

A General Divergence Modeling Strategy for Salient Object Detection

Xinyu Tian¹, Jing Zhang², and Yuchao Dai^{1✉}

¹ Northwestern Polytechnical University, Xi'an, China

² Australian National University, Canberra, Australia

Abstract. Salient object detection is subjective in nature, which implies that multiple estimations should be related to the same input image. Most existing salient object detection models are deterministic following a point to point estimation learning pipeline, making them incapable of estimating the predictive distribution. Although latent variable model based stochastic prediction networks exist to model the prediction variants, the latent space based on the single clean saliency annotation is less reliable in exploring the subjective nature of saliency, leading to less effective saliency “divergence modeling”. Given multiple saliency annotations, we introduce a general divergence modeling strategy via random sampling, and apply our strategy to an ensemble based framework and three latent variable model based solutions to explore the “subjective nature” of saliency. Experimental results prove the superior performance of our general divergence modeling strategy.

Keywords: Salient object detection · Divergence modeling.

1 Introduction

When viewing a scene, human visual system has the ability to selectively locate attention [1,2,3,4,5] on the informative contents, which locally stand out from their surroundings. This selection is usually performed in the form of a spatial circumscribed region, leading to the so-called “focus of attention” [4]. Itti *et al.* [1] introduced a general attention model to explain the human visual search strategies [6]. Specifically, the visual input is first decomposed into a group of topographic feature maps which they defined as the early representations. Then, different spatial locations compete for saliency within each topographic feature map, such that locations that locally stand out from their surrounding persist. Lastly, all the feature maps are fed into a master “saliency map”, indicating the topographical codes for saliency over the visual scene [4].

Following the above process of saliency selection, early saliency detection models focus on detecting the informative locations, leading to the eye fixation

Yuchao Dai (daiyuchao@nwpu.edu.cn) is the corresponding author. This work was supported in part by the NSFC (61871325). Code is available at https://npucvr.github.io/Divergence_SOD/.

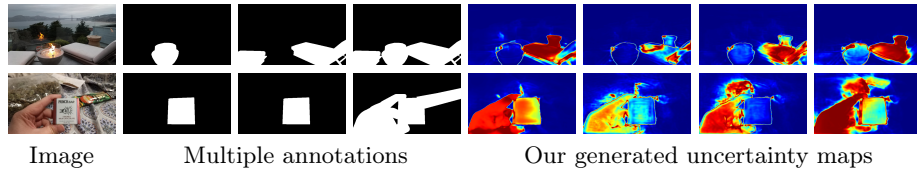


Fig. 1. Given “Multiple Annotations” of each training image, the proposed “Random Sampling” strategy based solutions aim to generate diverse predictions representing the subjective nature of saliency with reliable “generated uncertainty maps”.

prediction [7] task. [8] and [9] then extended the salient locations driven methods [1,4] and introduced the salient object detection task, which is a binary segmentation task aiming to identify the full scope of salient objects. In this way, “salient object” is defined as any item that is distinct from those around it. Many factors can lead something to be “salient”, including the stimulus itself that makes the item distinct, *i.e.* color, texture, direction of movement, and the internal cognitive state of observers, leading to his/her understanding of saliency.

As an important computer vision task, salient object detection is intrinsic to various tasks such as image cropping [10], object detection [11], semantic segmentation [12], video segmentation [13], action recognition [14], image caption generation [15] and semantic image labeling [16], where saliency models can be used to extract class-agnostic important areas in an image or a video sequence. Most of the existing salient object detection models intend to achieve the point estimation from input RGB image (or RGB-D image pair) to the corresponding ground truth (GT) saliency map, neglecting the less consistent saliency regions discarded while generating the binary ground truth maps via majority voting. Although [17] presented a generative model via conditional variational auto-encoder [18] to produce stochastic prediction at test time, the generated stochastic saliency maps have low diversity, where the difference of those stochastic predictions mainly along the object boundaries, making it less effective in explaining the “subjective nature” of saliency as shown in Fig. 3.

In this paper, we study “divergence modeling” for salient object detection, and propose a general strategy to achieve effective divergence modeling (see Fig. 1), representing the “subjective nature” of saliency. Given multiple saliency labels [19] of each image, we aim to generate one majority saliency map, reflecting the majority saliency attribute, and an uncertainty map, explaining the discarded less salient regions. Specifically, we adopt a separate “majority voting” module to regress the majority saliency map for deterministic performance evaluation. We also need to generate stochastic predictions given the multiple annotations. This can be achieved with multiple model parameters (*e.g.* deep ensemble [20]) or with extra latent variable to model the predictive distribution [17], leading to our “ensemble”-based framework and “latent variable model”-based solutions.

We observe that, training simultaneously with all the annotations participating into model updating is less effective in modeling the divergence of saliency,

as the model will be dominated by the majority annotation. The main reason is that there exists much less samples containing diverse annotations (around 20%) compared with samples with consistent annotations. To solve the above issue, we introduce a simple “random sampling” strategy. Specifically, within the “ensemble” framework, the network is randomly updated with one annotation. For the “latent variable model” based networks, the generated stochastic prediction is randomly compared with one annotation from the multiple annotations pool. The model is then updated with loss function based on both the deterministic prediction from the majority voting branch and one stochastic prediction from the latent variable branch. We have carried out extensive experiments and find that this simple “random sampling” based strategy works superbly in modeling the “divergence” of saliency, leading to meaningful uncertainty maps representing the “subjective nature” of saliency (see Fig. 1). We also verified the importance of multiple annotations in modeling saliency divergence, and further applied our random sampling strategy to other SOTA SOD models, which can also achieve diverse saliency predictions, demonstrating the flexibility of our approach.

Aiming at discovering the discarded less salient regions for human visual system exploration, we work on a general divergence modeling strategy for saliency detection. Our main contributions are: **1)** we introduce the first “random sampling” based divergence modeling strategy for salient object detection to produce reliable “object” level uncertainty map explaining human visual attention; **2)** we present an ensemble based framework and three latent variable model based solutions to validate our divergence modeling strategy; **3)** to maintain the deterministic performance, we also design an efficient majority voting branch to generate majority voting saliency maps without sampling at test time.

2 Related Work

Salient object detection: The main stream of salient object detection (SOD) models are fully-supervised [21,22,23,24,25,26,27,28,29], where most of them focus on effective high/low feature aggregation or structure-aware prediction [30,31,23,32]. With extra depth data, RGB-D SOD [33,34] models mainly focus on effective multimodal learning. In addition to above fully-supervised models, semi-supervised [35], weakly-supervised [36,37] SOD models have also been explored. We argue that most of the existing SOD models define saliency detection as a binary segmentation task in general, without exploring the subjective nature of saliency. Although [17] has taken a step further to model the predictive distribution via a conditional variational auto-encoder (CVAE) [18], the diversity of the stochastic predictions mainly distributes along object boundaries, making it less effective in discovering the discarded less salient regions.

