

Pixel-wise Energy-biased Abstention Learning for Anomaly Segmentation on Complex Urban Driving Scenes

Yu Tian^{1*}, Yuyuan Liu^{1*}, Guansong Pang², Fengbei Liu¹, Yuanhong Chen¹,
and Gustavo Carneiro¹

¹ Australian Institute for Machine Learning, University of Adelaide

² Singapore Management University

Abstract. State-of-the-art (SOTA) anomaly segmentation approaches on complex urban driving scenes explore pixel-wise classification uncertainty learned from outlier exposure, or external reconstruction models. However, previous uncertainty approaches that directly associate high uncertainty to anomaly may sometimes lead to incorrect anomaly predictions, and external reconstruction models tend to be too inefficient for real-time self-driving embedded systems. In this paper, we propose a new anomaly segmentation method, named pixel-wise energy-biased abstention learning (PEBAL), that explores pixel-wise abstention learning (AL) with a model that learns an adaptive pixel-level anomaly class, and an energy-based model (EBM) that learns inlier pixel distribution. More specifically, PEBAL is based on a non-trivial joint training of EBM and AL, where EBM is trained to output high-energy for anomaly pixels (from outlier exposure) and AL is trained such that these high-energy pixels receive adaptive low penalty for being included to the anomaly class. We extensively evaluate PEBAL against the SOTA and show that it achieves the best performance across four benchmarks. Code is available at <https://github.com/tianyu0207/PEBAL>.

1 Introduction

Recent advances in semantic segmentation have shown tremendous improvements on complex urban driving scenes [23]. Despite the accurate predictions on the inlier classes, the model fails to properly recognise anomalous objects that deviate from the training inlier distribution (col. 2 of Fig. 1). Addressing such failure cases is crucial to road safety for autonomous driving vehicles. For example, anomalies can be represented by unexpected objects in the middle of the road, such as a large rock or an unexpected animal that can be incorrectly predicted as a part of the road class, leading to potentially fatal traffic collisions.

Current methods [2, 4, 6, 11, 20, 30, 34, 40] to detect and segment anomalous objects in complex urban driving scenes tend to depend on classification uncertainty or image reconstruction. The association of high classification uncertainty with anomaly is intuitive, but it has a few caveats. For instance, classification uncertainty happens when samples are close to classification decision boundaries, but there is no guarantee that all anomalies will be close to classification boundaries.

*First two authors contributed equally to this work. GP is the corresponding author.




Fig. 1: **Anomaly segmentation overview.** From the **input image** (anomaly highlighted with a yellow box), the **initial prediction** shows the original segmentation results with anomalies classified as a one of the pre-defined inlier classes. **Anomaly predictions** by the previous SOTA **Meta-OoD** [6] and **our** method show an anomaly map with high scores (in yellow and red) for anomalous pixels, where our approach shows less false positive and false negative detections. Consequently, our method can detect small and distant anomalies (row 2) and blurry/unclear anomalies (rows 1, 3, 4) more accurately than Meta-OoD [6]. In our **final prediction**, anomalous pixels are coloured as cyan. **Some anomalies are small and blurred (e.g., row 2), so please zoom in the PDF for better visualisation.**

Furthermore, samples close to classification boundaries may not be anomalies at all, but just hard inlier samples. Hence, these uncertainty based methods may detect a large number of false positive and false negative anomalies. For example, Fig. 1 shows that the previous SOTA Meta-OoD [6] misses important anomalous pixels (all rows), while misclassifying anomalies (e.g., vegetation in rows 1, 2, 3), even with the use of the outlier exposure (OE) strategy [19]. In fact, the OE strategy maximises the uncertainty for proxy anomalies, which can cause the model to be more uncertain for all inlier classes and detect false positive anomalies (e.g., Meta-OoD mis-classifies trees or bush with high anomaly scores – Fig. 1 col 4). Reconstruction methods [11, 40] add an extra network to reconstruct the input images from the estimated segmentation, where differences are assumed to be anomalous. Not only does this approach depend on accurate segmentation results for precise reconstruction, but they also require an extra reconstruction network that is hard to train and inefficient to run in real-time self-driving embedded systems. Moreover, reconstruction methods that rely on a discrepancy module require re-training whenever the inlier segmentation model changes due to input distribution shift [11], limiting their applicability in real-world systems. Furthermore, previous approaches [2, 6, 11, 15, 20, 30] ignore a couple of important constraints for anomaly segmentation, namely smoothness (e.g., Meta-OoD fails to classify neighbouring anomaly pixels in Fig. 1, rows 1, 4) and sparsity (e.g.,

Meta-OoD incorrectly detects a large number of anomalous pixels—see yellow and red regions in Fig. 1, rows 1, 2, 3). Another common issue shared by previous methods [2,6,30] is that they usually rely on the re-training of the entire network for OE, which is inefficient and can also bias the classification towards outliers.

In this paper, we propose a new anomaly segmentation method, the pixel-wise energy-biased abstention learning (PEBAL), that directly learns a pixel-level anomaly class, in addition to the pre-defined inlier classes, to reject/abstain anomalous pixels that are dissimilar to any of the inlier classes. It is achieved by a joint optimisation of a novel pixel-wise anomaly abstention learning (PAL) and an energy based model (EBM) [14, 24, 31]. Particularly, abstention learning (AL) [32] was originally developed to learn an image-level anomaly class, which is significantly challenged by the pixel-wise anomaly segmentation task that requires pixel-level anomaly class learning. This is because the original AL model treats all pixel inputs equally with a single pre-defined fixed penalty factor to regularise the classification of anomalous pixels, while adaptive penalties are typically required for different pixels in a complex driving scene, e.g., pixels in small (distant) objects vs. large (near) objects, or centred pixels vs fringe pixels of objects. PEBAL is designed to address this issue by learning adaptive pixel-wise energy-based penalties, which automatically decreases the penalty for pixels that are likely to be anomalies. Hence, our model does not explore previously proposed uncertainty measures (e.g., entropy or softmax criteria) or image reconstruction, and instead, for the first time, explicitly learns a new pixel-wise anomaly class. The learned penalty factors are jointly optimised with EBM, resulting in a mutually beneficial optimisation of anomaly and inlier segmentation. Additionally, we impose smoothness and sparsity constraints to the learning of the anomaly segmentation by PEBAL, incorporating local and global dependencies into the pixel-wise penalty estimation and anomaly score learning. Finally, the training of PEBAL is efficient given that we only need to fine-tune the last block of the segmentation model to achieve accurate inference. To summarise, our contributions are the following:

