

UMAD: Unsupervised Mask-Level Anomaly Detection for Autonomous Driving

Daniel Bogdoll^{*1,2}

bogdoll@fzi.de

Noël Ollick^{*1}

noel.ollick@student.kit.edu

Tim Joseph²

joseph@fzi.de

Svetlana Pavlitska^{1,2}

pavlitska@fzi.de

J. Marius Zöllner^{1,2}

zoellner@fzi.de

¹ Karlsruhe Institute of Technology (KIT)
Karlsruhe, Germany

² FZI Research Center for Information
Technology
Karlsruhe, Germany

Abstract

Dealing with atypical traffic scenarios remains a challenging task in autonomous driving. However, most anomaly detection approaches cannot be trained on raw sensor data but require exposure to outlier data and powerful semantic segmentation models trained in a supervised fashion. This limits the representation of normality to labeled data, which does not scale well. In this work, we revisit unsupervised anomaly detection and present UMAD, leveraging generative world models and unsupervised image segmentation. Our method outperforms state-of-the-art unsupervised anomaly detection.

1 Introduction

Although great achievements have been made in autonomous driving, reacting to the unknown remains a significant challenge [6, 24]. Heidecker et al. [24] categorize anomalies into the sensor, content, and temporal layer: Anomalies in the sensor layer are related to sensory abnormalities, anomalies in the content layer regard abnormalities in single observations, such as atypical objects, and the temporal layer considers behavioral anomalies in the context of multiple frames.

Classically, anomaly detection is based on highly specialized methods, focusing on the content layer [14, 39, 40]. However, a perpendicular line of work tries to learn a more general understanding of the world. Generative world models have shown promising results in autonomous driving [9, 21, 26, 27, 53]. They embed sensory data into latent states, reconstruct observations based on those, and predict action-conditioned future states. For anomaly detection, however, they have not been utilized yet [2]. In this paper, we address to which extent world models and unsupervised image segmentation can be used for anomaly detection in autonomous driving and, contrary to many prior anomaly detection models in this

domain [5, 16, 22], propose an unsupervised anomaly detection method which does not rely on outliers in the training data. We present **Unsupervised Mask-Level Anomaly Detection for Autonomous Driving (UMAD)**, leveraging generative world models and segmentation models. In the experimental setup for this paper, UMAD utilizes the multimodal world model MUVO [9], which was trained on data from the CARLA simulator [17]. For refined localization of anomalies, UMAD leverages masks which are generated with the unsupervised image segmentation approach U2Seg [41]. We also provide experimental results for the zero-shot segmentation model SAM [30]. UMAD improves upon the current SotA by achieving an FPR₉₅ reduction of 36.90% on the AnVoX benchmark, setting a new baseline in unsupervised anomaly detection for autonomous driving.

2 Related Work

Recent trends in anomaly detection have shown that utilizing semantic segmentation models and including Out-of-Distribution (OOD) data during training achieves close-to-perfect benchmark results [5, 13, 24]. However, we argue that normality should be learned on raw sensory data and thus in an unsupervised fashion. Including anomalies during training poses the risk of missing anomalies in a never-ending open-world setting, and utilizing supervised semantic segmentation [2, 16, 39], bounding boxes [18, 36], or language [47] limits the definition of normality to labeled training data, which does not scale well. Here, we revisit the field of unsupervised anomaly detection and explore mask-level approaches.

Unsupervised Anomaly Detection. While modeling uncertainty of models on computer vision tasks in an unsupervised way has already been addressed [19, 20, 23, 29], these models were not evaluated on anomaly detection benchmarks, but regarding their eligibility to model uncertainty of general computer vision tasks. Since anomaly detection is not only relevant in autonomous driving, there are also unsupervised anomaly detection methods in other domains. For example, Zhou et al. [55] have developed an anomaly detection model on retinal images, e.g., for detecting retinal diseases or lesions, and Wang et al. [50] have evaluated their anomaly detection model on the MVTec AD dataset [9] for industrial inspection. Similarly, self-supervised detection methods exist in such a setting [28, 46, 52]. Others used the MNIST [3, 25, 32] or CIFAR [25, 31, 49] datasets which contain images of only small sizes for their evaluation.

In anomaly detection in the surveillance setting, there is also a trend towards supervision requiring labeled training data [18, 36]. However, there are two recent unsupervised methods. Abati et al. [1] have developed a novelty detection model that uses a deep autoencoder in combination with an autoregressive parametric density estimator, using real world data with the UCSD Ped2 [12] and the ShanghaiTech [38] datasets. Similar to Abati et al. [1], Park et al. [43] trained MNAD on datasets with images from the real world [12, 37, 38], which partly contain semantic classes that can also be found in autonomous driving, e.g., pedestrians, bicycles, and cars. They compare the reconstruction of an autoencoder to the initial input image by using the L2 distance and the peak signal-to-noise ratio (PSNR) in order to calculate abnormality scores.

