




# Reverse Attention for Salient Object Detection

Shuhan Chen<sup>(✉)</sup>, Xiuli Tan, Ben Wang, and Xuelong Hu

School of Information Engineering, Yangzhou University, Yangzhou, China  
c.shuhan@gmail.com, t.xiuli0214@gmail.com, wangben9503@163.com,  
xlhu@yzu.edu.cn

**Abstract.** Benefit from the quick development of deep learning techniques, salient object detection has achieved remarkable progresses recently. However, there still exists following two major challenges that hinder its application in embedded devices, low resolution output and heavy model weight. To this end, this paper presents an accurate yet compact deep network for efficient salient object detection. More specifically, given a coarse saliency prediction in the deepest layer, we first employ residual learning to learn side-output residual features for saliency refinement, which can be achieved with very limited convolutional parameters while keep accuracy. Secondly, we further propose reverse attention to guide such side-output residual learning in a top-down manner. By erasing the current predicted salient regions from side-output features, the network can eventually explore the missing object parts and details which results in high resolution and accuracy. Experiments on six benchmark datasets demonstrate that the proposed approach compares favorably against state-of-the-art methods, and with advantages in terms of simplicity, efficiency (**45 FPS**) and model size (**81 MB**).


**Keywords:** Salient object detection · Reverse attention  
Side-output residual learning

## 1 Introduction

Salient object detection, also known as saliency detection, aims to localize and segment the most conspicuous and eye-attracting objects or regions in an image. It is usually served as a pre-processing step to facilitate various subsequent high-level vision tasks, such as image segmentation [1], image captioning [2], and so on. Recently, with the quick development of deep convolutional neural networks (CNNs), salient object detection has achieved significant improvements over conventional hand-crafted feature based approaches. The emergence of fully convolutional neural networks (FCNs) [3] further pushed it to a new state-of-the-art due to its efficiency and end-to-end training. Such architecture also benefits other applications, *e.g.*, semantic segmentation [4], edge detection [5].

Albeit profound progresses have been made, there still exists two major challenges that hinder its applications in real-world, *e.g.*, embedded devices. One is


the low resolution of the saliency maps produced by FCNs based saliency models. Due to the repeated stride and pooling operations in CNN architectures, it is inevitable to lose resolution and difficult to refine, making it infeasible to locate salient objects accurately, especially for the object boundaries and small objects. The other is the heavy weight and large redundancy of the existing deep saliency models. As can be seen in Fig. 1, all the listed deep models are larger than 100 MB, which is too heavy for a pre-processing step to apply in subsequent high-level tasks, and also not memory efficient for embedded devices.



**Fig. 1.** Maximum F-Measure of recent deep CNN-based saliency detection models on ECSSD, including DS [6], ELD [7], DCL<sup>+</sup> [8], DHS [8], RFCN [9], NLDF [10], DSS<sup>+</sup> [11], MSRNet [12], Amulet [13], UCF [14], and ours (red circle). As can be seen that the proposed model is the only one less than 100 MB while achieves comparable performance with state-of-the-art methods. (Color figure online)

Diverse solutions have been explored to improve the resolution of the FCNs based prediction. Early works [8, 15, 16] usually combined it with an extra region or superpixel based stream to fuse their respective advantages at the expense of high time cost. Then, some simple yet effective structures are constructed to combine the complementary cues of shallow and deep CNN features, which capture low-level spatial details and high-level semantic information respectively, such as skip connections [12], short connections [11], dense connections [17], adaptive aggregation [13]. Such multi-level feature fusion schemes also play an important role in semantic segmentation [18, 19], edge detection [20], skeleton detection [21, 22]. Nevertheless, the existing archaic fusions are still incompetent for saliency detection under complex real-world scenarios, especially when dealing with multiple salient objects with diverse scales. In addition, some time consuming post-processing skills are also applied for refinement, *e.g.*, superpixel-based filter [23], fully connected conditional random field (CRF) [8, 11, 24]. However, to the best of our knowledge, there are no saliency detection networks explored considering both lightweight model and high accuracy.

To this end, we present an accurate yet compact deep salient object detection network which achieved comparable performance with state-of-the-art methods,



**Fig. 2.** Visual comparison of saliency maps produced by DSS [11] (top row), our method without (middle row) and with reverse attention (bottom row) in different side-outputs, respectively. As can be seen clearly that the resolutions of the saliency maps are improved gradually from deep to shallow side-outputs, and our reverse attention based side-output residual learning performs much better than short connections [11].

thus enables for real-time applications. In generally, more convolutional channels with large kernel size leads to better performance in salient object detection due to the large receptive field and model capacity to capture more semantic information, *e.g.*, there are 512 channels with kernel size  $7 \times 7$  in the last side-output of DSS [11]. In a different way, we introduce residual learning [25] into the architecture of HED [5], and regard salient object detection as a super-resolution reconstruction problem [26]. Given the low resolution prediction of FCNs, side-output residual features are learned to refine it step by step. Note that it can be achieved only using convolution with 64 channels and kernel size  $3 \times 3$  in each side-output, whose parameters are significant fewer than DSS.

