

Salient Objects in Clutter: Bringing Salient Object Detection to the Foreground

Deng-Ping Fan¹[0000-0002-5245-7518], Ming-Ming Cheng¹[0000-0001-5550-8758],
Jiang-Jiang Liu¹[0000-0002-1341-2763], Shang-Hua Gao¹[0000-0002-7055-2703],
Qibin Hou¹[0000-0002-8388-8708], and Ali Borji²[0000-0001-8198-0335]

¹ College of Computer Science, Nankai University, Tianjin, China

² CRCV, University of Central Florida, Orlando, Florida, U.S

<http://mmcheng.net/SOCBenchmark/>

Abstract. We provide a comprehensive evaluation of salient object detection (SOD) models. Our analysis identifies a serious design bias of existing SOD datasets which assumes that each image contains at least one clearly outstanding salient object in low clutter. The design bias has led to a saturated high performance for state-of-the-art SOD models when evaluated on existing datasets. The models, however, still perform far from being satisfactory when applied to real-world daily scenes. Based on our analyses, we first identify 7 crucial aspects that a comprehensive and balanced dataset should fulfill. Then, we propose a new high quality dataset and update the previous saliency benchmark. Specifically, our **SOC** (Salient Objects in Clutter) dataset, includes images with salient and non-salient objects from daily object categories. Beyond object category annotations, each salient image is accompanied by attributes that reflect common challenges in real-world scenes. Finally, we report attribute-based performance assessment on our dataset.

Keywords: Salient object detection · Saliency benchmark · Dataset · Attribute

1 Introduction

This paper considers the task of salient object detection (SOD). Visual saliency mimics the ability of the human visual system to select a certain subset of the visual scene. SOD aims to detect the most attention-grabbing objects in a scene and then extract pixel-accurate silhouettes of the objects. The merit of SOD lies in its applications in many other computer vision tasks including: visual tracking [4], image retrieval [14, 16], computer graphics [9], content aware image resizing [45], and weakly supervised semantic segmentation [18, 39, 40].

Our work is motivated by two observations. First, existing SOD datasets [2, 5, 10, 11, 23, 26, 29, 32, 43, 44] are flawed either in the data collection procedure or quality of the data. Specifically, most datasets assume that an image contains at least one salient object, and thus discard images that do not contain salient



Fig. 1. Sample images from our new dataset including *non-salient* object images (first row) and salient object images (rows 2 to 4). For salient object images, instance-level ground truth map (different color), object attributes (Attr) and category labels are provided. Please refer to the supplemental material for more illustrations of our dataset.

objects. We call this *data selection bias*. Moreover, existing datasets mostly contain images with a single object or several objects (often a person) in low clutter. These datasets do not adequately reflect the complexity of images in the real world where scenes usually contain multiple objects amidst lots of clutter. As a result, all top performing models trained on the existing datasets have nearly saturated the performance (*e.g.*, > 0.9 *F-measure* over most current datasets) but unsatisfactory performance on realistic scenes (*e.g.*, < 0.45 *F-measure* in Table 3). Because current models may be biased towards ideal conditions, their effectiveness may be impaired once they are applied to real world scenes. To solve this problem, it is important to introduce a dataset that reaches closer to realistic conditions.

Second, only the overall performance of the models can be analyzed over existing datasets. None of the datasets contains various attributes that reflect challenges in real-world scenes. Having attributes helps 1) gain a deeper insight into the SOD problem, 2) investigate the pros and cons of the SOD models, and 3) objectively assess the model performances over different perspectives, which might be diverse for different applications.

Considering the above two issues, we make two contributions. Our main contribution is the collection of a new high quality SOD dataset, named the **SOC**, Salient Objects in Clutter. To date, SOC is the largest instance-level SOD dataset and contains 6,000 images from more than 80 common categories. It differs from existing datasets in three aspects: 1) salient objects have category

annotation which can be used for new research such as weakly supervised SOD tasks, 2) the inclusion of non-salient images which make this dataset closer to the real-world scenes and more challenging than the existing ones, and 3) salient objects have attributes reflecting specific situations faced in the real-world such as *motion blur*, *occlusion* and *cluttered background*. As a consequence, our SOC dataset narrows the *gap* between existing datasets and the real-world scenes and provides a more realistic benchmark (see Fig. 1).

In addition, we provide a comprehensive evaluation of several state-of-the-art convolutional neural networks (CNNs) based models [8, 15, 17, 23, 24, 28, 31, 36, 38, 48–51]. To evaluate the models, we introduce three metrics that measure the region similarity of the detection, the pixel-wise accuracy of the segmentation, and the structure similarity of the result. Furthermore, we give an attribute-based performance evaluation. These attributes allow a deeper understanding of the models and point out promising directions for further research.

We believe that our dataset and benchmark can be very influential for future SOD research in particular for application-oriented model development. The entire dataset and analyzing tools will be released freely to the public.

2 Related Works

In this section, we briefly discuss existing datasets designed for SOD tasks, especially in the aspects including annotation type, the number of salient objects per image, number of images, and image quality. We also review the CNNs based SOD models.




Fig. 2. Previous SOD datasets only annotate the image by drawing (b) pixel-accurate silhouettes of salient objects. Different from (d) MS COCO object segmentation dataset [27] (Objects are not necessarily being *salient*), our work focuses on (c) segmenting *salient object instances*.

