

Model Calibration in Dense Classification with Adaptive Label Perturbation

Jiawei Liu Changkun Ye Shan Wang Ruikai Cui Jing Zhang Kaihao Zhang Nick Barnes
The Australian National University

Abstract

For safety-related applications, it is crucial to produce trustworthy deep neural networks whose prediction is associated with confidence that can represent the likelihood of correctness for subsequent decision-making. Existing dense binary classification models are prone to being over-confident. To improve model calibration, we propose Adaptive Stochastic Label Perturbation (ASLP) which learns a unique label perturbation level for each training image. ASLP employs our proposed Self-Calibrating Binary Cross Entropy (SC-BCE) loss, which unifies label perturbation processes including stochastic approaches (like DisturbLabel), and label smoothing, to correct calibration while maintaining classification rates. ASLP follows Maximum Entropy Inference of classic statistical mechanics to maximise prediction entropy with respect to missing information. It performs this while: (1) preserving classification accuracy on known data as a conservative solution, or (2) specifically improves model calibration degree by minimising the gap between the prediction accuracy and expected confidence of the target training label. Extensive results demonstrate that ASLP can significantly improve calibration degrees of dense binary classification models on both in-distribution and out-of-distribution data. The code is available on <https://github.com/Carlisle-Liu/ASLP>.

1. Introduction

Binary segmentation aims to differentiate foreground areas from the background in images. Its tasks include Salient Object Detection [86], Camouflaged Object Detection [13], Smoke Detection [79], etc. Performance in these tasks has been significantly advanced using the strong representation powers of Deep Neural Networks (DNNs). However, with complex structures and a tremendous number of parameters, DNNs are prone to over-fitting to training data and producing over-confident predictions in the real world [18]. Such issues can render the model predictions unreliable in decision making or utilisation in downstream tasks.

Recently, a growing body of literature has been pro-

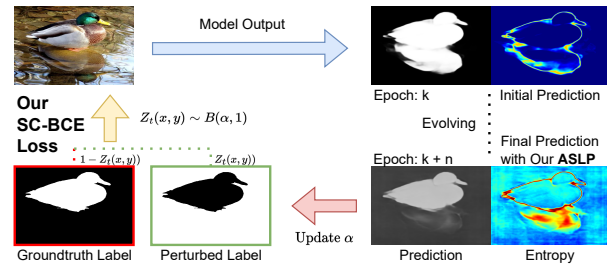


Figure 1: Applying Adaptive Label Perturbation during training can effectively moderate predictions at incorrect areas, highlighting them with high entropy values (red). $Z_t(x, y)$ is a sample-wise Bernoulli variable, parameterized by α , at the t^{th} iteration. After k iterations, update α using Eq. (5) to adjust the likelihood (or level) of label perturbation to increase entropy for incorrect predictions and so correct model calibration. The Perturbed Label (shown inverted) replaces the Groundtruth Label with probability α .

posed to address model mis-calibration problems in DNNs. They can be roughly categorised as: (1) post-hoc operations, such as temperature scaling [18], Platt scaling [54], etc., (2) training objective approaches [27], like MMCE [30], soft calibration objective [27], focal loss [45, 17], and (3) data/label augmentation techniques, e.g. label smoothing [46] and mixup [85]. We propose an Adaptive Label Perturbation which learns a unique label perturbation level for each training image. As illustrated in Fig. 1, training with Adaptive Stochastic Label Perturbation, a form of ALP, can effectively moderate incorrect predictions and highlight them with high entropy values.

Adaptive Label Perturbation employs our proposed Self-Calibrating Binary Cross Entropy (SC-BCE) loss, which unifies label perturbation processes including stochastic approaches (like DisturbLabel [75]), and label smoothing [60] to correct calibration while maintaining classification accuracy. SC-BCE loss is equivalent to a factored combination of (i) a BCE loss w.r.t. groundtruth label, and (ii) a BCE loss w.r.t. a uniform binary categorical distribution. The former enhances dense binary classification performance and the latter improves the model calibration degree. Our method can be connected to Maximum Entropy Inference [23] of classic statistical mechanics, to maximise prediction

entropy with respect to missing information while preserving the classification accuracy on known data.

The proposed Adaptive Label Perturbation (ALP) can approximate Maximum Entropy Inference [23] to maximise prediction entropy while preserving the ideal dense classification performance on known data. This represents a conservative solution that adopts classification accuracy as a proxy for known data and assumes maximum disorder on unknown data. We also present an alternative ALP solution that, instead, takes model calibration degree as a proxy for known data, using a calibration regulariser which constrains the expected confidence of individual supervision signal to not drop below the ideal accuracy on the validation set. This effectively minimises the gap between the distributions of prediction confidence and prediction accuracy, which is the source of model mis-calibration.

Our contributions can be summarised as: (i) We propose Adaptive Stochastic Label Perturbation that learns a sample-wise label perturbation level to improve model calibration; (ii) We present a Self-Calibrating Binary Cross Entropy loss that unifies label perturbation processes including stochastic approaches and label smoothing; (iii) Following Maximum Entropy Inference [23] we show that Adaptive Stochastic Label Perturbation (ASLP_{MEI}), can maximise the prediction entropy while preserving the ideal dense classification accuracy, and (iv) We present an alternative Adaptive Stochastic Label Perturbation (ASLP_{MC}) solution to maximise model calibration degree, which achieves state-of-the-art performance in terms of model calibration degree on both in-distribution and out-of-distribution data. We thoroughly evaluate our method on Salient Object Detection and demonstrate its effectiveness for Camouflaged Object Detection, Smoke Detection and Semantic Segmentation.

2. Related Works

Model Calibration: Model calibration methods can be roughly divided into three categories of approach: (1) post hoc processing (2) training object, and (3) input/label augmentation. The first category utilises a validation set to align the prediction confidence distribution with prediction accuracy distribution. It includes histogram binning [82], isotropic regression [83], Platt scaling [54, 49], Bayesian binning into quantiles [47], Dirichlet scaling [28], mix-n-match [87] and temperature scaling [18, 81].

The second category focuses on designing training objectives that (in)directly improve model calibration degree. Some methods address the in-continuity of expected calibration error, a widely adopted model calibration measure, and propose trainable calibration objectives like maximum mean calibration error [30], soft calibration objective [27]. Other works discover that certain existing training objects are beneficial to model calibration, *e.g.* Brier loss [4, 10], confidence penalty [52] and focal loss [45, 17].

The third category employs data or label augmentation techniques to regularise the prediction confidence distribution. Mixup [85, 62] explores the neighbourhood of training data through random interpolation of input images and associated labels to improve model calibration degree. Label smoothing [46] augments the one-hot training labels with softer versions to prevent the model being over-confident.

Salient Object Detection: Inspired by pioneering work [22], traditional Saliency Object Detection (SOD) methods rely on various heuristic priors with handcraft features to explore low-level cues [1, 25, 26, 38, 66]. However, these methods cannot cope with complex scenes because of the limited representation ability of handcrafted features [3]. Recently, deep learning based SOD methods broke the bottleneck of traditional methods due to the powerful capability of neural networks, achieving improved performance [6, 24, 55, 76, 37, 64, 68]. Early deep SOD methods use multi-layer perception to predict a map with a pixel-wise score for each image [71, 8]. These approaches rely on fixed fully connected layers and thus severely limit the ability of spatial information extraction. Later methods address this issue via using fully convolutional networks (FCNs) [40].

Most contemporary SOD methods are designed based on FCNs with various schemes to improve performance. One of the most popular strategies is to fuse multi-scale information extracted from different layers and aggregate them in the network [91, 74, 16, 51, 96, 90]. Attention modules are also applied to capture powerful multi-scale features via exploring relationship between local and global information [37, 53, 95, 21]. Training SOD networks using auxiliary supervision is also a popular strategy [70, 70, 61]. For example, the body map and detail map are utilized by [70] to help the network focus on center areas and edges, respectively. Skeleton [35] and uncertainty [61] are applied to the training processing due to their important roles in taking photos.

Noisy Label: Noisy labels refer to incorrect ground truth classes/values in classification/regression tasks. They arise from data collection or annotation processes, and exist commonly in real-world datasets [2]. Efforts are put to identify the noisy labels and exclude them from network training in various computer vision tasks, deeming their incorporation as harmful. [88] proposes a framework that learns from noisy labels, being a collection of predictions from classic SOD methods. The framework approximates the noise distribution in order to recover clean labels for model training.

