

Gradient-Induced Co-Saliency Detection

Zhao Zhang¹[0000-0002-1521-8163], Wenda Jin²[0000-0001-5964-4781],
Jun Xu¹[0000-0002-1602-538X], and Ming-Ming Cheng¹[0000-0001-5550-8758]

¹ TKLNDST, CS, Nankai University

² College of Intelligence and Computing, Tianjin University
`zzhang@mail.nankai.edu.cn; cmm@nankai.edu.cn`

Abstract. Co-saliency detection (Co-SOD) aims to segment the common salient foreground in a group of relevant images. In this paper, inspired by human behavior, we propose a gradient-induced co-saliency detection (GICD) method. We first abstract a consensus representation for a group of images in the embedding space; then, by comparing the single image with consensus representation, we utilize the feedback gradient information to induce more attention to the discriminative co-salient features. In addition, due to the lack of Co-SOD training data, we design a jigsaw training strategy, with which Co-SOD networks can be trained on general saliency datasets without extra pixel-level annotations. To evaluate the performance of Co-SOD methods on discovering the co-salient object among multiple foregrounds, we construct a challenging *CoCA* dataset, where each image contains at least one extraneous foreground along with the co-salient object. Experiments demonstrate that our GICD achieves state-of-the-art performance. Our codes and dataset are available at <https://mmcheng.net/gicd/>.

Keywords: Co-saliency detection, new dataset, gradient inducing, jigsaw training

1 Introduction

Co-Saliency Detection (Co-SOD) aims to discover the common and salient objects by exploring the inherent connection of multiple relevant images. It is a challenging computer vision task due to complex variations on the co-salient objects and backgrounds. As a useful task for understanding correlations in multiple images, Co-SOD is widely employed as a pre-processing step for many vision tasks, such as weakly-supervised semantic segmentation [44, 47], image surveillance [15, 30], and video analysis [18, 19], *etc.*

Previous researches study the Co-SOD problem from different aspects [6, 20, 23]. At the early stage, researchers explored the consistency among a group of relevant images using handcrafted features, *e.g.*, SIFT [5, 18], color and texture [13, 24], or multiple cues fusion [4], *etc.* These shallow features are not discriminative enough to separate co-salient objects in real-world scenarios. Recently, learning-based methods achieve encouraging Co-SOD performance by




Fig. 1. Human behavior inspired GICD. GIM is the gradient inducing module.

exploring the semantic connection within a group of images, via deep learning [23, 43], self-paced learning [17, 52], metric learning [16], or graph learning [20, 55], etc. However, these methods suffer from the inherent discrepancy in features, due to varying viewpoints, appearance, and positions of the common objects. How to better utilize the connections of relevant images is worth deeper investigation.

How do humans segment co-salient objects from a group of images? Generally, humans first browse the group of images, summarize the shared attributes of the co-salient objects with “general knowledge” [33], and then segment the common objects in each image with these attributes. This process is shown in Fig. 1. Inspired by human behavior, we design an end-to-end network with corresponding two stages. To obtain the shared attributes of the common objects as humans do, we calculate the consensus representation of multiple relevant images in a high-dimensional space with a learned embedding network. Once the consensus representation is obtained, for each image, we propose a Gradient Inducing Module (GIM) to imitate the human behavior of comparing a specific scene with the consensus description to feedback matching information.

In GIM, the similarity between the single and consensus representations can be measured first. As high-level convolutional kernels with different semantic awareness [36, 56], we can find out the kernels that are more related to the consensus representation and enhance them to detect co-salient objects. To this end, by partially back-propagating, we calculate the gradients of the similarity with respect to the top convolution layer as the feedback information. High gradient values mean corresponding kernels have a positive impact on the similarity results; thus, by assigning more weight to these kernels, the model will be induced to focus on the co-salient related features. Moreover, to better discriminate the co-salient object in each level of the top-down decoder, we propose an Attention Retaining Module (ARM) to connect the corresponding encoder-decoder pairs of our model. We call this two-stage framework with GIM and ARM as Gradient-Induced Co-saliency Detection (GICD) network. Experiments on benchmark datasets demonstrate the advantages of our GICD over previous Co-SOD methods.

Without sufficient labels, existing Co-SOD networks [23, 43, 48] are trained with semantic segmentation datasets, e.g., Microsoft COCO [27]. However, the annotated objects in segmentation datasets are not necessarily salient. In this




Fig. 2. Pipeline of our Gradient-Induced Co-saliency Detection (GICD) method. GIM denotes the Gradient Inducing Module, while ARM means the Attention Retaining Module. “ \bullet ”, “ \otimes ”, “ \oplus ”, and “ \odot ” represent the inner product, element-wise production, element-wise addition, and the sigmoid function, respectively.

paper, we introduce a novel jigsaw strategy to extend existing salient object detection (SOD) datasets, without extra pixel-level annotating, for training Co-SOD networks.

In addition, to better evaluate the Co-SOD methods’ ability of discovering co-salient object(s) among multiple foregrounds, most images in an evaluation dataset should contain at least one unrelated salient foreground except for the co-salient object(s). As can be seen in Fig. 3, this is ignored by the current Co-SOD datasets [2, 11, 45, 51]. To alleviate the problem, we meticulously construct a more challenging dataset, named *Common Category Aggregation (CoCA)*.