2.1 Datasets

Early datasets are either limited in the number of images or in their coarse annotation of salient objects. For example, the salient objects in datasets **MSRA-A** [29] and **MSRA-B** [29] are roughly annotated in the form of bounding boxes. **ASD** [1] and **MSRA10K** [11] mostly contain only one salient object in each image, while the **SED2** [2] dataset contains two objects in a single image

but contains only 100 images. To improve the quality of datasets, researchers in recent years started to collect datasets with multiple objects in relatively complex and cluttered backgrounds. These datasets include **DUT-OMRON** [44], **ECSSD** [43], **Judd-A** [5], and **PASCAL-S** [26]. These datasets have been improved in terms of annotation quality and the number of images, compared to their predecessors. Datasets **HKU-IS** [23], **XPIE** [41], and **DUTS** [37] resolved the shortcomings by collecting large amounts of pixel-wise labeled images (Fig. 2 (b)) with more than one salient object in images. However, they ignored the non-salient objects and did not offer instance-level (Fig. 2 (c)) salient objects annotation. Beyond these, researchers of [19] collected about 6k *simple background images* (most of them are pure texture images) to account for the non-salient scenes. This dataset is not sufficient to reflect real scenes as the real-world scenes are more complicated. The **ILSO** [22] dataset contains instance-level salient objects annotation but has boundaries roughly labeled as shown in Fig. 5 (a).

To sum up, as discussed above, existing datasets mostly focus on images with clear salient objects in simple backgrounds. Taking into account the aforementioned limitations of existing datasets, a more realistic dataset which contains realistic scenes with non-salient objects, textures “in the wild”, and salient objects with attributes, is needed for future investigations in this field. Such a dataset can offer deep insights into weaknesses and strengths of SOD models.

2.2 Models

We divide the state-of-the-art deep models for SOD based on the number of tasks.

Single-task models have the single goal of detecting the salient objects in images. In **LEGS** [36], local information and global contrast were separately captured by two different deep CNNs, and were then fused to generate a saliency map. In [51], Zhao *et al.* presented a multi-context deep learning framework (**MC**) for SOD. Li *et al.* [23] (**MDF**) proposed to use multi-scale features extracted from a deep CNNs to derive a saliency map. Li *et al.* [24] presented a deep contrast network (**DCL**), which not only considered the pixel-wise information but also fused the segment-level guidance into the network. Lee *et al.* [15] (**ELD**) considered both high-level features extracted from CNNs and hand-crafted features. Liu *et al.* [28] (**DHS**) designed a two-stage network, in which a coarse downsampled prediction map was produced. It is then followed by another network to refine the details and upsample the prediction map hierarchically and progressively. Long *et al.* [30] proposed a fully convolutional network (**FCN**) to make dense pixel prediction problem feasible for end-to-end training. **RFCN** [38] used a recurrent FCN to incorporate the coarse predictions as saliency priors and refined the generated predictions in a stage-wise manner. The **DISC** [8] framework was proposed for fine-grained image saliency computing. Two stacked CNNs were utilized to obtain coarse-level and fine-grained saliency maps, respectively. **IMC** [48] integrated saliency cues at different levels through FCN. It could efficiently exploit both learned semantic cues and higher-order region statistics for edge-accurate SOD. Recently, a deep architecture [17] with

Table 1. CNNs based SOD models. We divided these models into single-task (S-T) and multi-task (M-T). **Training Set:** MB is the MSRA-B dataset [29]. MK is the MSRA-10K [11] dataset. ImageNet dataset refers to [34]. D is the DUT-OMRON [44] dataset. H is the HKU-IS [23] dataset. P is the PASCAL-S [26] dataset. P2010 is the PASCAL VOC 2010 semantic segmentation dataset [12]. **Base Model:** VGGNet, ResNet-101, AlexNet, GoogleNet are base models. **FCN:** whether model uses the fully convolutional network. **Sp:** whether model uses superpixels. **Proposal:** whether model uses the object proposal. **Edge:** whether model uses the edge or contour information

	No	Model	Year	Pub	#Training	Training Set	Base Model	FCN	Sp	Proposal	Edge
S-T	1	LEGS [36]	2015	CVPR	3,340	MB + P	—	×	×	✓	×
	2	MC [51]	2015	CVPR	8,000	MK	GoogLeNet	×	✓	×	×
	3	MDF [23]	2015	CVPR	2,500	MB	—	×	✓	×	✓
	4	DCL [24]	2016	CVPR	2,500	MB	VGGNet	✓	✓	×	×
	5	ELD [15]	2016	CVPR	9,000	MK	VGGNet	×	✓	×	×
	6	DHS [28]	2016	CVPR	9,500	MK+D	VGGNet	×	×	×	×
	7	RFCN [38]	2016	ECCV	10,103	P2010	—	✓	×	×	✓
	8	DISC [8]	2016	TNNLS	9,000	MK	—	×	✓	×	×
	9	IMC [48]	2017	WACV	6,000	MK	ResNet-101	✓	✓	×	×
	10	DSS [17]	2017	CVPR	2,500	MB	VGGNet	✓	×	×	✓
	11	NLDF [31]	2017	CVPR	2,500	MB	VGGNet	✓	×	×	×
	12	AMU [49]	2017	ICCV	10,000	MK	VGGNet	✓	×	×	✓
	13	UCF [50]	2017	ICCV	10,000	MK	—	✓	×	×	×
M-T	1	DS [25]	2016	TIP	10,000	MK	VGGNet	✓	✓	×	×
	2	WSS [37]	2017	CVPR	456K	ImageNet	VGGNet	✓	✓	×	×
	3	MSR [22]	2017	CVPR	5,000	MB + H	VGGNet	✓	×	✓	✓

short connections (**DSS**) was proposed. Hou *et al.* added connections from high-level features to low-level features based on the HED [42] architecture, achieving good performance. **NLDF** [31] integrated local and global features and added a boundary loss term into standard cross entropy loss to train an end-to-end network. **AMU** [49] was a generic aggregating multi-level convolutional feature framework. It integrated coarse semantics and fine detailed feature maps into multiple resolutions. Then it adaptively learned to combine these feature maps at each resolution and predicted saliency maps with the combined features. **UCF** [50] was proposed to improve the robustness and accuracy of saliency detection. They introduced a reformulated dropout after specific convolutional layers to construct an uncertain ensemble of internal feature units. Also, they proposed reformulated dropout after an effective hybrid up-sampling method to reduce the checkerboard artifacts of deconvolution operators in the decoder network.

