

Residual Pattern Learning for Pixel-wise Out-of-Distribution Detection in Semantic Segmentation

Yuyuan Liu^{1†} Choubo Ding^{1†} Yu Tian² Guansong Pang³ Vasileios Belagiannis⁴
 Ian Reid¹ Gustavo Carneiro^{5,1}

¹ Australian Institute for Machine Learning, University of Adelaide ² Harvard University

³ Singapore Management University ⁴ Friedrich-Alexander-Universität Erlangen-Nürnberg ⁵ University of Surrey

<https://github.com/yyluo1/RPL>




Figure 1: **Visualisations for pixel-wise anomaly scores.** The top row shows real-world images along with **bounding boxes** around anomalous objects with respect to the training categories. Subsequent rows show the anomaly score map for different methods, namely SoftMax [14], PEBAL [35], and our approach. Note that our method successfully detect all anomalous objects without incorrectly deeming in-distribution pixels to be anomalous, while the previous SOTA PEBAL [35] either fails to detect the anomaly (column 2) or mis-detects inliers as anomalies (columns 3-5).

Abstract

Semantic segmentation models classify pixels into a set of known (“in-distribution”) visual classes. When deployed in an open world, the reliability of these models depends on their ability to not only classify in-distribution pixels but also to detect out-of-distribution (OoD) pixels. Historically, the poor OoD detection performance of these models has motivated the design of methods based on model re-training using synthetic training images that include OoD visual objects. Although successful, these re-trained methods have two issues: 1) their in-distribution segmentation accuracy may drop during re-training, and 2) their OoD detection

accuracy does not generalise well to new contexts outside the training set (e.g., from city to country context). In this paper, we mitigate these issues with: (i) a new residual pattern learning (RPL) module that assists the segmentation model to detect OoD pixels with minimal deterioration to inlier segmentation accuracy; and (ii) a novel context-robust contrastive learning (CoroCL) that enforces RPL to robustly detect OoD pixels in various contexts. Our approach improves by around 10% FPR and 7% AuPRC previous state-of-the-art in Fishyscapes, Segment-Me-If-You-Can, and RoadAnomaly datasets.

1. Introduction

Semantic segmentation is a fundamental computer vision task that classifies each image pixel into a set of in-distribution visual classes (or inliers) [11]. Despite its success in closed-set benchmarks, many real-world applications need to be able to cope with open-world scenarios, in which closed-world segmentation models will mistakenly estimate out-of-distribution (OoD) visual objects (or anomalies) as one of the in-distribution classes. For some use-cases, such as autonomous driving, such errors could be catastrophic [11, 42]. One possible way to mitigate this risk is based on the development of pixel-wise OoD detection methods that work together with an in-distribution segmentation model to produce an OoD map on top of the closed-set segmentation mask [14, 17, 30]. Generally, this can be achieved by measuring the in-distribution segmentation uncertainty using classification entropy [9], logits [17] or posterior distribution [14, 31] by *freezing* the inlier segmentation model. Even though those approaches do not jeopardize the inlier performance, they tend to produce low uncertainty for hard OoD pixels that share similar patterns with in-distribution objects, leading to unsatisfactory performance when dealing with complex scenes, as shown in Fig. 1, row 2 (labelled as 'Softmax').

Different from the OoD detection methods above, state-of-the-art (SOTA) OoD pixel detectors are based on methods that *re-train* the inlier segmentation models with OoD data [4, 10, 35]. Such OoD data is obtained from Outlier Exposure (OE) [15], which introduces additional OoD images [4] or synthetically mixes OoD objects to training images [10, 35]. Then, the closed-set model is re-trained to optimise the OE pixels via entropy maximisation [4] or abstention learning [35]. Such re-training boosts improves the pixel-wise OoD detection accuracy, but unlike the approaches that *freeze the inlier segmentation* [14, 17], it can worsen the in-distribution segmentation accuracy. For example, pixels from minority categories (e.g., “fence”, “sidewalks”) can be mis-classified as majority classes (e.g., “road”). **We argue that the ability to detect anomalies should be achieved with minimal detriment to the closed-set segmentation accuracy.** Another critical issue that affects OoD pixel detectors is their narrow context reliability [4, 9, 10, 17, 35], where context is defined by the in-distribution pixels that represent the scene surroundings, such as city (column 1) and country (columns 2-5) contexts in Fig. 1. This is because entropy or energy-based optimisation algorithms [4, 9, 35] train pixel-wise anomaly scores independently, but ignore the relationships between OoD and contexts, leading to over-confident minimisation of inlier uncertainty and a weak maximisation of outlier pixels. Therefore, an anomaly detector can incorrectly classify inliers with unfamiliar patterns as outliers in new contexts, or misclassify OoD pixels within small anomalies as inliers.

For example, although the SOTA PEBAL [35] (Fig. 1, row 3) barely detects the tiny anomaly in the city context of column 1, it ignores the outlier in column 2 and mis-detects anomalies in many inliers in columns 3-5, when the context changes from city to country. Given that such context changes are common and unpredictable in real-world scenarios, **it is important to develop a pixel-wise OoD detector that can generalise to various contexts.**

In this paper, we address the points above with a new pixel-wise OoD detection module, called Residual Pattern Learning (RPL), which consists of an external module attached to a **frozen** closed-set segmentation network. RPL is trained to learn the residual pattern of anomalies based on intermediate features and induce the (frozen) segmentation classifier to produce high uncertainty for the potentially anomalous regions. By not *re-training* the segmentation model, RPL guarantees high accuracy in detecting OoD pixels while minimising the impact on the closed-set segmentation performance. The effective training of RPL depends on our proposed positive energy loss that focuses only on the energy score of anomalies to deal with the imbalanced distribution of inlier and outlier samples in pixel-wise OoD detection. Moreover, we address open-world context robustness with our new context-robust contrastive learning (CoroCL), which optimises the pixel-wise embeddings by exploring the relationship between anomalies and inliers in different contexts. To summarise, our contributions are:

1. The residual pattern learning (RPL) module that induces the closed-set segmentation model to detect potential anomalies, which is a novel perspective for pixel-wise OoD detection;
2. A context-robust contrastive learning (CoroCL) that is designed to generalise the detection of OoD pixels to new contexts (see Fig. 1, row 4); and
3. A novel positive energy loss function that deals with the imbalanced distribution of inliers and outliers by focusing only on the energy score of anomalies.

Our approach is shown to be the most accurate method in various scene contexts, yielding improvements of around 10% FPR and 7% AuPRC over the SOTA results in Fishyscapes, Segment-Me-If-You-Can, and RoadAnomaly datasets. Our approach is also shown to be the most stable method in various scene contexts and to be easily integrated to other SOTAs, such as Meta-OoD [4] and PEBAL [35].