Uncertainty estimation: The goal of uncertainty estimation is to measure the confidence of model predictions. According to [38], the uncertainty of deep neural networks is usually divided into aleatoric uncertainty (or data uncertainty) and epistemic uncertainty (or model uncertainty). Aleatoric uncertainty describes the inherent randomness of the task (*e.g.* dataset annotation error or sensor

noise), which cannot be avoided. Epistemic uncertainty is caused by our limited knowledge about the latent true model, which can be reduced by using a larger or more diverse training dataset. The aleatoric uncertainty is usually estimated via a dual-head framework [39], which produces both task related prediction and the corresponding aleatoric uncertainty. Epistemic uncertainty aims to represent the model bias, which can be obtained via a Bayesian Neural Network (BNN) [40]. In practice, Monte Carlo Dropout (MC Dropout) [41] and Deep Ensemble [20,42,43] are two widely studied epistemic uncertainty estimation techniques.

Latent variable model based stochastic models: With extra latent variable involved, the latent variable models [44,18,45,46] can be used to achieve predictive distribution estimation. Currently, latent variable models have been used in numerous tasks, such as semantic segmentation[47], natural language processing[48], depth estimation[49], deraining[50], image deblurring[51], saliency detection[52,17], *etc.*. In this paper, we explore latent variable models for saliency divergence modeling with “random sampling” strategy.

Uniqueness of our solution: Different from the existing point estimation based salient object detection networks [30,53] which produce deterministic saliency maps, our method aims to generate diverse predictions representing human diverse perceptions towards the same scene. Although the existing generative model based saliency prediction networks such as [17] can produce stochastic predictions, those methods can mainly highlight the labeling noise, and are moderately effective in modeling predictive distribution, because the posterior of the latent variable z only depends on the single ground truth after majority voting. Building upon the newly released saliency dataset with multiple annotations [19] for each input image, we present a general “divergence modeling” strategy based on “random sampling”, where only the randomly selected annotation can participate into model updating, leading to reliable “object-level” uncertainty map (see Fig. 1) representing the “subjective nature” of saliency.

3 Our Method

Our training dataset is $D = \{x_i, \{y_i^j\}_{j=0}^M\}_{i=1}^N$, where x_i is the input RGB image, $\{y_i^j\}_{j=1}^M$ are the multiple saliency annotations for x_i , and y_i^0 (with $j = 0$) is the saliency map after majority voting. M and N are number of annotators and the size of the training dataset, respectively. i indexes the images, and j indexes the multiple annotations [19]. We aim to model the “subjective nature” of saliency with a general divergence modeling strategy as shown in Fig. 2.

Given the multiple annotations $\{y^j\}_{j=1}^M$ for image x (we omit image index i when there is no ambiguity), we introduce a random sampling based strategy to randomly select annotations for model updating. The main reason is that we have limited samples containing diverse annotations (around 20%), and training directly with all the annotations participating into model updating will generate a biased model, where samples with consistent annotations will dominate the training process (see Table 1 and Table 2), leading to less effective divergence modeling. Our “random sampling” strategy encourages the diverse annotations

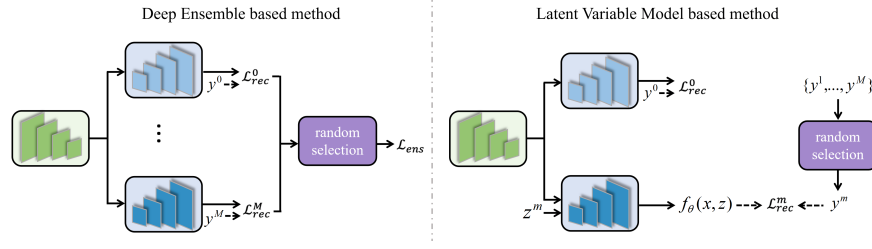


Fig. 2. The proposed strategy within the ensemble based framework (left) and the latent variable model based solutions (right). By randomly selecting one ground truth from the multiple annotations for model updating, the proposed strategy can better explore the contribution of multiple annotations for human visual system exploration.

to fully contribute to model updating. To extensively analyse the proposed strategy, we design an ensemble based framework and three latent variable based networks with the same saliency generator f_θ , except that the extra latent variable is used in the latent variable models.

Saliency generator: Our “saliency generator” f_θ is built upon the ResNet50 backbone [54], and we define the backbone features as $\{s_k\}_{k=1}^4$ of channel size 256, 512, 1024 and 2048 respectively. To relief the huge memory requirement and also obtain larger receptive field, we feed the backbone features to four different multi-scale dilated convolutional blocks [55] and obtain new backbone features $f_\theta^b(x) = \{s'_k\}_{k=1}^4$ of the same channel size $C = 32$. We then feed $\{s'_k\}_{k=1}^4$ to decoder from [56] $f_\theta^d(f_\theta^b(x))$ (or $f_\theta^d(f_\theta^b(x), z)$ for the latent variable model solutions) to generate saliency prediction.

Task-related loss function: The widely used loss functions for saliency detection include: 1) binary cross-entropy loss, 2) boundary IOU loss [57], and 3) structure-aware loss [30], which is a weighted combination of the above two loss functions. In this paper, for all the four related divergence modeling models, we adopt structure-aware loss function as the saliency reconstruction loss \mathcal{L}_{rec} .

3.1 Ensemble Framework

Deep ensemble [20] is a widely used method to generate multiple predictions, which uses different network structures for training to achieve multiple predictions with one single forward pass. In our case, instead of having multiple copies of mapping function from input image x to the diverse annotation y^j , we design $M + 1$ decoders with the same structure to achieve divergence modeling, where each output of the decoders is supervised by y^j (including the majority voting annotation y^0). The conventional loss definition for deep ensemble framework is computing the sum of the multiple predictions, or define the minimal loss of each decoder as the final loss to update model parameters. However, we find that the above strategies lead to less diverse predictions, making them less effective in modeling the subjective nature of saliency. In this paper, we first design an

ensemble structure with $M + 1$ decoders of the same structure $f_\theta^d(f_\theta^b(x))$, which are initialized differently. Then, we randomly select a decoder branch, and define its loss function as the final loss function for model updating (see Fig. 2). The loss function of our deep ensemble based model is then defined as:

$$\mathcal{L}_{ens} = \text{RS}_j[\{\mathcal{L}_{\text{rec}}^j\}_{j=0}^M], \quad (1)$$

where the $\text{RS}[\cdot]$ is the random selection operation and $\mathcal{L}_{\text{rec}}^m = \mathcal{L}(y_{\text{pred}}^m, y^m)$ is the reconstruction loss of prediction from the m_{th} decoder supervised by y^m .

3.2 Latent Variable Model based Networks

With extra latent variable z , the latent variable models [44,18,45,46], *e.g.* variational auto-encoder (VAE) [44,18], generative adversarial net (GAN) [45], alternating back-propagation (ABP) [46], are capable of modeling the predictive distribution by sampling z from the latent space. We then design three latent variable models with the proposed “random sampling” strategy to achieve effective divergence modeling. Note that, with extra conditional variable x , our VAE based framework is indeed a conditional variational auto-encoder (CVAE) [18].