Multi-task models at present include three methods, **DS**, **WSS**, and **MSR**. The **DS** [25] model set up a multi-task learning scheme for exploring the intrinsic correlations between saliency detection and semantic image segmentation, which shared the information in FCN layers to generate effective features for object perception. Recently, Wang *et al.* [37] proposed a model named **WSS** which developed a weakly supervised learning method using image-level tags for saliency detection. First, they jointly trained Foreground Inference Net (FIN) and FCN for image categorization. Then, they used FIN fine-tuned with iterative CRF to enforce spatial label consistency to predict the saliency map. **MSR** [22] was designed for both salient region detection and salient object contour detection,

integrated with multi-scale combinatorial grouping and a MAP-based [47] subset optimization framework. Using three refined VGG network streams with shared parameters and a learned attentional model for fusing results at different scales, the authors were able to achieve good results.

We benchmark a large set of the state-of-the-art CNNs based models (see Table 1) on our proposed dataset, highlighting the current issues and pointing out future research directions.

3 The Proposed Dataset

In this section, we present our new challenging SOC dataset designed to reflect the real-world scenes in detail. Sample images from SOC are shown in Fig. 1. Moreover, statistics regarding the categories and the attributes of SOC are shown in Fig. 4 (a) and Fig. 6, respectively. Based on the strengths and weaknesses of the existing datasets, we identify seven crucial aspects that a comprehensive and balanced dataset should fulfill.

1) Presence of Non-Salient Objects. Almost all of the existing SOD datasets make the assumption that an image contains at least one salient object and discard the images that do not contain salient objects. However, this assumption is an ideal setting which leads to *data selection bias*. In a realistic setting, images do not always contain salient objects. For example, some amorphous background images such as sky, grass and texture contain no salient objects at all [6]. The non-salient objects or background “stuff” may occupy the entire scene, and hence heavily constrain possible locations for a salient object. Xia *et al.* [41] proposed a state-of-the-art SOD model by judging what is or what is not a salient object, indicating that the non-salient object is crucial for reasoning about the salient object. This suggests that the non-salient objects deserve equal attention as the salient objects in SOD. Incorporating a number of images containing non-salient objects makes the dataset closer to real-world scenes, while becoming more challenging. Thus, we define the “*non-salient objects*” as images without salient objects or images with “stuff” categories. As suggested in [6, 41], the “stuff” categories including (a) densely distributed similar objects, (b) fuzzy shape, and (c) region without semantics, which are illustrated in Fig. 3 (a)-(c), respectively.



Fig. 3. Some examples of non-salient objects.

Based on the characteristics of non-salient objects, we collected 783 texture images from the DTD [21] dataset. To enrich the diversity, 2217 images including aurora, sky, crowds, store and many other kinds of realistic scenes were gathered from the Internet and other datasets [26,27,32,35]. We believe that incorporating enough non-salient objects would open up a promising direction for future works.

2) Number and Category of Images. A considerably large amount of images is essential to capture the diversity and abundance of real-world scenes. Moreover, with large amounts of data, SOD models can avoid over-fitting and enhance generalization. To this end, we gathered 6,000 images from more than 80 categories, containing 3,000 images with salient objects and 3,000 images without salient objects. We divide our dataset into training set, validation set and test set in the ratio of 6:2:2. To ensure fairness, the test set is not published, but with the *on-line testing* provided on our website³. Fig. 4 (a) shows the number of salient objects for each category. It shows that the “person” category accounts for a large proportion, which is reasonable as people usually appear in daily scenes along with other objects.

3) Global/Local Color Contrast of Salient Objects. As described in [26], the term “salient” is related to the global/local contrast of the foreground and background. It is essential to check whether the salient objects are easy to detect. For each object, we compute RGB color histograms for foreground and background separately. Then, χ^2 distance is utilized to measure the distance between the two histograms. The global and local color contrast distribution are shown in Fig. 4 (b) and (c), respectively. In comparison to ILSO, our SOC has more proportion of objects with low global color contrast and local color contrast.

4) Locations of Salient Objects. *Center bias* has been identified as one of the most significant biases of saliency detection datasets [3, 20, 26]. Fig. 4 (d) illustrates a set of images and their overlay map. As can be seen, although salient objects are located in different positions, the overlay map still shows that somehow this set of images is center biased. Previous benchmarks often adopt this incorrect way to analyze the location distribution of salient objects. To avoid this misleading phenomenon, we plot the statistics of two quantities r_o and r_m in Fig. 4 (e), where r_o and r_m denote how far an object center and the farthest (margin) point in an object are from the image center, respectively. Both r_o and r_m are divided by half image diagonal length for normalization so that $r_o, r_m \in [0, 1]$. From these statistics, we can observe that salient objects in our dataset do not suffer from center bias.

5) Size of Salient Objects. The size of an instance-level salient object is defined as the proportion of pixels in the image [26]. As shown in Fig. 4 (g), the size of salient objects in our SOC varies in a broader range, compared with the only existing instance-level ILSO [22] dataset. Also, medium-sized objects in SOC have a higher proportion.

