

Structure-measure: A New Way to Evaluate Foreground Maps

Ming-Ming Cheng · Deng-Ping Fan

Received: date / Accepted: date

Abstract Foreground map evaluation is crucial for gauging the progress of object segmentation algorithms, in particular in the field of salient object detection where the purpose is to accurately detect and segment the most salient object in a scene. Several measures (*e.g.*, Area-Under-the-Curve, F1-measure, Average Precision, *etc.*) have been used to evaluate the similarity between a foreground map and a ground-truth map. The existing measures are based on pixel-wise errors and often ignore the structural similarities. Behavioral vision studies, however, have shown that the human visual system is highly sensitive to structures in scenes. Here, we propose a novel, efficient (0.005s per image), and easy to calculate measure known as **S-measure** (structural measure) to evaluate foreground maps. Our new measure simultaneously evaluates region-aware and object-aware structural similarity between a foreground map and a ground-truth map. We demonstrate superiority of our measure over existing ones using 4 meta-measures on 5 widely-used benchmark datasets. Furthermore, we conduct a behavioral judgment study over a new database. Data from 45 subjects shows that on average they preferred the saliency maps chosen by our measure over the saliency maps chosen by the state-of-the-art measures. Our experimental results offer new insights into foreground map evaluation where current measures fail to truly examine the strengths and weaknesses of models. Code: <https://github.com/DengPingFan/S-measure>

Keywords S-measure · Structure measure · Foreground Maps · Evaluation · Salient Object Detection

Ming-Ming Cheng and Deng-Ping Fan are with the College of Computer Science, Nankai University, China. (Email: cmm@nankai.edu.cn; dengpfan@gmail.com)
A preliminary version of this work appeared at ICCV2017 [18].

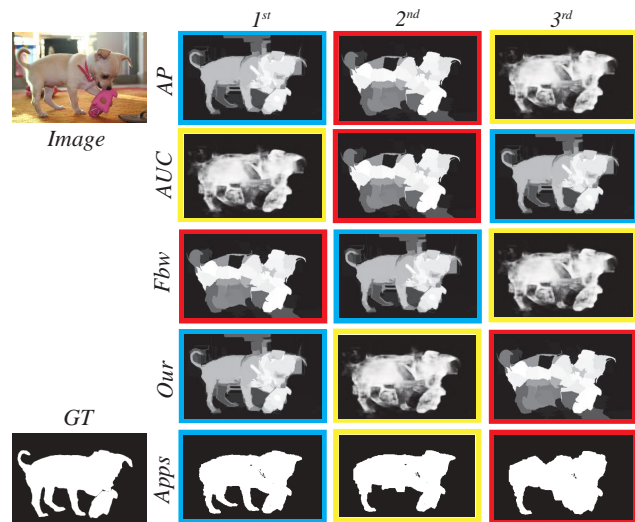


Figure 1 Inaccuracy of existing evaluation measures. We compare the ranking of saliency maps generated by three state-of-the-art (SOTA) salient object detection algorithms: DISC [14], MDF [43], and MC [83]. According to the segmentation application’s ranking (last row; see details in Apps-Sec. 5), the blue-border map ranks first, followed by the yellow- and red-border maps. The blue-border map captures the dog’s structure most accurately, with respect to the GT. The yellow-border map looks fuzzy although the overall outline of the dog is still present. The red-border map almost completely destroyed the structure of the dog. Surprisingly, all of the measures based on pixel-wise errors (first three rows) fail to rank the maps correctly. Our new measure (4^{th} row) ranks the three maps in the right order.

1 Introduction

The evaluation of a predicted foreground map (FM) against a ground-truth (GT) annotation map is crucial in evaluating and comparing various computer vision algorithm for applications such as **AR COPY**

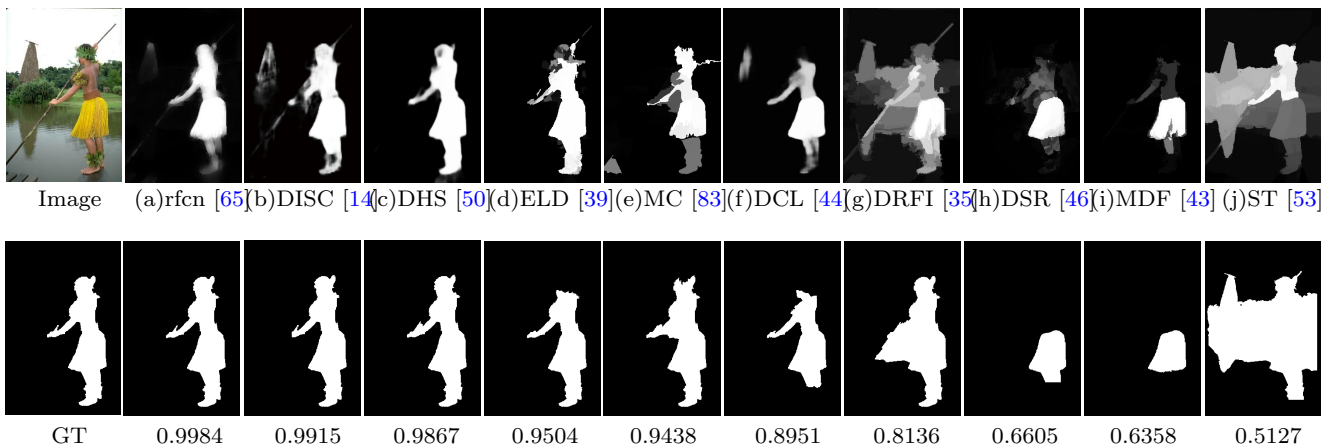


Figure 2 S-measure score ($\lambda = 0.5, K = 4$) for the outputs of SalCut [16] algorithm (2^{nd} row) when fed with inputs of 10 saliency maps (1^{st} row). The ranking results clearly indicate that our measure is good at capturing object structures and can provide a reliable ranking.

& PASTE¹ [62], object detection and recognition [6, 8, 31, 37, 63], video summarization [27], video compression [30, 32], image segmentation [69], content-based image retrieval [45, 48, 49], visual tracking, photo synthesis [13], image-text matching [87], image collection browsing [13], *etc.* As a specific example, here we focus on salient object detection [5–7, 10, 12, 29, 42, 51, 64, 66, 70, 75, 76, 78, 81], although the proposed measure is general and can be used for other purposes. It is worth noting that the salient object does not necessarily need to be the foreground object [9, 25].

The GT map is often the average of several manual annotations. Thus, it can be binary or non-binary. Similarly, the predicted foreground maps are either binary or non-binary. As a result, evaluation measures can be classified into two types:

1. **Binary map evaluation:** Common measures here include F_β -measure [3, 16, 52] and PASCAL’s VOC segmentation measure [17].
2. **Non-binary map evaluation:** Three traditional and popular measures here include Area Under the Curve (AUC), Precision-Recall (PR) curve, and Average Precision (AP) [17]. A newly released measure known as weighted F_β -measure (Fbw) [55] has been proposed to remedy flaws of AUC, PR and AP.

It is often desired to have the foreground map contain the entire structure of an object. Thus, evaluation measures are expected to tell which model generates a more complete object. For example, in Fig. 1 (first row) the blue-border map better captures the dog than the red-border map. In the latter case, shape of the the dog is drastically degraded to a degree that it is difficult to guess the object category from its segmentation map.

Surprisingly, all of the current evaluation measures fail to correctly rank these maps (in terms of preserving the structure).

We employed 10 state-of-the-art (SOTA) saliency detection algorithms to obtain 10 saliency maps (Fig. 2; 1^{st} row) and then fed these maps to the SalCut² [16] algorithm to generate corresponding binary maps (2^{nd} row). Finally, we used the proposed **S-measure** to rank these maps. A lower score for our measure corresponds to more destruction in the global structure of the man (columns e to j). This experiment clearly shows that our new measure emphasizes the entire structure of the object. In these ten binary maps (2^{nd} row), there are six maps with score below 0.95, *i.e.* with percentage 60%. Using the same threshold (0.95), we found that the proportions of destroyed images in four popular saliency datasets (*i.e.*, ECSSD [68], HKU-IS [43], PASCAL-S [47], and SOD [56]) are 66.80%, 67.30%, 81.82% and 83.03%, respectively. Using the F_β -measure to evaluate the binary maps, these proportions are 63.76%, 65.43%, 78.32% and 82.67%, respectively. This means that our measure is more stringent than the F_β -measure on object structure.