Mask-Level Anomaly Detection. A trend to improve anomaly detection methods is to use learned masks to generate instance-level detections. For detecting masks of anomalous instances in an image, the zero-shot Segment Anything Model (SAM) by Kirillov et al. [30] was quickly used for the localization of anomalies in images. Here, we give an overview of recent methods using segmentations during post-processing, as shown in Table 1.

Method	Supervision	Temporal	Multimodal	OOD Data	Extra Networks
SAA+ [10]	✓	—	✓	—	✓ [50, 52]
UGains [40]	✓	—	—	—	✓ [50, 59]
S2M [52]	✓	—	—	✓	✓ [50, 55, 55]
ClipSAM [53]	✓	—	✓	—	✓ [50, 49]
UMAD	—	✓	✓	—	✓ [9, 10]

Table 1: **Overview of mask-level anomaly detection methods.** The table shows methods that use segmentation masks for post-processing. Supervision refers to the necessity of labeled data during training. Temporality denotes the ability of a method to incorporate temporal context. Multimodal models utilize further modalities, such as text or lidar data, for anomaly detection. OOD data shows whether outliers were needed during training. Finally, all external networks are shown.

Segment Any Anomaly (SAA+) [10] utilizes pre-trained foundation models for mask-level anomaly detection without further training. They first employ Grounding DINO [54], which provides bounding boxes for regions defined by a prompt. To refine those box regions into masks, they utilize SAM [50]. Similarly, S2M [52] proposes bounding boxes that include anomalies, followed by SAM. Similar to many other anomaly detection models, they use outlier exposure during training. UGainS [40] uses the existing anomaly detection model Rejected by All (RbA) [59] in combination with SAM for localizing anomalous instances in the observation. Finally, ClipSAM [53] utilizes CLIP text and image encoders [42] to generate an initial anomaly mask and refine it with SAM.

Research Gap. In autonomous driving, recent trends have moved away from unsupervised anomaly detection [12, 16, 22], and benchmarks are saturated with near-perfect results [5, 13, 24]. While unsupervised anomaly detection methods from other domain methods are available, to the best of our knowledge, there exists no unsupervised anomaly detection model for autonomous driving. Similarly, the recent trend of mask-level anomaly detection methods works in a supervised manner. Thus, we see a clear need to revisit the field of unsupervised anomaly detection in order to use vast amounts of unlabeled data for training, as typically available in autonomous driving.

3 Method

As we have shown in Section 2 and Table 1, UMAD is the first unsupervised mask-level anomaly detection method in the context of autonomous driving, which means that UMAD can be trained purely based on unlabeled sensor recordings, that do not have to contain abnormal driving situations. As shown in Figure 1, first, UMAD takes multimodal data from several different sensors such as a camera and lidar sensor as input for a world model to reconstruct and predict future frames. Furthermore, semantic masks are derived from camera data. More details on the encoder-decoder architecture of the utilized world model MUVO can be found in [9].

For *visual differences*, a reconstruction of the current observation is compared to the accompanying sensor data frame based on multiple methodologies. For *temporal differences*, only multiple future predictions from the world model are compared. After a weighted fusion of the pixel-wise scores, the resulting anomaly map is refined based on the generated masks.

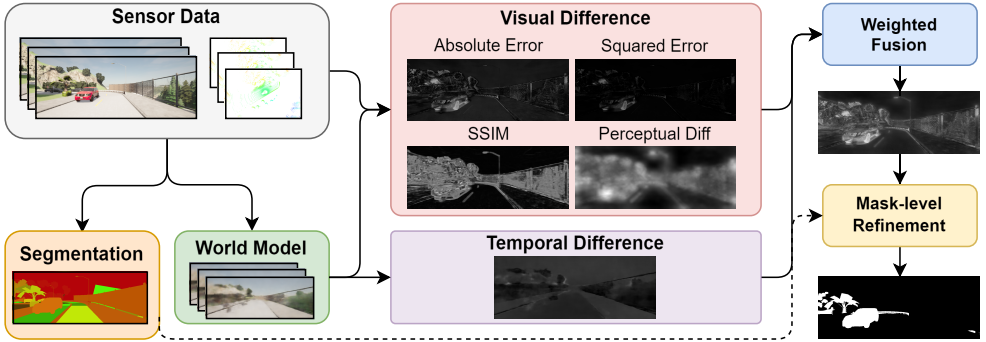


Figure 1: **Overview of UMAD.** First, multimodal sensor data is fed into a world model to reconstruct and predict frames, and semantic masks are derived from camera data. **Visual Differences** are used to compare a reconstruction of the current observation to the accompanying sensor data frame. The **Temporal Difference**, on the contrary, solely compares multiple future predictions from the world model. The pixel-wise scores are then fused and the resulting anomaly map is refined based on the generated masks.

