is that when we add HRSOD-training datasets (resized to low-resolution like  $352 \times 352$ ) in LRSCN, the performance in two high-resolution datasets HRSOD-TE and DAVIS-S can be further improved. However, this practice seems to be of little help in improving the performance in other three low-resolution datasets. A similar phenomenon can also be found in the performance of HRNet [52]. We think there could be some

image selection or data annotation biases between the high-resolution datasets and the low-resolution datasets, which cause this phenomenon.

**Qualitative Evaluation.** To exhibit the superiority of the proposed approach, Fig.7 shows representative examples of saliency maps generated by our approach and other state-of-the-art algorithms. As can be seen, with the help of LRSCN, our method can not only keep the wholeness of the salient object (row 3), but also accurately locate salient objects and suppress non-salient regions (row 5), compared to other methods. HRRN can help the model to restore accurate and complete boundaries of salient objects, which are more consistent with the GT boundaries. It can be clearly

Table 2. Ablation Studies of LRSCN.

Configurations	HRSOD-TE						DAVIS-S					
	$F_{max}$	$F_{\beta}$	$S_m$	MAE	BDE	$B_{\mu}$	$F_{max}$	$F_{\beta}$	$S_m$	MAE	BDE	$B_{\mu}$
Baseline+HRRN	0.900	0.880	0.896	0.034	65.732	0.842	0.899	0.887	0.919	0.021	31.201	0.678
Baseline+ME+HRRN	0.910	0.894	0.902	0.031	60.040	0.801	0.920	0.904	0.924	0.019	24.022	0.612
Baseline+CF+HRRN	0.909	0.892	0.900	0.030	59.028	0.804	0.918	0.905	0.923	0.018	23.988	0.605
Baseline+MECF+HRRN	0.913	0.898	0.909	0.029	54.377	0.766	0.928	0.915	0.929	0.016	20.010	0.578
Baseline+MECF+SGA+HRRN	<b>0.918</b>	<b>0.902</b>	<b>0.912</b>	<b>0.027</b>	<b>48.468</b>	<b>0.711</b>	<b>0.933</b>	<b>0.919</b>	<b>0.933</b>	<b>0.015</b>	<b>15.676</b>	<b>0.536</b>

Table 3. Ablation Studies of HRRN.

Ablation	Configurations	HRSOD-TE						DAVIS-S					
		$F_{max}$	$F_{\beta}$	$S_m$	MAE	BDE	$B_{\mu}$	$F_{max}$	$F_{\beta}$	$S_m$	MAE	BDE	$B_{\mu}$
Loss	Ours( $L_1 + L_{uncertainty}$ )	0.918	0.902	0.912	0.027	48.468	0.711	0.933	0.919	0.933	0.015	15.676	0.536
	Ours( $L_1$ )	0.907	0.896	0.908	0.029	53.891	0.780	0.921	0.909	0.927	0.016	18.014	0.622
	Ours-DH( $L_1 + L_{uncertainty}$ )	0.918	0.901	0.911	0.027	48.878	0.712	0.933	0.920	0.934	0.015	17.670	0.540
	Ours(LRSCN)	0.898	0.885	0.899	0.034	64.805	0.822	0.909	0.898	0.920	0.022	28.798	0.684
Architecture	Ours(LRSCN+HRRN)	0.918	0.902	0.912	0.027	48.468	0.711	0.933	0.919	0.933	0.015	15.676	0.536
	Ours(LRSCN+CRF)	0.905	0.896	0.897	0.029	60.521	0.797	0.920	0.907	0.918	0.018	24.65	0.665
	EGNet	0.883	0.814	0.888	0.044	73.500	0.896	0.886	0.794	0.897	0.030	37.369	0.799
	EGNet(+HRRN)	0.900	0.862	0.889	0.039	72.982	0.753	0.904	0.858	0.898	0.024	34.860	0.603
	EGNet(+CRF)	0.895	0.858	0.882	0.039	73.348	0.796	0.902	0.846	0.892	0.025	34.721	0.679
	CPD	0.876	0.829	0.887	0.039	72.886	0.824	0.898	0.822	0.903	0.025	38.669	0.703
	CPD(+HRRN)	0.891	0.855	0.888	0.036	72.796	0.734	0.893	0.859	0.907	0.022	35.294	0.564
	CPD(+CRF)	0.885	0.851	0.884	0.037	76.440	0.772	0.884	0.852	0.903	0.023	36.915	0.633
	BASNet(En-De)	0.873	0.827	0.888	0.039	70.944	0.824	0.852	0.802	0.880	0.039	48.309	0.703
	BASNet(En-De+HRRN)	0.895	0.858	0.892	0.036	67.191	0.719	0.875	0.832	0.882	0.036	45.922	0.566
	BASNet(En-De+RRM)	0.878	0.831	0.890	0.038	67.643	0.823	0.857	0.806	0.881	0.039	46.283	0.705
	BASNet(En-De+CRF)	0.888	0.850	0.886	0.036	77.235	0.758	0.867	0.825	0.878	0.037	46.423	0.629