6) High-Quality Salient Object Labeling. As also noticed in [17], training on the ECSSD dataset (1,000) allows to achieve better results than other

³ <http://dpfan.net/SOCBenchmark/>

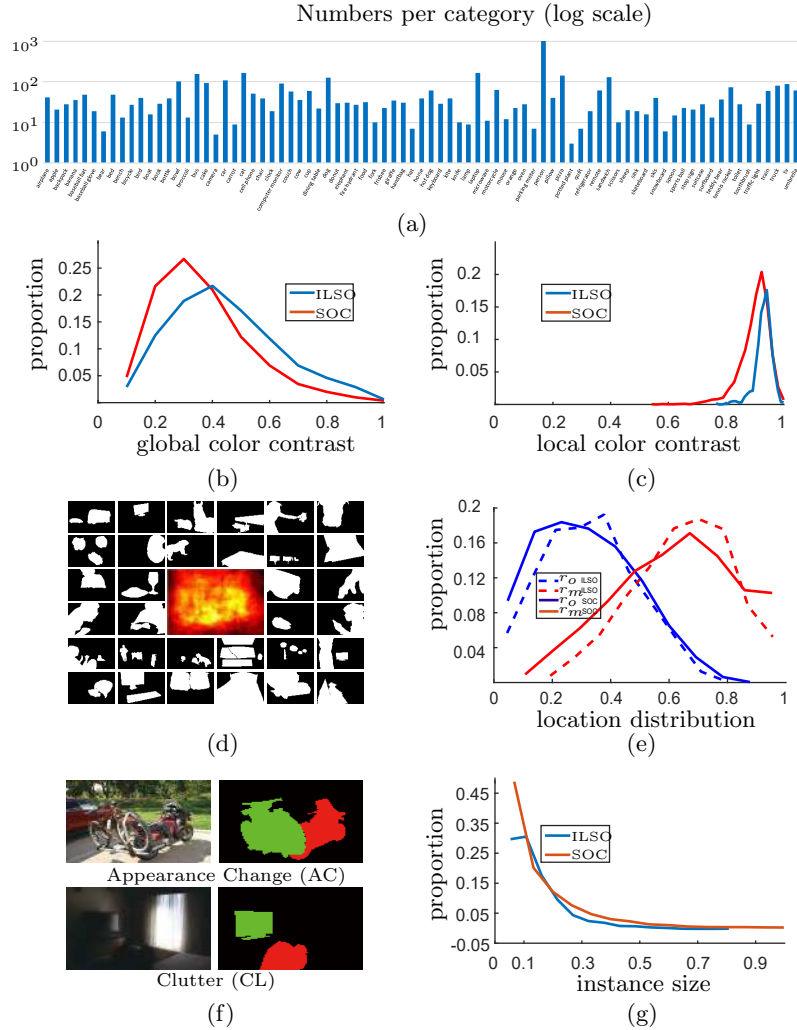


Fig. 4. (a) Number of annotated instances per category in our SOC dataset. (b, c) The statistics of global color contrast and local color contrast, respectively. (d) A set of saliency maps from our dataset and their overlay map. (e) Location distribution of the salient objects in SOC. (f) Attribute visual examples. (g) The distribution of instance sizes for the SOC and ILSO [22].

datasets (*e.g.*, MSRA10K, with 10,000 images). Besides the scale, dataset quality is also an important factor. To obtain a large amount of high quality images, we randomly select images from the MSCOCO dataset [27], which is a large-scale real-world dataset whose objects are labeled with polygons (*i.e.*, coarse labeling). High-quality labels also play a critical role in improving the accuracy of SOD models [1]. Toward this end, we relabel the dataset with pixel-wise annotations. Similar to famous SOD task oriented benchmark dataset-




Fig. 5. Compared with the recent new (a) instance-level ILSO dataset [22] which is labeled with discontinuous coarse boundaries, (c) MSCOCO dataset [27] which is labeled with polygons, our (b, d) SOC dataset is labeled with smooth fine boundaries.

Table 2. The list of salient object image attributes and the corresponding description. By observing the characteristics of the existing datasets, we summarize these attributes. Some visual examples can be found in Fig. 1 and Fig. 4 (f). For more examples, please refer to the supplementary materials

Attr	Description
AC	<i>Appearance Change.</i> The obvious illumination change in the object region.
BO	<i>Big Object.</i> The ratio between the object area and the image area is larger than 0.5.
CL	<i>Clutter.</i> The foreground and background regions around the object have similar color. We labeled images that their global color contrast value is larger than 0.2, local color contrast value is smaller than 0.9 with clutter images (see Sec. 3).
HO	<i>Heterogeneous Object.</i> Objects composed of visually distinctive/dissimilar parts.
MB	<i>Motion Blur.</i> Objects have fuzzy boundaries due to shake of the camera or motion.
OC	<i>Occlusion.</i> Objects are partially or fully occluded.
OV	<i>Out-of-View.</i> Part of object is clipped by image boundaries.
SC	<i>Shape Complexity.</i> Objects have complex boundaries such as thin parts (e.g., the foot of animal) and holes.
SO	<i>Small Object.</i> The ratio between the object area and the image area is smaller than 0.1.

s [1, 2, 11, 19, 22, 23, 29, 32, 37, 41, 43], we did not use the eye tracker device. We have taken a number of steps to provide the high-quality of the annotations. These steps include two stages: **In the bounding boxes (bboxes) stage**, (i) we ask 5 viewers to annotate objects with bboxes that they think are salient in each image. (ii) keep the images which majority (≥ 3) viewers annotated the same (the IOU of the bbox > 0.8) object. After the first stage, we have 3,000 salient object images annotated with bboxes. **In the second stage**, we further manually label the accurate silhouettes of the salient objects according to the bboxes. Note that we have 10 volunteers involved in the whole steps for cross-check the quality of annotations. In the end, we keep 3,000 images with high-quality, instance-level labeled salient objects. As shown in Fig. 5 (b,d), the boundaries of our object labels are precise, sharp and smooth. During the annotation process, we also add some new categories (e.g., *computer monitor, hat, pillow*) that are not labeled in the MSCOCO dataset [27].

