

UGainS: Uncertainty Guided Anomaly Instance Segmentation

Alexey Nekrasov¹, Alexander Hermans¹, Lars Kuhnert², and Bastian Leibe¹

¹ Visual Computing Institute, RWTH Aachen University, Aachen, Germany
nekrasov@vision.rwth-aachen.de

² Ford Motor Company, Dearborn, Michigan, USA

Abstract. A single unexpected object on the road can cause an accident or may lead to injuries. To prevent this, we need a reliable mechanism for finding anomalous objects on the road. This task, called anomaly segmentation, can be a stepping stone to safe and reliable autonomous driving. Current approaches tackle anomaly segmentation by assigning an anomaly score to each pixel and by grouping anomalous regions using simple heuristics. However, pixel grouping is a limiting factor when it comes to evaluating the segmentation performance of individual anomalous objects. To address the issue of grouping multiple anomaly instances into one, we propose an approach that produces accurate anomaly instance masks. Our approach centers on an out-of-distribution segmentation model for identifying uncertain regions and a strong generalist segmentation model for anomaly instances segmentation. We investigate ways to use uncertain regions to guide such a segmentation model to perform segmentation of anomalous instances. By incorporating strong object priors from a generalist model we additionally improve the per-pixel anomaly segmentation performance. Our approach outperforms current pixel-level anomaly segmentation methods, achieving an AP of 80.08% and 88.98% on the Fishyscapes Lost and Found and the RoadAnomaly validation sets respectively. Project page: <https://vision.rwth-aachen.de/ugains>

Keywords: Out-of-Distribution Detection · Anomaly Segmentation · Semantic Segmentation · Instance Segmentation

1 Introduction

Current semantic segmentation approaches [8,51,7] achieve impressive performance for known classes on driving datasets. However, in the context of autonomous driving, the real world is more complex than the test scenarios captured in today’s benchmarks [47]. In particular, a wide variety of objects, *e.g.* wild animals or debris on the road, can pose a threat to vehicles or passengers when encountered on the road. Since individual objects may occur very rarely, learning to reliably detect such potentially dangerous objects in the long tail of the data distribution is a major challenge. Even datasets with millions of images may only contain few frames with anomalous objects [29]. Since modern roads are built according to well-defined rules, it creates a model bias towards regular road context. To make matters worse, deep networks tend to make overconfident predictions and often fail on out-of-distribution cases [16,25]. Together, rare occurrences of




Fig. 1: Instance prediction of anomalous objects. We propose UGainS, an approach for instance segmentation of anomalous objects. UGainS combines a strong generalist segmentation model (SAM [26]), and a strong out-of-distribution method (RbA [39]). This allows us to predict an anomaly instance segmentation, refined uncertainty maps as well as an in-distribution semantic segmentation.

anomalous instances and model overconfidence, creates a need to address the problem of detecting anomalous road objects.

Most anomaly segmentation approaches focus on predicting a per-pixel outlier score to create an uncertainty map of the visible scene. Available approaches in this category are based on classification uncertainty [27,18,19], reconstruction errors [50,34], or they rely on data augmentation with synthetic outliers to regularize models [15,6,2].

In this paper, we propose an approach that addresses the anomaly instance segmentation problem. Our approach uses uncertainty maps to detect image structures that are not well captured by the given training data [39]. However, we then use the extracted outlier pixels as point guidance for a class-agnostic segmentation pipeline [26] to segment anomalous instances. The resulting out-of-distribution instance segmentations (see Fig. 1) can then be used for mining rare objects [23], which could be added to the training set to improve future model performance. When using such a point-guidance scheme, the key challenge is to decide how to sample and prioritize the guidance points in order to achieve a good coverage of the relevant anomalous regions with as few proposals as possible. We systematically explore this design space and present a detailed experimental evaluation of the different design choices.

Current anomaly segmentation benchmarks evaluate methods based on per-pixel or component-level metrics [5]. Per-pixel metrics make it possible to evaluate performance under large class imbalance, but they favor large objects. Component-level metrics are a proxy for object-level segmentation and evaluate the intersection between predicted and ground truth anomaly regions. To generate object instances, benchmarks often use the connected components algorithm in the metric formulation, which is a proxy for instance-level performance [5]. However, current benchmarks do not provide a clear way to evaluate instance prediction performance within anomaly regions and focus on finding the entire region. We argue that finding individual instances within large anomalous regions can be beneficial for downstream tasks such as object mining [29], tracking [35], or model introspection. Therefore we propose an anomaly instance segmentation evaluation measure (*instance AP* or *iAP*) and present experimental results using this measure, in addition to also evaluating on the standard per-pixel measures. Our proposed approach achieves state-of-the-art per-pixel performance on the Road-Anomaly and Fishyscapes Lost and Found benchmark datasets. In addition, we demon-

strate the utility of the extracted anomaly instances by visualizing cases where previous methods would not be able to recover separate anomalies.

Our contributions are summarized as follows: 1) We propose UGainS, a novel approach for predicting anomalous object instances based on point-guided segmentation with points sampled from an epistemic uncertainty map. 2) We evaluate the design choices of performing uncertainty guidance for a generalist segmentation model and provide detailed ablations to highlight the influence of the different steps of our approach. 3) UGainS achieves state-of-the-art anomaly segmentation performance on the Fishyscapes Lost and Found [3] validation set with 80.08% pixel-level average precision, and on the RoadAnomaly [34] dataset with 88.98% pixel-level average precision. Notably, UGainS improves per-pixel performance compared with previous anomaly segmentation methods, while at the same time extracting interpretable instance segmentation masks.