Differently from data augmentation techniques that are applied simultaneously to training samples and corresponding labels to generate more training data, one may artificially corrupt the label. We refer to this category of approaches as *label perturbation*, which includes label smoothing approaches [60], and DisturbLabel, Xie *et al.* [75]. In image classification, Xie *et al.* [75] shows that randomly replacing training labels with a prior distribution

leads to a regularising effect, preventing overfitting. Our work is different from [75] by employing labels corrupted to different scales to enhance the model calibration degrees for both in-distribution and out-of-distribution data. Further, in performing this, we assume noise that varies with different samples, making our method more adaptable.

3. Proposed Method

We first lay out the task setting in Sec. 3.1. Then we introduce our proposed Self-Calibrating Binary Cross Entropy loss in Sec. 3.2 and prove its connection to Maximum Entropy Inference [23] in Sec. 3.3. Lastly, we detail our major contribution - Adaptive Label Perturbation in Sec. 3.4.

3.1. Task Setting

Binary segmentation problems aim to differentiate between foreground object(s) and background. They can be formulated as a pixel-wise binary classification problem. Given an independent and identically distributed (i.i.d) training dataset $\mathcal{D}_{\text{TR}} = \{x_i, y_i\}_{i=1}^N$ drawn from an unknown joint distribution of training images and groundtruth labels $P(\mathcal{X}, \mathcal{Y})$, a neural network model parameterised by θ is employed to predict labels for an input image $x \in \mathcal{X}$: $f_\theta(x) \in (0, 1)^{1 \times H \times W}$. We use \hat{y} and $P_{\hat{y}}$ to denote the winning class and its associated probability respectively. The groundtruth label $y \in \{0, 1\}^{1 \times H \times W}$ represents the foreground pixels with “1” and background with “0”. In the following equations, we omit the spatial indexes H and W for simplicity. A perfectly calibrated model has $P(\hat{y} = y | P_{\hat{y}}) = P_{\hat{y}}, \forall P_{\hat{y}} \in (0, 1)$. That is, in the entire range of prediction probabilities, prediction with probability $P_{\hat{y}}$ has exactly $P_{\hat{y}}$ chance to be correct. The calibration degree of a model $f_\theta(\cdot)$ over a distribution \mathcal{D} is quantified with Expected Calibration Error (ECE), defined as $\mathbb{E}_{f_\theta(x)}[|P(\hat{y} = y | f_\theta(x)) - f_\theta(x)|]$.

3.2. Self-Calibrating Binary Cross Entropy Loss

We propose a Self-Calibrating Binary Cross Entropy (SC-BCE) loss that unifies Label Smoothing [60], DisturbLabel [75] and Stochastic Label Perturbation as:

$$\begin{aligned} & \mathcal{L}_{\text{SC-BCE}}(\theta, X, Y, \alpha, \beta) \\ &= \mathbb{E}_{x, y \in X, Y} \left[(1 - Z_t(x, y)) \cdot \mathcal{L}_{\text{BCE}}(\theta, x, y) \right. \\ & \quad \left. + Z_t(x, y) \cdot \mathcal{L}_{\text{BCE}}(\theta, x, p(y, \beta)) \right] \quad (1) \\ & \beta \in [0, 2], \quad \alpha \in [0, \frac{1}{\beta}], \end{aligned}$$

where $Z_t(x, y) \sim B(1, \alpha)$ follows a Bernoulli distribution with α probability to be 1 and $1 - \alpha$ chance to be 0, t denotes the training epoch, $\alpha \in [0, \frac{1}{\beta}]$ and $\beta \in [0, 2]$ are Label Perturbation Probability (LPP) and Label Perturbation Strength

(LPS) respectively, $p(y, \beta) = (1 - \beta) \cdot y + \frac{\beta}{2}, \beta \in [0, 2]$ is a perturbed label and $\mathcal{L}_{\text{BCE}}(\theta, x, y)$ ¹ is a Binary Cross Entropy (BCE) loss computed for training pair (x, y) . For $\alpha = 1$, the label perturbation equation follows the label smoothing equation for a binary label [60]. In the proposed SC-BCE loss, different label perturbations can be applied by setting (i) Label Smoothing [60]: $\alpha = 1$ and $\beta \in [0, 1)$, (2) DisturbLabel: $\beta = 1$ and $\alpha \in (0, 1)$, and (3) Stochastic Label Perturbation (SLP): $\beta \in (0, 2]$ and $\alpha \in (0, \frac{1}{\beta})$. For example, Hard Inversion (HI) that inverts the label category as shown in Fig. 2, can be stochastically applied by setting $\beta = 2$ and $\alpha \in (0, \frac{1}{2})$.

In the implementation of SLP, the supervision for an individual training image in each epoch is sampled by drawing from a Bernoulli distribution. That is, the individual supervision can take the form of the groundtruth label or perturbed label in each training iteration. The overall function of SLP can be connected to that of a smoothed label by taking expectation of the Bernoulli variable: $\mathbb{E}_{Z_t}[(1 - Z_t(y)) \cdot Y + Z_t(y) \cdot p(y, \beta)], \forall y \in Y$. Taking the expectation over the Bernoulli variation in each iteration is too expensive to implement in model training. Instead, following [75], the expectation of stochastically perturbed label is approximated by taking expectation across training epochs: $\mathbb{E}_{t \in T}[(1 - Z_t(y)) \cdot Y + Z_t(y) \cdot p(y, \beta)], \forall y \in Y$, where T is the total number of training epochs and $Z_t(y)$ is a variable drawn from a Bernoulli distribution for the t^{th} epoch.

3.3. Maximum Entropy Inference

Maximum Entropy Inference (MEI), assuming minimum distribution commitment in respect to missing information, was initially proposed by Jaynes [23]. That is, the probability distribution should have maximum Shannon entropy subject to the partially available information. Thus, in the complete absence of information, Shannon entropy for a binary prediction defined in Eq. (2) should be maximised:

$$\begin{aligned} \mathbb{H}(f_\theta(X)) &= \mathbb{E}_{x \in X} \left[-f_\theta(x) \cdot \log f_\theta(x) - \right. \\ & \quad \left. (1 - f_\theta(x)) \cdot \log(1 - f_\theta(x)) \right] \quad (2) \end{aligned}$$

For salient object detection and other binary segmentation problems, maximising Eq. (2) can be achieved with a binary uniform categorical distribution.

The proposed SC-BCE loss can be transformed into a factored combination of a BCE loss w.r.t. groundtruth label (the constraints of the data) and a BCE loss w.r.t. a binary uniform categorical distribution (See Appendix A.2 for

¹ $\mathcal{L}_{\text{BCE}}(\theta, x, y) = -y \cdot \log(f_\theta(x)) - (1 - y) \cdot \log(1 - f_\theta(x))$

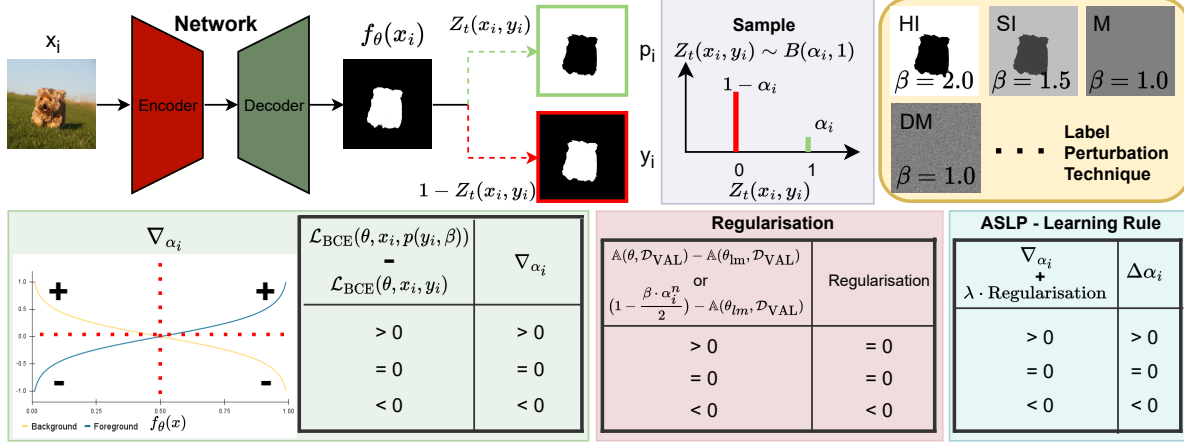


Figure 2: The method overview is comprised of model implementation (top) and Adaptive Stochastic Label Perturbation (ASLP) learning rule (bottom). In each iteration, the model uses a Bernoulli variable to sample a supervision, which can be a groundtruth label y_i or a perturbed label p_i and computes a sample-specific ∇_{α_i} based on the prediction. Regularisation is computed with Eq. 5 (ASLP_{MEI}) or Eq. 6 (ASLP_{MC}) on a validation set after each training epoch. The ASLP learning rule combines Grad- α_i and a factored regularisation to update label perturbation probability α_i for each individual sample.