To remedy the problem of existing measures (*i.e.*, low sensitivity to entire object structure), we present a structure-sensitive similarity measure based on two observations:

1. **Region perspectives:** Although it is difficult to describe the structure of a foreground map, we notice that the entire structure of an object can be well illustrated by combining structures of constituent object parts (regions).
2. **Object perspectives:** In high-quality foreground maps, the foreground regions contrast sharply with the

¹ <https://clipdrop.co/>

² <https://github.com/MCG-NKU/SalBenchmark/blob/master/Code/CmLib/Saliency/CmSalCut.cpp>

background regions. These regions usually have approximately uniform contrast distributions.

Consequently, the proposed structure measure consists of two parts, including a *region-aware structural similarity measure* and an *object-aware structural similarity measure*. The region-aware measure tries to capture the global object structure by combining the structural information of all the object parts. The structural similarity of regions has been well explored in the image quality assessment (IQA) community [67]. The object-aware similarity measure tries to compare global distributions of foreground and background regions in the foreground map and the GT map.

Our measure is compared against various existing measures including AP, AUC, PASCAL, Fbw, F_β -measure on several widely-used salient object detection benchmarks including ASD [1], SOD [56], ECSSD [68], PASCAL-S [47], and HKU-IS [43]. Extensive empirical investigations show that Structure-measure not only provides a reliable evaluation but also achieves significantly improved performance than current measures.

This work is an extension of our previous ICCV2017 version [18]. The major differences between these two versions include: (1) We extend the preliminary version to binary foreground map evaluation and provide a unified evaluation applicable to both binary and non-binary foreground maps. Our work offers new insights into foreground maps evaluation where current measures fail to examine the strengths and weaknesses of models fully. (2) We provide a set of new experiment to validate the efficiency, robustness, and extensibility of our proposed measure. These extension focus on non-binary maps' evaluation. Besides, we also give more details about our application-ranking framework in Appendices. (3) We build the several representative online Benchmark and model zoo of saliency detection, which integrates various publicly available saliency datasets with uniform input/output formats (*i.e.*, JPEG for image; PNG for GT). (4) We also provide Python³ and Matlab⁴ version code for existing benchmarking work which benefits many related tasks and our computer vision community. (5) We have made a lot of efforts to improve the presentations and organizations of our paper. First, several new figures are added or re-produced to better illustrate the meta-measure and key results of this work. Second, we have added several new sections to describe more details about the flaws of current measures, provide more theoretical details about our S-measure, and more theoretical details include region-aware (Sec. 4.1) and object-aware (Sec. 4.2) structure similarity evaluation.

The new content will allow readers to better understand our approach.

2 Current Evaluation Measures

Foreground maps can be generated by various algorithms (*e.g.*, for saliency detection or object segmentation). Saliency detection algorithms often generate non-binary maps, whereas object segmentation algorithms usually generate binary maps. As a result, the foreground maps can be divided as *non-binary* maps with values in the range [0,1] or *binary* maps with values either 0 or 1. Each map value denotes the probability of a specific pixel belonging to the foreground [55, 59].

2.1 Evaluation of Binary Maps

To evaluate a binary map, four values are computed from the prediction confusion matrix: True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN). These values are then used to compute three ratios: True Positive Rate (TPR) or Recall, False Positive Rate (FPR), and Precision:

$$Recall = TPR = \frac{TP}{TP + FN} \quad (1)$$

$$FPR = \frac{FP}{TN + FP} \quad (2)$$

$$Precision = \frac{TP}{TP + FP} \quad (3)$$

The Precision and Recall are combined to compute the traditional F_β -measure.

$$\begin{aligned} F_\beta &= \frac{(1 + \beta^2)Precision \times Recall}{\beta^2 \times Precision + Recall} \\ &= \frac{(1 + \beta^2)TP}{(1 + \beta^2)TP + \beta^2FN + FP} \end{aligned} \quad (4)$$

where β is a parameter to balance the accuracy and the recall (typically $\beta = 1$ leading to harmonic mean). Another commonly used binary map evaluation metric is the PASCAL measure:

$$PASCAL = \frac{TP}{TP + FN + FP} \quad (5)$$

³ <https://github.com/DengPingFan/eval-co-sod>

⁴ <https://dpfan.net/One-KeyEvaluation/>

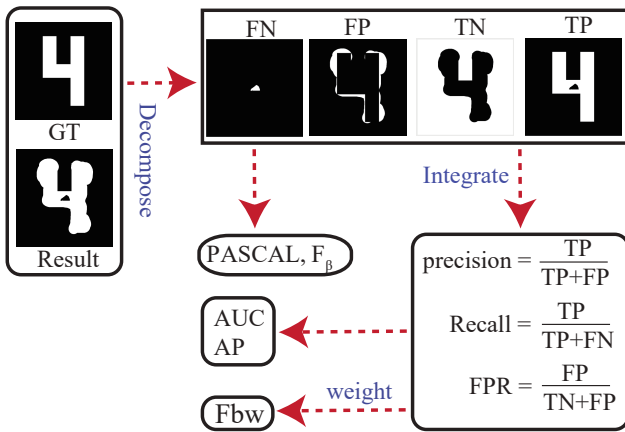


Figure 3 The current evaluation measures' framework. The AP, AUC and Fbw evaluation measures are computed in a similar way. They are all calculated in a pixel-wise manner and ignore the structural similarities.




Figure 4 A pixel-wise based evaluation example. In FM1, a black square falls inside the digit while in the FM2 it touches the boundary (1^{st} row). They are two binary maps (2^{nd} row) with the same TP, TN, FP and FN values. Visually, FM2 is favored over FM1 since FM1 destroys the foreground map's structure more drastically such that the the digit is hard to recognize. Current evaluation mean absolute error (MAE) measure is calculated in a pixel-wise manner and treat pixels independently. Hence, it ignores the structure of the foreground maps, thus it favor FM1 over FM2.

2.2 Evaluation of Non-Binary Maps

AUC and AP are two universally-agreed evaluation measures. Algorithms that produce non-binary maps apply three steps to evaluate the agreement between model predictions (non-binary maps) and human annotations (GT). First, multiple thresholds are applied to the non-binary map to get multiple binary maps. Second, these binary maps are compared to the GT to get a set of TPR (see Equ. (1)) & FPR (see Equ. (2)) values. These values are plotted in a 2D plot (a.k.a ROC curve), from which the AUC distills the area under the curve.

The AP measure is computed in a similar way. One can get a Precision (see Equ. (3)) & Recall (see Equ. (1)) curve by plotting Precision $p(r)$ as a function of Recall r . AP measure [17] is the average value of $p(r)$ over the evenly spaced x axis points from $r = 0$ to $r = 1$.

Recently, a measure called Fbw [55] has offered an intuitive generalization of the F_β measure. The authors of Fbw identified three causes of inaccurate evaluation of AP and AUC measures. To alleviate these flaws, they 1) extended the four basic quantities TP, TN, FP, and FN to non-binary values and, 2) assigned different weights (w) to different errors according to different location and neighborhood information.

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

While this measure improves upon other measures, sometimes it fails to correctly rank the foreground maps (see the 3rd row of the Fig. 1). In the next section, we will analyze why the current measures fail to rank these maps correctly.

3 Current Measures Are Pixel-Wise Based

Traditional measures (AP, AUC, PASCAL, F_β and Fbw) rely on two types of basic errors: FN, FP. Since these basic errors are calculated in a pixel-wise manner (see the Fig. 3), they cannot fully capture the structural information of foreground maps. However, foreground maps with fine structural details are often desired in several applications (*e.g.*, image retrieval, object detection and segmentation). Therefore, evaluation measures sensitive to foreground structures are favored. Unfortunately, the aforementioned measures fail to meet this expectation.

A contrived example is shown in Fig. 4 which contains two different types of foreground maps. In FM1, a black square falls inside the digit while in the FM2 it touches the boundary. In our opinion, FM1 is less favored than FM2 since it destroys the foreground map more drastically. However, the current pixel-level MAE (mean absolute error) measure favors FM1 over FM2. This seems to contradict our common sense.