2 Related Work

2.1 Pixel-Wise Anomaly Segmentation

The anomaly segmentation task has attracted attention in recent years, and is derived from the out-of-distribution classification task. Early approaches used probabilistic modeling directly to obtain outlier scores and did not require any special training procedures. These methods assume that anomalies should have high uncertainty values. Methods have approximated the epistemic uncertainty of predictions using deep ensembles [27], Monte-Carlo dropout [38], or directly using SoftMax probabilities [18]. Liang *et al.* [30] recently proposed to use Gaussian Mixture Models in a segmentation model to capture class conditional densities. In another line of work, methods focus on directly improving the predicted outlier scores [6,24].

Previous studies have relied on embedding space estimation [28,3] or direct classification of anomalous regions [3]. Reconstruction-based methods rely on re-synthesis of input images [34,33,10,50,40] and the assumption that anomalies cannot be reconstructed well. Other approaches train or fine-tune on images with artificial anomalies [15,14,2,18,5]. Following early methods derived from classification [31], recent methods focus on using adversarial examples on an unmodified in-distribution model [28,1]. While the majority of previous approaches use per-pixel prediction models [7], the most recent approach, Rejected by All (RbA), focuses on acquiring anomaly scores from a set of predicted mask-class pairs [39]. It assumes that pixels which are not covered by any masks correspond to anomalous objects. We follow this line of work and extend it to instance segmentation.

2.2 Anomaly Instance Segmentation

While the majority of anomaly segmentation approaches focus on a per-pixel anomaly classification, and thus ignore any instance information, several recent methods have attempted to address anomaly instance segmentation. MergeNet [17] uses two networks

to predict anomalous regions. LiDAR-guided small object segmentation focuses on retrieving small objects [44]. Xue *et al.* [53] uses occlusion edges to represent anomalous instances and a sliding window approach to generate object proposals. These methods use low-level heuristics, based on connected components or over-segmentation, to retrieve anomalous instances. The SegmentMeIfYouCan [5] benchmark suggests using component-level metrics to evaluate object segmentation performance, grouping anomalous pixels with connected components and evaluating component-level intersection over union. However, grouping multiple spatially close objects may hinder downstream tasks such as tracking, data collection, or model introspection. Treating multiple connected instances as a single object may furthermore be limiting for future advances in anomaly instance segmentation.

A few works tackle the problem of finding anomalous instances without such heuristics. OSIS [48] is an open-set instance segmentation work on 3D point clouds (which was adapted to images by Gasperini *et al.* [13]). It learns class-agnostic spatial instance embeddings that are clustered to predict an object. A more recent work, EOPSN [20], relies on existing unknown instances in the unannotated regions during training, which is a limiting factor for autonomous driving setups as training sets are often close to densely annotated. The concurrent Holistic Segmentation U3HS approach [13] also uses uncertainty scores to obtain anomaly instance segmentations, however U3HS uses explicit clustering of learned instance features. In contrast, we use an uncertainty-guided point sampling scheme with a generalist instance segmentation model to provide anomaly instance segmentation, which achieves a stronger instance segmentation performance in practice.

2.3 Promptable Segmentation Models

Recent advances in large-scale text-guided training for classification [22,42] are sparking interest in large-scale open-vocabulary and open-world segmentation [11,43,52]. XDecoder [55] addresses multiple segmentation tasks by combining the Mask2Former decoder [8] with text supervision. A recent extension, SEEM [56], extends the work to open-set segmentation using multiple prompt types. Our method uses the recent Segment Anything Model (SAM) [26]. Kirillov *et al.* propose the large-scale segmentation dataset SA-1B with class-agnostic mask annotations and SAM for interactive segmentation. SAM accepts prompts of different types, such as points, masks, boxes, or text. Several other works extend SAM to medical image segmentation [37,49], image inpainting [54], camouflaged object detection [45], industrial anomaly segmentation [4], and other applications [21]. The base SAM model is limited to prompted prediction and lacks query interaction for separate objects. As such it is difficult to use for general instance segmentation with a fixed set of classes. By first predicting general outlier regions, we can sample points and utilize SAM for anomaly instance segmentation though.

3 Method

UGainS combines a strong in-distribution segmentation model with a strong generalist segmentation model (see Fig. 2). At the core of our method is a Mask2Former [8]

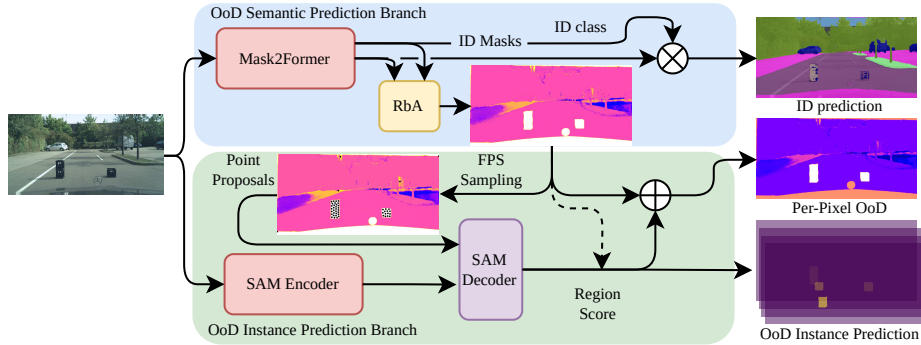


Fig. 2: Method Overview. UGainS consists of two main blocks: an OoD Semantic Prediction Branch (□) and OoD Instance Prediction Branch (□). We use the Mask2Former model [8] to get class and mask predictions for in-distribution (ID) classes. From these predictions we compute a per-pixel uncertainty map using RbA [39]. Using farthest point sampling (FPS) on the uncertainty map we generate point proposals for anomalous objects. These point proposals are then used as input for the Segment Anything Model (SAM) [26] to get anomaly instance predictions. The final prediction is an in-distribution semantic segmentation, instance-refined per-pixel anomaly scores, and a set of out-of-distribution anomaly instance masks.