- We propose the pixel-wise energy-biased abstention learning (PEBAL) that jointly optimises a novel pixel-wise anomaly abstention learning (PAL) and energy-based models (EBM) to learn adaptive pixel-level anomalies. PEBAL mutually reinforces PAL and EBM in detecting anomalies, enabling accurate segmentation of anomalous pixels without compromising the segmentation of inlier pixels (cols. 4,5 of Fig. 1).
- We introduce a new pixel-wise energy-biased penalty estimation, which can learn adaptive energy-based penalties to highly varying pixels in a complex driving scene, allowing a robust detection of small/distant and blurry anomalous objects (Fig. 1 row 2).
- We further refine our PEBAL training, using a novel smoothness and sparsity regularisation on anomaly scores to consider the local and global dependencies of the pixels, enabling the reduction of false positive/negative anomaly predictions.

We validate our approach on Fishyscapes leaderboard [4], and achieve SOTA classification accuracy on all relevant benchmarks. We also achieve the best classification results on LostandFound [36] and Road Anomaly [30] test sets, significantly surpassing other competing methods.

2 Related work

Uncertainty-based Anomaly Segmentation. Early uncertainty-based methods [18, 25, 28] focused on the estimation of image-level anomalies, but they tended to misclassify object boundaries as anomalies [20]. Jung et al. [20] mitigate this issue by iteratively replacing false anomalous boundary pixels with neighbouring non-boundary pixels that have low anomaly score. In [21, 22, 34], the boundary issue was tackled with a pixel-wise uncertainty estimated with MC dropout, but they showed a low pixel-wise anomaly detection accuracy [30]. Without fine-tuning using a proxy outlier dataset, uncertainty estimation may not be accurate enough to detect anomalies and can predict high uncertainty for challenging inliers or low uncertainty for outliers due to overconfident misclassification.

Reconstruction-based Anomaly Segmentation. Anomalies can also be segmented from the errors between the input image and its reconstruction obtained from its predicted segmentation map [1, 8, 10, 11, 16, 30, 39, 40]. Those approaches are challenged by the dependence on an accurate segmentation prediction, by the complexity of reconstruction models that usually require long training and inference processes, and also by the low quality of the reconstructed images.

Anomaly Segmentation via Outlier Exposure. Hendrycks et al. [19] propose the outlier exposure (OE) strategy that uses an auxiliary dataset of outliers that do not overlap with the real outliers/anomalies to improve the anomaly detection performance. This OE strategy uses outliers from ImageNet [2, 3, 38], void class of Cityscape [11] or COCO [6], where the expectation is that the model can generalise to unseen outliers. Maximising uncertainty for outliers using the OE strategy can lead to a deterioration of the segmentation of inliers [3, 38]. Another major drawback of OE methods is that they are trained using outlier images or objects without considering the fact that outliers are rare events that appear around inliers. Hence, the training contains a disproportionately high amount of outliers [6] that can bias the segmentation toward the anomaly class. We address this issue by respecting the anomaly detection assumption, where anomalous objects are rare, contribute to a small proportion of the training set, and appear around inliers.

Abstention Learning. The abstention learning mechanism [12] adds a “reserve” (i.e., anomaly) class that is predicted when the classification predictions for all inlier classes are not high enough. This method shows good performance in learning holistic image-level anomaly class with a single pre-defined penalty factor for the whole training set, but it fails to learn fine-grained pixel-level anomaly class as an adaptive pixel-wise penalty is required for highly varying pixel-level anomalies (see Table 5). We address this issue by learning a novel pixel-wise

energy-biased penalty estimator that is jointly trained with fine-grained abstention learning. It is worth noting that differently from uncertainty-based methods [4, 6, 17, 20] that assume anomaly even when the model is uncertain but confident, abstention learning requires all classes to have low confidence to predict the anomaly class.

Energy-based Models. EBM is trained such that inlier training samples have low energy, whereas non-training outlier samples (i.e., anomalies) are expected to have high energy [24]. This energy value can then be used to compute the probability of a sample to belong to the inlier distribution. Recently, EBMs are being implemented with deep learning models [14, 31, 35], and to learn them, it is necessary to compute the partition function, which is generally estimated with Markov Chain Monte Carlo (MCMC) [14], but this estimation cannot generate accurate high-resolution images. Hence, we follow the simpler idea of estimating the energy score with the *logsumexp* operator [14, 31], where we minimise the energy of inliers and use an OE strategy [19] to maximise the energy of outliers. Hence, we do not need to compute the partition function.

3 Method

We present our PEBAL in this section (see Fig. 2), where we first describe the dataset, then introduce abstention learning and EBM. Next, we present the loss function to train the model, followed by the training and inference procedures.

3.1 Training Set

We assume to have a set of inlier training images and annotations $\mathcal{D}^{in} = \{(\mathbf{x}_i, \mathbf{y}_i^{in})\}_{i=1}^{|\mathcal{D}^{in}|}$, where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{H \times W \times C}$ denotes an image with C colour channels, and $\mathbf{y}^{in} \in \mathcal{Y}^{in} \subset \{0, 1\}^{H \times W \times Y}$ denotes the inlier pixel level labels that can belong to Y classes. We also have a set of outlier images and annotations $\mathcal{D}^{out} = \{(\mathbf{x}_i, \mathbf{y}_i^{out})\}_{i=1}^{|\mathcal{D}^{out}|}$, where $\mathbf{y}^{out} \in \mathcal{Y}^{out} \subset \{0, 1\}^{H \times W \times (Y+1)}$ denotes the outlier pixel-level labels, with the class $Y + 1$ reserved for pixels belonging to the anomaly class. Note that similarly to previous papers [6], the types of anomalies in training set \mathcal{D}^{out} do not overlap with the anomalies to be found in the testing set.

