

Anomaly-Aware Semantic Segmentation via Style-Aligned OoD Augmentation

Dan Zhang^{1,2}, Kaspar Sakmann¹, William Beluch¹, Robin Huttmacher¹, Yumeng Li^{1,3}
¹Bosch Center for Artificial Intelligence ²University of Tübingen ³University of Siegen
 {dan.zhang2, kaspar.sakmann, william.beluch, robin.huttmacher, yumeng.li}@de.bosch.com

Abstract

Within the context of autonomous driving, encountering unknown objects becomes inevitable during deployment in the open world. Therefore, it is crucial to equip standard semantic segmentation models with anomaly awareness. Many previous approaches have utilized synthetic out-of-distribution (OoD) data augmentation to tackle this problem. In this work, we advance the OoD synthesis process by reducing the domain gap between the OoD data and driving scenes, effectively mitigating the style difference that might otherwise act as an obvious shortcut during training. Additionally, we propose a simple fine-tuning loss that effectively induces a pre-trained semantic segmentation model to generate a “none of the given classes” prediction, leveraging per-pixel OoD scores for anomaly segmentation. With minimal fine-tuning effort, our pipeline enables the use of pre-trained models for anomaly segmentation while maintaining the performance on the original task.

1. Introduction

Detecting unknown objects, referred to as Out-of-Distribution (OoD) objects, is vital in safety-critical applications such as autonomous driving. Anomaly segmentation, combined with standard semantic segmentation, offers precise per-pixel identification of OoD objects in addition to segmenting the in-distribution pixels of the training classes. The main challenge with anomaly segmentation lies in the vast number of potential OoD objects, far exceeding the limited number of training classes. Furthermore, in driving scenes, OoD objects frequently coexist with various known objects, a large difference compared to data in image classification which typically has a single salient object in the center of an image. Class imbalance poses yet another challenge in driving scenes: road pixels often comprise a large part of the training data and bias the network towards predicting road. Due to such imbalances in the training data, networks often exhibit a tendency to predict the majority class for OoD objects found on the road with high confidence, as illustrated in Fig. 1, a critical error that carries a

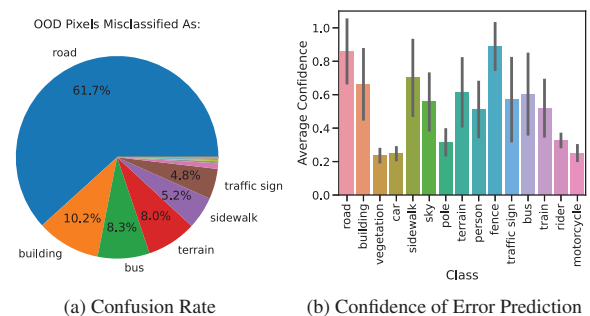


Figure 1: A standard semantic segmentation model, i.e., DeepLabv3+ [5] with a ResNet101 backbone, is trained on Cityscapes [8] and tested on the anomaly segmentation benchmark Fishyscapes Lost & Found [3]. The OoD pixels (anomalies) are frequently misclassified as “road” with very high confidence. This leads to the highest level of risk for autonomous driving.

high level of risk.

A particularly successful technique in the literature for anomaly segmentation is a form of outlier exposure [15]. While we do not have access to potential OoD scenarios at inference time, there are many ready-to-use data, different to the training data, which could serve as potential OoD exemplars during training. For instance, MS COCO [19] and ADE20K [36] are popular choices for training semantic segmentation on Cityscapes [8], e.g., in [4, 2, 30, 23, 25, 11]. Earlier work directly added these datasets to the training set [4], whereas subsequent work discovered that it is more effective to copy and paste objects from them into the training samples of the semantic segmenter, e.g., AnomalyMix in [30]. Prior works mostly focused on losses and/or architecture changes to incorporate the OoD data during training and fine-tuning, as naively using them can lead to performance degradation [23].

In this work, we discover the benefit of improving the OoD data with simple style transfer. In autonomous driving, the collected training data is often specialized by the camera



Figure 2: Visual examples of style transferring from Cityscapes [8] samples to MS COCO [19] objects, using an off-the-shelf ISSA model [17].

model, which varies across different vendors. The misalignment between OoD objects from natural images and driving scenes creates a shortcut cue at training, compromising the effectiveness of synthetic OoD data. We thus propose to first align the style before AnomalyMix. Specifically, we utilize ISSA [17] to transfer the driving scene data style to OoD objects; see examples in Fig. 2. It was initially proposed to improve domain generalization via style mixing among different training samples.

Equipped with the style-aligned OoD data, we further derive a novel fine-tuning loss to adapt the final classification head of a standard semantic segmentation model for anomaly segmentation, i.e., to output a per-pixel OoD score in addition to the predicted semantic label map. Specifically, we cast per-pixel multi-class classification in semantic segmentation as a set of one vs. the rest (OvR) binary classifications at fine-tuning, i.e., one for each class. The OvR loss allows the prediction of “none of the given classes”, which is the natural ground truth for unknown objects. By considering the top-K losses among the training classes, we specifically optimize for the hard cases and maintain comparably low logit responses from all classes. The loss has a simple form with no hard-to-tune hyperparameters.