Similar residual learning was also utilized in skeleton detection [21] and image super-resolution [27]. However, the performance is not satisfactory enough if we directly apply it for salient object detection due to its challenging. Since most of the existing deep saliency models are fine-tuned from image classification network, the fine-tuned network will unconsciously focus on the regions with high response values during residual learning as can be seen in Fig. 5, thus struggling to capture the residual details, *e.g.*, object boundaries and other undetected object parts. To solve it, we propose reverse attention to guide side-output residual learning in a top-down manner. Specifically, prediction of deep layer is upsampled then reversed to weight its neighbor shallow side-output feature, which quickly guides the network to focus on the undetected regions for residual capture, thus leads to better performance as seen in Fig. 2.

In summary, the contributions of this paper can be concluded as: (1) We introduce residual learning into the architecture of HED for salient object detection. With the help of the learned side-output residual features, the resolution of the saliency map can be improved gradually with much fewer parameters compared to the existing deep saliency networks. (2) We further propose reverse attention

to guide side-output residual learning. By erasing the current prediction, the network can discover the missing object parts and residual details effectively and quickly, which leads to significant performance improvement. (3) Benefit from the above two components, our approach consistently achieves comparable performance with state-of-the-art methods, and with advantages in terms of simplicity, efficiency (45 FPS) and model size (81 MB).

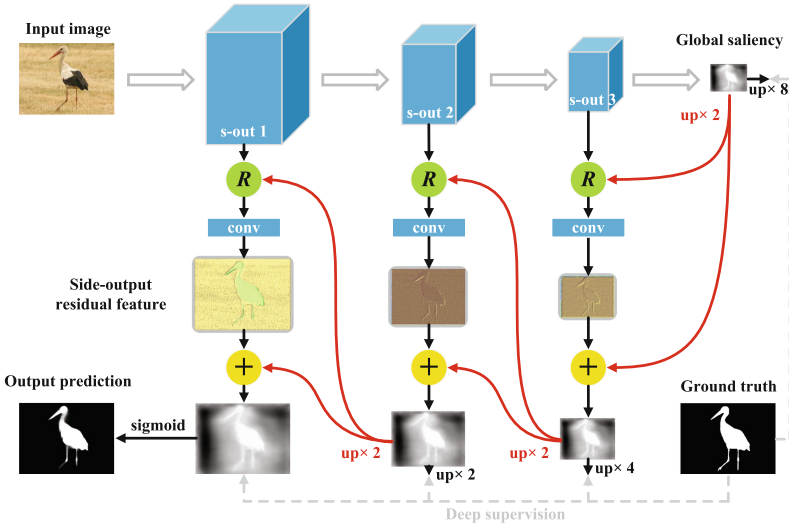
## 2 Related Work

There are plenty of saliency detection methods proposed in the past two decades. Here, we only focus on the recent state-of-the-art methods. Almost all of them are FCNs based and try to solve the common problem: how to produce saliency map with high resolution by using FCNs? Kuen *et al.* [28] applied recurrent unit into FCNs to iteratively refine each salient region. Hu *et al.* [23] extended a superpixel-based guided filter to be a layer in the network for boundary refinement. Hou *et al.* [11] designed short connections for multi-scale feature fusion, while in Amulet [13], multi-level convolutional features were aggregated adaptively. Luo *et al.* [10] proposed a multi-resolution grid structure to capture both local and global cues. In addition, a new loss function was introduced to penalize errors on the boundaries. Zhang *et al.* [14] further proposed a novel upsampling method to reduce the artifacts produced in deconvolution. Recently, dilated convolution [23] and dense connections [17] are further incorporated to obtain high resolution saliency map. There are also some progressive works to address the above issue in semantic segmentation. In [19], skip connections was proposed to refine object instances, while in [29], it was used to build a Laplacian pyramid reconstruction network for object boundary refinement.

Instead of fusing multi-level convolutional features as the above works, we try to learn residual feature for low resolution refinement. The idea of residual learning was first proposed by He *et al.* [25] for image classification. After that, it was widely applied in various applications. Ke *et al.* [21] learned side-output residual feature for accurate object symmetry detection. Kim *et al.* [27] built a very deep convolutional network based on residual learning for accurate image super-resolution.

Although it is natural to apply it for salient object detection, the performance is not satisfactory enough. To solve it, we introduce attention mechanism which is inspired from human perception process. By using top information to efficiently guide bottom-up feedforward process, it has achieved great success in many tasks. Attention model was designed to weight multi-scale features in [12, 30]. Residual attention module was started to generate deep attention-aware features for image classification in [31]. In ILSVRC 2017 Image Classification Challenge, Hu *et al.* [32] won the 1st place by constructing Squeeze-and-Excitation block for channel attention. Huang *et al.* [33] designed an attention mask to highlight the prediction of the reverse object class, which then be subtracted from the original prediction to correct the mistakes in the confusion area for semantic segmentation. Inspired but differed from it, we employ reverse attention in a

top-down manner to guide side-output residual learning. Benefit from it, we can learn more accurate residual details which leads to significant improvement.