7) Salient Objects with Attributes. Having attributes information regarding the images in a dataset helps objectively assess the performance of models over different types of parameters and variations. It also allows the inspection




Fig. 6. Left: Attributes distribution over the salient object images in our SOC dataset. Each number in the grids indicates the image number of occurrences. Right: The *dominant dependencies* among attributes base on the frequency of occurrences. Larger width of a link indicates higher probability of an attribute to other ones.

of model failures. To this end, we define a set of attributes to represent specific situations faced in the real-world scenes such as *motion blur*, *occlusion* and *cluttered background* (summarized in Table 2). Note that one image can be annotated with multiple attributes as these attributes are not exclusive.

Inspired by [33], we present the distribution of attributes over the dataset as shown in Fig. 6 Left. Type *SO* has the largest proportion due to accurate instance-level (*e.g.*, tennis racket in Fig. 2) annotation. Type *HO* accounts for a large proportion, because the real-world scenes are composed of different constituent materials. *Motion blur* is more common in video frames than still images, but it also occurs in still images sometimes. Thus, type *MB* takes a relatively small proportion in our dataset. Since a realistic image usually contains multiple attributes, we show the dominant dependencies among attributes based on the frequency of occurrences in the Fig. 6 Right. For example, a scene containing lots of heterogeneous objects is likely to have a large number of objects blocking each other and forming complex spatial structures. Thus, type *HO* has a strong dependency with type *OC*, *OV*, and *SO*.

4 Benchmarking Models

In this section, we present the evaluation results of the sixteen SOD models on our SOC dataset. Nearly all representative CNNs based SOD models are evaluated. However, since the codes of some models are not publicly available, we do not consider them here. In addition, most models are not optimized for non-salient objects detection. Thus, to be fair, we only use the test set of our SOC dataset to evaluate SOD models. We describe the evaluation metrics in Sec. 4.1. Overall model performance on SOC dataset is presented in Sec. 4.2 and summarized in Table 3, while the attribute level performance (*e.g.*, performance of the appearance changes) is discussed in Sec. 4.3 and summarized in Table 4. The evaluation scripts are publicly available, and on-line evaluation test is provided on our website.

4.1 Evaluation Metrics

In a supervised evaluation framework, given a predicted map M generated by a SOD model and a ground truth mask G , the evaluation metrics are expected to tell which model generates the best result. Here, we use three different evaluation metrics to evaluate SOD models on our SOC dataset.

Pixel-wise Accuracy ε . The region similarity evaluation measure does not consider the true negative saliency assignments. As a remedy, we also compute the normalized $([0,1])$ mean absolute error (MAE) between M and G , defined as:

$$\varepsilon = \frac{1}{W \times H} \sum_{x=1}^W \sum_{y=1}^H \|M(x, y) - G(x, y)\|, \quad (1)$$

where W and H are the width and height of images, respectively.

Region Similarity F . To measure how well the regions of the two maps match, we use the *F-measure*, defined as:

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

where $\beta^2 = 0.3$ is suggested by [1] to trade-off the *recall* and *precision*. However, the black (all-zero matrix) ground truth is not well defined in *F-measure* when calculating *recall* and *precision*. Under this circumstances, different foreground maps get the same result 0, which is apparently unreasonable. Thus, *F-measure* is not suitable for measuring the results of non-salient object detection.

However, both metrics of ε and F are based on pixel-wise errors and often ignore the structural similarities. Behavioral vision studies have shown that the human visual system is highly sensitive to structures in scenes [13]. In many applications, it is desired that the results of the SOD model retain the structure of objects.

Structure Similarity S . *S-measure* proposed by Fan *et al.* [13] evaluates the structural similarity, by considering both regions and objects. Therefore, we additionally use *S-measure* to evaluate the structural similarity between M and G . Note that the next overall performance we evaluated and analyzed are based on the *S-measure*.

4.2 Metric Statistics

To obtain an overall result, we average the scores of the evaluation metrics η ($\eta \in \{F, \varepsilon, S\}$), denoted by:

$$M_\eta(D) = \frac{1}{|D|} \sum_{I \in D} \bar{\eta}(I_i), \quad (3)$$

where $\bar{\eta}(I_i)$ is the evaluation score of the image I_i within the image dataset D .

Table 3. The performance of SOD models under three metrics. F stands for region similarity, ε is the mean absolute error, and S is the structure similarity. \uparrow stand for the higher the number the better, and vice versa for \downarrow . The evaluation results are calculated according to Eqn. (3) over our SOC dataset. $S_{all}, F_{all}, \varepsilon_{all}$ indicate the overall performance using the metric of S, F, ε , respectively. **Bold** for the best.

	Single-task												Multi-task			
Type	LEGS [36]	MC [51]	MDF [23]	DCL [24]	AMU [24]	RFCN [38]	DHS [28]	ELD [15]	DISC [8]	IMC [48]	UCF [50]	DSS [17]	NLDF [31]	DS [25]	WSS [37]	MSR [22]
$F_{all} \uparrow$.276	.291	.307	.339	.341	.435	.360	.317	.288	.352	.333	.341	.352	.347	.327	.380
$S_{all} \uparrow$.677	.757	.736	.771	.737	.814	.804	.776	.737	.664	.657	.807	.818	.779	.785	.819
$\varepsilon_{all} \downarrow$.230	.138	.150	.157	.185	.113	.118	.135	.173	.269	.282	.111	.104	.155	.133	.113

Single-task: For the single-task models, the best performing model on the entire SOC dataset (S_{all} in Table 3) is NLDF [31] ($M_S = 0.818$), followed by RFCN [38] ($M_S = 0.814$). MDF [23] and AMU [49] use edge cues to promote the saliency map but fail to achieve the ideal goal. Aiming at using the local region information of images, MC [51], MDF [23], ELD [15], and DISC [8] try to use superpixel methods to segment images into regions and then extract features from these regions, which is complex and time-consuming. To further improve the performance, UCF [50], DSS [17], NLDF [31], and AMU [49] utilize the FCN to improve the performance of SOD (S_{sal} in Table 4). Some other methods such as DCL [24] and IMC [48] try to combine superpixels with FCN to build a powerful model. Furthermore, RFCN [38] combines two related cues including edges and superpixels into FCN to obtain the good performance ($M_F = 0.435$, $M_S = 0.814$) over the overall dataset.