After fine-tuning, we leverage per-pixel OoD scores derived from the logit outputs of the semantic segmenter for anomaly segmentation. Besides the widely used maximum logit [13] and energy-based OoD score [30], a new per-pixel OoD score is investigated in this work. It subtracts the maximum logit by the minimum logit. We find such logit difference helpful to tell apart OoD pixels (comparably low logits across all classes) from uncertain in-distribution pixels (e.g., boundary pixels with low maximum logit, but much lower minimum logit).

Empirically, our proposal delivers impressive improvements on the anomaly segmentation benchmark Fishscapes Lost & Found [3]. On another real-world anomaly benchmark Road Anomaly [21], we hypothesize that the domain shift is a key bottleneck when adapting the semantic segmenter pre-trained only on Cityscapes [8]. While unknown objects and domain shifts are often separately tackled in the literature, they can jointly manifest in the real world. This co-occurrence is an interesting yet under-explored topic, awaiting for further investigation.

2. Related work

There are two prevalent avenues of approach to tackle anomaly segmentation. The first relies on reconstruction, e.g., [20, 31, 32], assuming that a poor reconstruction indicates something not well understood by the network. The other avenue adapts semantic segmentation models. Our work aligns with the latter approach. Being different to prior art such as [1, 23], we attempt to enable anomaly awareness without model architecture changes.

Synthetic OoD Augmentation Mixing up data samples, despite not necessarily appearing realistic, has proven to be an effective data augmentation strategy, e.g., CutOut [9], Mixup [34] and CutMix [33] for image classification, as well as copy-paste [10] and X-paste [35] on object detection and segmentation. For the anomaly segmentation task, prior work primarily focused on designing the loss and modifying the architecture to make use of the synthetic OoD data. Both [4] and [30] adopted a uniform label distribution as the ground-truth label for the OoD pixels, the latter also adding two extra free energy-based regularization terms and

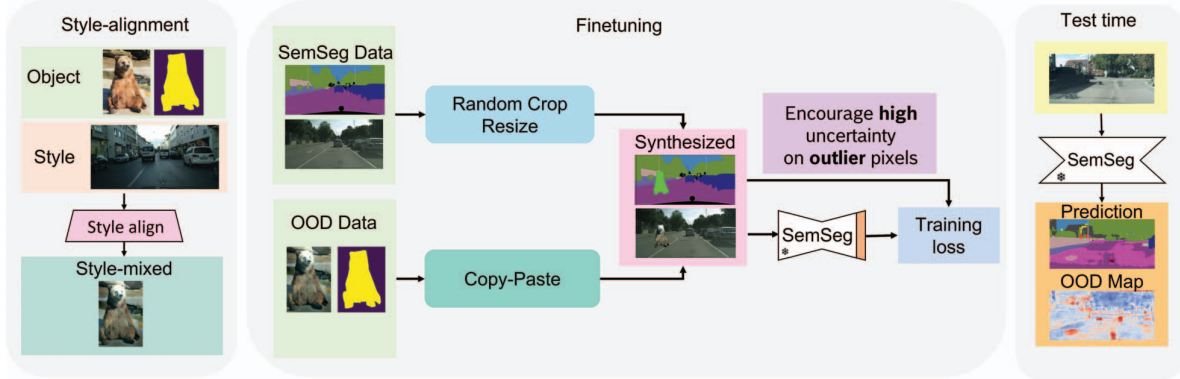


Figure 3: Our proposal begins by aligning the styles between the OoD objects and driving scenes, followed by random copy-pasting. The synthetic OoD data is mixed with the normal driving scenes at a pre-specified ratio. Only the final classification head is fine-tuned; the rest of the network is frozen. At inference, a single-channel OoD map is generated along with the predicted semantic label map. Each pixel of the OoD map carries a scalar score derived from model’s output.

the gambler loss [24]. DenseHybrid [11] also used a free energy-based loss with an additional binary classification head to classify between in-distribution (ID) and OoD pixels. The authors of [23] adopted a contrastive loss, using OoD data as negative examples in contrast to the ID pixels. Recent works [25, 12, 26, 28] discovered the benefit of using transformer-based models and performing mask level classification for segmentation. Building on top of Mask2Former [6], the authors of [25, 26, 28] used an OvR binary classification loss, while an ensemble over anomaly scores of mask-wide predictions was introduced in [12].

In this work, we focus on improving the OoD synthesis, while keeping the fine-tuning loss very simple. As the OoD proxy data often comes from a domain very different from driving scenes, the resulting style difference hinders the effectiveness of naive AnomalyMix. With the aid of style-aligned AnomalyMix, a simple variant of a OvR loss already leads to a strong anomaly segmentation performance.