derivation) as:

$$\begin{aligned} & \mathcal{L}_{\text{SC-BCE}}(\theta, X, Y, \alpha, \beta) \\ &= \mathbb{E}_{x, y \in X, Y} \left[(1 - \beta Z_t(x, y)) \cdot \mathcal{L}_{\text{BCE}}(\theta, x, y) \right. \\ & \quad \left. + \beta Z_t(x, y) \cdot \mathcal{L}_{\text{BCE}}(\theta, x, u) \right], \end{aligned} \quad (3)$$

where $Z_t(x, y) \sim B(1, \alpha)$, u is a binary uniform categorical distribution, and minimising the second term pushes the prediction distribution towards a uniform binary categorical distribution, equivalently maximising the inference entropy. Therefore, our proposed SC-BCE loss, a combination of a regular BCE loss and a BCE loss with a perturbed label, effectively performs a type of MEI. That is, the regular BCE loss component improves the model’s binary classification accuracy in the presence of information while the perturbed label maximizes prediction entropy with respect to missing information in order to close the gap between the available training data and the entire data distribution.

3.4. Adaptive Label Perturbation

Stochastic Label Perturbation (SLP) uses a single label perturbation probability and perturbation strength for the entire training dataset. However, this approach cannot adapt to predictive error that varies for different input images. To address this, we propose an Adaptive Stochastic Label Augmentation (ASLP) method to adjust the label augmentation probability for individual training samples. That is we allow the variable in Eq. (1) to be drawn from a per training image Bernoulli distribution with sample-specific label perturbation probability as:

$$Z_t(x, y) \sim B(1, \alpha_{x, y}), \forall x, y \in X, Y \quad (4)$$

where $\alpha_{x, y}$ is the label perturbation probability for sample (image-label) pair (x, y) . Initially, we set all label perturbation probabilities to $\{\alpha_i = 0\}_{i=1}^N$ and train a model with a regular BCE loss without label augmentation techniques, which is equivalent to $\mathcal{L}_{\text{SC-BCE}}(\theta, X, Y, \alpha = 0, \beta = 0)$. The trained model has weight θ_{lm} and its accuracy on the validation set, $\mathbb{A}(\theta_{\text{lm}}, \mathcal{D}_{\text{VAL}})$, is held as an ideal performance. Subsequently, we select a label perturbation technique and continue to train the model with SC-BCE loss with a learning rule to update the label perturbation probability for individual training samples.

We propose the learning rule for α , (ASLP_{MEI}) to approximate maximum entropy inference. The rule has two components: (1) $\nabla_{\alpha_i} = (2/\beta) \cdot \partial \mathbb{E}_{Z_t(x, y)} [\mathcal{L}_{\text{SC-BCE}}(X, Y, \theta, \{\alpha_i\}_{i=1}^N, \beta)] / \partial \alpha_i$ is the derivative of the expectation of SC-BCE over the Bernoulli variable w.r.t. α_i . We divide this by $\beta/2$ to ensure that different perturbation techniques (varying β values) have the same convergence speed (See derivation in Appendix A.3), and (2) Accuracy Regularization to encourage maintenance of prediction accuracy. The rule is:

$$\begin{aligned} \alpha_i^{n+1} &= \alpha_i^n + \eta \cdot (\nabla_{\alpha_i} + \lambda \cdot \text{Reg}_A), \quad \text{for } i = 1, \dots, N, \\ \nabla_{\alpha_i} &= \frac{2 \cdot (\mathcal{L}_{\text{BCE}}(\theta, x_i, p(y_i, \beta)) - \mathcal{L}_{\text{BCE}}(\theta, x_i, y_i))}{\beta}, \\ \text{Reg}_A &= \min \left(\frac{\mathbb{A}(\theta, \mathcal{D}_{\text{VAL}}) - \mathbb{A}(\theta_{\text{lm}}, \mathcal{D}_{\text{VAL}})}{\mathbb{A}(\theta_{\text{lm}}, \mathcal{D}_{\text{VAL}})}, 0 \right), \end{aligned} \quad (5)$$

where η and λ are hyperparameters controlling the updating pace of label perturbation probability and the regularisation strength respectively, $\mathbb{A}(\theta, \mathcal{D}_{\text{VAL}})$ and $\mathbb{A}(\theta_{\text{lm}}, \text{and } \mathcal{D}_{\text{VAL}})$ denote the current and ideal accuracy on the validation set separately. ∇_{α_i} aims to increase label perturbation proba-

bility to confident and correct samples and otherwise for incorrectly classified samples. For example, it returns a large positive value for correct predictions with small BCE loss value w.r.t. to groundtruth label y_i and large BCE loss value w.r.t. to perturbed label $p(y_i, \beta)$. The ‘‘Accuracy Regularisation’’ (Reg_A) is designed to reduce the overall perturbation probability if the accuracy on the validation set reduces to be below the local minima. It returns 0 if there is no accuracy drop on the validation set and a large decrease will overwhelm the ∇_{α_i} value and reduce the sample label perturbation probability. Intuitively, ASLP_{MEI} aims to construct a model that preserves the ideal classification accuracy while otherwise maximising the entropy [23]. Note that which particular examples are classified correctly are able to change, but the accuracy is constrained to remain the same. Intuitively, having a model that better captures ignorance may lead to changes in the treatment of test examples that are distant from training distribution. Note, however that adopting classification accuracy as proxy for known data and otherwise maximizing entropy is a conservative strategy and we find that it results in the model being significantly under-confident.

The model mis-calibration arises from the distribution mismatch between prediction confidence and prediction accuracy [45]. We offer an alternative model that uses the model calibration as proxy for known data and maximises the prediction entropy in respect to unknown data ASLP_{MC} . The learning rule replaces the ‘‘Accuracy Regularisation’’ in Eq. (5) with a ‘‘Calibration Regularisation’’ (Reg_C) as:

$$\text{Reg}_C = \min\left(\left(1 - \frac{\beta \cdot \alpha_i^n}{2}\right) - \mathbb{A}(\theta_{im}, \mathcal{D}_{\text{VAL}}), 0\right), \quad (6)$$

where $1 - (\beta \cdot \alpha_i^n / 2)$ denotes the expected confidence of the perturbed label (Derivation in Appendix A.4). For example, a foreground label ‘‘1’’ with 5% chance of being inverted to ‘‘0’’ has an expected confidence of 0.95. ‘‘Calibration Regularisation’’ constrains the expected confidence of perturbed label of each sample to not drop below the ideal classification accuracy on validation set, preventing the model from becoming under-confident. Note that we can also have an updating rule ALS_{MC} to learn per-image label perturbation strength (adaptive β and fixed $\alpha = 1$) (See Appendix A.5).

4. Experiments and Results

We verify the proposed method primarily on Salient Object Detection and also implement it for Camouflaged Object Detection, Smoke Detection and Semantic Segmentation tasks and report their results in the Appendices.

4.1. Implementation Details

Evaluation Metrics: We use Equal-Width Expected Calibration Error (ECE_{EW}) [18] and Equal-Width Overconfidence Error (OE_{EW}) [62] with 10 bins ($B = 10$) to

evaluate the model calibration degrees. Additionally, we adopt ECE_{EM} [48], $\text{ECE}_{\text{DEBIAS}}$ [29] and $\text{ECE}_{\text{SWEEP}}$ [56] to corroborate with the results of ECE_{EW} (See Appendix C).