In summary, our major contributions are as follows:

- **We propose a gradient-induced co-saliency detection (GICD) network for Co-SOD.** Specifically, we propose a gradient inducing module (GIM) to pay more attention to the discriminative co-salient features, and an attention retaining module (ARM) to keep the attention during the top-down decoding.
- **We present a jigsaw strategy to train Co-SOD models on general SOD datasets without extra pixel-level annotations,** to alleviate the problem of lacking Co-SOD training data.
- **We construct a challenging CoCA dataset with meticulous annotations,** providing practical scenarios to better evaluate Co-SOD methods.
- Experiments on the *CoSal2015* [51] and our *CoCA* datasets demonstrate that our GICD outperforms previous Co-SOD methods. Extensive ablation studies validate the effectiveness of our contributions.

2 Related Works

2.1 Co-Saliency Object Detection (Co-SOD)

Different from traditional salient object detection (SOD) task [8, 12, 14], Co-SOD aims to automatically segment the common salient objects in a group of relevant images. Early Co-SOD methods assume that the co-salient objects in multiple images share low-level consistency [52]. For instance, Li *et al.* [24] introduced a co-multi-layer graph by exploring color and texture properties. Fu *et al.* [13] explored the contrast, spatial, and corresponding cues to enforce global association constraint by clustering. Cao *et al.* [4] integrated multiple saliency cues by a self-adaptive weighting manner. Tsai *et al.* [39] extracted co-salient objects by solving an energy minimization problem over a graph.

Recently, many deep learning-based methods have been proposed to explore high-level features for the Co-SOD task [17, 51, 53]. These methods can be divided into two categories. One is a natural deep extension from traditional low-level consistency. It explores the high-level similarity to enhance the similar candidate regions among multiple images. For example, Zhang *et al.* [51] jointly investigated inter-group separability and intra-group consistency depending on high-level CNN features. Hsu *et al.* [17] proposed an unsupervised method by maximizing the similarity among multiple foregrounds and minimizing the similarity between foregrounds and backgrounds with graphical optimization. Jiang *et al.* [20] explored the superpixel-level similarity by intra- and inter-graph learning using the graph convolution network. Zhang *et al.* [53] proposed a mask-guided network to obtain coarse Co-SOD results and then refined the results by multi-label smoothing. The second category of deep methods is based on joint feature extracting. They often extract the common feature for a group of images, and then fuse it with each single image feature. For instance, Wei *et al.* [43] learn a shared feature for every five images with a group learning branch, then concatenate the shared feature with every single feature to get the final prediction. Li *et al.* [23] extend this idea with a sequence model to process variable length input. Wang *et al.* [40] and Zha *et al.* [48] learn a category vector for an image group to concatenate with each spatial position of a single image feature on multiple levels.

2.2 Co-SOD Datasets

Current Co-SOD datasets include mainly *MSRC* [45], *iCoseg* [2], *CoSal2015* [51], and *CoSOD3k* [11], etc. In Fig. 3, we show some examples of these datasets and our *CoCA* dataset. *MSRC* [45] is mainly for recognizing objects from images. In [13, 51], they select 233 images of seven groups from *MSRC-v1* for evaluating detection accuracy. *iCoseg* [2] contains 643 images of 38 groups in invariant scenes. In the above datasets, the co-salient objects are mostly the same in similar scenes, and consistent in appearance. *CoSal2015* [51] and *CoSOD3k* [11] are two large-scale datasets containing 2015 and 3316 images, respectively. In the two datasets, some target objects belong to the same category differ greatly in




Fig. 3. Examples of current popular datasets and our proposed *CoCA* dataset. In *CoCA*, except for the co-salient object(s), each image contains at least one extraneous salient object, which enables the dataset to better evaluate the models’ ability of discovering co-salient object(s) among multiple foregrounds.

appearance, which makes them more challenging datasets. However, these above datasets are not well-designed for evaluating the Co-SOD algorithms because they only have a single salient object in most images. Taking the athlete of the *iCoseg* in Fig. 3 as an example, although the athlete is co-salient in different images, these data can be easily processed by a SOD method because there is no other extraneous salient foreground interference. Although this awkwardness has been avoided in some groups in *CoSal2015* and *CoSOD3k*, it is not guaranteed in most cases. As discovering the co-salient object(s) among multiple foregrounds is the primary pursuit of a Co-SOD method in real-world applications [50], to evaluate this ability better, we construct a challenging *CoCA* dataset, where each image contains at least one extraneous salient object.

3 Proposed Method

Fig. 2 shows the flowchart of our gradient-induced co-saliency detection (GICD) network. Our backbone is the widely used Feature Pyramid Network (FPN) [26]. For the Co-SOD task, we incorporate it with two proposed modules: the gradient inducing module (GIM), and the attention retaining module (ARM). GICD detects co-salient objects in two stages. It first receives a group of images as input for exploring a consensus representation in a high-dimensional space with a learned embedding network. The representation describes the common patterns of the co-salient objects within the group. Then, it turns back to segment the co-salient object(s) for each sample. In this stage, for inducing the attention of the model on co-salient regions, we utilize GIM to enhance the features closely




Fig. 4. Visualization of high-level features induced by GIM. In the six small images on the left, the kernels above are not sensitive to the target object, while the kernels below are related to the target object. Their corresponding important values based on gradients are marked with green numbers on the top-left corners. The mean values of F_n^5 and induced \tilde{F}_n^5 are shown in the form of orange heat maps.

related to co-salient object by comparing single and consensus representation in the embedding space. In order to retain the attention during the top-down decoding, we use ARM to connect each encoder-decoder pairs. We train the GICD network with jigsaw training strategy, where the Co-SOD models can be trained on SOD dataset without extra pixel-level annotations.

