

“©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

RGB-D Saliency Detection via Cascaded Mutual Information Minimization

Jing Zhang¹ Deng-Ping Fan^{2,*} Yuchao Dai³ Xin Yu⁴ Yiran Zhong⁵ Nick Barnes¹ Ling Shao²
¹ Australian National University ² IIAI ³ Northwestern Polytechnical University
⁴ University of Technology Sydney ⁵ SenseTime

Abstract

Existing RGB-D saliency detection models do not explicitly encourage RGB and depth to achieve effective multi-modal learning. In this paper, we introduce a novel multi-stage cascaded learning framework via mutual information minimization to **explicitly** model the multi-modal information between RGB image and depth data. Specifically, we first map the feature of each mode to a lower dimensional feature vector, and adopt mutual information minimization as a regularizer to reduce the redundancy between appearance features from RGB and geometric features from depth. We then perform multi-stage cascaded learning to impose the mutual information minimization constraint at every stage of the network. Extensive experiments on benchmark RGB-D saliency datasets illustrate the effectiveness of our framework. Further, to prosper the development of this field, we contribute the largest ($7\times$ larger than NJU2K) COME15K dataset, which contains 15,625 image pairs with high quality polygon-/scribble-/object-/instance-/rank-level annotations. Based on these rich labels, we additionally construct four new benchmarks with strong baselines and observe some interesting phenomena, which can motivate future model design. Source code and dataset are available at https://github.com/JingZhang617/cascaded_rgbd_sod.

1. Introduction

Saliency detection models are trained to discover the regions of an image that attract human attention. Conventionally, saliency detection is performed mostly on RGB images only [37, 44, 22, 43, 36]. With the availability of depth data as shown in Table 1, RGB-D saliency detection [46, 35, 51, 53] attracts great attention. The extra depth data provides real-world geometric information, which is useful for scenarios when the foreground shares similar appearance to the background. Further, the robustness of depth sensors (e.g. Microsoft Kinect) against lighting changes can also benefit the saliency detection task.

As RGB and depth data capture different information




Figure 1. Comparison of saliency prediction of a state-of-the-art RGB-D saliency detection model, e.g. BBSNet [12], with ours.

about the same scene, existing RGB-D saliency detection models [35, 1, 3, 2, 53, 49, 30, 12, 19, 33, 51] focus on modeling the complementary information of the RGB image and depth data implicitly by using different fusion strategies. Three main fusion strategies have been widely studied: early fusion [38, 46], late fusion [41, 16, 36] and cross-level fusion [35, 1, 3, 2, 53, 49, 30, 12, 19, 33, 51, 25]. Although performance improvement can be achieved with effective fusion strategies, there are no constraints on the network design that force it to learn complementary information between the two modalities, and we cannot explicitly evaluate the contribution of depth data in those models [52].

As a multi-modal learning task, a trained model should maximize the joint entropy of different modalities within the network capacity. Maximizing the joint entropy is also equal to the minimization of mutual information, which prevents a network from focusing on redundant information. To *explicitly* model the complementary information between the RGB image and depth data, we introduce a multi-stage cascaded learning framework via mutual information minimization. Specifically, we introduce mutual information minimization as regularizer (as shown in Fig. 2) to achieve two main benefits: 1) explicitly modeling the redundancy between appearance features and geometric features; 2) effectively fusing appearance features and geometric features with the mutual information minimization constraint. The produced saliency maps in Fig. 1 illustrate effectiveness of our solution.

Furthermore, we find that there is no large-scale RGB-D saliency detection training set. In Table 1 we compare the widely used RGB-D saliency datasets, in terms of the size,

Table 1. Comparison with the widely used RGB-D datasets.

Dataset	Size	Type	Depth Source	Role
NJU2K[21]	1,985	Movie/Internet	FujiW3 camera + optical flow	Tr, Te
DUT [35]	1,200	Indoor/Outdoor	Light-field cameras	Tr, Te
NLPR [34]	1,000	Indoor/Outdoor	Microsoft Kinect	Tr, Te
SSB [32]	1,000	Internet	Stereo cameras	Te
SIP [11]	929	Person in outside	Huawei Mate10	Te
DES [7]	135	Indoor	Microsoft Kinect	Te
LFSD [27]	80	Indoor/Outdoor	Lytro Illum cameras	Te
Ours	15,625	Indoor/Outdoor	Holopix Social Platform	Tr, Te

types of data, the sources of depth data, and their roles (for training “Tr” or for testing “Te”) in RGB-D saliency detection. We note that the conventional training set for RGB-D saliency detection is a combination of samples from the NJU2K [21] dataset and NLPR [34], which includes only 2,200 image pairs in total. Although another 800 training images from the DUT dataset [35] can serve as the third part of the training set, the total number of training images is 3,000, which is not big enough, and may lead to biased model. In addition, we observe that there are similar backgrounds in the existing RGB-D saliency training set, *e.g.* more than 10% of the training dataset comes from the same scene with similar illumination conditions. The lack of diversity in the dataset may render models of poor generalization ability. At the same time, we also notice that the largest testing set [11] contains only 1,000 image pairs, which may not be enough to fully evaluate the overall performance of the deep RGB-D saliency detection models.

To provide an RGB-D saliency detection dataset for robust model training, and a sufficient size of testing data for model evaluation, we contribute the largest RGB-D saliency detection dataset, relabeled from Holo50K dataset [18], with 8,025 image pairs for training and 7,600 image pairs for testing. We provide not only binary annotations, but also annotations for stereoscopic saliency detection, scribble and polygon annotations for weakly supervised RGB-D saliency detection, instance-level RGB-D saliency annotations and RGB-D saliency ranking. Moreover, we contribute 5,000 unlabeled training images for semi-supervised or self-supervised RGB-D saliency detection.

