

Bergische Universität Wuppertal

Fakultät für Mathematik und Naturwissenschaften

Institute of Mathematical Modelling, Analysis and Computational Mathematics (IMACM)

Preprint BUW-IMACM 22/05

Robin Chan, Svenja Uhlemeyer, Matthias Rottmann und Hanno Gottschalk

Detecting and Learning the Unknown in Semantic Segmentation

February 18, 2022

http://www.imacm.uni-wuppertal.de

Detecting and Learning the Unknown in Semantic Segmentation

Robin Chan¹, Svenja Uhlemeyer¹, Matthias Rottmann¹, and Hanno Gottschalk¹

Abstract. Semantic segmentation is a crucial component for perception in automated driving. Deep neural networks (DNNs) are commonly used for this task and they are usually trained on a closed set of object classes appearing in a closed operational domain. However, this is in contrast to the open world assumption in automated driving that DNNs are deployed to. Therefore, DNNs necessarily face data that they have never encountered previously, also known as *anomalies*, which are extremely safety-critical to properly cope with.

In this work, we first give an overview about anomalies from an information-theoretic perspective. Next, we review research in detecting semantically unknown objects in semantic segmentation. We demonstrate that training for high entropy responses on anomalous objects outperforms other recent methods, which is in line with our theoretical findings. Moreover, we examine a method to assess the occurrence frequency of anomalies in order to select anomaly types to include into a model's set of semantic categories. We demonstrate that these anomalies can then be learned in an unsupervised fashion, which is particularly suitable in online applications based on deep learning.

1 Introduction

Recent developments in deep learning have enabled scientists and practitioners to advance in a broad field of applications that were intractable before. To this end, deep neural networks (DNNs) are mostly employed which are usually trained in a supervised fashion with closed-world assumption. However, when those algorithms are deployed to real-world applications, e.g., artificial intelligence (AI) systems used for perception in automated driving, they often operate in an open-world setting where they have to face diversity of the real world. Consequently, DNNs are likely exposed to data which is "unknown" to them and therefore possibly beyond their capabilities to process. For this reason, having methods at hand, that indicate when a DNN is operating outside of its learned domain to seek for human intervention, is of utmost importance in safety-critical applications.

¹Interdisciplinary Center for Machine Learning and Data Analytics (IZMD), University of Wuppertal, Germany Email: {rchan, suhlemeyer, rottmann, hanno.gottschalk}@uni-wuppertal.de

A generic term for such a task is *anomaly* detection, which is generally defined as recognizing when something departs from what is regarded as normal or common. More precisely, identifying anomalous examples during inference, i.e., new examples that are "extreme" in some sense as they lie in low density regimes or even outside of the training data distribution, is commonly referred to as *out-of-distribution (OoD)* or *novelty* detection in deep learning terminology. The latter is in close connection to the task of identifying anomalous examples in training data, which is contrarily known as *outlier* detection; a term originating from classical statistics to determine whether observational data is polluted. Those outlined notions are often used interchangeably in deep learning literature. Throughout this work, we will stick to the general term anomaly and only specify when distinguishing is relevant.

For the purpose of anomaly detection, plenty of methods, ranging from classical statistical ones have been developed in the past. For the nowadays most challenging computer vision tasks tackled by deep learning, where both the model weights and output are of high dimension (in the millions), specific approaches to anomaly detection are mandatory.

Classical methods such as density estimation fail due to the curse of dimensionality. Early approaches identify outliers based on the distance to their neighbors [KNT00; RRS00], i.e., they are looking for sparse neighborhoods. Other methods consider relative densities to handle clusters of different densities, e.g., by comparing one instance either to its *k*-nearest neighbors [Bre+00] or using an ε -neighborhood as reference set [Pap+03]. However, the concept of neighborhoods becomes meaningless in high dimensions [AHK01]. More advanced approaches for high-dimensional data compute outlier degrees based on angles instead of distances [KSZ08] or even identify lower-dimensional subspaces [AY01; Kri+09].