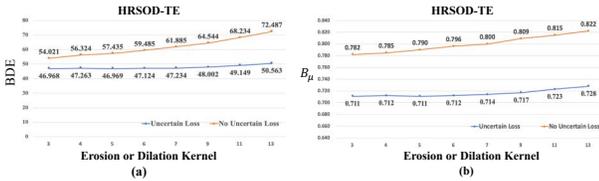


Figure 8. Training the network with noisy data.

observed that our method achieves impressive performance in all these cases, which indicates the effectiveness of disentangled framework and uncertainty loss. More comparison experiments can be found in section 2 of supplementary materials.

### 4.3. Ablation Studies

To validate the effectiveness of the proposed components of our method, we conduct a series of experiments on two high-resolution datasets with different settings under VGG-16 backbone. Specifically, we first verify the effectiveness of MECF and SGA in LRSCN. Then we validate the effectiveness of uncertainty loss in HRRN and the superiority of the proposed disentangled architecture.

**Ablation Studies of LRSCN.** To prove the effectiveness of MECF and SGA module, we report the quantitative comparison results of LRSCN with different architectures in Table.2. Baseline denotes that we conduct a experiment over on  $L_{LRSCN}$  with a pure U-Net architecture. We can see that only using ME or CF can already heavily improve the performance. A better performance has been achieved through the combination of these two architectures. Finally, performance can be further improved by SGA module, especially BDE and  $B_{\mu}$ , which means that SGA can help generate accurate trimap. While  $L_{saliency}$  is not our core innovation, more ablation studies about  $L_{saliency}$  can be found in section 5 of supplementary materials.

**Ablation Studies of HRRN.** In HRRN, uncertainty loss play a key role to estimate the accurate objects boundaries, so we first investigate the effectiveness of our proposed uncertainty loss. From Table.3, we can see that without uncertainty loss (Ours( $L_1$ )), the performance decreased a lot. Besides, when we add high-resolution HRSOD-Training

datasets in HRRN (Ours-DH( $L_1 + L_{uncertainty}$ )), the performance has no obvious improvement, which demonstrates that our network is not reliant on accurately annotated high-resolution images during training. To further demonstrate the effectiveness of our disentangled framework, we compare our HRRN with CRF [21], a widely used post-processing for saliency detection. Results in Table.3 show that our proposed method (Ours (LRSCN + HRRN)) outperforms CRF (Ours (LRSCN + CRF)) by a large margin. The same phenomenon can be found in the refinement of EGNet, CPD and BASNet (Their trimaps are generated with the corresponding saliency maps following Eq.1). Moreover, compared to RRM module proposed in BASNet, our HRRN can better improve performance. This ablation study demonstrates the superiority of HRRN within our novel disentangled framework. More analyses of the proposed disentangled framework can be found in section 3 of supplementary materials.

**Analysis of Uncertainty Loss** To further demonstrate that uncertainty loss can make network more robust to noisy data, we erase and dilate binary groundtruth maps of DUT-TR at the object boundaries with a random pixel number (3,4,5,6,7,9,11,13) to generate noisy training data. Then we train the network on these noisy data, BDE and  $B_{\mu}$  results on HRSOD-TE are reported in Fig.8(a) and Fig.8(b). When the erosion or dilation kernel ranges from 3 to 7, the network trained with uncertainty loss has a rather stable performance. With the increase in erosion or dilation kernels, even though the performance is dropped, training with uncertainty loss still yields a better model than training without uncertainty loss. This experiment further validates the effectiveness of uncertainty loss.