In general, given latent variable z , our latent variable model (see Fig. 2) based networks include two main parts, namely a deterministic prediction network $f_\theta(x)$ supervised by y^0 and a latent variable model $f_\theta(x, z)$ supervised by $\{y^j\}_{j \neq 0}$. The former is the same as the saliency generator with single head in the ensemble framework. For the latter, we first tile z in spatial dimension, and then concatenate it with s'_4 channel-wise, leading to sz'_4 of channel size $C + K$, where K is the dimension of the latent variable. We set it as $K = 32$ for all the three latent variable models. To adapt to the saliency decoder, we feed sz'_4 to another 3×3 convolutional layer to obtain a new feature nz'_4 of channel size C . In this way, the stochastic prediction $f_\theta(x, z)$ is obtained via: $f_\theta^d(\{s'_k\}_{k=1}^3, nz'_4)$.

CVAE Solution: A CVAE [18] is a conditional directed graph model, which includes three variables, the input x or conditional variable that modulates the prior on Gaussian latent variable z and generates the output prediction y . Two main modules are included in a conventional CVAE based framework: a generator model $f_\theta(x, z)$, which is a saliency generator in this paper, and an inference model $q_\theta(z|x, y)$, which infers the latent variable z with image x and annotation y as input. Learning a CVAE framework involves approximation of the true posterior distribution of z with an inference model $q_\theta(z|x, y)$, with the loss function as:

$$\mathcal{L}_{cvae} = \underbrace{\mathbb{E}_{z \sim q_\theta(z|x, y)}[-\log p_\theta(y|x, z)]}_{\mathcal{L}_{\text{rec}}} + D_{KL}(q_\theta(z|x, y) \parallel p_\theta(z|x)), \quad (2)$$

where the first term is the reconstruction loss and the second is the Kullback-Leibler divergence of prior distribution $p_\theta(z|x)$ and posterior distribution $q_\theta(z|x, y)$. The saliency generator has been discussed, and we then introduce the details about the prior and posterior distribution models.

Prior net and posterior net construction: We have almost the same structure for the prior net and the posterior net, where the only difference is that the input channel of the first convolutional layer within the prior net is 3 (channel of the RGB image x) and that within the posterior net is 4 (channel of the concatenation of x and y). We construct our prior and posterior net with five convolution layers of the same kernel size (4×4) and two fully connected layers. The channel size of the convolutional layers are $C' = C, 2*C, 4*C, 8*C, 8*C$ ($C = 32$, which is the same as the channel reduction layers for the backbone features), and we have a batch normalization layer and a ReLU activation layer after each convolutional layer. The two fully connected layers are used to map the feature map to vectors representing the mean μ and standard deviation σ of the latent variable z . The latent variable z is then obtained via the reparameterization trick: $z = \mu + \sigma \odot \epsilon$, where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$.

Divergence modeling with CVAE: Recall that the posterior distribution is defined as $q_\theta(z|x, y)$. Given the multiple annotations $\{y^j\}_{j \neq 0}$, we randomly select one annotation y^m , and use it to define the posterior as $q_\theta(z|x, y^m)$. In this way, loss function for the CVAE based divergence modeling model is defined as:

$$\mathcal{L}_{cvae} = \mathcal{L}_{\text{rec}}^m + D_{KL}(q_\theta(z|x, y^m) \parallel p_\theta(z|x)). \quad (3)$$

GAN Solution: Within the GAN based framework, we design an extra fully convolutional discriminator g_β following [58], where β is the parameter set of the discriminator. Two different modules (the saliency generator f_θ and the discriminator g_β in our case) play the minimax game in GAN based framework:

$$\min_{f_\theta} \max_{g_\beta} V(g_\beta, f_\theta) = E_{(x,y) \sim p_{\text{data}}(x,y)} [\log g_\beta(y|x)] + E_{z \sim p(z)} [\log(1 - g_\beta(f_\theta(x, z)))], \quad (4)$$

where $p_{\text{data}}(x, y)$ is the joint distribution of training data, $p(z)$ is the prior distribution of the latent variable z , which is usually defined as $p(z) = \mathcal{N}(0, \mathbf{I})$. In practice, we define loss function for the generator as the sum of a reconstruction loss \mathcal{L}_{rec} , and an adversarial loss $\mathcal{L}_{\text{adv}} = \mathcal{L}_{ce}(g_\beta(f_\theta(x, z)), 1)$, which is $\mathcal{L}_{\text{gen}} = \mathcal{L}_{\text{rec}} + \lambda \mathcal{L}_{\text{adv}}$, where the hyper-parameter λ is tuned, and empirically we set $\lambda = 0.1$ for stable training. \mathcal{L}_{ce} is the binary cross-entropy loss. The discriminator g_β is trained via loss function as: $\mathcal{L}_{\text{dis}} = \mathcal{L}_{ce}(g_\beta(f_\theta(x, z)), 0) + \mathcal{L}_{ce}(g_\beta(y), 1)$. Similar to CVAE, for each iteration of training, we randomly select an annotation y^m , which will be treated as y for both generator and discriminator updating. In this way, with the proposed random sampling strategy for divergence modeling, generator loss and discriminator loss can be defined as:

$$\begin{aligned} \mathcal{L}_{\text{gen}} &= \mathcal{L}_{\text{rec}}^m + \lambda \mathcal{L}_{\text{adv}}, \\ \mathcal{L}_{\text{dis}} &= \mathcal{L}_{ce}(g_\beta(f_\theta(x, z)), 0) + \mathcal{L}_{ce}(g_\beta(y^m), 1). \end{aligned} \quad (5)$$

ABP Solution: Alternating back-propagation [46] was introduced for learning the generator network model. It updates the latent variable and network parameters in an EM-manner. Firstly, given the network prediction with the current

parameter set, it infers the latent variable by Langevin dynamics based MCMC [59], which is called “Inferential back-propagation” in [46]. Secondly, given the updated latent variable z , the network parameter set is updated with gradient descent. [46] calls it “Learning back-propagation”. Following the previous variable definitions, given the training example (x, y) , we intend to infer z and learn the network parameter θ to minimize the reconstruction loss as well as a regularization term that corresponds to the prior on z .

Different from the CVAE solution, where extra inference model $q_\theta(z|x, y)$ is used to approximate the posterior distribution of the latent variable z . ABP [46] samples z directly from its posterior distribution with a gradient-based Monte Carlo method, namely Langevin Dynamics [59]:

$$z_{t+1} = z_t + \frac{s^2}{2} \left[\frac{\partial}{\partial z} \log p_\theta(y, z_t|x) \right] + s\mathcal{N}(0, \mathbf{I}), \quad (6)$$

where $z_0 \sim \mathcal{N}(0, \mathbf{I})$, and the gradient term is defined as:

$$\frac{\partial}{\partial z} \log p_\theta(y, z|x) = \frac{1}{\sigma^2} (y - f_\theta(x, z)) \frac{\partial}{\partial z} f_\theta(x, z) - z. \quad (7)$$

t is the time step for Langevin sampling, s is the step size, σ^2 is variance of the inherent labeling noise. Empirically we set $s = 0.1$ and $\sigma^2 = 0.3$ in this paper.