3.2 Pixel-wise Energy-biased Abstention Learning (PEBAL)

The PEBAL model is denoted by

$$p_{\theta}(y|\mathbf{x})_{\omega} = \frac{\exp(f_{\theta}(y; \mathbf{x})_{\omega})}{\sum_{y' \in \{1, \dots, Y+1\}} \exp(f_{\theta}(y'; \mathbf{x})_{\omega})}, \quad (1)$$

where θ is the model parameter, ω indexes a pixel in the image lattice Ω , $p_{\theta}(y|\mathbf{x})_{\omega}$ represents the probability of labelling pixel ω with $y \in \{1, \dots, Y + 1\}$, and $f_{\theta}(y; \mathbf{x})_{\omega}$ is the logit for class y at pixel ω .

To train the model in (1), we formulate a cost function that jointly trains PAL and EBM to classify anomalous pixels. An important training hyper-parameter




Fig. 2: **PEBAL**. The pixel-wise anomaly abstention (PAL) loss ℓ_{pal} learns to abstain the prediction of outlier pixels from \mathbf{x}^{out} containing OE objects (i.e., cyan coloured masks) and calibrate the logit of inlier classes (i.e., reduction of the inlier logits) from both inlier image \mathbf{x}^{in} and outlier image \mathbf{x}^{out} . The EBM loss ℓ_{ebm} pushes the free energy E_θ to low values for inlier pixels and pulls that to high values for outlier pixels, where a regularisation loss ℓ_{reg} enforces the smoothness and sparsity constraints on the energy maps. Such EBM learning reduces the logit of inlier classes to share similar values at the same time, facilitating the ℓ_{pal} learning. Then, the pixel-wise penalty a_ω associated with the abstention class at position ω is estimated to bias the penalty to be low for outlier pixels and high for inlier pixels, which in turn encourages high free energy for anomalies and enforces ℓ_{pal} to abstain the anomalous pixels.

for PAL is the penalty to abstain from the classification into one of the inlier classes in $\{1, \dots, Y\}$ —this penalty is generally tuned to a single value for all training samples through model selection (e.g., cross validation) [32]. Instead of treating this as a tunable hyper-parameter, we propose the use of EBM (defined below in (4)) to automatically estimate this penalty during the training process for each pixel within each training image. More specifically, the cost function to train the PEBAL model in (1) is:

$$\begin{aligned} \ell(\mathcal{D}^{in}, \mathcal{D}^{out}, \theta) = & \\ & \sum_{(\mathbf{x}, \mathbf{y}^{in}) \in \mathcal{D}^{in}} (\ell_{pal}(\theta, \mathbf{y}^{in}, \mathbf{x}, E_\theta(\mathbf{x})) + \lambda \ell_{ebm}^{in}(E_\theta(\mathbf{x})) + \ell_{reg}(E_\theta(\mathbf{x}))) + \\ & \sum_{(\mathbf{x}, \mathbf{y}^{out}) \in \mathcal{D}^{out}} (\ell_{pal}(\theta, \mathbf{y}^{out}, \mathbf{x}, E_\theta(\mathbf{x})) + \lambda \ell_{ebm}^{out}(E_\theta(\mathbf{x})) + \ell_{reg}(E_\theta(\mathbf{x}))). \end{aligned} \quad (2)$$

where $\ell_{pal}(\cdot)$ denotes the PAL loss defined as

$$\ell_{pal}(\theta, \mathbf{y}, \mathbf{x}, E_\theta(\mathbf{x})) = - \sum_{\omega \in \Omega} \log \left(f_\theta(y_\omega; \mathbf{x})_\omega + \frac{f_\theta(Y+1; \mathbf{x})_\omega}{a_\omega} \right), \quad (3)$$

with $y_\omega \in \{1, \dots, Y\}$ for \mathbf{y}^{in} , $y_\omega \in \{1, \dots, Y+1\}$ for \mathbf{y}^{out} , and a_ω denotes the pixel-wise penalty associated with abstaining from the classification of the inlier classes. The minimisation of the loss in (3) will abstain from classifying outlier pixels into one of the inlier classes, where a pixel is estimated to be an outlier with a_ω . Before formulating a_ω , let us define the inlier free energy at pixel ω , which is denoted by $E_\theta(\mathbf{x})_\omega$ and computed with the *logsumexp* operator as

follows [14, 24, 31]:

$$E_{\theta}(\mathbf{x})_{\omega} = -\log \sum_{y \in \{1, \dots, Y\}} \exp(f_{\theta}(y; \mathbf{x})_{\omega}). \quad (4)$$

The pixel-wise penalty associated with abstaining from the classification of the inlier classes is defined by

$$a_{\omega} = (-E_{\theta}(\mathbf{x})_{\omega})^2, \quad (5)$$

which means that the larger the a_{ω} (i.e., low inlier free energy, so the sample is an inlier), the higher the loss to abstain from classifying into one of the Y classes, and low value of a_{ω} (i.e., high free inlier energy, which means an outlier sample) implies a lower loss to abstain from classifying one of the Y classes. Also in (2), $\ell_{ebm}^{in}(\cdot)$ (weighted by hyper-parameter λ) represents the EBM loss that pushes the inlier free energy in (4) for samples in \mathcal{D}^{in} to low values, with

$$\ell_{ebm}^{in}(E_{\theta}(\mathbf{x})) = \sum_{\omega \in \Omega} (\max(0, E_{\theta}(\mathbf{x})_{\omega} - m_{in}))^2, \quad (6)$$

representing the loss of having inlier samples with free energy larger than threshold m_{in} , and

$$\ell_{ebm}^{out}(E_{\theta}(\mathbf{x})) = \sum_{\omega \in \Omega} (\max(0, m_{out} - E_{\theta}(\mathbf{x})_{\omega}))^2, \quad (7)$$

denoting the loss of having outlier samples with inlier free energy smaller than threshold m_{out} , where the margin losses in (6) and (7) effectively create an energy gap between normal and abnormal pixels. The last term to define in (2) is the inlier free energy regularisation loss to enforce that anomalous pixels are sparse and pixel anomaly classification is smooth (i.e., anomalous pixels tend to have anomalous neighbouring pixels), which is defined as

$$\ell_{reg}(E_{\theta}(\mathbf{x})) = \sum_{\omega \in \Omega} \beta_1 |E_{\theta}(\mathbf{x})_{\omega} - E_{\theta}(\mathbf{x})_{\mathcal{N}(\omega)}| + \beta_2 |E_{\theta}(\mathbf{x})_{\omega}|, \quad (8)$$

where β_1 and β_2 are hyper-parameters that weight the contributions of the smoothness and sparsity and sparsity regularisations, and $\mathcal{N}(\omega)$ denotes neighbouring pixels in horizontal and vertical directions.