Our main contributions are: 1) We design a multi-stage cascaded learning framework via mutual information minimization for RGB-D saliency detection to “explicitly” model redundancy between the RGB image and depth data. 2) The mutual information minimization regularizer can be easily extend to other multi-modal learning pipelines to model the redundancy of multiple modalities. 3) We contribute the largest RGB-D saliency detection dataset, with a 15,625 labeled set and a 5,000 unlabeled set to achieve fully-/weakly-/un-supervised RGB-D saliency detection. 4) We present new benchmarks for RGB-D saliency detection, and introduce baseline models for stereoscopic and weakly supervised RGB-D saliency detection.

2. Related Work

2.1. RGB-D saliency detection models

For RGB-D saliency detection, one of the main focuses is to explore the complementary information between the RGB image and the depth data. The former provides appearance information of a scene, while the latter introduces geometric information. Depending on how information from these two modalities is fused, existing RGB-D saliency detection models can be divided into three categories: early-fusion models [38, 46], late-fusion models [41, 16, 36] and cross-level fusion models [35, 1, 3, 2, 53, 49, 30, 12, 19, 33, 51, 25, 26]. The first solution directly concatenates the RGB image with its depth, while the late fusion models treat each mode (RGB and depth) separately, and then fusion is achieved at the output layer. The above two solutions perform multi-modal fusion at the input or output layer, while the cross-level fusion models fuse RGB and depth in the feature space. Specifically, features of an RGB image and depth are gradually fused in different level of the network [33, 26, 12, 25, 31, 5, 48, 30, 49]. Although the existing methods fuse the RGB image and depth data for multi-modal learning, none of them *explicitly* illustrate how the network achieve effective multi-modal learning. We propose a cross-level fusion model as shown in Fig. 2. By designing the “Mutual information regularizer”, we aim to reduce redundancy of appearance features and geometric features for effective multi-modal learning.

2.2. Multi-modal learning with RGB-D dataset

The basic assumption for multi-modal learning is that there exists both common and diverse information in the separate modalities. For the RGB-D dataset, the RGB image and depth data share similar semantic information, which can be defined as the common parts. The RGB image encodes the appearance information, including intensity or color of the object, while the depth data encodes geometric information, showing the relative geometric localization of the objects. The difference between appearance information and geometric information is the diverse part of these two modalities. The main focus to achieve multi-modal learning for RGB-D data is through using different fusion strategies [42, 4, 31, 24], *e.g.* early fusion, late fusion or cross-level fusion. Different from conventional solutions, we introduce a multi-stage cascaded learning framework via mutual information minimization to reduce the feature redundancy of each modal. Although mutual information maximization [29, 39] is widely used in representation learning to produce a representation that is similar to the input, we take the mutual information minimization as a regularizer to reduce the feature redundancy for effective multi-modal learning.

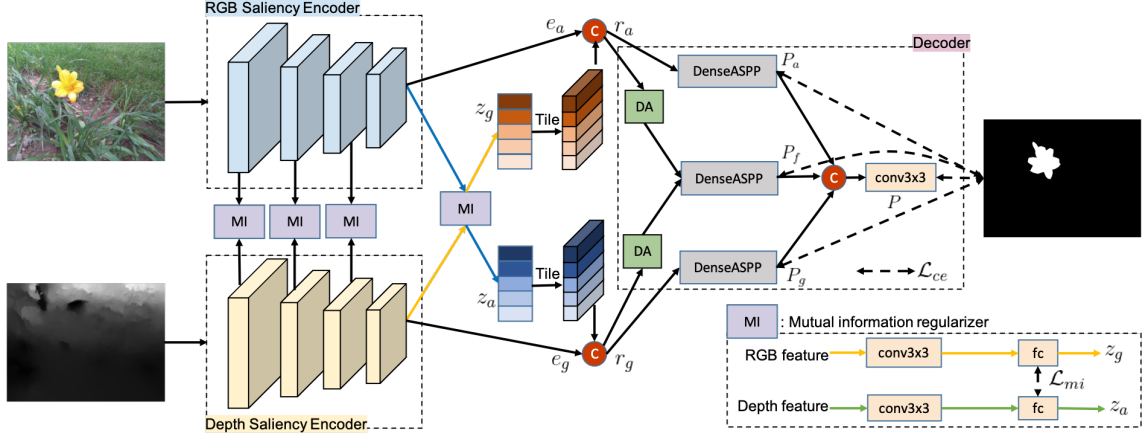


Figure 2. Overview of the proposed multi-stage cascaded learning framework for RGB-D saliency detection. We feed the RGB image and depth to the saliency encoders to extract saliency feature of each mode with the mutual information regularizer term to push the features to be different from each other. Then, we fuse the lower dimensional feature of each mode (z_a and z_g) with raw image feature (e_a and e_g) to effectively model the complementary information of each mode and obtain our final prediction P . The ‘‘DenseASPP’’ module is the dense atrous spatial pyramid pooling module from [45], and ‘‘DA’’ is the dual attention module from [14].

2.3. RGB-D saliency datasets

The widely used RGB-D saliency detection datasets include NJU2K [21], NLPR [34], SSB [32], DES [7], LFS [27], SIP [11], DUT [35], *etc.*, as shown in Table 1. The typical training dataset is the combination of 1,485 images from NJU2K [21] and 700 images from NLPR [34]. Piao *et al.* [35] introduces the DUT dataset, with 800 images for training and 400 images for testing. To prosper the RGB-D saliency detection task, we introduce the largest RGB-D saliency detection training and testing dataset, which will be introduced in Section 4.

3. Proposed CMINet

We introduce a multi-stage cascaded learning framework in Fig. 2 to explicitly model complementary information for RGB-D saliency detection.

3.1. Saliency encoder

We denote our training dataset as $T = \{X_i, Y_i\}_{i=1}^N$, where i indexes the images and N is the size of the training set, X_i and Y_i are the input RGB-D image pair and its corresponding ground-truth (GT) saliency map. We feed the training image pairs (RGB image I and depth D) to the saliency encoder, as illustrated in Fig. 2, to extract appearance features $f_{\alpha_a}(I)$ and geometric features $f_{\alpha_g}(D)$ respectively, where α_a and α_g are the parameters of our RGB saliency encoder and depth saliency encoder respectively.

