BiconNet: An Edge-preserved Connectivity-based Approach for Salient Object Detection

Ziyun Yang^a, Somayyeh Soltanian-Zadeh^a, Sina Farsiu^{a,b,c,d,*}

^aDepartment of Biomedical Engineering, Duke University, Durham, 27708, NC, USA ^bDepartment of Ophthalmology, Duke University Medical Center, Durham, 27710, NC, USA ^cDepartment of Electrical & Computer Engineering, Duke University, Durham, 27708, NC, USA

^dDepartment of Computer Science, Duke University, Durham, 27708, NC, USA

Abstract

Salient object detection (SOD) is viewed as a pixel-wise saliency modeling task by traditional deep learning-based methods. A limitation of current SOD models is insufficient utilization of inter-pixel information, which usually results in imperfect segmentation near edge regions and low spatial coherence. As we demonstrate, using a saliency mask as the only label is suboptimal. To address this limitation, we propose a connectivity-based approach called bilateral connectivity network (BiconNet), which uses connectivity masks together with saliency masks as labels for effective modeling of inter-pixel relationships and object saliency. Moreover, we propose a bilateral voting module to enhance the output connectivity map, and a novel edge feature enhancement method that efficiently utilizes edge-specific features. Through comprehensive experiments on five benchmark datasets, we demonstrate that our proposed method can be plugged into any existing state-of-the-art saliency-based SOD framework to improve its performance with negligible parameter increase.

Keywords: Salient object detection, Visual saliency, Connectivity modeling, Deep learning, Edge modeling

^{*}Corresponding author

Email addresses: ziyun.yang@duke.edu (Ziyun Yang), s.soltanian.zadeh@duke.edu (Somayyeh Soltanian-Zadeh), sina.farsiu@duke.edu (Sina Farsiu)

1. Introduction

As a fundamental task in computer vision, salient object detection (SOD) plays an essential role in image scene understanding [1] and has been applied to different tasks, such as weakly supervised semantic segmentation [2, 3], visual tracking [4], scene analysis [5, 6], video processing [7] and medical image analysis [8]. Convolutional neural networks (CNNs) have greatly promoted the develop-

- ment of SOD due to their capacity to extract multi-level semantic information. Most current CNN-based SOD models [9, 10] view the problem as a pixel-level saliency classification task; i.e., their only goal is to assign a saliency score to
- individual pixels. Despite promising results, these models are limited by insuf-10 ficient utilization of edge information, and insufficient attention to inter-pixel relationships. These problems together can result in blurred edges or low spatial coherence (i.e., have inconsistent saliency predictions for neighboring pixels that share similar spatial features), as Fig. 1.



(a) Original Image

(b) Ground Truth



(d) MINet + BiconNet

Figure 1: An example of insufficient modeling of pixel-wise relationship and structural information. MINet [11] results in both blurred edges (green box) and spatial inconsistency problems (red box). However, our model (MiNet + BiconNet) results in sharper edges and uniformly highlighted predictions near the boundaries.

15

The edge problem has been alleviated somewhat by adding edge information into networks using extra supervision flows [12, 13, 14], but there is still room for impactful improvement. First, edge features represent only a small fraction of the image; using an extra path for edge supervision is still likely to provide insufficient information while generating redundant non-structural fea-

tures [11]. Second, the extra flows result in extra computational cost, making 20 these methods less efficient.

The problem of low spatial coherence due to insufficient attention to interpixel relationships has been addressed by using post-processing methods such as conditional random fields (CRF) to refine the output prediction [9, 15, 16].

However, these methods usually result in low processing speed. Some studies [12, 17] proposed that spatial coherence can be enhanced by adding complementary edge information. Other studies [11, 18] have suggested that the incoherence is caused by scale variation of multi-level features, and have proposed using multi-scale fusion to alleviate the problem.

Another approach to solve these problems is to remodel SOD with new informative labels. Traditional salient masks used as training labels treat all pixels inside a salient object equally and independently; as a result, they lack information about inter-pixel relationships and fundamentally make edges hard to detect. Therefore, using saliency masks as the only training label is a subopti-

- ³⁵ mal choice. In the label decoupling framework (LDF) for SOD [19], traditional salient masks were decoupled into a location-aware detailed map and body map, which were used as auxiliary labels for training. However, these new labels required specifically designed extra supervision flows and were not proved to be compatible with other existing models. Another group [20] introduced the con-
- ⁴⁰ nectivity mask—a multi-channel mask exhibiting connectivity of each pixel with its neighboring pixels—as the CNN label. Although the connectivity mask is an inter-pixel relation-aware label, this method completely replaces saliency prediction with pixel connectivity modeling, and therefore does not effectively utilize the original saliency information. In addition, the method ignores the inherent
- ⁴⁵ properties of this new label, making the results less promising. We propose that the design and effective utilization of an informative label which is *compatible with any existing method* can efficiently improve the performance of existing models.

Inspired by this concept, we developed a novel connectivity-based SOD framework called the Bilateral Connectivity Network (BiconNet) as shown in Fig. 2. BiconNet consists of four parts: a connectivity-based SOD backbone, a bilateral voting (BV) module, a region-guided channel aggregation (RCA) module, and a bilateral connectivity (Bicon) loss function. To model inter-pixel relationships, we first replace the backbone's label with a connectivity mask.

Then, to enhance the spatial coherence between neighboring pixels, we use a 55 BV module to obtain a more representative connectivity map called the Bicon map. After this step, we generate two single-channel saliency maps, with edge information emphasized, via an RCA module. Finally, we propose the Bicon loss function to further emphasize edge features and spatial consistency for final salient object detection.

60



Figure 2: The overview of BiconNet, which contains a backbone, an 8-channel connectivity fully connected layer, a BV module, and an RCA module. Note that we can directly get edge information from the ground truth connectivity map and use it for highlighting the edge-specific features in the RCA module.

BiconNet exhibits three advantages: First, by changing the CNN's intermediate goal to predicting pixel-wise connectivity, inter-pixel relation modeling has become one of the network's tasks. Thus, BiconNet can focus more attention on inter-pixel relationships. Second, based on the inherent property of connectivity masks, edge regions can be located directly from ground truth, 65 which are then emphasized in the final output for network training via the RCA module. Compared to other edge-based methods [12, 21], this is a more efficient way to aggregate edge features. Third and most importantly, since BiconNet changes only the output layer of the backbones and all other modules

(BV and RCA) are trained after it, BiconNet can be built on any saliency-based 70 SOD framework without changing the framework's original design (e.g., internal structure and loss functions), and will improve its performance.

In summary, there are three main contributions of this work:

- We propose a connectivity-based SOD framework called BiconNet to explicitly model pixel connectivity, enhance edge modeling, and preserve spatial coherence of salient regions. BiconNet can be *easily plugged into any existing SOD model* with neglectable parameter increases.
- We propose an efficient, connectivity-based edge feature extraction method that can directly emphasize the edge-specific information from the network output. We also introduce a new loss function, Bicon loss, to further enhance the utilization of the edge features and preserve the spatial consistency of the output.
- We build BiconNets with backbones of seven state-of-the-art SOD models. By comparing these BiconNets with the corresponding baselines, we show that our model outperforms the latter models on five widely used benchmarks using different evaluation metrics.