In deep-learning-driven computer vision applications, novelties are typically regarded as more relevant than outliers. In semantic segmentation, i.e., pixel-level image classification, novelty detection may even refer to a number of sub-tasks. On the one hand, we might be concerned with the detection of semantically anomalous objects. This is also known as anomaly segmentation in the case of semantic segmentation. On the other hand, we also might be concerned with the detection of changed environmental conditions that are novel. The latter may be effects of a domain shift and include change in weather, time of day, seasonality, location and time. In this work, we focus only on semantically novel objects as anomalies.

In general, an important capability of AI systems is to identify the unknown. However, when striving for improved self-reflection capabilities, anomaly detection is not sufficient. Another important capability for real-world deployment of AI systems is to realize that some specific concept appears over and over again and potentially constitutes a new (or novel) object class. Incremental learning refers to the task of learning new classes, but, especially in semantic segmentation, mostly in a strictly supervised or semi-supervised fashion where data for the new class is labeled with ground truth [MZ19; Cer+20]. This is accompanied by an enormous data collection and annotation effort. In contrast to supervised incremental learning, humans may recognize a novelty of a given class that appears over and over again very well, such that in the end a single feedback might be sufficient to assign a name to a novel class. Ideally, this is accomplished in an unsupervised manner, like [HZ21] for image classification.

In this work, we first introduce anomaly detection from an information-based perspective in Section 2. We provide theoretical evidence that the entropy is a suitable quantity for anomaly detection, particularly in semantic segmentation. In Section 3, we review recent developments in the fields of anomaly detection and unsupervised learning of new classes. We give an overview on existing methods, both in the context of image classification and semantic segmentation. In this setting, we examine an approach to train semantic segmentation DNNs for high entropy on anomaly data in Section 4. We compare this approach against other established and recent state-of-the-art anomaly segmentation methods, and empirically show the effectiveness of entropy maximization in identifying unknown objects. Lastly, we investigate an unsupervised learning technique for novel object classes in Section 5. In this light, we further provide an outlook how the latter approach can be combined with entropy maximization to handle the unknown at run time in automated driving.

2 Anomaly Detection Using Information and Entropy

Anomaly detection is a common routine in any data analysis task. Before training a statistical model on data, the data should be investigated whether the underlying distribution generating the data is polluted by anomalies. In this context, anomalies can generally be understood as samples that do not fit into a distribution. Such anomalous samples can, e.g., be generated in the data recording process either by extreme observations, by errors in recording and transmission, or by the fusion of datasets that use different systems of units. Most common for the detection of anomalies in statistics is the inspection of maximum and minimum values for each feature, or simple univariate visualization via box-whisker plots or histograms.

More sophisticated techniques are applied in multivariate anomaly detection. Here, anomalous samples do not necessarily have to contain extreme values for single features, but rather an untypical combination of them. One of the application areas for multivariate anomaly detection is, e.g., fraud detection.

In both outlined cases, an anomaly $z \in \mathbb{R}^d$ can be qualified as an observation that occurs at a location of extremely low density of the underlying distribution p(z) or, equivalently, has an exceptionally high value of the information

$$I(\mathbf{z}) = -\log p(\mathbf{z}) \,. \tag{1}$$

Here, two problems occur: First, it is generally not specified what is considered as exceptionally high. Second, p(z) and thereby I(z) are generally unknown. Regarding the latter issue, however, the estimate $\hat{I}(z) = -\log \hat{p}(z)$ can be used which in turn relies on estimating $\hat{p}(z)$ from data associated to the probability density function p(z). Estimation approaches for $\hat{p}(z)$ can be distinguished between parametric and non-parametric ones.