We build the saliency encoders upon the ResNet50 network [17], which includes four convolutional stages $\{s^1, s^2, s^3, s^4\}$. We add one additional convolutional layer of kernel size 3×3 after each $s^c \in \{s^c\}_{c=1}^4$ to reduce the channel dimension of s^c to $C = 32$, and obtain feature maps

$\{e^1, e^2, e^3, e^4\}$. The final output of the RGB saliency encoder module is $e_a = \{e_a^1, e_a^2, e_a^3, e_a^4\}$, and that of the depth saliency encoder is $e_g = \{e_g^1, e_g^2, e_g^3, e_g^4\}$. Note that, the RGB saliency encoder and depth saliency encoder share the same network structure but not weights.

3.2. Feature embedding

Given the output $e_a = \{e_a^1, e_a^2, e_a^3, e_a^4\}$ from the RGB saliency encoder and $e_g = \{e_g^1, e_g^2, e_g^3, e_g^4\}$ from the depth saliency encoder, we aim to map both the RGB feature and depth feature to a lower-dimensional feature space for feature embedding. Specifically, we propose a multi-stage cascaded learning strategy to perform the complementary learning at each stage of the network. For the lower stages, we feed the RGB feature $\{e_a^c\}_{c=1}^3$ and the depth feature $\{e_g^c\}_{c=1}^3$ to two different 3×3 convolutional layers (‘‘conv3x3’’ in Fig. 2) to obtain feature maps of channel size $4 * C$ for both the RGB branch and depth branch. Then we adopt two fully connected layers (‘‘fc’’ in Fig. 2) to map the feature map of channel size $4 * C$ to two different lower-dimensional feature vectors $\{z_a^c\}_{c=1}^3$ and $\{z_g^c\}_{c=1}^3$ of size $K = 6$ respectively. The complementary learning related loss (which will be introduced in Section 3.3 and 3.5) is adopted to reduce the feature redundancy of RGB and depth at lower stages. At the highest stage, we first tile the lower dimensional feature vector z_a^4 and z_g^4 in spatial dimension. Then, we concatenate them with raw image feature¹ of the other mode to obtain the $4 * C + K$ channel size feature map r_a and r_g for the RGB branch and depth branch respectively.

¹We define the concatenation of $\{e_a^c\}_{c=1}^4$ as the raw RGB feature and the concatenation of $\{e_g^c\}_{c=1}^4$ as raw depth feature




Figure 3. Annotations of our new RGB-D saliency detection datasets: (a) the RGB image, (b) depth data and (c) binary ground truth, (d) instance level annotation, (e) ranking based annotation, (f) scribble annotation and (g) polygon annotation. Our diverse annotations will facilitate developing different fully/weakly supervised RGB-D saliency detection.

3.3. Multi-modal learning

After obtaining the feature embedding z_a and z_g for the RGB image and depth data, we introduce a mutual information minimization regularizer to explicitly reduce the redundancy between these two modalities. Our basic assumption is that a good appearance saliency feature and geometric saliency feature pair should carry both common parts (semantic related) and different attributes (domain related). Mutual information M_I is used to measure the difference between the entropy terms:

$$M_I(z_a, z_g) = H(z_a) + H(z_g) - H(z_a, z_g), \quad (1)$$

where $H(\cdot)$ is the entropy, $H(z_a)$ and $H(z_g)$ are marginal entropies, and $H(z_a, z_g)$ is the joint entropy of z_a and z_g . Intuitively, we have the Kullback–Leibler divergence (KL) of the two latent variable (or the conditional entropies) as:

$$KL(z_a||z_g) = H_{z_g}(z_a) - H(z_a), \quad (2)$$

$$KL(z_g||z_a) = H_{z_a}(z_g) - H(z_g), \quad (3)$$

where $H_{z_g}(z_a) = -\sum_x z_a(x) \log z_g(x)$ is the cross-entropy. We then sum Eq. 1, Eq. 2 and Eq. 3, and obtain:

$$M_I(z_a, z_g) = H_{z_g}(z_a) + H_{z_a}(z_g) - H(z_a, z_g) - (KL(z_a||z_g) + KL(z_g||z_a)). \quad (4)$$

Given the RGB image and depth data, $H(z_a, z_g)$ is non-negative, then minimizing the mutual information can be achieved by minimizing: $\mathcal{L}_{mi} = (H_{z_g}(z_a) + H_{z_a}(z_g)) - (KL(z_a||z_g) + KL(z_g||z_a))$. Intuitively, $M_I(z_a, z_g)$ measures the reduction of uncertainty in z_a when z_g is observed, or vice versa. As a multi-modal learning task, each mode should learn some new attributes of the task from other modalities. By minimizing $M_I(z_a, z_g)$, we can effectively explore the complementary attributes of both modalities. Note that, although KL loss term was used in [46] as distribution similarity measure, we use it to measure mode similarity for multi-modal learning.

3.4. Saliency decoder

With the mutual information as a regularizer, we achieve the feature redundancy constraint at the lower stages of the network, and obtain the refined RGB saliency feature r_a and refined depth saliency feature r_g at the highest stage. We then adopt one DenseASPP [45] module after r_a to obtain the RGB saliency prediction P_a with multi-scale context information. Similarly, we can obtain the depth saliency prediction P_g . The saliency decoder f_γ (“Decoder” in Fig. 2) takes the refined saliency features r_a, r_g , as well as the RGB saliency prediction P_a and depth saliency prediction P_g as input to produce our final prediction P , where γ is the parameter set of the saliency decoder. Specifically, we add a position attention module and a channel attention module [14] after r_a and r_g to obtain $da(r_a)$ and $da(r_g)$ respectively with discriminative features highlighted. Then we concatenate $da(r_a)$ and $da(r_g)$, and feed it to the DenseASPP [45] module to obtain our saliency prediction P_f . To further fuse information from both modes, we concatenate P_a, P_g and P_f channel-wise, and feed it to a 3×3 convolutional layer to achieve our final prediction P .

