Random Walk on Pixel Manifolds for Anomaly Segmentation of Complex Driving Scenes

Zelong Zeng¹[®] and Kaname Tomite¹

SenseTime Japan, Tokyo, Japan {zengzelong,tomite}@sensetime.jp

Abstract. In anomaly segmentation for complex driving scenes, stateof-the-art approaches utilize anomaly scoring functions to calculate anomaly scores. For these functions, accurately predicting the logits of inlier classes for each pixel is crucial for precisely inferring the anomaly score. However, in real-world driving scenarios, the diversity of scenes often results in distorted manifolds of pixel embeddings in the space. This effect is not conducive to directly using the pixel embeddings for the logit prediction during inference, a concern overlooked by existing methods. To address this problem, we propose a novel method called Random Walk on Pixel Manifolds (RWPM). RWPM utilizes random walks to reveal the intrinsic relationships among pixels to refine the pixel embeddings. The refined pixel embeddings alleviate the distortion of manifolds, improving the accuracy of anomaly scores. Our extensive experiments show that RWPM consistently improve the performance of the existing anomaly segmentation methods and achieve the best results ¹.

Keywords: Autonomous vehicles \cdot Anomaly segmentation \cdot Semantic segmentation

1 Introduction

Semantic segmentation [9-11,13,39] plays a crucial role in the field of automated driving, which can aid the automated driving system in perceiving and understanding the surrounding environment, thereby facilitating the system make accurate decision. Typically, the semantic segmentation models are trained on closed-datasets with fixed semantic categories, showing good recognition capabilities for inlier classes (*e.g.*, roads, vehicles, etc.). However, these models often fail to effectively detect potential outlier class objects (*e.g.*, animals, fallen tires, etc.) in real-world scenarios. Due to the critical importance of correctly detecting outlier class targets for the safety of autonomous driving, there is a growing focus on addressing this task, namely, anomaly segmentation.

The standard paradigm for anomaly segmentation is assigning an anomaly score to each pixel of an input image, where a higher score indicates a greater probability that the pixel belongs to an outlier. Thus, the way used for estimating the anomaly score is crucial for anomaly segmentation methods. In

¹ Code is available at: https://github.com/ZelongZeng/RWPM





Fig. 1: A toy example about manifolds of pixel embeddings. The dashed circles indicate regions with high predicted logit to the corresponding inlier class. For the pixel in these regions, its anomaly score is small. The prototype is the vector of network weights of the classifier for the corresponding inlier class. (i) An ideal manifold structure. All inlier pixels are within the corresponding high logit region, whereas all outlier pixels are outside the areas. The anomaly score derived from the anomaly scoring function effectively discriminates between them. (ii) The manifolds are affected by the diversity of data, causing some inlier pixels to deviate from the corresponding high logit region, while some outlier pixels approach the regions. This results in false positives/negatives in the anomaly scores. However, pixels of the same class are still in the same manifold, indicating that the manifold structure can be utilized to reveal the intrinsic relationships between pixels. (iii) Our RWPM utilizes random walks to capture the manifold structure to diffuse and update pixel embeddings. The embeddings within the same manifold tend to become more similar after the updating. Fig 3 and Fig 4 respectively present the visualization of embedding distributions and the qualitative results of real examples, consistent with the description of this toy example.

early approaches, the anomaly score is often given by estimating the prediction uncertainty [16, 21, 27, 29], or by comparing the discrepancy between the original image and its reconstructed image generated based on the semantic segmentation predictions [14, 15, 31, 35]. Different from early methods, recent methods [12, 19, 32, 34, 36, 37] adopt Outlier Exposure (OE) strategy [22], which directly learn a pixel-level anomaly class to detect anomalous pixels that deviate from inlier classes. Specifically, these methods propose anomaly scoring functions to infer the anomaly scores and, based on the functions, leverage auxiliary outlier data to fine-tune a semantic segmentation model. The auxiliary outlier data is typically generated by extracting outlier-class objects from other datasets (e.g., COCO [30]) and incorporating them into driving scene images [34, 36, 37]. Furthermore, some of them [19, 34, 36] adopt mask-based segmentation networks instead of pixel-based segmentation networks, achieving state-of-the-art (SOTA) performance in anomaly segmentation. In this paper, we concentrate on the current methods using OE strategy. These methods indicate superior performance to the early methods using prediction uncertainty or image reconstruction.

The anomaly scoring functions in OE methods infer anomaly scores based on the inlier class logit predictions from a semantic segmentation model. Specifically, the logic of these functions can be represented as follows: when the predicted logits of a pixel embedding are small across all inlier classes, indicating that the pixel is dissimilar to all inlier classes and is an outlier with high probability, the function yields a higher anomaly score. Therefore, accurately measuring the logits of each pixel is crucial for precisely evaluating the anomaly score. However, in practical open-world driving scenarios, various environmental conditions (such as lighting, road materials, etc.) and a diverse range of unknown class objects are present. During the inference phase, the diversity of driving scenes often results in distorted manifolds of pixel embedding in the embedding space, causing pixel embeddings to deviate from the ideal distribution. This phenomenon has been validated in previous works of metric learning [3, 25, 40, 41]. Since the logits are calculated from the inner product (*i.e.*, cosine distance) of the pixel embeddings and the vectors of classifiers' network weights, the distorted manifolds are not conducive to directly using the pixel embeddings for the logit prediction. However, current anomaly segmentation methods ignore this issue and directly use the pixel embeddings for the logit prediction in the inference phase, so the results often fail to accurately reflect the anomaly score for the corresponding pixels, causing high false positive/negative rate (see Fig. 1(i) and Fig. 1(ii)).