regularized to avoid predicting objects in unknown regions [39], and the Segment Anything Model (SAM) [26]. The semantic branch predicts masks for known semantic classes, such as cars or pedestrians, and avoids prediction in regions occupied by out-of-distribution objects. Based on the Reject by All (RbA) score definition [39] these unknown regions will have higher uncertainty values. We obtain point proposals for the out-of-distribution branch by sampling points in these regions. For each sampled point, the Segment Anything Model generates an instance mask. We average the anomaly scores under the mask regions to obtain the final prediction. To this end, we propose to tackle the problem of anomaly segmentation as segmenting *everything* we know and prompting to segment *everywhere* where we do not know. In this section, we provide a detailed description of each component of our approach.

OoD Semantic Prediction Branch. With the out-of-distribution (OoD) semantic prediction branch we predict in-distribution (ID) semantic segmentation and OoD per-pixel scores (see Fig. 2, □). Here we use the RbA adaptation of the Mask2Former model. The Mask2Former model consists of a backbone with a pixel decoder and a transformer decoder. For an image \mathbf{x} , the Swin-B [36] backbone produces latent features at a coarse scale, and in combination with the pixel decoder, produces an L -layer feature pyramid $\{f_i(\mathbf{x})\}_{i=1}^L$. The transformer decoder processes these multi-resolution C -dimensional features together with the N learned object queries $\mathbf{Q}_0 \in \mathbb{R}^{N \times C}$. The object queries are sequentially updated by cross-attending the multi-resolution features in a round-robin fashion. The final set of object queries \mathbf{Q}_i is then multiplied with the high-resolution feature map $f_L(\mathbf{x})$, generating $\mathcal{M} \in \mathbb{R}^{N \times H \times W}$ binary masks for an image of the size $H \times W$. In addition, class probabilities $\mathcal{L} \in \mathbb{R}^{N \times (K+1)}$ are predicted from queries, where $(K+1)$ are closed-set classes with an additional class to signal a query deacti-

vation. During training we use the RbA approach, which regularizes Mask2Former to predict no in-distribution class or mask in an anomalous region. RbA uses a per-pixel negative sum of mask predictions to score anomalous regions, where every pixel in a mask is assigned a sum of predicted class posterior probabilities:

$$U(\mathbf{x}) = - \sum_{k=1}^K \sum_{n=1}^N \mathcal{L}_{k,n}(\mathbf{x}) \cdot \mathcal{M}_{k,n}(\mathbf{x}).$$

This results in per-pixel RbA scores $U(\mathbf{x}) \in [-N, 0]^{H \times W}$, where a low score indicates that a pixel is contained in one or multiple predicted masks for in-distribution classes. A high score indicates that the pixel belongs to no masks or all masks have a very low confidence for the pixel. We use the RbA method because it predicts low anomaly scores on semantic boundaries, *i.e.*, road-to-sidewalk boundaries (Fig. 1), and since it generally has a strong anomaly segmentation performance. However, in principle we could use any model that predicts a per-pixel anomaly score. For a detailed overview of Mask2Former and RbA we refer the interested reader to the respective papers [8,39].

OoD Instance Prediction Branch. To get instance predictions we use an instance prediction branch that is based on semantic outlier scores (see Fig. 2, \square). A naive approach to obtain instances would be to apply the connected components algorithm on a binary mask from outlier scores. However, as discussed earlier, this approach does not work when objects have an overlap or an intersection with known classes (see Fig. 3). We address this problem with a novel approach using a segmentation model promptable with point proposals. To convert the per-pixel anomaly predictions into instances, we prompt SAM to predict instance masks for a set of sampled anomalous points. However, there is no obvious way to create a prompt from a noisy and partially incorrect anomaly mask. While we could use all anomaly points as prompts, this would result in increased computational complexity and diminishing results (see Sec. 4.2). We therefore resort to farthest point sampling (FPS) which we run on the 2D coordinates of the thresholded anomaly points. (Ablations on hyperparameters and other sampling strategies can be found in Section 4.2.)

We then prompt SAM with the resulting anomaly point proposals to predict instance masks. SAM consists of an image encoder, a prompt encoder, and a lightweight mask decoder. The image encoder is a plain ViT-H [12] which extracts latent dense image features. During the prompting, we obtain a prompt embedding for each sampled coordinate, comprised of a positional and learned embedding. Finally, the lightweight mask decoder from SAM processes image features together with prompt embeddings to predict object masks. The mask decoder maps each prompt to a hierarchy of three masks with foreground probabilities. Based on the predicted foreground probability, the best of the three masks is selected to remove the ambiguity. To provide each mask with anomaly scores, we use an average of the uncertainty values in the mask region. After obtaining individual scores, we remove duplicate masks using non-maximum suppression (NMS).

Since we obtain high-quality object masks, we further combine them with the initial per-pixel anomaly predictions. Similar to the computation of the per-pixel anomaly scores in the semantic branch, we obtain the anomaly scores as a sum of individual

scores. We add these new instance-based anomaly scores to the existing RbA scores to produce new per-pixel scores, which can be seen as a form of RbA ensembling. Effectively, we improve the per-pixel scores by using strong object priors from the instance predictions. In summary, UGainS predicts an in-distribution semantic predictions, a set of out-of-distribution anomaly instance masks, and an improved per-pixel anomaly score.