3.3 Training and Inference

Training. An important point of the training process is how to setup the inlier and outlier datasets \mathcal{D}^{in} and \mathcal{D}^{out} . A recently published paper [6] carefully selects images to be included in \mathcal{D}^{out} by making sure that the segmentation labels presented in those images do not overlaps with the inlier labels. In particular for [6], \mathcal{D}^{in} has images and annotations from Cityscape and \mathcal{D}^{out} has images and annotations from COCO [29]. We argue that there are two issues with this strategy to form \mathcal{D}^{out} , which are: 1) the selected COCO images generally only

contain anomalous pixel labels, leading to unstable training of the outlier losses (i.e., second summation in (2)) given the exclusive presence of the anomaly class (in effect, this becomes a one-class segmentation problem); 2) re-training the model with images containing only anomalous pixels removes the semantic context of inlier pixels when training for the outlier losses, which can deteriorate the segmentation accuracy of the inlier labels.

To mitigate these issues, we form \mathcal{D}^{out} using a novel extension based on Cut-Mix and CutPaste [27, 41], which we refer to as AnomalyMix. AnomalyMix cuts the anomalous objects from an outlier dataset (e.g., COCO) using its labelled masks and paste them into the images of the inlier dataset (e.g., CitySpace), where we label the pixels of the anomalous object with the class $Y + 1$ – these images are then inserted into \mathcal{D}^{out} . AnomalyMix addresses the two issues above because the outlier images now contain a combination of inlier and outlier pixels, allowing a balanced learning and keeping the visual context of inlier labels when training for the outlier losses. Furthermore, AnomalyMix can form a potentially infinite number of training images for \mathcal{D}^{out} given the range of transformations to be applied to the cut objects and the locations of the inlier images that the objects can be pasted. Previous papers [11, 20] argue that re-training the whole segmentation model can jeopardise the segmentation accuracy for the inlier classes. Furthermore, such re-training requires a long training time, leading to inefficient optimisation. In this work, we propose to **fine-tune only the final classification block** using the loss in (2), instead of re-training the whole segmentation model. Besides being efficient, this fast fine-tuning keeps the segmentation accuracy of the model in the original dataset used for pre-training the model. Furthermore, an interesting side-effect of our training is that the cost function in (2) will calibrate the segmentation prediction for the inlier classes. This happens because the terms $\ell_{pal}(\cdot)$, $\ell_{ebm}^{in}(\cdot)$ and $\ell_{ebm}^{out}(\cdot)$ jointly constrain the maximisation of logits and naturally calibrate classification confidence (See supplementary material).

Inference. During inference, pixel-wise anomaly detection is performed by computing the inlier free energy score $\mathbb{E}_{\theta}(\mathbf{x})_{\omega}$ from (4) for each pixel position ω given a test image \mathbf{x} and inlier segmentation is obtained from the inlier classes from the PEBAL model in (1). Following [20], we also apply a Gaussian smoothing kernel to produce the final energy map.

4 Experiment

4.1 Datasets

LostAndFound [36] is one of the first publicly available urban driving scene anomaly detection datasets containing real-world anomalous objects. The dataset has an official testing set containing 1,203 images with small obstacles in front of the cars, collecting from 13 different street scenes, featuring 37 different types of anomalous objects with various sizes and material.

Fishyscapes [4] is a high-resolution dataset for anomaly estimation in semantic segmentation for urban driving scenes. The benchmark has an online testing set

that is entirely unknown to the methods. The dataset is composed by two data sources: Fishyscapes LostAndFound that contains a set of real road anomalous objects [36] and a blending-based Fishyscapes Static dataset. The Fishyscapes LostAndFound validation set consists of 100 images from the aforementioned LostAndFound dataset with refined labels and the Fishyscapes Static validation set contains 30 images with the blended anomalous objects from Pascal VOC [13]. For all datasets, we select the checkpoints based on the results on the public validation sets, but submitted our code and checkpoints to the benchmark website to be evaluated on their hidden test sets.

Road Anomaly [30] contains real-world road anomalies in front of the vehicles. The dataset has 60 images from the Internet, containing unexpected animals rocks, cones and obstacles. Unlike the LostAndFound and Fishyscapes, this dataset contains abnormal objects with various scales and sizes, making it even more challenging.

4.2 Implementation Details

Following [5,6], we use DeepLabv3+ [7] with WideResnet38 trained by Nvidia [42] and ResNet101 from [20] as the backbone of our segmentation models. The training details of those models can be found in their original papers or our supplementary material. The models are trained on Cityscapes [9] training set. For our PEBAL fine-tuning, we empirically set the m_{in} and m_{out} in Eq. 6 and Eq. 7 as -12 and -6, respectively. The weights β_1 and β_2 in Eq. 8 are set to $5e - 4$ and $3e - 6$ [37], λ in Eq. 2 to 0.1, and the weight of ℓ_{ebm} to 0.1, respectively. Note that those hyper-parameters are selected at the first training epoch to normalise loss values to a similar scale. We also show our model can obtain consistently SOTA results regardless of the selection of hyper-parameters in the supplementary material. Our training consists of fine-tuning the final classification block of the model for 20 epochs. We use the same resolution of random crop as in [42], and use Adam with a learning rate of $1e^{-5}$. The batch size is set to 16. Following [6], for our AnomalyMix augmentation, we randomly sample 297 images as training data from the remaining COCO images that do not contain objects in Cityscapes or our anomaly validation/testing sets and randomly apply AnomalyMix to mix them into the Cityscape training images, following Chan et al. [6].

4.3 Evaluation Measures

Following [4,6,11,20], we compute the the area under receiver operating characteristics (AUROC), average precision (AP), and the false positive rate at a true positive rate of 95% (FPR95) to validate our approach. For Fishyscapes public leaderboard, we use AP and FPR95 to compare with other methods, same as their website.