2. Related Work

Semantic segmentation refers to pixel-wise classification methods that are useful for applications like autonomous driving [11, 42] or scene understanding [16]. Influenced

by fully convolutional network (FCN) [29], current approaches provide better segmentation details by maintaining the high-level image representations [19,34,41], or construct a feature pyramid to merge the multi-scale representations to learn global image contexts [24,44]. DeepLab methods [5–7] utilise dilated convolution to filter the incoming features and enlarge the receptive field. Our model is based on DeepLabV3+ [7] to enable a fair comparison with recently proposed anomaly detection approaches.

Pixel-wise OoD Detection methods are based on approaches that *freeze* [9,17] or *re-train* the segmentation model [4,10,35]. Most of the early anomaly detection approaches [30] freeze the segmentation model, and OoD pixels are classified based on posterior distribution measures (e.g., entropy) [14,20] or on distance measures (e.g., Mahalanobis) [21]. Those methods have low computational cost and do not affect the original inlier segmentation accuracy, but show inaccurate OoD detection. Recently, Synboost [9] improved OoD detection accuracy with an extra FCN network to classify anomalies based on the segmentation model’s output and a generative model [37]. Unfortunately, the incorrect information from segmentation output introduces confirmation bias to the training process, leading to poor generalisation. Instead of working with the segmentation model’s output, our RPL module is integrated with the segmentation network. Given the richer information content of intermediate features, compared with the output logit [9] or classification probabilities [30], our approach can detect potential anomalies more effectively.

Promising results are shown by approaches that re-train the segmentation model [15]. Meta-OoD [4] introduces OoD images from an outlier dataset (e.g., COCO [25]) to the inlier training set to boost the entropy results in outlier regions. To better simulate anomalies, PEBAL [35], Dense-hybrid [10] and [1] randomly crop anomalous objects from the outlier exposure (OE) datasets into inlier images, where PEBAL achieves SOTA performance with pixel-wise abstention learning [28] and energy regularisation [27]. However, these re-training approaches can degrade the closed-set segmentation accuracy, which motivated us to introduce a method that produces pixel-wise anomalous scores with minimal deterioration to the inlier segmentation performance.

Supervised Contrastive Learning [18,22,38] has shown promising results for classification tasks by leveraging labelled data to learn well structured feature representations. These methods build an additional projector to extract image embeddings and push apart embedding pairs from different classes while pulling together same class pairs. Such optimisation allows the model to explore the relationship among the input representations, instead of only relying on individual samples’ classification results for training [36,38]. Motivated by this, our proposed Context-robust

Contrastive Learning (CoroCL) aims to produce consistent representations for RPL by understanding the relationship between the pixel-wise contexts and anomaly embeddings.

3. Methodology

Before we describe our approach and the training process, we introduce our dataset for the pixel-wise OoD detection. We assume to have the inlier dataset (e.g., Cityscapes [8]) with $\mathcal{D}^{in} = \{(\mathbf{x}_i^{in}, \mathbf{y}_i^{in})\}_i^{|\mathcal{D}^{in}|}$, where $\mathbf{x}^{in} \in \mathcal{X} \subset \mathbb{R}^{H \times W \times 3}$ is the input image with resolution $H \times W$, and $\mathbf{y}^{in} \in \mathcal{Y}^{in} \subset \{1, \dots, C\}^{H \times W}$ is the segmentation map, with C closed-set categories. We also have an outlier dataset (e.g., COCO [25]) with $\mathcal{D}^{out} = \{(\mathbf{x}_i^{out}, \mathbf{y}_i^{out})\}_i^{|\mathcal{D}^{out}|}$, where $\mathbf{x}^{out} \in \mathcal{X}$, and $\mathbf{y}^{out} \in \mathcal{Y}^{out} \subset \{0, P\}^{H \times W}$ is a binary mask distinguishing foreground and background objects, where $P > C$ to ensure the OoD are distinguishable from the inlier categories. Following previous anomaly detection approaches [10,35], we define the Outlier Exposure (OE) that grabs anomalous regions in \mathcal{D}^{out} and copies them to the inlier images from \mathcal{D}^{in} during training without any overlap policies, as follows:

$$\begin{aligned} \mathbf{x}^{oe} &= (1 - \mathbf{m}) \odot \mathbf{x}^{in} + \mathbf{m} \odot \mathbf{x}^{out} \\ \mathbf{y}^{oe} &= (1 - \mathbf{m}) \odot \mathbf{y}^{in} + \mathbf{m} \odot \mathbf{y}^{out}, \end{aligned} \quad (1)$$

where $\mathbf{m} = \mathbb{1}(\mathbf{y}^{out} = P)$ ($\mathbb{1}(\cdot)$ is an indicator function) is a binary map that is equal to 1 for the outlier pixels, but equal to 0, otherwise. This allows us to define the OE (mix-content) dataset with $\mathcal{D}^{oe} = \{(\mathbf{x}_i^{oe}, \mathbf{y}_i^{oe}, \mathbf{m}_i)\}_i^{|\mathcal{D}^{oe}|}$.

3.1. Residual Pattern Learning (RPL) for Pixel-wise Anomaly Objects

As depicted in Fig. 2, the original segmentation model consists of a fully convolutional network (e.g., ResNet [12]) denoted by a backbone $f_{\phi_{fcn}} : \mathcal{X} \rightarrow \mathcal{Z}$, which transforms the image to a representation $\mathbf{z} \in \mathcal{Z} \subset \mathbb{R}^{H' \times W' \times C'}$, and the feature extraction block (e.g., ASPP [7]) defined by $f_{\phi_{aspp}} : \mathcal{Z} \rightarrow \mathcal{K}$ embeds \mathbf{z} into the space of $\mathcal{K} \subset \mathbb{R}^K$. Then, the segmentation head $f_{\phi_{seg}} : \mathcal{K} \rightarrow [0, 1]^{H \times W \times C}$ maps from \mathcal{K} to C segmentation channels (more details are shown in Supplementary Sec.A), producing the inlier segmentation:

$$\tilde{\mathbf{y}} = f_{\phi_{seg}}(f_{\phi_{aspp}}(f_{\phi_{fcn}}(\mathbf{x}))). \quad (2)$$

Re-training approaches [4,10,35] optimise the segmentation network to recognise the anomaly patterns via the entropy [4,10] or energy [35] from $f_{\phi_{seg}}(\cdot)$, as depicted in Fig. 3b. Such re-training yields promising results by boosting anomaly uncertainty, but they also shift the inlier’s decision boundaries and produce mis-classified inlier predictions. On the other hand, as depicted in Fig. 3a, approaches that *freeze the inlier segmentation* [9,14,17,30] and use segmentation output information (e.g., logits [17],