3.1 Learning Consensus Representation

Given a group of images $\mathcal{I} = \{I_n\}_{n=1}^N$, to locate the co-salient object(s) in each image, we should first know what patterns the co-salient objects have based on prior knowledge. To this end, we propose to learn a consensus representation with a pre-trained embedding network, for the co-salient objects of the image group \mathcal{I} . Deep classifiers can be naturally utilized for representation learning [34], where the prior knowledge of semantic attribute can be transformed from the parameters pre-trained on ImageNet [7]. In this case, we employ a pre-trained classified network $\mathcal{F}(\cdot)$, such as VGG-16, as our embedding network by removing the softmax layer. It first extracts the representation $e_n = \mathcal{F}(I_n) \in \mathbb{R}^d$ of each image I_n , where d is the dimension of the last full connection layer. The consensus representation e^\dagger can be calculated by $e^\dagger = \text{Softmax}\left(\sum_{n=1}^N e_n\right)$, to describe the common attributes of this image group.

3.2 Gradient Inducing Module

After obtaining the consensus representation e^\dagger of the group \mathcal{I} , for each image, we focus on how to find the discriminative features that match the consensus description. As demonstrated in [36, 56], high-level convolutional layers naturally possess semantic-specific spatial information. We denote the five convolutional feature blocks of $\mathcal{F}(\cdot)$ as $\{F^1, F^2, \dots, F^5\}$. In Fig. 4, we show the feature maps of the last convolutional layer F^5 . The input image (1st column) contains a pocket watch and blue gloves, and the convolutional kernels focus on different regions (2nd to 4th columns). If assigning more importance to these kernels which closely concern about the co-salient objects, the model will tend to segment the co-salient objects (pocket watch) by decoding the induced features. As indicated

in [36], the discriminability of features in neural networks can be measured by the gradient obtained by optimizing objectives. Therefore, we propose a gradient inducing module (GIM) for enhancing the discriminative feature by exploring the feedback gradient information. As the encoder of our FPN backbone shares the fixed parameters with the consensus embedding network, it can also embed each image into the same space as consensus representation e^\dagger . For the extracted representation e_n of the n -th image, the similarity c_n between e_n and its consensus representation e^\dagger can be defined by inner product, *i.e.*, $c_n = e_n^\top e^\dagger$. Then we compute the positive gradient G_n flowing back into the last convolutional layer $F^5 \in \mathbb{R}^{w \times h \times c}$ to select discriminative features in F_n^5 , specifically,

$$G_n = \text{ReLU} \left(\frac{\partial c_n}{\partial F_n^5} \right) \in \mathbb{R}^{w \times h \times c}. \quad (1)$$

In this partial backpropagation, the positive gradient G_n reflects the sensitivity of the corresponding position to the final similarity score; that is, increasing activation value with a larger gradient will make the specific representation e_n more consistent with the consensus one e^\dagger . Therefore, the importance of a convolution kernel for a particular object can be measured by the mean of its feature gradients. Specifically, the channel-wise importance values can be calculated by global average pooling (GAP), namely $w_n = \text{GAP}(G_n) = \frac{1}{wh} \sum_i \sum_j G_n$, where $i = 1, \dots, w$ and $j = 1, \dots, h$. Once obtaining the weight, we can induce the high-level feature F_n^5 by assigning the importance value to each kernel $\tilde{F}_n^5 = F_n^5 \otimes w_n$, where \otimes denotes the element-wise production. As shown in Fig. 4, we visualize the mean heat-maps of F_n^5 and \tilde{F}_n^5 . without our GIM module, the kernels will averagely focus on both objects. One can see that the kernels more relevant to the co-salient category have higher gradient weights (marked with green numbers), and the attention of the network has shifted to the co-salient object(s) after gradient inducing.

3.3 Attention Retaining Module

In GIM, the high-level features have been induced by the gradient. However, top-down decoder is built upon the bottom-up backbone, and the induced high-level features will be gradually diluted when transmitted to lower layers. To this end, we propose an attention retaining module (ARM) to connect the corresponding encoder-decoder pairs of our GICD network. As shown in Fig. 2, for each ARM, the feature of encoder used for skip-connection is guided by the higher-level prediction. Through top-down iterative reminding, the network will focus the detail recovery of the co-salient regions without being interfered by other irrelevant objects. We take the channel-wise mean of \tilde{F}_n^5 as the first low-resolution guide map S_n^5 , and reduce \tilde{F}_n^5 to feature P_n^5 containing 64 channels. The decoding process with ARM is as follows:

$$\begin{cases} \tilde{F}_n^i = (S_n^{i+1}) \uparrow \odot F_n^i \\ P_n^i = \mathcal{E}^i \left((P_n^{i+1}) \uparrow + \mathcal{R}^i \left(\tilde{F}_n^i \right) \right), i \in \{4, 3, 2, 1\}, \\ S_n^i = \mathcal{D}^i (P_n^i), \end{cases} \quad (2)$$