In practice, we sample z_0 from $\mathcal{N}(0, \mathbf{I})$, and update z via Eq. 6 by running $T = 5$ steps of Langevin sampling [59]. The final z_T is then used to generate saliency prediction in our case. Similar to the other two latent variable models, we randomly select an annotation y^m from $\{y^j\}_{j \neq 0}$, and use it to infer the latent variable z . The loss function for ABP with our random sampling strategy for saliency detection is then defined as: $\mathcal{L}_{abp} = \mathcal{L}_{\text{rec}}^m$.

3.3 Prediction Generation

Majority Voting Regression: To simulate the majority voting in the saliency labeling process, we add a deterministic majority voting prediction branch to the latent variable models, which share the same structure as our single ensemble network. We define its loss function as $\mathcal{L}_{mj} = \mathcal{L}_{\text{rec}}^0$. In this way, in addition to the stochastic prediction based loss functions for each latent variable model (\mathcal{L}_{cvae} , \mathcal{L}_{gen} , \mathcal{L}_{abp}), we further add the deterministic loss function \mathcal{L}_{mj} to each of them. At test time, prediction from the majority voting regression branch is defined as our deterministic prediction for performance evaluation.

Obtaining the Uncertainty: Given the ensemble based framework, the predictive uncertainty is defined as entropy of the mean prediction [39]: $U_p = \mathbb{H}[\text{mean}_{j \in [1, M]} \{f_\theta^{dj}(f_\theta^b(x))\}]$, where f_θ^{dj} is the j_{th} decoder. For the latent variable models, the predictive uncertainty is defined as $U_p = \mathbb{H}[\mathbb{E}_{p_\theta(z|x)}(p_\theta(y|x, z))]$. For both models, the aleatoric uncertainty is defined as the mean entropy of multiple predictions. In practice, we usually use Monte Carlo integration as approximation of the intractable expectation operation. The epistemic uncertainty is then defined as $U_e = U_p - U_a$.

4 Experimental Results

Dataset: We use the COME dataset [19] for training as it’s the only large training dataset (with 8,025 images) containing multiple annotations for each image. For each image, there are five annotations ($\{y^j\}_{j=1}^M$ with $M = 5$) from five different annotators. The ground truth after majority voting is y^0 . The testing images include 1) DUTS testing dataset [60], 2) ECSSD [61], 3) DUT [62], 4) HKU-IS [63], 5) COME-E [19], 6) COME-H [19], where each image in COME-E [19] and COME-H [19] datasets has 5 annotations from 5 different annotators.

Evaluation metrics: We use Mean Absolute Error (\mathcal{M}), mean F-measure (F_β) and mean E-measure (E_ξ) [64] for deterministic performance evaluation. We also present an uncertainty based mean absolute error to estimate the divergence modeling ability of each model. Details about these metrics will be introduced in the supplementary material.

Implementation details: We train our models in Pytorch with ResNet-50 backbone, which is initialized with weights trained on ImageNet, and other newly added layers are initialized by default. We resize all the images and ground truth to 352×352 for both training. The maximum epoch is 50. The initial learning rate is 2.5×10^{-5} . The whole training takes 9 hours, 11 hours with batch size 8 on one NVIDIA GTX 2080Ti GPUs for deep ensemble model and latent variable models respectively. The inference speed is around 20 images/second.

4.1 Divergence Modeling Performance

As a divergence modeling strategy, we aim to produce both accurate deterministic predictions and reliable uncertainty. We design the following experiments to explain the divergence modeling ability of our proposed strategy.

Firstly, we train two baseline models with only the majority voting ground truth y^0 (“Base_M”) and with all the annotations (“Base_A”). The results are shown in Table 1. For the former, the training dataset is $D1 = \{x_i, y_i^0\}_{i=1}^N$, and for the latter, the training dataset is $D2 = \{\{x_i, y_i^j\}_{j=0}^M\}_{i=1}^N$, where $\{x_i\}_{j=0}^M = x_i$. Then, we train all four frameworks with dataset $D2$, and report their performance as “DE_A”, “GAN_A”, “VAE_A” and “ABP_A” in Table 1 respectively. With the same network structures, we change the learning strategy by randomly select an annotation to learn our divergence modeling models, and the results are shown as “DE_R”, “GAN_R”, “VAE_R” and “ABP_R” in Table 1 respectively. In this way, for each iteration, we only use one divergent annotation to update model parameters of the stochastic prediction branch. Specifically, as shown in Fig. 2, for “DE_R”, we randomly pick up one decoder branch and an annotation from $\{y^j\}_{j=0}^M$, and use their loss function to update the whole network. For the latent variable model based frameworks, the deterministic majority voting branch is updated with the majority voting ground truth, and the stochastic prediction branch is updated with randomly picked annotation $\{y^j\}_{j=1}^M$ as supervision.

The better performance of training with the randomly selected annotations (the “_R” models) compared with training with all the annotation simultaneously (the “_A” models) illustrates the effectiveness of our strategy. We also

Table 1. Performance of the proposed strategy within the ensemble and the latent variable model based frameworks.

	DUTS [60]	ECSSD [61]	DUT [62]	HKU-IS [63]	COME-E [19]	COME-H [19]
	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$
Base_A	.811 .897 .045	.903 .931 .044	.741 .846 .059	.887 .935 .036	.852 .905 .049	.807 .859 .080
Base_M	.824 .911 .041	.916 .946 .038	.754 .864 .055	.900 .950 .032	.877 .927 .040	.838 .882 .070
DE_A	.819 .906 .042	.910 .939 .040	.744 .855 .057	.895 .940 .034	.870 .920 .043	.831 .876 .073
DE_R	.816 .900 .043	.916 .944 .038	.738 .843 .060	.892 .939 .036	.869 .920 .043	.832 .880 .072
GAN_A	.803 .895 .045	.906 .939 .040	.736 .850 .060	.879 .931 .038	.855 .912 .046	.816 .867 .076
GAN_R	.817 .906 .041	.914 .944 .039	.741 .849 .059	.895 .946 .033	.873 .926 .040	.834 .882 .070
VAE_A	.798 .888 .049	.910 .940 .039	.727 .843 .068	.890 .937 .036	.853 .909 .048	.813 .865 .078
VAE_R	.822 .899 .042	.913 .936 .042	.723 .824 .062	.895 .937 .036	.880 .924 .041	.838 .879 .071
ABP_A	.777 .866 .052	.888 .916 .049	.720 .829 .063	.853 .900 .046	.835 .886 .054	.794 .844 .085
ABP_R	.805 .896 .047	.911 .938 .044	.734 .846 .060	.885 .937 .039	.864 .911 .050	.824 .863 .083

Table 2. Mean absolute error of the generated uncertainty maps. The results in each block from left to right are models trained with all the annotations simultaneously (“_A” in Table 1), only the majority annotation (“_M” in Table 3), random sampling based on three (Table 4) and five annotations pool (“_R” in Table 1).