In this paper, we propose a novel method, called Random Walk on Pixel Manifolds (RWPM), to alleviate the effect of data diversity on the manifolds. RWPM utilizes random walks to capture the manifold structure of the pixel embeddings to diffuse and update the embedding of each pixel. Since the structure of data manifolds implicitly contains the intrinsic relationship among data points (see Fig. 1(ii)) [3,24,40,43], using random walks to measure the similarity between pixels on manifolds (*i.e.*, distance on manifolds rather than Euclidean distance) better reflects the intrinsic relationships between pixels. Consequently, diffusing and updating pixel embeddings based on the similarity on manifolds results in high similarity among pixels on the same manifold and low similarity among pixels on different manifolds. In the embedding space, this manifests as pixel embeddings on the same manifold forming more compact clusters, thereby mitigating the distortion effect of data diversity on the manifolds, which facilitates the logit prediction (see Fig. 1(iii)). Furthermore, since the typically large number of pixels in an image, utilizing all pixels for the random walks results in significant memory and computational overhead, which damages the deployment and real-time execution of our RWPM. To address this issue, we introduce the Partial Random Walk strategy for RWPM, reducing memory consumption and enhancing the operational efficiency of RWPM. Note that, RWPM is proposed to use in the inference phase, thus no any additional training is required. Additionally, RWPM can be directly integrated into existing anomaly segmentation framework without the need for any extra modification to the network structure. Extensive experiments in various anomaly segmentation benchmarks of road scenes, namely Fishyscapes [5], Road Anomaly [31], and SMIYC [6], demonstrate that RWPM consistently improves the performance of existing methods, achieving SOTA results. To summarise, our contributions are the following:

 We introduce that the diversity of driving scenes results in distorted manifolds of pixel embedding, thereby affecting the accuracy of anomaly scoring

functions in inferring anomaly scores. This problem is ignored by existing anomaly segmentation methods.

- We propose the Random Walk on Pixel Manifolds (RWPM) that utilizes random walks to reveal the intrinsic relationships among pixels to refine the pixel embeddings. This way alleviate the effect on manifolds caused by the driving scene diversity.
- We propose the Partial Random Walk strategy for RWPM to reduce the memory consumption and improve the operational efficiency.

2 Related Work

Anomaly Segmentation. Anomaly segmentation aims to detect the outlier class object in some specific scenes, such as complex driving scenes [6, 13, 31]. Existing anomaly segmentation can be broadly divided into three categories: (a) Uncertainty-based methods, (b) Reconstruction-based methods and (c) Out*lier Exposure methods.* Uncertainty-based methods assume that outlier samples result in low-confidence predictions. Based on this assumption, they focus on measuring the pixel-wise anomaly score by estimating the prediction uncertainty through various ways, such as ensembles [27], Bayesian estimation [16], maximum softmax probability [21, 29], and logit [20, 26]. Since segmentation models are trained on close-set, they may misclassify unseen classes with high confidence, leading uncertainty-based methods fail to detect anomalies. Reconstructionbased methods detect anomaly pixel by comparing the discrepancy between the original image and its reconstructed image generated based on the semantic segmentation predictions [14, 15, 31, 35]. These methods are challenged by the accuracy of segmentation prediction and the difficulties of training. Recent methods [7, 12, 19, 34, 36, 37] introduce the Outlier Exposure (OE) strategy [22] that utilizes auxiliary outlier data and anomaly scoring functions to learn models. The auxiliary outlier data is the image from other datasets (e.q., COCO [30]), or cutting the outlier objects from other datasets and pasting them into the inlier scenes, and the anomaly scoring functions are used to calculate the anomaly score for each pixel. Compared with previous works, OE strategy can reach better performance. More recently, some OE methods [19,34,36] adopt mask-based segmentation networks [10, 11], instead of pixel-based segmentation networks [8, 9], as their backbones. Based on the powerful capability of mask-based networks in segmentation tasks, these methods significantly reduces false detection rates, achieving SOTA results. In this paper, we only concentrate on the OE methods. The anomaly scoring functions employed in OE methods are based on the predicted logits of inlier classes. However, during the inference phase, the diversity of driving scenes results in distorted manifolds of pixel embeddings. As a result, the scoring functions may fail to use the embeddings to accurately calculate logits, thereby affects the accuracy of measuring anomaly scores. Current OE methods ignore this problem. Our proposed RWPM addresses this problem by utilizing random walks on pixel manifolds to reveal the intrinsic relationships among pixels and refine the pixel embeddings. To the best of our knowledge, RWPM is the first work to propose and resolve this problem.

5

Mining on Data Manifolds. Mining on data manifolds, also known as diffusion process, is a well-established technique commonly used in retrieval tasks [3, 4, 25, 40, 43]. These works utilize random walks on data manifolds to capture the data manifold structure for revealing the intrinsic relationship between data points. The data manifold is constructed by a graph, where each node represents a sample and each edge presents the connection between a node and its neighborhoods, weighted proportionally according to their distance. For retrieval tasks, the images of the same object may be shown in different conditions (e.q., lightingand angles). This diversity often leads to the manifold structure is not conducive to accurately ranking using distance metrics directly [25,40]. Thus, diffusion process is used to diffuse the state (typically initialized to 1) of a query image on the data manifolds. If a sample is on the same data manifold with the query (the distance on manifolds is close), then more energy of the state will flow to that sample. Therefore, the final state of all samples can be seen as the accurate ranking scores. Moreover, diffusion process is also adopted in semi-supervised learning [24, 42] to propagate the label of labeled samples to unlabeled samples. Similarly, some unsupervised segmentation works [17, 23] utilize diffusion process to propagate the labels of manually annotated pixels to the unlabeled pixels. These existing works utilize random walks on data manifolds for ranking or label propagation. Different from these works, our target is to solve the problem occurring in pixel manifolds for anomaly scoring functions. We propose innovatively using random walks to diffuse and update pixel embeddings, eliminating the effect of distorted manifolds directly at the embedding-level, thereby improving the accuracy of subsequent processing.