Fig. 5. Visualization of attention retaining with ARM. The first row shows the multiple level intermediate features with (w/) our ARM, and the second row shows salient maps without (w/o) ARM. The prediction w/ ARM (second column, up) is more accurate than that w/o ARM (second column, down), since our ARM pays stronger attention to the co-salient regions.

where $(\cdot) \uparrow$ is the up-sampling operation. $\mathcal{R}^i(\cdot)$ consists of two convolutional layers, and reduces the enhanced features \tilde{F}_n^i to 64 channels. $\mathcal{E}^i(\cdot)$ is the corresponding two convolutional layers, with 64 kernels, in decoder. $\mathcal{D}^i(\cdot)$ is applied for deep supervision, and outputs a prediction by two convolutional layers followed by a sigmoid layer. The last S_n^1 is the final output.

To validate the effectiveness of our ARM, in Fig. 5, we show the intermediate features in different levels of the decoder with ARM (1st row) and without ARM (2nd row). We observe that, through GIM, both locate the co-salient object (*i.e.*, Teddy bear) successfully, while our GICD w/o ARM is gradually interfered by other salient objects during upsampling and produces inaccurate detection results. These results show that our ARM can effectively hold the attention on the co-salient objects in relevant images.

3.4 Jigsaw Training Strategy

Strategy. One important problem in Co-SOD task is that current SOD datasets, *e.g.*, DUTS [41] and MSRA-B [29], are not suitable for the training of Co-SOD networks. The reasons are two-fold: 1) they do not have class information, so it is impossible to train models in groups; 2) most samples in them only contain one salient foreground object. It is difficult to enable the network to distinguish the co-salient object(s) among multiple foreground objects. Recent Co-SOD methods [23, 40, 43] are trained on semantic segmentation datasets [27]. This suffers from two problems: 1) the label of such a semantic segmentation dataset is relatively rough, so the ability of recovery details of the trained network is not ideal, which cannot meet the accuracy requirements of downstream tasks; 2) the objects in such datasets are not necessarily salient. To alleviate these problems, we design a jigsaw strategy to transform SOD datasets into suitable training data for Co-SOD models: Step 1: we employ a classifier [31] to classify every SOD dataset into multiple categories since it has no category information. Step 2: we splice the samples of one category with the samples of other ones to form a new jigsaw, as shown in Fig. 6. This step is to ensure that an input image contains not only a co-salient foreground but also extraneous foreground objects.



Fig. 6. Demo of jigsaw training. A sample (cat), together with the samples from other categories, constitute new jigsaws for training.

Through the above steps, existing SOD datasets can be seamlessly utilized to train Co-SOD networks without additional pixel-level annotations.

Loss function. Considering the most important goal for co-saliency detection is to find the position of the common foreground objects correctly, we employ the soft intersection over union (IoU) loss [25, 35] for GICD, specifically,

$$\mathcal{L}(S, G) = 1 - \frac{\sum_c S(c) G(c)}{\sum_c [S(c) + G(c) - S(c) G(c)]}, \quad (3)$$

where S is the prediction, and G denotes ground-truth. c represents each pixel position in the image. The loss function of our model can be expressed as

$$L_{total} = \sum_{n=1}^N \sum_{i=1}^4 \mathcal{L}(S_n^i, G_n). \quad (4)$$

4 Proposed CoCA Dataset

Construction guidelines. We construct our *CoCA* dataset under four guidelines. **G1:** each image should contain at least one extraneous foreground, excluding the co-salient object(s). **G2:** in each image group, the co-salient objects are better to be different. **G3:** the dataset needs to be misaligned with the categories of the common training set, to explore the ability of the model on handling unseen categories. The guideline **G1** reflects whether or not the model can detect the co-salient objects, rather than only segmenting the foreground and background. **G2** can evaluate whether the model is robust to the intra-group differences. **G3** ensures that the model can be evaluated for its ability to detect co-salient objects from unknown categories robustly.

Construction procedures. With the above guidelines, we collect images from pixabay³. We divide them into 80 categories, covering everyday indoor and outdoor scenes. It is worth noting that these categories are outright staggered with

³ <https://pixabay.com>

Microsoft COCO [27], which is often used for the training of co-saliency models [23, 40, 43]. Most importantly, with manual screening, the images in our dataset include at least one extraneous salient object, excluding the co-salient object(s). We provide four levels of annotations: class level, bounding box level, object level, and instance level. The high-quality object-level annotations are applicable to the co-saliency detection task in this paper. Different levels of annotations of our dataset corresponds to different tasks, such as co-localization [21, 38], few-shot object segmentation [49, 54], and instance co-segmentation [37].

Dataset statistics. Our *CoCA* dataset consists of 80 categories with 1295 images. As shown in Fig. 3, these images are challenging in occlusion, clutter background, extraneous object interference *etc*. The number of images in each category is different, varying from 8 to 40. This diversity is helpful in evaluating the ability of the model for different image set sizes. The number of co-salient instances in an image is also diverse. 336 images have more than two co-salient instances. The diversity of the number of instances can help to evaluate the robustness of the model to multi-object scenarios.