A more realistic example is shown in Fig. 5. The blue-border map here better captures the pyramid than the red-border map, because the latter offers a fuzzy detection map that mostly highlights the top part of the pyramid while ignoring the rest. From an application standpoint (3^{rd} row, the output of the SalCut algorithm fed with saliency maps; ranked by our measure, *i.e.*, the 2^{nd} row), the blue-border map offers a complete shape of the pyramid. In practice, this situation is very common. Thus, if the evaluation measure cannot capture the

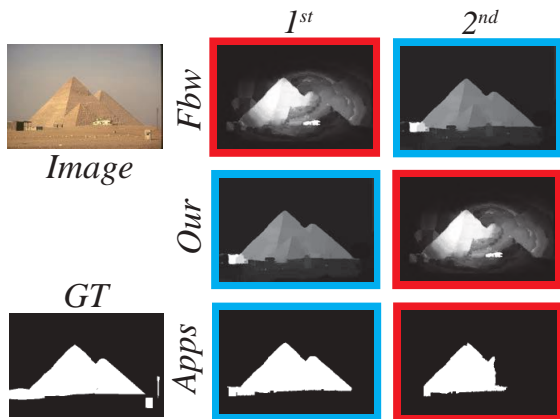


Figure 5 A pixel-wise based evaluation example. Two foreground maps are generated by two saliency detection algorithms DSR [46], and ST [53]. According to the application’s ranking and our user-study (Apps-Sec. 5; last row), the blue-border map does the best, followed by the red-border map. Since Fbw measure does not account for the structural similarity, it results in the complete reverse ranking. Our measure (2^{nd} row) correctly ranks the blue-border map as higher.

structural object information, it will not be able provide reliable information for model selection in applications.

4 Proposed Measure

In this section, we introduce our new measure to evaluate foreground maps. In the image quality assessment (IQA) field, a measure known as structural similarity measure (SSIM) [67] has been widely used to capture the structural similarity of the original image and a test image.

Let $x = \{x_i, i = 1, 2, \dots, N\}$ and $y = \{y_i, i = 1, 2, \dots, N\}$ be the FM and GT pixel values, respectively. The \bar{x} , \bar{y} , σ_x , σ_y are the mean and standard deviations of x and y , respectively. σ_{xy} is the covariance between the two. The SSIM is formulated as the product of three comparison terms including *luminance*, *contrast*, and *structure*:

$$ssim = \frac{2\bar{x}\bar{y} + C_1}{(\bar{x})^2 + (\bar{y})^2 + C_1} \cdot \frac{2\sigma_x\sigma_y + C_2}{\sigma_x^2 + \sigma_y^2 + C_2} \cdot \frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3} \quad (7)$$

where, the constants C_1, C_2 , and C_3 are set to very small values to avoid instability when denominator (*e.g.*, $\bar{x}^2 + \bar{y}^2$) is very close to zero in each component.

In Equ. (7), the first two terms denote the luminance comparison and contrast comparison, respectively. The closer the two (*i.e.*, \bar{x} and \bar{y} , or σ_x and σ_y), the closer the comparison (*i.e.*, luminance or contrast) to 1. The structures of the objects in an image are independent of the luminance that is affected by illumination and the reflectance. So the design of a structure comparison




Figure 6 The framework of our S-measure.

formula should be independent of luminance and contrast. SSIM [67] associates two unit vectors $(x - \bar{x})/\sigma_x$ and $(y - \bar{y})/\sigma_y$ with the structure of the two images. Since the correlation between these two vectors is equivalent to the correlation coefficient between x and y , the formula of structure comparison is denoted by the third term in Equ. (7).

To build salient object detection or object segmentation algorithms, researchers are often more concerned about the foreground object structures. Thus, our proposed structure measure combines both region-aware and object-aware structural similarities. The region-aware structural similarity performs similar to [67], which aims to capture “object-part” structure information without any special concern regarding how complete is the foreground. The object-aware structural similarity, on the other hand, is designed to mainly capture the structure information of “object-holistic” which focus on the complete object.

4.1 Region-Aware Structural Similarity Measure

Here we explain how to measure region-aware similarity. The region-aware similarity is designed to assess the object-part structure similarity in FM against the GT map. We first divide each of the FM and GT maps into four blocks using a horizontal and a vertical cut-off lines that intersect at the centroid of the GT foreground. Then, the subimages are divided recursively as in [38]. The total number of blocks is denoted as K . An example is shown in Fig. 6. The region-aware similarity $ssim(k)$ of each block is computed independently using Equ. (7). We assign a different weight (w_k) to each block proportional to the GT foreground region that this block covers. Thus, the region-aware structural similarity measure can be formulated as:

$$S_r = \sum_{k=1}^K w_k \times ssim(k). \quad (8)$$

Our investigation shows that the proposed S_r can well describe the object-part similarity between a FM

and a GT map. We also tried to directly use SSIM to assess the similarity between FM and GT at the image level or in the sliding window fashion as mentioned in [67]. These approaches fail to capture region structure similarities.

4.2 Object-Aware Structural Similarity Measure

Dividing the foreground map into blocks helps evaluate the object-part structural similarity. However, the region-aware measure (S_r) can not well account for the object similarity in a holistic way. For high-level vision tasks such as salient object detection, the evaluation of the object-level similarity is crucial. To achieve this goal, we propose a novel method to assess the foreground and background separately. Since, the GT maps usually have important characteristics including **sharp foreground-background contrast** and **uniform distribution**, the predicted FM is expected to possess these properties. This helps easily distinguish foreground from the background. We design our object-aware structural similarity measure with respect to these characteristics.

4.2.1 Sharp Foreground-Background Contrast

Our first observation is that the foreground region of the GT map usually contrasts sharply with the background region. We employ a formulation that is similar with the luminance component of SSIM, to measure how close the mean probability is between the foreground region of FM and the foreground region of GT. Let x_{FG} and y_{FG} represent the probability values of foreground region of FM and GT, respectively. \bar{x}_{FG} and \bar{y}_{FG} denote the means of x_{FG} and y_{FG} , respectively. The foreground comparison can be represented as:

$$O_{FG} = \frac{2\bar{x}_{FG}\bar{y}_{FG}}{(\bar{x}_{FG})^2 + (\bar{y}_{FG})^2}. \quad (9)$$

Equ. (9) has several appealing properties:

- Swapping the value of \bar{x}_{FG} and \bar{y}_{FG} , O_{FG} will not change the result,
- The range of O_{FG} is $[0,1]$,
- If and only if $\bar{x}_{FG} = \bar{y}_{FG}$, then $O_{FG} = 1$, and
- The closer the two maps, the closer the O_{FG} to 1 (the most important property).

These properties make Equ. (9) suitable for our purpose.

4.2.2 Uniform Distribution

Our second observation is that the foreground and background regions of the GT maps usually have uniform distributions. So, it is important to assign a higher score to a FM with the object being uniformly highlighted (*i.e.*, similar values across the entire object; see the Fig. 5). If the variability of the foreground values in the FM is high, then the distribution will not be uniform.

In probability theory and statistics, the coefficient of variation defined as the ratio of the standard deviation to the mean (σ_x/\bar{x}) is a standard measure of dispersion of a probability distribution. Here, we use it to represent the dispersion of the FM. In other words, the coefficient of variation is used to compute the dissimilarity between FM and GT distributions. According to Equ. (9), the total dissimilarity between FM and GT at object-level can be written as:

$$D_{FG} = \frac{(\bar{x}_{FG})^2 + (\bar{y}_{FG})^2}{2\bar{x}_{FG}\bar{y}_{FG}} + \lambda \times \frac{\sigma_{x_{FG}}}{\bar{x}_{FG}}, \quad (10)$$

where λ is a constant to balance the two terms. Then, the similarity between FM and GT at object level can be formulated as:

$$\begin{aligned} S'_{FG} &= \frac{1}{D_{FG}} \\ &= \frac{2\bar{x}_{FG}\bar{y}_{FG}}{(\bar{x}_{FG})^2 + (\bar{y}_{FG})^2 + 2\lambda \times \bar{y}_{FG} \times \sigma_{x_{FG}}}. \end{aligned} \quad (11)$$

Since in practice the mean probability of the GT foreground is exactly 1 ($\bar{y}_{FG} = 1$), the similarity between FM and GT in object level can be rewritten as:

$$S_{FG} = \frac{2\bar{x}_{FG}}{(\bar{x}_{FG})^2 + 1 + 2\lambda \times \sigma_{x_{FG}}}. \quad (12)$$

To compute background comparison S_{BG} , we regard the background as the complementary component of foreground by subtracting the FM and GT maps from 1 (change 1 to the maximum value of GT when GT is a non-binary map) as shown in Fig. 6. Then, S_{BG} can be similarly defined as:

$$S_{BG} = \frac{2\bar{x}_{BG}}{(\bar{x}_{BG})^2 + 1 + 2\lambda \times \sigma_{x_{BG}}}. \quad (13)$$

Let μ be the ratio of the foreground area in GT to the image area (*width* \times *height*). The final object-aware structural similarity measure can then be written as:

$$S_o = \mu \times S_{FG} + (1 - \mu) \times S_{BG}. \quad (14)$$

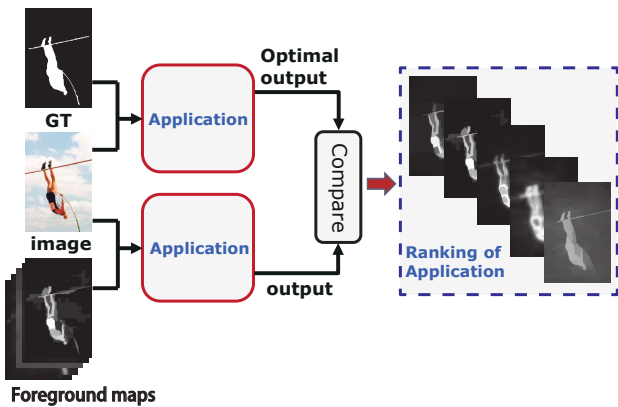


Figure 7 Meta-Measure 1: Application Ranking. To rank foreground maps according to an application, we compare the output obtained when using the GT, to the output when using the FM. The more similar a FM is to the GT map, the closer its application’s output should be to the GT output.