3 Proposed Method

3.1 Preliminaries

We denote an image space $\mathcal{X} \subset \mathbb{R}^{3 \times H \times W}$, where H and W are the height and width, respectively. During the inference phase, the goal for an anomaly segmentation model f is to map an input image $\boldsymbol{x} \in \mathcal{X}$ to an anomaly space, *i.e.*, $f : \boldsymbol{x} \mapsto \mathcal{Z} \subset \mathbb{R}^{H \times W}$. Specifically, an anomaly segmentation model typically consists of the backbone encoder-decoder F and the classifier C. For an input image \boldsymbol{x} , firstly, its pixel embedding map is encoded as $\boldsymbol{p} = F(\boldsymbol{x}) \in \mathbb{R}^{d \times H \times W}$, where d is the dimension of the pixel embedding. Then, its logit map \boldsymbol{l} is predicted by $\boldsymbol{l} = C(F(\boldsymbol{x})) \in \mathbb{R}^{K \times H \times W}$, where K is the number of inlier classes. Note that, depending on the backbone network, the size of the embedding map and logit map may sometimes be smaller than the original image. For the sake of convenience in this paper, we do not make a special distinction. Finally, a scoring function is used to calculate the anomaly score map $\boldsymbol{z} \in \mathcal{Z}$ based on the logit map \boldsymbol{l} .

For anomaly scoring functions, their logic can be represented as follows: when the predicted logits of a pixel embedding are small across all inlier classes, indicating that the pixel is dissimilar to all inlier classes and is likely an outlier, the function generates a higher anomaly score. For example, let's define $l_{h,w}(k)$



Fig. 2: Overview. This figure illustrates the application of RWPM to an existing anomaly segmentation framework during the inference phase. The red dashed box highlights the RWPM part. First, the encoder-decoder is used to extract the pixel embedding map of the input image. The pixel embedding map are subsequently partitioned into n^2 sub-maps. For each sub-map, we update its pixel embeddings by using random walks to obtain a refined sub-map. Finally, the refined sub-maps are concatenated to form the refined embedding map, which are then input into the subsequent network structure. Notably, RWPM can be directly integrated into existing frameworks without requiring extra training or changes to the network structure.

is the logit of a pixel $\boldsymbol{x}_{h,w}$ for class k. The scoring function in PEBAL [37] is $E_{h,w} = -\log \sum_{k \in \{1...K\}} \exp(\boldsymbol{l}_{h,w}(k))$. Obviously, when the anomaly score $E_{h,w}$ is large, it indicates that the logits of the pixel across all inlier classes are small, *i.e.*, the pixel $\boldsymbol{x}_{h,w}$ is dissimilar from all inlier classes (see Section S.2 of the supplementary material for the analysis of other anomaly scoring functions).

3.2 Graph Construction

We begin by constructing a graph to represent the manifolds of pixel embeddings. In our approach, we utilize an affinity matrix based on cosine similarity to characterize the relationships between pixels. Specifically, we first use the encoder-decoder component of an anomaly segmentation model to extract the pixel embeddings $\boldsymbol{p} \in \mathbb{R}^{d \times H \times W}$. For ease of description and computation, we reshape the \boldsymbol{p} as $\boldsymbol{p}^r \in \mathbb{R}^{HW \times d}$, where each row represents a pixel's embedding. We then normalize each row of \boldsymbol{p}^r to obtain $\hat{\boldsymbol{p}}^r$, *i.e.*, $\hat{\boldsymbol{p}}^r_i = \boldsymbol{p}^r_i / || \boldsymbol{p}^r_i ||$. The affinity matrix is define as $\mathbf{W} \in \mathbb{R}^{HW \times HW}$, where each component is obtained by:

$$\mathbf{W}_{ij} = \begin{cases} \left\langle \widehat{\boldsymbol{p}}_i^r, \widehat{\boldsymbol{p}}_j^r \right\rangle & i \neq j \\ 0 & i = j \end{cases}, \tag{1}$$

where $\langle \cdot, \cdot \rangle$ denotes inner product, $\mathbf{W}_{ii} = 0$ is guaranteed to circumvent self-loops in the graph represented by \mathbf{W} .

To mitigate the influence of noise, locally constrained random walk methods [24, 40, 41] adopt local constraints to construct the graph, *i.e.*, each point is only connected to its locally nearest neighbors in the graph. Determining the nearest neighbors for each point traditionally involves applying nearest neighbor search algorithms, such as the k-NN algorithm. However, in anomaly segmentation tasks, the size of the graph is large and the graph is constructed online, making the search algorithms impractical. To benefit from the local constraints, we utilize a softmax function to normalize each row of \mathbf{W} to obtain locally constrained graph \mathbf{S} :

$$\mathbf{S}_{ij} = \begin{cases} \frac{\exp(\mathbf{W}_{ij}/\tau)}{\sum_{i=1, i\neq j}^{HW} \exp(\mathbf{W}_{ij}/\tau)} & i \neq j\\ 0 & i = j \end{cases},$$
(2)

where τ is a temperature hyper-parameter, we set its value less than 1.0 to remove the effect of non-neighbor points. Compared to typical search algorithms, Eq. 2 can impose local constraints more efficiently and effectively (see Section S.4 of the supplementary material), and keep the scale of embedding in subsequent random walks.

3.3 Random Walk Process

Now we have constructed the graph \mathbf{S} , we can start the random walks. Different from previous mining on data manifolds works utilizing random walks to propagate the states of known samples to other unknown samples, we use a random walk to propagate the embedding between pixels on manifolds.

We first define $\mathbf{m}^t \in \mathbb{R}^{HW \times d}$ as the updated pixel embeddings in the *t*-th step of the random walks. \mathbf{m}^0 is initialized as p^r . Then, each iteration step of the random walks is described as:

$$\mathbf{m}^{t+1} = \alpha \mathbf{S} \mathbf{m}^t + (1 - \alpha) \mathbf{m}^0, \, \alpha \in (0, 1), \tag{3}$$

where α is the probability continuing the random walk from the current state \mathbf{m}^t , while $(1 - \alpha)$ represents the probability of restarting from the initial state \mathbf{m}^0 . At each iteration, guided by the manifold structure, each pixel diffuses its embedding to other pixels while also receiving embedding from other pixels for updating its own embedding. The updated pixel embedding can be seen as a weighted ensemble of other pixel embeddings. If pixels are similar on manifolds, their embeddings will gradually become similar. Mathematically speaking, the whole iteration process can be written as a closed-from as follows [43]:

$$\mathbf{m}^{\infty} = (1 - \alpha) \left(\mathbf{I} - \alpha \mathbf{S} \right)^{-1} \mathbf{m}^{0}, \tag{4}$$

where $\mathbf{I} \in \mathbb{R}^{HW \times HW}$ is an identity matrix. Finally, \mathbf{m}^{∞} is used as the refined pixel embeddings for the subsequent anomaly segmentation process.