	DE	GAN	VAE	ABP
COME-E	.097 - .094 .092	.085 .095 .086 .082	.089 .093 .088 .086	.108 .112 .110 .105
COME-H	.128 - .125 .124	.118 .124 .119 .114	.123 .124 .121 .119	.142 .146 .140 .136

show the generated uncertainty maps in Fig. 1 as “Our generated uncertainty maps”, where from left to right are uncertainty map from our ensemble solution, GAN solution, VAE solution and ABP solution. We observe reliable uncertainty maps with all our solutions, where uncertainty of the consistent salient regions usually distributes along object boundaries, and we produce uniform uncertainty activation for the regions with less consistent saliency agreement.

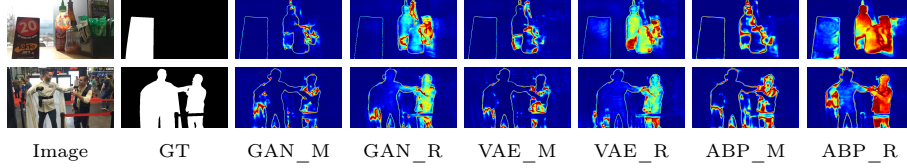
To quantitatively analyse the effectiveness of uncertainty maps with the proposed strategy, we compute mean absolute error of the estimated predictive uncertainty and the predictive uncertainty of the corresponding testing dataset, and show performance in Table 2. Given the multiple annotations, we first compute entropy of the mean annotation, and normalize it to the range of [0,1]. We then define it as the ground truth predictive uncertainty. The predictive uncertainty of the four frameworks can be obtained similarly. Table 2 shows that our random sampling based strategy leads to lower uncertainty based mean absolute error compared with updating the model with all the multiple annotations, which quantitatively validate our strategy.

4.2 Uncertainty Comparison with Single Annotation based Models

Existing generative model based saliency detection models [17] are based on single version ground truth. Although [17] generated multiple pseudo ground truth, the training pipeline is the same as our experiment “_A” in Table 1, where the model simply learns from more pair of samples with dataset $D2 = \{\{x_i^j, y_i^j\}_{j=0}^M\}_{i=1}^N$. To compare uncertainty maps based on the single annotation (y^0 in our case) and multiple annotations ($\{y^j\}_{j=1}^M$), we train the three latent

Table 3. Performance of the latent variable models for divergence modeling with only majority voting ground truth as supervision.

	DUTS [60]	ECSSD [61]	DUT [62]	HKU-IS [63]	COME-E [19]	COME-H [19]
	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$	$F_\beta \uparrow E_\xi \uparrow \mathcal{M} \downarrow$
Base_M	.824 .911 .041	.916 .946 .038	.754 .864 .055	.900 .950 .032	.877 .927 .040	.838 .882 .070
GAN_M	.817 .905 .042	.913 .945 .038	.747 .855 .058	.893 .944 .035	.870 .923 .042	.831 .881 .071
VAE_M	.817 .905 .042	.913 .945 .038	.744 .852 .058	.893 .943 .035	.873 .924 .042	.831 .880 .072
ABP_M	.801 .893 .045	.908 .942 .040	.732 .840 .060	.886 .937 .038	.862 .916 .045	.825 .874 .074

**Fig. 3.** Uncertainty maps of latent variable models with single majority voting GT (“_M”) and multiple ($M = 5$) diverse GT using our divergence modeling strategy (“_R”).

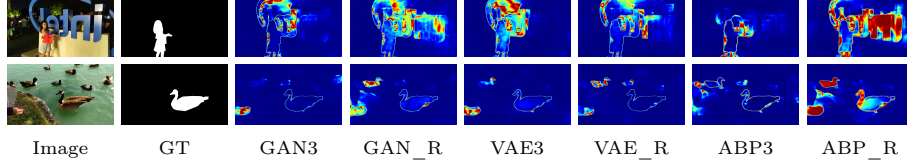
variable models with single majority voting ground truth y^0 , and show model performance in Table 3. Similarly, we find that the majority voting ground truth based latent variable models lead to comparable performance as the base model “Base_M” trained with y^0 . Then, we compare the uncertainty map difference of the two settings in Fig. 3. Different from our strategy with uniform uncertainty activation within the less confident region, the majority voting ground truth based model usually generate uncertainty distributing along object boundaries, making it less effective in modeling the “subjective nature” of saliency, which explains the necessity of multiple annotations for saliency divergence modeling.

4.3 Uncertainty w.r.t. Diversity of Annotations

In this paper, we investigate in diversity modeling ability of ensemble based model and latent variable models with multiple annotations ($M = 5$) for each input image. We further analyse the uncertainty estimation techniques with less annotations, and show model performance in Table 4, where we train our models with three different annotations. Compared with training with five different annotations, we observe similar deterministic performance when training with three annotations. Then, we visualize the produced uncertainty maps in Fig. 4. The more reliable uncertainty map of using more annotations shows that better latent space exploration can be achieved with more diverse annotations. The main reason is more annotations can provide network with more diverse targets that can be learned, making the predictions more diverse and uncertainty more focus on potentially less salient objects. We also show the mean absolute error of the estimated predictive uncertainty of using three annotations as supervision in Table 2. The slightly increased mean absolute uncertainty error further explains the necessity of more diverse annotations for saliency divergence modeling.

Table 4. Performance of our strategy based models with three diverse annotations.

	DUTS [60]			ECSSD [61]			DUT [62]			HKU-IS [63]			COME-E [19]			COME-H [19]		
	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$	$F_\beta \uparrow$	$E_\xi \uparrow$	$\mathcal{M} \downarrow$
DE3	.803	.896	.045	.913	.946	.037	.731	.841	.059	.891	.942	.035	.862	.918	.043	.827	.878	.072
GAN3	.821	.905	.041	.919	.946	.037	.738	.843	.059	.897	.942	.034	.875	.925	.041	.835	.880	.071
VAE3	.824	.909	.040	.920	.948	.036	.747	.854	.057	.898	.946	.034	.876	.926	.040	.836	.881	.070
ABP3	.781	.880	.051	.906	.938	.041	.714	.835	.066	.878	.931	.040	.852	.909	.048	.815	.865	.079

**Fig. 4.** Uncertainty map comparison w.r.t. the number of annotations for each image.**Table 5.** Performance comparison with benchmark SOD models trained on the COME training dataset. Red and blue indicate the best and the second best respectively.