UMAD first uses the world model to generate a reconstruction of the current frame. This reconstruction is then compared to the ground truth sensory data from the camera sensor of the autonomous vehicle. While UMAD only uses camera data, the world model is grounded and conditioned by further sensor modalities, planned actions, and the provided route. To compute *visual differences*, we employ several image comparison methods. The absolute error Δ_{ABS} and mean squared error Δ_{MSE} are calculated for each pixel individually and measure the differences in the r, g, b color channels of the reconstruction and the sensory image. In Eq. 1 and 2, \tilde{z}_x denotes a ground truth value and $\hat{z}_{\hat{x}}$ a predicted value.

$$\Delta_{ABS} = \frac{|\tilde{r}_x - \hat{r}_{\hat{x}}| + |\tilde{g}_x - \hat{g}_{\hat{x}}| + |\tilde{b}_x - \hat{b}_{\hat{x}}|}{3} \quad (1)$$

$$\Delta_{MSE} = \frac{(\tilde{r}_x - \hat{r}_{\hat{x}})^2 + (\tilde{g}_x - \hat{g}_{\hat{x}})^2 + (\tilde{b}_x - \hat{b}_{\hat{x}})^2}{3} \quad (2)$$

Contrary to this, the difference based on the Structural Similarity Index Δ_{SSIM} [6] compares two images based on their structure by utilizing batches of multiple proximate pixels. We compare sliding window patches between the ground truth x and the prediction \hat{x} . In Eq. 3, μ denotes means and σ (co)variances. The constants κ_1 and κ_2 are added for numerical stability [28, 6].

$$\Delta_{SSIM} = \frac{(2\mu_x\mu_{\hat{x}} + \kappa_1)(2\sigma_{x\hat{x}} + \kappa_2)}{(\mu_x^2 + \mu_{\hat{x}}^2 + \kappa_1)(\sigma_x^2 + \sigma_{\hat{x}}^2 + \kappa_2)} \quad (3)$$

Finally, perceptual difference Δ_{PD} [6] is an image comparison method that leverages a pre-trained deep convolutional network to extract features and compare two images pixel-wise based on their content. Similar to Di Biase et al. [6], we utilize weights which are pre-trained on the ImageNet [15] dataset. In Eq. 4, F^i denotes the i -th layer of a VGG network, and M and N refer to elements and layers, respectively.

$$\Delta_{PD} = \sum_{i=1}^N \frac{1}{M_i} \|F^i(x) - F^i(\hat{x})\|_1 \quad (4)$$

For *temporal differences* Δ_{TD} , we compare multiple predictions of the world model to each other. The temporal difference is calculated by comparing prior predictions for the current time step to each other. For this, the mean of the absolute errors between n past predictions \hat{z}_{t-i} for time t and the current reconstruction \hat{z}_t is computed, as shown in Eq. 5.

$$\Delta_{TD} = \frac{1}{n} \left(\sum_{i=1}^n \Delta_{ABS}(\hat{z}_{t-i}, \hat{z}_t) \right) \quad (5)$$

All difference maps are then normalized and can be fused by assigning weights $w_i \in [0, 1]$ to each of them. While the resulting anomaly map assigns anomaly scores to each pixel in the image, it does not classify instances in an observation as anomalous. For this, we refine the scores with instance masks to generate mask-level anomaly scores. By utilizing an image segmentation approach for mask generation, UMAD iterates through each predicted mask and calculates average instance-wise anomaly scores. More details can be found in [42].

4 Experiments

While there are common anomaly benchmarks in the context of autonomous driving, such as Fishyscapes [6] or SegmentMeIfYouCan [43], they are limited to camera data and do not contain data on actions, e.g., steering wheel angle, or additional sensory data. Among existing anomaly detection benchmarks [8], the recent AnoVox anomaly detection benchmark [44] is the only benchmark containing multimodal sensory data and action data of the ego-vehicle.

Benchmark. The AnoVox benchmark [44] contains both static and temporal behavioral anomalies. Here, we only generate a subset with static anomalies, i.e., unexpected objects on the road. It includes images, lidar point clouds, routemaps, panoptic segmentation maps and was created using the CARLA simulator. For evaluation, we generated 16 abnormal driving scenarios with 200 frames each using the provided framework for generating a small-size dataset with anomalies comparable to current benchmarks. The scenarios take place in different towns under different weather conditions and contain static anomalies, e.g., an object or an animal standing on the street, as depicted in Figure 2. The dataloader for the world model samples each 10^{th} frame, i.e., every second.