Multi-task: Different from models mentioned above, MSR [22] detects the instance-level salient objects using three closely related steps: estimating saliency maps, detecting salient object contours, and identifying salient object instances. It creates a multi-scale saliency refinement network that results in the highest performance (S_{all}). Other two multi-task models DS [25] and WSS [37] utilize the segmentation and classification results simultaneously to generate the saliency maps, obtaining a moderate performance. It is worth mentioning that although WSS is a weakly supervised multi-task model, it still achieves comparable performance to other single-task, fully supervised models. So, the weakly-supervised and multi-task based models can be promising future directions.

4.3 Attributes-based Evaluation

We assign the salient images with attributes as discussed in Sec. 3 and Table 2. Each attribute stands for a challenging problem faced in the real-world scenes. The attributes allow us to identify groups of images with a dominant feature (*e.g.*, presence of clutter), which is crucial to illustrate the performance of SOD models and to relate SOD to application-oriented tasks. For example, sketch2photo application [7] prefers models with good performance on big objects, which can be identified by attributes-based performance evaluation methods.

Table 4. Attributes-based performance on our SOC salient objects sub-dataset. For each model, the score corresponds to the average structure similarity M_S (in Sec. 4.1) over all datasets with that specific attribute (*e.g.*, CL). The higher the score the better the performance. **Bold** for the best. The average salient-object performance S_{sal} is presented in the first row using the structure similarity S . The symbol of $^+$ and $^-$ indicates *increase* and *decrease* compared to the average (S_{sal}) result, respectively

Attr	Single-task										Multi-task					
	LEGS [36]	MC [51]	MDF [23]	DCL [24]	AMU [24]	RFCN [38]	DHS [28]	ELD [15]	DISC [8]	IMC [48]	UCF [50]	DSS [17]	NLDF [31]	DS [25]	WSS [37]	MSR [22]
S_{sal}	.607	.619	.610	.705	.705	.709	.728	.664	.629	.679	.678	.698	.714	.719	.676	.748
<i>AC</i>	.625	.631	.614	.734	.736	.744	.745	.673	.644	.702	.714	.726	.737	.764	.691	.789
<i>BO</i>	.509	.490	.461 $^-$.610	.569	.540	.590	.576	.517	.701$^+$.636	.496 $^-$.568	.685	.566	.667
<i>CL</i>	.620	.635	.566	.699	.708	.714	.743	.658	.635	.696	.704	.677 $^-$.713	.729	.678	.756
<i>HO</i>	.666	.666	.648	.745	.755	.759	.766	.706	.681	.715	.744	.748	.755	.756	.707	.777
<i>MB</i>	.543 $^-$.603	.615	.693	.706	.715	.722	.639	.600	.689	.682	.695	.685	.711	.641	.757
<i>OC</i>	.609	.617	.608	.708 $^+$.725$^+$.711	.716	.658	.630	.672	.701 $^+$.689	.709	.725 $^+$.672	.740
<i>OV</i>	.548	.584	.568	.699	.708$^+$.687	.706	.637	.573	.693 $^+$.685 $^+$.665	.688	.722 $^+$.624	.743
<i>SC</i>	.608	.620	.669 $^+$.738	.731	.735	.763	.688	.653	.690	.722 $^+$.746 $^+$.745	.724	.677	.773
<i>SO</i>	.573 $^-$.601	.621	.691	.685	.698	.713	.644	.614	.648 $^-$.650	.696 $^-$.703	.696	.659	.730

Results. In Table 4, we show the performance on subsets of our dataset characterized by a particular attribute. Due to space limitation, in the following parts, we only select some representative attributes for further analysis. More details can be found in the supplementary material.

Big Object (BO) scenes often occur when objects are in a close distance with the camera, in which circumstances the tiny text or patterns would always be seen clearly. In this case, the models which prefer to focus on local information will be misled seriously, leading to a considerable (*e.g.*, 28.9% loss for DSS [17], 20.8% loss for MC [51] and 23.8% loss for RFCN [38]) loss of performance.

However, the performance of IMC [48] model goes up for a slight margin of 3.2% instead. After taking a deeper look of the pipeline of this model, we came up a reasonable explanation. IMC uses a coarse predicted map to express semantics and utilizes over-segmented images to supplement the structural information, achieving a satisfying result on type *BO*. However, over-segmented images cannot make up the missing details, causing 4.6% degradation of performance on the type of *SO*.

Small Object (SO) is tricky for all SOD models. All models encounter performance degradation (*e.g.*, from DSS [17] -0.3% to LEGS [36] -5.6%), because *SOs* are easily ignored during down-sampling of CNNs. DSS [17] is the only model that has a slight decrease of performance on type *SO*, while it has the biggest (28.9%) loss of performance on type *BO*. MDF [23] uses multi-scale superpixels as the input of network, so it retains the details of small objects well. However, due to the limited size of superpixels, MDF can not efficiently sense the global semantics, causing a big failure on type *BO*.