4.4 Comparison on Anomaly Segmentation Benchmarks

Comparison on LostAndFound. Table 1 shows the result on the testing set of LostAndFound. Notably, our approach surpasses the previous baseline approaches (i.e., MSP [17], Mahalanobis [26], Max Logit [18] and Entropy [18])

Table 1: Anomaly segmentation results on **LostAndFound** testing set, with **WideResnet38** backbone. All methods use the same segmentation models. * indicate that the model requires additional learnable parameters. † indicates that the results are obtained from the official code with our WideResnet38 backbone.

Methods	AUC ↑	AP ↑	FPR ₉₅ ↓
MSP [17]	85.49	38.20	18.56
Mahalanobis [26]	79.53	42.56	24.51
Max Logit [18]	94.52	65.45	15.56
Entropy [18]	86.52	50.66	16.95
Energy [31]	94.45	66.37	15.69
Meta-OoD [6]	97.95	71.23	5.95
†SML [20]	88.05	25.89	44.48
†SynBoost* [11]	98.38	70.43	4.89
Deep Gambler [32]	98.67	72.73	3.81
Ours	99.76	78.29	0.81

by 10% to 40% AP, and 13% to 22% FPR95, respectively. When compared with previous SOTA approaches such as SynBoost [11], SML [20] and Meta-OoD [6], we improve the AP performance by a large margin (15% to 40%), and decrease the FPR95 by about 5% to 70%. This illustrates the robustness and effectiveness on detecting small and distant anomalous objects given that the dataset contains mostly real-world small objects. Our PEBAL also improves the EBM baseline [31] and the AL baseline based on Deep Gambler [32]. This demonstrates that a simple adaptation of AL and EBM is not enough to enable accurate pixel-wise anomaly detection. Previous SOTA SML [20] aims to balance the inlier class-wise discrepancy on prediction scores, which is disadvantageous for measuring performance on LostAndFound test set since there may be no classes in the evaluation other than the road class (i.e., most of the inlier classes within LF test set is road class), thus leading to significant performance variations between LostAndFound and Fishyscapes. It is worth noting that our approach achieves 1.03% FPR95, significantly reducing the false positive pixels, improving the chances of applying it to real-world applications.

Comparison on Fishyscapes Leaderboard. Table 2 shows the leaderboard results on the test set of Fishyscapes LostAndFound and Fishyscapes Static. Following [20], we compared the methods based on whether they require re-training of the entire segmentation network, adding the extra network, or utilising the OoD data. We achieve the SOTA performance by a large margin on Fishyscapes leaderboard when compared with the previous methods except [2] (Static) that rely on an inefficient re-training segmentation model, extra learnable parameters, and extra OoD training data. Without re-training the entire network or adding extra learnable parameters, our approach can work efficiently to surpass previous SOTA competing approaches that fall into the same category by about 13% to 42% on LostAndFound and 40% to 50% AP on Static. Such significant improvements indicate the generalisation ability of our proposed PEPAL on detecting a wide variety of unseen abnormalities (i.e., of different size, type, scene, and distance) substantially reducing false negative and positive pixels. Moreover,

Table 2: Comparison with previous approaches on **Fishyscapes Leaderboard**. We achieve a new state-of-the-art performance among the approaches that require extra OoD data, and without re-training the segmentation networks and extra networks on Fishyscapes Leaderboard.

Models	re-training	Extra Network	OoD Data	FS LostAndFound		FS Static	
				AP \uparrow	FPR95 \downarrow	AP \uparrow	FPR95 \downarrow
Discriminative Outlier Detection Head [2]	✓	✓	✓	31.31	19.02	96.76	0.29
MSP [17]	✗	✗	✗	1.77	44.85	12.88	39.83
Entropy [18]	✗	✗	✗	2.93	44.83	15.41	39.75
SML [20]	✗	✗	✗	31.05	21.52	53.11	19.64
kNN Embedding - density [4]	✗	✗	✗	3.55	30.02	44.03	20.25
Bayesian Deeplab [34]	✓	✗	✗	9.81	38.46	48.70	15.05
Density - Single-layer NLL [4]	✗	✓	✗	3.01	32.9	40.86	21.29
Density - Minimum NLL [4]	✗	✓	✗	4.25	47.15	62.14	17.43
Image Resynthesis [30]	✗	✓	✗	5.70	48.05	29.6	27.13
OoD Training - Void Class	✓	✗	✓	10.29	22.11	45.00	19.40
Dirichlet Deeplab [33]	✓	✗	✓	34.28	47.43	31.30	84.60
Density - Logistic Regression [4]	✗	✓	✓	4.65	24.36	57.16	13.39
SynBoost [11]	✗	✓	✓	43.22	15.79	72.59	18.75
Ours	✗	✗	✓	44.17	7.58	92.38	1.73

Table 3: Anomaly segmentation results on **Fishyscapes validation sets** (LostAndFound and Static), and the **Road Anomaly testing set**, with **WideResnet38** backbone. * indicate that the model requires additional learnable parameters. \dagger indicates that the results are obtained from the official code with our WideResnet38 backbone. Best and second best results in bold.

Methods	FS LostAndFound			FS Static			Road Anomaly		
	AUC \uparrow	AP \uparrow	FPR ₉₅ \downarrow	AUC \uparrow	AP \uparrow	FPR ₉₅ \downarrow	AUC \uparrow	AP \uparrow	FPR ₉₅ \downarrow
MSP [17]	89.29	4.59	40.59	92.36	19.09	23.99	67.53	15.72	71.38
Max Logit [17]	93.41	14.59	42.21	95.66	38.64	18.26	72.78	18.98	70.48
Entropy [18]	90.82	10.36	40.34	93.14	26.77	23.31	68.80	16.97	71.10
Energy [31]	93.72	16.05	41.78	95.90	41.68	17.78	73.35	19.54	70.17
Mahalanobis [26]	96.75	56.57	11.24	96.76	27.37	11.7	62.85	14.37	81.09
Meta-OoD [11]	93.06	41.31	37.69	97.56	72.91	13.57	-	-	-
\dagger Synboost* [11]	96.21	60.58	31.02	95.87	66.44	25.59	81.91	38.21	64.75
\dagger SML [20]	94.97	22.74	33.49	97.25	66.72	12.14	75.16	17.52	70.70
Deep Gambler [32]	97.82	31.34	10.16	98.88	84.57	3.39	78.29	23.26	65.12
Ours	98.96	58.81	4.76	99.61	92.08	1.52	87.63	45.10	44.58

it is worth noting that PEBAL reduces the amount of false positive pixels to 7.58 and 1.73 FPR on the two datasets. This result is publicly available on the Fishyscapes website.