Experimental Setup. UMAD requires both an unsupervised world model and an unsupervised segmentation model. Among all published world models [9, 21, 26, 27, 53], MUVO was the only one with code and weights publicly available during time of writing, thus we selected it for our experiments. MUVO is a multimodal world model that uses camera and lidar data and is capable of reconstructing observations in both spaces. Available MUVO weights were trained on a large dataset which was created using the CARLA simulator [47]. It was trained on different driving scenarios in different towns, under different weather conditions, and at different times of the day. The training dataset of MUVO does not contain anomalies and thus establishes the baseline for typical behavior in the context of anomaly detection. Since AnoVox was also generated with CARLA, retraining MUVO for our approach was therefore not necessary.

For image segmentation, all prior works shown in Table 1 utilized the Segment Anything Model [40]. However, SAM was trained in a supervised manner, limiting the use of large-scale, unlabeled datasets as typically available in autonomous driving. Thus, we opted for U2Seg for unsupervised image segmentation [41]. U2Seg is an unsupervised image segmentation model that is capable of generating panoptic segmentation masks by using self-supervised learning and clustering. It would have been beneficial to train U2Seg exclusively on the target domain, but as it was trained on the entirety of ImageNet [42], we lacked the necessary resources and used a provided checkpoint.

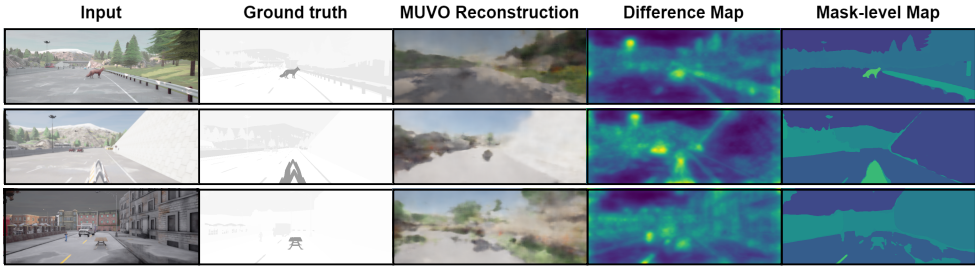


Figure 2: **Exemplary Detections.** The first columns show the input image and the corresponding ground truth. MUVO reconstructions are utilized to generate difference maps, which are finally refined to mask-level maps. Masks are generated by the unsupervised segmentation model U2Seg. The first two rows show positive cases, while the third row shows a failure case.

Baseline. As described in Section 2, there are only two recent anomaly detection models that are trained in a fully unsupervised manner. While both models demonstrate similar performances, Abati et al. [43] do only provide inference, but not training code for their approach. Thus, we decided to evaluate our approach against MNAD [43]. MNAD provides code both for prediction and reconstruction tasks but focuses on frame-wise evaluations. We were able to reproduce the experimental results of Park et al. [43].

For our evaluation, we trained MNAD on the dataset that was used to train MUVO but sampled each 100th frame from it, resulting in 2,725 frames. This ensures that MNAD was trained on images from the same towns, with the same driving conditions, and thus with the same semantic structure as MUVO. The sampling was necessary to prevent overfitting, as MUVO was trained on a much larger dataset. UCSD Ped2, however, which was originally used by Park et al. [43], only contains 2,550 images. The data sampling thus allows a dataset size which is comparable to the one used to train MNAD in the experimental setup by Park et al. [43]. Following Park et al. [43], we trained for 60 epochs. Contrary to our approach, MNAD only localizes anomalies as an intermediate step and uses additional metrics for their final frame-wise score. Based on these intermediate reconstructions, we use the L2 distance to compute pixel-wise anomaly scores.

5 Evaluation

For evaluating UMAD and MNAD, we computed the Average Precision (AP), the False Positive Rate at 95 % True Positive Rate (FPR₉₅), and the Area under the Receiver Operating

Characteristic curve (AUROC), as they are common metrics in anomaly detection benchmarks for autonomous driving [6, 13]. All results can be found in Table 2.

Ablation Studies. In order to better understand our method, we perform a set of ablation studies. First, next to utilizing U2Seg, we were interested in the possible performance gains of using SAM [30] or ground truth masks. SAM is a zero-shot image segmentation model that is also used by SotA anomaly detection models. While SAM was trained with labeled data, it performs well in the context of zero-shot inference. On the other hand, we also wanted to understand the effects of not refining our anomaly map with masks at all.

Second, rather than averaging all anomaly scores per mask, we were interested in whether picking the maximum value, inspired by Liu et al. [56], impacts the performance. On a similar note, we were also interested in picking only the mask with the highest anomaly score, neglecting the rest.