**Fig. 3.** The overall architecture of the proposed network. Here, only three side-outputs are listed for illustration. “ $R$ ” denotes the proposed reverse attention block that is illustrated in Fig. 4. As can be seen, the residual error decreases along the stacking orientation with the supervision both on the input and output of the residual unit (yellow circle). (Color figure online)

### 3 Proposed Method

In this section, we first describe the overall architecture of the proposed deep salient object detection network, and then present the details of the main components one by one, which are corresponding to side-output residual learning and top-down reverse attention respectively.

#### 3.1 Architecture

The proposed network is built upon the HED [5] architecture and chooses VGG-16 [34] as backbone. We use the layers up to “pool5” and select {conv1\_2, conv2\_2, conv3\_3, conv4\_3, conv5\_3} as side-outputs, which have strides of {1, 2, 4, 8, 16} pixels with respect to the input image respectively. We first reduce the dimension of “pool5” into 256 by convolution with kernel size  $1 \times 1$ , and then add three convolutional layers with  $5 \times 5$  kernels to capture global saliency. Since the resolution of the global saliency map is only  $1/32$  of the input image, we further learn residual feature in each side-output to improve its resolution gradually. In specifically,  $D$  convolutional layers with  $3 \times 3$  kernels and 64 channels are

stacked for residual learning. The reverse attention block is embedded before side-output residual learning. The prediction of the shallowest side-output is fed into a sigmoid layer for final output. The overall architecture is shown in Fig. 3 and complete configurations are outlined in Table 1.

**Table 1.** The configurations of the proposed network.  $(n, k \times k) \times D$  denotes stacking  $D$  convolutional layers with channel number  $(n)$  and kernel size  $(k)$ , and ReLU layer is added for nonlinear transformation.

Side output 1-5	Global saliency
$(64, 1 \times 1)$	$(256, 1 \times 1)$
$\{(64, 3 \times 3), \text{ReLU}\} \times D$	$\{(256, 5 \times 5), \text{ReLU}\} \times 3$
$(1, 3 \times 3)$	$(1, 1 \times 1)$

### 3.2 Side-Output Residual Learning

As we know, deep layers of network capture high-level semantic information but messy details, while it is opposite for shallow ones. Based on this observation, multi-level features fusion is a common choice to capture their complementary cues, however, it will degrade the confident prediction of deep layers when combining with shallow ones. In this paper, we implement it in a different yet more efficient way by employing residual learning to remedy the errors between the predicted saliency maps and the ground truth. Specifically, the residual feature is learned by applying deep supervision both on the input and output of the designed residual unit, which is illustrated in Fig. 3. Formally, given the upsampled input saliency map  $S_{i+1}^{up}$  by a factor 2 in side-output stage  $i + 1$ , and the residual feature  $R_i$  learned in side-output stage  $i$ , then the deep supervision can be formulated as:

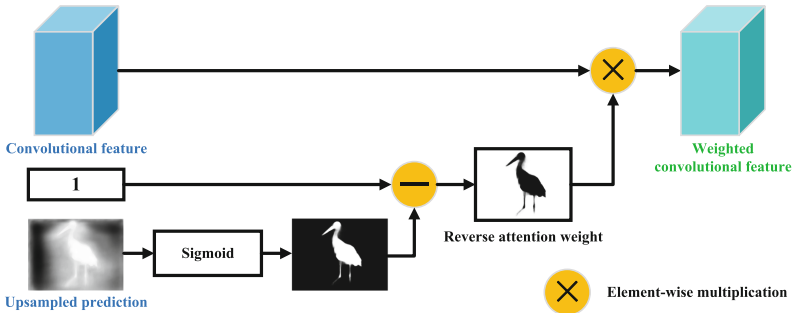
$$\begin{cases} \{S_{i+1}\}^{up \times 2^{i+1}} \approx G \\ \{S_{i+1}^{up} + R_i\}^{up \times 2^i} = \{S_i\}^{up \times 2^i} \approx G \end{cases}, \quad (1)$$

where  $S_i$  is the output of the residual unit and  $G$  is ground truth,  $up \times 2^i$  denotes the upsample operation by a factor  $2^i$ , which is implemented by the same bilinear interpolation with HED [5].

Such a learning objective inherits the following good property. The residual units establish shortcut connections between the predictions from different scales and the ground truth, which makes it easier to remedy their errors with higher scale adaptability. Generally, the error between the input and output of the residual unit is fairly small based on the same supervision, thus can be learned more easily with fewer parameters and iterations. To the extreme, the error is approximately equal to zero if the prediction is close enough to the ground truth. As a result, the constructed network can be very efficient and lightweight.