2. Related Work

75

80

85

Earlier SOD methods [22, 23, 24] mostly utilized hand-crafted features to detect salient regions. These methods cannot effectively capture high-level semantic information from data, and are ineffective when dealing with complex scenes in images. CNN-based models have recently become the main choice for SOD due to their multi-level feature extraction ability. However, in earlier CNN-based SOD models [9, 25], erroneous predictions were usually made near the salient edges, and low spatial coherence occurred in the middle of the

salient region or near the edges. There are three ways to solve these problems: multi-scale feature aggregation models, edge-enhanced models, and problem remodeling methods.

2.1. Multi-scale Feature Aggregation Models

One reason for the problems described above is that detailed features can be diluted as the CNN becomes deeper. To utilize saliency features more efficiently, one solution is to aggregate multi-scale information. Hou et al. [9] demonstrated that using short connections between different layers helped aggregate multi-scale features. Chen et al. [26] proposed a model that can aggregate low-level detailed features, high-level semantic features, and global context features to learn the relationship between different salient regions. Qin et al. [10] proposed a nested network that uses Residual U-blocks to extracted multi-scale

features. Li et al. [27] extracted saliency features from three different scales of the images and aggregated them for final detection. Pang et al. [11] extracted effective multi-scale features from two interaction modules and preserved the spatial consistency of intra-class units. Although effective, these methods usually require extra computational power for the frequent feature aggregations between different layers.

2.2. Edge-enhanced Models

- To preserve edge information, edge-enhanced models intentionally generate ¹¹⁵ extra edge features for training. Zhao et al. [12] built another supervision flow for the edge features, which were fused with the salient features at the final stages of their network. Liu et al. [21] extracted the edge features from another edge detection dataset and used these for joint training with saliency detection. Qin et al. [28] added a refinement module after their encoder-decoder ¹²⁰ structure to refine the boundary details. Zhang et al. [29] proposed a boundary localization module to extract structural information. Wu et al. [13] exploited the logical interrelation between the edge map and saliency map and proposed a bidirectional framework to refine both tasks. Zhou et al. [14] proposed an
- ¹²⁵ models show the effectiveness of adding edge features for saliency detection, but they usually generate redundant features and are computationally expensive

approach that interactively fuse edge features and saliency features. These

since they add extra supervision flows for the edge path. In our work, the edge information is used in a more efficient way, as shown in Fig. 3.



Figure 3: Different edge-based models: (a) edge cue models [12, 21]; (b) interactive edge models [13, 14]; (c) BiconNet. Both (a) and (b) need to include at least one extra flow branch for the edge features. In contrast, BiconNet can directly receive the edge location from the connectivity ground truth and then emphasize the edge-specific information in the output via a simple RCA module.

2.3. Problem Remodeling Methods

Compared to the above models which focus on the internal structure of the network, an efficient way to solve the SOD problem is to rethink the task and remodel it using more informative labels. Wei et al. [19] decoupled the ground truth label into a body map and a detail map according to the location of object edges and used three supervision flows for training. However, the authors did not demonstrate a general way to utilize these labels in an existing framework. In addition, although these labels worked well in detecting the salient edges, they were not inter-pixel relation-aware. Kampffmeyer et al. [20] replaced the saliency labels with connectivity masks and illustrated improvements achieved by this change. This approach, called ConnNet, remodeled the problem of SOD

¹⁴⁰ by converting the saliency prediction task into several sub-tasks of foreground connectivity prediction. However, this method did not fully utilize the information of the connectivity mask. In addition, the method is incompatible with many saliency evaluation metrics as it does not predict a single-channel saliency probability map. We propose a method to overcome these problems, described in the next sections.

3. Proposed Method

145

3.1. Framework Overview

Our framework, BiconNet, consists of four parts: a connectivity-based SOD backbone, a BV module, an RCA module, and a Bicon loss function. For the backbone, we can use any existing saliency-based SOD framework. An overview of our method is shown in Fig. 2.

3.2. Connectivity Vector/Mask

Given an existing SOD backbone, our first step is to replace its single-channel saliency map output with an 8-channel connectivity map by changing its fully connected layers and to replace its label with the connectivity mask. In the next step, we will introduce connectivity vectors and masks/maps.

A connectivity [30] vector of a pixel is a multi-entry binary vector used to indicate whether the pixel is connected to its neighboring pixels. In the 8-neighbor system, given a pixel at coordinates (x, y), we use an 8-entry connectivity vector to represent the unidirectional connectivity with its neighbors in the square area of $[x \pm 1, y \pm 1]$ with every entry representing one specific direction. Given a

binary saliency mask G_S with size $H \times W$, by deriving the connectivity vector for every pixel in G_S , we obtain an 8-channel mask G_C with size $H \times W \times 8$ called the connectivity mask (Fig. 4). The i_{th} channel of G_C (G_{Ci}) represents

¹⁶⁵ if the original pixels on G_S are connected with their neighboring pixels at the i_{th} directions (e.g., upper left if i = 1 using row-major order). In this work, as in [20], We define connectedness only for the adjacent salient pixels. For better understanding, we call the discrete ground truths as connectivity masks G_C and the network's continuous outputs C as connectivity maps.

170

160

We show that learning a connectivity mask G_C provides three main advantages over a binary segmentation mask G_S . First, compared to G_S where every



Figure 4: Visualization of converting a saliency mask to a connectivity mask. The dashed box on G_S shows the 8-neighbor region of the selected pixel. G_S will be boundary-mirrored if needed. Every channel of G_C represents pixel connectivity at a certain direction.

entry only indicates the saliency of the current pixel, G_C focuses more on the mutual relationship between its pixels. Second, G_C itself contains more structural information (such as edges) than G_S . Specifically, in G_C , the elements

of the connectivity vector for an edge pixel are always a mixture of zeros and ones, whereas internal foreground pixels have all-ones connectivity vectors and background pixels have all-zeros connectivity vectors (Fig. 5). We call this property the turbidity of the edge connectivity vectors. Thus, given a ground truth connectivity vector of a pixel, we can always determine whether it is an edge pixel simply by checking the zero and one distribution of the vector. As shown in future sections, this property is important as it provides an efficient way to utilize edge information. Third, besides showing the connectivity of saliency pixels, every entry of G_C also reflects the connection direction. Thus, G_C is a structure- and inter-pixel relationship-aware label.

185 3.3. Bilateral Voting Module

For every two neighboring pixels in G_S , there is a unique element pair in G_C representing the unidirectional connectivity between them. We call such a pair in G_C a connectivity pair. To be specific, consider a pair of neighboring pixels on



Connectivity vector of edge pixel E: (0,0,0,0,0,1,1,1) M: (0,0,0,0,0,0,0) N: (1,1,1,1,1,1,1)

Figure 5: The turbidity property for edge pixels. M, N, and E represent pixels in the background (M), inside the salient region (N), and at the edge (E). Only pixel E has a mixture of zeros and ones in its connectivity vector, whereas M and N have all-zeros and all-ones connectivity vectors, respectively.