Datasets: The proposed methods are trained with the DUTS-TR [63] training dataset. It is divided into a training set $|\mathcal{D}_{\text{TR}}| = 9,553$ and validation set $|\mathcal{D}_{\text{VAL}}| = 1,000$. We use six testing datasets, including DUTS-TE [63], DUT-OMRON [80], SOD [44], PASCAL-S [34], ECSSD [78], HKU-IS [33], to evaluate the model calibration degree.

Compared Methods: We compare with both SOD models and model calibration methods in terms of model calibration degrees. The SOD models include: MSRNet [32], SRM [65], Amulet [92], BMPM [91], DGRL [67], PAGR [93], PiCANet [37], CPD [73], BASNet [55], EGNet [94], AFNet [15], PoolNet [36], GCPANet [7], MINet [51], F³Met [69], EBMGSOD [89], ICON [97], EDN [72]. We evaluate ECE on their published prediction results, or results produced with their released model weights. We also compare with model calibration methods include: Temperature Scaling (TS) [18], Brier Loss [4], MMCE [30], Label Smoothing [46], Mixup [62], Focal Loss [45] and AdaFocal [17] implemented on our baseline model.

Baseline Model Structure: We implement our method in the Pytorch framework. Our model has a simple U-Net structure, comprising of a ResNet50 encoder [20] and a decoder, where the former is initialised with ImageNet-pretrained weights and the later by default. We also experiment with VGG16 [58] and Swin transformer [39] encoders and report their results in Appendix I.

Label Perturbation Techniques: We experiment with four different label perturbation strategies: (1) Hard Inversion $T(Y, \alpha, \beta = 1.0)$, (2) Soft Inversion (SI) $T(Y, \alpha, \beta = 0.75)$, (3) Moderation (M) $T(Y, \alpha, \beta = 0.5)$ and Dynamic Moderation (DM) $T(Y, \alpha, \beta = 0.5)$ with additional Gaussian noise. See Appendix H.1 for implementation details.

Training Details: Each model is trained with an Adam optimiser for 30 epochs. The learning rate is initialised to 2.5×10^{-5} , and decays by a factor of 0.9 for each epoch after the 10th epoch. All training images are scaled to 384×384 . Basic data augmentation techniques, including random flipping, random translation and random cropping, are applied.

Hyperparameters: The reported model calibration results associated with ASLP are obtained by setting $\eta = 0.002$ and $\lambda = 2,000$. We set the number of bins to $B = 10$ for ECE and OE evaluation metrics.

4.2. Model Calibration Degree Performance

Tab. 1 presents the calibration degree of existing SOD models, existing model calibration methods and our proposed technique on the six SOD testing datasets. Our proposed ASLP_{MC} , designed to optimise the model calibration degree, achieves the best ECE performances on all testing datasets. In addition, ASLP_{MC} also obtains the second-best

Table 1: Salient object detection model calibration degree benchmark. Results are evaluated in with ECE_{EW} and OE_{EW} with 10 bins (units in %). See Appendix C for evaluations with ECE_{EM} [48], ECE_{DEBIAS} [29] and ECE_{SWEEP} [56].

Methods	Year	DUTS-TE [63]		DUT-OMRON [80]		PASCAL-S [34]		SOD [44]		ECSSD [78]		HKU-IS [33]		
		ECE ↓	OE ↓	ECE ↓	OE ↓	ECE ↓	OE ↓	ECE ↓	OE ↓	ECE ↓	OE ↓	ECE ↓	OE ↓	
SOD Methods	MSRNet [32]	2017	2.57	2.34	3.32	3.16	3.44	3.23	6.42	6.14	0.97	0.94	0.92	0.87
	SRM [65]	2017	4.02	3.72	4.19	3.96	4.88	4.59	9.93	9.58	2.53	2.35	1.86	1.72
	Amulet [92]	2017	5.67	5.28	5.84	5.49	5.76	5.43	10.03	9.59	2.56	2.42	1.98	1.87
	BMPM [91]	2018	3.74	3.52	4.52	4.37	4.88	4.68	8.16	7.93	1.95	1.89	1.58	1.53
	DGRL [67]	2018	4.12	3.86	4.41	4.21	5.01	4.77	8.44	8.20	2.13	2.02	1.63	1.53
	PAGR [93]	2018	4.04	3.79	5.14	4.96	5.64	5.37	12.17	11.87	2.84	2.70	1.62	1.54
	PiCANet [37]	2018	5.12	4.90	4.84	4.70	8.14	7.92	10.50	10.30	3.48	3.39	2.55	2.47
	CPD [73]	2019	3.97	3.78	4.20	4.06	5.37	5.17	9.65	9.39	2.29	2.19	1.99	1.90
	BASNet [55]	2019	5.00	4.86	4.93	4.83	6.50	6.36	10.40	10.27	2.74	2.70	2.30	2.26
	EGNet [94]	2019	3.33	3.14	3.66	3.50	5.42	5.19	8.04	7.79	1.98	1.88	1.47	1.40
	AFNet [15]	2019	3.95	3.74	4.25	4.09	5.06	4.84	8.15	8.02	2.38	2.27	1.87	1.78
	PoolNet [36]	2019	3.33	3.12	3.86	3.70	5.32	5.07	8.14	7.87	2.00	1.90	1.82	1.75
	GCPANet [7]	2020	3.18	2.99	3.99	3.84	4.16	3.97	7.05	6.88	1.61	1.54	1.27	1.21
	MINet [51]	2020	3.65	3.48	4.45	4.29	4.94	4.75	8.01	7.89	2.13	2.03	1.74	1.65
	F ³ Met [69]	2020	3.67	3.50	4.25	4.10	4.85	4.67	7.95	7.78	2.26	2.16	1.92	1.83
	EBMGSOD [89]	2021	3.45	3.29	4.11	3.95	4.79	4.61	7.48	7.30	2.14	2.05	1.79	1.70
	ICON [97]	2021	2.89	2.76	3.84	3.71	4.08	3.95	6.70	6.55	1.56	1.49	1.38	1.32
	PFSNet [43]	2021	2.94	2.72	3.95	3.81	4.45	4.27	7.59	7.39	2.41	2.25	2.06	1.96
EDN [72]	2022	3.62	3.47	4.02	3.90	4.89	4.74	8.81	8.66	2.20	2.13	1.65	1.58	
Model Calibration Methods	Brier Loss [4]	1950	2.77	2.58	3.55	3.38	3.90	3.70	6.40	6.16	1.37	1.30	1.04	0.99
	Temperature Scaling [18]	2017	2.53	2.34	3.18	3.03	3.56	3.36	6.32	6.05	0.96	0.93	0.83	0.70
	MMCE [30]	2018	2.86	2.67	3.56	3.41	4.00	3.81	6.85	6.63	1.41	1.35	1.18	1.13
	Label Smoothing [46]	2019	2.00	1.79	2.89	2.71	3.04	2.83	5.97	5.69	0.83	0.68	0.82	0.47
	Mixup [62]	2019	2.45	2.25	3.41	3.23	3.13	2.99	5.82	5.70	1.41	0.18	3.83	0.05
	Focal Loss [45]	2020	2.25	2.08	3.10	2.82	3.40	3.13	6.21	5.98	1.41	1.03	1.24	0.77
	AdaFocal [17]	2022	1.61	1.41	2.31	1.84	2.53	2.27	5.88	5.47	1.63	0.79	1.35	0.52
Our Methods	ASLP _{MC}	2023	1.40	1.22	1.99	1.83	2.31	2.10	5.50	5.17	0.48	0.20	0.79	0.17
	ASLP _{MEI}	2023	27.9	0.01	26.0	0.00	26.1	0.00	22.4	0.00	29.9	0.00	30.5	0.00