5 Experiments

5.1 Implementation Details

We train our GICD network on the training set of *DUTS* [41] with our jigsaw training strategy. These samples are classified into 291 groups, which contains 8250 images with removing the noisy samples. Each sample will be combined with others to form three jigsaws as supplementary samples, as shown in Fig. 6; thus, the candidate training data is quadrupled. In each training epoch, we randomly select at most 20 samples from each group. The Adam optimizer [22] is used with an initial learning rate of 0.0001, $\beta_1 = 0.9$, and $\beta_2 = 0.99$. The learning rate is divided by 10 at the 50-th epoch. We train our GICD for 100 epochs in total. To accommodate the input images with our FPN backbone (VGG network), we resize them to 224×224 during the training and test stage, and the output saliency maps are resized back to the original size for evaluating. Our GICD is implemented in PyTorch [32], and runs at ~ 55 FPS on an NVIDIA GeForce RTX 2080Ti.

5.2 Evaluation Datasets and Metrics

Datasets. We employ two challenging datasets to evaluate the performance of various methods. The first dataset is *CoSal2015* [51]. In some image groups, *e.g.* baseball, it is challenging in the interference of extraneous salient objects. The other is our *CoCA*, where most images possess more than one irrelevant salient objects besides the co-salient target.

Metrics. We employ five widely used metrics as suggested by [17, 50, 53]: mean F-measure (F_{avg}) [1], maximum F-measure (F_{max}) [3], Precision-Recall (PR) curve, S-measure (S_α) [9], and mean E-measure (E_ξ) [10].

	CBCD [13]	GW [42]	CSMG [53]	RCAN [23]	BASNet [35]	PoolNet [28]	SCRN [46]	GICD Ours	
CoSal2015	$F_{\text{avg}} \uparrow$	0.378	0.639	0.721	0.670	0.778	0.768	0.755	0.835
	$F_{\text{max}} \uparrow$	0.547	0.706	0.787	0.764	0.791	0.785	0.783	0.844
	$S_\alpha \uparrow$	0.550	0.744	0.776	0.779	0.822	0.823	0.817	0.844
	$E_\xi \uparrow$	0.516	0.727	0.763	0.742	0.841	0.836	0.822	0.883
CoCA	$F_{\text{avg}} \uparrow$	0.230	0.358	0.390	0.360	0.398	0.394	0.394	0.504
	$F_{\text{max}} \uparrow$	0.313	0.408	0.499	0.422	0.408	0.404	0.413	0.513
	$S_\alpha \uparrow$	0.523	0.602	0.627	0.616	0.592	0.602	0.612	0.658
	$E_\xi \uparrow$	0.535	0.615	0.606	0.614	0.600	0.616	0.625	0.701

Table 1. Quantitative comparisons of mean F-measure [1] (F_{avg}), maximum F-measure [3] (F_{max}), S-measure [9] (S_α), and mean E-measure [10] (E_ξ) by our GICD and other methods on the *CoSal2015* [51] and *CoCA* datasets. “ \uparrow ” means that the higher the numerical value, the better the model performance.

5.3 Comparison with State-of-the-Arts

Comparison methods. We compare our GICD with seven state-of-the-art methods, including four Co-SOD method: RCAN [23], CSMG [53], GW [42], and CBCD [13], as well as three SOD mehtods: BASNet (ResNet-34) [35], PoolNet (ResNet-50) [28], and SCRN (ResNet-50) [46].

Quantitative evaluation. In the Tab. 1, we illustrate the quantitative results of our GICD and other state-of-the-art methods on the *CoSal2015* and our *CoCA* datasets. As can be seen, our GICD achieves better performance. The results show some interesting phenomena. On the *CoSal2015*, the SOD methods outperform most Co-SOD methods except GICD. The reason is that a large part of the images in *CoSal2015* have only one salient object, which can be solved by SOD algorithms. The advantages of Co-SOD algorithms cannot be fully reflected on this data, and these detail-oriented SOD methods easily surpass their performance. However, in our newly proposed *CoCA* dataset, this phenomenon is no longer obvious, because the salient objects in an image contain many objects that are not co-salient. This is why our *CoCA* dataset is more suitable for evaluating Co-SOD algorithms. Nevertheless, our GICD still surpasses the SOD methods on *CoSal2015*. It brings 11.4% improvement in terms of mean F-measure compared with the best Co-SOD method, 5.7% improvement compared with the SOD method. In our *CoCA* dataset, GICD brings 3.1% improvement in terms of S-measure compared with the best Co-SOD method, 4.6% improvement compared with the SOD method. Seen from Fig. 7, our method also outperforms other methods on the PR curve and F-measure curve. The trend of the curves demonstrates our method is less affected by the threshold, because it predicts the result with high confidence. This can avoid the problem of how to select the appropriate threshold in the subsequent practical applications.




Fig. 7. Precision-Recall (PR) and F-measure curves of our GICD and seven state-of-the-art methods on the *CoSal2015* and *CoCA* datasets. The node on each PR curve denotes the precision and recall value used for calculating maximum F-measure.