3.5. Objective function

We adopt the binary cross-entropy loss \mathcal{L}_{ce} as our objective function to train our multi-stage cascaded learning framework, where the complementary constraint, as indicated in Eq. (1), pushes the saliency feature distribution of the RGB image to be apart from that of the depth data. Our final objective function is:

$$\mathcal{L} = \mathcal{L}_{ce}(P, Y) + \lambda_1 \mathcal{L}_{ce}(P_f, Y) + \lambda_2 \mathcal{L}_{ce}(P_a, Y) + \lambda_3 \mathcal{L}_{ce}(P_g, Y) + \lambda \sum_{c=1}^4 \mathcal{L}_{mi}(z_a^c, z_g^c), \quad (5)$$

and empirically we set $\lambda_1 = 0.8, \lambda_2 = 0.6, \lambda_3 = 0.4$. As the range of the \mathcal{L}_{mi} is 10 times larger than that of the \mathcal{L}_{ce} , we set its loss weight as $\lambda = 0.1$ for balanced learning.

4. COME15K Dataset

As shown in Table 1, the existing RGB-D saliency detection training dataset is not big enough, which may lead to models of poor generalization ability. Further, as the training dataset is a combination of samples from NJU2K [21] and NLPR dataset [34], different splits of the training set often lead to inconsistent performance evaluation. Lastly, the small size of testing dataset may fail to fully evaluate the RGB-D saliency detection models. To boost RGB-D saliency detection, we contribute the largest RGB-D saliency detection dataset. We provide binary annotation, instance level annotation, ranking based annotation, weak annotation as shown in Fig. 3. The detailed analysis of the dataset is introduced in the supplementary material.

4.1. Dataset annotation

Our new COME15K dataset is based on Holo50K [18], which is a stereo dataset, including scenarios from both indoor and outdoor. We first filter² the Holo50K dataset and then obtain 16,000 stereo image pairs for labelling (the candidate labeled set) and another 5,000 image pairs as the unlabeled set. Note that the stereo pairs in Holo50K dataset are directly captured by a stereo camera without rectification, we use a modified version of a SOTA off-the-shelf stereo matching algorithm [55] to compute the depth for both the candidate labeled set and unlabeled set with the left-right view images as input.

To provide annotations for the candidate labeled set, we firstly ask five “coarse” annotators to label salient regions of each image (only the right view image is used) with scribble annotations. Secondly, the “fine” annotators will segment the full scopes of salient objects and provide instance-level annotations. Thirdly, we perform “majority voting” to obtain the binary GT saliency maps for our RGB-D saliency detection task. Note that, we delete those samples with no common salient regions, and obtain our final labeled dataset of size 15,625. Further, based on the scribble annotations and instance-level saliency maps, we rank each saliency instance according to the initial scribble annotations to form our RGB-D saliency ranking dataset.

We also provide weak annotations for weakly-supervised RGB-D saliency detection, including scribble annotations and polygon annotations. We define the majority of the scribble annotations from multiple coarse annotators as the scribble annotations of our dataset. Specifically, we first obtain the instance with the majority of scribble. Then, we define the scribble on the majority instance as our scribble annotation. We label the majority salient instance with polygons to form our polygon based annotations.

²We deleted the violent images.

4.2. Dataset splitting

We divide the labeled set into one training set with 8,025 samples and two different testing sets of size 4,600 and 3,000 respectively, namely the “Normal” and the “Difficult” sets. The training dataset is generated by randomly selecting 8,025 images from the labeled set. For the testing datasets, we intend to introduce two sets of different difficulties. Specifically, we rank the RGB images based on both global and interior contrast, and denote samples with low global contrast and high interior contrast as the difficult samples³. Then we have a pool of 1,800 difficult samples D_d and 5,800 normal samples D_n . We random select 30% samples from D_d and 70% samples from D_n to obtain our “Normal” testing set, and the others as “Difficult” set.

5. Experiments

We compare our method CMINet with existing RGB-D saliency detection models, and report the performance in Table 2 & 3. Furthermore, we retrain the state-of-the-art RGB-D saliency detection models on our new training dataset, and provide the performance of those models on our testing dataset in Table 6.

5.1. Setup

Dataset: For fair comparisons with existing RGB-D saliency detection models, we follow the conventional training setting, in which the training set is a combination of 1,485 images from the NJU2K dataset [21] and 700 images from the NLPR dataset [34]. We then test the performance of our model and competing models on the NJU2K testing set, NLPR, testing set LFS [27], DES [7], SSB [32] SIP [11] and DUT [35] testing set.

Metrics: We evaluate the performance of the models on four golden evaluation metrics, *i.e.*, Mean Absolute Error (\mathcal{M}), Mean F-measure (F_β), Mean E-measure (E_ξ) [10] and S-measure (S_α) [9], which are explained in detail in the supplementary materials.

Training details: Our model is implemented using the *Pytorch* library. The two saliency encoders share the same network structure, and are initialized with ResNet50 [17] trained on ImageNet, and the other newly added layers are randomly initialized. We resize all the images and ground truth to the same spatial size of 352×352 pixels. We set the maximum epoch as 100, and initial learning rate as $5e-5$. We adopt the “step” learning rate decay policy, and set the decay size as 80 and decay rate as 0.1. The whole training takes 4.5 hours with batch size 5 on an NVIDIA GeForce RTX 2080Ti GPU for the conventional training (NJU2K-train+NLPR-train) dataset, and 16 hours with our new training (COME15K-train) dataset.

³Details about image global and interior contrast are introduced in the supplementary materials.

Table 4. Performance of the extra experiments.