 G_S : M at location (x, y) and N at (x+a, y+b) $a, b \in \{0, \pm 1\}$. We can obtain the ¹⁹⁰ unidirectional connectivity from M to N from the value of $G_{Cj}(x, y)$, where G_{Cj} is the channel that represents the relative direction from N to M. For example, if N is located lower-right of M, then j = 8 (row-major order). Similarly, the connectivity from N to M can be found at $G_{C(9-j)}(x+a, y+b) = G_{C1}(x+1, y+1)$. We call the two elements $G_{Cj}(x, y)$ and $G_{C(9-j)}(x+a, y+b)$ a connectivity pair

of M and N. The same concept is also defined for the output connectivity map C, where every two neighboring pixels in the salient map have a unique connectivity pair in C representing the *probability* of the unidirectional connection. Fig. 6 shows an example of this case when a = b = 1.

The concepts of saliency and connectedness are closely related and mutually convertible: If two pixels are connected, they are salient. Two pixels of G_S are considered as connected (salient) if and only if both elements of its connectivity pair agree with this connection, i.e., if and only if $G_{Cj}(x,y) = G_{C(9-j)}(x + a, y + b) = 1$. We call this the *discrete* bilateral connectivity agreement, which reveals the bidirectional property of pixel connections and shows the importance of mutual impacts between neighboring pixels.

From this agreement, we know theoretically that the two elements from a connectivity pair should have the same connection probability to each other. However, in practice, connectivity pairs of the network's *continuous* outputs (i.e., the connectivity maps C) rarely satisfy this agreement. These disagree²¹⁰ ments result in spatial inconsistencies. To model the neighboring dependency and preserve the spatial consistency, we propose a novel connectivity-enhancement module called bilateral voting (BV) module.



Figure 6: Illustration of how connectivity pair is defined and how a Bicon map is generated via bilateral voting (BV) when a = 1 and b = 1. In the predicted Conn map (middle) the two entries $C_1(2,2)$ and $C_8(1,1)$ represent two predicted unidirectional connectivity probabilities of the highlighted neighboring pixels at (1,1) and (2,2) in the saliency map (left). After the BV module, the generated Bicon map is the bidirectional representation of the connectivity for pixels in the saliency map.

Given a connectivity map output C, the goal of the BV module is to generate another connectivity map that satisfies the bilateral connectivity agreement. To do so, we first extract all of the connectivity pairs. Then, we multiply the two elements in every connectivity pair and assign the resulting value to both elements, yielding a new connectivity map \tilde{C} . This process is shown in Fig. 6 and is formulated as:

$$\widetilde{C}_{j}(x,y) = \widetilde{C}_{9-j}(x+a,y+b)$$

$$= C_{j}(x,y) \times C_{9-j}(x+a,y+b),$$
(1)

where the subscript j means the j_{th} channel, $a, b \in \{0, \pm 1\}$. The logic behind the formula is that we can get the bidirectional pixel connection probability by multiplying every two elements of a connectivity pair, as each represents a unidirectional connectivity probability relative to the other. Since this continuous process is similar to making the discrete bilateral agreement, we call it bilateral voting. We name the new map \tilde{C} the bilateral connectivity map (Bicon map), and the original output C the Conn map. In the later sections, we will show that the BV module is important both at the training and inference phases.

3.4. Region-guided Channel Aggregation

225

After BV module, we obtain an enhanced 8-channel map \widetilde{C} with every channel representing the bidirectional probability of pixel connection at a specific ²³⁰ direction. In the previous sections, we pointed out that pixel connectivity is the sufficient and necessary condition of pixel saliency for neighboring pixels. Therefore, the overall probability of a pixel being connected with its neighbors reflects its saliency. To obtain a single-channel map representing saliency, we propose a region-guided channel aggregation (RCA) module to summarize the directional connectivity information in the eight channels of \widetilde{C} into a single-channel output \widetilde{S} using a function f. The generic form is written as:

$$\widetilde{S}(x,y) = f\{\widetilde{C}_i(x,y)\}_{i=1}^8 \quad , \tag{2}$$

where f is an adaptive aggregating operation that varies with location (x, y), \tilde{S} represents the aggregated overall probability of current pixel being salient. This process can also be interpreted as applying a function f over every predicted connectivity vector in \tilde{C} to obtain an overall connection probability for the corresponding pixel. Here we define two types of f:

Global aggregation. A simple way to aggregate the values from different directions into a single value is to average them. By doing this, we will obtain a single-channel map with every pixel representing the average connection probability to its neighbors. In this case, f is the averaging operation for all locations. We call the resultant map the global map, denoted as \tilde{S}_{alobal} :

$$\widetilde{S}_{global}(x,y) = Mean\{\widetilde{C}_i(x,y)\}_{i=1}^8 \quad . \tag{3}$$

Edge-guided aggregation. As mentioned, the edge pixels are the only pixels that have zero-one ground truth connectivity vectors. This property

yields two advantages. First, we can generate ground truth edge masks by sim-

ply searching the location of zero-one vectors in connectivity masks. Based on these ground truth edge masks, we can locate and decouple the edge regions and non-edge regions from the output. Second, we can encourage the network to learn this special data representation based on prior knowledge about the turbidity of edge vectors. Due to the imbalance between edge pixels and non-edge

- pixels, the network intends to make uniform predictions among all directions; i.e., predicting connectivity vectors as all ones or all zeros. An edge pixel, since it is part of the salient region, is more likely to be predicted as an internal foreground pixel with an all-ones connectivity vector. This is the main reason for blurring edges: it is difficult for networks to learn to discriminate edge pixels
- from other salient pixels. To emphasize the difference between these two types of pixels, we want the networks to pay extra attention to the likely misclassified direction of a predicted edge connectivity vector; i.e., directions that are in fact not connected to the current edge pixel. As for non-edge pixels, since they have all-ones or all-zeros ground truth connectivity vectors, we want the network to uniformly focus on all directions. To this end, we designed a region-adaptive

aggregation method for these two regions:

270

$$\widetilde{S}_{decouple}(x,y) = \begin{cases} 1 - \min\{\widetilde{C}_i(x,y)\}_{i=1}^8 & (x,y) \in P_{edge}, \\ Mean\{\widetilde{C}_i(x,y)\}_{i=1}^8 & (x,y) \notin P_{edge}, \end{cases}$$
(4)

where $\tilde{S}_{decouple}$ is called the edge-decoupled map and P_{edge} is the set of edge pixels. For the edge part, we emphasized the most likely disconnected directions by finding the minimum values of the predicted connectivity vectors. This design is highly correlated with the loss function, which is discussed in the next section.

So far, we have generated two single-channel maps: the global map \tilde{S}_{global} and the edge-decoupled map $\tilde{S}_{decouple}$ as shown in Fig. 2. $\tilde{S}_{decouple}$ is used for learning the edge-specific information; \tilde{S}_{global} is a more robust representation of salient objects and will be used as the final saliency prediction during inference.

275 3.5. Bicon Loss

Our loss function is defined as:

$$L_{bicon} = L_{decouple} + L_{con_const} + L_{opt} \quad . \tag{5}$$

We call this hybrid loss the Bicon loss L_{bicon} , where $L_{decouple}$ is the edgedecoupled loss, L_{con_const} is the connectivity consistency loss, and L_{opt} is the optional loss. We define the loss terms in the following sections.