3.4 Partial Random Walk

When the size of the image is large, the size of the relevant matrices used for the calculation of random walks becomes extremely huge as well. For instance, when the image size is (512×1024) , the size of matrix S reaches (524288×524288) . This significantly increases the memory consumption and computational time of RWPM, which is not conducive to online operation. Unfortunately, in the field

of autonomous driving, the image size of driving scenes is generally large. To address this problem, we propose Partial Random Walk, which consists of the following two aspects: *embedding map partitioning* and *limited iteration*.

To improve efficiency, we first propose embedding map partitioning strategy to reduce the size of the pixel embedding map used for random walks. Specifically, for a pixel embedding map $\boldsymbol{p} \in \mathbb{R}^{d \times H \times W}$, we equally divide it into $n \times n$ sub-maps $\left\{ \boldsymbol{p}^{(1)}, \cdots, \boldsymbol{p}^{(n^2)} \right\} \in \mathbb{R}^{d \times \frac{H}{n} \times \frac{W}{n}}$. Then, for each sub-map in $\left\{ \boldsymbol{p}^{(1)}, \cdots, \boldsymbol{p}^{(n^2)} \right\}$, we follow the steps of Eq. 1 and Eq. 2 to construct their sub-graphs $\left\{ \boldsymbol{S}_{\boldsymbol{p}^{(1)}}, \cdots, \boldsymbol{S}_{\boldsymbol{p}^{(n^2)}} \right\} \in \mathbb{R}^{d \times \frac{HW}{n^2} \times \frac{HW}{n^2}}$. We also attempted other methods to reduce the number of pixel embeddings, such as resize or superpixel [1]. Experimental results indicate that embedding map partitioning is the optimal solution among them (see Section 4.2).

Next, for each sub-map $p^{(i)}$, we perform random walks with the limited iteration strategy in their corresponding sub-graphs $S_{p^{(i)}}$. That is, we use Eq. 3 for a small number of iterations instead of using the closed-form Eq. 4. We found that a few iterations by using Eq. 3 are enough to yield high quality results. The time complexity of Eq. 3 with T iterations for each sub-map is $O(Td(\frac{HW}{n^2})^2)$, while the time complexity of Eq. 4 is $O((\frac{HW}{n^2})^3) + O(d(\frac{HW}{n^2})^2)$. Since Td is much smaller than $\frac{HW}{n^2}$, the limited iteration strategy can significantly improve the efficiency of RWPM. After the random walks, we obtain the refined pixel embeddings sub-map $\mathbf{m}_{p^{(i)}}^T$ for each sub-map $p^{(i)}$. Finally, we concatenate the refined sub-maps based on their positions to obtain the final refined pixel embedding map. Additionally, When n is large (observed experimentally as n > 2), due to spatial constraints, the baseline values of anomaly scores outputted on each sub-map may vary. Hence, a calibration is performed. For adjacent refined sub-maps $p^{(i)'}$ and $p^{(j)'}$, we compute the average anomaly scores I and J of all pixels on their adjacent edges, respectively. Subsequently, using $p^{(i)'}$ as the reference, we calculate the ratio $\frac{I}{J}$, and then multiply all anomaly scores on sub-map $p^{(j)'}$ by this ratio to achieve the calibration.

Fig. 2 shows the workflow of the proposed RWPM.

4 Experiment

Dataset: We evaluate our approach on three standard benchmark datasets for anomaly segmentation, *i.e.*, Fishyscapes Lost&Found [5], Road Anomaly [31] and Segment Me If You Can (SMIYC) [6]. Fishyscapes Lost&Found has 100 validation images. The domain of this dataset is similar to that of Cityscapes [13] dataset. Road Anomaly dataset contains 60 images of real-world road scenes. Unlike the Fishyscapes Lost&Found dataset, the images in Road Anomaly contain various classes and sizes of anomaly objects, making it more challenging. SMIYC consists of two subsets, Anomaly track and Obstacle track. The Anomaly track focuses on detecting large anomalous objects, while the Obstacle track emphasizes detecting small anomalous objects on the road. Anomaly track contains 10 validation images and 100 online test images, while Obstacle track contains 30 validation images and 327 online test images. The Road Anomaly and SMIYC datasets both exhibit high domain shifts compared to the Cityscapes dataset.

Evaluation Metrics: Following the previous works, we evaluate our method on the standard metrics for anomaly segmentation. On the Fishyscapes Lost&Found and Road Anomaly, we adopt the the area under receiver operating characteristics (AUROC), average precision (AP), and the false positive rate at a true positive rate of 95% (FPR95). On the SMIYC, in addition to reporting per-pixel metrics such as AP and FPR95, we also calculate results for component-level metrics [6] such as the component-wise intersection over union (sIoU), positive predictive value (PPV), and mean F1.

Implementation Details: To demonstrate the effectiveness of RWPM for anomaly segmentation methods, we selected representative methods from both pixel-based and mask-based networks. For pixel-based methods, we chose PE-BAL [37] and Balanced Energy [12]. These methods all utilize DeepLabv3+ [9] with WideResnet38 as the backbone and were trained on the Cityscapes dataset and COCO OE images [37]. For mask-based methods, we selected RbA [34] and Mask2Anomaly [36]. These methods use Mask2Former [10] with Swin-L [33] as the backbone and were also trained on the Cityscapes and COCO OE images. All experiments are conducted on a single NVIDIA RTX A6000 GPU. To ensure efficient implementation on a single GPU for all experiments, we adopt embedding map partitioning and limited iteration in all experiments. Specifically, for pixel-based methods, we set the partitioning parameter n = 4. For mask-based methods, as their embedding map size is smaller, we set n = 2. Additionally, for the random walk, if not specifically emphasized, we set the number of iterations T = 20/T = 5 for the Road Anomaly/the other datasets and the transition probability $\alpha = 0.99$ for all datasets. The τ is set as 0.01. It is important to note that our method is used in the inference phase. In other words, RWPM is directly applied to the pre-trained models of other existing anomaly segmentation methods, thus requiring no additional training.