Experimental Results. Here, we present our findings on the performance of UMAD compared to the MNAD baseline, as well as our ablation studies. Since the visual differences and the temporal difference are normalized, they can be individually weighted and combined in order to form an anomaly map. This process is done in the Weighted Fusion Model. In the following, we also evaluate the impact of the single difference metrics and their combinations.

Since MNAD does not use masks, we first compare the pixel-wise L2 distance of MNAD to the similarly calculated squared error of our model on the raw pixel-wise output without masks. The experimental results indicate that utilizing world models instead of autoencoders is beneficial for anomaly detection in autonomous driving: With UMAD in the setting of using the *squared error* as visual difference, the AP is 7.03% and the FPR₉₅ 78.49%. The AP is thus 10.36% higher and the FPR₉₅ 12.41% lower than in the evaluation of MNAD. Even better results can be achieved when using the *perceptual difference* for visual difference: Here, UMAD achieves by far the highest average precision, lowest FPR₉₅, and highest AUROC in the pixel-wise setup without masks: AP is 70.49% higher and FPR₉₅ 63.27% lower than in the experimental results for MNAD. When using masks that are generated with the unsupervised image segmentation model U2Seg and the perceptual difference as visual difference, it is possible to achieve an AP that is 196.39% higher than the AP in the evaluation of MNAD, indicating that using masks for anomaly detection is highly beneficial. Using a combination of the mean squared error, the SSIM, the perceptual difference and the temporal difference, the FPR₉₅ is 36.90% lower than in the evaluation of MNAD.

When using the zero-shot image segmentation approach SAM, which is also used as an image segmentation approach in prior anomaly detection models, it is possible to further improve our experimental results. With the perceptual difference as visual difference in the setup, AP in this setup is 18.93%, FPR₉₅ is 52.77% smaller and AUROC 16.01% higher than in the respective results for MNAD. To evaluate the full potential of utilizing masks for anomaly detection, we furthermore evaluated UMAD with masks from a ground truth instance segmentation map. This setup achieved by far the best experimental results, again showing the huge potential of leveraging masks in anomaly detection: The best AP score with this experimental setup is 29.90%, the best FPR₉₅ is 16.93%, and the best AUROC is 83.18%.

In the prior experimental setups, the average anomaly score of the masks was used for evaluation. Interestingly, we found that the perceptual difference is not suitable for anomaly detection when assigning the maximum anomaly score to masks rather than their average score. Generally, however, substituting the average anomaly score per instance with the maximum score, did not achieve better results. Worst results are achieved when only con-