Qualitative results. In Fig. 8, we show some saliency maps produced by GICD and other compared methods for intuitive comparison. The samples we illustrate are challenging because the salient objects in each input include not only co-salient object(s) but also interference from other extraneous foreground(s). This is also reflected in the prediction results of SOD algorithms, which over segmented many unrelated regions. From the overall results, our GICD has high confidence in the prediction maps, even at the edge, while most of the other methods suffer from uncertain regions. Back to specific examples, baseball is the most challenging subset of the *CoSal2015* [51], because it varies greatly in size across images and is interfered by other salient objects. Results show that our method successfully handles the tiny size and the occlusions. In *CoCA*, boots class faces the interference of background color, and strawberry class has multiple segmentation targets. Nevertheless, GICD locates the target object accurately.

5.4 Ablation Study

In order to explore the contribution and mechanism of gradient inducing module (GIM), attention retaining module (ARM), and jigsaw training (JT) to our GICD network, we evaluated all possible combinations of the three candidates. Note that, these three candidates are interdependent and are not recommended to be used alone. As shown in Tab. 2, “A” is the baseline without JT, GIM, and ARM. It is actually a SOD model because it does not take into account any relationship between images.

Effectiveness of GIM. GIM is the core module of our GICD. With GIM, the variants of “C”, “E”, “G”, and our GICD can be regarded as a Co-SOD network; while, without GIM, the variants of “A”, “B”, “D”, and “F” have limited ability on discovering co-salient objects. By directly applying GIM, the variant “C” shifts the attention to co-salient objects in the high-level features, and achieves a certain performance improvement compared to the baseline. However, in this case, the training set is ill-posed for co-saliency detection task without JT, and the attention will be disturbed during the decoding without the help of ARM. These factors limit its performance. By introducing JT or ARM (variants “E”, “G” and our GICD), the effect of GIM is further enhanced.




Fig. 8. Visual comparison of our GICD with 7 state-of-the-arts (4 Co-SOD methods and 3 SOD methods) on the *CoSal2015* [51] and our *CoCA* datasets.

Effectiveness of ARM. ARM plays a role in retaining high-level prediction information during the top-down decoding. As shown in variant “D”, using ARM alone does not improve Co-SOD performance. The reason is that, without inducing by GIM, the prediction in high-level is actually the salient objects rather than co-salient objects. When cooperating with GIM in variant “G”, although trained on ill-posed data, it still compulsorily keeps the inducing information of GIM to an extend; thus, “G” achieves significantly better performance than the variant “C”. “E” is a variant with our GIM and ARM modules. As shown in Fig. 5, without ARM, it is easy to be interfered by irrelevant foreground when recovering object details. Therefore, its performance is inferior to our GICD.

Effectiveness of JT. The jigsaw training (JT) helps turn SOD datasets into Co-SOD ones, and serves as a useful strategy for training Co-SOD networks. In Tab. 2, without GIM, the variant model “B” and “F” are SOD models, not Co-SOD ones. Since no interactive cue between images is considered, a SOD model trained on Co-SOD dataset is unable to discover group-wise co-salient

Variant	Candidate			CoCA				CoSal2015 [51]			
	JT	GIM	ARM	$F_{\text{avg}} \uparrow$	$F_{\text{max}} \uparrow$	$S_{\alpha} \uparrow$	$E_{\xi} \uparrow$	$F_{\text{avg}} \uparrow$	$F_{\text{max}} \uparrow$	$S_{\alpha} \uparrow$	$E_{\xi} \uparrow$
A				0.420	0.430	0.601	0.627	0.788	0.800	0.818	0.852
B	✓			0.424	0.430	0.602	0.655	0.750	0.759	0.782	0.821
C		✓		0.446	0.462	0.618	0.643	0.809	0.824	0.833	0.868
D			✓	0.429	0.437	0.607	0.628	0.800	0.809	0.829	0.860
E	✓	✓		0.470	0.478	0.631	0.689	0.795	0.803	0.808	0.850
F	✓		✓	0.436	0.442	0.612	0.654	0.762	0.770	0.795	0.832
G		✓	✓	0.471	0.480	0.636	0.667	0.826	0.835	0.845	0.879
GICD	✓	✓	✓	0.504	0.513	0.658	0.701	0.835	0.844	0.844	0.883

Table 2. Ablation study of the proposed GICD on the *CoCA* and *CoSal2015* datasets. The candidates are jigsaw training (JT), gradient inducing module (GIM), and attention retaining module (ARM). Note that, the variants “A”, “B”, “D”, and “F”, without GIM, are actually SOD models rather than Co-SOD models. The experiments reflect the interaction mechanism of our three contributions.

connections, and the generated JT labels will bring meaningless predictions in this ill-posed scene; therefore, the JT does not work in these cases. When working with GIM in variant “E”, JT improves the effect in the challenging *CoCA* dataset. Similarly, this improvement can also be seen through the comparison between our GICD and variant “G”.

In summary, our three contributions of the GIM, ARM, and JT candidates are mutually reinforced for better co-saliency detection performance, as validated through comprehensive experiments.