4 Experiments

Datasets. We use the Cityscapes dataset [9] for in-distribution training. We use 2975 training and 500 validation images with 19 labeled classes with instance and semantic labels.

To validate the quality of anomaly predictions we use the Fishyscapes Lost and Found (FS L&F) [3] and RoadAnomaly [34] datasets. Fishyscapes L&F [3] is a filtered and re-annotated version of the Lost and Found dataset [41]. The original Lost and Found dataset (L&F) was created by placing anomalous objects on the road in front of a vehicle. Lost and Found contains both semantic and instance annotations for anomalous objects, as well as a *road* annotation; it labels the rest as *ignore*. However, it contains noisy annotations, *e.g.* labeling bicycles and children as anomalies [3]. Thus, the FS L&F dataset has become a standard benchmark in the anomaly segmentation community. Fishyscapes L&F contains 100 images in the validation set and 275 images in the test set. However, the FS L&F dataset does not originally contain instance labels. To evaluate instance segmentation performance, we re-label objects in the validation set by applying connected components on semantic annotations. This works for FS L&F, since objects in the dataset are far apart. For comparability with pixel-level evaluation, we only use the newly generated instance for the instance-level evaluation.

The RoadAnomaly [34] dataset is an earlier version of the RoadAnomaly21 dataset, which is part of the SegmentMeIfYouCan [5] benchmark. SegmentMeIfYouCan does not provide annotations for instance masks in the validation set. The RoadAnomaly dataset contains 60 images with semantic labels and partial instance annotations, as such we only evaluate pixel-level metrics on this dataset.

Dataset-Specific Adjustments. The RbA method has strong responses for the car hood and rectification artifacts at the image boundaries, since these areas are not labeled in the Cityscapes dataset. We create a dataset-specific ignore mask and do not sample points from the masked regions. While this is a heuristics, we argue such an ignore masks can easily be constructed during a calibration process and does not harm the generality of our method. Furthermore, common evaluation protocols exclude predictions from ignore regions, so we do not gain an unfair advantage over other methods.

Evaluation Metrics. To evaluate the out-of-distribution performance we use the standard community metrics *average precision* (AP) and *false positive rate at 95% true positive rate* (FPR_{95}). To distinguish these pixel-level metrics from our later instance-level metrics, we prepend a p: pAP and p FPR_{95} . Evaluating anomalous instances on the other hand is not trivial. One approach is to use a component-level metric, such as the sIoU [5] evaluation protocol. In the sIoU formulation, the connected components




Fig. 3: Groups of anomalous pixels as instances. Current evaluation of methods do not distinguish object groups and individual objects to evaluate the quality of object retrieval. We argue that for purposes such as tracking, data mining, or model introspection, it is important to recover individual instances rather than groups of objects. Methods or metrics that rely on connected components are therefore not well suited for instance segmentation in practice.

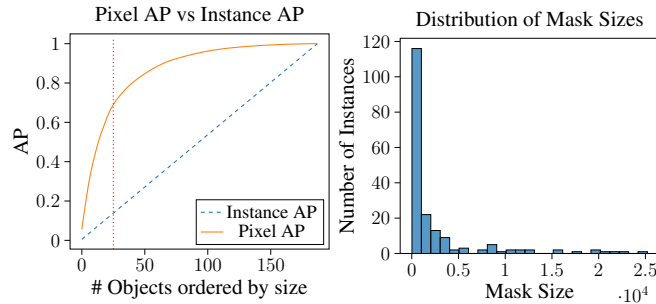


Fig. 4: Pixel-level vs Instance-level AP. The commonly used pixel-level AP metric favors large objects. We sort the 188 instances in Fishyscapes L&F in descending order of size and compute both pixel-level and instance-level metrics. By correctly predicting only 25 objects, we can achieve a pAP of 69.03%, while only getting an iAP of 13.82%. Since the majority of the objects in the dataset are small (as seen on the right), it is important to use a metric that does not favor large objects near to the ego-vehicle.

method creates instances from predicted uncertainty maps and then scores them using the intersection-over-union metric. However, it does not penalize grouping of objects or prediction that overlaps with an *ignore* region. We argue that it is still important to segment individual instances for downstream tasks, such as data mining or model introspection. If an anomalous object appears in front of another anomalous region, the computation of component-level metrics will ignore grouping (see Fig. 3). For example, in the Cityscapes dataset tunnels or bridges are labeled as a *void* class, resulting in high anomaly scores. Therefore, to evaluate unknown instance segmentation we use the common AP and AP50 from the Cityscapes evaluation protocol, here referred to as iAP and iAP50. We argue that the instance-level AP metric is a harder metric since objects of all sizes are treated the same, while the pixel-level AP is dominated by large objects. This is visually demonstrated in Figure 4. Unlike the default Cityscapes evaluation, we evaluate any instances larger than 10 pixels, using all objects in the Fishyscapes L&F dataset for evaluation.

Models and Training. We follow the setup by Nayal *et al.* [39] for training and fine-tuning the Mask2Former. Initially, we train the model on the Cityscapes dataset. Then, we fine-tune the model for 5000 iterations on Cityscapes images with pasted COCO

Table 1: Comparison to State-of-the-Art. We report pixel-level scores for UGainS on the Road Anomaly and the Fishyscapes Lost and Found datasets. As can be seen, we obtain state-of-the-art performances, both with and without additional OoD training data. (*: RbA scores based on our re-implementation.)