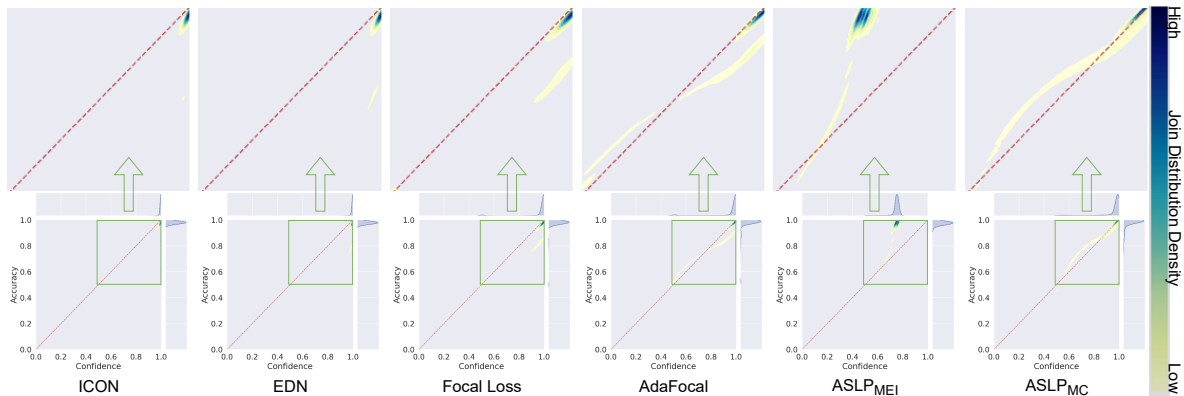


Figure 3: Joint distribution of prediction confidence (horizontal axis) and prediction accuracy (vertical axis) on the DUTS-TE dataset. A perfectly calibrated model has an identical confidence distribution and accuracy distribution, denoted as the oracle (diagonal red dotted line). The joint distribution of a better calibrated model is more aligned with the oracle line, especially its high density area. See Appendix D for results of other methods and on other testing datasets.

OE performances on all six testing datasets, outperformed only by our ASLP_{MEI}. On the other hand, ASLP_{MEI}, though almost eliminates the over-confidence issue completely, is significantly mis-calibrated on the six testing datasets. This can be attributed to it being significantly under-confident rather than over-confident. The observed performances of

ASLP_{MEI} are in accordance with its design - assuming minimum distribution commitment with respect to missing information. That is, in the presence of limited training data, to maximise the prediction entropy while maintaining the prediction accuracy for in-distribution data.

Fig. 3 presents the joint distribution of prediction con-

confidence and prediction accuracy of some of the best calibrated methods and our proposed technique on DUTS-TE dataset (See Appendix D for other testing datasets and other methods). Existing SOD methods produce extremely confident predictions whose confidence scores are nearly 100% for the majority of samples. On the contrary, the prediction accuracy is on average lower than prediction confidence, resulting in the model being over-confident. On the other hand, existing model calibration methods are generally more calibrated than the SOD methods which in general do not strive to improve model calibration degree.

Our proposed ASLP_{MC} produces the most calibrated model whose joint distribution is closer to the oracle than those of the compared calibration methods and SOD methods. AdaFocal [17] produces the second most calibrated model on DUTS-TE. However, the majority of their joint distribution (blue high density area) is slightly to the right bottom of the oracle line, making their model slightly less calibrated. Whereas the high density area of our joint distribution is well aligned with the oracle line, showing ASLP_{MC} is more calibrated. Despite the small difference on the plot, ASLP_{MC} improves over AdaFocal by more than 10% in terms of ECE scores on DUTS-TE dataset. We can also observe that ASLP_{MEI} is significantly under-confident for in-distribution data with its joint distribution being at the top-left side of the oracle line. Its prediction confidences are limited to between 70% and 80% while the prediction accuracies are generally above 90%.

4.3. Model Calibration Degrees on Out-of-Distribution Dataset

We compare our proposed method with existing model calibration methods in terms of model calibration degrees on Out-of-Distribution data. We consider texture images, where salient objects are completely absent, as OoD samples for the SOD task. We use Describable Texture Dataset [9] to evaluate the model calibration degrees on OoD samples. Fig. 4 shows the joint distribution of prediction confidence and prediction accuracy of various model calibration methods and our proposed techniques. It can be seen that the baseline model produces extremely confident predictions for OoD data. However, its accuracy is only 41.88%, worse than a uniform prior in a binary classification task. We also observe that Temperature Scaling does not calibrate the model under data distribution shift in accordance with literature [50, 45]. Our ASLP_{MC}, being the most calibrated for in-distribution data, is also more calibrated on the OoD samples than the existing model calibration methods by a large margin as shown in Tab. 2. On the other hand, our ASLP_{MEI} is more successful in handling OoD data. It is the most calibrated on OoD data, with a larger proportion of the distribution aligned with the oracle line. As shown in Tab. 2, it outperforms existing model calibration methods in

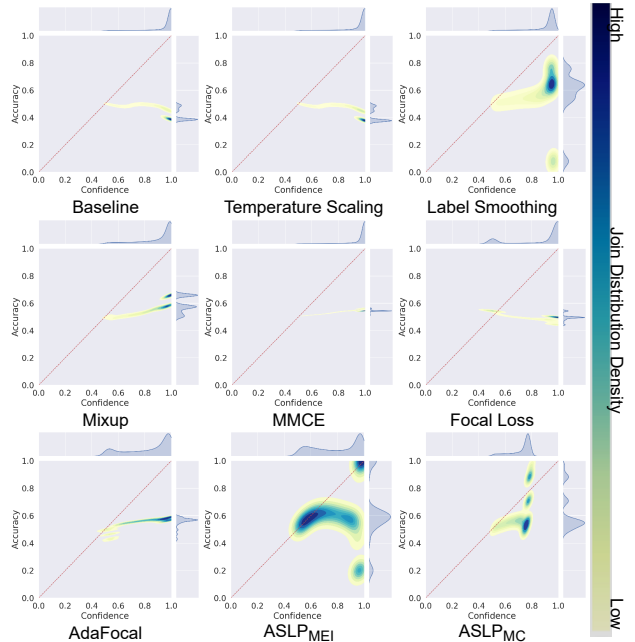


Figure 4: Joint distribution of prediction confidence (horizontal axis) and prediction accuracy (vertical axis) on the Describable Texture Dataset [9], of compared model calibration methods and our proposed ASLP_{MC} and ASLP_{MEI}.

terms of both ECE and OE by significant margins. This can be attributed to its minimum distribution assumption in the presence of limited training data.

Table 2: Model calibration methods and our ASLP_{MC} and ASLP_{MEI} are evaluated on the Out-of-Distribution dataset, Describable Texture Dataset [9], in terms of ECE_{EW} and OE_{EW} with 10 bins, and Accuracy (ACC).

Method	Evaluation (%)		
	ECE ↓	OE ↓	ACC ↑
Baseline	52.36	51.05	41.88
Brier Loss [4]	38.85	37.18	53.62
Temperature Scaling [18]	51.95	50.46	41.59
Label Smoothing [46]	37.22	35.48	55.41
MMCE [30]	40.64	39.67	54.39
Mixup [62]	31.07	29.10	58.71
Focal Loss [45]	40.01	38.43	49.71
AdaFocal [17]	27.55	25.07	55.39
ASLP _{MC}	18.31	16.37	61.93
ASLP _{MEI}	13.43	8.40	62.47

4.4. Discussion

Adaptive Stochastic Label Perturbation: We study the effect of ASLP on ECE and OE and present the experimental results in Tab. 3. It shows that ASLPs significantly outperforms the baseline model, “B”, which does not adopt any model calibration measures. In addition, ASLPs also outperforms their static counterparts (SLPs) which use a single

Table 3: Ablation: Effect of Stochastic Label Perturbation (SLP) and Adaptive Stochastic Label Perturbation (ASLP) with different label perturbation techniques on the model calibration degrees evaluated on Expected Calibration Error (ECE) and Over-confidence Error (OE). The proposed ASLP is generalised to an Adaptive Label Smoothing (ALS) technique that adaptively tunes the label softening scale (β_{ada}).