280

285

Edge-decoupled loss. Binary cross entropy (BCE) [31] is one of the most widely used loss functions, and is defined as:

$$L_{bce}(S,G) = -\sum_{(x,y)} \left[G(x,y) \log(S(x,y)) + (1 - G(x,y)) \log(1 - S(x,y)) \right], \quad (6)$$

where $G(x, y) \in \{0, 1\}$ is the ground truth label of pixel (x, y) and S(x, y) is the prediction. BCE loss is a pixel-wise loss function that considers each pixel equally; thus, it does not consider inter-pixel relationships when the traditional saliency maps are used as the loss input [11, 32, 28]. We propose that this problem can be alleviated with use of a more informative and spatial relation-aware input. To this end, we used $\tilde{S}_{decouple}$ as the input of BCE. Although BCE is still calculating the loss independently for every unit, single units carry information

about their intrinsic saliency scores and the region-based connectivity. Based on Eq. 4, this loss is formulated as:

$$L_{decouple} = L_{bce}(\tilde{S}_{decouple}, G_S)$$

$$= \begin{cases} L_{bce}(1 - \min\{\tilde{C}_i(x, y)\}_{i=1}^8, G_S(x, y)) & (x, y) \in P_{edge}, \\ L_{bce}(mean\{\tilde{C}_i(x, y)\}_{i=1}^8, G_S(x, y)) & (x, y) \notin P_{edge}, \end{cases}$$
(7)

where $G_S(x, y) \in \{0, 1\}$ is the saliency ground truth of pixel (x, y), indicating whether the pixel is salient. Specifically, we can derive the edge part as:

$$L_{bce}(1 - \min\{\tilde{C}_{i}(x, y)\}_{i=1}^{8}, G_{S}(x, y))$$

= $L_{bce}(1 - \min\{\tilde{C}_{i}(x, y)\}_{i=1}^{8}, \mathbf{1})$
= $L_{bce}(\min\{\tilde{C}_{i}(x, y)\}_{i=1}^{8}, \mathbf{0})$. (8)

For the edge pixels, our goal is to make the network learn the sparse representation of the turbid edge vectors. As discussed previously, the edge pixels ²⁹⁵ are most likely to be predicted as internal salient pixels that have all ones in their connectivity vectors. Thus, a feasible way to learn the turbidity is to force the minimum value of the edge connectivity vector to be zero; i.e., we want the network to only focus on the disconnected direction for edge pixels. For the non-edge pixels, since they all have all-zeros or all-ones connectivity vectors, our goal is to make the average value across channels to be close to their labels. Namely, we want the network to put uniform weights on all directions.

Connectivity consistency loss. The connectivity consistency loss is the weighted sum of BCE losses applied to both the original Conn map (C) and the Bicon map (\tilde{C}) . It is defined as:

$$L_{con_const} = \omega_1 L_{conmap} + \omega_2 L_{bimap}$$

= $\omega_1 L_{bce}(C, G_C) + \omega_2 L_{bce}(\widetilde{C}, G_C),$ (9)

- where G_C is the corresponding ground truth 8-channel connectivity mask with every element $G_{Ci}(x, y) \in \{0, 1\}$, specifying whether a pixel at location (x, y)is connected to its c = 8 neighboring pixels. ω_1 and ω_2 are weighting factors. The first term, L_{conmap} , is designed for preserving spatial consistency. For the second term, L_{bimap} , since the bidirectional connection probability in \tilde{C} is
- exponentially correlated with the original unidirectional probability, it usually generates larger loss on hard pixels [18], such as edge pixels, while generating a smaller loss on easy pixels. The exception is background pixels, where a small loss will be generated no matter what background is. Therefore, there is a tradeoff between edge enhancement and background dilution in this loss term. To effectively utilize L_{bimap} , we assign it a lower weight. For all of our experiments,

we set $\omega_1 = 0.8, \omega_2 = 0.2$ unless otherwise noted.

Optional loss. As mentioned above, the BV and RCA modules together with the Bicon loss can be inserted into any existing saliency-based backbone to form the BiconNet architecture. Some existing studies [28, 11] have proposed specific loss functions with their network architectures. To maintain the integrity of these backbones, we apply the same loss function in these papers as our third term:

$$L_{opt} = L_{orig}(\widetilde{S}_{global}, G_S), \tag{10}$$

where $L_{orig}(\cdot)$ is the loss function defined in the original backbone's paper, \widetilde{S}_{global} is the global map. Note that L_{opt} is an optional loss term and will be applied according to the selection of backbones.

3.6. Inference

To obtain the single-channel saliency probability map in the inference stage of BiconNet, we first pass the output Conn map C through the BV module to get the Bicon map \tilde{C} . Then, we aggregate the channels with the averaging operation to get the global map \tilde{S}_{global} . Finally, we use \tilde{S}_{global} as the predicted saliency map, as shown in Fig. 2.

4. Experiments

4.1. Datasets and Evaluation Metrics

We evaluated our model on five frequently used SOD benchmark datasets: ³³⁵ HKU-IS [27] with 4,447 images, DUTS [33] with 10,553 images for training (DUTS-TR) and 5,019 for testing (DUTS-TE), ECSSD [34] with 1,000 images, PASCAL-S [35] with 850 images, and DUT-OMRON [36] with 5,168 images. For the evaluation metrics, we adopted the mean absolute error (MAE) [37], F-measure (F_{β}) [38], and E-measure (E_m) [39]. For the F-measure, we used the mean F-measure, F_{ave} , which is generated by thresholding the prediction map using an adaptive value equal to twice the mean of the prediction and is correlated with spatial consistency of the prediction [40].

4.2. Experiment Setup and Implementation Details

Model Setup. We adopted seven state-of-the-art models as both baselines and backbones to form the BiconNets: PoolNet [21], CPD-R [40], EGNet [12], F3Net [18], GCPANet [26], ITSD [14], MINet [11]. We replaced all of their saliency prediction layers with 8-channel fully-connected layers, followed by our BV and RCA modules. We used Bicon Loss as the loss function for all models. For the models with deep supervision mechanisms such as [12, 26], we replaced
³⁵⁰ all of the fully-connected layers with our connectivity layer followed by BV and RCA. For the extra edge supervision flows in [12, 14], we only replaced their edge labels with our connectivity-based edge labels generated by zero-one vector searching as discussed in Section 3.4 for consistency.

Implementation Details. We used the released official codes of the backbones for training both the baselines and the BiconNets. For baselines, we trained all of them from scratch, strictly following the instructions on their websites and the hyperparameter setting in their original papers. For the BiconNets, we used the same data pre-processing tricks as the corresponding baselines. For the hyperparameters, we only changed the starting learning rate (about 40% of the baselines') and the batch size for our BiconNets, as in Table 1. The rest of hyperparameters were the same as the baselines'. We implemented all our experiments in Pytorch 1.4.0 [41] using an NVIDIA RTX 2080Ti GPU. The

Table 1: The starting learning rate and batch size of BiconNet with different backbones.

code is available at: https://github.com/Zyun-Y/BiconNets.