	DUTS [60]			ECSSD [61]			DUT [62]			HKU-IS [63]			COME-E [19]			COME-H [19]		
Method	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}
SCRN [65]	.805	.883	.047	.911	.937	.040	.738	.836	.059	.887	.929	.038	.871	.918	.043	.830	.873	.073
ITSD [66]	.824	.907	.042	.916	.945	.038	.772	.872	.053	.895	.944	.033	.870	.922	.043	.835	.880	.072
MINet[23]	.820	.902	.041	.920	.947	.037	.751	.852	.053	.898	.945	.032	.877	.924	.040	.840	.884	.068
LDF [53]	.818	.902	.043	.919	.947	.036	.754	.856	.055	.902	.948	.031	.874	.924	.041	.836	.882	.070
GateNet[24]	.820	.898	.042	.920	.943	.038	.749	.847	.056	.897	.939	.035	.876	.922	.041	.839	.880	.070
PAKRN [67]	.836	.908	.042	.920	.947	.036	.765	.863	.055	.895	.937	.037	.893	.932	.040	.849	.886	.070
SAMNet[28]	.700	.816	.076	.844	.883	.069	.689	.812	.072	.817	.879	.060	.796	.858	.073	.766	.818	.104
DCN[68]	.829	.904	.041	.923	.946	.037	.765	.861	.051	.905	.945	.032	.885	.927	.039	.843	.880	.071
PFSNet[69]	.795	.885	.053	.919	.947	.037	.739	.847	.065	.897	.944	.034	.867	.917	.044	.831	.879	.072
Ours	.824	.911	.041	.916	.946	.038	.754	.864	.055	.900	.950	.032	.877	.927	.040	.838	.882	.070

4.4 Performance Comparison with SOTA SOD Models

We compare performance of our baseline model (“Base_M” in Table 1) with state-of-the-art (SOTA) SOD models and show model performance in Table 5, where “Ours” represents our baseline model trained with the majority ground truth. As existing salient object detection models are usually trained with DUTS [60] training dataset, we re-train those models with COME training dataset [19], where the majority voting ground truth is used as supervision. Table 5 shows comparable performance of our baseline model compared with the SOTA SOD models. Note that, the focus of our paper is achieving “divergence modeling” with informative uncertainty maps (see Fig. 1), discovering the “less salient regions” that are discarded while preparing the majority ground truth maps. Although existing models [17] can produce stochastic predictions, the generated uncertainty maps fail to discover the “object level” uncertainty (see Fig. 3). Differently, our solution can generate more reliable uncertainty maps (see Fig. 3 and Table 2), which is more informative in explaining human visual system.

Table 6. Performance of applying the proposed strategy to SOTA SOD model, where “-E”, “-G”, “-V” and “-A” indicate the corresponding deep ensemble, GAN, VAE and ABP based model. Note that, we perform random sampling within the models.

Method	DUTS [60]			ECSSD [61]			DUT [62]			HKU-IS [63]			COME-E [19]			COME-H [19]		
	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}	F_β	E_ξ	\mathcal{M}
PFSNet[69]	.795	.885	.053	.919	.947	.037	.739	.847	.065	.897	.944	.034	.867	.917	.044	.831	.879	.072
-E	.791	.890	.047	.904	.936	.039	.741	.858	.062	.892	.946	.035	.865	.928	.042	.839	.882	.071
-G	.830	.908	.040	.924	.948	.036	.755	.855	.055	.900	.945	.033	.889	.932	.037	.847	.887	.066
-V	.829	.902	.040	.919	.943	.038	.733	.832	.058	.897	.939	.034	.891	.930	.038	.845	.882	.069
-A	.791	.883	.053	.915	.948	.037	.739	.856	.061	.898	.951	.034	.866	.924	.043	.831	.886	.072

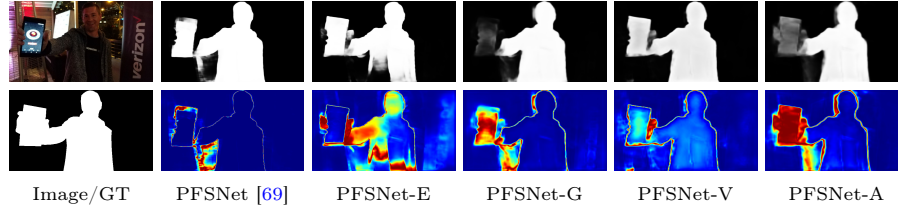


Fig. 5. Uncertainty maps of SOTA SOD model (PFSNet [69] in particular) with the proposed general divergence modeling strategy, where “-E”, “-G”, “-V” and “-A” are introduced in Table 6. The first column show image and the majority voting ground truth, and from the second column to the last one, we show prediction (top) and the corresponding predictive uncertainty (bottom).

Applying Our Solution to SOTA SOD Model: We apply our strategy to one SOTA SOD model, namely PFSNet [69], and show the deterministic performance in Table 6, and the uncertainty maps in Fig. 5. Table 6 shows that the proposed divergence modeling strategy can keep the deterministic performance unchanged, which is consistent with our experiments in Table 3. Besides it, with the proposed “random sampling” based divergence modeling strategy, we observe more reliable “object level” uncertainty maps (“PFSNet-{E,G,V,A}” in Fig. 5) compared with the original Softmax based uncertainty map³ (second column of Fig. 5), which further explains the superiority our strategy in applying to existing SOTA SOD models. We have also applied our strategy to other SOTA SOD models [23,24] and achieved comparable deterministic performance with reliable uncertainty estimation, which further explains superiority of our approach.

4.5 Comparison with Relative Saliency Ranking

Similar to our divergence modeling strategy, relative saliency ranking [71,70,72,73] aims to explore human perception system for better saliency understanding, while differently, it achieves this via inferring the saliency levels, and we present

³ To obtain the “Softmax uncertainty” of deterministic prediction, we compute the entropy of the prediction, which is then defined as the corresponding uncertainty.

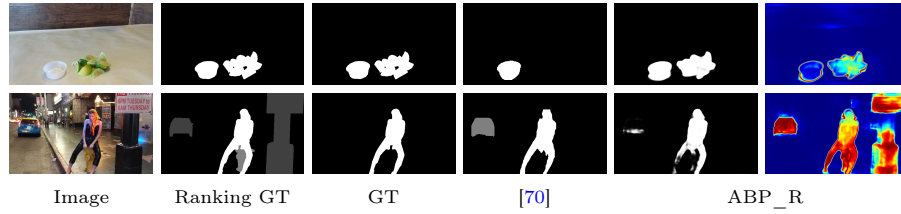


Fig. 6. Performance comparison with saliency ranking model [70], where “GT” is the binary saliency GT after majority voting, and the two predictions of “ABP_R” represent the regressed majority saliency maps and the generated uncertainty maps.

a general divergence modeling strategy to discover the less salient objects. To clearly explain the difference of saliency ranking and our divergence modeling strategy, we visualize the ranking ground truth maps, the ranking predictions of [70], and our generated uncertainty maps via divergence modeling in Fig. 6. Note that we run the provided model of [70] to generate the ranking predictions.

Due to the inconsistent labeling of different ranking dataset and the difficulty in generating precise ranking annotations, we argue that the ranking predictions may not be always reliable (see [70] in the fourth column of Fig. 6). Differently, our general divergence modeling strategy suffers no such issue, which is general and more effective in discovering the less salient regions (see “ABP_R” in Fig. 6). However, the explicit saliency level exploration of the relative saliency task can be beneficial for our general divergence modeling strategy to have a better understanding about human perception, which will be our future work.