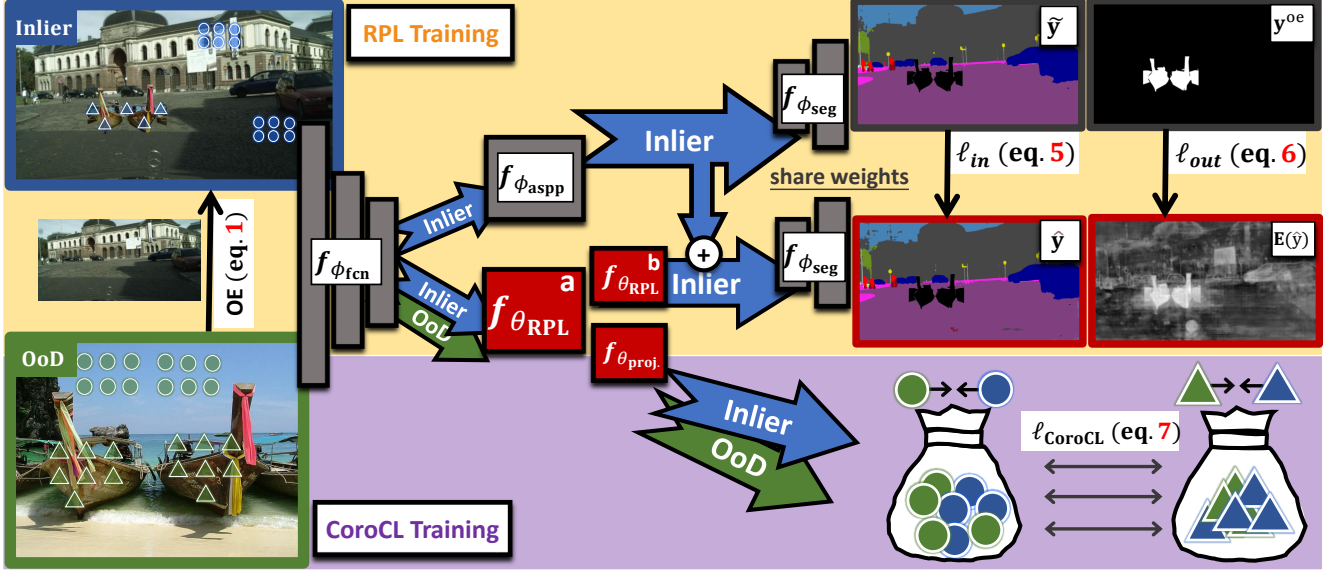


Figure 2: The **RPL module** is trained to approximate the inlier segmentation \tilde{y} with $\hat{y} = f_{\phi_{\text{seg}}}(\cdot)$, which uses as input the model’s intermediate features and RPL’s output, and to find the OoD pixels in y^{oe} . **The training of CoroCL** pulls together embedding pairs that both belong to in-distribution (●, ●) or OoD pixels (▲, ▲), while pushing in-distribution embeddings apart from OoD embeddings.

prediction [14] or entropy [9]) have shown poor performance when dealing with hard outliers. Our RPL in Fig. 3c is designed to address the issues above by placing an external residual module, denoted by $f_{\theta_{\text{rpl}}} : \mathcal{Z} \rightarrow \mathcal{K}$, between the FCN layers $f_{\phi_{\text{fcn}}}(\cdot)$ and the segmentation head $f_{\phi_{\text{seg}}}(\cdot)$. Note that the false positive anomaly detection (i.e., inliers located within the uncertain regions) is a side effect of the production of OoD masks that affects all OoD detectors and is orthogonal to the problem of inlier segmentation accuracy.

The RPL result is added to the ASPP output and decoded by the segmentation head with

$$\hat{y} = f_{\phi_{\text{seg}}}\left(f_{\phi_{\text{aspp}}}\left(f_{\phi_{\text{fcn}}}(\mathbf{x})\right) + f_{\theta_{\text{rpl}}}\left(f_{\phi_{\text{fcn}}}(\mathbf{x})\right)\right). \quad (3)$$

Our training differs from previous SOTA methods given that we **freeze the entire segmentation model** (i.e., ϕ_{fcn} , ϕ_{aspp} and ϕ_{seg}) during the training that only optimises the RPL block θ_{rpl} . Therefore, the decision boundary for in-distribution categories is always fixed, so the accuracy of the close-set semantic segmentation prediction, i.e., \tilde{y} from (2), will suffer minimal degradation.

3.2. RPL Training

Intuitively, RPL is trained to approximate the inlier pixel classification from the segmentation model and at the same time to detect OoD pixels. Hence, the training depends on the closed-set segmentation \tilde{y} from (2) and on our model segmentation \hat{y} from (3) to optimise the following function:

$$\ell_{\text{RPL}}(\mathcal{D}^{\text{oe}}, \theta_{\text{rpl}}) = \ell_{\text{in}}(\mathcal{D}^{\text{oe}}, \theta_{\text{rpl}}) + \alpha \times \ell_{\text{out}}(\mathcal{D}^{\text{oe}}, \theta_{\text{rpl}}), \quad (4)$$

where α weights the contribution of the second loss function (we set $\alpha = 0.05$ following [26] for all the experiments). The first term in (4) is the loss function for RPL to recognise the inlier features, as defined below:

$$\ell_{\text{in}}(\mathcal{D}^{\text{oe}}, \theta_{\text{rpl}}) = \sum_{(\mathbf{x}^{\text{oe}}, \mathbf{y}^{\text{oe}}, \mathbf{m}) \in \mathcal{D}^{\text{oe}}} \sum_{\omega \in \Omega} (1 - \mathbf{m}(\omega)) \left(\ell_{\text{ce}}(\text{OneHot}(\tilde{y}(\omega)), \hat{y}(\omega)) + \ell_{\text{reg}}(H(\tilde{y}(\omega)), H(\hat{y}(\omega))) \right), \quad (5)$$

where Ω is the image lattice of size $H \times W$, \mathbf{m} represents the anomaly binary mask from (1) at pixel ω , $\text{OneHot}(\tilde{y})$ returns the one-hot representation of \tilde{y} , $\ell_{\text{ce}}(\cdot)$ denotes the cross-entropy loss, H represents the Shannon entropy, and the dis-similarity regularisation is defined as $\ell_{\text{reg}}(a, b) = \left\| \frac{(a-b)}{t} \right\|_2^2$. The second term in (4) is our proposed positive energy that maximises the energy at OoD pixels, as in:

$$\ell_{\text{out}}(\mathcal{D}^{\text{oe}}, \theta_{\text{rpl}}) = \sum_{(\mathbf{x}_i^{\text{oe}}, \mathbf{y}_i^{\text{oe}}, \mathbf{m}_i) \in \mathcal{D}^{\text{oe}}} \sum_{\omega \in \Omega} \max(-\mathbf{m}_i(\omega)E(\hat{y}_i(\omega)), 0), \quad (6)$$

where $E(x) = -\log \sum_{i=1}^C \exp(x_i)$, and C is the number of inlier categories. The positive energy in (6) has two advantages over previous OoD energy-based methods for detecting OoD pixels [27, 35]: 1) the hinge loss in [27, 35] depends on two hard-to-tune hyperparameters to minimise the inlier energy and maximise the outlier energy, while our loss in (6) does not have any hyperparameters; 2) given that