The Mahalanobis distance [Mah36] is the best known parametric method for anomaly detection which is based on information of the multivariate normal distribution N. In fact, if $z \sim N(\mu, \Sigma)$ with mean $oldsymbol{\mu} \in \mathbb{R}^d$ and positive definite covariance matrix $oldsymbol{\Sigma} \in \mathbb{R}^{d imes d}$, then

$$I(\mathbf{z}) = -\log\left(\frac{1}{(2\pi)^{d/2}(\det \boldsymbol{\Sigma})^{1/2}}\exp\left(-\frac{1}{2}(\mathbf{z}-\boldsymbol{\mu})^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{z}-\boldsymbol{\mu})\right)\right)$$
(2)

$$= \frac{d}{2}\log(2\pi) + \frac{1}{2}\log(\det \Sigma) + \frac{1}{2}(\mathbf{z} - \boldsymbol{\mu})^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{z} - \boldsymbol{\mu}) = \frac{1}{2}(\mathsf{d}_{\boldsymbol{\Sigma}}(\mathbf{z}, \boldsymbol{\mu}))^{2} + c, \qquad (3)$$

where

$$\mathsf{d}_{\Sigma}(\mathbf{z},\boldsymbol{\mu}) := \sqrt{(\mathbf{z}-\boldsymbol{\mu})^{\mathsf{T}} \Sigma^{-1} (\mathbf{z}-\boldsymbol{\mu})} \tag{4}$$

denotes the Mahalanobis distance. The estimation $\hat{I}(\mathbf{z})$ is obtained by replacing $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ by the arithmetic mean $\hat{\boldsymbol{\mu}}$ and the empirical covariance matrix $\hat{\boldsymbol{\Sigma}}$, respectively, and likewise $d_{\boldsymbol{\Sigma}}(\mathbf{z}, \boldsymbol{\mu})$ by the empirical Mahalanobis distance $d_{\hat{\boldsymbol{\Sigma}}}(\mathbf{z}, \hat{\boldsymbol{\mu}})$.

In contrast, non-parametric techniques of anomaly detection rely on non-parametric techniques to estimate p(z). Here, a large variety of methods from histograms, kernel estimators and many others exist [Kle09]. We note, however, that the non-parametric estimation of densities and information generally suffers from the curse of dimensionality. To alleviate the latter issue in anomaly detection, estimation of information is often combined with techniques of dimensionality reduction, such as, e.g., principal component analysis [HTF07] or autoencoders [SY14].

When using non-linear dimensionality reduction with autoencoders, densities obtained in the latent space depend on the encoder and not only on the data itself. This points towards a general problem in anomaly detection. If p(z) is the density of a random quantity z and $z' = \phi(z)$ is an equivalent encoding of the data z using a bijective and differentiable mapping $\phi : \mathbb{R}^d \mapsto \mathbb{R}^d$, the *change of variables* formula [Rud87; Asa+21]

$$p(\mathbf{z}') = p(\mathbf{z}) \cdot |\det (\mathbf{J}_{\mathbf{z}'} \mathbf{z})| = p(\mathbf{z}) \cdot |\det (\mathbf{J}_{\mathbf{z}} \boldsymbol{\phi}^{-1}(\mathbf{z}))|$$
(5)

implies that the information of \mathbf{z}' is

$$I(\mathbf{z}') = -\log\left(p(\mathbf{z})\right) - \log\left(\left|\det\left(\mathbf{J}_{\mathbf{z}}\boldsymbol{\phi}^{-1}(\mathbf{z})\right)\right|\right) , \qquad (6)$$

where $\mathbf{J}_{\mathbf{z}}\phi^{-1}(\mathbf{z})$ denotes the Jacobian matrix of the inverse function ϕ^{-1} . Thus, whenever a high value of $I(\mathbf{z})$ indicates an anomaly, there always exists another equivalent representation of the data \mathbf{z}' , where the information $I(\mathbf{z}')$ is low. In other words, if \mathbf{z} is remote from other instances \mathbf{z}_j of a dataset and therefore considered an anomaly, there will be a transformation $\mathbf{z}' = \phi(\mathbf{z})$ that brings \mathbf{z}' right into the center of the data $\mathbf{z}'_j = \phi(\mathbf{z})_j$. In fact, via the Rosenblatt transformation [Ros52] any representation \mathbf{z} of the data can be expressed via a representation $\mathbf{z}' = \phi(\mathbf{z})$ where $I(\mathbf{z}')$ is constant over all data points. This stresses the importance to understand, that an anomaly always refers to probability *and* encoding of the data \mathbf{z} . This is true for both the original data and its approximated lower-dimensional representation.