Per-pixel OoD Score It is a common practice to segment OoD objects based on per-pixel OoD scores, which can be understood as a form of predictive uncertainty on the given training classes. Maximum logit [13] has been shown as a better choice than maximum Softmax response, and the authors of [16] improved the latter by correcting the frequent class bias. Free energy was adopted by [30], and DenseHybrid [11] further subtracted it by the logit generated by its additional classification head.

Our fine-tuning method can work with different OoD scores derived from the logits of the semantic segmentation model. Besides the widely used ones (i.e., maximum logit and energy), we additionally introduce a new OoD score,

i.e., Max-Min Logit. It computes the difference between the maximum and minimum logit, which helps to better distinguish uncertain ID pixels from OoD pixels, reducing false positive rates (FPRs).

3. Method

As shown in Fig. 3, our method consists of style-aligned AnomalyMix for fine-tuning a pre-trained model with a top-K OvR loss, and a single-channel OoD map derived from the predictive logits.

3.1. Style-aligned AnomalyMix

Similar to AnomalyMix in [30], we extract OoD objects from an off-the-shelf dataset, then copy and paste them into the training samples of the semantic segmentation model. Specifically, we only take non-occluded objects from MS COCO and randomly paste them into driving scenes (e.g., Cityscapes) throughout fine-tuning. However, it is quite obvious to see from Fig. 2 that MS COCO and Cityscapes have their own domain-specific styles. The network can make use of this style difference to learn anomaly segmentation, which is not a generalizable solution in the real world. Thanks to the advances in generative models, we can transfer the style of training samples to the OoD objects before mixing them together, i.e., aligning the style before AnomalyMix as shown in Fig. 3. Technically, we take a pre-trained ISSA model [17], which consists of a StyleGAN2-based generator and a masked noise encoder for GAN inversion. The ISSA encoder can respectively extract the style and content of a given exemplar. By mixing up the style and content from different exemplars, the ISSA generator can transfer the style of one exemplar to the other. While the model was pre-trained only on Cityscapes, it successfully

transfers the Cityscapes style to unseen MS COCO objects.

3.2. Fine-tuning Loss

We focus on fine-tuning only the classification head of a pre-trained model, aiming to preserve the integrity of the backbone that is shared with other perception tasks in the autonomous driving system. Under this constraint, grouping the diverse OoD objects into an additional OoD class is suboptimal and can potentially harm the recognition of known objects. While the Softmax-based cross entropy loss over the given training classes is adopted at training, we propose to re-purpose the logit of each class for parameterizing the probability of “being” vs. “not being” in that class. This effectively turns multi-class classification into a set of binary OvR classifications. On the OoD pixels, all class logits should then be supervised with the ground-truth of “none of any given classes”. The concurrent work [25] utilized such a loss for both training a Mask2Former based segmenter, and fine-tuning it on synthetic OoD data. Here, we alternatively minimize the top-K OvR losses across the classes

$$\mathcal{L}_{\text{ood}} = \frac{1}{K|\mathcal{N}_{\text{ood}}|} \sum_{i \in \mathcal{N}_{\text{ood}}} \sum_{k \in \mathcal{S}_{\text{topK}}(i)} -\log \sigma(-s\lambda_{i,k}), \quad (1)$$

where σ is the sigmoid-function, \mathcal{N}_{ood} is the index set of OoD pixels, $\lambda_{i,k}$ is one of the top-K largest logits $\mathcal{S}_{\text{topK}}(i)$ on the pixel i , and s is a hyperparameter to control the slope of the gradients of the sigmoid-function with respect to the logit $\lambda_{i,k}$. Compared to simply averaging the per-class loss over all classes, the top-K variant focuses on improving the worst cases, e.g., the logit of the frequent “road” class often has a larger response than that of other classes on OoD pixels. Thus, it helps to reduce the logit difference across different classes on the OoD pixels. Besides the loss on the OoD pixels, the total fine-tuning loss is a weighted sum $\mathcal{L}_{\text{all}} = \mathcal{L}_{\text{id}} + \gamma\mathcal{L}_{\text{ood}}$ with the original training loss on the in-distribution (ID) pixels.

3.3. Per-pixel OoD score

After fine-tuning, we resort to per-pixel OoD scores for anomaly segmentation. They are derived from the logits generated by the semantic segmenter. The logit response of an input can be regarded as the negative distance of that sample to the corresponding class, i.e., larger logit value indicates closer to the class prototype (which is the final layer weight vector). In the literature, maximum logit response and negative free energy are top-performing scores. The latter is essentially a smoothed version of the former. Besides using them, we additionally consider the difference between the maximum and minimum logit. In semantic segmentation, there are often many ambiguous in-distribution pixels. For instance, the logit responses of boundary pixels are not as high as that of center pixels in segments. While

the maximum logits of boundary pixels may not be sufficiently high, their minimum logits are still much smaller. Therefore, using the gap can help to better tell OoD pixels apart from uncertain in-distribution pixels.