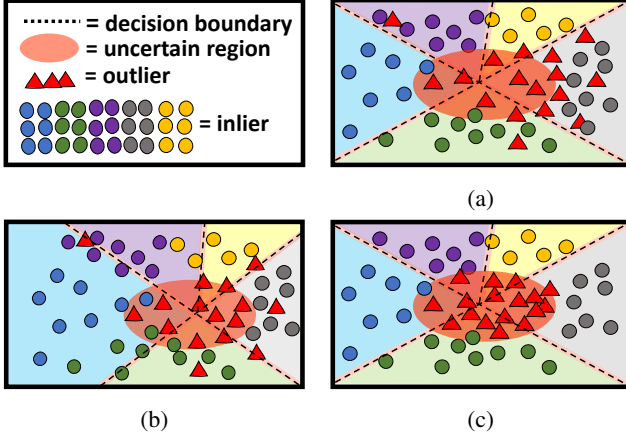


Figure 3: Sketch anomaly detection models. (a) Approaches [9, 17] that freeze the inlier segmentation model maintain inlier categories segmentation accuracy, but fail when hard OoD pixels share patterns with inliers. (b) Re-training approaches [4, 10, 35] fine-tune inlier decision boundaries to detect anomalies but can mis-classify inliers. (c) Our RPL pushes the OoD pixels toward the uncertainty region by minimising the impact to the inlier decision boundary. Note that the inliers inside the uncertain region are the false positive anomalies detected by OoD methods.

there are many more in-distribution than OoD pixels, the hinge loss in [27, 35] has a poor convergence for anomalous pixels, whereas our loss in (6) mitigates such weak optimisation for outliers by exclusively focusing on them.

3.3. Training RPL with Context-robust Contrastive Learning (CoroCL)

Most OoD detection methods fail to distinguish between unfamiliar inlier pixels and potentially OoD pixels when the context changes because those approaches lack understanding of the relationship between contexts and the anomalous objects. With the goal of exploring such relationship, our proposed CoroCL (as displayed in Fig. 2) involves a new contrastive learning loss function and a projection layer, **built on top of the RPL block**, represented by $f_{\theta_{\text{proj}}} : \mathcal{K} \rightarrow \mathbb{R}^{I \times R}$ that takes the output from RPL and project it to a space of I embeddings of R dimensions. To explain the projection layer, we need to subdivide the RPL model as $\theta_{\text{rpl}} = \{\theta_{\text{rpl}}^a, \theta_{\text{rpl}}^b\}$, where θ_{rpl}^a denotes the RPL main layers and θ_{rpl}^b represents the last convolutional layers of θ_{rpl} , with the proposed projection layer receiving the features from $f_{\theta_{\text{rpl}}^a}(\cdot)$. CoroCL trains the main layers of the RPL module and projection layer to pull the embeddings belonging to the same category (inliers or OoD) closer together and push away the embeddings from different categories. We target the pixel-wise embedding learning [38, 39], but to enable the recognition of anomalies across different contexts, our approach uses the samples via both OE (mix-content)

dataset \mathcal{D}^{oe} and the outlier dataset \mathcal{D}^{out} . Let us assume that we have the embeddings from the OE dataset denoted by $\mathcal{R}^{oe} = \{\mathbf{r} | \mathbf{r} = f_{\theta_{\text{proj}}}(f_{\theta_{\text{rpl}}^a}(f_{\phi_{\text{fcn}}}(\mathbf{x}^{oe}))), (\mathbf{x}^{oe}, \mathbf{y}^{oe}, \mathbf{m}) \in \mathcal{D}^{oe}\}$ that combines anomalies with the training set domain, and the embeddings from the outlier dataset represented by $\mathcal{R}^{out} = \{\mathbf{r} | \mathbf{r} = f_{\theta_{\text{proj}}}(f_{\theta_{\text{rpl}}^a}(f_{\phi_{\text{fcn}}}(\mathbf{x}^{out}))), (\mathbf{x}^{out}, \mathbf{y}^{out}) \in \mathcal{D}^{out}\}$.

Given that validation images are mainly based on the driving scenes (e.g., \mathcal{D}^{in}), we use \mathcal{R}^{oe} to form the anchor set $\mathcal{A} = \{\mathbf{r}_i | \mathbf{r}_i \sim \mathbf{r}, \mathbf{r} \in \mathcal{R}^{oe}, \mathbf{r}_i \in \mathbb{R}^R\}$ by random sampling \mathcal{R}^{oe} to extract a subset of the embeddings from $\mathbf{r} \in \mathcal{R}^{oe}$ (more details is shown in Supplementary Sec.E). To make the model robust to various contexts, we form the contrastive set for CoroCL with $\mathcal{C} = \{\mathbf{r}_i | \mathbf{r}_i \sim \mathbf{r}, \mathbf{r} \in (\mathcal{R}^{oe} \cup \mathcal{R}^{out}), \mathbf{r}_i \in \mathbb{R}^R\}$ by random sampling \mathcal{R}^{oe} and \mathcal{R}^{out} to extract a subset of the embeddings from $\mathbf{r} \in (\mathcal{R}^{oe} \cup \mathcal{R}^{out})$. CoroCL is defined by:

$$\ell_{\text{CoroCL}}(\mathcal{D}^{oe}, \mathcal{D}^{out}, \theta_{\text{proj}}, \theta_{\text{rpl}}^a) = \sum_{\mathbf{r}_i \in \mathcal{A}} \sum_{\mathbf{p} \in \mathcal{P}(\mathbf{r}_i)} -\log \frac{\exp(\mathbf{r}_i \cdot \mathbf{p} / \tau)}{\exp(\mathbf{r}_i \cdot \mathbf{p} / \tau) + \sum_{\mathbf{n} \in \mathcal{N}(\mathbf{r}_i)} \exp(\mathbf{r}_i \cdot \mathbf{n} / \tau)}, \quad (7)$$

where $\mathcal{P}(\mathbf{r}_i) = \{\mathbf{p} | \mathbf{p} = \mathbf{r}_j, \mathbf{r}_j \in \mathcal{C}, \mathbf{m}_i = \mathbf{m}_j\}$ is the set of positive embeddings for the anchor embedding, with the class (inlier or outlier) of the anchor denoted by \mathbf{m}_i and of the positive embedding represented by \mathbf{m}_j , $\mathcal{N}(\mathbf{r}_i) = \{\mathbf{n} | \mathbf{n} = \mathbf{r}_j, \mathbf{r}_j \in \mathcal{C}, \mathbf{m}_i \neq \mathbf{m}_j\}$ is the set of negative embeddings for the anchor embedding, with the class (inlier or outlier) of the anchor denoted by \mathbf{m}_i and of the negative embedding represented by \mathbf{m}_j .