Backbone	PoolNet	CPD-R	EGNet	F3Net	GCPANet	ITSD	MINet
Start Lr	$2e^{-4}$	$3.5e^{-5}$	$2e^{-5}$	0.0018	0.01	0.005	0.0018
Batch Size	10	10	10	16	16	8	32

4.3. Comparison with State-of-the-art Methods

365

370

Quantitative Comparison. To compare our method and the baselines, we list all experiments and their results in Table 2. As the results show, the absolute majority of our results (98/105) show better or the same performance compared to the corresponding baselines. Our method also achieved most of the best overall results (14/15) (marked with \dagger). The results also indicate that our model can make a uniform prediction on the salient regions and preserve

spatial consistency of the input more effectively than the baseline.

Table 2: Quantitative evaluation. seven methods were tested among five benchmark datasets. The mean F-measure (F_{ave}) , mean absolute error (MAE), and E-measure (E_m) were used to evaluate the results. \uparrow indicates that higher is better. We highlight the better result between every baseline and its BiconNet in red. We denote the best result of a column with a \dagger superscript, the second best one with a * superscript.

Model	1	HKU-I	s	DUT-TE		DUT-OMRON			PASCAL-S			ECSSD			
inodor	$\mathbf{F}_{ave}\uparrow$	MAE↓	$E_m \uparrow$	$\mathbf{F}_{ave}\uparrow$	MAE↓	$E_m \uparrow$	$\mathbf{F}_{ave}\uparrow$	MAE↓	$E_m \uparrow$	$\mathbf{F}_{ave}\uparrow$	MAE↓	$E_m \uparrow$	$\mathbf{F}_{ave}\uparrow$	MAE ↓	$E_m \uparrow$
$PoolNet_{19}$ [21]	0.885	0.038	0.941	0.787	0.047	0.876	0.728	0.061	0.851	0.787	0.085	0.833	0.904	0.045	0.919
${\rm PoolNet} + {\rm Bicon}$	0.909	0.034	0.950	0.826	0.042	0.902	0.759	0.057	0.866	0.812	0.072	0.853	0.916	0.040	0.925
$CPD\text{-}R_{19}~[40]$	0.888	0.034	0.946	0.788	0.044	0.886	0.737	0.056	0.863	0.783	0.071	0.848	0.892	0.038	0.925
CPD-R + Bicon	0.905	0.034	0.952	0.806	0.044	0.895	0.750	0.056	0.867	0.794	0.069	0.857	0.898	0.039	0.925
$EGNet_{19}$ [12]	0.900	0.031	0.952	0.804	0.038	0.894	0.750	0.053	0.867	0.794	0.073	0.847	0.905	0.037	0.927
EGNet + Bicon	0.917	0.031	0.954	0.842*	0.037*	0.912*	0.770	0.050 [†]	0.868	0.821	0.067	0.863*	0.922	0.037	0.930 [†]
$F3Net_{20}$ [18]	0.914	0.031	0.953	0.828	0.039	0.896	0.749	0.055	0.853	0.830	0.062	0.857	0.924	0.037	0.926
F3Net + Bicon	0.915	0.029	0.954	0.835	0.038	0.899	0.765	0.051*	0.863	0.830	0.062*	0.855	0.927	0.034 [†]	0.929*
$\operatorname{GCPANet}_{20}[26]$	0.896	0.032	0.950	0.812	0.038	0.892	0.743	0.056	0.856	0.812	0.063*	0.845	0.913	0.035	0.924
GCPANet + Bicor	0.918*	0.032	0.954	0.834	0.040	0.901	0.762	0.055	0.863	0.838*	0.061†	0.858	0.929*	0.036	0.929*
ITSD_{20} [14]	0.900	0.030	0.952	0.806	0.041	0.891	0.752	0.058	0.862	0.800	0.067	0.850	0.903	0.034^\dagger	0.925
ITSD + Bicon	0.908	0.029	0.952	0.838	0.038	0.905	0.774*	0.053	0.874*	0.831	0.064	0.857	0.920	0.035*	0.926
MINet ₂₀ [11]	0.916	0.026†	0.956*	0.838	0.035†	0.903	0.762	0.053	0.870	0.830	0.064	0.858	0.926	0.035*	0.924
$\mathrm{MINet} + \mathrm{Bicon}$	0.923 [†]	0.028*	0.957 [†]	0.856 [†]	0.035 [†]	0.915 [†]	0.778 [†]	0.051*	0.875 [†]	0.846 [†]	0.061 [†]	0.868 [†]	0.933 [†]	0.036	0.929*

Qualitative Evaluation. Representative examples of our qualitative analyses are shown in Fig. 7. Compared to baselines, our model can predict sharper boundaries and uniformly highlight salient regions in various challenging scenarios, including small objects (rows 4 and 7), complex background (rows 1, 3, 9 and 10) and foreground (rows 2 and 11), multiple objects (rows 5, 8 and 10), and interfering objects in the background (row 13).

4.4. Ablation Study

375

In this section, we study the effectiveness of different components of our ³⁸⁰ model. The experiments in this section were trained on the DUT-TR dataset and tested on DUT-OMRON and HKU-IS. For a fair comparison, all experiments use GCPANet [26] as backbone. The overall ablation study results are listed in Table 3.



Figure 7: Visual comparisons of different models.

- Connectivity modeling. We explore the role of the connectivity prediction strategy using two experiments. First, we used the original GCPANet as our baseline, denoted as *Base* (Exp. 1). Then, we replaced its output layers with 8-channel connectivity prediction layers and used connectivity masks instead of the saliency masks as our ground truth. We denote this connectivity version of the baseline as *Conn*. For the loss function, we used the multi-channel BCE loss L_{conmap} for the output Conn map *C*. This second experiment, denoted as Exp. 2 in Table 3, is very similar with ConnNet proposed in [20]. We used channel averaging at testing to get the single-channel saliency maps for evaluation. As seen in Table 3, the results did not improve compared to Exp. 1, which follows our key hypothesis that completely replacing saliency modeling
- ³⁹⁵ with connectivity modeling is not sufficient for modeling the saliency region.

Bilateral voting mechanism. Next, we studied the proposed BV module,

Exp		Model							HKU-IS		DUT-OMRON		
	Base	Conn	L_{conmap}	BV	L_{global_bce}	$L_{decouple}$ (RCA)	L _{con-const}	$F_{ave} \uparrow$	MAE↓	E_m \uparrow	$F_{ave} \uparrow$	MAE↓	E_m \uparrow
1	\checkmark							0.896	0.032	0.950	0.743	0.056	0.856
2		\checkmark	\checkmark					0.899	0.033	0.949	0.738	0.058	0.854
3		\checkmark	\checkmark	\checkmark	\checkmark			0.911	0.031	0.951	0.750	0.057	0.853
4		\checkmark	\checkmark	\checkmark		\checkmark		0.916	0.033	0.951	0.760	0.057	0.860
5		\checkmark		\checkmark		\checkmark	\checkmark	0.918	0.032	0.954	0.762	0.055	0.863

Table 3: Ablation study on HKU-IS and DUT-OMRON datasets

which is important both at training and testing phases. The BV module helps the training in two ways: first, it provides an enhanced connectivity map \tilde{C} for the RCA module; second, in the connectivity consistency loss term, it generates

- the input for L_{bimap} , which is a position-aware loss. To simplify the experiment and avoid interference, we tested only the first part in this subsection. Based on *Conn*, we first conducted the bilateral voting on the output Conn map *C* and got the Bicon map \widetilde{C} . Then, we computed the global map \widetilde{S}_{global} by averaging among channels of \widetilde{C} . For the loss term, we calculated the BCE loss on both the
- global map $(L_{global,bce})$ and the Conn map (L_{conmap}) . This process is shown as Exp. 3 of Table 3. As seen, inclusion of the BV module improved the F_{ave} , indicating that the BV module can enhance the spatial consistency of the output predictions.