As a side remark, autoencoders designed from neural networks have been very successfully applied in anomaly detection. Encoder and decoder networks possess the universal approximation property [Cyb89]. Furthermore, common training losses like the reconstruction error are invariant under a

change of the representation on latent spaces. Therefore, additional insights seem to be required to explain the empirical success of anomaly detection with autoencoders which is, however, not the scope of this work.

Another way of looking at the issue of anomaly detection in the context of different representations of same data is an explicit choice of a reference measure. This reference measure represents to which extent, or how likely, data is contaminated by potential anomalies. Suppose we can associate the probability density $p^{anom}(z)$ to the reference measure, then we can base anomaly detection on the quotient of densities, i.e., the odds $\frac{p(z)}{p^{anom}(z)}$, and apply a threshold whenever this ratio is low or, equivalently, when the relative information

$$I^{\text{rel}}(\mathbf{z}) := -\log\left(\frac{p(\mathbf{z})}{p^{\text{anom}}(\mathbf{z})}\right) = I(\mathbf{z}) - I^{\text{anom}}(\mathbf{z})$$
(7)

is high. We note that the relative information is independent under changes of the representation $\mathbf{z}' = \phi(\mathbf{z})$ as the $-\log |\det(\mathbf{J}_{\mathbf{z}}\phi^{-1}(\mathbf{z}))|$ term from Equation (6) occurs once with positive sign in $I(\mathbf{z}')$ and once with negative sign in $-I^{\text{anom}}(\mathbf{z}')$ and therefore cancels. Thus, the choice of a reference measure and the choice of a representation for the data is largely equivalent.

In practical situations, $p^{anom}(z)$ is often represented by some data $\{z_i^{anom}\}_{i \in \mathcal{T}'}$ that are either simulated or drawn from some data source of known anomalies. A binary classifier $\hat{p}(anom|z)$ can then be trained on basis of proper data $\{z_i\}_{i \in \mathcal{T}}$ and anomalous data $\{z_i^{anom}\}_{i \in \mathcal{T}'}$. The assumed prior probability p(anom) for anomalies, i.e., the degree of contamination, acts as a threshold for the estimated odds. Equivalently, the estimate of the relative information

$$\hat{I}^{\text{rel}}(\mathbf{z}) = -\log\left(\frac{\hat{p}(\mathbf{z})}{\hat{p}^{\text{anom}}(\mathbf{z})}\right) \stackrel{\text{Bayes' Theorem}}{=} -\log\left(\frac{\hat{p}(\text{non-anom}|\mathbf{z})p(\mathbf{z})}{p(\text{non-anom})} \cdot \frac{p(\text{anom})}{\hat{p}(\text{anom}|\mathbf{z})p(\mathbf{z})}\right)$$
(8)

$$= -\log\left(\frac{1 - p(anom|\mathbf{z})}{\hat{p}(anom|\mathbf{z})} \cdot \frac{p(anom)}{1 - p(anom)}\right)$$
(9)

$$= -\log\left(\frac{1 - \hat{p}(anom|\mathbf{z})}{\hat{p}(anom|\mathbf{z})}\right) - \log\left(\frac{p(anom)}{1 - p(anom)}\right)$$
(10)

$$= -\log\left(\frac{1-\hat{p}(anom|\mathbf{z})}{\hat{p}(anom|\mathbf{z})}\right) + c \tag{11}$$

with the prior log-odds $c = -\log\left(\frac{p(anom)}{1-p(anom)}\right)$ being a parameter controlling the threshold for the binary classifier $\hat{p}(anom|z)$.