4.3 Structure Measure

Having region-aware and object-aware structural similarity evaluation definitions, we can formulate the final measure as,

$$S = \alpha \times S_o + (1 - \alpha) \times S_r, \quad (15)$$

where $\alpha \in [0, 1]$. We set $\alpha = 0.5$ in our implementation to assign equal contribution to both region similarity and object similarity. Using this measure to evaluate the three foreground maps in Fig. 1, we can correctly rank the maps consistent with the application rank and human rank.

5 Experiments

In order to assess the quality of our new measure, we utilized 4 meta-measures proposed by Margolin *et al.* [55] and 1 meta-measure (human judgments) proposed by us. These meta-measures are used to assess the quality of evaluation measures [61]. To conduct fair comparisons, the 4 meta-measures are computed on the ASD (a.k.a ASD1000) dataset [1]. The non-binary foreground maps (5000 maps in total) were generated using five saliency detection algorithms including CA [28], CB [34], RC [16], PCA [54], and SVO [11] (binary maps are achieved by feeding non-binary maps to the SalCut [16]).

Setting and Runtime. We assign $\lambda = 0.5$ and $K = 4$ in all experiments as [18]. We also test our measure on the ASD1000 dataset using a single CPU machine. The average run time for a single image is 0.0053 second.

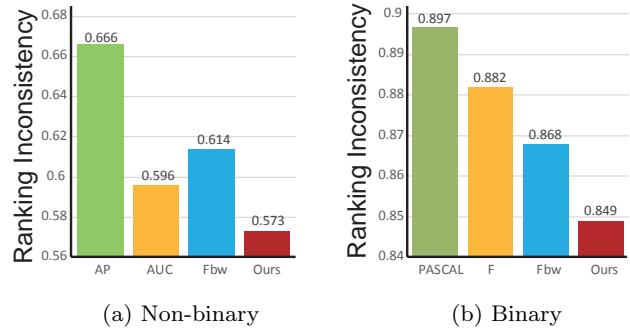


Figure 8 Meta-Measure 1 - results. The ranking correlation of an evaluation measure to that given by the SalCut segmentation (Non-binary evaluation) and image retrieval (Binary evaluation) application. We used 1-spearman’s rho as the results presentation. The lower the score, the better an evaluation measure is in term of predicting the preference of the application. Our measure achieves a better performance over other evaluation measure.

5.1 Meta-Measure 1: Application Ranking

The first meta-measure examines the ranking correlation of the evaluation measure to that of an application that uses foreground maps [55]. We assume that the GT map is the optimal input for the application (the top path in Fig. 7). Then, given a foreground map, we compare the application’s output (the bottom path in Fig. 7) to that of the GT output. The closer the saliency map is to the GT, the closer its application output should be to the GT output. We compare the ranking result by each binary and non-binary evaluation measure: AP, AUC, Fbw, PASCAL, F_β and ours, to the ranking result by the application.

The work in [55] has examined three applications: object detection, segmentation and image retrieval. Here, we use the SalCut [16] method (for non-binary) and image retrieval (for binary) application (see Appendix Sec. 7) to compute this meta-measure⁵.

We utilize the 1-Spearman’s ρ measure [4] to evaluate the ranking accuracy of the measures, where a lower values indicates better ranking consistency. The score of 0 indicates that the evaluation measure ranked the saliency maps identically to that of the application. The score of 2 indicates that the evaluation measure ranked the foreground maps in a complete reverse order. Comparison between different measures (AP, AUC, Fbw, Ours) is shown in Fig. 8 (a), which indicates that our structure measure produces the best ranking consistency among other alternative methods. According to the example shown in Fig. 1, all of the current non-binary

⁵ We follow the same experimental settings with Fbw [55] for a fair comparison. Note that Fbw only provides the retrieval application, we can not achieve the other two application details.




Figure 9 Meta-Measure 2: State-of-the-Art vs. Generic. An evaluation measure should give the FM generated by the SOTA method (c) a higher score than the generic map (d) that does not consider the content of the image. Unfortunately, all of the current evaluation measures give the map in (d) a higher score than (c). Only our measure correctly ranks the SOTA result higher.

measures fail to rank the foreground maps correctly. Our measure correctly ranks these maps. In the case of binary maps, S-measure also offers a 5.35%, 3.74%, 2.19%, improvement over the PASCAL, F_β , and Fbw measure, which score 0.897, 0.882, and 0.868, respectively, compared to 0.849 by our measure.

5.2 Meta-Measure 2: SOTA vs. Generic

Our second meta-measure is that a measure should prefer the output achieved by a SOTA algorithm over generic baseline maps (*e.g.*, centered Gaussian map, see Fig. 9(d)) that discard the image content. A good evaluation measure should rank the SM generated by a SOTA model higher than a generic map.

We count the number of times a generic map scored higher than the mean score generated by the five SOTA models (CA [28], CB [34], RC [16], PCA [54], SVO [11]). The mean score provides an indication of model robustness. Results are shown in Fig. 10. The lower the value here, the better the measure is. Over 1000 images, our measure has only 11 errors (*i.e.*, generic winning over the s.t.a) for non-binary maps. Meanwhile, the AP and AUC measures are very poor and make significantly more mistakes. Our measure also offers a large improvement over the PASCAL, F_β , and Fbw.

5.3 Meta-Measure 3: Ground-Truth Switch

The third meta-measure specifies that a good SM should not obtain a higher score when switching to a wrong GT map. In Margolin *et al.* [55], a SM is considered as “good” when it scores at least 0.5 out of 1 (when compared to the original GT map). Using this threshold (0.5), top 41.8% of the total 5000 maps were deemed as “good” ones. For a fair comparison, we follow Margolin




Figure 10 Meta-Measure 2 - results. The percentage of times that an evaluation measure ranked a generic map (non-binary circle or binary centered gaussian map) higher than the FM generated by the SOTA model. The lower the score, the better the evaluation measure is. Our measure achieves the best performance.




Figure 11 Meta-Measure 3: Ground-Truth Switch. The score of a FM should decrease when using a wrong Switched GT as the reference. However, both AUC and AP gave the map in (b) a higher score when using (d) instead of (c) as the reference GT map. Using our measure, the score of (b) appropriately decreased when switching to random ground-truth (d).

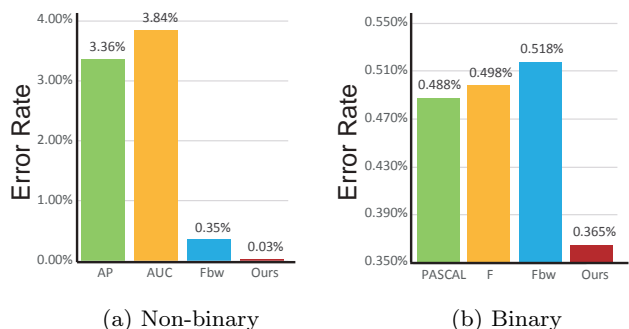


Figure 12 Meta-Measure 3 - results. The percentage of times (tested on 1000 ASD dataset) that an evaluation measure assigned a higher score when using an incorrect GT map. The lower the score, the better the measure is. Our measure achieves significant improvement over other measures in both non-binary and binary maps.

et al. to select the same percentage of “good” maps. For each of the 1000 images, 100 random GT switches were tested. We then counted the percentage of times that a measure increased a saliency map’s score when an incorrect GT map was used.