To test the effectiveness of the BV module at the testing phase, based on ⁴¹⁰ Exp. 3, we tested the output both with and without the BV module. As seen in Table 4 and in Fig. 8, all three metrics have been improved after we applied the BV module to the testing phase.

	1								
Test Method		nku-15		DUI-OMRON					
	$\mathbf{F}_{ave}\uparrow$	$\mathrm{MAE}{\downarrow}$	$E_m \uparrow$	$\mathbf{F}_{ave}\uparrow$	$\mathrm{MAE}{\downarrow}$	$E_m \uparrow$			
Without BV	0.889	0.033	0.945	0.732	0.061	0.849			
With BV	0.911	0.031	0.951	0.750	0.057	0.853			

Table 4: Different testing methods based on Exp. 3.



Figure 8: Visualization of outputs at different stages of BiconNet. As representative examples, for the Conn map C and Bicon map \tilde{C} , we show only the first channel. The predicted Bicon map has much higher spatial coherence than the Conn map.

The edge decoupling mechanism. In this subsection, we study the proposed edge decoupling mechanism, which is the key innovation of the RCA ⁴¹⁵ module. Based on Exp. 3, we computed the edge-decoupled map $\tilde{S}_{decouple}$ from Bicon map \tilde{C} via the RCA module and replaced the loss with $L_{decouple}$ and L_{conmap} , respectively. This experiment is denoted as Exp. 4 in Table 3. As seen, the F_{ave} and E_m values increased. This result shows that the RCA module effectively utilized the extracted edge features.

The connectivity consistency loss. To test the effectiveness of the connectivity consistency loss, we replaced L_{conmap} with $0.8 \times L_{conmap} + 0.2 \times L_{bimap}$; i.e., the L_{con_const} in Exp. 4. Thus, the total loss function for this experiment is $L_{con_const} + L_{decouple}$. For this complete BiconNet model with backbone GC-PANet (Exp. 5 in Table 3), all three metrics improved, which demonstrates the ability of the connectivity consistency loss to improve the results.

Additionally, to illustrate the different effects of L_{conmap} and L_{bimap} in L_{con_const} , we conducted another set of experiments based on Exp. 5, using different weights for these two terms. The results are shown in Fig. 9, where 10 experiments are plotted with ω_2 as the x-axis ($\omega_1 = 1 - \omega_2$). When we introduced L_{bimap} and gradually increased its weight ω_2 (from left to right), we observed that F_{ave} and E_m increased while MAE decreased at the beginning ($\omega_2 \leq 0.2$). Then, when L_{bimap} had a larger weight, the overall performance decreased. The best performance was achieved at $\omega_1 = 0.8, \omega_2 = 0.2$. This result is consistent with our assumption that there is a tradeoff between edge enhancement and background dilution when using L_{bimap} . We also visualized the two loss terms L_{bimap} and L_{conmap} in Fig. 10 to further demonstrate this idea.



Figure 9: Training the network with different ω_1 and ω_2 on the (a) HKU-IS and (b) DUT-OMRON datasets. The x-axis represents the value for ω_2 ($\omega_1 = 1 - \omega_2$). The best performance was achieved at $\omega_2 = 0.2$ (dashed red line).



Figure 10: Comparison between L_{bimap} and L_{conmap} . L_{bimap} can generate much larger loss on hard pixels such as the edges of different objects while generating smaller loss on background pixels.

4.5. Model Size and Testing Speed

We list the model size and testing speed of our method in Table 5. To make fair comparisons, we reported the frame per second (FPS) processing speed with images of size 320×320 pixels for all listed models. Since we only changed the output fully-connected (FC) layers of the backbones, the increase in the parameters and testing time is directly related to the number of FC layers in the backbones. Thus, for those models with deep supervision mechanisms

(multiple FC layers, such as GCPANet and EGNet), the increase is more than those using shallow supervisions. However, even for the deep supervised models, the increase of model size is still insignificant and the computational cost of our BiconNet is nearly identical to those of the baselines. Therefore, our method is efficient and can improve existing frameworks with neglectable increase in computational cost.

Table 5: Comparison of model size and testing speed between BiconNet and the corresponding baselines.

	PoolNet		CPD-R GO		GCP	GCPANet I		F3Net EG		EGNet		ITSD		MINet	
	Base	Bicon	Base	Bicon	Base	Bicon	Base	Bicon	Base	Bicon	Base	Bicon	Base	Bicon	
Number of Parameters(M)	68.26	68.24	47.85	47.85	111.69	111.85	25.54	25.56	67.06	67.12	26.47	26.47	115.69	115.69	
Testing Speed (FPS)	49	49	55	53	38	34	64	63	60	53	47	44	45	43	

4.6. Compatibility Analysis

455

In section 4.3, we showed that BiconNet is compatible with existing SOD frameworks in their *entirety*. Here, we investigate the compatibility of BiconNet with *individual modules* that have a similar function (i.e., enhancement of spatial coherence and edge modeling), such as inter-pixel consistency/edge-aware loss functions and CRF.

4.6.1. With Pixel Relationship/Edge-Aware Loss Functions

As illustrative examples to show the compatibility of BiconNet with the state-of-the-art loss functions, we considered two loss functions here: the Consistency enhanced Loss (CEL) [11] (which can enhance the inter-pixel relationship) and Adaptive ConTour (ACT) [14] (which can improve the edge modeling). In each case, we compared the baselines and BiconNets with and without the loss functions (Table 6). Again, the networks that included the BiconNet outperformed the baselines. We also note that the performance of ITSD didn't significantly improve when added with ACT, while the combination of ACT and BiconNet had a more pronounced positive impact on performance.

4.6.2. With CRF

CRF is a widely used post-processing method that can enhance the interpixel relationship of the prediction and has been applied in SOD [15, 9, 16]. ⁴⁷⁰ Using GCPANet as our baseline, we added a fully connected CRF at the end of both GCPANet and GCPANet + Bicon for testing (Table 7). The results of GCPANet + CRF show that F_{ave} and MAE both improved while E_m decreased compared to GCPANet.Similar results were observed in GCPANet + Bicon + CRF. Nonetheless, GCPANet + Bicon + CRF outperformed GCPANet + CRF, ⁴⁷⁵ suggesting that BiconNet is compatible with CRF.

When added to any model, CRF usually significantly increases the computational cost. However, the results show that our model (GCPANet + Bicon) can achieve comparable results with GCPANet + CRF (the 2nd and 3rd rows in Table 7) without significantly compromising speed.