Method	NJU2K [21]				SSB [32]				DES [7]				NLPR [34]				LFSD [27]				SIP [11]			
	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\epsilon \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\epsilon \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\epsilon \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\epsilon \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\epsilon \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\epsilon \uparrow$	$\mathcal{M} \downarrow$
Base	.910	.900	.935	.035	.890	.870	.917	.043	.926	.915	.959	.018	.920	.898	.942	.024	.842	.835	.880	.077	.879	.876	.917	.049
K3	.928	.908	.947	.032	.909	.892	.939	.036	.934	.922	.964	.018	.925	.904	.956	.022	.869	.845	.898	.067	.885	.879	.919	.047
K32	.924	.909	.944	.033	.908	.894	.941	.036	.938	.923	.966	.017	.927	.906	.959	.021	.856	.853	.900	.065	.885	.878	.921	.046
SS	.926	.913	.943	.034	.914	.882	.942	.036	.946	.927	.968	.017	.932	.896	.954	.021	.861	.852	.896	.067	.885	.879	.925	.046
W0	.918	.907	.944	.033	.892	.877	.923	.042	.934	.924	.964	.017	.924	.900	.945	.023	.843	.836	.881	.076	.884	.878	.916	.048
W1	.919	.909	.946	.032	.905	.886	.937	.037	.938	.927	.971	.016	.923	.903	.956	.022	.857	.853	.891	.071	.887	.882	.921	.045
P_f	.925	.908	.945	.033	.908	.887	.939	.036	.946	.925	.965	.016	.938	.907	.962	.023	.862	.845	.896	.068	.889	.886	.927	.045
S_{rgb}	.898	.890	.930	.040	.899	.876	.924	.042	.891	.883	.920	.028	.908	.885	.932	.031	.817	.807	.853	.095	.860	.865	.905	.056
$S_{rgb\delta}$.915	.901	.932	.037	.903	.878	.931	.039	.920	.908	.942	.021	.914	.893	.943	.026	.850	.841	.886	.071	.876	.870	.910	.051
CMINet	.939	.925	.956	.032	.921	.895	.959	.034	.953	.926	.970	.015	.941	.909	.964	.019	.877	.860	.911	.064	.894	.887	.933	.044

performance of the proposed solution of using mutual information as regularizer for redundancy constraining.

The dimension of the feature space: We set the dimension of lower-dimensional feature embedding (z_a and z_g) as $K = 6$. To test the impact of feature dimensions on the network performance, we set $K = 3$ and $K = 32$, and report their performance as “K3” and “K32” respectively in Table 4. The experimental results demonstrates that our model achieves relative stable performance with different dimensions of the lower-dimensional feature, while the current dimension with $K = 6$ works the best.

The structure of the “Mutual information regularizer” module as shown in Fig. 2: As discussed in Section 3.2, the “Mutual information regularizer” module is composed of one 3×3 convolutional layers and one fully connected layer. One may also achieve this directly from the output of the saliency encoder. Specifically, we can feed the RGB feature and depth feature to two fully connected layers to obtain z_a and z_g respectively. In Table 4, we report the performance of our model with this simple setting, marked as “SS”. We observe performance decreases, which indicates the necessity of introducing more nonlinearity to effectively extract the feature representation of each mode.

The weight of the mutual information regularizer: The weight λ of the mutual information regularization term controls the level of complementary information. We set $\lambda = 0.1$ in this paper to achieve balanced training. We then test how the model performs with different λ , and set $\lambda = 0$ and $\lambda = 1$ respectively. We show the performance of those variants in Table 4, denoted by “W0” and “W1”. The inferior performance of “W0” indicates the effectiveness of our complementary information modeling strategy. Further, compared with our performance in Table 2, we observe relatively worse performance of “W1”, which inspires us to further investigate finding an optimal weight for the mutual information regularizer.

5.4. Discussion

The effectiveness of mutual information minimization as regularizer: We computed the mean absolute cosine similarities of the highest stage feature embeddings (z_a^4 and z_g^4) for “W0” from Table 4 (without mutual information minimization as regularizer) and ours, which are

$\text{cosine}(z_{a,g}^{M0}) = 0.90$ and $\text{cosine}(z_{a,g}^{Ours}) = 0.11$ on the NLPR testing dataset. This clearly shows the advantage of our solution in extracting a less correlated feature for each mode. We visualize the learned feature embedding in the supplementary materials.

The merging strategy: In this paper, we produce four different saliency maps as intermediate outputs, including the saliency prediction from both the RGB branch (P_a) and depth branch (P_g), the feature embedding fusion branch (P_f), and our final prediction P by fusing P_a , P_g and P_f . As P_f has already been included in the complementary information of z_a and z_g , we define our final prediction as P_f without the final fusion to obtain P . The performance is shown in Table 4 “ P_f ”. We observe inferior performance of “ P_f ” compared with our final prediction. The main reason is that z_a and z_g are high-level feature embeddings of the RGB image and depth data, which mainly capture the semantic information. The direct merging of z_a and z_g will generate saliency prediction with less structure accuracy.

The contribution of depth data: Saliency detection can be achieved merely with the RGB image. As discussed in Section 1, depth introduces useful geometric information for saliency detection. To verify this conclusion, we train our model (including only the encoder and decoder in Fig. 2) with and without depth as input⁵. The performance is shown in Table 4 as “ $S_{rgb\delta}$ ” and “ S_{rgb} ” respectively. The superior performance of “ $S_{rgb\delta}$ ” compared with “ S_{rgb} ” explains contribution of the depth data for saliency detection. We also show examples explaining how depth contributes to saliency detection in the supplementary materials.

Depth generation: We generate our COME15K dataset with Holo50K [18], where the stereo pairs are not strictly rectified, which may cause severe matching failures even with state-of-the-art stereo algorithms [6]. To solve this issue, we relax the horizontal search in stereo algorithms to both horizontal and vertical search but only keep the horizontal displacements as the stereo disparities. We use a modified stereo matching algorithm [40] to generate the disparities / depth in our dataset. Further, stereo cameras are widely used in mobile devices, which makes it easier to obtain depth information for both indoor and outdoor.

⁵The model with depth is an early-fusion model, where depth and RGB image are concatenated at input layer

Table 5. Performance of the weakly supervised saliency detection baselines.

Method	NJU2K[21]				SSB[32]				NLPR [34]				SIP [11]				COME15K-Normal				COME15K-Difficult			
	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$
Scribble	.823	.806	.869	.080	.820	.803	.884	.073	.820	.737	.863	.058	.815	.793	.888	.076	.802	.780	.856	.082	.767	.749	.812	.115
Polygon	.847	.827	.896	.065	.853	.831	.913	.056	.848	.789	.899	.043	.846	.822	.909	.060	.827	.805	.884	.065	.786	.774	.841	.096

Table 6. Performance on the test sets of our new COME15K.