The Fig. 12 shows the results. The lower the score, the higher capability to match to the correct GT. Our measure performs the best about 10 times better compared to the second best measure. This is due to the fact that our measure captures the object structural similarity between a FM and a GT map. Our measure will assign a lower value to the “good” FM when using a random selected GT since the object structure has changed in the random GT.

5.4 Meta-Measure 4: Annotation Errors

The fourth meta-measure specifies that an evaluation measure should not be sensitive to slight errors/inaccuracies in the manual annotation of the GT boundaries. To perform this meta-measure, we make a slightly modified GT map by using morphological operations. An example is shown in Fig. 13. While the two ground truth maps in (b) & (c) are slightly different, a good measure should not switch the ranking between the two foreground maps (d) & (e), when using (b) or (c) as the reference.

We use the 1-Spearman’s ρ measure to examine the ranking correlation before and after annotation errors were introduced. The lower the score, the more robust an evaluation measure is to annotation errors [55]. Results are shown in Fig. 14. Our measure outperforms both the AP and the AUC but is not the best. Inspecting this reason, we realized that it is not always the case that the lower the score, the better an evaluation measure is. It is that sometimes “slight” inaccurate manual annotations can change the structure of the GT map, which in turn can change the rank.

We examined the effect of the structure change more carefully. Major structural changes often correspond to continuous large regions in the difference map between ground truth and its morphologically changed version. We used the sum of corroded version of the difference map as a measure of major structure change and to sort all the ground truth images.

Among top 10% least changed images, our measure and Fbw have the same MM4 scores (both of them are 0). When the topology of ground truth map does not change, our measure and Fbw preserve the original ranking. This can be seen in the example in Fig. 15. While ground truth maps (GT and Morphologic GT) differ slightly, both Fbw and our measure preserve the ranking order of the two saliency maps, depending on the used ground truth map.

For top 10% most changed images, we asked 3 users to judge whether the ground truth images have major structure changes. 95 out of 100 ground truth images were considered to have major structure changes (*e.g.*,

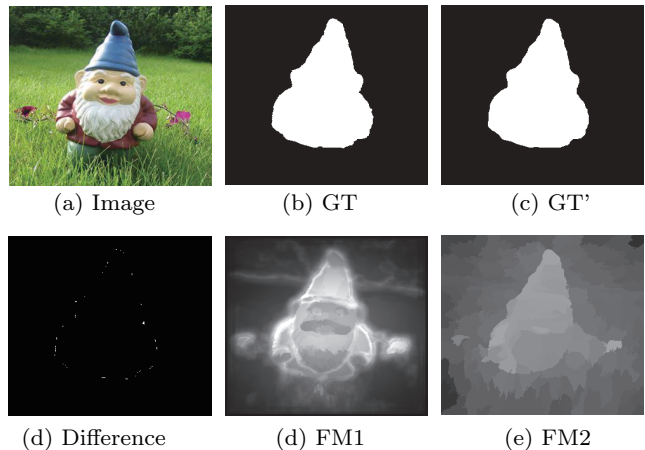


Figure 13 Meta-measure 4: Annotation Errors. An evaluation measure should not be sensitive to slight changes in the manual annotation of the GT boundaries. While GT (b) & GT’ (c) are almost identical, some measures switched the ranking order of the two foreground maps (d) & (e), depending on the different GT used. Our measure consistently ranked (d) higher than (e). Best viewed in color.

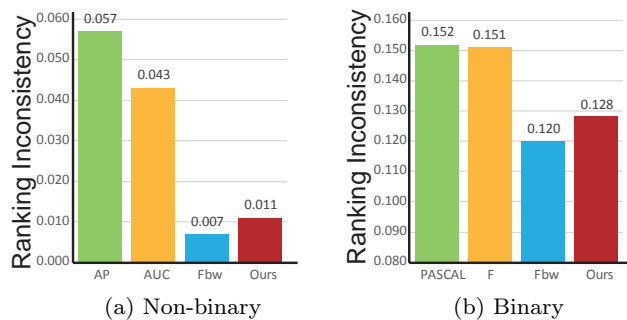


Figure 14 Meta-Measure 4 - results. The ranking correlation of an evaluation measure under small manual annotation inaccuracies. We use the 1-Spearman’s Rho measure to present the results. The lower the score, the better.

small bar, thin legs, slender foot and minute lines in each group; see Fig. 16), for which we believe that keeping the same ranks is not good. Fig. 17 is a typical example. When we use the GT map as the reference, Fbw and our measure rank the two maps properly. However, when using Morphologic GT as the reference, ranking results are different. Clearly, the blue-border SM is visually and structurally more similar to the Morphologic GT map than the red-border SM. A good measure should rank the blue-border SM higher than red-border SM. So the ranking of these two maps should be changed. While the Fbw measure fails to meet this end, our measure gives the correct order.

Above-mentioned analysis suggests that this meta-measure is not very reliable. Therefore, we do not include it in our comparison.




Figure 15 Structural Unchanged Case. Both of ours and the Fbw measure are not sensitive to inaccuracies (structural unchanged) in manual annotations of the GT boundaries.




Figure 16 Structural Changed Examples. The first row shows the GT maps. The second row shows morphologically changed versions. We observe significant structural changes.

5.5 Further Comparison

The results in Fig. 8, Fig. 10, and Fig. 12 show that our measure achieves the best performance using 3 meta-measures over the ASD1000 dataset. However, a good evaluation measure should perform well over almost all datasets. To demonstrate the robustness of our measure, we further use 10 SOTA algorithms for salient object detection to perform experiments on another 4 widely-used benchmark datasets.

Foreground Maps Collection. We used 10 SOTA algorithms including 3 traditional models (*i.e.*, ST [53], DRFI [35], and DSR [46]) and 7 deep learning based models (DCL [44], rfcn [65], MC [83], MDF [43], DISC [14], DHS [50], and ELD [39]) to generate the binary and non-binary foreground maps. Binary maps are obtained by thresholding the non-binary maps using image dependent adaptive thresholding method in [1].




Figure 17 Structural Changed Case. The ranking of an evaluation measure should be sensitive to the structural changes. Surprisingly, the current best measure (Fbw) does not account for structural changes. Using our measure, we rank the maps correctly. Best viewed on screen.

Benchmark Datasets. The 4 widely-used datasets include PASCAL-S [47], ECSSD [68], HKU-IS [43], and SOD [56]. PASCAL-S contains 850 challenging images, which have multiple objects in high background clutter. ECSSD contains 1000 semantically meaningful but structurally complex images. HKU-IS is another large dataset that contains 4445 large scale images. Most of the images in this dataset contain more than one salient object with low contrast. Finally, we also evaluate our measure over the SOD dataset, which is a subset of the BSDS dataset. It contains a relatively small number of images (300), but with multiple complex objects.


Results. Non-binary and binary maps' quantitative comparison results are shown in Tab. 1. Our measure performs the best according to the first meta-measure for both binary and non-binary maps evaluation. This indicates that our measure is more useful for applications than other measures.

For the evaluation results (binary and non-binary) in MM2, our measure performs better than the existing four measures (AP, AUC, F, PASCAL) with a large margin. The results on two easier datasets (ECSSD and HKU-IS) show that our measure and Fbw perform on par for both binary and non-binary maps.

According to meta-measure 3, our measure reduces the non-binary error rate by 67.62%, 44.05%, 17.81%, 69.23% on PASCAL, ECSSD, SOD and HKU-IS, respectively compared to the second ranked measure. For binary maps, our measure also reduces the error rate

Table 1 Non-binary (N-binary) & Binary maps’ quantitative comparison with current measures on 3 meta-Measures. The best result is highlighted in **bold**. MM:meta-Measure.