Table 0. Compatibility analysis with different loss funct									
Model	I	OUT-TE		DUT-OMRON					
	$\mathbf{F}_{ave}\uparrow$	MAE↓	$E_m \uparrow$	$\mathbf{F}_{ave}\uparrow$	MAE↓	$E_m \uparrow$			
ITSD w/o ACT	0.805	0.041	0.898	0.750	0.059	0.862			
+Bicon	0.830	0.041	0.902	0.763	0.059	0.865			
ITSD w/ ACT	0.806	0.041	0.891	0.752	0.058	0.862			
+Bicon	0.838	0.038	0.905	0.774	0.053	0.874			
$\mathrm{MINet}~\mathbf{w/o}~\mathrm{CEL}$	0.801	0.036	0.901	0.749	0.053	0.868			
+Bicon	0.846	0.037	0.910	0.766	0.053	0.870			
MINet $\mathbf{w}/$ CEL	0.838	0.035	0.903	0.762	0.053	0.870			
+Bicon	0.856	0.035	0.915	0.778	0.051	0.875			

Table 6: Compatibility analysis with different loss functions.

	Table 7. Compatibility and testing speed analysis with Citr.									
	Model	I	DUT-TE		DUT-OMRON					
		$\mathbf{F}_{ave}\uparrow$	$\mathrm{MAE}{\downarrow}$	$E_m \uparrow$	$\mathbf{F}_{ave}\uparrow$	$\mathrm{MAE}{\downarrow}$	$E_m \uparrow$			
	GCPANet	0.896	0.032	0.950	0.743	0.056	0.856			
	GCPANet + Bicon	0.918	0.032	0.954	0.762	0.055	0.863			
	GCPANet + CRF	0.920	0.029	0.947	0.763	0.053	0.840			
G	CPANet + CRF + Bicon	0.928	0.029	0.950	0.775	0.051	0.856			

Table 7: Compatibility and testing speed analysis with CRF.

480 5. Conclusion

In this study, we examined the spatial inconsistency and blurred edge issues of general salient object detection methods. To overcome these problems, we proposed a connectivity-based approach called BiconNet. We first showed that the connectivity mask is a more structure- and inter-pixel relation-aware label than a single-channel saliency mask. To utilize this informative label, we proposed a BV module to enhance the spatial consistency of the output and an RCA module to extract the edge features. Then, we trained the model with a novel Bicon loss. Extensive experiments demonstrated the advantages of our method over state-of-the-art algorithms. Finally, we demonstrated the efficiency of our model as it can improve existing SOD frameworks with a neglectable increase in computational cost.

Although this work demonstrated significant advances in the field of SOD, there are still properties of the connectivity mask worth exploiting in future work. For example, a weakness of our approach is that we only considered the ⁴⁹⁵ inter-class connectivity for the single-class segmentation problem. When dealing with the multi-class segmentation task, our method is expected to further benefit from modeling the intra-class relationship between connectivity masks. We envision that our connectivity-based approach to the image segmentation problem can be adopted by us and others in a wide range of applications, includion example, and the intra-class relation task of a point of the single-class segmentation. images.

Acknowledgments

This work was supported by the National Institutes of Health (R01 EY030124 and P30 EY005722), Hartwell Foundation Postdoctoral Fellowship, and Re-⁵⁰⁵ search to Prevent Blindness Unrestricted Grant to Duke University.

References

- A. Borji, M.-M. Cheng, H. Jiang, J. Li, Salient object detection: A benchmark, IEEE Trans. Image Process 24 (12) (2015) 5706-5722. doi: 10.1109/TIP.2015.2487833.
- [2] W. Shimoda, K. Yanai, Distinct class-specific saliency maps for weakly supervised semantic segmentation, in: Proceedings of European Conference on Computer Vision, Springer, 2016, pp. 218–234. doi:https://doi.org/ 10.1007/978-3-319-46493-0_14.
 - [3] W. Luo, M. Yang, W. Zheng, Weakly-supervised semantic segmentation with saliency and incremental supervision updating, Pattern Recognit. 115 (2021) 107858. doi:https://doi.org/10.1016/j.patcog.2021.107858.
 - [4] P. Zhang, W. Liu, D. Wang, Y. Lei, H. Wang, H. Lu, Non-rigid object tracking via deep multi-scale spatial-temporal discriminative saliency maps, Pattern Recognit. 100 (2020) 107130. doi:https://doi.org/10.1016/j. patcog.2019.107130.

520

515

- [5] L. Itti, C. Koch, E. Niebur, A model of saliency-based visual attention for rapid scene analysis, IEEE Trans. Pattern Anal. Mach. Intell. 20 (11) (1998) 1254–1259. doi:10.1109/34.730558.
- [6] B. Schauerte, B. Kühn, K. Kroschel, R. Stiefelhagen, Multimodal saliencybased attention for object-based scene analysis, in: Proceedings of
- 525

IEEE/RSJ International Conference on Intell. Robots and Systems, IEEE, 2011, pp. 1173–1179. doi:10.1109/IROS.2011.6095124.

- Y. Fang, C. Zhang, X. Min, H. Huang, Y. Yi, G. Zhai, C.-W. Lin, Devsnet: Deep video saliency network using short-term and long-term cues, Pattern Recognit. 103 (2020) 107294. doi:https://doi.org/10.1016/j.patcog. 2020.107294.
- [8] E. Ahn, L. Bi, Y. H. Jung, J. Kim, C. Li, M. Fulham, D. D. Feng, Automated saliency-based lesion segmentation in dermoscopic images, in: Proceedings of Annual International Conference of the IEEE Engineering in Medicine and Biology Society, IEEE, 2015, pp. 3009–3012. doi: 10.1109/EMBC.2015.7319025.
- Q. Hou, M.-M. Cheng, X. Hu, A. Borji, Z. Tu, P. H. S. Torr, Deeply supervised salient object detection with short connections, IEEE Trans. Pattern Anal. Mach. Intell. 41 (4) (2018) 815–828. doi:10.1109/TPAMI. 2018.2815688.
- [10] X. Qin, Z. Zhang, C. Huang, M. Dehghan, O. R. Zaiane, M. Jagersand, U2-net: Going deeper with nested u-structure for salient object detection, Pattern Recognit. 106 (2020) 107404. doi:https://doi.org/10.1016/j. patcog.2020.107404.
- ⁵⁴⁵ [11] Y. Pang, X. Zhao, L. Zhang, H. Lu, Multi-scale interactive network for salient object detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2020, pp. 9413–9422. doi:10. 1109/CVPR42600.2020.00943.
- [12] J.-X. Zhao, J.-J. Liu, D.-P. Fan, Y. Cao, J. Yang, M.-M. Cheng, Egnet:
 Edge guidance network for salient object detection, in: Proceedings of the IEEE International Conference on Computer Vision, IEEE, 2019, pp. 8779– 8788.