Methods	Perturbation Params			DUTS-TE [63]		DUT-OMRON [80]		PASCAL-S [34]		SOD [44]		ECSSD [78]		HKU-IS [33]	
	α	β	e	ECE ↓	OE ↓	ECE ↓	OE ↓	ECE ↓	OE ↓	ECE ↓	OE ↓	ECE ↓	OE ↓	ECE ↓	OE ↓
Baseline (“B”)	0	0	0	3.48	3.29	4.17	4.02	4.60	4.41	7.42	7.17	1.93	1.86	1.64	1.59
SLP _{HI} ^{$\alpha=0.01$}	0.01	1.0	✗	2.21	1.84	2.96	2.78	3.11	2.82	6.09	5.80	1.03	0.68	1.01	0.53
SLP _{SI} ^{$\alpha=0.02$}	0.02	0.75	✗	2.25	2.05	3.00	2.82	3.05	2.83	6.40	6.09	0.93	0.84	0.87	0.60
SLP _M ^{$\alpha=0.03$}	0.03	0.5	✗	2.24	2.03	3.17	2.97	3.41	3.20	6.26	5.97	0.83	0.77	0.96	0.81
SLP _{DM} ^{$\alpha=0.03$}	0.03	0.5	✓	2.29	2.09	3.00	2.83	3.47	3.24	6.72	6.43	1.13	1.04	0.96	0.80
LS ^{$\beta=0.03$}	1.0	0.03	✗	2.20	1.99	3.09	2.91	3.24	3.03	6.27	5.99	1.03	0.78	0.92	0.67
ASLP _{MC} ^{HI}	α_{ada}	1.0	✗	1.40	1.22	1.99	1.83	2.31	2.10	5.50	5.17	0.48	0.20	0.79	0.17
ASLP _{MC} ^{SI}	α_{ada}	0.75	✗	1.51	1.29	2.14	1.95	2.29	2.07	5.12	4.80	0.61	0.34	0.84	0.22
ASLP _{MC} ^M	α_{ada}	0.5	✗	1.47	1.27	1.87	1.80	2.37	2.13	5.63	5.29	0.51	0.23	0.80	0.20
ASLP _{MC} ^M	α_{ada}	0.5	✓	1.64	1.20	1.94	1.75	2.03	1.81	4.14	3.84	0.80	0.42	0.87	0.42
ALS	1.0	β_{ada}	✗	1.46	1.25	2.07	1.87	2.30	2.10	5.44	5.18	0.61	0.25	0.81	0.32

α for the entire dataset. This can be attributed to the approach modelling variance of noise with input image.

Generalisation of Adaptive Label Smoothing: We generalise the proposed ASLP to label smoothing, developing an Adaptive Label Smoothing (ALS) that fixes the label perturbation probability to 100%, akin to the label smoothing technique [46], and tunes a smoothing factor for each training sample. As shown in Tab. 3, ALS effectively reduces the ECE and OE scores over its static version LS ^{$\beta=0.03$} , and achieves similar performances with ASLP_{MC} approaches on the six testing datasets. It manifests that our proposed ASLP can be generalised onto other label perturbation techniques as a measure to calibrate the SOD models.

Compatibility with SOTA SOD Models We retrain EBMGSOD [89], ICON [97] and EDN [72] with the proposed ASLP_{MC} and find significant improvements in terms of model calibration degrees without compromising their classification abilities (See Appendix E).

Compatibility with Different Backbones: We demonstrate that our proposed method is also compatible with VGG16 [58] and Swin transformer [39] backbones. See Appendix I for details.

Effectiveness in Other Dense Binary Classification Tasks: Experiments on Smoke Detection [79] and Camouflaged Object Detection [14] demonstrate that our method can be generalised to dense binary classification tasks to improve model calibration degrees. See Appendix F.

Generalisation to Multi-Class Segmentation task: Experiments on Semantic Segmentation [11] demonstrate that our method can also be generalised to dense multi-class classification tasks. See Appendix G.

4.5. Hyperparameters

Static Stochastic Label Perturbation: Tab. 3 shows that, under a small label perturbation probability, the four label perturbation techniques can alleviate the model over-confidence issues of the baseline model, “B”, on the six

testing datasets. They also achieve similar results to Label Smoothing [46], setting $\beta = 0.03$ and $\alpha = 1$. Each SLP has a wide range of effective label perturbation probabilities that improves model calibrations (See Appendix H.2 Tab. 13), and these improvements do not sacrifice the model’s classification performance (see Appendix H.3). Larger values of the label perturbation probability eventually lead to increasing ECE scores as the model transitions to being under-confident (see Appendix H.2).

Updating Rate η : ASLP_{MC} models trained with $\eta \in [0.0002, 0.005]$ are generally stable, producing similar calibration degrees and classification performances. Values smaller than 0.001 require longer training and high values lead to sub-optimal results (See Appendix J).

Regularisation Strength λ : λ spanning between 500 and 10,000 are optimal. Very high values for λ can lead to oscillation resulting in poor performance (See Appendix J).

5. Conclusion

This work first introduces a Self-Calibrating Binary Cross Entropy loss that unifies label perturbation processes including stochastic approaches and label smoothing to improve model calibration while preserving classification accuracy. We further propose an Adaptive Stochastic Label Perturbation that learns a unique label perturbation level for individual training image. Following Maximum Entropy Inference, ASLP adopts classification / calibration as proxy for known data and maximises the prediction entropy with respect to missing data. The proposed ASLP_{MC} improves model calibration degrees on both in-distribution samples and out-of-distribution samples, without negatively impacting classification performance. The approach can be easily applied to different models, which we demonstrate with several SOTA models. It is also demonstrated to be effective on a semantic segmentation task and other binary tasks.

Acknowledgments. This research was in-part supported by the ANU-Optus Bushfire Research Center of Excellence.