The overall loss function to train the RPL module is:

$$\theta_{\text{rpl}}^*, \theta_{\text{proj}}^* = \arg \min_{\theta_{\text{rpl}}, \theta_{\text{proj}}} \ell_{\text{RPL}}(\mathcal{D}^{oe}, \theta_{\text{rpl}}) + \ell_{\text{CoroCL}}(\mathcal{D}^{oe}, \mathcal{D}^{out}, \theta_{\text{rpl}}^a, \theta_{\text{proj}}), \quad (8)$$

where $\theta_{\text{RPL}}^a \subset \theta_{\text{RPL}}$ denotes the main layers of θ_{RPL} . Note that the loss function (8) is applied only to the RPL module and projection layer instead of the original segmentation model, so that our model is trained to recognise anomaly patterns in various contexts with minimal impact to the performance of closed-set segmentation categories.

Inference. During inference, the inlier segmentation map is represented by $\hat{\mathbf{y}}$ from (2), and the OoD pixels are estimated from the energy score $E(\cdot)$ of (6) on the segmentation map $\hat{\mathbf{y}}$ from (3), for each pixel position ω of a test image \mathbf{x} . We also apply Gaussian smoothing to produce the final energy map, following previous works [17, 35].

4. Experiments

In this section, we first explain the experimental setup, then compare our method with the current SOTA, present the inlier segmentation accuracy, and show the ablation study for RPL and CoroCL.

Table 1: **Comparison with SOTA methods on Fishyscapes and SMIYC validation sets.** All approaches are based on the **DeepLabv3+** architecture with WiderResNet38 backbone. Best results are in **boldface**, and the worst results from previous SOTAs caused by the lack of context robustness are marked in **cyan**.

Methods	Fishyscapes (validation set)						SMIYC (validation set)					
	Static			L&F			Anomaly			Obstacle		
	FPR ↓	AuPRC ↑	AUROC ↑	FPR ↓	AuPRC ↑	AUROC ↑	FPR ↓	AuPRC ↑	F1* ↑	FPR ↓	AuPRC ↑	F1* ↑
Maximum softmax [14] [baseline]	23.31	26.77	93.14	10.36	40.34	90.82	60.2	40.4	42.6	3.8	43.4	53.7
Mahalanobis [21] [baseline]	11.7	27.37	96.76	11.24	56.57	96.75	86.4	22.5	31.7	26.1	25.9	27.7
SML [17] [ICCV'21]	12.14	66.72	97.25	33.49	22.74	94.97	84.13	21.68	28.00	91.31	18.60	28.39
Synboost [9] [CVPR'21]	25.59	66.44	95.87	31.02	60.58	96.21	30.9	68.8	65.6	2.8	81.4	73.2
Meta-OoD [4] [ICCV'21]	13.57	72.91	97.56	37.69	41.31	93.06	17.43	80.13	74.3	0.41	94.14	88.4
DenseHybrid [10] [ECCV'22]	4.17	76.23	99.07	5.09	69.79	99.01	52.65	61.08	53.72	0.71	89.49	81.05
PEBAL [35] [ECCV'22]	1.52	92.08	99.61	4.76	58.81	98.96	36.74	53.10	57.99	7.92	10.45	22.10
RPL+CoroCL [Ours]	0.85	92.46	99.73	2.52	70.61	99.39	7.18	88.55	82.90	0.09	96.91	91.75

Table 2: **Comparison with SOTA on RoadAnomaly validation set.** All approaches are based on the **DeepLabv3+** architecture and best results are in **boldface**.

Method	RoadAnomaly		
	FPR ↓	AuPRC ↑	AuROC ↑
Maximum softmax [14] [baseline]	68.15	22.38	75.12
Gambler [28] [NIPS'19]	48.79	31.45	85.45
SynthCP [40] [ECCV'20]	64.69	24.86	76.08
Synboost [9] [CVPR'21]	59.72	41.83	85.23
SML [17] [ICCV'21]	49.74	25.82	81.96
GMMSeg [23] [NIPS'22]	47.90	34.42	84.71
PEBAL [35] [ECCV'22]	44.58	45.10	87.63
RPL+CoroCL [Ours]	17.74	71.60	95.72

Training set. We establish our experiments using the urban driving dataset **Cityscapes** [8], which consists of 2, 975 images for training, 500 for validation and 1, 525 for testing. Every image in the dataset has resolution $2, 048 \times 1, 024$, and there are 19 classes in total. Following previous works [4, 35], we rely on **COCO** [25] for the outlier exposure (OE). COCO contains objects captured in various contexts and images have at least 480 pixels for height or width. We only consider COCO samples that have visual classes that do not overlap with Cityscapes classes, which allowed us to form the dataset \mathcal{D}^{out} with 46, 751 images.

Validation sets. We evaluate our approach on multiple benchmarks. **Fishyscapes** [2] has high-resolution images for anomaly detection, with two validation subsets: FS-L&F and FS-Static. FS-L&F consists of 1000 images from the LostAndFound dataset [33] with refined labels, and FS-Static contains 30 images with synthetic anomalous objects. **RoadAnomaly** [26] is another large anomaly validation set that contains real-world anomalous objects with 60 images from the Internet. **Segment-If-You-Can (SMIYC)** [3] has two subsets, including *AnomalyTrack* that contains large anomaly objects on different contexts (partially overlapped with [26]), and *ObstacleTrack* which contains small obstacles on the road ahead. There are 10 and 30 images for validation, and 100 and 327 images for online test in AnomalyTrack and ObstacleTrack, respectively. In addition, Chan et al. [3] eliminated the mislabeled normal frames from LostAndFound [33] and proposed the SMIYC-L&F to simu-

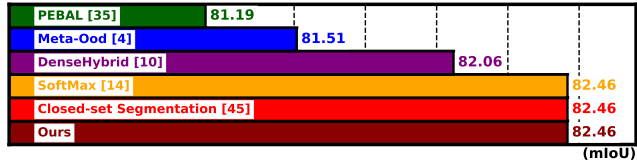


Figure 4: Cityscapes **test set** results based on **closed-set segmentation model** (i.e., without considering the OoD detection) with **single-scale** sliding evaluation.

late real-world driving scene situations.