### 3.3 Top-Down Reverse Attention

Although it is natural and straightforward to learn residual details for saliency refinement, it is not easy for the network to capture them accurately without extra supervision, which will result in unsatisfactory detection. Since most of the existing saliency detection networks are fine-tuned from image classification networks which are only responsive to small and sparse discriminative object parts, it obviously deviates from the requirement of the saliency detection task that needs to explore dense and integral regions for pixel-wise prediction. To mitigate this gap, we propose a reverse attention based side-output residual learning approach for expanding object regions progressively. Starting with a coarse saliency map generated in the deepest layer with high semantic confidence but low resolution, our proposed approach guides the whole network to sequentially discover complement object regions and details by erasing the current predicted salient regions from side-output features, where the current prediction is upsampled from its deeper layer. Such a top-down erasing manner can eventually refine the coarse and low resolution prediction into a complete and high resolution saliency map with these explored regions and details, see Fig. 4 for illustration.



**Fig. 4.** Illustration of the proposed reverse attention block, whose input and output are highlighted in blue and green respectively. (Color figure online)


Given the side-output feature  $T$  and reverse attention weight  $A$ , then the output attentive feature can be produced by their element-wise multiplication, which can be formulated as:

$$F_{z,c} = A_z \cdot T_{z,c}, \quad (2)$$

where  $z$  and  $c$  denote the spatial position of the feature map and the index of the feature channel, respectively. And the reverse attention weight in side-output stage  $i$  is simply generated by subtracting the upsampled prediction of side-output  $i + 1$  from one, which is computed as below:

$$A_i = 1 - \text{Sigmoid}(S_{i+1}^{up}). \quad (3)$$

Figure 5 shows some visual examples of the learned residual feature to illustrate the effectiveness of the proposed reverse attention. As can be seen, the proposed network well captured the residual details near object boundaries with the help of reverse attention. While without reverse attention, it learned some redundant features inside object which is helpless for saliency refinement.



**Fig. 5.** Visualization of residual features in different side-outputs of the proposed network without (the first row) and with reverse attention (the second row). From left to right are saliency map, the last convolutional feature from side output 1 to 4, respectively. After applying our reverse attention, the proposed network well captured spatial details near object boundaries which is beneficial for saliency refinement, especially in shallow layers. Best viewed in color. (Color figure online)

### 3.4 Supervision

As shown in Fig. 3, deep supervision is applied to each side-output stage as did in [5, 11]. Each side-output produces a loss term  $\mathcal{L}_{side}$  which is defined as below:

$$\mathcal{L}_{side}(I, G, W, w) = \sum_{m=1}^M \ell_{side}^{(m)}(I, G, W, w^{(m)}), \quad (4)$$

where  $M$  regards to the total side-output numbers including global saliency,  $W$  denotes the collection of all standard network layer parameters,  $I$  and  $G$  refer to the input image and the corresponding ground truth respectively. Each side-output layer is regarded as a pixel-wise classifier with the corresponding weights  $w$  which is represented by

$$w = (w^{(1)}, w^{(2)}, \dots, w^{(M)}). \quad (5)$$

Here,  $\ell_{side}^{(m)}$  represents the image-level class-balanced cross-entropy loss function [5] of the  $m$ th side output, which is computed by the following formulation:

$$\begin{aligned} \ell_{side}^{(m)}(I, G, W, w^{(m)}) = & - \sum_{z=1}^{|I|} G(z) \log Pr(G(z) = 1 | I(z); W, w^{(m)}) \\ & + (1 - G(z)) \log Pr(G(z) = 0 | I(z); W, w^{(m)}), \end{aligned} \quad (6)$$



where  $Pr(G(z) = 1|I(z); W, w^{(m)})$  represents the probability of the activation value at location  $z$  in the  $m$ th side output,  $z$  is the spatial coordinate. Different with HED [5] and DSS [11], there is no fusion layer included in our approach. The output of the first side-output is used as our final prediction after a sigmoid layer in the testing stage.

### 3.5 Difference to Other Networks

Though shares the same name, the proposed network significantly differs from reverse attention network [33], which applied reverse attention to weight the prediction that is not associated with a target class, in this way to amplify the reverse-class response in the confused region, thus can help the original branch make correct prediction. While in our approach, the usage of reverse attention is totally different. It is used to erase the confident prediction from deep layer, which can guide the network to explore the missing object regions and details effectively.

There are also some significant differences with other residual learning based architectures, *e.g.*, side-output residual network (SRN) [21], and Laplacian reconstruction network (LRN) [29]. In SRN, the residual feature is learned from each side-output of VGG-16 directly, while in this paper, it is learned after reverse attention that is applied to guide residual learning. The main difference with LRN lies in the usage of the weight mask, which is used to weight the learned side-output features for boundary refinement in LRN, in contrast, we apply it before side-output feature learning for guidance. In addition, the weight mask in LRN is generated from the edge of deep prediction which will miss some object regions due to its low resolution, while in this paper, we apply it to focus on all the undetected regions for saliency refinement, which not only refines object boundaries well but also highlights object regions more completely.