## 5. Conclusions

In this paper, we argue that there are two difficult and inherently different problems in high-resolution SOD. From this perspective, we propose a novel deep learning framework to disentangle the high-resolution SOD into two tasks: LRSCN and HRRN. LRSCN can identify the definite salient, background and uncertain regions at low-resolution with sufficient semantics. While HRRN can accurately refining the saliency value of pixels in the uncertain region to preserve a clear object boundary at high-resolution with limited GPU memory. We also make the earliest efforts to introduce the uncertainty into SOD network training, which empower HRRN to learn rich details without using any high-resolution training datasets. Extensive evaluations on high-resolution datasets and popular benchmark datasets not only verify the superiority of our method but also demonstrate the importance of disentanglement for SOD. We believe our novel disentanglement view in this work can contribute to other high-resolution computer vision tasks in the future.

## References

- [1] Ali Borji, Ming-Ming Cheng, Qibin Hou, Huaizu Jiang, and Jia Li. Salient object detection: A survey. *Computational Visual Media*, 5(2):117–150, 2019.
- [2] Ali Borji, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. Salient object detection: A benchmark. *IEEE TIP*, 24(12):5706–5722, 2015.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2019.
- [4] Shuhan Chen, Xiuli Tan, Ben Wang, and Xuelong Hu. Reverse attention for salient object detection. In *ECCV(2018)*, volume 11213 of *Lecture Notes in Computer Science*, pages 236–252. Springer, 2018.
- [5] Zuyao Chen, Qianqian Xu, Runmin Cong, and Qingming Huang. Global context-aware progressive aggregation network for salient object detection. In *AAAI*, pages 10599–10606. AAAI Press, 2020.
- [6] Ming-Ming Cheng, Niloy J. Mitra, Xiaolei Huang, Philip H. S. Torr, and Shi-Min Hu. Global contrast based salient region detection. *IEEE TPAMI*, 37(3):569–582, 2015.
- [7] Pieter-Tjerk de Boer, Dirk P. Kroese, Shie Mannor, and Reuven Y. Rubinfeld. A tutorial on the cross-entropy method. *Ann. Oper. Res.*, 134(1):19–67, 2005.
- [8] Zijun Deng, Xiaowei Hu, Lei Zhu, Xuemiao Xu, Jing Qin, Guoqiang Han, and Pheng-Ann Heng. R<sup>3</sup>net: Recurrent residual refinement network for saliency detection. In *IJCAI*, pages 684–690. ijcai.org, 2018.
- [9] Deng-Ping Fan, Ming-Ming Cheng, Jiangjiang Liu, Shanghua Gao, Qibin Hou, and Ali Borji. Salient objects in clutter: Bringing salient object detection to the foreground. In *ECCV (15)*, volume 11219 of *Lecture Notes in Computer Science*, pages 196–212. Springer, 2018.
- [10] Deng-Ping Fan, Ming-Ming Cheng, Yun Liu, Tao Li, and Ali Borji. Structure-measure: A new way to evaluate foreground maps. In *ICCV*, pages 4558–4567(2017). IEEE, 2017.
- [11] Jordi Freixenet, Xavier Muñoz, David Raba, Joan Martí, and Xavier Cufí. Yet another survey on image segmentation: Region and boundary information integration. In *ECCV (3)*, volume 2352 of *Lecture Notes in Computer Science*, pages 408–422. Springer, 2002.
- [12] Shanghua Gao, Yong-Qiang Tan, Ming-Ming Cheng, Chengze Lu, Yunpeng Chen, and Shuicheng Yan. Highly efficient salient object detection with 100k parameters. *CoRR*, abs/2003.05643, 2020.
- [13] Shang-Hua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip Torr. Res2net: A new multi-scale backbone architecture. *IEEE TPAMI*, 2020.
- [14] Junwei Han, Dingwen Zhang, Gong Cheng, Nian Liu, and Dong Xu. Advanced deep-learning techniques for salient and category-specific object detection: A survey. *IEEE Signal Process. Mag.*, 35(1):84–100, 2018.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778. IEEE, 2016.
- [16] Qibin Hou, Ming-Ming Cheng, Xiaowei Hu, Ali Borji, Zhuowen Tu, and Philip H. S. Torr. Deeply supervised salient object detection with short connections. *IEEE TPAMI*, 41(4):815–828, 2019.
- [17] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE TPAMI*, 20(11):1254–1259, 1998.
- [18] Zhong Ji, Haoran Wang, Jungong Han, and Yanwei Pang. Saliency-guided attention network for image-sentence matching. In *ICCV*, pages 5753–5762. IEEE, 2019.
- [19] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*, pages 5574–5584, 2017.
- [20] Dominik A. Klein and Simone Frintrop. Center-surround divergence of feature statistics for salient object detection. In *ICCV*, pages 2214–2219(2011). IEEE, 2011.
- [21] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, pages 109–117, 2011.
- [22] Guanbin Li and Yizhou Yu. Visual saliency based on multi-scale deep features. In *CVPR*, pages 5455–5463. IEEE Computer Society, 2015.
- [23] Jiawei Liu, Zheng-Jun Zha, Xierong Zhu, and Na Jiang. Co-saliency spatio-temporal interaction network for person re-identification in videos. In *IJCAI*, pages 1012–1018. ijcai.org, 2020.
- [24] Nian Liu and Junwei Han. Dhsnet: Deep hierarchical saliency network for salient object detection. In *CVPR*, pages 678–686. IEEE, 2016.
- [25] Tie Liu, Zejian Yuan, Jian Sun, Jingdong Wang, Nanning Zheng, Xiaou Tang, and Heung-Yeung Shum. Learning to detect a salient object. *IEEE TPAMI*, 33(2):353–367, 2011.
- [26] Yi Liu, Qiang Zhang, Dingwen Zhang, and Jungong Han. Employing deep part-object relationships for salient object detection. In *ICCV*, pages 1232–1241. IEEE, 2019.
- [27] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.
- [28] Youwei Pang, Xiaoqi Zhao, Lihe Zhang, and Huchuan Lu. Multi-scale interactive network for salient object detection. In *CVPR*, pages 9410–9419. IEEE, 2020.
- [29] Chao Peng, Xiangyu Zhang, Gang Yu, Guiming Luo, and Jian Sun. Large kernel matters - improve semantic segmentation by global convolutional network. In *CVPR*, 2017.
- [30] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus H. Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, pages 724–732. IEEE, 2016.
- [31] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jägersand. Basnet: Boundary-aware salient object detection. In *CVPR*, pages 7479–7489. Computer Vision Foundation / IEEE, 2019.
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI (3)*, pages 234–241. Springer, 2015.
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.