Implementation Details. Our experiments are based on the DeepLabv3+ architecture with WiderResNet38 as the backbone. We present the results with other backbones (e.g., resnet101, resnet50) in Supplementary Sec.D. Following [3, 4, 10, 35], our training is based on the pre-trained checkpoints from Cityscapes [8], where WiderResNet38 is from [45]. During training, we set the initial learning rate to be $7.5e^{-5}$ and utilise poly learning rate decay with $(1 - \frac{iter}{max_iter})^{0.9}$ and 40 epochs in total. We set temperature hyper-parameters $t = 1$ for ℓ_{reg} in (5) and $\tau = 0.10$ in (7) for all the experiments. Our RPL block $f_{\theta_{ipl}^a}(\cdot)$ is formed with the ASPP layers, followed by a layer $f_{\theta_{ipl}^b}(\cdot)$ that expands the channels to be the same as the input of segmentation head $f_{\phi_{seg}}(\cdot)$ (from 256 to 304). The projection layer $f_{\theta_{proj}}(\cdot)$ for contrastive learning is represented by a one-layer perceptron. Supplementary Sec.A provides the details about the RPL architecture, and Supplementary Sec.C contains more implementation details.

Number of parameters in OoD detector. The total number of RPL parameters for training is **30.89M**, which is significantly lower than the parameters in Meta-OoD [4] and DenseHybrid [10], which are $\sim 137.11M$.

Evaluation metrics. Following [4, 17, 35], we compute the receiver operating characteristics (AuROC), the area under precision recall curve (AuPRC), false positive rate at a true positive rate (FPR) of 95% and F1 score to evaluate our approach. We utilise mean Intersection Over Union (mIoU) to measure inlier segmentation performance, following the common practice in segmentation [5, 38, 39, 44].

We evaluate approaches for all the validation/test sets with one consistent model to simulate the real-world chal-

Table 3: **Comparison with SOTAs** on Fishyscapes and SMIYC official test benchmarks^{1,2}. All approaches are based on the **DeepLabv3+** architecture with WiderResNet38 backbone. Best results are in **boldface**, and the worst results caused by the lack of context robustness are marked in **cyan**. Additional experiments with methods fine-tuned with extra training sets (e.g., Vistas [32], Wilddash2 [43]) are in Supplementary Sec.B. * denotes the results reported in the paper.

Methods	Fishyscapes (test)				SMIYC (test)				Overall	
	Static		L&F		Anomaly		Obstacle			
	FPR ↓	AuPRC ↑	FPR ↓	AuPRC ↑	FPR ↓	AuPRC ↑	FPR ↓	AuPRC ↑	FPR ↓	AuPRC ↑
Resynthesis [26] _[ICCV*19]	27.13	29.6	48.05	5.70	25.93	52.28	4.70	37.71	26.45	31.32
Embedding [2] _[ICV*19]	20.25	44.03	30.02	3.55	70.76	37.52	46.38	0.82	41.85	21.48
Synboost [9] _[CVPR*19]	18.75	72.59	15.79	43.22	61.86	56.44	3.15	71.34	24.89	60.90
Meta-OoD [4] _[ICCV*21]	8.55	86.55	35.14	29.96	15.00	85.47	0.75	85.07	14.86	71.76
DenseHybrid [10] _[ECCV*22]	5.51	72.27	6.18	43.90	62.25	42.05	6.02	80.79	19.99	59.75
GMMSeg [23]* _[NIPS*22]	15.96	76.02	6.61	55.63	-	-	-	-	-	-
PEBAL [35] _[ECCV*22]	1.73	92.38	7.58	44.17	40.82	49.14	12.68	4.98	15.70	47.67
RPL+CoroCL [Ours]	0.52	95.96	2.27	53.99	11.68	83.49	0.58	85.93	3.76	79.84

linging scenario. All methods are based on the same segmentation architecture for fair competition.

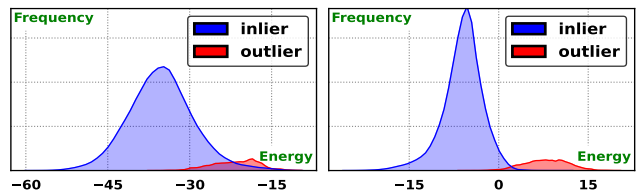
4.1. Comparison with SOTA Methods

In this section, we compare our approach with others for both inlier and pixel-wise OoD detection. We first evaluate our approach on all **validation sets** with AP and FPR in Tab. 1, where we follow [35] and measure AuROC on Fishyscapes [2] and the F1 score on SMIYC [3] based on DeepLabv3+ architecture. The results show that the previous SOTA methods [4, 10, 35] do not produce consistent performance in various contexts, where PEBAL [35] performs well in Fishyscapes [2], but fail in the SMIYC-Anomaly and SMIYC-Obstacle from [3]. In contrast, Meta-OoD [4] yields stable results in SMIYC [3], while its performance degrades in Fishyscapes [2]. Our results are stable across all validation sets, achieving SOTA performance for all the measurements. Notably, we surpass the previous SOTA PEBAL [35] with 0.67%, 2.24% FPR and 0.38%, 11.8% AuPRC in FS-Static and FS-L&F, respectively. Meanwhile, we outperform Meta-OoD [4] with 10.25%, 0.32% FPR and 8.42%, 2.77% AuPRC in SMIYC [3]. Tab. 2 compares our results with SOTA methods on RoadAnomaly, which is one of the most challenging datasets. Our approach shows the best performance with improvements of 10.54%, 9.24%, in FPR and AuPRC comparing with PEBAL [35].

To further show the efficacy of our approach, we test it on the official black-box **test sets** in Tab. 3, which shows the results on Fishyscapes¹ and SMIYC². We achieve the best performance by a large margin on the Fishyscapes-Static compared with the previous SOTAs. Specifically, our results is 1.21% FPR, 3.58% AuPRC better than PEBAL, and 4.99% FPR, 23.69% AuPRC better than Densehybrid in FS-Static. Our approach also yields around 4% to 6% FPR and 10% AuPRC improvements in FS-L&F, demon-

¹<https://fishyscapes.com/results>

²<https://segmentmeifyoucan.com/leaderboard>



(a) Hinge Energy Loss [27] (b) Positive Energy Loss (ours)

Figure 5: **Energy distribution** of hinge energy loss (a) and positive energy loss (b) on SMIYC (obstacle) validation set.

strating its effectiveness in detecting unseen anomalies in urban driving situations. However, our approach shows a slightly worse (1.6%) AuPRC, compared with GMM-Seg [23], which can be explained by its external GMM. Otherwise, our results are substantially better than GMM-Seg, particularly on RoadAnomaly (see Tab. 2). Our result also reaches SOTA on the SMIYC test set, except for the slightly worse AuPRC (in SMIYC-Anomaly) compared with Meta-OoD [4], which failed in FS-L&F with 35.14% FPR and 29.96% AuPRC. Also, our approach’s lowest FPR in SMIYC shows its effectiveness. Considering the overall average FPR and AP results, our approach outperforms previous SOTA methods [4, 35] by a significant margin.

The experiments based on fine-tuning with the extra training sets (including Vistas [32], Wilddash2 [43]) are shown in the Supplementary Sec.B, where our method shows better results than the SOTAs under same setup.