## 4 Experiments

### 4.1 Experimental Setup

The proposed network is built on the top of the implementations of HED [5] and DSS [11], and trained through the publicly available Caffe [35] library. The whole network is trained end-to-end using full-resolution images and optimized by stochastic gradient descent method. The hyper-parameters are set as below: batch size (1), iter\_size (10), the momentum (0.9), the weight decay ( $5e-4$ ), learning rate is initialized as  $1e-8$  and decreased by 10% when the training loss reaches a flat, the training iteration number (10K). All these parameters were fixed during the following experiments. The source code will be released<sup>1</sup>.

We comprehensively evaluated our method on six representative datasets, including MSRA-B [36], HKU-IS [37], ECSSD [38], PASCAL-S [39], SOD [40], and DUT-OMRON [41], which contain 5000, 4447, 1000, 850, 300, 5168 well

<sup>1</sup> <http://shuhanchen.net>.

annotated images, respectively. Among them, PASCAL-S and DUT-OMRON are more challenging than the others. To guarantee a fair comparison with the existing approaches, we utilize the same training sets as in [8, 10, 11, 42] and test all of the datasets with the same model. Data augmentation is also implemented the same with [10, 11] to reduce the over-fitting risk, which increased by 2 times through horizontal flipping.

Three standard and widely agreed metrics are used to evaluate the performance, including Precision-Recall (PR) curve, F-measure, and the Mean Absolute Error (MAE). Pairs of precision and recall values are calculated by comparing the binary saliency maps with the ground truth to plot the PR curve, where the thresholds are in the range of [0, 255]. The F-measure is adopted to measure the overall performance, which is defined as the weighted harmonic mean of precision and recall:

$$F_{\beta} = (1 + \beta^2) \frac{\text{Precision} \times \text{Recall}}{\beta^2 \text{Precision} + \text{Recall}}, \quad (7)$$

where  $\beta^2$  is set to 2 to emphasize the precision over recall as suggested in [43]. Only the maximum F-Measure is reported here to show the best performance a detector can achieve. Given the normalized saliency map  $S$  and ground truth  $G$ , the MAE score is calculated by their average per-pixel difference:

$$\text{MAE} = \frac{1}{H \times W} \sum_{x=1}^H \sum_{y=1}^W |S(x, y) - G(x, y)|, \quad (8)$$

where  $W$  and  $H$  are the width and height of the saliency map, respectively.

## 4.2 Ablation Studies

Before comparing with the state-of-the-art methods, we first evaluate the influence of different design options (the depth  $D$ ), the effectiveness of the proposed side-output residual learning and reverse attention in this section.

**Depth  $D$ .** We make a experiment to see how the depth  $D$  affects the performance by varying it from 1 to 3. The results on PASCAL-S and DUT-OMRON are shown in Table 2. As can be seen that the best performance is obtained when  $D=2$ . Therefore, we set it as 2 in the following experiments.

**Side-Output Residual Learning.** To investigate the effectiveness of the side-output residual learning, we separately evaluate the performance of each side-output prediction and show in Table 3. We can find that the performance is gradually improved by combing more side-output residual features.

**Table 2.** Performance comparison with different numbers of  $D$ .

	PASCAL-S		DUT-OMRON	
	$F_\beta$	MAE	$F_\beta$	MAE
$D = 1$	0.830	<b>0.100</b>	0.776	0.067
$D = 2$	<b>0.834</b>	0.104	<b>0.786</b>	<b>0.062</b>
$D = 3$	0.824	0.106	0.778	0.064

**Table 3.** Performance comparison with different side-output predictions.

	PASCAL-S		DUT-OMRON	
	$F_\beta$	MAE	$F_\beta$	MAE
Side-output 5	0.817	0.111	0.755	0.071
Side-output 4	0.827	0.106	0.776	0.065
Side-output 3	0.831	<b>0.104</b>	0.785	<b>0.062</b>
Side-output 2	0.832	<b>0.104</b>	<b>0.786</b>	<b>0.062</b>
Side-output 1	<b>0.834</b>	<b>0.104</b>	<b>0.786</b>	<b>0.062</b>

**Reverse Attention.** As illustrated in Fig. 5, the network well located at the object boundaries with the help of reverse attention. Here, we perform a detailed comparison using F-measure and MAE scores which are reported in Table 4. From the results, we can get the following observations: (1) Without reverse attention, our performance is similar to the state-of-the-art method DSS (without CRF-based post-processing), which indicates its large redundancy. (2) After applying reverse attention, the performance is improved by a large margin, specifically, we obtained an average of 1.4% gain in terms of F-measure and 0.5% decrease for MAE score, which clearly demonstrates its effectiveness.


### 4.3 Performance Comparison with State-of-the-art

We compare the proposed method with 10 state-of-the-art ones, including 9 recent CNN-based approaches, DCL<sup>+</sup> [8], DHS [44], SSD [45], RFCN [9], DLS [23], NLDF [10], DSS and DSS<sup>+</sup> [11], Amulet [13], UCF [14], and one conventional top approach, DRFI [42], where symbol “+” indicates that the network includes CRF-based post-processing. Note that all the saliency maps of the above methods are produced by running source codes or pre-computed by the authors, and ResNet based methods are not included for fair comparison.