Method	OoD Data	Extra Network	Road Anomaly		FS L&F	
			pAP \uparrow	FPR $_{95}\downarrow$	pAP \uparrow	FPR $_{95}\downarrow$
MSP [18]	\times	\times	20.59	68.44	6.02	45.63
Mahalanobis [28]	\times	\times	22.85	59.20	27.83	30.17
SML [24]	\times	\times	25.82	49.74	36.55	14.53
GMMSeg [30]	\times	\times	57.65	44.34	50.03	12.55
SynthCP [50]	\times	\checkmark	24.86	64.69	6.54	45.95
RbA [39]	\times	\times	<u>78.45</u>	<u>11.83</u>	<u>60.96</u>	<u>10.63</u>
RbA*	\times	\times	74.78	17.83	59.51	11.34
UGainS (Ours)	\times	\checkmark	81.32	11.59	70.90	10.38
SynBoost [10]	\checkmark	\checkmark	38.21	64.75	60.58	31.02
Maximized Entropy [6]	\checkmark	\times	–	–	41.31	37.69
PEBAL [46]	\checkmark	\times	44.41	37.98	64.43	6.56
DenseHybrid [15]	\checkmark	\times	–	–	63.8	6.1
RbA [39]	\checkmark	\times	85.42	6.92	70.81	<u>6.30</u>
RbA*	\checkmark	\times	<u>87.12</u>	12.06	<u>75.61</u>	7.46
UGainS (Ours)	\checkmark	\checkmark	88.98	<u>10.42</u>	80.08	6.61

[32] objects as simulated anomalies. Differently to RbA, we fine-tune the full transformer decoder and use the entire COCO [32] dataset. We do not fine-tune or alter the SAM model.

4.1 Results

Comparison to State-of-the-art. To validate the overall performance of our approach, we first compare it on a pixel-level to existing anomaly segmentation methods. We report performance on the FS L&F and RoadAnomaly datasets according to the common protocol, evaluating both a model without fine-tuning, as well as a fine-tuned version. These are shown above and below the dashed line in Table 1, respectively. On both the RoadAnomaly and Fishyscapes Lost and Found datasets we outperform other approaches in the pixel-level average precision (pAP) metric, with and without the use of additional training data. We attribute this performance to strong object priors from our OoD instance prediction branch.

Lower and Upper Performance Bounds of UGainS. Since no other method has been properly evaluated for anomaly instance segmentation on the FS L&F or RoadAnomaly datasets using a standard average precision metric, we unfortunately cannot compare our method to other works, but we can create two baselines using the strong SAM model. For the first baseline, we can obtain instance predictions using the perfect sampling to define an upper boundary. For the second, we can get instance predictions from SAM based on the simple dense grid sampling approach to get instance masks.

Table 2: Upper and lower bound comparison. We show how our method compares to two important baselines, sampling points in an oracle way from the ground truth and simply using dense grid sampling as proposed by SAM. As expected we cannot match the oracle sampling, however, we do outperform the simple dense sampling strategy significantly.

Method	N samples	GT Data	Fishyscapes L&F			
			pAP \uparrow	FPR $_{95}\downarrow$	iAP \uparrow	iAP50 \uparrow
Center Sampling (Oracle)	1 per object	✓	8.39	100.0	49.94	77.56
+ RbA[39] scoring	1 per object	✓	70.46	48.55	36.18	60.60
FPS Sampling + RbA Scoring	50 points	✗	63.70	99.91	6.12	9.20
+ NMS	50 points	✗	73.32	99.54	29.75	48.35
+ RbA ensemble	50 points	✗	80.08	6.61		
SAM [26] dense + RbA scoring	1024 points	✗	73.85	5.15	17.02	24.60

These two baselines can be seen as an upper and lower bound on the performance of our approach.

The upper part of Table 2 shows the upper-bound performance when we use ground-truth annotations to sample a single point per anomaly object. Note that such an oracle setting can only be considered for finding theoretical bounds and should not be considered as an actual performance of our approach. While this setting achieves a strong instance AP of nearly 50%, the pAP score is low. To address this, instead of using the IoU scores predicted by SAM, we use the average of the underlying RbA scores. This significantly improves the pixel AP, but also reduces the instance-level AP (iAP) by assigning a lower score to anomalies that were not well captured by RbA.

The other end of the spectrum would be to use SAM with the default automatic mask generation module, which samples 32×32 points in a dense grid and predicts three mask for each point, followed by a series of post-processing steps [26]. Since this sampling approach generates masks for the majority of objects, masks require anomaly scoring. Similar to the upper-bound baseline, we assign the average uncertainty score of the pixels inside a mask, since SAM does not provide any type of anomaly score out of the box. While, such an approach can be used in practice, it achieves a significantly lower iAP compared to our approach, indicating that sampling proposal points for the mask prediction in an informed way is beneficial (see Tab. 2). In addition, while low scores of dense sampling could be addressed by sampling a denser grid, the computational complexity increases quadratically.