4.1 Effectiveness of RWPM

First, to validate that RWPM can improve the performance of existing anomaly segmentation methods, we integrate our approach with four recent representative methods, namely PEBAL [37], Balanced Energy [12], Mask2Anomaly [36] and RbA [34]. PEBAL and Balanced Energy employ pixel-based networks as their backbones, whereas Mask2Anomaly and RbA utilize mask-based networks as their backbones. We evaluate the results on both the Fishyscapes Lost&Found validation set and Road Anomaly test set. The experimental results are reported in Table 1. These results demonstrate that our proposed RWPM consistently and significantly enhances the performance of existing anomaly segmentation methods, across different backbone architectures. Additional results and parameters details are shown in Section S.3 of the supplementary material.

 $^{^{2}}$ We carefully reproduce the experiments by using the official code.

Table 1: Comparison with strong and representative anomaly segmentation baselines across different backbone Architectures (Arch). **Bold** denotes the better results between using RWPM and not using RWPM with the same anomaly segmentation baselines. [†] indicates that use the calibration in sub-map concatenation as described in Section 3.4.

$Benchmark \rightarrow$		Fishysca	apes Lost	Found	Ro	ad Anoma	aly
Method ↓	Arch	AuROC↑	$AP\uparrow$	$FPR95\downarrow$	AuROC [↑]	$AP\uparrow$	$FPR95\downarrow$
$\begin{array}{l} \text{PEBAL} [37] \\ \text{PEBAL} + \text{RWPM}^{\dagger} \ (\text{Ours}) \end{array}$	Pixel-based	98.96	58.81	4.77	87.63	45.10	44.58
	Pixel-base	99.20	66.85	3.68	89.48	50.29	36.81
Balanced Energy [12]	Pixel-based	99.03	67.07	2.93	88.31	41.48	41.46
Balanced Energy + $RWPM^{\dagger}$ (Ours)	Pixel-base	99.29	72.95	2.38	90.51	48.26	32.10
Mask2Anomaly ² [36]	Mask-based	95.47	65.27	7.79	96.54	80.04	13.95
Mask2Anomaly + RWPM (Ours)	Mask-base	95.48	65.38	7.33	97.44	80.09	7.45
RbA [34]	Mask-based	98.62	70.81	6.30	97.99	85.42	6.92
RbA + RWPM (Ours)	Mask-base	98.82	71.16	6.12	98.04	87.34	5.27

Table 2: Comparison with state-of-the-art anomaly segmentation methods across different backbone Architectures (Arch). **Bold** and <u>underline</u> denote the best and the second best results, respectively.

$\operatorname{Benchmark} \rightarrow$		Fishys	capes L&F	Road	Anomaly	Anom	aly Track	Obsta	cle Track	Av	erage
Method \downarrow	Arch	$AP\uparrow$	$FPR95\downarrow$	$AP\uparrow$	$\mathrm{FPR95}{\downarrow}$	$AP\uparrow$	$\mathrm{FPR95}{\downarrow}$	$AP\uparrow$	$\mathrm{FPR95}{\downarrow}$	$AP\uparrow$	$\mathrm{FPR95}{\downarrow}$
Synboost [15] (CVPR'21)	Pixel-based	60.58	31.02	41.84	59.72	56.44	61.86	71.34	3.15	57.46	38.94
SML [26] (ICCV'21)	Pixel-based	22.74	33.49	25.82	49.47	46.8	39.5	3.4	36.8	24.69	39.82
Meta-OOD [7] (ICCV'21)	Pixel-based	41.31	37.69	48.84	31.77	85.47	15.00	85.07	0.75	65.13	21.30
Learning Embedding [5](IJCV'21)	Pixel-based	4.65	24.36	-	-	37.52	70.76	0.82	46.38	-	-
Void Classifier [5](IJCV'21)	Pixel-based	10.29	22.11	-	-	36.61	63.49	10.44	41.54	-	-
GMMSeg-DL [28] (NeurIPS'22)	Pixel-based	43.47	13.11	34.42	47.90	-	-	-	-	-	-
DenseHybrid [18] (ECCV'22)	Pixel-based	69.79	5.09	31.39	63.97	77.96	9.81	87.08	<u>0.24</u>	65.56	19.78
PEBAL [37] (ECCV'22)	Pixel-based	58.81	4.77	45.10	44.58	49.14	40.82	4.98	12.68	39.58	25.71
Balanced Energy [12] (CVPR'23)	Pixel-based	67.07	2.93	41.48	41.46	-	-	-	-	-	-
RPL-CoroCL [32] (ICCV'23)	Pixel-based	70.61	2.52	71.61	17.74	83.49	11.68	85.93	0.58	77.91	8.13
Mask2Anomaly [36] (ICCV'23)	Mask-based	69.44	9.22	80.04	13.95	88.62	14.57	<u>93.10</u>	0.20	82.87	9.49
M2F-EAM [19] (CVPRW'23)	Mask-based	52.03	20.51	66.67	13.42	76.3	93.9	66.9	17.9	65.48	36.43
RbA [34] (ICCV'23)	Mask-based	70.81	6.30	85.42	<u>6.92</u>	90.86	11.59	91.85	0.46	84.73	<u>6.33</u>
$\mathbf{RDA} + \mathbf{RWPM}$ (Ours)	wask-based	11.10	0.12	01.34	5.27	92.00	10.15	95.30	0.28	00.00	5.46

Second, to further demonstrate the effectiveness of our RWPM, we compare it with other state-of-the-art (SOTA) methods. Specifically, we adopt the combination of RbA with RWPM as our instance approach for the comparison. In addition to the Fishyscapes and Road Anomaly datasets, we also evaluate our approach on a more challenging dataset, *i.e.*, SMIYC online test set (Anomaly track and Obstacle track). SMIYC is characterized by its high domain shift and diversity of objects. Table 2 presents the results. The results indicate that our method can consistently achieve high performance across various datasets, while the previous SOTA methods fail to do so. We also compare all approaches by computing the average results of AP and FPR95. The results shows that our approach outperforms the previous methods and achieve the best results.