**Quantitative Evaluation.** The results of quantitative comparison with state-of-the-art methods are reported in Table 4 and Fig. 7. We can clearly observe that our approach significantly outperforms the competing methods both in terms of F-measure and MAE scores, especially on the challenging datasets (*e.g.*, DUT-OMRON). For PR curves, we also achieved comparable performance with state-of-the-arts except at high level of recall (recall > 0.9). In comparison to the top method, DSS<sup>+</sup>, which uses a CRF-based post-processing step to



**Fig. 6.** Visual comparisons with the existing methods in some challenging cases: complex scenes, low contrast, and multiple (small) salient objects. (Color figure online)



(a) ECSSD (b) SOD (c) DUT-OMRON

**Fig. 7.** Comparison of precision-recall curves on different datasets.

**Table 4.** Quantitative comparison with state-of-the-art methods on six benchmark datasets. Each cell (from up to down) contains max F-measure (higher better), and MAE (lower better). The top two results are highlighted in **red** and **green** respectively. “RA” denotes the proposed reverse attention, and “MK” is MSRA-10K [46], the other abbreviations are the initials of each dataset mentioned in the paper. Note that the number of images listed here are including the augmented ones.

	Training		MSRA-B	HKU-IS	ECSSD	PASCAL-S	SOD	DUT-OMRON
	Dataset	#Images						
DRFI [42]	MB	2.5k	0.851	0.775	0.784	0.690	0.699	0.664
			0.123	0.146	0.172	0.210	0.223	0.150
DCL <sup>+</sup> [8]	MB	2.5k	0.918	0.907	0.898	0.810	0.831	0.757
			0.047	0.048	0.071	0.115	0.131	0.080
DHS [44]	MK+D	9.5k × 12	-	0.892	0.905	0.824	0.823	-
			-	0.052	0.061	<b>0.094</b>	0.127	-
SSD [45]	MB	2.5k	0.902	-	0.865	0.774	0.793	0.754
			0.160	-	0.193	0.220	0.222	0.193
RFCN [9]	MK	10k	-	0.894	0.889	0.829	0.799	0.744
			-	0.088	0.109	0.133	0.169	0.111
DLS [23]	MK	10k	-	0.835	0.852	0.753	-	0.687
			-	0.070	0.088	0.132	-	0.090
NLDF [10]	MB	2.5k × 2	0.911	0.902	0.903	0.826	0.837	0.753
			0.048	0.048	0.065	0.099	0.123	0.080
Amulet [13]	MK	10k × 8	-	0.899	0.914	0.832	0.795	0.743
			-	0.050	0.061	0.100	0.144	0.098
UCF [14]	MK	10k × 8	-	0.888	0.902	0.818	0.805	0.730
			-	0.061	0.071	0.116	0.148	0.120
DSS [11]	MB	2.5k × 2	0.920	0.900	0.908	0.826	0.834	0.764
			0.043	0.050	0.063	0.102	0.126	0.072
DSS <sup>+</sup> [11]	MB	2.5k × 2	<b>0.929</b>	<b>0.916</b>	<b>0.919</b>	<b>0.835</b>	<b>0.843</b>	<b>0.781</b>
			<b>0.034</b>	<b>0.040</b>	<b>0.055</b>	<b>0.095</b>	<b>0.122</b>	<b>0.063</b>
Ours w/o RA	MB	2.5k × 2	0.919	0.898	0.905	0.818	0.839	0.762
			0.042	0.049	0.063	0.106	0.126	0.071
Ours	MB	2.5k × 2	<b>0.931</b>	<b>0.913</b>	<b>0.918</b>	<b>0.834</b>	<b>0.844</b>	<b>0.786</b>
			<b>0.036</b>	<b>0.045</b>	<b>0.059</b>	0.104	<b>0.124</b>	<b>0.062</b>

refine the resolution, nevertheless, our approach still attains nearly identical (or better) performance across the board. It also needs to point out that the existing methods used different training datasets and data augmentation strategies, which caused an unfair comparison. Nevertheless, we still perform much better that clearly shows the superiority of the proposed approach. And we also believe that further performance gain can be obtained by using larger training dataset with more augmented training images, which is beyond the scope of this paper.

**Qualitative Evaluation.** We also show some visual results of some representative images to exhibit the superiority of the proposed approach in Fig. 6, including complex scenes, low contrast between salient object and background, multiple (small) salient objects with diverse characteristics (*e.g.*, size, color).

**Table 5.** Average execution time comparison with other methods on ECSSD.

	DHS	DSS	NLDF	UCF	Amulet	Ours
Times(s)	0.026	0.048	0.048	0.168	0.080	<b>0.022</b>

Taking all the cases into account, it can be observed clearly that our approach not only highlights the salient regions correctly with less false detection but also produces sharp boundaries and coherent details (*e.g.*, the *mouth* of the bird in the 4th row of Fig. 6). It is also interesting to note that the proposed method even corrected some false labeling in the ground truth, *e.g.*, the left *horn* in the 7th row of Fig. 6. Nevertheless, we still obtain unsatisfactory results in some challenging cases, taking the last row of Fig. 6 for example, to segment all the salient objects completely is still very difficult for the existing methods.