Metric	UCNet	JLDCF	A2dele	DMRA	CPFP	S2MA	CoNet	BBS-Net	CMINet
	[46]	[15]	[36]	[35]	[53]	[30]	[19]	[12]	Ours
<i>Normal</i>									
$S_\alpha \uparrow$.894	.894	.833	.782	.795	.877	.820	.902	.915
$F_\beta \uparrow$.883	.875	.835	.744	.716	.829	.796	.879	.893
$E_\xi \uparrow$.929	.919	.882	.812	.801	.881	.850	.923	.941
$\mathcal{M} \downarrow$.036	.042	.060	.105	.104	.059	.082	.039	.033
<i>Difficult</i>									
$S_\alpha \uparrow$.822	.845	.787	.743	.770	.828	.779	.853	.867
$F_\beta \uparrow$.814	.832	.795	.724	.704	.789	.774	.834	.852
$E_\xi \uparrow$.859	.870	.838	.775	.776	.836	.813	.876	.893
$\mathcal{M} \downarrow$.079	.075	.092	.137	.131	.092	.113	.071	.064

5.5. New benchmarks on COME15K

We provide a new benchmark of state-of-the-art models trained with our new training dataset, and show performance in Table 6. Further, with our rich annotations as shown in Fig. 3, we discuss another three benchmarks for fully/weakly-supervised learning.

Benchmark #1: Re-train existing RGB-D saliency models with our new training dataset. We split our testing dataset to a moderate-level testing set (“Normal”) and a hard testing set (“Difficult”) with 4,600 image pairs and 3,000 pairs respectively. To test how existing RGB-D saliency detection models perform on our new testing sets, we re-train existing RGB-D saliency detection models with our new training set, and show their performance on the new testing sets in Table 6. The performance gap between those existing techniques illustrates effectiveness of our dataset in both model learning and evaluation.

Benchmark #2: Stereo saliency detection. As our RGB-D saliency dataset is constructed on a stereo dataset [18], we directly train a stereo image pair based saliency object detection model, where the depth is implicitly instead of explicitly obtained from the stereo image pairs. Although there exist some stereoscopic saliency detection models [23, 8, 32, 13], all of them take both the RGB image and depth as input. In this paper, similar to [50], we design a real⁶ stereoscopic saliency detection model, and provide a baseline to manifest the potential of our dataset for stereoscopic saliency detection. We train our stereoscopic saliency detection model on our new training dataset, where the left-right view images are taken as inputs and the GT saliency maps in the right view images are used as supervision. The same encoder and decoder as in Fig. 2 are adopted for our stereoscopic saliency detection model. Specifically, we implicitly model the geometric information by using a cost volume between the saliency encoder for the right and left view images. The performance is shown in Table 7.

⁶We define the stereoscopic saliency models taking only the left and right view images as input as the “real” stereoscopic saliency models.

Table 7. Performance of the stereo saliency detection baseline.

Metric	NJU2K[21]		NJU400[20]		COME15K-Normal		COME15K-Difficult		
	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$S_\alpha \uparrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	
	.874	.851	.056	.882	.851	.044	.874	.855	.047
	.825	.812	.080						

We explained the architecture and the other stereo saliency datasets in detail in the supplementary material.

Benchmark #3 and #4: Scribble/Polygon as supervision. For scribble supervision, we follow [47] and use the smoothness loss and an auxiliary edge detection branch as a constraint to maintain structure information in the prediction. We train our scribble supervised RGB-D saliency detection model by concatenating RGB and depth in the input layer, and feed the concatenated feature to one 3×3 convolutional layer to adapt the model in [47]. Performance of the scribble annotation based baseline model is shown in Table 5 “Scribble”. The polygon label is generated after majority voting. Fig. 3 (g) shows that the polygon label covers a larger area with better structure information than scribbles. We train directly with polygon annotations as pseudo labels by adopting our model in Fig. 2, and provide performance of this baseline model in Table 5 “Polygon”.

Benchmark analysis: Our RGB-D saliency benchmark in Table 6 shows the superior performance of our method. Furthermore, the gap between state-of-the-art models illustrates the effectiveness of our new testing dataset in model evaluation. Our stereoscopic saliency benchmark in Table 7 introduces another solution to implicitly use the geometric information. Our two weakly-supervised baselines in Table 5 provide new options for weakly-supervised RGB-D saliency detection.

6. Conclusion

We proposed a multi-stage cascaded learning based RGB-D saliency detection framework that explicitly models complementary information between RGB images and depth data. By minimizing the mutual information between these two modalities during training, our model focuses on the diverse parts of each mode rather than the redundant information. In this fashion, our model is able to exploit the multi-modal information more effectively. Further, we introduced the largest RGB-D saliency detection dataset with five types of annotations to prosper the development of fully-/weakly-/un-supervised RGB-D saliency detection tasks. Four new benchmarks on 7 datasets and our new dataset demonstrate the superiority of our model compared to the existing RGB-D saliency detection techniques.