W_{ABS}	W_{MGE}	W_{SSIM}	W_{Pcr}	W_{temp}	$AP \uparrow$	$FPR_{0.5} \downarrow$	$AUROC \uparrow$	$AP \uparrow$	$FPR_{0.5} \downarrow$	$AUROC \uparrow$
					Ground truth			SAM		
1	0	0	0	0	17.68	35.56	65.23	13.72	50.58	65.16
0	1	0	0	0	19.05	38.92	63.61	13.77	52.22	64.93
0	0	1	0	0	19.77	21.26	79.03	11.43	46.79	68.26
0	0	0	1	0	29.90	16.93	83.18	18.93	42.32	71.88
0	0	0	0	1	11.41	52.70	49.15	7.11	69.26	47.72
0	0	0	0	0	<u>27.50</u>	<u>17.81</u>	<u>82.47</u>	<u>17.11</u>	44.01	70.83
0	0	0	0	0	26.21	18.16	82.07	16.02	44.55	70.88
0	0	0	0	0	25.52	20.67	79.73	<u>17.11</u>	<u>43.83</u>	<u>71.74</u>
0	0	0	0	0	18.85	22.88	77.73	12.85	46.44	70.20
0	0	0	0	0	25.28	18.53	81.83	14.30	45.18	69.85
0	0	0	0	0	24.34	19.39	81.27	14.74	44.87	69.95
0	0	0	0	0	22.28	21.43	79.05	16.15	45.28	70.88
0	0	0	0	0	17.25	24.15	76.92	12.60	48.12	68.42
0	0	0	0	0	23.47	19.04	81.34	15.55	44.35	71.12
0	0	0	0	0	22.38	19.71	80.74	14.52	45.02	70.21
5	5	5	5	5						
					U2Seg			Max. Value		
1	0	0	0	0	14.04	60.20	59.55	19.00	59.68	40.68
0	1	0	0	0	14.54	60.98	59.93	18.86	59.52	40.76
0	0	1	0	0	12.17	58.44	62.88	10.87	67.03	33.30
0	0	0	1	0	18.88	56.74	64.77	17.26	57.68	42.55
0	0	0	0	1	9.02	68.89	54.44	11.01	74.23	25.97
0	0	0	0	0	17.70	56.76	65.09	<u>20.97</u>	<u>52.01</u>	48.57
0	0	0	0	0	17.13	<u>56.73</u>	65.50	18.71	52.63	47.85
0	0	0	0	0	<u>17.99</u>	57.07	<u>65.47</u>	21.91	51.83	<u>48.49</u>
0	0	0	0	0	13.64	58.31	63.97	18.51	60.24	40.52
0	0	0	0	0	17.08	56.54	65.13	16.69	56.77	43.76
0	0	0	0	0	16.35	56.92	65.18	15.44	57.88	42.90
0	0	0	0	0	17.15	57.09	64.97	18.44	56.67	43.68
0	0	0	0	0	12.17	58.44	62.88	16.05	63.27	37.25
0	0	0	0	0	17.16	56.84	62.88	19.86	53.36	47.22
0	0	0	0	0	16.38	57.06	65.04	20.01	56.84	43.82
5	5	5	5	5						
					No Mask			Single Mask		
1	0	0	0	0	6.80	78.19	60.19	5.04	93.57	50.54
0	1	0	0	0	7.03	78.49	60.68	5.04	93.57	50.53
0	0	1	0	0	4.72	50.87	73.02	5.83	92.83	51.12
0	0	0	1	0	10.86	32.91	79.51	<u>10.40</u>	88.49	53.26
0	0	0	0	1	4.09	73.37	53.05	5.06	93.57	50.59
0	0	0	0	0	<u>9.51</u>	<u>37.37</u>	53.05	12.66	86.31	54.48
0	0	0	0	0	9.29	38.99	<u>78.70</u>	8.88	89.93	52.59
0	0	0	0	0	9.42	42.34	76.24	8.83	89.94	52.54
0	0	0	0	0	6.93	52.40	72.32	5.05	93.56	50.64
0	0	0	0	0	8.29	39.37	77.51	8.88	89.93	52.57
0	0	0	0	0	8.14	40.26	77.17	10.37	<u>88.48</u>	<u>53.35</u>
0	0	0	0	0	8.50	44.02	75.05	7.30	91.39	51.79
0	0	0	0	0	6.16	53.28	71.14	4.30	94.29	50.26
0	0	0	0	0	8.83	40.07	77.62	8.83	89.93	52.57
0	0	0	0	0	8.11	41.12	76.69	8.07	90.66	52.18
5	5	5	5	5						
MNAD []					6.37	89.61	61.96	—	—	—

Table 2: **Evaluation Results.** We show evaluation results for six settings of our model and MNAD with **best** and **second-best** results highlighted. The experimental results for a setup comprising a ground truth image segmentation map, segmentation maps generated with U2Seg and SAM, a setup where the instance-wise maximum anomaly score is selected instead of the average anomaly score, a setup without an image segmentation map, and a setup where only the instance with the highest anomaly score is left in the anomaly map are depicted. All evaluation metrics in %.

sidering the instance with the highest anomaly score. We found that often not the anomalous object, but a different object in the observation has the highest anomaly score. This then results in completely ignoring the abnormal object.

6 Conclusion

In this work, we presented UMAD, the first fully unsupervised anomaly detection model for autonomous driving which utilizes generative world models and is capable of combining anomaly scores and image segmentation approaches for masked anomaly detection. We find that utilizing world models in combination with image segmentation approaches is highly beneficial for anomaly detection in autonomous driving. Furthermore, we demonstrate that perceptual difference, compared to other approaches, is highly suitable for generating reconstruction errors in generative anomaly detection models. UMAD sets a new baseline in unsupervised anomaly detection for autonomous driving by achieving a FPR_{95} reduction of 36.90% on the challenging AnoVox benchmark.

Limitations and Outlook. Since the reconstruction quality of the MUVO world model was, under some circumstances, highly fluctuating and affecting the anomaly detection performance, we are interested in evaluating a more recent world model [21] in the future. Also, since we lacked the computational resources to train U2Seg on our target domain, a domain shift exists. Once unsupervised segmentation models become less compute-intensive, we are interested in ablating the effects of such a domain shift. Furthermore, the perceptual difference utilizes weights which are pre-trained on the ImageNet [15] dataset. Despite achieving promising results using the perceptual difference as visual difference, we are interested in evaluating whether further improvements can be achieved when the VGG network for the perceptual difference is trained on the same dataset which is used to train the underlying world model. Finally, we only performed experiments in a simulated environment. In the future, we want to apply the approach to real-world driving scenarios.

Acknowledgment

This work results from the just better DATA project supported by the German Federal Ministry for Economic Affairs and Climate Action (BMWK), grant number 19A23003H.