5 Conclusion

As a one-to-many mapping task, we argue that salient object detection deserves better exploration with effective divergence modeling strategies. Our proposed “divergence modeling” strategy aims to explore those “discarded” salient regions for better human visual system understanding. Specifically, given multiple saliency annotations for each training image, we aim to generate one majority saliency map, and an uncertainty map, explaining the discarded less salient regions. To achieve this, we propose a general random sampling based strategy and apply it to an ensemble based framework and three latent variable models. Our extensive experimental results show that multiple annotations are necessary for reliable saliency divergence modeling, and our simple random sampling based strategy works effectively in modeling the divergence of saliency. We also apply the proposed strategy to other SOTA SOD models to further explain the flexibility of our model and effectiveness of our strategy. Note that, our saliency divergence modeling approach relies on multiple annotations, which can lead to expensive labeling costs. We will further explore fixation data[74] from multiple observers to achieve both cheaper annotation and more reliable saliency divergence supervision.

References

1. Itti, L., Koch, C., Niebur, E.: A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **20** (1998) 1254–1259
2. Tsotsos, J.K., Culhane, S.M., Wai, W.Y.K., Lai, Y., Davis, N., Nuflo, F.: Modeling visual attention via selective tuning. *Artificial intelligence* **78** (1995) 507–545
3. Olshausen, B.A., Anderson, C.H., Van Essen, D.C.: A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *Journal of Neuroscience* **13** (1993) 4700–4719
4. Koch, C., Ullman, S.: Shifts in selective visual attention: towards the underlying neural circuitry. *Human neurobiology* **44** (1985) 219–27
5. Robinson, D., Petersen, S.: The pulvinar and visual salience. *Trends in Neurosciences* **15** (1992) 127–132
6. Treisman, A.M., Gelade, G.: A feature-integration theory of attention. *Cognitive Psychology* **12** (1980) 97–136
7. Huang, X., Shen, C., Boix, X., Zhao, Q.: Salicon: Reducing the semantic gap in saliency prediction by adapting deep neural networks. In: *IEEE International Conference on Computer Vision (ICCV)*. (2015) 262–270
8. Liu, T., Yuan, Z., Sun, J., Wang, J., Zheng, N., Tang, X., Shum, H.Y.: Learning to detect a salient object. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **33** (2010) 353–367
9. Achanta, R., Hemami, S., Estrada, F., Sussstrunk, S.: Frequency-tuned salient region detection. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2009) 1597–1604
10. Wang, W., Shen, J., Ling, H.: A deep network solution for attention and aesthetics aware photo cropping. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **41** (2018) 1531–1544
11. Zhang, D., Meng, D., Zhao, L., Han, J.: Bridging saliency detection to weakly supervised object detection based on self-paced curriculum learning. In: *International Joint Conference on Artificial Intelligence (IJCAI)*. (2016) 3538–3544
12. Wang, W., Zhou, T., Yu, F., Dai, J., Konukoglu, E., Van Gool, L.: Exploring cross-image pixel contrast for semantic segmentation. In: *IEEE International Conference on Computer Vision (ICCV)*. (2021) 7303–7313
13. Wang, W., Shen, J., Yang, R., Porikli, F.: Saliency-aware video object segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **40** (2017) 20–33
14. Sharma, S., Kiros, R., Salakhutdinov, R.: Action recognition using visual attention. In: *Conference on Neural Information Processing Systems (NeurIPS) Workshop*. (2015)
15. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: *International Conference on Machine Learning (ICML)*. Volume 37. (2015) 2048–2057
16. Oh, S.J., Benenson, R., Khoreva, A., Akata, Z., Fritz, M., Schiele, B.: Exploiting saliency for object segmentation from image level labels. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2017) 4410–4419
17. Zhang, J., Fan, D.P., Dai, Y., Anwar, S., Sadat Saleh, F., Zhang, T., Barnes, N.: Uc-net: Uncertainty inspired RGB-D saliency detection via conditional variational autoencoders. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2020) 8582–8591

18. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. In: Conference on Neural Information Processing Systems (NeurIPS). Volume 28. (2015) 3483–3491
19. Zhang, J., Fan, D.P., Dai, Y., Yu, X., Zhong, Y., Barnes, N., Shao, L.: Rgb-d saliency detection via cascaded mutual information minimization. In: IEEE International Conference on Computer Vision (ICCV). (2021) 4338–4347
20. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: Conference on Neural Information Processing Systems (NeurIPS). Volume 30. (2017)
21. Qin, X., Zhang, Z., Huang, C., Gao, C., Dehghan, M., Jagersand, M.: Basnet: Boundary-aware salient object detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019) 7479–7489
22. Wang, W., Zhao, S., Shen, J., Hoi, S.C., Borji, A.: Salient object detection with pyramid attention and salient edges. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019) 1448–1457
23. Pang, Y., Zhao, X., Zhang, L., Lu, H.: Multi-scale interactive network for salient object detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2020) 9413–9422
24. Zhao, X., Pang, Y., Zhang, L., Lu, H., Zhang, L.: Suppress and balance: A simple gated network for salient object detection. In: European Conference on Computer Vision (ECCV). (2020) 35–51
25. Qin, X., Zhang, Z., Huang, C., Dehghan, M., Zaiane, O.R., Jagersand, M.: U2-Net: Going deeper with nested u-structure for salient object detection. Pattern Recognition (PR) **106** (2020) 107404
26. Tang, L., Li, B., Zhong, Y., Ding, S., Song, M.: Disentangled high quality salient object detection. In: IEEE International Conference on Computer Vision (ICCV). (2021) 3580–3590
27. Li, A., Zhang, J., Lv, Y., Liu, B., Zhang, T., Dai, Y.: Uncertainty-aware joint salient object and camouflaged object detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2021) 10071–10081
28. Liu, Y., Zhang, X.Y., Bian, J.W., Zhang, L., Cheng, M.M.: Samnet: Stereoscopically attentive multi-scale network for lightweight salient object detection. IEEE Transactions on Image Processing (TIP) **30** (2021) 3804–3814
29. Siris, A., Jiao, J., Tam, G.K., Xie, X., Lau, R.W.: Scene context-aware salient object detection. In: IEEE International Conference on Computer Vision (ICCV). (2021) 4156–4166
30. Wei, J., Wang, S., Huang, Q.: F³net: Fusion, feedback and focus for salient object detection. In: AAAI Conference on Artificial Intelligence (AAAI). (2020) 12321–12328
31. Chen, Z., Xu, Q., Cong, R., Huang, Q.: Global context-aware progressive aggregation network for salient object detection. In: AAAI Conference on Artificial Intelligence (AAAI). (2020) 10599–10606
32. Liu, J.J., Hou, Q., Cheng, M.M., Feng, J., Jiang, J.: A simple pooling-based design for real-time salient object detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019) 3917–3926
33. Zhou, T., Fu, H., Chen, G., Zhou, Y., Fan, D.P., Shao, L.: Specificity-preserving rgb-d saliency detection. In: IEEE International Conference on Computer Vision (ICCV). (2021) 4681–4691
34. Sun, P., Zhang, W., Wang, H., Li, S., Li, X.: Deep rgb-d saliency detection with depth-sensitive attention and automatic multi-modal fusion. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2021) 1407–1417