Occlusions (OC) scenes in which objects are partly obscured. Thus, it requires SOD models to capture global semantics to make up for the incomplete

information of objects. To do so, DS [25] & AMU [49] made use of the multi-scale features in the down-sample progress to generate a fused saliency map; UCF [50] proposed an uncertain learning mechanism to learn uncertain convolutional features. All these methods try to get saliency maps containing both global and local features. Unsurprisingly, these methods have achieved pretty good results on type *OC*. Based on the above analyses, we also find that these three models perform very well on the scenes requiring more semantic information like type *AC*, *OV* and *CL*.

Heterogeneous Object (*HO*) is a common attribute in nature scenes. The performance of different models on type *HO* gets some improvement to their average performances respectively, all fluctuating from 3.9% to 9.7%. We suspect this is because type *HO* accounts for a significant proportion of all datasets, objectively making models more fitting to this attribute. This result in some degree confirms our statistics in Fig. 6.

5 Discussion and Conclusion

To our best knowledge, this work presents the currently largest scale performance evaluation of CNNs based salient object detection models. Our analysis points out a serious *data selection bias* in existing SOD datasets. This design bias has lead to state-of-the-art SOD algorithms almost achieve saturated high performance when evaluated on existing datasets, but are still far from being satisfactory when applied to real-world daily scenes. Based on our analysis, we first identify 7 important aspects that a comprehensive and balanced dataset should fulfill. We firstly introduces a high quality SOD dataset, **SOC**. It contains salient objects from daily life in their natural environments which reaches closer to realistic settings. The SOC dataset will evolve and grow over time and will enable research possibilities in multiple directions, *e.g.*, salient object subitizing [46], instance level salient object detection [22], weakly supervised based salient object detection [37], *etc.* Then, a set of attributes (*e.g.*, *Appearance Change*) is proposed in the attempt to obtain a deeper insight into the SOD problem, investigate the pros and cons of the SOD algorithms, and objectively assess the model performances over different perspectives/requirements. Finally, we report attribute-based performance assessment on our SOC dataset. The results open up promising future directions for model development and comparison.

Acknowledgments

This research was supported by NSFC (NO. 61620106008, 61572264), the national youth talent support program, Tianjin Natural Science Foundation for Distinguished Young Scholars (NO. 17JCJQJC43700), Huawei Innovation Research Program.

References

1. Achanta, R., Hemami, S., Estrada, F., Susstrunk, S.: Frequency-tuned salient region detection. In: CVPR. pp. 1597–1604. IEEE (2009)
2. Alpert, S., Galun, M., Basri, R., Brandt, A.: Image Segmentation by Probabilistic Bottom-Up Aggregation and Cue Integration. In: CVPR. pp. 1–8. IEEE (2007)
3. Borji, A., Cheng, M.M., Jiang, H., Li, J.: Salient object detection: A benchmark. IEEE TIP **24**(12), 5706–5722 (2015)
4. Borji, A., Itti, L.: State-of-the-art in visual attention modeling. IEEE TPAMI **35**(1), 185–207 (2013)
5. Borji, A., Sihite, D.N., Itti, L.: Salient object detection: a benchmark. In: ECCV. pp. 414–429. Springer (2012)
6. Caesar, H., Uijlings, J., Ferrari, V.: COCO-Stuff: Thing and Stuff Classes in Context. In: CVPR. IEEE (2018)
7. Chen, T., Cheng, M.M., Tan, P., Shamir, A., Hu, S.M.: Sketch2photo: Internet image montage. ACM Transactions on Graphics (TOG) **28**(5), 124 (2009)
8. Chen, T., Lin, L., Liu, L., Luo, X., Li, X.: DISC: Deep image saliency computing via progressive representation learning. IEEE transactions on neural networks and learning systems **27**(6), 1135–1149 (2016)
9. Cheng, M.M., Hou, Q.B., Zhang, S.H., Rosin, P.L.: Intelligent Visual Media Processing: When Graphics Meets Vision. Journal of Computer Science and Technology **32**(1), 110–121 (2017)
10. Cheng, M.M., Mitra, N.J., Huang, X., Hu, S.M.: Salienshape: group saliency in image collections. The Visual Computer **30**(4), 443–453 (2014)
11. Cheng, M.M., Mitra, N.J., Huang, X., Torr, P.H.S., Hu, S.M.: Global contrast based salient region detection. IEEE TPAMI **37**(3), 569–582 (2015)
12. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2010 (VOC2010) Results
13. Fan, D.P., Cheng, M.M., Liu, Y., Li, T., Borji, A.: Structure-measure: A New Way to Evaluate Foreground Maps. In: ICCV. pp. 4548–4557. IEEE (2017)
14. Fan, D.P., Gong, C., Cao, Y., Ren, B., Cheng, M.M., Borji, A.: Enhanced-alignment Measure for Binary Foreground Map Evaluation. In: International Joint Conference on Artificial Intelligence (IJCAI). pp. 698–704 (2018)
15. Gayoung, L., Yu-Wing, T., Junmo, K.: Deep Saliency with Encoded Low level Distance Map and High Level Features. In: CVPR. IEEE (2016)
16. He, J., Feng, J., Liu, X., Cheng, T., Lin, T.H., Chung, H., Chang, S.F.: Mobile product search with bag of hash bits and boundary reranking. In: CVPR. pp. 3005–3012. IEEE (2012)
17. Hou, Q., Cheng, M.M., Hu, X., Borji, A., Tu, Z., Torr, P.: Deeply supervised salient object detection with short connections. IEEE TPAMI (2018)
18. Hou, Q., Dokania, P.K., Massiceti, D., Wei, Y., Cheng, M.M., Torr, P.H.S.: Bottom-Up Top-Down Cues for Weakly Supervised Semantic Segmentation. In: EMM-CVPR. IEEE (2017)
19. Jiang, H., Cheng, M.M., Li, S.J., Borji, A., Wang, J.: Joint Salient Object Detection and Existence Prediction. Front. Comput. Sci (2018)
20. Judd, T., Durand, F., Torralba, A.: A benchmark of computational models of saliency to predict human fixations. In: MIT Technical Report (2012)
21. Lazebnik, S., Schmid, C., Ponce, J.: A sparse texture representation using local affine regions. IEEE TPAMI **27**(8), 1265–1278 (2005)