References

- [1] Davide Abati, Angelo Porrello, Simone Calderara, and Rita Cucchiara. Latent Space Autoregression for Novelty Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [2] Jan Ackermann, Christos Sakaridis, and Fisher Yu. Maskomaly:Zero-Shot Mask Anomaly Segmentation. In *British Machine Vision Conference (BMVC)*, 2023.
- [3] Jinwon An and Sungzoon Cho. Variational Autoencoder based Anomaly Detection using Reconstruction Probability. *Special Lecture on IE*, 2015.
- [4] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad — a comprehensive real-world dataset for unsupervised anomaly detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [5] Hermann Blum, Paul-Edouard Sarlin, Juan Nieto, Roland Siegwart, and Cesar Cadena. The Fishyscapes Benchmark: Measuring Blind Spots in Semantic Segmentation. *International Journal of Computer Vision (IJCV)*, 2021.
- [6] Daniel Bogdoll, Maximilian Nitsche, and J. Marius Zollner. Anomaly detection in autonomous driving: A survey. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2022.
- [7] Daniel Bogdoll, Lukas Bosch, Tim Joseph, Helen Gremmelmaier, Yitian Yang, and J. Marius Zöllner. Exploring the Potential of World Models for Anomaly Detection in Autonomous Driving. In *IEEE Symposium Series on Computational Intelligence (SSCI)*, 2023.
- [8] Daniel Bogdoll, Svenja Uhlemeyer, Kamil Kowol, and J. Marius Zöllner. Perception datasets for anomaly detection in autonomous driving: A survey. In *IEEE Intelligent Vehicles Symposium (IV)*, 2023.
- [9] Daniel Bogdoll, Yitian Yang, and J. Marius Zöllner. MUVO: A Multimodal Generative World Model for Autonomous Driving with Geometric Representations. *arXiv:2311.11762*, 2023.
- [10] Daniel Bogdoll, Iramm Hamdard, Lukas Namgyu Rößler, Felix Geisler, Muhammed Bayram, Felix Wang, Jan Imhof, Miguel de Campos, Anushervon Tabarov, Yitian Yang, Hanno Gottschalk, and J. Marius Zöllner. Anovox: A benchmark for multimodal anomaly detection in autonomous driving. *arXiv:2405.07865*, 2024.
- [11] Yunkang Cao, Xiaohao Xu, Chen Sun, Yuqi Cheng, Zongwei Du, Liang Gao, and Weiming Shen. Segment Any Anomaly without Training via Hybrid Prompt Regularization. *arXiv:2305.10724*, 2023.
- [12] Antoni B. Chan and Nuno Vasconcelos. Modeling, Clustering, and Segmenting Video with Mixtures of Dynamic Textures. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 30, 2008.

- [13] Robin Chan, Krzysztof Lis, Svenja Uhlemeyer, Hermann Blum, Sina Honari, Roland Siegwart, Pascal Fua, Mathieu Salzmann, and Matthias Rottmann. SegmentMeIfYouCan: A Benchmark for Anomaly Segmentation. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [14] Anja Delić, Matej Grcić, and Siniša Šegvić. Outlier detection by ensembling uncertainty with negative objectness. *arXiv:2402.15374*, 2024.
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [16] Giancarlo Di Biase, Hermann Blum, Roland Siegwart, and Cesar Cadena. Pixel-wise Anomaly Detection in Complex Driving Scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [17] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Conference on Robot Learning (CoRL)*, 2017.
- [18] Jianwu Fang, Jiahuan Qiao, Jie Bai, Hongkai Yu, and Jianru Xue. Traffic Accident Detection via Self-Supervised Consistency Learning in Driving Scenarios. *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, 23, 2022.
- [19] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In *International Conference on Machine Learning (ICML)*, 2016.
- [20] Yarin Gal, Jiri Hron, and Alex Kendall. Concrete Dropout. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [21] Shenyuan Gao, Jiazhi Yang, Li Chen, Kashyap Chitta, Yihang Qiu, Andreas Geiger, Jun Zhang, and Hongyang Li. Vista: A Generalizable Driving World Model with High Fidelity and Versatile Controllability. *arXiv:2405.17398*, 2024.
- [22] Matej Grcić, Petra Bevandić, and Siniša Šegvić. DenseHybrid: Hybrid Anomaly Detection for Dense Open-Set Recognition. In *European Conference on Computer Vision ECCV*, 2022.
- [23] Fredrik K. Gustafsson, Martin Danelljan, and Thomas B. Schon. Evaluating scalable bayesian deep learning methods for robust computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2020.
- [24] Florian Heidecker, Jasmin Breitenstein, Kevin Rösch, Jonas Löhdefink, Maarten Bieshaar, Christoph Stiller, Tim Fingscheidt, and Bernhard Sick. An Application-Driven Conceptualization of Corner Cases for Perception in Highly Automated Driving. In *IEEE Intelligent Vehicles Symposium (IV)*, 2021.
- [25] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations (ICLR)*, 2017.