**Execution Time.** Finally, we investigate the efficiency of our method, and conduct all the experiments on a single NVIDIA TITAN Xp GPU for fair comparison. It only takes less than 2 h to train our model, for comparison, DSS needs about 6 h. We also compared the average execution time with other five leading CNN-based methods on ECSSD. As can be seen from Table 5, our approach is much faster than all the competing methods. Therefore, considering both in visual quality and efficiency, our approach is the best choice for real-time applications up to now.

## 5 Conclusions

As a low-level pre-processing step, salient object detection has great applicability in various high-level tasks yet remains not being well solved, which mainly lies on the following two aspects: low resolution output and heavy model weight. In this paper, we presented an accurate yet compact deep network for efficient salient object detection. Instead of directly learning multi-scale saliency features in different side-output stages, we employ residual learning to learn side-output residual features for saliency refinement. Based on it, the resolution of the global saliency map generated by the deepest convolutional layer was improved gradually with very limited parameters. We further propose reverse attention to guide such side-output residual learning in a top-down manner. Benefit from it, our network learned more accurate residual features, which leads to significant performance improvement. Extensive experimental results demonstrate that the proposed approach performs favorably against state-of-the-art ones both in quantitative and qualitative comparisons, which enables it a better choice for further real-world applications, and also makes it a great potential to apply in other end-to-end pixel-level prediction tasks. Nevertheless, the global saliency branch and backbone (VGG-16) network still contain large redundancy, which will be further explored by introducing handcrafted saliency prior and learning from scratch in our future work.

**Acknowledgments.** This work was supported by the Natural Science Foundation of China (No. 61502412), Natural Science Foundation for Youths of Jiangsu Province (No. BK20150459), Foundation of Yangzhou University (No. 2017CXJ026).

## References

1. Wei, Y., et al.: STC: a simple to complex framework for weakly-supervised semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(11), 2314–2320 (2017)
2. Xu, K., et al.: Show, attend and tell: neural image caption generation with visual attention. In: *ICML*, pp. 2048–2057 (2015)
3. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *CVPR*, pp. 3431–3440 (2015)
4. Dai, J., He, K., Li, Y., Ren, S., Sun, J.: Instance-sensitive fully convolutional networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016 Part VI. LNCS*, vol. 9910, pp. 534–549. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46466-4\\_32](https://doi.org/10.1007/978-3-319-46466-4_32)
5. Xie, S., Tu, Z.: Holistically-nested edge detection. In: *ICCV*, pp. 1395–1403 (2015)
6. Li, X., et al.: DeepSaliency: multi-task deep neural network model for salient object detection. *IEEE Trans. Image Proc.* **25**(8), 3919–3930 (2016)
7. Lee, G., Tai, Y.W., Kim, J.: Deep saliency with encoded low level distance map and high level features. In: *CVPR*, pp. 660–668 (2016)
8. Li, G., Yu, Y.: Deep contrast learning for salient object detection. In: *CVPR*, pp. 478–487 (2016)
9. Wang, L., Wang, L., Lu, H., Zhang, P., Ruan, X.: Saliency detection with recurrent fully convolutional networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016 Part IV. LNCS*, vol. 9908, pp. 825–841. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46493-0\\_50](https://doi.org/10.1007/978-3-319-46493-0_50)
10. Luo, Z., Mishra, A., Achkar, A., Eichel, J., Li, S., Jodoin, P.M.: Non-local deep features for salient object detection. In: *CVPR*, pp. 6593–6601 (2017)
11. Hou, Q., Cheng, M.M., Hu, X., Borji, A., Tu, Z., Torr, P.: Deeply supervised salient object detection with short connections. In: *CVPR*, pp. 5300–5309 (2017)
12. Li, G., Xie, Y., Lin, L., Yu, Y.: Instance-level salient object segmentation. In: *CVPR*, pp. 247–256 (2017)
13. Zhang, P., Wang, D., Lu, H., Wang, H., Ruan, X.: Amulet: aggregating multi-level convolutional features for salient object detection. In: *ICCV*, pp. 202–211 (2017)
14. Zhang, P., Wang, D., Lu, H., Wang, H., Yin, B.: Learning uncertain convolutional features for accurate saliency detection. In: *ICCV*, pp. 212–221 (2017)
15. Chen, T., Lin, L., Liu, L., Luo, X., Li, X.: DISC: deep image saliency computing via progressive representation learning. *IEEE Trans. Neural Netw. Learn. Syst.* **27**(6), 1135–1149 (2016)
16. Tang, Y., Wu, X.: Saliency detection via combining region-level and pixel-level predictions with CNNs. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016 Part VIII. LNCS*, vol. 9912, pp. 809–825. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46484-8\\_49](https://doi.org/10.1007/978-3-319-46484-8_49)
17. Xiao, H., Feng, J., Wei, Y., Zhang, M.: Deep salient object detection with dense connections and distraction diagnosis. *IEEE Trans. Multimedia* (2018)
18. Olaf, R., Philipp, F., Thomas, B.: U-net: convolutional networks for biomedical image segmentation. In: *MICCAI*, pp. 234–241 (2015)