Furthermore, in addition to pixel-level metrics, we also utilize componentlevel metrics [6] to evaluate performance on the SMIYC online test set. Compared to pixel-level metrics, component-level metrics provide a better evaluation of the detection results for all anomalous objects in the scene. We report the results in Table 3. The experimental results show that RWPM significantly improves the performance of RbA in component-level metrics, surpassing other existing SOTA methods by a large margin. This indicates that RWPM can assist models in better detecting anomalous objects and achieving more precise localization of these anomalies [6].

Table 3: Component level evaluation: Comparison with state-of-the-art anomaly segmentation methods on the SMIYC online test set.
 Bold denote the best results.

$\operatorname{Benchmark} \rightarrow$			Anomaly Track			Obstacle Track		
$\mathrm{Method} \downarrow$	Arch	$sIoU \uparrow$	$\mathrm{PPV}\uparrow$	mean F1 \uparrow	sIoU \uparrow	$\mathrm{PPV}\uparrow$	mean F1 \uparrow	
Synboost [15] (CVPR'21)	Pixel-based	34.68	17.81	9.99	44.28	41.75	37.57	
SML [26] (ICCV'21)	Pixel-based	26.00	24.70	12.20	5.10	13.30	3.00	
Meta-OOD [7] (ICCV'21)	Pixel-based	49.31	39.51	28.72	47.87	63.64	48.51	
Void Classifier [5](IJCV'21)	Pixel-based	21.14	22.13	6.49	6.34	20.27	5.41	
Learning Embedding [5](IJCV'21)	Pixel-based	33.86	20.54	7.90	35.64	2.87	2.31	
JSRNet [38] (ICCV'21)	Pixel-based	20.20	29.27	13.66	18.55	24.46	11.02	
DenseHybrid [18] (ECCV'22)	Pixel-based	54.17	24.13	31.08	45.74	50.10	50.72	
PEBAL [37] (ECCV'22)	Pixel-based	38.88	27.20	14.48	29.91	7.55	5.54	
RPL-CoroCL [32] (ICCV'23)	Pixel-based	49.76	29.96	30.16	52.61	56.65	56.69	
Mask2Fromer [36] (ICCV'23)	Mask-based	25.20	18.20	15.30	5.00	21.90	4.80	
Maskomaly [2] (BMVC'23)	Mask-based	55.4	51.6	50.0	-	-	-	
RbA [34] (ICCV'23)	Mask-based	55.69	52.14	46.80	58.36	58.78	60.85	
RbA + RWPM (Ours)	Mask-based	57.00	61.25	58.44	58.89	72.51	69.85	

4.2 Ablation Study

If not specifically emphasized, all the results reported in this section are conducted on a RTX A6000 GPU.

The effect of the partitioning parameter n and calibration: We first test the effect of the embedding map partitioning parameter n and the calibration, and we report the results in Table 4. FPS denotes the frames-per-second results. Row n = 1 displays the RWPM results without utilizing the embedding map partitioning strategy, indicating significant computational and memory overhead. However, the partitioning strategy (n = 2, 4, 8) significantly improves the running speed and reduces memory consumption while maintaining effectiveness. Additionally, as n increases, the running efficiency also improves, albeit with a slight decrease in performance. However, when n > 2, the calibration can enhance the performance, particularly for the FPR95 metrics.

The effect of iteration number T: Table 5 shows that, compared to using Eq. 4 ($T = \infty$), utilizing the limited iteration strategy ($T = 5 \sim 100$) can lead to higher efficiency, resulting in even better performance for RWPM. Additionally, as the number of iterations increases, performance tends to improve, but running efficiency (FPS) decreases accordingly.

Different strategies for reducing the number of pixel embedding: We also attempt other strategies to reduce the number of pixel embedding for RWPM, such as resizing the embedding map or use SLIC [1] algorithm to extract superpixels. Table 6(Left) presents the results. The resize strategy achieves good performance but requires preserving a larger map size. Therefore, it fails

$\operatorname{Benchmark} \rightarrow$	Road Anomaly			
	$\mathbf{AP}\uparrow$	$\mathbf{FPR95}\downarrow$	$\mathbf{FPS}\uparrow$	GPU Memory Usage↓
RbA w/o RWPM	85.42	6.92	11.12	3.49GiB
n=1 (two GPUs)	87.90	5.17	0.32	74.96GiB
n=2 $n=2^{\dagger}$	87.34 87.18	5.27 5.25	2.04 2.04	7.24GiB 7.24GiB
$\substack{n=4\\n=4^{\dagger}}$	87.05 87.17	6.31 5.32	4.35 4.35	3.49GiB 3.49GiB
$\substack{n=8\\n=8^{\dagger}}$	86.67 86.91	6.42 5.41	6.26 6.26	3.55GiB 3.55GiB

Table 4: The effect of the partitioning parameter n and calibration. We set T = 20 for all experiments.[†] indicates that use the calibration in sub-map concatenation.

Table 5: The effect of iteration number T. We set $n = 2$ for all exp	periments.
--	------------

$\operatorname{Benchmark} \rightarrow$	Road Anomaly				
	$\mathbf{AP}\uparrow$	$\mathbf{FPR95} \downarrow$	$\mathbf{FPS}\uparrow$	GPU Memory Usage \downarrow	
RbA w/o RWPM	85.42	6.92	11.12	3.49GiB	
$T=\infty$	87.51	5.24	0.29	8.88GiB	
T=5	86.89	5.98	3.70	7.24GiB	
T=10	87.08	5.77	2.78	7.24GiB	
T=20	87.34	5.27	2.04	7.24GiB	
T=50	87.36	5.08	0.92	7.24GiB	
T=100	87.54	5.02	0.50	7.24GiB	

to substantially reduce the number of pixel embeddings, resulting in significant GPU memory usage. On the other hand, the superpixel strategy requires a considerable amount of runtime for the superpixel extraction. In comparison, our proposed partitioning strategy can achieve good results with higher efficiency and lower resource consumption. This makes it more suitable for applications on practical platforms.