4. Experiments

In this work, we use semantic segmentation models pre-trained on Cityscapes [8], specifically DeepLabv3+ [5] with either a Wide ResNet38 or ResNet101 backbone. We use the mIoU (%) as the in-distribution performance metric and compute the area under receiver operating characteristics (AUC), average precision (AP), and the false positive rate at a true positive rate of 95% (FPR95) to validate our approach. The in-distribution dataset is Cityscapes [8]; two real-world anomaly segmentation benchmarks, Fishyscapes Lost & Found [3] and Road Anomaly [21] are investigated. For fine-tuning, the default loss configuration is with $\gamma = 0.05/0.01$ (WideResNet38/ResNet101), $K = 5$, and $s = 2$. We adopt the AdamW optimizer with an initial learning rate of 10^{-5} , polynomial decay at the rate 0.9, and 20 fine-tuning epochs in total. The mixing probability for style-aligned AnomalyMix is set as 0.1. The OoD loss is computed based on the network output before the final bilinear upsampling operation in DeepLabv3+.

4.1. Comparison with Prior Art

The target setting is to adapt an existing semantic segmentation model for anomaly segmentation without architecture changes. Thus, our comparison is with methods that use the same architectures, and can preserve the original mIoU on the Cityscapes validation set. In this case, DeepLabv3+ with the two types of backbones are the most common choices in the literature. As we can see from Table 1, our method improves over the baselines while preserving the in-distribution performance mIoU. As the methods (in the first block of the table) focused on deriving OoD scores from the pre-trained models, their mIoU stays exactly identical to the original performance. However, such post-hoc OoD score derivation does not lead to strong anomaly segmentation performance, as the pre-trained models are over-confident and uncalibrated as shown in Fig. 1. Both Meta-OOD [4] and PEBAL [30] trained the model with OoD data, thus outperform the post-hoc methods. Compared to them, we better preserve the original mIoU performance. Notably, on AP, which measures the precision of localizing the OoD pixels, we outperform them by a large margin. Moreover, our method is compatible with different OoD scores. Compared to prior methods that also used Max Logit and Energy, our fine-tuning method greatly improved their performance. Among the three OoD scores, the new score, i.e., Max-Min Logit, is better at reducing FPRs, but slightly worse at increasing APs. It can be an interesting future step to fuse multiple

DeepLabv3+ Method	w. WideResNet38 backbone				w. ResNet101 backbone			
	Cityscapes mIoU \uparrow	Fishyscapes L & F AUC \uparrow	AP \uparrow	FPR95 \downarrow	Cityscapes mIoU \uparrow	Fishyscapes L & F AUC \uparrow	AP \uparrow	FPR95 \downarrow
Max Softmax Pred. [13]		89.29	4.59	40.59		86.99	6.02	45.63
Max Logit (ML)[13]		93.41	14.59	42.21		92.00	18.77	38.13
Entropy [14]	90.62	90.82	10.36	40.34	80.50	88.32	13.91	44.85
Energy [22]		93.72	16.05	41.78		93.50	25.79	32.26
Standardized ML [16]		94.97	22.74	33.49		96.88	36.55	14.53
Meta-OOD [4]	89.00	93.06	41.31	37.69	-	-	-	-
PEBAL [30]	89.12	98.96	58.81	4.76	-	99.09	59.83	6.49
Ours (Max Logit)		98.71	71.94	6.42		98.45	67.35	9.36
Ours (Energy)	90.39	98.79	70.87	5.88	80.50	98.58	69.93	8.38
Ours (Max-Min Logit)		98.87	70.84	5.52		98.83	66.32	5.74

Table 1: Comparison of our proposal with prior art on the Fishyscapes Lost & Found [3]. Note that our comparison scope lies in the line of methods aiming to adapt a standard semantic segmentation model for anomaly segmentation without major architecture changes and extra (sub)networks. The pre-trained model checkpoints and the baseline numbers are taken from the source specified in PEBAL [30].

OoD Score	w/o. Style Align.			w. Style Align.		
	AUC \uparrow	AP \uparrow	FPR95 \downarrow	AUC \uparrow	AP \uparrow	FPR95 \downarrow
Max Softmax Pred.	94.82	32.32	20.76	+1.55	+18.84	-2.74
Entropy	96.21	47.14	19.76	+1.13	+16.37	-2.85
Max Logit	97.84	51.79	12.61	+0.61	+15.56	-3.25
Energy	98.02	52.32	11.92	+0.56	+17.61	-3.54
Max - Min. Logit	98.24	45.02	9.14	+0.59	+21.30	-3.40

Table 2: Style-alignment before AnomalyMix introduces consistent gains across different OoD scores in all metrics. Here, we use DeepLabv3+ (ResNet101) and Fishyscapes Lost & Found [3].

good OoD scores together for further improvements.

4.2. Ablation Study

We first evaluate the benefit from using style alignment before AnomalyMix. As shown in Table 2, style alignment is beneficial in all cases, independent of the choice of the OoD score and in all OoD evaluation metrics.