If specifying what is an exceptionally high value for the information $I(\mathbf{z})$ or relative information $I^{rel}(\mathbf{z})$, the distinction between the detection of outliers in the training data and the detection of novelties during inference has to be taken into account. In outlier detection, observations, which have high (relative) information but which are in agreement with the extreme value of the (relative) information

$$I^{\max} = \max_{i \in \mathcal{T}} I(\mathbf{z}_i) \text{ or } I^{\operatorname{rel}\max} = \max_{i \in \mathcal{T}} I^{\operatorname{rel}}(\mathbf{z}_i),$$
(12)

are usually intentionally not eliminated. An outlier z for the level of significance $0 < \alpha < 1$ can then be detected using the condition

$$P_{\{\mathbf{z}_i\}_{i\in\mathcal{T}}}(I^{\max} > I(\mathbf{z})) \le \alpha \text{ or } P_{\{\mathbf{z}_i\}_{i\in\mathcal{T}}}(I^{\operatorname{rel}\max} > I^{\operatorname{rel}}(\mathbf{z})) \le \alpha.$$
(13)

Note again that the distribution of $I^{rel}(\mathbf{z}_j)$ has to be estimated to derive the associated distribution for the extreme values, see, e.g., [DF07], and also $I^{rel}(\mathbf{z})$ requires the estimation $\hat{p}(\mathbf{z})$ or $\hat{p}^{anom}(\mathbf{z})$. Therefore, a quantification of the epistemic uncertainty is essential for a proper outlier detection. Given the already mentioned problems of density estimation in high dimension, epistemic uncertainties may play a major role, unless a massive amount of data is available.

For the case of novelty detection taking place at inference, a comparison of the information of a single instance $I^{\text{rel}}(\mathbf{z})$ with the usual distribution of information $P_{\mathbf{z}_i}$ seems to be in order, which leads to the novelty criterion for level of significance $0 < \alpha < 1$

$$P_{\mathbf{z}_i}(I(\mathbf{z}_i) > I(\mathbf{z})) \le \alpha \text{ or } P_{\mathbf{z}_i}(I^{\text{rel}}(\mathbf{z}_i) > I^{\text{rel}}(\mathbf{z})) \le \alpha.$$
(14)

As a variant to this criterion, $I^{\text{rel}}(\mathbf{z}_i)$ could also be replaced by the extreme value statistics over the number of inferences alleviating the problem of generating false novelties by multiple testing. What has been stated on the necessity to quantify the epistemic uncertainty for the case of outlier detection equally applies for novelty detection.

While anomaly detection is generally seen as a sub-field of unsupervised learning, some specific effects occur in the case of novelty detection in supervised learning. During the phase of inference, the data $\mathbf{z} = (y, \mathbf{x})$ contain an unobserved component $y \in S$, which, e.g., represent the instance's label in a classification problem for the classes contained in S. Using the decomposition $p(\mathbf{z}) = p(y, \mathbf{x}) = p(y|\mathbf{x})p(\mathbf{x})$, one obtains the (relative) information from

$$I(\mathbf{z}) = I(y|\mathbf{x}) + I(\mathbf{x}), \text{ or } I^{\text{rel}}(\mathbf{z}) = I(y|\mathbf{x}) + I^{\text{rel}}(\mathbf{x}) - I^{\text{anom}}(y|\mathbf{x}),$$
(15)

where $I(y|\mathbf{x}) = -\log(p(y|\mathbf{x})), I^{\text{anom}}(y|\mathbf{x}) = -\log(p^{\text{anom}}(y|\mathbf{x}))$ is the conditional information on the right hand side. Often, for the data of the reference measure $p^{\text{anom}}(\mathbf{z})$, the labels are not contained in S. In this case, one uses a non-informative conditional distribution $p^{\text{anom}}(y|\mathbf{x}) = \frac{1}{|S|}$. If this is done, the last term of Equation (15) becomes a constant that can be integrated into a threshold parameter.