22. Li, G., Xie, Y., Lin, L., Yu, Y.: Instance-Level Salient Object Segmentation. In: CVPR. pp. 247–256. IEEE (2017)
23. Li, G., Yu, Y.: Visual saliency based on multiscale deep features. In: CVPR. pp. 5455–5463. IEEE (2015)
24. Li, G., Yu, Y.: Deep Contrast Learning for Salient Object Detection. In: CVPR. pp. 478–487. IEEE (2016)
25. Li, X., Zhao, L., Wei, L., Yang, M.H., Wu, F., Zhuang, Y., Ling, H., Wang, J.: DeepSaliency: Multi-task deep neural network model for salient object detection. IEEE TIP **25**(8), 3919–3930 (2016)
26. Li, Y., Hou, X., Koch, C., Rehg, J.M., Yuille, A.L.: The secrets of salient object segmentation. In: CVPR. pp. 280–287. IEEE (2014)
27. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: ECCV. pp. 740–755. Springer (2014)
28. Liu, N., Han, J.: DHSNet: Deep Hierarchical Saliency Network for Salient Object Detection. In: CVPR. pp. 678–686. IEEE (2016)
29. Liu, T., Sun, J., Zheng, N., Tang, X., Shum, H.Y.: Learning to Detect A Salient Object. In: CVPR. pp. 1–8. IEEE (2007)
30. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. pp. 3431–3440. IEEE (2015)
31. Luo, Z., Mishra, A.K., Achkar, A., Eichel, J.A., Li, S., Jodoin, P.M.: Non-local deep features for salient object detection. In: CVPR. vol. 2, p. 7 (2017)
32. 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. vol. 2, pp. 416–423. IEEE (2001)
33. Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: CVPR. pp. 724–732. IEEE (2016)
34. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. IJCV **115**(3), 211–252 (2015)
35. Wang, J., Jiang, H., Yuan, Z., Cheng, M.M., Hu, X., Zheng, N.: Salient Object Detection: A Discriminative Regional Feature Integration Approach. IJCV **123**(2), 251–268 (2017)
36. Wang, L., Lu, H., Ruan, X., Yang, M.H.: Deep networks for saliency detection via local estimation and global search. In: CVPR. pp. 3183–3192. IEEE (2015)
37. Wang, L., Lu, H., Wang, Y., Feng, M., Wang, D., Yin, B., Ruan, X.: Learning to detect salient objects with image-level supervision. In: CVPR. pp. 136–145. IEEE (2017)
38. Wang, L., Wang, L., Lu, H., Zhang, P., Ruan, X.: Saliency detection with recurrent fully convolutional networks. In: ECCV. pp. 825–841. Springer (2016)
39. Wei, Y., Feng, J., Liang, X., Cheng, M.M., Zhao, Y., Yan, S.: Object Region Mining with Adversarial Erasing: A Simple Classification to Semantic Segmentation Approach. In: CVPR. IEEE (2017)
40. Wei, Y., Liang, X., Chen, Y., Shen, X., Cheng, M.M., Feng, J., Zhao, Y., Yan, S.: STC: A simple to complex framework for weakly-supervised semantic segmentation. IEEE TPAMI (2017)
41. Xia, C., Li, J., Chen, X., Zheng, A., Zhang, Y.: What is and what is not a salient object? learning salient object detector by ensembling linear exemplar regressors. In: CVPR. IEEE (2017)

42. Xie, S., Tu, Z.: Holistically-nested edge detection. In: ICCV. pp. 1395–1403. IEEE (2015)
43. Yan, Q., Xu, L., Shi, J., Jia, J.: Hierarchical saliency detection. In: CVPR. pp. 1155–1162. IEEE (2013)
44. Yang, C., Zhang, L., Lu, H., Ruan, X., Yang, M.H.: Saliency detection via graph-based manifold ranking. In: CVPR. pp. 3166–3173. IEEE (2013)
45. Zhang, G.X., Cheng, M.M., Hu, S.M., Martin, R.R.: A Shape-Preserving Approach to Image Resizing. *Computer Graphics Forum* **28**(7), 1897–1906 (2009)
46. Zhang, J., Ma, S., Sameki, M., Sclaroff, S., Betke, M., Lin, Z., Shen, X., Price, B., Mech, R.: Salient object subitizing. In: CVPR. pp. 4045–4054. IEEE (2015)
47. Zhang, J., Sclaroff, S., Lin, Z., Shen, X., Price, B., Mech, R.: Unconstrained salient object detection via proposal subset optimization. In: CVPR. pp. 5733–5742. IEEE (2016)
48. Zhang, J., Dai, Y., Porikli, F.: Deep Salient Object Detection by Integrating Multi-level Cues. In: Winter Conference on Applications of Computer Vision (WACV). pp. 1–10. IEEE (2017)
49. 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)
50. Zhang, P., Wang, D., Lu, H., Wang, H., Yin, B.: Learning Uncertain Convolutional Features for Accurate Saliency Detection. In: ICCV. pp. 212–221 (2017)
51. Zhao, R., Ouyang, W., Li, H., Wang, X.: Saliency detection by multi-context deep learning. In: CVPR. pp. 1265–1274. IEEE (2015)