Type	Measure	PASCAL-S [47]			ECSSD [68]			SOD [56]			HKU-IS [43]		
		MM1	MM2(%)	MM3(%)	MM1	MM2(%)	MM3(%)	MM1	MM2(%)	MM3(%)	MM1	MM2(%)	MM3(%)
N-binary	AP	0.452	12.1	5.50	0.449	9.70	3.32	0.504	9.67	7.69	0.518	3.76	1.25
	AUC	0.449	15.8	8.21	0.436	12.1	4.18	0.547	14.0	8.27	0.519	7.02	2.12
	Fbw	0.365	7.06	1.05	0.401	3.00	0.84	0.384	16.3	0.73	0.498	0.36	0.26
	Ours	0.320	4.59	0.34	0.312	3.30	0.47	0.349	9.67	0.60	0.424	0.34	0.08
Binary	F	0.757	21.65	1.62	0.619	12.5	1.53	0.760	25.3	1.81	0.593	4.54	0.41
	PASCAL	0.905	19.41	1.59	0.787	11.0	1.54	0.911	24.0	1.81	0.786	3.62	0.41
	Fbw	0.802	13.76	1.72	0.675	7.50	1.54	0.814	19.0	1.42	0.665	2.16	0.58
	Ours	0.655	12.50	1.11	0.523	7.30	1.29	0.707	15.0	1.56	0.483	2.27	0.32

**Figure 18** The rank distance between Fbw and our measure. The (a)-(c) are the three datasets used to compute the rank distance between Fbw and our S-measure. The y axis of the plot is the number of the images. The x axis is the rank distance.

by 62.86%, 52.38%, 10.96%, 61.54% on PASCAL, ECSSD, SOD and HKU-IS, respectively compared to the second ranked measure. This indicates that our measure has higher capacity to capture the structural similarity between FM and GT maps.

Overall, our measure wins in the majority of cases indicating that it is more robust than other measures.

5.6 Meta-Measure 5: Human Judgments

Here, we propose a new meta-measure to evaluate foreground evaluation measures. This meta-measure specifies that the map ranking according to an evaluation measure should highly agree with the human ranking. It is argued that “a human being is the best judge to evaluate the output of any segmentation algorithm” [58]. However, subjective evaluation over all images of a dataset is impractical due to time and monetary costs. To the best of our knowledge, there is no such visual similarity evaluation database available in the object segmentation domain that meets these requirements. Here, we focus on the non-binary maps to collect such a database.

Stimuli. The source foreground maps are sampled from three large scale datasets: PASCAL-S, ECSSD, and HKU-IS. As mentioned above, we use 10 SOTA saliency models to generate the maps for each dataset.

Therefore, we have 10 foreground maps for each image. We use Fbw and our measure to evaluate the 10 maps and then pick the first ranked map according to each measure. If the two measures choose the same map, their rank distance is 0. If one measure ranks a map first, but the other ranks the same map in the n -th place, then their rank distance is $|n - 1|$. Fig. 18-20 (a), (b) and (c) show the histogram of rank distances between the two measures. The blue-box is the number of images for each rank distance. Some maps with rank distance greater than 0 are chosen as candidates for our user study.

User Study. We randomly selected 100 pairs of maps from the three datasets. The top panel in Fig. 21 (b) shows one example trial with the best map according to our measure on the left, and the best map according to the Fbw on the far right. The user is asked to choose the map she thinks resembles the most with the GT map. In this example, these two maps are obviously different making the user decide easily. In another example (bottom panel in Fig. 21 (b)), the two maps are too similar making it difficult for the used to choose the one closet to the GT. Therefore, we avoid showing such cases to the subjects. Finally, we are left with a stimulus set of size 50 pairs. We developed a mobile phone app (see Fig. 21 (a)) to conduct the user study.




Figure 19 The rank distance between AP and our measure. The (a)-(c) are the three datasets used to compute the rank distance between AP and our S-measure. The y axis of the plot is the number of the images. The x axis is the rank distance.





Figure 20 The rank distance between AUC and our measure. The (a)-(c) are the three datasets used to compute the rank distance between AUC and our S-measure. The y axis of the plot is the number of the images. The x axis is the rank distance.



(a) the mobile phone app (b) the one example trial

Figure 21 Our user study platform.

We collected data from 45 viewers who were naive to the purpose of the experiment. Viewers had normal or corrected vision. (age distribution is 19-29 years old; education from undergraduate to Ph.D; from 10 different majors such as history, medicine and finance; 25 males and 20 females)

Results. Results (Fbw vs. our measure) are shown in Fig. 22. The percentage of trials (averaged over subjects) in which a viewer preferred the map chosen by our measure is 63.69%. We used the same procedure to




Figure 22 Results of our user study (Fbw & S-measure). The x axis is the viewer id. The y axis shows the percentage of the trials in which a viewer preferred the map chosen by our measure.

conduct two additional user studies (AP vs. our measure, AUC vs. our measure). The results are 72.11% and 73.56%, respectively. This indicates that our measure correlates better with human judgments.

5.7 Saliency Model Comparison

Establishing that our S-measure offers a better way to evaluate foreground maps, here we compare 10 SOTA models on 4 datasets (PASCAL-S, ECSSD, HKU-IS, and SOD). Fig. 23 shows the rank of the 10 models. According to our measure, top three models in order are




Figure 23 Ranking of 10 saliency models using our new measure. The y axis shows the average score on each dataset (PASCAL-S [47], ECSSD [68], HKU-IS [43], SOD [56]).

dhsnet, DCL, and rfcn. Moreover, we also establish many representative online benchmarks (1.<http://dpfan.net/socbenchmark>; 2.<http://dpfan.net/d3netbenchmark>; 3.<http://dpfan.net/cosod3k/>; 4.<http://dpfan.net/camouflage/>) to compared our S-measure with other measures.

6 Ablation Study

To investigate the contribution of each part in our S-measure, we further conduct the ablation study on a new human ranking dataset [20]. The FMDatabase⁶ consists of 185 color images and 555 ranked maps. Similar to meta-measure1, we also utilize the 1-Spearman’s ρ metric to evaluate the ranking performance of the measures.

Table 2 Ablation studies of human ranking (Using FMDatabase-IJCAI’18) in terms of 1 - Spearman’s measure.

Settings	1 - Spearman’s measure
Object-aware	0.195
Region-aware	0.142
S-measure (Ours)	0.140
Fbw [55]	0.149
SSIM [67]	0.223

As shown in Tab. 2, we observe that our S-measure outperforms other settings (*i.e.*, object-aware, region-aware) on FMDatabase. It clearly shows that only region-level or object-level structural similarity cannot provide stable evaluation performance. Since the object-aware structure similarity mainly focusing on assessing the property of sharp foreground-background contrast. It more like a global evaluation. On the other hand, the region-aware links to the local evaluation which based on window-level statistics. Compared with existing two classical metrics (*i.e.*, Fbw and SSIM), we also found that our metric achieve the best results.

⁶ <http://dpfan.net/e-measure/>



Figure 24 Failure Case. Our S-measure ranked (d) Generic higher than (c) rfcn due to these prediction map without obvious structure.

7 Discussion and Conclusion

In this paper, we analyzed the current saliency evaluation measures based on pixel-wise errors and showed that they ignore the structural similarities. We then presented a new structural similarity measure known as **S-measure** which simultaneously evaluates region-aware and object-aware structural similarities between a saliency map and a ground-truth map. Our measure is based on two important characteristics: 1) sharp foreground-background contrast, and 2) uniform saliency distribution. Further, the proposed measure is efficient and easy to calculate. Experimental results on 5 datasets demonstrate that our measure performs better than the current measures including AP, AUC, and Fbw. Finally, we conducted a behavioral judgment study over a database of 100 saliency maps and 50 ground-truth maps. Data from 45 subjects shows that on average they preferred the saliency maps chosen by our measure over the saliency maps chosen by the Fbw.

All metrics are double-edged swords. Generally, it’s hard to argue which measure is the best one. These measures are deeply coupled to the actual applications, *e.g.*, some applications may favor the correctness of important regions while some may prefer the continuity in local structures. We observe a few failure cases where a prediction map without (or less) object structure will achieve a higher S-measure score. For example, as shown in Fig. 24, the proposed S-measure does not work well in this situation in which the object of the GT (b) without a clear structure. Consequently, to evaluate the foreground maps, we need to assess whether it performs well on multi metrics at the same time.

In summary, our measure offers new insights into foreground map evaluation where current measures fail to examine the strengths and weaknesses of models fully. For now, we have found that the saliency (*e.g.*, RGB SOD [15, 74, 82, 86], RGB-D SOD [23, 26, 73, 84, 85], RGB-T SOD [80], light field SOD [36, 60, 77], VSOD [33], 360 SOD, SID [41], Saliency Ranking [2], Co-SOD [22, 24, 79], and HR SOD [71]) community has begun to widely adopt this measure even in the camouflaged object detection [19, 57, 72, 72] and medical image segmentation [21].




Figure 25 Integration of the image with its foreground map. (a) & (c) are the images (Imgs) with its GT and FM. (b) & (d) are the combined images.