Next, we compare different choices of K for our method in Table 3. The baseline PEBAL and OvR essentially considers all classes, performing worse than our top- K OvR variant. This reveals the benefit of focusing on improving the worst cases, i.e., the top- K confident predicted classes. Among all three OoD scores, $K = 5$ is generally better than $K = 3$ and $K = 7$, and they all outperform the two competitive baselines. Overall, the performance is insensitive to the value of K , making it an easy-to-tune hyperparameter.

Finally, we compare different fine-tuning losses. The two baselines, Uniform and Energy, respectively indicate minimizing the cross entropy with a uniform label distribution, and maximizing the negative Log-Sum-Exp of the

Method	K	AUC \uparrow	AP \uparrow	FPR95 \downarrow
PEBAL [30]	-	99.09	59.83	6.49
OvR (Max Logit)	-	97.70	52.24	12.97
OvR (Energy)	-	97.95	59.96	12.09
OvR (Max-Min Logit)	-	98.52	59.19	7.51
Ours (Max Logit)	3	98.34	60.86	10.50
	5	98.45	67.35	9.36
	7	98.12	63.88	11.40
Ours (Energy)	3	98.48	64.07	9.73
	5	98.58	69.93	8.38
	7	98.28	68.30	10.47
Ours (Max-Min Logit)	3	98.79	58.59	6.37
	5	98.83	66.32	5.74
	7	98.69	66.56	6.43

Table 3: Ablation on the choice of K for our proposal. The fine-tuning loss of PEBAL and OvR involves all classes, whereas only the top- K ones are considered in **Ours**. Here, we use DeepLabv3+ (ResNet101) and Fishyscapes Lost & Found [3].

multi-class logits. In addition, we compare with the OvR loss averaged over all classes. Compared to them, our top- K OvR loss does not involve all multi-class logits at each optimization step, but rather focuses on the hardest cases, i.e., top- K high logit responses to the OoD pixel. As shown in Fig. 4, our loss is the only one that excels at both AP and FPR95. Notably, our fine-tuning loss leads to superior performance using all three OoD scores. As all these scores reflect some form of ‘‘uncertainty’’ on OoD pixels, the consistent performance improvement indicates our fine-tuning




Figure 4: Comparison of different fine-tuning losses, using DeepLabv3+ (ResNet101) and Fishyscapes Lost & Found. The proposed top-K OvR loss improves the AP and FPR95 when combining with the three top-performing OoD scores.

process is effective at inducing a “none of the given classes” prediction on anomalies.

4.3. Visual Results on Fishyscapes Lost & Found

Fig. 5 visualizes the semantic label map and per-pixel OoD score produced by PEBAL [30] and our method (Max-Min Logit). The training loss of PEBAL has multiple terms. In addition to the uniform label prediction and energy loss, the gambler loss [24] introduced an extra class in prediction. It is quite interesting to observe that some anomalies are classified as the extra class (in green). However, as also pointed out by the authors, it is not yet a reliable prediction for all cases, and the per-pixel OoD score is more informative about the anomalies as we can observe in Fig. 5. Compared to PEBAL, our method uses a single top-K OvR loss without the extra class. The per-pixel OoD score has a strong response to the anomalies, while the false positives, such as paintings on the road or sidewalk, are less frequent than that in PEBAL.

4.4. Domain Shift in Road Anomaly

Table 4 shows the results on the Road Anomaly [21] benchmark. While our method still noticeably outperforms the others, all solutions are still far from delivering satisfactory results, with APs below 50% and FPRs above 35%. To gain some insights on the performance gap between Fishyscapes Lost and Found [3] and Road Anomaly [21], Fig. 6 shows some interesting visual examples. First, we can clearly notice the domain shift from Cityscapes [8], including not only the style change, but also geographic shift, and object part ambiguity. For instance, the first example is closer to Cityscapes and both approaches handle it well. In the next example, we observe false positives from unfamiliar fences. More interestingly, it is arguable whether the umbrella is an outlier or a part of person in this case. The overall scene is also quite different to the Cityscapes examples, which were collected in Europe. For

Method	AUC ↑	AP ↑	FPR95 ↓
Max Softmax Pred. [13]	67.53	15.72	71.38
Max Logit (ML)[13]	72.78	18.98	70.48
Entropy [14]	68.80	16.97	71.10
Energy [22]	73.35	19.54	70.17
Standardized ML [16]	75.16	17.52	70.70
SynBoost [2]	81.91	38.21	64.75
PEBAL [30]	87.63	45.10	44.58
Ours (Max Logit)	89.04	43.82	39.04
Ours (Energy)	89.65	46.39	38.09
Ours (Max-Min Logit)	87.15	38.38	43.40

Table 4: Comparison of our proposal with prior art on Road Anomaly [21], using DeepLabv3+ (WideResNet38). Note that our comparison scope lies in the line of methods aiming to adapt a standard semantic segmentation model for anomaly segmentation without major architecture changes and extra (sub)networks. The pre-trained model checkpoints and the baseline numbers are taken from the source specified in PEBAL [30].