- [34] Shaoyue Song, Hongkai Yu, Zhenjiang Miao, Jianwu Fang, Kang Zheng, Cong Ma, and Song Wang. Multi-spectral salient object detection by adversarial domain adaptation. In *AAAI*, pages 12023–12030. AAAI Press, 2020.
- [35] Bo Wang, Quan Chen, Min Zhou, Zhiqiang Zhang, Xiaogang Jin, and Kun Gai. Progressive feature polishing network for salient object detection. In *AAAI*, pages 12128–12135. AAAI Press, 2020.
- [36] Jingdong Wang, Huaizu Jiang, Zejian Yuan, Ming-Ming Cheng, Xiaowei Hu, and Nanning Zheng. Salient object detection: A discriminative regional feature integration approach. *IJCV*, 123(2):251–268, 2017.
- [37] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *CVPR*, pages 3796–3805(2017). IEEE, 2017.
- [38] T. Wang, A. Borji, L. Zhang, P. Zhang, and H. Lu. A stage-wise refinement model for detecting salient objects in images. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4039–4048, 2017.
- [39] Tiantian Wang, Yongri Piao, Huchuan Lu, Xiao Li, and Lihe Zhang. Deep learning for light field saliency detection. In *ICCV*, pages 8837–8847. IEEE, 2019.
- [40] Tiantian Wang, Lihe Zhang, Huchuan Lu, Chong Sun, and Jinqing Qi. Kernelized subspace ranking for saliency detection. In *ECCV(2016)*, volume 9912 of *Lecture Notes in Computer Science*, pages 450–466. Springer, 2016.
- [41] Tiantian Wang, Lihe Zhang, Shuo Wang, Huchuan Lu, Gang Yang, Xiang Ruan, and Ali Borji. Detect globally, refine locally: A novel approach to saliency detection. In *CVPR*, pages 3127–3135. IEEE, 2018.
- [42] Wenguan Wang, Qiuxia Lai, Huazhu Fu, Jianbing Shen, and Haibin Ling. Salient object detection in the deep learning era: An in-depth survey. *CoRR*, abs/1904.09146, 2019.
- [43] Wenguan Wang, Jianbing Shen, Ruigang Yang, and Fatih Porikli. Saliency-aware video object segmentation. *IEEE TPAMI*, 40(1):20–33, 2018.
- [44] Wenguan Wang, Shuyang Zhao, Jianbing Shen, Steven C. H. Hoi, and Ali Borji. Salient object detection with pyramid attention and salient edges. In *CVPR*, pages 1448–1457(2019), 2019.
- [45] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004.
- [46] Jun Wei, Shuhui Wang, and Qingming Huang. F3net: Fusion, feedback and focus for salient object detection. *CoRR*, abs/1911.11445, 2019.
- [47] Jun Wei, Shuhui Wang, Zhe Wu, Chi Su, Qingming Huang, and Qi Tian. Label decoupling framework for salient object detection. In *CVPR*, pages 13022–13031. IEEE, 2020.
- [48] Zhe Wu, Li Su, and Qingming Huang. Cascaded partial decoder for fast and accurate salient object detection. In *CVPR*, pages 3907–3916. Computer Vision Foundation / IEEE, 2019.
- [49] Zhe Wu, Li Su, and Qingming Huang. Stacked cross refinement network for edge-aware salient object detection. In *ICCV*, pages 7263–7272. IEEE, 2019.
- [50] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, pages 5987–5995. IEEE, 2017.