The performance of the inlier categories is another essential measurement of a model’s efficacy, so we show the **in-distribution segmentation accuracy** on Cityscapes in Fig. 4, following [9, 10, 17]. The results based on **single-scale sliding evaluation** and show that PEBAL [35] decreases 1.27% mIoU, and Meta-OoD [4] decreases 0.95% mIoU, while our approach and other OoD methods that freeze networks (i.e., Synboost [9]) maintain the same accuracy for inlier categories.

Table 4: **Ablation study for our approach.** The datasets Static, L&F are from Fishyscapes validation set and Anomaly, Obstacle are from SMIYC. The first two rows show baseline results from the hinge energy loss [35] and the entropy maximisation [4]. PE denotes the positive energy loss from (6), DS denotes the dis-similarity regularisation ℓ_{reg} in (5). CoroCL is for the context-robust contrastive learning, and we directly optimise the RPL output $f_{\theta_{\text{rpl}}}$ (instead of (3)) in the bottom row.

Entropy [4]	Energy [35]	PE	DS	CoroCL	Static		L&F		Anomaly		Obstacle		RoadAnomaly	
					FPR	AuPRC	FPR	AuPRC	FPR	AuPRC	FPR	AuPRC	FPR	AuPRC
✓					1.78	86.88	5.04	52.10	27.59	60.76	2.50	80.85	28.63	49.22
	✓				1.65	89.08	6.64	51.47	28.23	70.18	1.57	71.40	34.39	52.66
		✓			1.55	88.05	4.52	56.85	26.58	73.66	0.51	92.36	31.96	57.28
		✓	✓		1.30	91.16	3.79	63.72	25.65	76.43	0.42	93.25	30.66	63.02
		✓	✓	✓	0.85	92.46	2.52	70.61	7.18	88.55	0.09	96.91	17.74	71.61
		✓			6.92	54.31	17.18	32.57	29.69	68.99	0.95	83.08	33.72	57.99

Table 5: FP detection in Cityscape validation set based on threshold $t = 0.0$. The results are averaged over the images, each with $\approx 1.8 \times 10^6$ pixels (excluding the *ignore* labels).

RPL($\alpha = .05$)+CoroCL	RPL($\alpha = .05$)	RPL ($\alpha = .10$)	RPL ($\alpha = .15$)
FP=1949.306	FP=4360.534	FP=10784.576	FP=19837.211

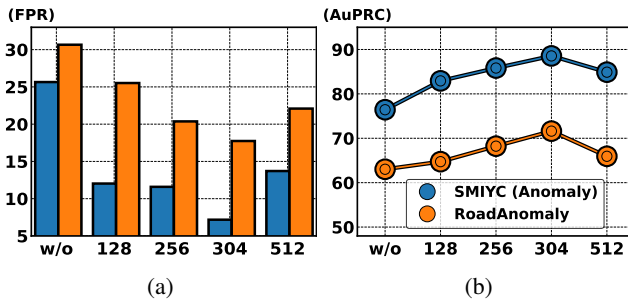


Figure 6: **The effect of the embedding depth** on the SMIYC-Anomaly and RoadAnomaly datasets, where x-axis is the depth number. Note that 304 is the output of the RPL from $f_{\theta_{\text{rpl}}}$ in (3).

4.2. Ablation Studies

Tab. 4 shows the ablation study of the RPL module, where Entropy [4] is the entropy maximisation (first row), and Energy [35] denotes the hinge energy loss (second row). In RPL optimisation, we test the positive energy (PE) loss from (6), the dis-similarity (DS) regularisation in (5), and the Context-robust Contrastive Learning (CoroCL) in (7). In general, the entropy-based loss [4] and the hinge-based energy loss [35] yield unsatisfactory performance for all the validation sets, where the former provides low AuPRC and the latter fails in small obstacle situations (SMIYC-obstacle). Our proposed PE loss improves 0.48% FPR, 4.75% AuPRC and 1.99% FPR, 11.51% AuPRC in FS-L&F and SMIYC-Obstacle by comparing it with Entropy optimisation [4]. After adding the DS regularisation, the results improve for all validation sets, where AuPRC has large improvement, for example, it increases 3.11%, 6.87% AuPRC in FS-Static and FS-L&F, respectively. In the row before the last, CoroCL shows 18.47% FPR, 12.12% AuPRC improvements in SMIYC-Anomaly and 12.92% FPR, 8.59%

Table 6: **Our RPL brings improvements** to previous SO-TAs in Fishyscapes and SMIYC validation sets.

Methods	Static		L&F		Anomaly		Obstacle	
	FPR	AuPRC	FPR	AuPRC	FPR	AuPRC	FPR	AuPRC
Meta-Ood [4]	13.57	72.91	37.69	41.31	17.43	80.13	0.41	94.14
RPL+Meta-Ood	2.34	87.98	7.52	55.25	12.91	82.08	0.24	94.47
PEBAL [35]	1.52	92.08	4.76	58.81	36.74	53.10	7.92	10.15
RPL+PEBAL	0.78	93.28	3.74	63.99	17.89	78.53	3.12	84.23

AuPRC in RoadAnomaly. Lastly, we show the results of the direct optimisation of RPL as an independent classifier (i.e., without adding the RPL output back to the intermediate features of the segmentation model). Compared with direct optimisation (the bottom row), the result with PE (the 3rd row) shows that our RPL can more effectively detect potential anomalies (e.g., improves 12.66% FPR in L&F).

The RPL module is optimised with the positive energy loss from (5), which we claim to be more robust to imbalanced distribution of inliers and outliers than the hinge loss [27], resulting in better Ood pixel sensitivity detection. This is shown in Fig. 5 that compares the energy distribution obtained from the hinge loss (a), and our positive energy loss (b) using our RPL module on the SMIYC-obstacle dataset, where the positive energy loss yields better separation and clustering of energy distribution.

Tab. 4 shows that CoroCL provides more improvements on SMIYC-Anomaly and RoadAnomaly, so we study it further. Recall that CoroCL works with RPL’s intermediate features, so it is important to investigate the impact of embedding depth and projector architecture. Fig. 6 shows that the best performance is observed when the embedding depth is **304**, which is better than 512 by around 5% in AuPRC for both datasets. This result enables us to conclude that when the output dimension of the projector $f_{\theta_{\text{proj}}}(\cdot)$ equals that of the RPL block $f_{\theta_{\text{rpl}}}(\cdot)$, we have a more effective optimisation. Tab. 5 shows the pixel-wise FP anomaly detection with energy threshold $t = 0.0$, where we can note that more FPs occur if we set larger α during training, as the model will be more sensitive to the outliers, and CoroCL successfully reduces the FPs for the inlier scenes.