References

- [1] Hao Chen and Youfu Li. Progressively complementarity-aware fusion network for RGB-D salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3051–3060, 2018. [1](#), [2](#), [6](#)
- [2] Hao Chen and Youfu Li. Three-stream attention-aware network for RGB-D salient object detection. *IEEE T. Image Process.*, pages 2825–2835, 2019. [1](#), [2](#), [6](#)
- [3] Hao Chen, Youfu Li, and Dan Su. Multi-modal fusion network with multi-scale multi-path and cross-modal interactions for RGB-D salient object detection. *Pattern Recognit.*, 86:376–385, 2019. [1](#), [2](#), [6](#)
- [4] Lin-Zhuo Chen, Zheng Lin, Ziqin Wang, Yong-Liang Yang, and Ming-Ming Cheng. Spatial information guided convolution for real-time rgbd semantic segmentation. *IEEE T. Image Process.*, 30:2313–2324, 2021. [2](#)
- [5] Shuhan Chen and Yun Fu. Progressively guided alternate refinement network for rgb-d salient object detection. In *Eur. Conf. Comput. Vis.*, 2020. [2](#)
- [6] Xuilian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Hongdong Li, Tom Drummond, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching. In *Adv. Neural Inform. Process. Syst.*, volume 33, pages 22158–22169, 2020. [7](#)
- [7] Yupeng Cheng, Huazhu Fu, Xingxing Wei, Jiangjian Xiao, and Xiaochun Cao. Depth enhanced saliency detection method. In *ACM ICIMCS*, pages 23–27, 2014. [2](#), [3](#), [5](#), [6](#), [7](#)
- [8] Runmin Cong, Jianjun Lei, Changqing Zhang, Qingming Huang, Xiaochun Cao, and Chunping Hou. Saliency detection for stereoscopic images based on depth confidence analysis and multiple cues fusion. *IEEE Signal Processing Letters*, 23(6):819–823, 2016. [8](#)
- [9] Deng-Ping Fan, Ming-Ming Cheng, Yun Liu, Tao Li, and Ali Borji. Structure-measure: A new way to evaluate foreground maps. In *Int. Conf. Comput. Vis.*, pages 4548–4557, 2017. [5](#)
- [10] Deng-Ping Fan, Ge-Peng Ji, Xuebin Qin, and Ming-Ming Cheng. Cognitive vision inspired object segmentation metric and loss function. *SCIENTIA SINICA Informationis*, 2021. [5](#), [6](#)
- [11] Deng-Ping Fan, Zheng Lin, Zhao Zhang, Menglong Zhu, and Ming-Ming Cheng. Rethinking RGB-D salient object detection: Models, datasets, and large-scale benchmarks. *IEEE T. Neural Netw. Learn. Syst.*, 2020. [2](#), [3](#), [5](#), [7](#), [8](#)
- [12] Deng-Ping Fan, Yingjie Zhai, Ali Borji, Jufeng Yang, and Ling Shao. Bbs-net: Rgb-d salient object detection with a bifurcated backbone strategy network. In *Eur. Conf. Comput. Vis.*, 2020. [1](#), [2](#), [6](#), [8](#)
- [13] Yuming Fang, Junle Wang, Manish Narwaria, Patrick Le Callet, and Weisi Lin. Saliency detection for stereoscopic images. *IEEE T. Image Process.*, 23:1–6, 11 2013. [8](#)
- [14] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3146–3154, 2019. [3](#), [4](#)
- [15] Keren Fu, Deng-Ping Fan, Ge-Peng Ji, and Qijun Zhao. Jldcf: Joint learning and densely-cooperative fusion framework for rgb-d salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. [6](#), [8](#)
- [16] Junwei Han, Hao Chen, Nian Liu, Chenggang Yan, and Xuelong Li. CNNs-based RGB-D saliency detection via cross-view transfer and multiview fusion. *IEEE T. Cybern.*, pages 3171–3183, 2018. [1](#), [2](#), [6](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. [3](#), [5](#)
- [18] Yiwen Hua, Puneet Kohli, Pritish Uplavikar, Anand Ravi, Saravana Gunaseelan, Jason Orozco, and Edward Li. Holopix50k: A large-scale in-the-wild stereo image dataset. In *CVPR Workshop on Computer Vision for Augmented and Virtual Reality*, 2020. [2](#), [5](#), [7](#), [8](#)
- [19] Wei Ji, Jingjing Li, Miao Zhang, Yongri Piao, and Huchuan Lu. Accurate rgb-d salient object detection via collaborative learning. In *Eur. Conf. Comput. Vis.*, 2020. [1](#), [2](#), [6](#), [8](#)
- [20] Ran Ju, Ling Ge, Wenjing Geng, Tongwei Ren, and Gangshan Wu. Depth saliency based on anisotropic center-surround difference. In *IEEE Int. Conf. Image Process.*, pages 1115–1119, 2014. [8](#)
- [21] Ran Ju, Yang Liu, Tongwei Ren, Ling Ge, and Gangshan Wu. Depth-aware salient object detection using anisotropic center-surround difference. *Signal Processing: Image Communication*, 38:115 – 126, 2015. [2](#), [3](#), [5](#), [7](#), [8](#)
- [22] Shuhui Wang Jun Wei and Qingming Huang. F3net: Fusion, feedback and focus for salient object detection. In *AAAI Conf. Art. Intell.*, 2020. [1](#)
- [23] Haksun Kim, Sanghoon Lee, and Alan Bovik. Saliency prediction on stereoscopic videos. *IEEE T. Image Process.*, 23:1476–90, 04 2014. [8](#)
- [24] Chongyi Li, Runmin Cong, Sam Kwong, Junhui Hou, Huazhu Fu, Guopu Zhu, Dingwen Zhang, and Qingming Huang. Asif-net: Attention steered interweave fusion network for rgb-d salient object detection. *IEEE T. Cybern.*, 51(1):88–100, 2021. [2](#)
- [25] Chongyi Li, Runmin Cong, Yongri Piao, Qianqian Xu, and Chen Change Loy. Rgb-d salient object detection with cross-modality modulation and selection. In *Eur. Conf. Comput. Vis.*, 2020. [1](#), [2](#), [6](#)
- [26] Gongyang Li, Zhi Liu, Linwei Ye, Yang Wang, and Haibin Ling. Cross-modal weighting network for rgb-d salient object detection. In *Eur. Conf. Comput. Vis.*, 2020. [2](#)
- [27] Nianyi Li, Jinwei Ye, Yu Ji, Haibin Ling, and Jingyi Yu. Saliency detection on light field. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2806–2813, 2014. [2](#), [3](#), [5](#), [7](#)
- [28] Fangfang Liang, Lijuan Duan, Wei Ma, Yuanhua Qiao, Zhi Cai, and Laiyun Qing. Stereoscopic saliency model using contrast and depth-guided-background prior. *Neurocomputing*, 275:2227–2238, 2018. [6](#)
- [29] Ralph Linsker. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988. [2](#)
- [30] Nian Liu, Ni Zhang, and Junwei Han. Learning selective self-mutual attention for rgb-d saliency detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. [1](#), [2](#), [6](#), [8](#)
- [31] Ao Luo, Xin Li, Fan Yang, Zhicheng Jiao, Hong Cheng, and Siwei Lyu. Cascade graph neural networks for rgb-d salient object detection. In *Eur. Conf. Comput. Vis.*, 2020. [2](#)