Comparison on Fishyscapes validation sets and Road Anomaly. In Tables 3 and 4, we compare our approach on the Fishyscapes validation sets and Road Anomaly using two different backbones. Our model outperforms the previous methods by a large margin on all three benchmarks, regardless of the backbones and their segmentation accuracy. To verify the applicability of our method, except for the modern WideResnet38 backbone, we use a ResNet101 DeepLabv3+ to investigate the performance in terms of the size of the architecture and its inlier segmentation accuracy. The results demonstrate that our approach is applicable to a wide-range of segmentation models, indicating the effectiveness of PEBAL to adapt to real-world systems.

Table 4: Anomaly segmentation results on **Fishyscapes validation sets** (LostAndFound and Static), and the **Road Anomaly testing set**, with **Resnet101** backbone. * indicate that the model requires additional learnable parameters. † indicates that the results are obtained from the official code with our Resnet101 backbone. Best and second best results in bold.

Methods	FS LostAndFound			FS Static			Road Anomaly		
	AUC ↑	AP ↑	FPR ₉₅ ↓	AUC ↑	AP ↑	FPR ₉₅ ↓	AUC ↑	AP ↑	FPR ₉₅ ↓
MSP [17]	86.99	6.02	45.63	88.94	14.24	34.10	73.76	20.59	68.44
Max Logit [17]	92.00	18.77	38.13	92.80	27.99	28.50	77.97	24.44	64.85
Entropy [18]	88.32	13.91	44.85	89.99	21.78	33.74	75.12	22.38	68.15
Energy [31]	93.50	25.79	32.26	91.28	31.66	37.32	78.13	24.44	63.36
†SynthCP* [40]	88.34	6.54	45.95	89.9	23.22	34.02	76.08	24.86	64.69
†Synboost* [11]	94.89	40.99	34.47	92.03	48.44	47.71	85.23	41.83	59.72
SML [20]	96.88	36.55	14.53	96.69	48.67	16.75	81.96	25.82	49.74
Deep Gambler [32]	97.19	39.77	12.41	97.51	67.69	15.39	85.45	31.45	48.79
Ours	99.09	59.83	6.49	99.23	82.73	6.81	92.51	62.37	28.29

Moreover, our fine-tuning sacrifices only marginally the inlier segmentation accuracy (i.e., 0.2% - 0.7% mIoU on Cityscapes) for both backbones, achieving good performance on both inlier and anomaly segmentation. We present details of all inlier segmentation models (i.e., Cityscapes training setup and mIoU), and include more experimental results of other DeepLabv3+ checkpoints in supplementary material.

Remarks – Superior Performance on Challenging Benchmarks. Each dataset has different challenges. For example, the LostAndFound testing set considers only drivable areas with homogeneous normal scenes (i.e., road) and limited categories of abnormalities (i.e., road obstacles), leading to a relatively less challenging benchmark on which most methods can obtain good AUC performance, as shown in Tables 1, 3 and 4. On the contrary, Fishyscapes and Road-Anomaly contain large number of heterogeneous inlier and outlier pixels from diverse classes, leading to significantly more difficult testbeds than the LostAndFound testing set. Furthermore, Fishyscapes and RoadAnomaly contain domain shift compared with Cityscapes (e.g., both datasets contain different scenes than Cityscapes) and have different types/sizes of OoD objects. Most existing SOTA methods work ineffectively on these two datasets due to those challenges, while our adaptive pixel-level anomaly class learning helps our model effectively detect these challenging inlier and outlier pixels in the aforementioned heterogeneous and domain-shifted scenes, yielding substantial improvements (i.e., 20% to 50%) to previous approaches, as shown in Tables 2, 3 and 4.

4.5 Ablation Study

Table 5 shows the contribution of each component of our PEBAL on the LostAndFound testing set. All modules are trained with COCO OE images using AnomalyMix. Adding an extra OoD class to learn the OE training samples with entropy maximisation (**EM**) is our baseline (first row). To justify the effectiveness of our proposed joint training, we show the results using energy-based

Table 5: Ablation studies for anomaly segmentation on **LostAndFound**, with **WideResnet38** backbone, where all proposed modules are trained with COCO OE images with AnomalyMix. EM denotes the baseline method that adds an extra OoD class to learn the OE training samples with entropy maximisation (first row).

EM	ℓ_{ebm}	ℓ_{pal}	ℓ_{reg}	AUC \uparrow	AP \uparrow	FPR ₉₅ \downarrow
✓				96.88	69.02	8.03
	✓			97.88	70.24	8.92
		✓		98.67	72.73	3.81
	✓	✓		99.63	77.19	1.19
	✓	✓	✓	99.76	78.29	0.81

Table 6: The performance comparison of our approach on Fishyscapes benchmark w.r.t different **diversity of OE classes** (mean results over six random seeds), in terms of AP and FPR95.

Class Per.	FS LostAndFound		FS Static	
	AP \uparrow	FPR ₉₅ \downarrow	AP \uparrow	FPR ₉₅ \downarrow
1%	53.57 \pm 3.74	6.97 \pm 1.98	85.84 \pm 1.01	3.05 \pm 0.97
5%	52.16 \pm 3.88	6.58 \pm 1.95	90.57 \pm 1.75	1.93 \pm 0.52
10%	55.14 \pm 3.02	5.78 \pm 1.59	91.37 \pm 1.28	1.64 \pm 0.58
25%	55.48 \pm 3.32	5.98 \pm 1.27	91.28 \pm 1.94	1.77 \pm 0.18
50%	56.69 \pm 2.57	5.32 \pm 1.16	91.88 \pm 0.71	1.62 \pm 0.05
75%	57.86 \pm 2.83	5.11 \pm 1.69	91.85 \pm 0.56	1.63 \pm 0.09

models (ℓ_{ebm} without ℓ_{pal}) and pixel-wise abstention (ℓ_{pal} with pre-defined fixed penalty). Both outperform the baselines (AP=70.2, FPR=8.9 and AP=72.7, FPR=3.8 vs. AP=69, FPR=8.03), while our proposed joint training ($\ell_{ebm} + \ell_{pal}$) obtains 77.19% of AP and 1.19% of FPR, improving over each module by 4% to 7%. This indicates the effectiveness of our joint training and the significance of our proposed PAL with learnable adaptive energy-based penalties a_ω . Finally, the smoothness and sparsity regularisation losses stabilise the training and further improve the performance.