Our RPL module can also be deployed to SOTA methods to improve their performance. Tab. 6 shows improve-

ments brought by RPL for PEBAL [35] and Meta-OoD [4] on Fishycapes and SMIYC datasets. For instance, RPL improves Meta-OoD [4] by around 15% AuPRC on Static and L&F in Fishycapes and by over 70% AuPRC for PEBAL [35] on SMIYC-Obstacle.

5. Conclusion and Discussion

In this paper, we introduced the residual pattern learning (RPL) block that induces the closed-set segmentation model to become highly uncertain for potential anomalous regions. Compared with recent re-training approaches [4, 10, 35], our RPL detects potential anomalies more effectively, while causing minimal impact to the closed-set segmentation accuracy [9, 13, 17]. Our positive energy loss enables the detection of small anomalies with a more effective optimisation than the one used in previous hinge-loss energy optimisation [35]. Furthermore, our proposed CoroCL is shown to be robust to context changes between training and testing images, avoiding the massive mis-detections observed in previous anomaly detectors [4, 9, 10, 17, 35]. A limitation of our approach is that the proposed optimisation tries to approximate the closed-set segmentation results, without further reducing its inlier entropy. Future work will focus on how to reduce such inlier uncertainty without affecting the sensitivity of the OoD pixels.

6. Acknowledgement

We acknowledge the support provided by the Australian Research Council through grant FT190100525, which has been instrumental in the completion of this work.

References

- [1] Petra Bevandić, Ivan Krešo, Marin Oršić, and Siniša Šegvić. Simultaneous semantic segmentation and outlier detection in presence of domain shift. In *Pattern Recognition: 41st DAGM German Conference, DAGM GCPR 2019, Dortmund, Germany, September 10–13, 2019, Proceedings 41*, pages 33–47. Springer, 2019.
- [2] Hermann Blum, Paul-Edouard Sarlin, Juan Nieto, Roland Siegwart, and Cesar Cadena. The fishyscapes benchmark: Measuring blind spots in semantic segmentation. *International Journal of Computer Vision*, 129(11):3119–3135, 2021.
- [3] Robin Chan, Krzysztof Lis, Svenja Uhlemeyer, Hermann Blum, Sina Honari, Roland Siegwart, Mathieu Salzmann, Pascal Fua, and Matthias Rottmann. Segmentmeifyoucan: A benchmark for anomaly segmentation. *arXiv preprint arXiv:2104.14812*, 2021.
- [4] Robin Chan, Matthias Rottmann, and Hanno Gottschalk. Entropy maximization and meta classification for out-of-distribution detection in semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5128–5137, 2021.
- [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 transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [6] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [9] Giancarlo Di Biase, Hermann Blum, Roland Siegwart, and Cesar Cadena. Pixel-wise anomaly detection in complex driving scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16918–16927, 2021.
- [10] Matej Grcić, Petra Bevandić, and Siniša Šegvić. Densehybrid: Hybrid anomaly detection for dense open-set recognition. *arXiv preprint arXiv:2207.02606*, 2022.
- [11] Shijie Hao, Yuan Zhou, and Yanrong Guo. A brief survey on semantic segmentation with deep learning. *Neurocomputing*, 406:302–321, 2020.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [13] Dan Hendrycks, Steven Basart, Mantas Mazeika, Mohammadreza Mostajabi, Jacob Steinhardt, and Dawn Song. Scaling out-of-distribution detection for real-world settings. *arXiv preprint arXiv:1911.11132*, 2019.
- [14] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- [15] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018.
- [16] Jie Hong, Weihao Li, Junlin Han, Jiyang Zheng, Pengfei Fang, Mehrtash Harandi, and Lars Petersson. Goss: Towards generalized open-set semantic segmentation. *arXiv preprint arXiv:2203.12116*, 2022.
- [17] Sanghun Jung, Jungsoo Lee, Daehoon Gwak, Sungha Choi, and Jaegul Choo. 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*, pages 15425–15434, 2021.
- [18] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.

- [19] Ivan Krešo, Josip Krapac, and Siniša Šegvić. Efficient ladder-style densenets for semantic segmentation of large images. *IEEE Transactions on Intelligent Transportation Systems*, 22(8):4951–4961, 2020.
- [20] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- [21] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018.
- [22] Tianhong Li, Peng Cao, Yuan Yuan, Lijie Fan, Yuzhe Yang, Rogerio S Feris, Piotr Indyk, and Dina Katabi. Targeted supervised contrastive learning for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6928, 2022.
- [23] Chen Liang, Wenguan Wang, Jiaxu Miao, and Yi Yang. Gmmseg: Gaussian mixture based generative semantic segmentation models. *arXiv preprint arXiv:2210.02025*, 2022.
- [24] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [26] Krzysztof Lis, Krishna Nakka, Pascal Fua, and Mathieu Salzmann. Detecting the unexpected via image resynthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2152–2161, 2019.
- [27] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems*, 2020.
- [28] Ziyin Liu, Zhikang Wang, Paul Pu Liang, Russ R Salakhutdinov, Louis-Philippe Morency, and Masahito Ueda. Deep gamblers: Learning to abstain with portfolio theory. *Advances in Neural Information Processing Systems*, 32, 2019.
- [29] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [30] Jishnu Mukhoti and Yarin Gal. Evaluating bayesian deep learning methods for semantic segmentation. *arXiv preprint arXiv:1811.12709*, 2018.
- [31] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [32] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *Proceedings of the IEEE international conference on computer vision*, pages 4990–4999, 2017.
- [33] Peter Pinggera, Sebastian Ramos, Stefan Gehrig, Uwe Franke, Carsten Rother, and Rudolf Mester. Lost and found: detecting small road hazards for self-driving vehicles. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1099–1106. IEEE, 2016.
- [34] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [35] Yu Tian, Yuyuan Liu, Guansong Pang, Fengbei Liu, Yuanhong Chen, and Gustavo Carneiro. Pixel-wise energy-biased abstention learning for anomaly segmentation on complex urban driving scenes. *arXiv preprint arXiv:2111.12264*, 2021.
- [36] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020.
- [37] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [38] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring cross-image pixel contrast for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7303–7313, 2021.
- [39] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3024–3033, 2021.
- [40] Yingda Xia, Yi Zhang, Fengze Liu, Wei Shen, and Alan Yuille. Synthesize then compare: Detecting failures and anomalies for semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [41] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *European conference on computer vision*, pages 173–190. Springer, 2020.
- [42] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE access*, 8:58443–58469, 2020.
- [43] Oliver Zendel, Katrin Honauer, Markus Murschitz, Daniel Steinger, and Gustavo Fernandez Dominguez. Wilddash-creating hazard-aware benchmarks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 402–416, 2018.
- [44] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [45] Yi Zhu, Karan Sapra, Fitsum A Reda, Kevin J Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relax-

ation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8856–8865, 2019.