530

535

540

- [13] Z. Wu, L. Su, Q. Huang, Stacked cross refinement network for edge-aware salient object detection, in: Proceedings of IEEE Conference on Computer
- 555 Vision and Pattern Recognition, IEEE, 2019, pp. 7264–7273. doi:10. 1109/ICCV.2019.00736.
 - [14] H. Zhou, X. Xie, J.-H. Lai, Z. Chen, L. Yang, Interactive two-stream decoder for accurate and fast saliency detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2020, pp. 9141–9150. doi:10.1109/CVPR42600.2020.00916.
 - [15] N. Liu, J. Han, M. Yang, Picanet: Learning pixel-wise contextual attention for saliency detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2018, pp. 3089–3098. doi:10. 1109/CVPR.2018.00326.
- 565 [16] G. Li, Y. Yu, Deep contrast learning for salient object detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2016, pp. 478–487. doi:10.1109/CVPR.2016.58.
 - [17] Z. Luo, A. Mishra, A. Achkar, J. Eichel, S. Li, P. Jodoin, Non-local deep features for salient object detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2017, pp. 6593–6601. doi:10.1109/CVPR.2017.698.
 - [18] J. Wei, S. Wang, Q. Huang, F³net: Fusion, feedback and focus for salient object detection, in: Proceedings of Association for the Advancement of Artificial Intelligence, Vol. 34, 2020, pp. 12321–12328. doi:10.1609/aaai. v34i07.6916.

575

570

560

[19] J. Wei, S. Wang, Z. Wu, C. Su, Q. Huang, Q. Tian, Label decoupling framework for salient object detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2020, pp. 13025– 13034. doi:10.1109/CVPR42600.2020.01304.

- [20] M. Kampffmeyer, N. Dong, X. Liang, Y. Zhang, E. P. Xing, Connnet: A long-range relation-aware pixel-connectivity network for salient segmentation, IEEE Trans. Image Process 28 (5) (2018) 2518-2529. doi: 10.1109/TIP.2018.2886997.
- [21] J.-J. Liu, Q. Hou, M.-M. Cheng, J. Feng, J. Jiang, A simple pooling-based design for real-time salient object detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2019, pp. 3917–3926. doi:10.1109/CVPR.2019.00404.
 - [22] M.-M. Cheng, J. Warrell, W.-Y. Lin, S. Zheng, V. Vineet, N. Crook, Efficient salient region detection with soft image abstraction, in: Proceedings of IEEE International Conference on Computer Vision, IEEE, 2013, pp. 1529–1536. doi:10.1109/ICCV.2013.193.
 - [23] S. S. Naqvi, W. N. Browne, C. Hollitt, Salient object detection via spectral matting, Pattern Recognit. 51 (2016) 209-224. doi:https://doi.org/10. 1016/j.patcog.2015.09.026.
- ⁵⁹⁵ [24] J. Harel, C. Koch, P. Perona, Graph-based visual saliency, in: Advances in Neural Information Processing Systems, Curran Associates, Inc., 2007, pp. 545-552. doi:https://doi.org/10.7551/mitpress/7503.001.0001.
 - [25] L. Wang, H. Lu, X. Ruan, M.-H. Yang, Deep networks for saliency detection via local estimation and global search, in: Proceedings of IEEE Conference
- on Computer Vision and Pattern Recognition, IEEE, 2015, pp. 3183–3192.
 doi:10.1109/CVPR.2015.7298938.
 - [26] Z. Chen, Q. Xu, R. Cong, Q. Huang, Global context-aware progressive aggregation network for salient object detection, in: Proceedings of Association for the Advancement of Artificial Intelligence, Vol. 34, 2020, pp. 10599–10606. doi:10.1609/aaai.v34i07.6633.
- 605

590

[27] G. Li, Y. Yu, Visual saliency based on multiscale deep features, in: Pro-

ceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2015, pp. 5455–5463. doi:10.1109/CVPR.2015.7299184.

[28] X. Qin, Z. Zhang, C. Huang, C. Gao, M. Dehghan, M. Jagersand, Basnet: Boundary-aware salient object detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2019, pp. 7479–7489. doi:10.1109/CVPR.2019.00766.

610

615

620

- [29] Q. Zhang, Y. Shi, X. Zhang, Attention and boundary guided salient object detection, Pattern Recognit. 107 (2020) 107484. doi:https://doi.org/ 10.1016/j.patcog.2020.107484.
- [30] R. C. Gonzalez, R. E. Woods, Digital Image Processing (3rd Edition), Prentice-Hall, Inc., 2006.
- [31] P.-T. De Boer, D. P. Kroese, S. Mannor, R. Y. Rubinstein, A tutorial on the cross-entropy method, Annals of Operations Res. 134 (1) (2005) 19–67. doi:https://doi.org/10.1007/s10479-005-5724-z.
- [32] K. Zhao, S. Gao, W. Wang, M.-M. Cheng, Optimizing the f-measure for threshold-free salient object detection, in: Proceedings of the IEEE International Conference on Computer Vision, IEEE, 2019, pp. 8849–8857. doi:10.1109/ICCV.2019.00894.
- [33] L. Wang, H. Lu, Y. Wang, M. Feng, D. Wang, B. Yin, X. Ruan, Learning to detect salient objects with image-level supervision, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2017, pp. 136–145. doi:10.1109/CVPR.2017.404.
- [34] Q. Yan, L. Xu, J. Shi, J. Jia, Hierarchical saliency detection, in: Proceed ⁶³⁰ ings of IEEE Conference on Computer Vision and Pattern Recognition,
 IEEE, 2013, pp. 1155–1162. doi:10.1109/CVPR.2013.153.
 - [35] Y. Li, X. Hou, C. Koch, J. M. Rehg, A. L. Yuille, The secrets of salient object segmentation, in: Proceedings of IEEE Conference on Computer

30

Vision and Pattern Recognition, IEEE, 2014, pp. 280–287. doi:10.1109/CVPR.2014.43.

- [36] C. Yang, L. Zhang, H. Lu, X. Ruan, M.-H. Yang, Saliency detection via graph-based manifold ranking, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2013, pp. 3166–3173. doi:10.1109/CVPR.2013.407.
- ⁶⁴⁰ [37] F. Perazzi, P. Krähenbühl, Y. Pritch, A. Hornung, Saliency filters: Contrast based filtering for salient region detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2012, pp. 733–740. doi:10.1109/CVPR.2012.6247743.
- [38] R. Achanta, S. Hemami, F. Estrada, S. Susstrunk, Frequency-tuned salient
 region detection, in: Proceedings of IEEE Conference on Computer Vision
 and Pattern Recognition, IEEE, 2009, pp. 1597–1604. doi:10.1109/CVPR.
 2009.5206596.
 - [39] D.-P. Fan, C. Gong, Y. Cao, B. Ren, M.-M. Cheng, A. Borji, Enhancedalignment measure for binary foreground map evaluation, in: Proceedings of International Joint Conference on Artificial Intelligence, Elsevier, 2018, p. 698–704. doi:10.24963/ijcai.2018/97.
 - [40] Z. Wu, L. Su, Q. Huang, Cascaded partial decoder for fast and accurate salient object detection, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2019, pp. 3907–3916. doi:10. 1109/CVPR.2019.00403.
- 655

[41] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, Pytorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems, Curran Associates, Inc., 2019, pp. 8026–8037. doi:arXiv:1912.01703.

635

650

660