In-distribution segmentation performance on Cityscapes validation set: We report the results on the Cityscapes validation set in Table 6(Right). All methods use Mask2Former [10] as their backbones. We observe that most SOTA methods affect the in-distribution segmentation performance of Mask2Former, such as RbA (drop from 82.25 to 81.93). However, RWPM can alleviate this problem. Additionally, even for the SOTA segmentation method (Mask2Former), RWPM can further enhance its in-distribution segmentation performance.

Table 6: Ablation Tables: (Left) shows the results of different strategies for reducing the number of pixel embedding. We find that partitioning is the best. (Right) shows the segmentation performance on the Cityscapes validation set. RWPM can also improve the segmentation performance of in-distribution.

${\rm Benchmark} {\rightarrow}$	Road Anomaly						
Strategy↓	AP↑	FPR95↓	$\mathbf{FPS}\uparrow$	Memory			
Resize (0.9 full size) Resize (0.8 full size)	87.80 87.24	5.39 5.44	1.41 2.15	23.82GiB 15.98GiB			
Superpixel (30000 pixels) Superpixel (20000 pixels)	87.51 87.03	5.50 5.98	0.19 0.32	11.46GiB 6.31GiB			
Partitioning $(n = 2)$	87.34	5.27	2.04	7.24 GiB			

$\operatorname{Benchmark} \rightarrow$	Cityscapes
Method↓	mIoU↑
PEBAL-Mask [34]	75.32
DenseHybrid-Mask [34]	80.27
RbA [34]	81.93
RbA+RWPM(Ours)	82.16
Mask2Former [10]	82.25
Mask2Former+RWPM(Ours)	82.43

4.3 Visualization of Pixel Embedding Distribution

To investigate the effect of RWPM in the embedding space, we visualized the distribution of pixel embeddings by using t-SNE. We selected an image from the Road Anomaly dataset (as shown in Fig 3(a)). Due to the gravel road surface in this image, there is a significant domain gap between the road in this image and the roads in the training set, resulting in many false positives on the road (as shown in the PEBAL results in Fig 3(a)). Then, we randomly selected 300 inlier pixels and 150 outlier pixels from the red dashed box region of this image for visualization. Fig 3(b) and Fig 3(c) respectively show the distributions of these pixels in the space without and with applying RWPM. In Fig 3(b) and Fig 3(c), deep blue circles represent inlier pixels (road), light blue circles represent outlier pixels (boar), and X symbols represent the inlier class prototype (the vector of network weights of classifier). From Fig 3(b), we observe that without using RWPM, due to the influence of data diversity, the distribution of inliers and outliers is distorted. Consequently, it is difficult for an anomaly scoring function to distinguish between inlier and outlier pixels based on their logits. However, from Fig 3(c), we find that RWPM can significantly alleviate the problem. Specifically, the pixel embeddings of the same category form more compact clusters, enabling the anomaly scoring function to clearly discriminate between inliers and outliers based on the logit. This is consistent with the description in the toy sample in Fig 1. Note that, in the top left corner of Fig 3(c), some inlier points become very close to outlier points after applying RWPM. This is because these points are located in the boundaries of anomalous objects, making it very difficult to discriminate them.



Fig. 3: The visualization of pixel embedding distribution. All embeddings are extracted from PEBAL model. We use cosine distance as the distance metric. We observed that RWPM can optimize the distribution of pixel embeddings in the space.

4.4 Qualitative Result

Finally, to visually demonstrate the effectiveness of RWPM, we present qualitative results. Specifically, we compare the results of RbA on the Road Anomaly dataset without and with RWPM. The qualitative results are illustrated in Fig 4.

Fig 4 shows that RWPM significantly improves the performance of RbA, reducing both the false positive rate and false negative rate. For example, in the first row, RbA generates some false positive results on the road and some false negative results on the manhole region (third column of the first row). However, after applying RWPM, these false positive/negative results are effectively eliminated (fourth column of the first row). Additionally, RWPM assists the RbA model to accurately detect and localize each anomalous object at the component level, such as more precise edge detection. This observation aligns with the results of the component-level metrics presented in Table 3.



Fig. 4: Qualitative results: The first column displays the input image, while the second column shows its corresponding ground truth. The third column displays the anomaly score results generated by RbA, and the fourth column presents the anomaly score results after applying RWPM to RbA. Yellow indicates high anomaly scores.

5 Conclusions and Discussions

In this paper, we propose a simple yet effective method, called Random Walk on Pixel Manifolds (RWPM), to improve the performance of existing anomaly segmentation methods. RWPM is a post-processing method that can be directly integrated into existing anomaly detection frameworks without requiring additional training or modifications to the network structure. RWPM mitigates the impact of data diversity on manifold structure by diffusing and updating pixel embeddings on pixel manifolds using random walk. This process enables anomaly scoring functions to obtain more accurate estimations of anomaly scores. Additionally, we propose Partial Random Walk strategy, which significantly improve the efficiency of RWPM.

While RWPM achieves significant results in most cases, it is limited when the pre-trained model itself exhibits significant inaccuracies. To address this limitation, we plan to enhance the model's generalization ability during training to accurately detect a broader range of diverse and unknown anomalies.

Acknowledgements

This work was supported by SenseTime Japan.