Appendices

Image Retrieval Application. We use the publicly available content based image retrieval system LIRE [40] as our application. Firstly, we generate a combined image as Fig. 25 (a)-(d). For each combined GT image (e.g., GT_1, \dots, GT_n ; n denotes the total number of images), we use LIRE to extract the CEDD feature and then search a list of 100 most similar images $GT_{lst-i} = \{G_{1-i}, \dots, G_{100-i}\}$. LIRE also assigns the 100 images with score ($GT_{score-i} = \{G_{s1-i}, \dots, G_{s100-i}\}$) which indicates the similarity value. Accordingly, for each FM_{comb} image we can obtain the 100 most similarity images FM_{lst-i} and corresponding score sets $FM_{score-i}$. Finally, let $Q_i = \{GT_{lst-i} \cap FM_{lst-i}\}$. We search FM_k which equals to G_i in the FM_{lst-i} . If FM_k exists, we record the index k and the corresponding score F_{sk-i} . The similarity S_i of each FM assigned by LIRE is:

$$S_i = \begin{cases} F_{sk-i} + \frac{1}{k} + \frac{\|Q_i\|}{100}, & G_i \in Q_i \\ \frac{\|Q_i\|}{100}, & otherwise \end{cases} \quad (16)$$

Acknowledgments

This research was supported by the National Key Research and Development Program of China under Grant No. 2018AAA0100400, NSFC (61922046), and S&T innovation project from Chinese Ministry of Education.

References

- Achanta, R., Hemami, S., Estrada, F., Susstrunk, S.: Frequency-tuned salient region detection. In: IEEE Conf. Comput. Vis. Pattern Recog., pp. 1597–1604 (2009)
- Amirul Islam, M., Kalash, M., Bruce, N.D.: Revisiting salient object detection: Simultaneous detection, ranking, and subitizing of multiple salient objects. In: IEEE Conf. Comput. Vis. Pattern Recog., pp. 7142–7150 (2018)
- Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. IEEE Trans. Pattern Anal. Mach. Intell. **33**(5), 898–916 (2011)
- Best, D., Roberts, D.: Algorithm as 89: the upper tail probabilities of spearman’s rho. Journal of the Royal Statistical Society. Series C (Applied Statistics) **24**(3), 377–379 (1975)
- Borji, A.: What is a salient object? a dataset and a baseline model for salient object detection. IEEE Trans. Image Process. **24**(2), 742–756 (2015)
- Borji, A., Cheng, M.M., Jiang, H., Li, J.: Salient object detection: A benchmark. IEEE Trans. Image Process. **24**(12), 5706–5722 (2015)
- Borji, A., Itti, L.: State-of-the-art in visual attention modeling. IEEE Trans. Pattern Anal. Mach. Intell. **35**(1), 185–207 (2013)
- Borji, A., Sihite, D., Itti, L.: Quantitative analysis of human-model agreement in visual saliency modeling: A comparative study. IEEE Trans. Image Process. **22**(1), 55–69 (2013)
- Borji, A., Sihite, D.N., Itti, L.: What stands out in a scene? a study of human explicit saliency judgment. Vis. Res. **91**(0), 62–77 (2013)
- Bylinskii, Z., Judd, T., Borji, A., Itti, L., Durand, F., Oliva, A., Torralba, A.: Mit saliency benchmark (2015) (2015)
- Chang, K.Y., Liu, T.L., Chen, H.T., Lai, S.H.: Fusing generic objectness and visual saliency for salient object detection. In: Int. Conf. Comput. Vis., pp. 914–921 (2011)
- Chen, H., Li, Y.f.: Progressively complementarity-aware fusion network for rgb-d salient object detection. In: IEEE Conf. Comput. Vis. Pattern Recog. (2018)
- Chen, T., Cheng, M.M., Tan, P., Shamir, A., Hu, S.M.: Sketch2photo: internet image montage. ACM Trans. Graph. **28**(5), 124 (2009)
- Chen, T., Lin, L., Liu, L., Luo, X., Li, X.: Disc: Deep image saliency computing via progressive representation learning. IEEE T. Neural Netw. Learn. Syst. **27**(6), 1135–1149 (2016)
- Chen, Z., Xu, Q., Cong, R., Huang, Q.: Global context-aware progressive aggregation network for salient object detection. In: AAAI Conf. Art. Intell. (2020)
- Cheng, M., Mitra, N.J., Huang, X., Torr, P.H., Hu, S.: Global contrast based salient region detection. IEEE Trans. Pattern Anal. Mach. Intell. **37**(3), 569–582 (2015)
- Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. Int. J. Comput. Vis. **88**(2), 303–338 (2010)
- Fan, D.P., Cheng, M.M., Liu, Y., Li, T., Borji, A.: Structure-measure: A new way to evaluate foreground maps. In: Int. Conf. Comput. Vis., pp. 4548–4557 (2017)
- Fan, D.P., Ji, G.P., Cheng, M.M., Shao, L.: Concealed object detection. arXiv preprint arXiv:2102.10274 (2021)

20. Fan, D.P., Ji, G.P., Qin, X., Cheng, M.M.: Cognitive vision inspired object segmentation metric and loss function. *SSI* (2021). DOI 10.1360/SSI-2020-0370
21. Fan, D.P., Ji, G.P., Zhou, T., Chen, G., Fu, H., Shen, J., Shao, L.: Pragnet: Parallel reverse attention network for polyp segmentation. In: *Med. Image. Comput. Comput. Assist. Interv.*, pp. 263–273 (2020)
22. Fan, D.P., Li, T., Lin, Z., Ji, G.P., Zhang, D., Cheng, M.M., Fu, H., Shen, J.: Re-thinking co-salient object detection. *IEEE Trans. Pattern Anal. Mach. Intell.* (2021)
23. Fan, D.P., Lin, Z., Zhang, Z., Zhu, M., Cheng, M.M.: Rethinking rgb-d salient object detection: Models, data sets, and large-scale benchmarks. *IEEE T. Neural Netw. Learn. Syst.* (2021)
24. Fan, Q., Fan, D.P., Fu, H., Tang, C.K., Shao, L., Tai, Y.W.: Group collaborative learning for co-salient object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2021)
25. Feng, D., Barnes, N., You, S., McCarthy, C.: Local background enclosure for rgb-d salient object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 2343–2350 (2016)
26. Fu, K., Fan, D.P., Ji, G.P., Zhao, Q., Shen, J., Zhu, C.: Siamese network for rgb-d salient object detection and beyond. *IEEE Trans. Pattern Anal. Mach. Intell.* (2021)
27. Ghosh, J., Lee, Y.J., Grauman, K.: Discovering important people and objects for egocentric video summarization. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 1346–1353 (2012)
28. Goferman, S., Zelnik-Manor, L., Tal, A.: Context-aware saliency detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **34**(10), 1915–1926 (2012)
29. Gorji, S., Clark, J.: Going from image to video saliency: Augmenting image salience with dynamic attentional push. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
30. Guo, C., Zhang, L.: A novel multiresolution spatiotemporal saliency detection model and its applications in image and video compression. *IEEE Trans. Image Process.* **19**(1), 185–198 (2010)
31. Islam, M.A., Kalash, M., D. B. Bruce, N.: Revisiting salient object detection: Simultaneous detection, ranking, and subitizing of multiple salient objects. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
32. Itti, L.: Automatic foveation for video compression using a neurobiological model of visual attention. *IEEE Trans. Image Process.* **13**(10), 1304–1318 (2004)
33. Ji, Y., Zhang, H., Jie, Z., Ma, L., Wu, Q.J.: Casnet: A cross-attention siamese network for video salient object detection. *IEEE T. Neural Netw. Learn. Syst.* (2020)
34. Jiang, H., Wang, J., Yuan, Z., Liu, T., Zheng, N., Li, S.: Automatic salient object segmentation based on context and shape prior. In: *Brit. Mach. Vis. Conf.*, p. 9 (2011)
35. Jiang, H., Wang, J., Yuan, Z., Wu, Y., Zheng, N., Li, S.: Salient object detection: A discriminative regional feature integration approach. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 2083–2090 (2013)
36. Jiang, Y., Zhou, T., Ji, G.P., Fu, K., Zhao, Q., Fan, D.P.: Light field salient object detection: A review and benchmark. *arXiv preprint arXiv:2010.04968* (2020)
37. Kanan, C., Cottrell, G.: Robust classification of objects, faces, and flowers using natural image statistics. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 2472–2479 (2010)
38. Lazebnik, S., Schmid, C., Ponce, J.: Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, vol. 2, pp. 2169–2178 (2006)
39. Lee, G., Tai, Y.W., Kim, J.: Deep saliency with encoded low level distance map and high level features. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 660–668 (2016)
40. Lew, M.S., Sebe, N., Djeraba, C., Jain, R.: Content-based multimedia information retrieval: State of the art and challenges. *ACM T Multim. Comput.* **2**(1), 1–19 (2000)
41. Li, G., Xie, Y., Lin, L., Yu, Y.: Instance-level salient object segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 2386–2395 (2017)
42. Li, G., Xie, Y., Wei, T., Lin, L.: Flow guided recurrent neural encoder for video salient object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
43. Li, G., Yu, Y.: Visual saliency based on multiscale deep features. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 5455–5463 (2015)
44. Li, G., Yu, Y.: Deep contrast learning for salient object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 478–487 (2016)
45. Li, L., Jiang, S., Zha, Z., Wu, Z., Huang, Q.: Partial-duplicate image retrieval via saliency-guided visually matching. *IEEE Trans. Multimedia* **20**(3), 13–23 (2013)
46. Li, X., Lu, H., Zhang, L., Ruan, X., Yang, M.H.: Saliency detection via dense and sparse reconstruction. In: *Int. Conf. Comput. Vis.*, pp. 2976–2983 (2013)
47. Li, Y., Hou, X., Koch, C., Rehg, J.M., Yuille, A.L.: The secrets of salient object segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 280–287 (2014)
48. Liu, G., Fan, D.: A model of visual attention for natural image retrieval. In: *IEEE Inf. Sci. Cloud Comput. Companion*, pp. 728–733 (2013)
49. Liu, G.H., Yang, J.Y., Li, Z.: Content-based image retrieval using computational visual attention model. *Pattern Recog.* **48**(8), 2554–2566 (2015)
50. Liu, N., Han, J.: Dhsnet: Deep hierarchical saliency network for salient object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 678–686 (2016)
51. Liu, N., Han, J., Yang, M.H.: Picanet: Learning pixel-wise contextual attention for saliency detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
52. Liu, T., Yuan, Z., Sun, J., Wang, J., Zheng, N., Tang, X., Shum, H.Y.: Learning to detect a salient object. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(2), 353–367 (2011)
53. Liu, Z., Zou, W., Le Meur, O.: Saliency tree: A novel saliency detection framework. *IEEE Trans. Image Process.* **23**(5), 1937–1952 (2014)
54. Margolin, R., Tal, A., Zelnik-Manor, L.: What makes a patch distinct? In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 1139–1146 (2013)
55. Margolin, R., Zelnik-Manor, L., Tal, A.: How to evaluate foreground maps? In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 248–255 (2014)
56. 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: *Int. Conf. Comput. Vis.*, vol. 2, pp. 416–423 (2001)
57. Mei, H., Ji, G.P., Wei, Z., Yang, X., Wei, X., Fan, D.P.: Camouflaged object segmentation with distraction mining. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2021)
58. Pal, N.R., Pal, S.K.: A review on image segmentation techniques. *Pattern Recog.* **26**(9), 1277–1294 (1993)
59. Peng, H., Li, B., Xiong, W., Hu, W., Ji, R.: Rgb-d salient object detection: a benchmark and algorithms. In: *Eur. Conf. Comput. Vis.*, pp. 92–109 (2014)
60. Piao, Y., Rong, Z., Zhang, M., Lu, H.: Exploit and replace: An asymmetrical two-stream architecture for versatile light field saliency detection. In: *AAAI Conf. Art. Intell.*, pp. 11865–11873 (2020)