- [32] Yuzhen Niu, Yujie Geng, Xueqing Li, and Feng Liu. Leveraging stereopsis for saliency analysis. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 454–461, 2012. [2](#), [3](#), [5](#), [7](#), [8](#)
- [33] Youwei Pang, Lihe Zhang, Xiaoqi Zhao, and Huchuan Lu. Hierarchical dynamic filtering network for rgb-d salient object detection. In *Eur. Conf. Comput. Vis.*, 2020. [1](#), [2](#), [6](#)
- [34] Houwen Peng, Bing Li, Weihua Xiong, Weiming Hu, and Rongrong Ji. Rgb-d salient object detection: a benchmark and algorithms. In *Eur. Conf. Comput. Vis.*, pages 92–109, 2014. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [35] Yongri Piao, Wei Ji, Jingjing Li, Miao Zhang, and Huchuan Lu. Depth-induced multi-scale recurrent attention network for saliency detection. In *Int. Conf. Comput. Vis.*, 2019. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)
- [36] Yongri Piao, Zhengkun Rong, Miao Zhang, Weisong Ren, and Huchuan Lu. A2dele: Adaptive and attentive depth distiller for efficient rgb-d salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. [1](#), [2](#), [6](#), [8](#)
- [37] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. Basnet: Boundary-aware salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. [1](#)
- [38] Liangqiong Qu, Shengfeng He, Jiawei Zhang, Jiandong Tian, Yandong Tang, and Qingxiong Yang. RGBD salient object detection via deep fusion. *IEEE T. Image Process.*, 26(5):2274–2285, 2017. [1](#), [2](#), [6](#)
- [39] M. Tschannen, J. Djolonga, P. K. Rubenstein, S. Gelly, and M. Lucic. On mutual information maximization for representation learning. In *Int. Conf. Learn. Represent.*, 2020. [2](#)
- [40] Jianyuan Wang, Yiran Zhong, Yuchao Dai, Kaihao Zhang, Pan Ji, and Hongdong Li. Displacement-invariant matching cost learning for accurate optical flow estimation. In *Adv. Neural Inform. Process. Syst.*, 2020. [7](#)
- [41] Ningning Wang and Xiaojin Gong. Adaptive fusion for RGB-D salient object detection. *arXiv:1901.01369*, 2019. [1](#), [2](#), [6](#)
- [42] Weiyue Wang and Ulrich Neumann. Depth-aware cnn for rgb-d segmentation. In *Eur. Conf. Comput. Vis.*, 2018. [2](#)
- [43] Wenguan Wang, Jianbing Shen, Ming-Ming Cheng, and Ling Shao. An iterative and cooperative top-down and bottom-up inference network for salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. [1](#)
- [44] Zhe Wu, Li Su, and Qingming Huang. Stacked cross refinement network for edge-aware salient object detection. In *Int. Conf. Comput. Vis.*, 2019. [1](#)
- [45] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3684–3692, 2018. [3](#), [4](#)
- [46] Jing Zhang, Deng-Ping Fan, Yuchao Dai, Saeed Anwar, Fatemeh Sadat Saleh, Tong Zhang, and Nick Barnes. Uncnet: Uncertainty inspired rgb-d saliency detection via conditional variational autoencoders. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. [1](#), [2](#), [4](#), [6](#), [8](#)
- [47] Jing Zhang, Xin Yu, Aixuan Li, Peipei Song, Bowen Liu, and Yuchao Dai. Weakly-supervised salient object detection via scribble annotations. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. [8](#)
- [48] Miao Zhang, Sun Xiao Fei, Jie Liu, Shuang Xu, Yongri Piao, and Huchuan Lu. Asymmetric two-stream architecture for accurate rgb-d saliency detection. In *Eur. Conf. Comput. Vis.*, 2020. [2](#)
- [49] Miao Zhang, Weisong Ren, Yongri Piao, Zhengkun Rong, and Huchuan Lu. Select, supplement and focus for rgb-d saliency detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. [1](#), [2](#)
- [50] Qiudan Zhang, Xu Wang, Shiqi Wang, Shikai Li, Sam Kwong, and Jianmin Jiang. Learning to explore intrinsic saliency for stereoscopic video. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2019. [8](#)
- [51] Zhao Zhang, Zheng Lin, Jun Xu, Wenda Jin, Shao-Ping Lu, and Deng-Ping Fan. Bilateral attention network for rgb-d salient object detection. *arXiv preprint arXiv:2004.14582*, 2020. [1](#), [2](#), [6](#)
- [52] Jiawei Zhao, Yifan Zhao, Jia Li, and Xiaowu Chen. Is depth really necessary for salient object detection? In *Proceedings of the 28th ACM International Conference on Multimedia*, 2020. [1](#)
- [53] Jia-Xing Zhao, Yang Cao, Deng-Ping Fan, Ming-Ming Cheng, Xuan-Yi Li, and Le Zhang. Contrast prior and fluid pyramid integration for rgb-d salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. [1](#), [2](#), [6](#), [8](#)
- [54] Xiaoqi Zhao, Lihe Zhang, Youwei Pang, Huchuan Lu, and Lei Zhang. A single stream network for robust and real-time rgb-d salient object detection. In *Eur. Conf. Comput. Vis.*, 2020. [6](#)
- [55] Yiran Zhong, Charles Loop, Wonmin Byeon, Stan Birchfield, Yuchao Dai, Kaihao Zhang, Alexey Kamenev, Thomas Breuel, Hongdong Li, and Jan Kautz. Displacement-invariant cost computation for efficient stereo matching. *arXiv preprint arXiv:2012.00899*, 2020. [5](#)