6 Conclusions

In this paper, inspired by the mechanism of how human behaves on the Co-SOD task, we proposed an end-to-end Gradient-Induced Co-saliency Detection (GICD) method. In GICD, the gradient information, which highlights the discrimination of features, is generated from the comparison between single and consensus representations. Induced by the gradient, GICD pays more attention to discriminative convolutional kernels, enabling our model to locate the co-salient regions. Due to the lack of Co-SOD training data, we designed a novel jigsaw training strategy, with which we trained Co-SOD models on a general SOD dataset without extra pixel-level annotations. In addition, we constructed a challenging *CoCA* dataset for Co-SOD evaluation, to prosper the subsequent research on exploring real-world Co-SOD scenarios.

Acknowledgements. Ming-Ming Cheng is the corresponding author. Zhao Zhang and Wenda Jin are the joint first authors. This research was supported by Major Project for New Generation of AI under Grant No. 2018AAA0100400, NSFC (61922046), Tianjin Natural Science Foundation (18ZXZNGX00110), and the Fundamental Research Funds for the Central Universities, Nankai University (63201169).

References

1. Achanta, R., Hemami, S., Estrada, F., Susstrunk, S.: Frequency-tuned salient region detection. In: CVPR. pp. 1597–1604 (2009)
2. Batra, D., Kowdle, A., Parikh, D., Luo, J., Chen, T.: iCoseg: Interactive co-segmentation with intelligent scribble guidance. In: CVPR. pp. 3169–3176. IEEE (2010)
3. Borji, A., Cheng, M.M., Jiang, H., Li, J.: Salient object detection: A benchmark. IEEE TIP **24**(12), 5706–5722 (2015)
4. Cao, X., Tao, Z., Zhang, B., Fu, H., Feng, W.: Self-adaptively weighted co-saliency detection via rank constraint. IEEE TIP **23**(9), 4175–4186 (2014)
5. Chang, K.Y., Liu, T.L., Lai, S.H.: From co-saliency to co-segmentation: An efficient and fully unsupervised energy minimization model. In: CVPR 2011. pp. 2129–2136. IEEE (2011)
6. Chen, H.T.: Preattentive co-saliency detection. In: ICIP. pp. 1117–1120. IEEE (2010)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255. Ieee (2009)
8. Fan, D.P., Cheng, M.M., Liu, J.J., Gao, S.H., Hou, Q., Borji, A.: Salient objects in clutter: Bringing salient object detection to the foreground. In: ECCV. pp. 186–202 (2018)
9. Fan, D.P., Cheng, M.M., Liu, Y., Li, T., Borji, A.: Structure-measure: A new way to evaluate foreground maps. In: ICCV. pp. 4548–4557 (2017)
10. Fan, D.P., Gong, C., Cao, Y., Ren, B., Cheng, M.M., Borji, A.: Enhanced-alignment measure for binary foreground map evaluation. In: IJCAI. pp. 698–704 (2018)
11. Fan, D.P., Lin, Z., Ji, G.P., Zhang, D., Fu, H., Cheng, M.M.: Taking a deeper look at the co-salient object detection. In: CVPR (2020)
12. Fan, D.P., Zhai, Y., Borji, A., Yang, J., Shao, L.: BBS-Net: RGB-D salient object detection with a bifurcated backbone strategy network. In: ECCV (2020)
13. Fu, H., Cao, X., Tu, Z.: Cluster-based co-saliency detection. IEEE TIP **22**(10), 3766–3778 (2013)
14. Gao, S.H., Tan, Y.Q., Cheng, M.M., Lu, C., Chen, Y., Yan, S.: Highly efficient salient object detection with 100k parameters. In: ECCV (2020)
15. Gao, Z., Xu, C., Zhang, H., Li, S., de Albuquerque, V.H.C.: Trustful internet of surveillance things based on deeply-represented visual co-saliency detection. IEEE Internet of Things Journal (2020)
16. Han, J., Cheng, G., Li, Z., Zhang, D.: A unified metric learning-based framework for co-saliency detection. IEEE TCSVT **28**(10), 2473–2483 (2017)
17. Hsu, K.J., Tsai, C.C., Lin, Y.Y., Qian, X., Chuang, Y.Y.: Unsupervised CNN-based co-saliency detection with graphical optimization. In: ECCV. pp. 485–501 (2018)
18. Jerripothula, K.R., Cai, J., Yuan, J.: CARS: Co-saliency activated tracklet selection for video co-localization. In: ECCV. pp. 187–202. Springer (2016)
19. Jerripothula, K.R., Cai, J., Yuan, J.: Efficient video object co-localization with co-saliency activated tracklets. IEEE TCSVT **29**(3), 744–755 (2018)
20. Jiang, B., Jiang, X., Zhou, A., Tang, J., Luo, B.: A unified multiple graph learning and convolutional network model for co-saliency estimation. In: ACM Multimedia. pp. 1375–1382 (2019)
21. Joulin, A., Tang, K., Fei-Fei, L.: Efficient image and video co-localization with frank-wolfe algorithm. In: ECCV. pp. 253–268. Springer (2014)

22. Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. In: ICLR (2015)
23. Li, B., Sun, Z., Tang, L., Sun, Y., Shi, J.: Detecting robust co-saliency with recurrent co-attention neural network. In: IJCAI. pp. 818–825 (2019)
24. Li, H., Ngan, K.N.: A co-saliency model of image pairs. IEEE TIP **20**(12), 3365–3375 (2011)
25. Li, Z., Chen, Q., Koltun, V.: Interactive image segmentation with latent diversity. In: CVPR. pp. 577–585 (2018)
26. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR. pp. 2117–2125 (2017)
27. Lin, T.Y., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755 (2014)
28. Liu, J.J., Hou, Q., Cheng, M.M., Feng, J., Jiang, J.: A simple pooling-based design for real-time salient object detection. In: CVPR. pp. 3917–3926 (2019)
29. Liu, T., Yuan, Z., Sun, J., Wang, J., Zheng, N., Tang, X., Shum, H.Y.: Learning to detect a salient object. IEEE TPAMI **33**(2), 353–367 (2010)
30. Luo, Y., Jiang, M., Wong, Y., Zhao, Q.: Multi-camera saliency. IEEE TPAMI **37**(10), 2057–2070 (2015)
31. Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., van der Maaten, L.: Exploring the limits of weakly supervised pretraining. In: ECCV. pp. 181–196 (2018)
32. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An imperative style, high-performance deep learning library. In: NeurIPS. pp. 8024–8035 (2019)
33. Plaut, D.C.: Graded modality-specific specialisation in semantics: A computational account of optic aphasia. Cognitive neuropsychology **19**(7), 603–639 (2002)
34. Qi, H., Brown, M., Lowe, D.G.: Low-shot learning with imprinted weights. In: CVPR. pp. 5822–5830 (2018)
35. Qin, X., Zhang, Z., Huang, C., Gao, C., Dehghan, M., Jagersand, M.: Basnet: Boundary-aware salient object detection. In: CVPR. pp. 7479–7489 (2019)
36. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: Visual explanations from deep networks via gradient-based localization. In: ICCV. pp. 618–626 (2017)
37. Sun, H., Zhen, X., Zheng, Y., Yang, G., Yin, Y., Li, S.: Learning deep match kernels for image-set classification. In: CVPR. pp. 6240–6249 (2017)
38. Tang, K., Joulin, A., Li, L.J., Fei-Fei, L.: Co-localization in real-world images. In: CVPR. pp. 1464–1471 (2014)
39. Tsai, C.C., Li, W., Hsu, K.J., Qian, X., Lin, Y.Y.: Image co-saliency detection and co-segmentation via progressive joint optimization. IEEE TIP **28**(1), 56–71 (2018)
40. Wang, C., Zha, Z.J., Liu, D., Xie, H.: Robust deep co-saliency detection with group semantic. AAAI **33**, 8917–8924 (2019)
41. 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 (2017)
42. Wei, L., Zhao, S., Bourahla, O.E.F., Li, X., Wu, F.: Group-wise deep co-saliency detection. In: IJCAI. pp. 3041–3047 (2017)
43. Wei, L., Zhao, S., Bourahla, O.E.F., Li, X., Wu, F., Zhuang, Y.: Deep group-wise fully convolutional network for co-saliency detection with graph propagation. IEEE TIP (2019)

44. 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* **39**(11), 2314–2320 (2016)
45. Winn, J., Criminisi, A., Minka, T.: Object categorization by learned universal visual dictionary. In: *ICCV*. pp. 1800–1807 (2005)
46. Wu, Z., Su, L., Huang, Q.: Stacked cross refinement network for edge-aware salient object detection. In: *ICCV*. pp. 7264–7273 (2019)
47. Zeng, Y., Zhuge, Y., Lu, H., Zhang, L.: Joint learning of saliency detection and weakly supervised semantic segmentation. In: *ICCV*. pp. 7223–7233 (2019)
48. Zha, Z., Wang, C., Liu, D., Xie, H., Zhang, Y.: Robust deep co-saliency detection with group semantic and pyramid attention. *IEEE TNNLS* pp. 1–11 (2020)
49. Zhang, C., Lin, G., Liu, F., Yao, R., Shen, C.: CANet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In: *CVPR*. pp. 5217–5226 (2019)
50. Zhang, D., Fu, H., Han, J., Borji, A., Li, X.: A review of co-saliency detection algorithms: Fundamentals, applications, and challenges. *ACM TIST* **9**(4), 1–31 (2018)
51. Zhang, D., Han, J., Li, C., Wang, J., Li, X.: Detection of co-salient objects by looking deep and wide. *IJCV* **120**(2), 215–232 (2016)
52. Zhang, D., Meng, D., Han, J.: Co-saliency detection via a self-paced multiple-instance learning framework. *IEEE TPAMI* **39**(5), 865–878 (2016)
53. Zhang, K., Li, T., Liu, B., Liu, Q.: Co-saliency detection via mask-guided fully convolutional networks with multi-scale label smoothing. In: *CVPR*. pp. 3095–3104 (2019)
54. Zhang, X., Wei, Y., Yang, Y., Huang, T.: SG-One: Similarity guidance network for one-shot semantic segmentation (2018)
55. Zheng, X., Zha, Z.J., Zhuang, L.: A feature-adaptive semi-supervised framework for co-saliency detection. In: *ACM Multimedia*. pp. 959–966 (2018)
56. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: *CVPR*. pp. 2921–2929 (2016)