61. Pont-Tuset, J., Marques, F.: Measures and meta-measures for the supervised evaluation of image segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 2131–2138 (2013)
62. Qin, X., Fan, D.P., Huang, C., Diagne, C., Zhang, Z., Sant’Anna, A.C., Suárez, A., Jagersand, M., Shao, L.: Boundary-aware segmentation network for mobile and web applications. *arXiv preprint arXiv:2101.04704* (2021)
63. Rutishauser, U., Walther, D., Koch, C., Perona, P.: Is bottom-up attention useful for object recognition? In: *IEEE Conf. Comput. Vis. Pattern Recog.*, vol. 2, pp. II–37 (2004)
64. Tiantian, W., Zhang, L., Lu, H., Borji, A.: Detect globally, refine locally: A novel approach to saliency detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
65. Wang, L., Wang, L., Lu, H., Zhang, P., Ruan, X.: Saliency detection with recurrent fully convolutional networks. In: *Eur. Conf. Comput. Vis.*, pp. 825–841 (2016)
66. Wang, W., Shen, J., Dong, X., Borji, A.: Saliency object detection driven by fixation prediction. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
67. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)
68. Xie, Y., Lu, H., Yang, M.H.: Bayesian saliency via low and mid level cues. *IEEE Trans. Image Process.* **22**(5), 1689–1698 (2013)
69. Yu, Q., Xie, L., Wang, Y., Zhou, Y., K. Fishman, E., L. Yuille, A.: Recurrent saliency transformation network: Incorporating multi-stage visual cues for small organ segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
70. Zeng, Y., Lu, H., Zhang, L., Feng, M., Borji, A.: Learning to promote saliency detectors. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
71. Zeng, Y., Zhang, P., Zhang, J., Lin, Z., Lu, H.: Towards high-resolution salient object detection. In: *Int. Conf. Comput. Vis.*, pp. 7234–7243 (2019)
72. Zhai, Q., Li, X., Yang, F., Chen, C., Cheng, H., Fan, D.P.: Mutual graph learning for camouflaged object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2021)
73. Zhang, J., Fan, D.P., Dai, Y., Anwar, S., Saleh, F., Aliakbarian, S., Barnes, N.: Uncertainty inspired rgb-d saliency detection. *IEEE Trans. Pattern Anal. Mach. Intell.* (2021)
74. Zhang, J., Xie, J., Barnes, N.: Learning noise-aware encoder-decoder from noisy labels by alternating back-propagation for saliency detection. In: *Eur. Conf. Comput. Vis.* (2020)
75. Zhang, J., Zhang, T., Dai, Y., Harandi, M., Hartley, R.: Deep unsupervised saliency detection: A multiple noisy labeling perspective. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
76. Zhang, L., Dai, J., Lu, H., He, Y., Wang, G.: A bi-directional message passing model for salient object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
77. Zhang, M., Li, J., Wei, J., Piao, Y., Lu, H.: Memory-oriented decoder for light field salient object detection. In: *Adv. Neural Inform. Process. Syst.*, pp. 898–908 (2019)
78. Zhang, P., Wang, D., Lu, H., Wang, H., Ruan, X.: Amulet: Aggregating multi-level convolutional features for salient object detection. In: *Int. Conf. Comput. Vis.* (2017)
79. Zhang, Q., Cong, R., Hou, J., Li, C., Zhao, Y.: Coad-net: Collaborative aggregation-and-distribution networks for co-salient object detection. In: *Adv. Neural Inform. Process. Syst.* (2020)
80. Zhang, Q., Huang, N., Yao, L., Zhang, D., Shan, C., Han, J.: Rgb-t salient object detection via fusing multi-level cnn features. *IEEE Trans. Image Process.* **29**, 3321–3335 (2019)
81. Zhang, X., Wang, T., Qi, J., Lu, H., Wang, G.: Progressive attention guided recurrent network for salient object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018)
82. Zhao, J.X., Liu, J.J., Fan, D.P., Cao, Y., Yang, J., Cheng, M.M.: Egnnet: Edge guidance network for salient object detection. In: *Int. Conf. Comput. Vis.*, pp. 8779–8788 (2019)
83. Zhao, R., Ouyang, W., Li, H., Wang, X.: Saliency detection by multi-context deep learning. In: *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 1265–1274 (2015)
84. Zhao, X., Pang, Y., Zhang, L., Lu, H., Zhang, L.: Suppress and balance: A simple gated network for salient object detection. In: *Eur. Conf. Comput. Vis.* (2020)
85. Zhou, T., Fan, D.P., Cheng, M.M., Shen, J., Shao, L.: Rgb-d salient object detection: A survey. *Comput. Vis. Media* pp. 1–33 (2021)
86. Zhuge, M., Fan, D.P., Liu, N., Zhang, D., Xu, D., Shao, L.: Saliency object detection via integrity learning. *arXiv preprint arXiv:2101.07663* (2021)
87. Zhuge, M., Gao, D., Fan, D.P., Jin, L., Chen, B., Zhou, H., Qiu, M., Shao, L.: Kaleido-bert: Vision-language pre-training on fashion domain. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2021)