- [51] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via graph-based manifold ranking. In *CVPR*, pages 3166–3173. IEEE Computer Society, 2013.
- [52] Yi Zeng, Pingping Zhang, Zhe L. Lin, Jianming Zhang, and Huchuan Lu. Towards high-resolution salient object detection. In *ICCV*, pages 7233–7242. IEEE, 2019.
- [53] Han Zhang, Ian J. Goodfellow, Dimitris N. Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 7354–7363. PMLR, 2019.
- [54] Jing Zhang, Deng-Ping Fan, Yuchao Dai, Saeed Anwar, Fatemeh Sadat Saleh, Tong Zhang, and Nick Barnes. Uc-net: Uncertainty inspired RGB-D saliency detection via conditional variational autoencoders. In *CVPR*, pages 8579–8588. IEEE, 2020.
- [55] Jing Zhang, Jianwen Xie, and Nick Barnes. Learning noise-aware encoder-decoder from noisy labels by alternating back-propagation for saliency detection. *CoRR*, abs/2007.12211, 2020.
- [56] Jing Zhang, Xin Yu, Aixuan Li, Peipei Song, Bowen Liu, and Yuchao Dai. Weakly-supervised salient object detection via scribble annotations. In *CVPR*, pages 12543–12552. IEEE, 2020.
- [57] Pingping Zhang, Dong Wang, Huchuan Lu, Hongyu Wang, and Xiang Ruan. Amulet: Aggregating multi-level convolutional features for salient object detection. In *ICCV*, pages 202–211, 2017.
- [58] Pingping Zhang, Dong Wang, Huchuan Lu, Hongyu Wang, and Baocai Yin. Learning uncertain convolutional features for accurate saliency detection. In *ICCV*, pages 212–221. IEEE, 2017.
- [59] Xiaoning Zhang, Tiantian Wang, Jinqing Qi, Huchuan Lu, and Gang Wang. Progressive attention guided recurrent network for salient object detection. In *CVPR*, pages 714–722. IEEE Computer Society, 2018.
- [60] Jiaying Zhao, Jiangjiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. Egnnet: Edge guidance network for salient object detection. In *ICCV*, pages 8778–8787. IEEE, 2019.
- [61] Kai Zhao, Shanghua Gao, Wenguan Wang, and Ming-Ming Cheng. Optimizing the f-measure for threshold-free salient object detection. In *ICCV*, pages 8848–8856. IEEE, 2019.
- [62] Xiaoqi Zhao, Youwei Pang, Lihe Zhang, Huchuan Lu, and Lei Zhang. Suppress and balance: A simple gated network for salient object detection. *CoRR*, abs/2007.08074, 2020.
- [63] Huajun Zhou, Xiaohua Xie, Jian-Huang Lai, Zixuan Chen, and Lingxiao Yang. Interactive two-stream decoder for accurate and fast saliency detection. In *CVPR*, pages 9138–9147. IEEE, 2020.
- [64] Yanzhao Zhou, Xin Wang, Jianbin Jiao, Trevor Darrell, and Fisher Yu. Learning saliency propagation for semi-supervised instance segmentation. In *CVPR*, pages 10304–10313. IEEE, 2020.

- [65] Ziqi Zhou, Zheng Wang, Huchuan Lu, Song Wang, and Meijun Sun. Multi-type self-attention guided degraded saliency detection. In *AAAI*, pages 13082–13089. AAAI Press, 2020.