To compare our approach, we look at the individual steps and indicate how they compare to the two upper and lower bounds. In the simplest case, where we sample only 50 points from the RbA pixel-level predictions, we cannot match the dense predictions of SAM. However, since we typically sample multiple points from the same object, a significant improvement can be gained with a simple non-maximum suppression step, which is also applied to dense point sampling. At this point, we can obtain an iAP that is significantly better than dense sampling, although there is still a gap to the oracle performance that could be attributed to incorrect sampling. As a final step, we ensemble the average uncertainty scores under each anomaly mask with the initial RbA scores,

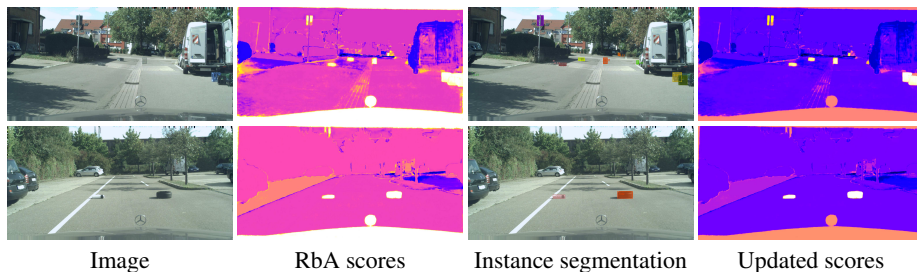


Fig. 5: Qualitative results on Fishyscapes L&F. The third column shows the instance segmentations generated using RbA predictions in the second column. Note the prominent anomalous objects in the refined per-pixel anomaly predictions in the fourth column.



Fig. 6: Qualitative results on CODA. Our model also performs well even under a larger domain shift. We are able to segment the typical unknown dangerous objects on the road, but it also reveals other instances for which the model has no proper semantic knowledge, *e.g.* bridges.

significantly increasing the final pixel-level AP of our approach, while also being able to predict meaningful instance segmentation of the anomalous objects in the scene.

Qualitative Results. We show qualitative results of our approach on the Fishyscapes L&F dataset in Figure 5. It can be seen how our approach segments instances from the pixel-level anomaly score predictions. In addition, UGainS can further refine the initial pixel-level predictions using instance predictions. In Figure 6 we show additional qualitative results on the more challenging CODA dataset [29]. Here we show that our method can capture interesting objects that are not well-covered in the training set. This ability is valuable for model introspection, providing insight into instances that the model may struggle with.

4.2 Ablations

In this section, we ablate some of our design choices. We focus on how pixel-level performance, the proposal sampling method, and the number of sampled points affect instance-level performance.

Evaluation of Other Per-Pixel Methods. Although we get the best performance with the RbA approach, we want to note that we can get instance segmentation and per-pixel score improvements for other anomaly segmentation methods as well. Table 3 shows the performance differences and anomaly instance segmentation performance of our approach for several other per-pixel anomaly segmentation methods. In all cases, we can improve the per-pixel AP by quite a large margin, and we observe that higher

Table 3: UGainS improvements of several anomaly segmentation methods. Our method notably improves pixel-level AP results of several methods in addition to providing anomaly instance segmentations.

Method	UGainS	Fishyscapes L&F				
		pAP \uparrow	FPR ₉₅ \downarrow	iAP \uparrow	iAP50 \uparrow	
Max Softmax [18]	\times	4.58	6.14 \rightarrow +5.73	40.59	—	—
	\checkmark	10.31		46.43	5.61	9.11
Mahalanobis [28]	\times	28.75	2.76 \rightarrow +1.51	9.52	—	—
	\checkmark	30.26		13.27	2.69	4.18
Maximized Entropy [6]	\times	41.31	6.36 \rightarrow +6.67	37.67	—	—
	\checkmark	47.98		37.86	12.16	18.72
DenseHybrid [15]	\times	69.78	1.63 \rightarrow +1.41	5.08	—	—
	\checkmark	71.19		5.77	27.49	44.07
RbA[39]	\times	75.61	4.16 \rightarrow +4.47	7.46	—	—
	\checkmark	80.08		6.61	29.75	48.35




Fig. 7: Visual comparison of different anomaly point sampling methods. Bright areas correspond to high RbA scores and black dots represent the sampled points. K-means estimates points mostly in the center of uncertain regions, while FPS generates most points at the boundary. Random sampling randomly selects points from the high uncertainty regions and thus often misses smaller objects.

pixel-level AP generally leads to higher instance-level AP. Since none of the other approaches achieve a similar or higher resulting instance-level AP, this indicates that all tested methods still have problems with small/far objects.

Sampling methods. In Table 4, we evaluate farthest point sampling (FPS), K-means, and random sampling from uncertainty regions. We observe that K-means predicts points closer to the centroids of the objects (see Fig. 7). In contrast, FPS generates proposals at boundaries and achieves higher recall. Due to the maximum recall achieved with the same number of points, we chose farthest point sampling as the best option. We also evaluate random sampling, but samples favor larger objects, resulting in a higher point-level precision, but a significantly lower recall compared to FPS and K-means.

Number of samples. In Table 5 we evaluate the impact of the number of sampled point proposal used to generating instance masks. Since the FS L&F dataset typically contains several instances per scene, we observe diminishing returns as the number of samples increases. With increased number of samples, most of the predicted masks end

Table 4: Performance impact of the proposal sampling methods. On the left we report the typical metrics, whereas on the right we additionally report metrics of the sampled points, *i.e.* how many points fall on an anomalous object (precision) and how many anomalous objects are sampled at least once (recall). FPS and K-means performs roughly on par, while random sampling performs worse due to a lower recall on the point level.

Sampling method	Full Metrics				Point Level	
	pAP \uparrow	FPR $_{95}\downarrow$	iAP \uparrow	iAP $_{50}\uparrow$	Precision	Recall
Random Sampling	79.81	5.97	25.82	40.45	57.22	60.11
K-means	79.92	6.60	29.04	46.23	47.51	70.74
FPS	80.08	6.61	29.75	48.35	35.30	71.81

Table 5: Performance impact of the number of sampled points. Considering there are often a few larger anomalous instances per frame, we see a diminishing return after sampling more than 50 points. While the recall of smaller instances might go up, we are also bound to segment false positive instances, thus reducing the precision.