References

- [1] Radhakrishna Achanta, Sheila Hemami, Francisco Estrada, and Sabine Susstrunk. Frequency-tuned salient region detection. In *CVPR*, pages 1597–1604. IEEE, 2009. [2](#)
- [2] Görkem Algan and Ilkay Ulusoy. Image classification with deep learning in the presence of noisy labels: A survey. *Knowledge-Based Systems*, 215:106771, 2021. [2](#)
- [3] Ali Borji, Ming-Ming Cheng, Qibin Hou, Huaizu Jiang, and Jia Li. Salient object detection: A survey. *Computational visual media*, 5(2):117–150, 2019. [2](#)
- [4] GLENN W BRIER. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950. [2](#), [5](#), [6](#), [7](#), [20](#), [21](#), [22](#)
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4):834–848, 2017. [29](#)
- [6] Shuhan Chen, Xiuli Tan, Ben Wang, and Xuelong Hu. Reverse attention for salient object detection. In *ECCV*, pages 234–250, 2018. [2](#)
- [7] Zuyao Chen, Qianqian Xu, Runmin Cong, and Qingming Huang. Global context-aware progressive aggregation network for salient object detection. In *AAAI*, volume 34, pages 10599–10606, 2020. [5](#), [6](#), [20](#), [21](#), [22](#)
- [8] Ming-Ming Cheng, Niloy J Mitra, Xiaolei Huang, Philip HS Torr, and Shi-Min Hu. Global contrast based salient region detection. *IEEE TPAMI*, 37(3):569–582, 2014. [2](#)
- [9] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *CVPR*, pages 3606–3613, 2014. [7](#), [18](#), [19](#), [36](#), [37](#), [38](#), [39](#), [40](#), [41](#), [42](#), [43](#), [44](#)
- [10] Morris H DeGroot and Stephen E Fienberg. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 32(1-2):12–22, 1983. [2](#)
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010. [8](#), [29](#)
- [12] Deng-Ping Fan, Cheng Gong, Yang Cao, Bo Ren, Ming-Ming Cheng, and Ali Borji. Enhanced-alignment measure for binary foreground map evaluation. In *IJCAI*, pages 698–704, 2018. [18](#), [27](#), [28](#), [34](#)
- [13] Deng-Ping Fan, Ge-Peng Ji, Guolei Sun, Ming-Ming Cheng, Jianbing Shen, and Ling Shao. Camouflaged object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2020. [1](#)
- [14] Deng-Ping Fan, Ge-Peng Ji, Guolei Sun, Ming-Ming Cheng, Jianbing Shen, and Ling Shao. Camouflaged object detection. In *CVPR*, pages 2777–2787, 2020. [8](#), [28](#)
- [15] Mengyang Feng, Huchuan Lu, and Errui Ding. Attentive feedback network for boundary-aware salient object detection. In *CVPR*, pages 1623–1632, 2019. [5](#), [6](#), [20](#), [21](#), [22](#)
- [16] Shang-Hua Gao, Yong-Qiang Tan, Ming-Ming Cheng, Chengze Lu, Yunpeng Chen, and Shuicheng Yan. Highly efficient salient object detection with 100k parameters. In *ECCV*, pages 702–721. Springer, 2020. [2](#)
- [17] Arindam Ghosh, Thomas Schaaf, and Matthew R. Gormley. Adafocal: Calibration-aware adaptive focal loss. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *NeurIPS*, 2022. [1](#), [2](#), [5](#), [6](#), [7](#), [20](#), [21](#), [22](#)
- [18] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Int. Conf. Mach. Learn.*, pages 1321–1330. PMLR, 2017. [1](#), [2](#), [5](#), [6](#), [7](#), [17](#), [20](#), [21](#), [22](#)
- [19] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhansu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, pages 991–998. IEEE, 2011. [29](#)
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [5](#)
- [21] Xiaowei Hu, Chi-Wing Fu, Lei Zhu, Tianyu Wang, and Pheng-Ann Heng. Sac-net: Spatial attenuation context for salient object detection. *IEEE TCSVT*, 31(3):1079–1090, 2020. [2](#)
- [22] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE TPAMI*, 20(11):1254–1259, 1998. [2](#)
- [23] Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957. [1](#), [2](#), [3](#), [5](#)
- [24] Wei Ji, Xi Li, Lina Wei, Fei Wu, and Yueting Zhuang. Context-aware graph label propagation network for saliency detection. *IEEE TIP*, 29:8177–8186, 2020. [2](#)
- [25] Huaizu Jiang, Jingdong Wang, Zejian Yuan, Yang Wu, Nanning Zheng, and Shipeng Li. Salient object detection: A discriminative regional feature integration approach. In *CVPR*, pages 2083–2090, 2013. [2](#)
- [26] Zhuolin Jiang and Larry S Davis. Submodular salient region detection. In *CVPR*, pages 2043–2050, 2013. [2](#)
- [27] Archit Karandikar, Nicholas Cain, Dustin Tran, Balaji Lakshminarayanan, Jonathon Shlens, Michael C Mozer, and Becca Roelofs. Soft calibration objectives for neural networks. In *NeurIPS*, volume 34, pages 29768–29779, 2021. [1](#), [2](#)
- [28] Meelis Kull, Miquel Perello Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter Flach. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. *NeurIPS*, 32, 2019. [2](#)
- [29] Ananya Kumar, Percy S Liang, and Tengyu Ma. Verified uncertainty calibration. In *NeurIPS*, volume 32, 2019. [5](#), [6](#), [17](#), [22](#)
- [30] Aviral Kumar, Sunita Sarawagi, and Ujjwal Jain. Trainable calibration measures for neural networks from kernel mean embeddings. In *ICLR*, pages 2805–2814. PMLR, 2018. [1](#), [2](#), [5](#), [6](#), [7](#), [20](#), [21](#), [22](#)
- [31] Trung-Nghia Le, Tam V Nguyen, Zhongliang Nie, Minh-Triet Tran, and Akihiro Sugimoto. Anabran network for camouflaged object segmentation. *Comput. Vis. Image Unders.*, 184:45–56, 2019. [28](#)
- [32] Guanbin Li, Yuan Xie, Liang Lin, and Yizhou Yu. Instance-level salient object segmentation. In *CVPR*, pages 2386–2395, 2017. [5](#), [6](#), [20](#), [21](#), [22](#)
- [33] Guanbin Li and Yizhou Yu. Visual saliency based on multi-scale deep features. In *CVPR*, pages 5455–5463, 2015. [5](#), [6](#), [8](#), [18](#), [20](#), [21](#), [22](#), [27](#), [30](#), [34](#)

- [34] Yin Li, Xiaodi Hou, Christof Koch, James M. Rehg, and Alan L. Yuille. The secrets of salient object segmentation. In *CVPR*, pages 280–287, 2014. 5, 6, 8, 18, 20, 21, 22, 27, 30, 34
- [35] Jiang-Jiang Liu, Qibin Hou, and Ming-Ming Cheng. Dynamic feature integration for simultaneous detection of salient object, edge, and skeleton. *IEEE TIP*, 29:8652–8667, 2020. 2
- [36] Jiang-Jiang Liu, Qibin Hou, Zhi-Ang Liu, and Ming-Ming Cheng. Poolnet+: Exploring the potential of pooling for salient object detection. *IEEE TPAMI*, pages –, 2022. 5, 6, 18, 20, 21, 22
- [37] Nian Liu, Junwei Han, and Ming-Hsuan Yang. Picanet: Learning pixel-wise contextual attention for saliency detection. In *CVPR*, pages 3089–3098, 2018. 2, 5, 6, 18, 20, 21, 22
- [38] Tie Liu, Zejian Yuan, Jian Sun, Jingdong Wang, Nanning Zheng, Xiaouo Tang, and Heung-Yeung Shum. Learning to detect a salient object. *IEEE TPAMI*, 33(2):353–367, 2010. 2
- [39] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *CVPR*, pages 10012–10022, 2021. 5, 8, 34
- [40] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 2
- [41] Michal Lukasik, Srinadh Bhojanapalli, Aditya Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? In *Int. Conf. Mach. Learn.*, pages 6448–6458. PMLR, 2020. 14
- [42] Yunqiu Lv, Jing Zhang, Yuchao Dai, Aixuan Li, Bowen Liu, Nick Barnes, and Deng-Ping Fan. Simultaneously localize, segment and rank the camouflaged objects. In *CVPR*, pages 11591–11601, 2021. 28
- [43] Mingcan Ma, Changqun Xia, and Jia Li. Pyramidal feature shrinking for salient object detection. In *AAAI*, volume 35, pages 2311–2318, 2021. 6, 20, 21, 22
- [44] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001. 5, 6, 8, 18, 20, 21, 22, 27, 30, 34
- [45] Jishnu Mukhoti, Viveka Kulharia, Amartya Sanyal, Stuart Golodetz, Philip Torr, and Puneet Dokania. Calibrating deep neural networks using focal loss. In *NeurIPS*, volume 33, pages 15288–15299, 2020. 1, 2, 5, 6, 7, 20, 21, 22
- [46] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In *NeurIPS*, volume 32, 2019. 1, 2, 5, 6, 7, 8, 14, 20, 21, 22
- [47] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *AAAI*, 2015. 2
- [48] Khanh Nguyen and Brendan O’Connor. Posterior calibration and exploratory analysis for natural language processing models. In *ACL Conf. Emp. Meth. Natu. Langu. Proc.*, pages 1587–1598, Lisbon, Portugal, Sept. 2015. Association for Computational Linguistics. 5, 6, 17
- [49] Alexandru Niculescu-Mizil and Rich Caruana. Predicting good probabilities with supervised learning. In *Int. Conf. Mach. Learn.*, pages 625–632, 2005. 2
- [50] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In *NeurIPS*, volume 32, 2019. 7
- [51] Youwei Pang, Xiaoqi Zhao, Lihe Zhang, and Huchuan Lu. Multi-scale interactive network for salient object detection. In *CVPR*, pages 9413–9422, 2020. 2, 5, 6, 20, 21, 22
- [52] Gabriel Pereyra, George Tucker, Jan Chorowski, Lukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. In *ICLR*, 2017. 2
- [53] Yongri Piao, Wei Ji, Jingjing Li, Miao Zhang, and Huchuan Lu. Depth-induced multi-scale recurrent attention network for saliency detection. In *ICCV*, pages 7254–7263, 2019. 2
- [54] John Platt et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74, 1999. 1, 2
- [55] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. Basnet: Boundary-aware salient object detection. In *CVPR*, pages 7479–7489, 2019. 2, 5, 6, 20, 21, 22
- [56] Rebecca Roelofs, Nicholas Cain, Jonathon Shlens, and Michael C Mozer. Mitigating bias in calibration error estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 4036–4054. PMLR, 2022. 5, 6, 17, 21
- [57] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 17
- [58] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5, 8, 34
- [59] Przemysław Skurowski, Hassan Abdulameer, J Błaszczuk, Tomasz Depta, Adam Kornacki, and P Kozieł. Animal camouflage analysis: Chameleon database. *Unpublished manuscript*, 2(6):7, 2018. 28
- [60] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. 1, 2, 3
- [61] Lv Tang, Bo Li, Yijie Zhong, Shouhong Ding, and Mofei Song. Disentangled high quality salient object detection. In *ICCV*, pages 3580–3590, 2021. 2
- [62] Sunil Thulasidasan, Gopinath Chennupati, Jeff A Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In *NeurIPS*, volume 32, 2019. 2, 5, 6, 7, 20, 21, 22