35. Lv, Y., Liu, B., Zhang, J., Dai, Y., Li, A., Zhang, T.: Semi-supervised active salient object detection. *Pattern Recognition (PR)* **123** (2021) 108364
36. Piao, Y., Wang, J., Zhang, M., Lu, H.: Mfnet: Multi-filter directive network for weakly supervised salient object detection. In: *IEEE International Conference on Computer Vision (ICCV)*. (2021) 4136–4145
37. Zhang, J., Yu, X., Li, A., Song, P., Liu, B., Dai, Y.: Weakly-supervised salient object detection via scribble annotations. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2020) 12546–12555
38. Zhang, J., Dai, Y., Xiang, M., Fan, D.P., Moghadam, P., He, M., Walder, C., Zhang, K., Harandi, M., Barnes, N.: Dense uncertainty estimation. *arXiv preprint arXiv:2110.06427* (2021)
39. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: *Conference on Neural Information Processing Systems (NeurIPS)*. Volume 30. (2017)
40. Depeweg, S., Hernández-Lobato, J.M., Doshi-Velez, F., Udluft, S.: Decomposition of uncertainty in bayesian deep learning for efficient and risk-sensitive learning. In: *International Conference on Machine Learning (ICML)*. (2018) 1184–1193
41. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: *International Conference on Machine Learning (ICML)*. (2016) 1050–1059
42. Chitta, K., Feng, J., Hebert, M.: Adaptive semantic segmentation with a strategic curriculum of proxy labels. *arXiv preprint arXiv:1811.03542* (2018)
43. Osband, I., Blundell, C., Pritzel, A., Van Roy, B.: Deep exploration via bootstrapped dqn. In: *Conference on Neural Information Processing Systems (NeurIPS)*. Volume 29. (2016) 4026–4034
44. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: *International Conference on Learning Representations (ICLR)*. (2014)
45. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Conference on Neural Information Processing Systems (NeurIPS)*. Volume 27. (2014) 2672–2680
46. Han, T., Lu, Y., Zhu, S.C., Wu, Y.N.: Alternating back-propagation for generator network. In: *AAAI Conference on Artificial Intelligence (AAAI)*. (2017)
47. Li, Z., Togo, R., Ogawa, T., Haseyama, M.: Variational autoencoder based unsupervised domain adaptation for semantic segmentation. In: *IEEE International Conference on Image Processing (ICIP)*. (2020) 2426–2430
48. Cotterell, R., Eisner, J.: Probabilistic typology: Deep generative models of vowel inventories. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*. (2017) 1182–1192
49. Puscas, M.M., Xu, D., Pilzer, A., Sebe, N.: Structured coupled generative adversarial networks for unsupervised monocular depth estimation. In: *2019 International Conference on 3D Vision (3DV)*. (2019) 18–26
50. Du, Y., Xu, J., Zhen, X., Cheng, M.M., Shao, L.: Conditional variational image deraining. *IEEE Transactions on Image Processing (TIP)* **29** (2020) 6288–6301
51. Ramakrishnan, S., Pachori, S., Gangopadhyay, A., Raman, S.: Deep generative filter for motion deblurring. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*. (2017) 2993–3000
52. Liu, Z., Tang, J., Xiang, Q., Zhao, P.: Salient object detection for rgb-d images by generative adversarial network. *Multimedia Tools and Applications* **79** (2020) 25403–25425

53. Wei, J., Wang, S., Wu, Z., Su, C., Huang, Q., Tian, Q.: Label decoupling framework for salient object detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2020) 13025–13034
54. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016) 770–778
55. Yang, M., Yu, K., Zhang, C., Li, Z., Yang, K.: Denseaspp for semantic segmentation in street scenes. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018) 3684–3692
56. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2020)
57. Luo, Z., Mishra, A., Achkar, A., Eichel, J., Li, S., Jodoin, P.M.: Non-local deep features for salient object detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2017) 6609–6617
58. Hung, W.C., Tsai, Y.H., Liou, Y.T., Lin, Y.Y., Yang, M.H.: Adversarial learning for semi-supervised semantic segmentation. In: British Machine Vision Conference (BMVC). (2018)
59. Neal, R.: Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo* (2011)
60. Wang, L., Lu, H., Wang, Y., Feng, M., Wang, D., Yin, B., Ruan, X.: Learning to detect salient objects with image-level supervision. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2017) 136–145
61. Yan, Q., Xu, L., Shi, J., Jia, J.: Hierarchical saliency detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2013) 1155–1162
62. Yang, C., Zhang, L., Lu, H., Ruan, X., Yang, M.H.: Saliency detection via graph-based manifold ranking. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2013) 3166–3173
63. Li, G., Yu, Y.: Visual saliency based on multiscale deep features. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2015) 5455–5463
64. Fan, D.P., Gong, C., Cao, Y., Ren, B., Cheng, M.M., Borji, A.: Enhanced-alignment measure for binary foreground map evaluation. *International Joint Conference on Artificial Intelligence (IJCAI)* (2018) 698–704
65. Wu, Z., Su, L., Huang, Q.: Stacked cross refinement network for edge-aware salient object detection. In: IEEE International Conference on Computer Vision (ICCV). (2019) 7264–7273
66. Zhou, H., Xie, X., Lai, J.H., Chen, Z., Yang, L.: Interactive two-stream decoder for accurate and fast saliency detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2020) 9141–9150
67. Xu, B., Liang, H., Liang, R., Chen, P.: Locate globally, segment locally: A progressive architecture with knowledge review network for salient object detection. In: AAAI Conference on Artificial Intelligence (AAAI). (2021) 3004–3012
68. Wu, Z., Su, L., Huang, Q.: Decomposition and completion network for salient object detection. *IEEE Transactions on Image Processing (TIP)* **30** (2021) 6226–6239
69. Ma, M., Xia, C., Li, J.: Pyramidal feature shrinking for salient object detection. In: AAAI Conference on Artificial Intelligence (AAAI). (2021) 2311–2318
70. Liu, N., Li, L., Zhao, W., Han, J., Shao, L.: Instance-level relative saliency ranking with graph reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2021)

71. Islam, M.A., Kalash, M., Bruce, N.D.: Revisiting salient object detection: Simultaneous detection, ranking, and subitizing of multiple salient objects. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018) 7142–7150
72. Siris, A., Jiao, J., Tam, G.K., Xie, X., Lau, R.W.: Inferring attention shift ranks of objects for image saliency. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2020) 12133–12143
73. Tian, X., Xu, K., Yang, X., Du, L., Yin, B., Lau, R.W.: Bi-directional object-context prioritization learning for saliency ranking. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2022) 5882–5891
74. Xu, J., Jiang, M., Wang, S., Kankanhalli, M.S., Zhao, Q.: Predicting human gaze beyond pixels. *Journal of vision* **14** (2014)