# Points	pAP \uparrow	FPR $_{95}\downarrow$	iAP \uparrow	iAP $_{50}\uparrow$
10	76.73	5.89	26.81	41.33
20	78.83	<u>5.90</u>	30.82	49.14
50	80.08	6.61	<u>29.75</u>	<u>48.35</u>
100	78.96	7.01	27.16	45.39
200	78.99	7.09	25.77	43.82
300	78.55	7.16	24.44	41.98

up getting removed by non-maximum suppression, and the approach tends to predict multiple small boundary regions increasing the number of false positives.

5 Conclusion

We have proposed a novel approach for uncertainty-guided segmentation of anomalous instances. By combining anomaly object scores with initial per-pixel anomaly scores our approach shows consistent improvement on the task of anomaly segmentation. We show that UGainS achieves state-of-the-art average precision performance for pixel-level predictions on the Fishyscapes L&F and RoadAnomaly datasets. While anomaly segmentation on existing datasets seems to slowly saturate, we find that the more challenging task of anomaly instance segmentation is far from this point. Our method provides anomaly instance masks that can be used for model introspection and collection of rare instances and results can further be improved both through better anomaly point proposals and better instance segmentation.

Acknowledgements. This project was partially funded by the Ford University Alliance Program (project FA0347). Compute resources were granted by RWTH Aachen under project ‘supp0003’. We thank Ali Athar, George Lydakis, Idil Esen Zulfikar, Fabian Remper and the reviewers for helpful feedback.

References

1. Besnier, V., Bursuc, A., Picard, D., Briot, A.: Triggering Failures: Out-Of-Distribution detection by learning from local adversarial attacks in Semantic Segmentation. In: International Conference on Computer Vision (ICCV) (2021)
2. Bevandić, P., Krešo, I., Oršić, M., Šegvić, S.: Simultaneous Semantic Segmentation and Outlier Detection in Presence of Domain Shift. In: German Conference on Pattern Recognition (GCPR) (2019)
3. 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**(11), 3119–3135 (2021)
4. Cao, Y., Xu, X., Sun, C., Cheng, Y., Du, Z., Gao, L., Shen, W.: Segment Any Anomaly without Training via Hybrid Prompt Regularization. *arXiv preprint arXiv:2305.10724* (2023)
5. Chan, R., Lis, K., Uhlemeyer, S., Blum, H., Honari, S., Siegwart, R., Fua, P., Salzmann, M., Rottmann, M.: SegmentMeIfYouCan: A Benchmark for Anomaly Segmentation. In: Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (2021)
6. Chan, R., Rottmann, M., Gottschalk, H.: Entropy Maximization and Meta Classification for Out-Of-Distribution Detection in Semantic Segmentation. In: International Conference on Computer Vision (ICCV) (2021)
7. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In: European Conference on Computer Vision (ECCV) (2018)
8. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention Mask Transformer for Universal Image Segmentation. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
9. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The Cityscapes Dataset for Semantic Urban Scene Understanding. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
10. Di Biase, G., Blum, H., Siegwart, R., Cadena, C.: Pixel-wise Anomaly Detection in Complex Driving Scenes. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
11. Ding, Z., Wang, J., Tu, Z.: Open-Vocabulary Universal Image Segmentation with MaskCLIP. In: International Conference on Machine Learning (ICML) (2023)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In: International Conference on Learning Representations (ICLR) (2021)
13. Gasperini, S., Marcos-Ramiro, A., Schmidt, M., Navab, N., Busam, B., Tombari, F.: Holistic Segmentation. *arXiv preprint arXiv:2209.05407* (2022)
14. Grcić, M., Bevandić, P., Kalafatić, Z., Šegvić, S.: Dense anomaly detection by robust learning on synthetic negative data. *arXiv preprint arXiv:2112.12833* (2021)
15. Grcić, M., Bevandić, P., Šegvić, S.: DenseHybrid: Hybrid Anomaly Detection for Dense Open-set Recognition. In: European Conference on Computer Vision (ECCV) (2022)
16. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: International Conference on Machine Learning (ICML) (2017)
17. Gupta, K., Javed, S.A., Gandhi, V., Krishna, K.M.: MergeNet: A Deep Net Architecture for Small Obstacle Discovery. In: International Conference on Robotics and Automation (ICRA) (2018)
18. Hendrycks, D., Gimpel, K.: A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. In: International Conference on Learning Representations (ICLR) (2018)