4.6 Outlier Samples and Computational Efficiency

Outlier Diversity and Efficiency. In Table 6, we randomly select 1%, 5%, 10%, 25% 50%, and 75% of COCO classes as the OE data during training and compute the mean results over six different random seeds. We achieve consistent AP and FPR performance regardless of the number of COCO classes used during the training on Fishyscapes. It is also worth noting that our approach can effectively learn the PEBAL model using **only one class** (1% in Table 6) of outlier data, which selects some of the irrelevant classes of COCO objects that are not possible to be found on road in real life (e.g., dining table, laptop, and clock). The results indicate that our model can consistently achieve SOTA performance on Fishyscapes without a careful selection of OE classes, demonstrating the robustness of our approach under diverse outlier classes. We also investigate the outlier sample efficiency of our model w.r.t smaller OE training sets with a fixed 100% COCO classes (80 classes) on Fishyscapes in Table 7, and we achieve consistently good performance regardless the number of outlier

Table 7: The performance comparison of our approach on Fishyscapes benchmark w.r.t different **amount of OE training samples** (mean results over six random seeds), in terms of AP and FPR95.

Train Size	FS LostAndFound		FS Static	
	AP \uparrow	FPR ₉₅ \downarrow	AP \uparrow	FPR ₉₅ \downarrow
5%	54.32 \pm 1.89	5.77 \pm 2.38	89.11 \pm 1.52	2.23 \pm 0.65
10%	56.28 \pm 1.05	4.66 \pm 1.36	90.02 \pm 0.57	1.67 \pm 0.28
25%	56.18 \pm 1.69	4.81 \pm 1.44	91.23 \pm 0.95	1.63 \pm 0.22
50%	57.34 \pm 1.19	4.75 \pm 1.32	91.29 \pm 0.92	1.67 \pm 0.17

training samples. All those experiments show the applicability of our PEBAL to real-world autonomous driving systems.

Computational Efficiency. We compare the computational efficiency of our PEBAL with previous SOTA Meta-OoD [6] and Synboost [11] in terms of the trainable parameters, training time and mean inference time per image, on an NVIDIA3090. As PEBAL requires the fine-tuning of the final classification block, it has only 1.3M parameters and each training epoch takes about 12 minutes, which is significantly less than the re-training approach Meta-OoD that has 137.1M parameters and each training epoch takes about 26 minutes, and the reconstruction based approach Synboost that takes about 33 minutes to train a epoch of its re-synthesis and dissimilarity networks with 157.3M parameters. Moreover, our method also has a much faster mean inference time of 0.55s compared to 0.85s of Meta-OoD and 1.95s of Synboost. Those results suggest the practicability of our model in real-world self-driving systems.

5 Conclusions and Discussions

We proposed a simple yet effective approach, named Pixel-wise Energy-biased Abstention Learning (PEBAL), to fine-tune the last block of a segmentation model to detect unexpected road anomalies. The approach introduces a non-trivial training that jointly optimises a novel pixel-wise abstention learning and an energy-based model to learn an adaptive pixel-wise anomaly class, in which a new pixel-wise energy-biased penalty estimation method is proposed to improve the precision and robustness to detect small and distant anomalous objects. The resulting model significantly reduces the false positive and false negative detected anomalies, compared with previous SOTA methods. The results on four benchmarks demonstrate the accuracy and robustness of our approach to detect anomalous objects regardless of the amount or diversity of exposed training outliers. Despite the remarkable performance on most datasets, PEBAL is not as effective on the most challenging dataset, Road Anomaly, that contains significantly more diverse and realistic anomalous objects. We plan to further enhance the generalisation of our model to accurately detect more unknown, diverse anomalies. ¹

¹Supported by Australian Research Council through grants DP180103232 and FT190100525.

References

1. Baur, C., Wiestler, B., Albarqouni, S., Navab, N.: Deep autoencoding models for unsupervised anomaly segmentation in brain mr images. In: International MICCAI Brainlesion Workshop. pp. 161–169. Springer (2018) [4](#)
2. Bevandi, P., et al.: Simultaneous semantic segmentation and outlier detection in presence of domain shift. In: German Conference on Pattern Recognition. pp. 33–47. Springer (2019) [1](#), [2](#), [3](#), [4](#), [10](#), [11](#)
3. Bevandić, P., Krešo, I., Oršić, M., Šegvić, S.: Discriminative out-of-distribution detection for semantic segmentation. arXiv preprint arXiv:1808.07703 (2018) [4](#)
4. Blum, H., Sarlin, P.E., Nieto, J., Siegwart, R., Cadena, C.: The fishyscapes benchmark: Measuring blind spots in semantic segmentation. arXiv preprint arXiv:1904.03215 (2019) [1](#), [4](#), [5](#), [8](#), [9](#), [11](#)
5. Chan, R., Lis, K., Uhlemeyer, S., Blum, H., Honari, S., Siegwart, R., Salzmann, M., Fua, P., Rottmann, M.: Segmentmeifyoucan: A benchmark for anomaly segmentation. NeurIPS (2021) [9](#)
6. Chan, R., Rottmann, M., Gottschalk, H.: Entropy maximization and meta classification for out-of-distribution detection in semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5128–5137 (2021) [1](#), [2](#), [3](#), [4](#), [5](#), [7](#), [9](#), [10](#), [13](#)
7. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 801–818 (2018) [9](#)
8. Chen, Y., Tian, Y., Pang, G., Carneiro, G.: Deep one-class classification via interpolated gaussian descriptor (2021) [4](#)
9. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016) [9](#)
10. Creusot, C., Munawar, A.: Real-time small obstacle detection on highways using compressive rbm road reconstruction. In: 2015 IEEE Intelligent Vehicles Symposium (IV). pp. 162–167. IEEE (2015) [4](#)
11. Di Biase, G., Blum, H., Siegwart, R., Cadena, C.: Pixel-wise anomaly detection in complex driving scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16918–16927 (2021) [1](#), [2](#), [4](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#)
12. El-Yaniv, R., et al.: On the foundations of noise-free selective classification. Journal of Machine Learning Research **11**(5) (2010) [4](#)
13. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. International journal of computer vision **88**(2), 303–338 (2010) [9](#)
14. Grathwohl, W., Wang, K.C., Jacobsen, J.H., Duvenaud, D., Norouzi, M., Swersky, K.: Your classifier is secretly an energy based model and you should treat it like one. arXiv preprint arXiv:1912.03263 (2019) [3](#), [5](#), [6](#)
15. Grcić, M., Bevandić, P., Šegvić, S.: Dense anomaly detection by robust learning on synthetic negative data. arXiv preprint arXiv:2112.12833 (2021) [2](#)
16. Haldimann, D., Blum, H., Siegwart, R., Cadena, C.: This is not what i imagined: Error detection for semantic segmentation through visual dissimilarity. arXiv preprint arXiv:1909.00676 (2019) [4](#)