the last two examples, the rocks on the road should be detected as anomalies, whereas many small stones on the road could be just “road”. To separate both cases, the model will need to have a higher-level concept of the scene if the stone road was not shown during training. Therefore, to tackle the Road Anomaly benchmark, our hypothesis is that focusing on unknown objects with the assumption of no domain shift is suboptimal. It can be interesting to combine anomaly segmentation with domain generalization techniques, e.g., ISSA [18], RobustNet [7], and StyleLess [29]. Moreover, the energy-based OoD score appears to be more robust than Max Logit and Max-Min Logit under the domain shift. Different to the observation on Fishyscapes [3], subtracting the minimum logit is no longer effective to reduce false positives. We hypothesize that logits are generally less reliable

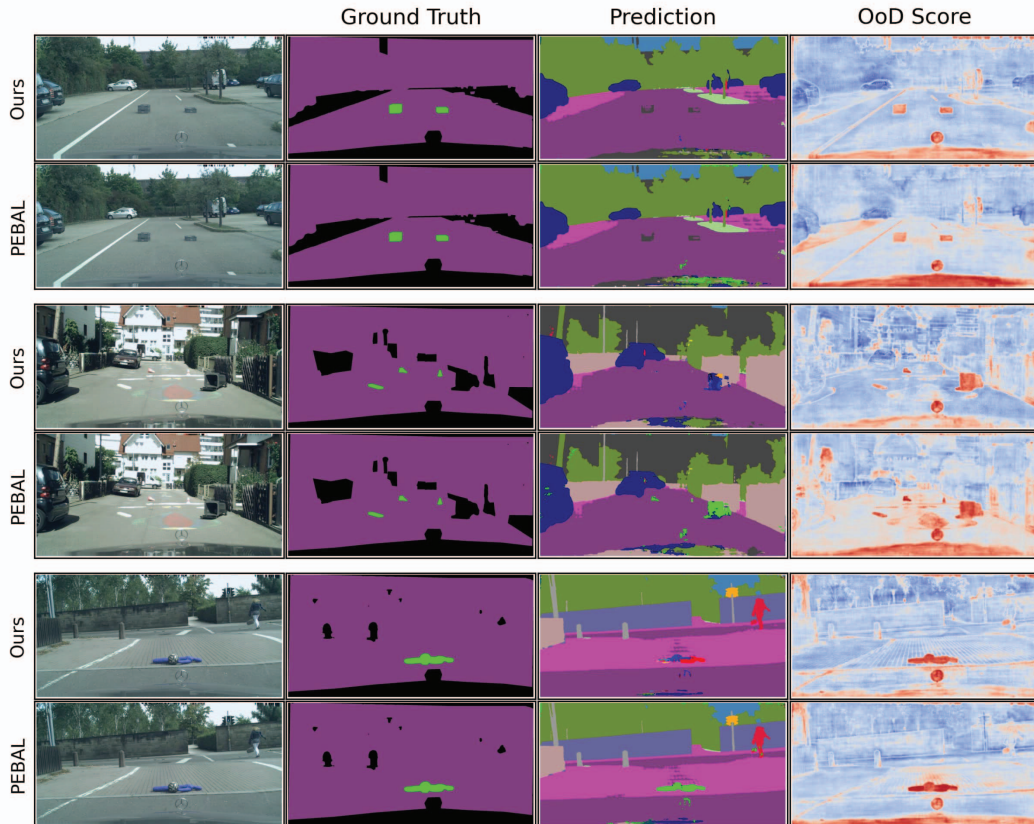


Figure 5: Visual examples using PEBAL [30] and our method for semantic segmentation and anomaly segmentation on Fishyscapes [3]. Our pixel-wise OoD score has less false positive responses to drawings on the road and side walk curb-stone, while PEBAL can map some anomalies into its extra class prediction (e.g., dummy person, trash bin), but not always consistently (e.g., boxes and white stick).

under the domain shift, and simple logit processing is ineffective to reduce false positives caused by domain shifts.

Notably, several recent works, i.e., [23] and [25, 28, 26], achieved promising progresses, even though they lay beyond the scope of fine-tuning a standard semantic segmentation model. The former added an extra network to Deeplabv3+, and exploited contrastive learning for regularizing the feature space. Contrastive learning is a successful representation learning technique that improves OoD generalization in different tasks. The latter ones built on top of Mask2Former [6], which is not only transformer-based, but also changes from per-pixel multi-class classification to mask proposals and mask-level classification. Moreover, it is beneficial to pre-train the model on multi-domain driving scene data such as Mapillary [27] plus Cityscapes [8], indicating the value of improved domain generalization achievable by more diverse source data at training.

5. Conclusion

We proposed a simple fine-tuning process that demonstrates impressive performance gains on the anomaly segmentation benchmark Fishyscapes Lost & Found. The fine-tuned semantic segmentation model preserves the performance on the original task, and additionally generates high-quality per-pixel OoD scores for anomaly segmentation in situations where there is no major domain shift from training/fine-tuning to testing. Notably, our results indicate the value of improving the existing synthetic OoD synthesis process extensively adopted in previous studies. While our proposed style alignment has mitigated the synthetic-to-real gap, it remains an ongoing challenge to completely close the gap. We hope that our findings will inspire further exploration in this direction. Additionally, it is highly interesting to consider the domain shift together with anomalies, e.g., combining with domain generalization techniques for tackling the Road Anomaly challenge.