- [26] Anthony Hu. *Neural World Models for Computer Vision*. PhD thesis, University of Cambridge, 2023.
- [27] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. GAIA-1: A Generative World Model for Autonomous Driving. *arXiv:2309.17080*, 2023.
- [28] Jieliang Jiang, Jiale Zhu, Muhammad Bilal, Yan Cui, Neeraj Kumar, Ruihan Dou, Feng Su, and Xiaolong Xu. Masked swin transformer unet for industrial anomaly detection. *IEEE Transactions on Industrial Informatics*, 19, 2022.
- [29] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [30] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment Anything. *arXiv:2304.02643*, 2023.
- [31] Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images, 2009.
- [32] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-Based Learning Applied to Document Recognition. In *Proceedings of the IEEE*, 1998.
- [33] Shengze Li, Jianjian Cao, Peng Ye, Yuhang Ding, Chongjun Tu, and Tao Chen. ClipSAM: CLIP and SAM Collaboration for Zero-Shot Anomaly Segmentation. *arXiv:2401.12665*, 2024.
- [34] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision (ECCV)*, 2024.
- [35] 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. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [36] Zhian Liu, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. A Hybrid Video Anomaly Detection Framework via Memory-Augmented Flow Reconstruction and Flow-Guided Frame Prediction. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [37] Cewu Lu, Jianping Shi, and Jiaya Jia. Abnormal Event Detection at 150 FPS in MATLAB. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2013.
- [38] Weixin Luo, Wen Liu, and Shenghua Gao. A Revisit of Sparse Coding Based Anomaly Detection in Stacked RNN Framework. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [39] Nazir Nayal, Misra Yavuz, João F. Henriques, and Fatma Güney. RbA: Segmenting Unknown Regions Rejected by All. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.

- [40] Alexey Nekrasov, Alexander Hermans, Lars Kuhnert, and Bastian Leibe. UGainS: Uncertainty Guided Anomaly Instance Segmentation. In *German Conference on Pattern Recognition (GCPR)*, 2023.
- [41] Dantong Niu, Xudong Wang, Xinyang Han, Long Lian, Roei Herzig, and Trevor Darrell. Unsupervised Universal Image Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [42] Noël Ollick. Camera-based Anomaly Detection with Generative World Models. Bachelor thesis, Karlsruhe Institute of Technology (KIT), 2024.
- [43] Hyunjong Park, Jongyoun Noh, and Bumsub Ham. Learning Memory-guided Normality for Anomaly Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [44] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning (ICML)*, 2021.
- [45] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2015.
- [46] Eli Schwartz, Assaf Arbelle, Leonid Karlinsky, Sivan Harary, Florian Scheidegger, Sivan Doherty, and Raja Giryes. MAEDAY: MAE for few-and zero-shot Anomaly-Detection. *Computer Vision and Image Understanding*, 241, 2024.
- [47] Beiwen Tian, Mingdao Liu, Huan-ang Gao, Pengfei Li, Hao Zhao, and Guyue Zhou. Unsupervised road anomaly detection with language anchors. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [48] 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 *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [49] Ha Son Vu, Daisuke Ueta, Kiyoshi Hashimoto, Kazuki Maeno, Sugiri Pranata, and Sheng Mei Shen. Anomaly Detection with Adversarial Dual Autoencoders. *arXiv:1902.06924*, 2019.
- [50] Lu Wang, Dongkai Zhang, Jiahao Guo, and Yuexing Han. Image anomaly detection using normal data only by latent space resampling. *Applied Sciences*, 2020.
- [51] Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing (TIP)*, 13, 2004.
- [52] Vitjan Zavrtnik, Matej Kristan, and Danijel Škočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112, 2021.

- [53] Lunjun Zhang, Yuwen Xiong, Ze Yang, Sergio Casas, Rui Hu, and Raquel Urtasun. Learning unsupervised world models for autonomous driving via discrete diffusion. *International Conference on Learning Representations (ICLR)*, 2024.
- [54] Wenjie Zhao, Jia Li, Xin Dong, Yu Xiang, and Yunhui Guo. Segment Every Out-of-Distribution Object. *arXiv:2311.16516*, 2023.
- [55] Kang Zhou, Yuting Xiao, Jianlong Yang, Jun Cheng, Wen Liu, Weixin Luo, Zaiwang Gu, Jiang Liu, and Shenghua Gao. Encoding structure-texture relation with p-net for anomaly detection in retinal images. In *European Conference on Computer Vision (ECCV)*, 2020.