17. Hendrycks, D., Basart, S., Mazeika, M., Mostajabi, M., Steinhardt, J., Song, D.: Scaling out-of-distribution detection for real-world settings. arXiv preprint arXiv:1911.11132 (2019) [5](#), [9](#), [10](#), [11](#), [12](#)
18. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. arXiv preprint arXiv:1610.02136 (2016) [4](#), [9](#), [10](#), [11](#), [12](#)
19. Hendrycks, D., Mazeika, M., Dietterich, T.: Deep anomaly detection with outlier exposure. arXiv preprint arXiv:1812.04606 (2018) [2](#), [4](#), [5](#)
20. Jung, S., Lee, J., Gwak, D., Choi, S., Choo, J.: Standardized max logits: A simple yet effective approach for identifying unexpected road obstacles in urban-scene segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15425–15434 (2021) [1](#), [2](#), [4](#), [5](#), [8](#), [9](#), [10](#), [11](#), [12](#)
21. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? arXiv preprint arXiv:1703.04977 (2017) [4](#)
22. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. arXiv preprint arXiv:1612.01474 (2016) [4](#)
23. Lateef, F., Ruichek, Y.: Survey on semantic segmentation using deep learning techniques. Neurocomputing **338**, 321–348 (2019) [1](#)
24. LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., Huang, F.: A tutorial on energy-based learning. Predicting structured data **1**(0) (2006) [3](#), [5](#), [6](#)
25. Lee, K., Lee, H., Lee, K., Shin, J.: Training confidence-calibrated classifiers for detecting out-of-distribution samples. arXiv preprint arXiv:1711.09325 (2017) [4](#)
26. Lee, K., Lee, K., Lee, H., Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. Advances in neural information processing systems **31** (2018) [9](#), [10](#), [11](#)
27. Li, C.L., Sohn, K., Yoon, J., Pfister, T.: Cutpaste: Self-supervised learning for anomaly detection and localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9664–9674 (2021) [8](#)
28. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. arXiv preprint arXiv:1706.02690 (2017) [4](#)
29. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014) [7](#)
30. Lis, K., Nakka, K., Fua, P., Salzmann, M.: Detecting the unexpected via image resynthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2152–2161 (2019) [1](#), [2](#), [3](#), [4](#), [9](#), [11](#)
31. Liu, W., Wang, X., Owens, J.D., Li, Y.: Energy-based out-of-distribution detection. arXiv preprint arXiv:2010.03759 (2020) [3](#), [5](#), [6](#), [10](#), [11](#), [12](#)
32. Liu, Z., Wang, Z., Liang, P.P., Salakhutdinov, R.R., Morency, L.P., Ueda, M.: Deep gamblers: Learning to abstain with portfolio theory. Advances in Neural Information Processing Systems **32**, 10623–10633 (2019) [3](#), [6](#), [10](#), [11](#), [12](#)
33. Malinin, A., Gales, M.: Predictive uncertainty estimation via prior networks. arXiv preprint arXiv:1802.10501 (2018) [11](#)
34. Mukhoti, J., Gal, Y.: Evaluating bayesian deep learning methods for semantic segmentation. arXiv preprint arXiv:1811.12709 (2018) [1](#), [4](#), [11](#)
35. Nijkamp, E., Hill, M., Zhu, S.C., Wu, Y.N.: Learning non-convergent non-persistent short-run mcmc toward energy-based model. arXiv preprint arXiv:1904.09770 (2019) [5](#)

36. Pinggera, P., Ramos, S., Gehrig, S., Franke, U., Rother, C., Mester, R.: Lost and found: detecting small road hazards for self-driving vehicles. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1099–1106. IEEE (2016) [4](#), [8](#)
37. Sultani, W., Chen, C., Shah, M.: Real-world anomaly detection in surveillance videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6479–6488 (2018) [9](#)
38. Vandenhende, S., Georgoulis, S., Proesmans, M., Dai, D., Van Gool, L.: Revisiting multi-task learning in the deep learning era. arXiv preprint arXiv:2004.13379 **2** (2020) [4](#)
39. Vojir, T., Šipka, T., Aljundi, R., Chumerin, N., Reino, D.O., Matas, J.: Road anomaly detection by partial image reconstruction with segmentation coupling. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15651–15660 (2021) [4](#)
40. Xia, Y., Zhang, Y., Liu, F., Shen, W., Yuille, A.L.: Synthesize then compare: Detecting failures and anomalies for semantic segmentation. In: European Conference on Computer Vision. pp. 145–161. Springer (2020) [1](#), [2](#), [4](#), [12](#)
41. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6023–6032 (2019) [8](#)
42. Zhu, Y., Sapra, K., Reda, F.A., Shih, K.J., Newsam, S., Tao, A., Catanzaro, B.: Improving semantic segmentation via video propagation and label relaxation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8856–8865 (2019) [9](#)