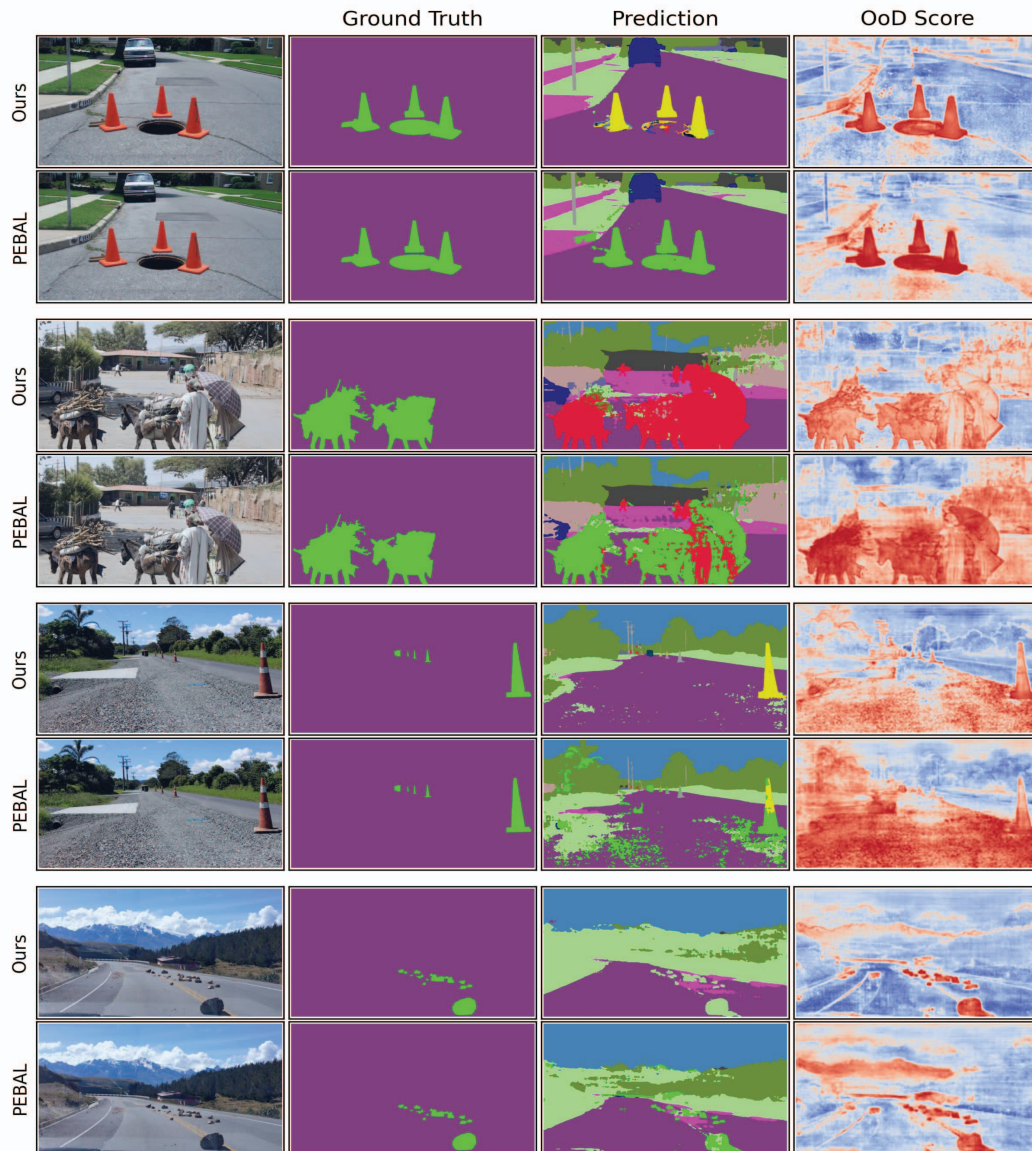


Figure 6: Visual examples using PEBAL [30] and our method for semantic segmentation and anomaly segmentation on Road Anomaly [21], which is clearly more challenging than Fishyscapes due to the larger domain gap. We can observe a large amount of false positives generated by both methods, where our method is slightly better than PEBAL. It is also worth mentioning that some of them are actually arguable anomalies.