References

- Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: Slic superpixels compared to state-of-the-art superpixel methods. IEEE transactions on pattern analysis and machine intelligence 34(11), 2274–2282 (2012)
- Ackermann, J., Sakaridis, C., Yu, F.: Maskomaly: Zero-shot mask anomaly segmentation. arXiv preprint arXiv:2305.16972 (2023)
- Bai, S., Bai, X., Tian, Q., Latecki, L.J.: Regularized diffusion process for visual retrieval. In: Proceedings of the AAAI conference on artificial intelligence. vol. 31 (2017). https://doi.org/10.1609/aaai.v31i1.11216
- Bai, S., Zhou, Z., Wang, J., Bai, X., Jan Latecki, L., Tian, Q.: Ensemble diffusion for retrieval. In: Proceedings of the IEEE International conference on computer vision. pp. 774–783 (2017)
- Blum, H., Sarlin, P.E., Nieto, J., Siegwart, R., Cadena, C.: The fishyscapes benchmark: Measuring blind spots in semantic segmentation. International Journal of Computer Vision 129, 3119–3135 (2021). https://doi.org/10.1007/s11263-021-01511-6
- Chan, R., Lis, K., Uhlemeyer, S., Blum, H., Honari, S., Siegwart, R., Fua, P., Salzmann, M., Rottmann, M.: Segmentmeifyoucan: A benchmark for anomaly segmentation. arXiv preprint arXiv:2104.14812 (2021)
- 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)
- Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: 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)
- 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). https: //doi.org/10.1007/978-3-030-01234-2_49
- Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1290–1299 (2022). https://doi.org/10.1109/CVPR52688.2022.00135
- Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. Advances in Neural Information Processing Systems 34, 17864–17875 (2021)
- Choi, H., Jeong, H., Choi, J.Y.: Balanced energy regularization loss for out-ofdistribution detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15691–15700 (2023)
- 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). https://doi.org/10.1109/CVPR.2016. 350

- 16 Z. Zeng et al.
- 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)
- 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). https://doi.org/10. 1109/CVPR46437.2021.01664
- Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: international conference on machine learning. pp. 1050–1059. PMLR (2016)
- Grady, L.: Random walks for image segmentation. IEEE transactions on pattern analysis and machine intelligence 28(11), 1768–1783 (2006)
- Grcić, M., Bevandić, P., Šegvić, S.: Densehybrid: Hybrid anomaly detection for dense open-set recognition. In: European Conference on Computer Vision. pp. 500–517. Springer (2022)
- Grcić, M., Šarić, J., Šegvić, S.: On advantages of mask-level recognition for outlieraware segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2936–2946 (2023)
- Hendrycks, D., Basart, S., Mazeika, M., Zou, A., Kwon, J., Mostajabi, M., Steinhardt, J., Song, D.: Scaling out-of-distribution detection for real-world settings. arXiv preprint arXiv:1911.11132 (2019)
- Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-ofdistribution examples in neural networks. In: International Conference on Learning Representations (2017), https://openreview.net/forum?id=Hkg4TI9x1
- 22. Hendrycks, D., Mazeika, M., Dietterich, T.: Deep anomaly detection with outlier exposure. arXiv preprint arXiv:1812.04606 (2018)
- Huang, Z., Li, S.: From difference to similarity: A manifold ranking-based hyperspectral anomaly detection framework. IEEE Transactions on Geoscience and Remote Sensing 57(10), 8118–8130 (2019)
- Iscen, A., Tolias, G., Avrithis, Y., Chum, O.: Mining on manifolds: Metric learning without labels. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7642-7651 (2018). https://doi.org/10.1109/CVPR. 2018.00797
- Iscen, A., Tolias, G., Avrithis, Y., Furon, T., Chum, O.: Efficient diffusion on region manifolds: Recovering small objects with compact cnn representations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2077-2086 (2017). https://doi.org/10.1109/CVPR.2017.105
- 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)
- Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. Advances in neural information processing systems **30** (2017)
- Liang, C., Wang, W., Miao, J., Yang, Y.: Gmmseg: Gaussian mixture based generative semantic segmentation models. Advances in Neural Information Processing Systems 35, 31360–31375 (2022)
- Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. In: International Conference on Learning Representations (2018), https://openreview.net/forum?id=H1VGkIxRZ

- 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: Computer Vision-ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13. pp. 740-755. Springer (2014). https://doi.org/10.1007/978-3-319-10602-1_48
- 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). https://doi.org/10.1109/ICCV.2019.00224
- 32. Liu, Y., Ding, C., Tian, Y., Pang, G., Belagiannis, V., Reid, I., Carneiro, G.: Residual pattern learning for pixel-wise out-of-distribution detection in semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1151–1161 (2023)
- 33. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
- Nayal, N., Yavuz, M., Henriques, J.F., Güney, F.: Rba: Segmenting unknown regions rejected by all. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 711–722 (2023)
- 35. Ohgushi, T., Horiguchi, K., Yamanaka, M.: Road obstacle detection method based on an autoencoder with semantic segmentation. In: Proceedings of the Asian Conference on Computer Vision (2020)
- Rai, S.N., Cermelli, F., Fontanel, D., Masone, C., Caputo, B.: Unmasking anomalies in road-scene segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4037–4046 (2023)
- Tian, Y., Liu, Y., Pang, G., Liu, F., Chen, Y., Carneiro, G.: Pixel-wise energybiased abstention learning for anomaly segmentation on complex urban driving scenes. In: European Conference on Computer Vision. pp. 246–263. Springer (2022). https://doi.org/10.1007/978-3-031-19842-7_15
- 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)
- Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. Advances in Neural Information Processing Systems 34, 12077–12090 (2021)
- Yang, F., Hinami, R., Matsui, Y., Ly, S., Satoh, S.: Efficient image retrieval via decoupling diffusion into online and offline processing. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 9087–9094 (2019). https://doi. org/10.1609/aaai.v33i01.33019087
- Yang, F., Wang, Z., Xiao, J., Satoh, S.: Mining on heterogeneous manifolds for zero-shot cross-modal image retrieval. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 12589–12596 (2020)
- Zhou, D., Bousquet, O., Lal, T., Weston, J., Schölkopf, B.: Learning with local and global consistency. Advances in neural information processing systems 16 (2003)
- Zhou, D., Weston, J., Gretton, A., Bousquet, O., Schölkopf, B.: Ranking on data manifolds. Advances in neural information processing systems 16 (2003)