- [63] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *CVPR*, pages 136–145, 2017. [5](#), [6](#), [8](#), [18](#), [20](#), [21](#), [22](#), [27](#), [30](#), [34](#)
- [64] Linzhao Wang, Lijun Wang, Huchuan Lu, Pingping Zhang, and Xiang Ruan. Saliency detection with recurrent fully convolutional networks. In *ECCV*, pages 825–841. Springer, 2016. [2](#)
- [65] Tiantian Wang, Ali Borji, Lihe Zhang, Pingping Zhang, and Huchuan Lu. A stagewise refinement model for detecting salient objects in images. In *ICCV*, Oct 2017. [5](#), [6](#), [20](#), [21](#), [22](#)
- [66] Tiantian Wang, Lihe Zhang, Huchuan Lu, Chong Sun, and Jinqing Qi. Kernelized subspace ranking for saliency detection. In *ECCV*, pages 450–466. Springer, 2016. [2](#)
- [67] Tiantian Wang, Lihe Zhang, Shuo Wang, Huchuan Lu, Gang Yang, Xiang Ruan, and Ali Borji. Detect globally, refine locally: A novel approach to saliency detection. In *CVPR*, pages 3127–3135, 2018. [5](#), [6](#), [20](#), [21](#), [22](#)
- [68] 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 *CVPR*, pages 5968–5977, 2019. [2](#)
- [69] Jun Wei, Shuhui Wang, and Qingming Huang. F³net: fusion, feedback and focus for salient object detection. In *AAAI*, volume 34, pages 12321–12328, 2020. [5](#), [6](#), [20](#), [21](#), [22](#)
- [70] Jun Wei, Shuhui Wang, Zhe Wu, Chi Su, Qingming Huang, and Qi Tian. Label decoupling framework for salient object detection. In *CVPR*, pages 13025–13034, 2020. [2](#)
- [71] Yichen Wei, Fang Wen, Wangjiang Zhu, and Jian Sun. Geodesic saliency using background priors. In *ECCV*, pages 29–42. Springer, 2012. [2](#)
- [72] Yu-Huan Wu, Yun Liu, Le Zhang, Ming-Ming Cheng, and Bo Ren. Edn: Salient object detection via extremely-downsampled network. *IEEE TIP*, 2022. [5](#), [6](#), [8](#), [18](#), [20](#), [21](#), [22](#), [27](#)
- [73] Zhe Wu, Li Su, and Qingming Huang. Cascaded partial decoder for fast and accurate salient object detection. In *CVPR*, pages 3907–3916, 2019. [5](#), [6](#), [20](#), [21](#), [22](#)
- [74] Zhe Wu, Li Su, and Qingming Huang. Stacked cross refinement network for edge-aware salient object detection. In *ICCV*, pages 7264–7273, 2019. [2](#)
- [75] Lingxi Xie, Jingdong Wang, Zhen Wei, Meng Wang, and Qi Tian. Disturblabel: Regularizing cnn on the loss layer. In *CVPR*, June 2016. [1](#), [2](#), [3](#)
- [76] Binwei Xu, Haoran Liang, Ronghua Liang, and Peng Chen. Locate globally, segment locally: A progressive architecture with knowledge review network for salient object detection. In *AAAI*, volume 35, pages 3004–3012, 2021. [2](#)
- [77] Yi Xu, Yuanhong Xu, Qi Qian, Hao Li, and Rong Jin. Towards understanding label smoothing. *arXiv preprint arXiv:2006.11653*, 2020. [14](#)
- [78] Qiong Yan, Li Xu, Jianping Shi, and Jiaya Jia. Hierarchical saliency detection. In *CVPR*, pages 1155–1162, 2013. [5](#), [6](#), [8](#), [18](#), [20](#), [21](#), [22](#), [27](#), [30](#), [34](#)
- [79] Siyuan Yan, Jing Zhang, and Nick Barnes. Transmission-guided bayesian generative model for smoke segmentation. In *AAAI*, pages 3009–3017, 2022. [1](#), [8](#), [28](#)
- [80] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via graph-based manifold ranking. In *CVPR*, pages 3166–3173. IEEE, 2013. [5](#), [6](#), [8](#), [18](#), [20](#), [21](#), [22](#), [27](#), [30](#), [34](#)
- [81] Yaodong Yu, Stephen Bates, Yi Ma, and Michael Jordan. Robust calibration with multi-domain temperature scaling. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *NeurIPS*, 2022. [2](#)
- [82] Bianca Zadrozny and Charles Elkan. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *Int. Conf. Mach. Learn.*, volume 1, pages 609–616. Citeseer, 2001. [2](#)
- [83] Bianca Zadrozny and Charles Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *ACM SIGKDD Int. Conf. Knowledge Disc. Data Min.*, pages 694–699, 2002. [2](#)
- [84] Chang-Bin Zhang, Peng-Tao Jiang, Qibin Hou, Yunchao Wei, Qi Han, Zhen Li, and Ming-Ming Cheng. Delving deep into label smoothing. *IEEE TIP*, 30:5984–5996, 2021. [14](#)
- [85] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. [1](#), [2](#)
- [86] Jing Zhang, Deng-Ping Fan, Yuchao Dai, Saeed Anwar, Fatemeh Saleh, Sadegh Aliakbarian, and Nick Barnes. Uncertainty inspired rgb-d saliency detection. *IEEE TPAMI*, 2021. [1](#)
- [87] Jize Zhang, Bhavya Kailkhura, and T Yong-Jin Han. Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning. In *Int. Conf. Mach. Learn.* PMLR, 2020. [2](#)
- [88] Jing Zhang, Jianwen Xie, and Nick Barnes. Learning noise-aware encoder-decoder from noisy labels by alternating back-propagation for saliency detection. In *ECCV*, pages 349–366. Springer, 2020. [2](#)
- [89] Jing Zhang, Jianwen Xie, Nick Barnes, and Ping Li. Learning generative vision transformer with energy-based latent space for saliency prediction. In *NeurIPS*, 2021. [5](#), [6](#), [8](#), [20](#), [21](#), [22](#), [27](#)
- [90] Jing Zhang, Xin Yu, Aixuan Li, Peipei Song, Bowen Liu, and Yuchao Dai. Weakly-supervised salient object detection via scribble annotations. In *CVPR*, pages 12546–12555, 2020. [2](#)
- [91] Lu Zhang, Ju Dai, Huchuan Lu, You He, and Gang Wang. A bi-directional message passing model for salient object detection. In *CVPR*, pages 1741–1750, 2018. [2](#), [5](#), [6](#), [20](#), [21](#), [22](#)
- [92] Pingping Zhang, Dong Wang, Huchuan Lu, Hongyu Wang, and Xiang Ruan. Amulet: Aggregating multi-level convolutional features for salient object detection. In *ICCV*, pages 202–211, 2017. [5](#), [6](#), [18](#), [20](#), [21](#), [22](#)
- [93] Xiaoning Zhang, Tiantian Wang, Jinqing Qi, Huchuan Lu, and Gang Wang. Progressive attention guided recurrent network for salient object detection. In *CVPR*, pages 714–722, 2018. [5](#), [6](#), [20](#), [21](#), [22](#)
- [94] Jia-Xing Zhao, Jiang-Jiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. Egnet: Edge guidance network for salient object detection. In *ICCV*, pages 8779–8788, 2019. [5](#), [6](#), [20](#), [21](#), [22](#)

- [95] Ting Zhao and Xiangqian Wu. Pyramid feature attention network for saliency detection. In *CVPR*, pages 3085–3094, 2019. 2
- [96] Xiaoqi Zhao, Youwei Pang, Lihe Zhang, Huchuan Lu, and Lei Zhang. Suppress and balance: A simple gated network for salient object detection. In *ECCV*, pages 35–51. Springer, 2020. 2
- [97] Mingchen Zhuge, Deng-Ping Fan, Nian Liu, Dingwen Zhang, Dong Xu, and Ling Shao. Salient object detection via integrity learning. *IEEE TPAMI*, 2022. 5, 6, 8, 20, 21, 22, 27