19. Hendrycks, D., Mazeika, M., Dietterich, T.: Deep Anomaly Detection with Outlier Exposure. In: International Conference on Learning Representations (ICLR) (2019)
20. Hwang, J., Oh, S.W., Lee, J.Y., Han, B.: Exemplar-Based Open-Set Panoptic Segmentation Network. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
21. Ji, W., Li, J., Bi, Q., Liu, T., Li, W., Cheng, L.: Segment Anything Is Not Always Perfect: An Investigation of SAM on Different Real-world Applications. In: Conference on Computer Vision and Pattern Recognition Workshop (CVPR'W) (2023)
22. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y., Li, Z., Duerig, T.: Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In: International Conference on Machine Learning (ICML) (2021)
23. Jiang, C.M., Najibi, M., Qi, C.R., Zhou, Y., Anguelov, D.: Improving the Intra-class Long-tail in 3D Detection via Rare Example Mining. In: European Conference on Computer Vision (ECCV) (2022)
24. 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: International Conference on Computer Vision (ICCV) (2021)
25. Kendall, A., Gal, Y.: What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In: Neural Information Processing Systems (NIPS) (2017)
26. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment Anything. In: International Conference on Computer Vision (ICCV) (2023)
27. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In: Neural Information Processing Systems (NIPS) (2017)
28. Lee, K., Lee, K., Lee, H., Shin, J.: A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks. In: Neural Information Processing Systems (NIPS) (2018)
29. Li, K., Chen, K., Wang, H., Hong, L., Ye, C., Han, J., Chen, Y., Zhang, W., Xu, C., Yeung, D.Y., Liang, X., Li, Z., Xu, H.: CODA: A Real-World Road Corner Case Dataset for Object Detection in Autonomous Driving. In: European Conference on Computer Vision (ECCV) (2022)
30. Liang, C., Wang, W., Miao, J., Yang, Y.: GMMSeg: Gaussian Mixture based Generative Semantic Segmentation Models. In: Neural Information Processing Systems (NIPS) (2022)
31. Liang, S., Li, Y., Srikant, R.: Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks. In: International Conference on Learning Representations (ICLR) (2018)
32. Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft COCO: Common Objects in Context. In: European Conference on Computer Vision (ECCV) (2014)
33. Lis, K., Honari, S., Fua, P., Salzmann, M.: Detecting Road Obstacles by Erasing Them. arXiv preprint arXiv:2012.13633 (2021)
34. Lis, K., Nakka, K., Fua, P., Salzmann, M.: Detecting the Unexpected via Image Resynthesis. In: International Conference on Computer Vision (ICCV) (2019)
35. Liu, Y., Zulfikar, I.E., Luiten, J., Dave, A., Ramanan, D., Leibe, B., Ošep, A., Leal-Taixé, L.: Opening up Open-World Tracking. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
36. 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: Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
37. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment Anything in Medical Images. arXiv preprint arXiv:2304.12306 (2023)

38. Mukhoti, J., Gal, Y.: Evaluating Bayesian Deep Learning Methods for Semantic Segmentation. arXiv preprint arXiv:1811.12709 (2019)
39. Nayal, N., Yavuz, M., Henriques, J.F., Güney, F.: RbA: Segmenting Unknown Regions Rejected by All. In: International Conference on Computer Vision (ICCV) (2023)
40. Ohgushi, T., Horiguchi, K., Yamanaka, M.: Road Obstacle Detection Method Based on an Autoencoder with Semantic Segmentation. In: Asian Conference on Computer Vision (ACCV) (2020)
41. Pinggera, P., Ramos, S., Gehrig, S., Franke, U., Rother, C., Mester, R.: Lost and Found: Detecting Small Road Hazards for Self-Driving Vehicles. In: International Conference on Intelligent Robots and Systems (IROS) (2016)
42. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning Transferable Visual Models From Natural Language Supervision. In: International Conference on Machine Learning (ICML) (2021)
43. Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., Lu, J.: DenseCLIP: Language-Guided Dense Prediction with Context-Aware Prompting. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
44. Singh, A., Kamireddypalli, A., Gandhi, V., Krishna, K.M.: LiDAR guided Small obstacle Segmentation. In: International Conference on Intelligent Robots and Systems (IROS) (2020)
45. Tang, L., Xiao, H., Li, B.: Can SAM Segment Anything? When SAM Meets Camouflaged Object Detection. arXiv preprint arXiv:2304.04709 (2023)
46. Tian, Y., Liu, Y., Pang, G., Liu, F., Chen, Y., Carneiro, G.: Pixel-wise Energy-biased Abstention Learning for Anomaly Segmentation on Complex Urban Driving Scenes. In: European Conference on Computer Vision (ECCV) (2022)
47. Valdenegro-Toro, M.: I Find Your Lack of Uncertainty in Computer Vision Disturbing. In: Conference on Computer Vision and Pattern Recognition Workshop (CVPR'W) (2021)
48. Wong, K., Wang, S., Ren, M., Liang, M., Urtasun, R.: Identifying Unknown Instances for Autonomous Driving. In: Conference on Robot Learning (CoRL) (2019)
49. Wu, J., Zhang, Y., Fu, R., Fang, H., Liu, Y., Wang, Z., Xu, Y., Jin, Y.: Medical SAM Adapter: Adapting Segment Anything Model for Medical Image Segmentation. arXiv preprint arXiv:2304.12620 (2023)
50. Xia, Y., Zhang, Y., Liu, F., Shen, W., Yuille, A.: Synthesize then Compare: Detecting Failures and Anomalies for Semantic Segmentation. In: European Conference on Computer Vision (ECCV) (2020)
51. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. In: Neural Information Processing Systems (NIPS) (2021)
52. Xu, J., De Mello, S., Liu, S., Byeon, W., Breuel, T., Kautz, J., Wang, X.: GroupViT: Semantic Segmentation Emerges from Text Supervision. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
53. Xue, F., Ming, A., Zhou, Y.: Tiny Obstacle Discovery by Occlusion-Aware Multilayer Regression. *IEEE Transactions on Image Processing (TIP)* **29**, 9373–9386 (2020)
54. Yu, T., Feng, R., Feng, R., Liu, J., Jin, X., Zeng, W., Chen, Z.: Inpaint Anything: Segment Anything Meets Image Inpainting. arXiv preprint arXiv:2304.06790 (2023)
55. Zou, X., Dou, Z.Y., Yang, J., Gan, Z., Li, L., Li, C., Dai, X., Behl, H., Wang, J., Yuan, L., Peng, N., Wang, L., Lee, Y.J., Gao, J.: Generalized Decoding for Pixel, Image, and Language. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
56. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Wang, J., Wang, L., Gao, J., Lee, Y.J.: Segment Everything Everywhere All at Once. arXiv preprint arXiv:2304.06718 (2023)