References

- [1] Victor Besnier, Andrei Bursuc, David Picard, and Alexandre Briot. Triggering failures: Out-of-distribution detection by learning from local adversarial attacks in semantic segmentation. In *ICCV*, 2021. 2
- [2] Giancarlo Di Biase, Hermann Blum, Roland Siegwart, and César Cadena. Pixel-wise anomaly detection in complex driving scenes. In *CVPR*, 2021. 1, 6
- [3] Hermann Blum, Paul-Edouard Sarlin, Juan Nieto, Roland Siegwart, and Cesar Cadena. The fishyscapes benchmark: Measuring blind spots in semantic segmentation. *IJCV*, 2021. 1, 2, 4, 5, 6, 7
- [4] Robin Chan, Matthias Rottmann, and Hanno Gottschalk. Entropy maximization and meta classification for out-of-distribution detection in semantic segmentation. In *ICCV*,

2021. 1, 2, 4, 5
- [5] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 1, 4
- [6] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. 2022. 3, 7
- [7] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *CVPR*, 2021. 6
- [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 *CVPR*, 2016. 1, 2, 4, 6, 7
- [9] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 2
- [10] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D. Cubuk, Quoc V. Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, 2021. 2
- [11] Matej Grcic, Petra Bevandic, and Sinisa Segvic. Densehybrid: Hybrid anomaly detection for dense open-set recognition. In *ECCV*, 2022. 1, 3
- [12] Matej Grcic, Josip Šarić, and Siniša Šegvić. On advantages of mask-level recognition for outlier-aware segmentation. 2023. 3
- [13] Dan Hendrycks, Steven Basart, Mantas Mazeika, Andy Zou, Joe Kwon, Mohammadreza Mostajabi, Jacob Steinhardt, and Dawn Song. Scaling out-of-distribution detection for real-world settings. *ICML*, 2022. 2, 3, 5, 6
- [14] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017. 5, 6
- [15] Dan Hendrycks, Mantas Mazeika, and Thomas G. Dietterich. Deep anomaly detection with outlier exposure. In *ICLR*, 2019. 1
- [16] 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 *ICCV*, 2021. 3, 5, 6
- [17] Yumeng Li, Dan Zhang, Margret Keuper, and Anna Khoreva. Intra- & extra-source exemplar-based style synthesis for improved domain generalization. *IJCV*, 2023. 2, 3
- [18] Yumeng Li, Dan Zhang, Margret Keuper, and Anna Khoreva. Intra-source style augmentation for improved domain generalization. In *WACV*, 2023. 6
- [19] 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 *ECCV*, 2014. 1, 2
- [20] Krzysztof Lis, Krishna Nakka, Pascal Fua, and Mathieu Salzmann. Detecting the unexpected via image resynthesis. In *ICCV*, 2019. 2
- [21] Krzysztof Lis, Krishna Kanth Nakka, Pascal Fua, and Mathieu Salzmann. Detecting the unexpected via image resynthesis. In *ICCV*, 2019. 2, 4, 6, 8
- [22] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *NeurIPS*, 2020. 5, 6
- [23] Yuyuan Liu, Choubo Ding, Yu Tian, Guansong Pang, Vasileios Belagiannis, Ian Reid, and Gustavo Carneiro. Residual pattern learning for pixel-wise out-of-distribution detection in semantic segmentation. *arXiv preprint arXiv:2211.14512*, 2022. 1, 2, 3, 7
- [24] Ziyin Liu, Zhikang Wang, Paul Pu Liang, Russ R Salakhutdinov, Louis-Philippe Morency, and Masahito Ueda. Deep gamblers: Learning to abstain with portfolio theory. In *NeurIPS*, 2019. 3, 6
- [25] Nazir Nayal, Misra Yavuz, João F. Henriques, and Fatma Güney. Rba: Segmenting unknown regions rejected by all. In *ICCV*, 2023. 1, 3, 4, 7
- [26] Alexey Nekrasov, Alexander Hermans, Lars Kuhnert, and Bastian Leibe. Ugains: Uncertainty guided anomaly instance segmentation. *arXiv preprint arXiv:2308.02046*, 2023. 3, 7
- [27] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017. 7
- [28] Shyam Nandan Rai, Fabio Cermelli, Dario Fontanel, Carlo Masone, and Barbara Caputo. Unmasking anomalies in road-scene segmentation. In *ICCV*, 2023. 3, 7
- [29] Julien Rebut, Andrei Bursuc, and Patrick Pérez. Styleless layer: Improving robustness for real-world driving. In *IROS*, 2021. 6
- [30] 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. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *ECCV*, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [31] Tomas Vojir, Tomáš Šipka, Rahaf Aljundi, Nikolay Chumerin, Daniel Olmeda Reino, and Jiri Matas. Road anomaly detection by partial image reconstruction with segmentation coupling. In *ICCV*, 2021. 2
- [32] Yingda Xia, Yi Zhang, Fengze Liu, Wei Shen, and Alan L Yuille. Synthesize then compare: Detecting failures and anomalies for semantic segmentation. In *ECCV*, 2020. 2
- [33] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 2
- [34] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 2
- [35] Hanqing Zhao, Dianmo Sheng, Jianmin Bao, Dongdong Chen, Dong Chen, Fang Wen, Lu Yuan, Ce Liu, Wenbo Zhou, Qi Chu, et al. X-paste: Revisit copy-paste at scale with clip and stablediffusion. *arXiv preprint arXiv:2212.03863*, 2022. 2
- [36] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. 1