19. Pinheiro, P.O., Lin, T.-Y., Collobert, R., Dollár, P.: Learning to refine object segments. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016 Part I. LNCS, vol. 9905, pp. 75–91. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46448-0\\_5](https://doi.org/10.1007/978-3-319-46448-0_5)
20. Liu, Y., Yao, J., Li, L., Lu, X., Han, J.: Learning to refine object contours with a top-down fully convolutional encoder-decoder network. In: ArXiv e-prints (2017)
21. Ke, W., Chen, J., Jiao, J., Zhao, G., Ye, Q.: SRN: side-output residual network for object symmetry detection in the wild. In: CVPR, pp. 302–310 (2017)
22. Shen, W., Zhao, K., Jiang, Y., Wang, Y., Bai, X., Yuille, A.: DeepSkeleton: learning multi-task scale-associated deep side outputs for object skeleton extraction in natural images. *IEEE Trans. Image Proc.* **26**(11), 5298–5311 (2017)
23. Hu, P., Shuai, B., Liu, J., Wang, G.: Deep level sets for salient object detection. In: CVPR, pp. 2300–2309 (2017)
24. Krähenbühl, P., Koltun, V.: Efficient inference in fully connected CRFs with Gaussian edge potentials. In: NIPS, pp. 109–117 (2011)
25. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR, pp. 770–778 (2016)
26. Lai, W.S., Huang, J.B., Ahuja, N., Yang, M.H.: Deep Laplacian pyramid networks for fast and accurate super-resolution. In: CVPR, pp. 624–632 (2017)
27. Kim, J., Kwon Lee, J., Mu Lee, K.: Accurate image super-resolution using very deep convolutional networks. In: CVPR, pp. 1646–1654 (2016)
28. Kuen, J., Wang, Z., Wang, G.: Recurrent attentional networks for saliency detection. In: CVPR, pp. 3668–3677 (2016)
29. Ghiasi, G., Fowlkes, C.C.: Laplacian pyramid reconstruction and refinement for semantic segmentation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016 Part III. LNCS, vol. 9907, pp. 519–534. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46487-9\\_32](https://doi.org/10.1007/978-3-319-46487-9_32)
30. Chen, L.C., Yang, Y., Wang, J., Xu, W., Yuille, A.L.: Attention to scale: Scale-aware semantic image segmentation. In: CVPR, pp. 3640–3649 (2016)
31. Wang, F., et al.: Residual attention network for image classification. In: CVPR, pp. 6450–6458 (2017)
32. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: ArXiv e-prints (2017)
33. Huang, Q., et al.: Semantic segmentation with reverse attention. In: BMVC (2017)
34. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ArXiv e-prints (2014)
35. Jia, Y., Shelhamer, E., et al.: Caffe: convolutional architecture for fast feature embedding. In: ACM Multimedia, pp. 675–678 (2014)
36. Liu, T., et al.: Learning to detect a salient object. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(2), 353–367 (2011)
37. Li, G., Yu, Y.: Visual saliency detection based on multiscale deep cnn features. *IEEE Trans. Image Proc.* **25**(11), 5012–5024 (2016)
38. Shi, J., Yan, Q., Xu, L., Jia, J.: Hierarchical image saliency detection on extended CSSD. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(4), 717–729 (2016)
39. Li, Y., Hou, X., Koch, C., Rehg, J.M., Yuille, A.L.: The secrets of salient object segmentation. In: CVPR, pp. 280–287 (2014)
40. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: ICCV, pp. 416–423 (2001)
41. Yang, C., Zhang, L., Lu, H., Ruan, X., Yang, M.H.: Saliency detection via graph-based manifold ranking. In: CVPR, pp. 3166–3173 (2013)



42. Jiang, H., Wang, J., Yuan, Z., Wu, Y., Zheng, N., Li, S.: Salient object detection: A discriminative regional feature integration approach. In: CVPR. 2083–2090 (2013)
43. Borji, A., Cheng, M.M., Jiang, H., Li, J.: Salient object detection: a benchmark. *IEEE Trans. Image Proc.* **24**(12), 5706–5722 (2015)
44. Liu, N., Han, J.: DHSNet: deep hierarchical saliency network for salient object detection. In: CVPR, pp. 678–686 (2016)
45. Kim, J., Pavlovic, V.: A shape-based approach for salient object detection using deep learning. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016 Part IV. LNCS, vol. 9908, pp. 455–470. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46493-0\\_28](https://doi.org/10.1007/978-3-319-46493-0_28)
46. Cheng, M.M., Mitra, N.J., Huang, X., Torr, P.H., Hu, S.M.: Global contrast based salient region detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **37**(3), 569–582 (2015)