The (relative) information cannot be computed without knowing y. Therefore, the conditional expectation is used as an unbiased estimate, yielding the expected information

$$EI(\mathbf{x}) = \mathbb{E}_{y \sim p(y|\mathbf{x})}(I^{\text{rel}}(\mathbf{z})) = E(\mathbf{x}) + I^{\text{rel}}(\mathbf{x}) + b^{\text{rel}},$$
(16)

where $E(\mathbf{x}) = \sum_{y \in S} p(y|\mathbf{x})I(y|\mathbf{x})$ is the expected information, or entropy, of the conditional distribution $p(y|\mathbf{x})$ and b^{rel} is zero for the information and equal to $-\log(|\mathcal{S}|)$ for the relative information with non-informative conditional distribution $p^{\text{anom}}(y|\mathbf{x})$. Note that $E(\mathbf{x})$ is bounded by $\log(|\mathcal{S}|)$. Therefore, under normal circumstances, the term $I^{\text{rel}}(\mathbf{x})$ will outweigh $E(\mathbf{x})$ by far. However, in problems like semantic segmentation, each component of \mathbf{x} is assigned a label from \mathcal{S} . This implies solving $|\mathcal{I}|$ classification problems, where \mathcal{I} denotes the pixel space of x, thus the maximum value for $E(\mathbf{x})$ yields $|\mathcal{I}| \log(|\mathcal{S}|)$.

Therefore, the first term in Equation (16) contains significant contributions, especially in situations where $|\mathcal{I}|$ is large. The second term, $I^{\text{rel}}(\mathbf{x})$ loses importance under the hypothesis that the probability of the inputs \mathbf{x} does not vary greatly. Despite this hypothesis could be supported by fair sampling strategies, it requires further critical evaluation. But at least to a significant part, the expected information as an anomaly measure with regard to instance \mathbf{x} is given by a dispersion measure, namely the entropy of the conditional probability. As the entropy can be well estimated using a supervised machine learning approach to estimate $\hat{\mathbf{p}}(y|\mathbf{x})$ from the data $\{\mathbf{z}_j\}_{j\in\mathcal{T}}$, this part of the information is well accessible in contrast to $I^{\text{rel}}(\mathbf{x})$, which requires density estimation in high dimension.

Lastly in this section, let us give a remark on the role of anomaly data $\{\mathbf{z}_{j}^{\text{anom}}\}_{j\in\mathcal{T}'} = \{\mathbf{x}_{j}^{\text{anom}}\}_{j\in\mathcal{T}'}$. If such data is available, it is desirable to train the machine learning model $\hat{p}(y|\mathbf{x})$ to produce high values for $E(\mathbf{x}_{j}^{\text{anom}})$ so that the tractable part of the expected information $EI(\mathbf{x})$ shows good separation properties. This requirement can be inserted to the loss function, as it has been proposed in [HAB19; Hen+19] for classification. In fact, as the entropy $E(\mathbf{x})$ is maximized by the uniform (non-informative) label distribution $p(y|\mathbf{x}_{j}^{\text{anom}}) = 1/|\mathcal{S}|$, the aforementioned loss will favor this prediction on anomalous inputs $\{\mathbf{x}_{j}^{\text{anom}}\}_{j\in\mathcal{T}'}$. In this work, in Section 4, we will investigate this approach for the computer vision task of semantic segmentation, after having reviewed related work based on deep learning in the following section.

3 Related Work

After the introduction to anomaly detection from a theoretical point of view, we now turn to anomaly detection in deep learning. In this section, we review research in the direction of detecting and learning unknown objects in semantic segmentation.

3.1 Anomaly Detection in Semantic Segmentation

An emerging body of works explores the detection of anomalous inputs on image data, where the task is more commonly referred to as *anomaly* or *out-of-distribution* (*OoD*) *detection*. Anomaly detection was first tackled in the context of image classification by introducing post-processing techniques applied to softmax probabilities to adjust the confidence values produced by a classification model [HG17; Lee+18a; LLS18; HAB19; MH20]. These methods have proven to successfully lower confidence scores for anomalous inputs at image-level, which is why they were also adapted to anomaly detection in semantic segmentation [ACS19; Blu+19], i.e., to *anomaly segmentation* by treating each single pixel in an image as a potential anomaly. Although those methods represent good baselines, they usually do not generalize well to segmentation, e.g., due to the high prediction uncertainties at object boundaries. The latter problem can, however, be mitigated by using segment-wise prediction quality

estimates [Rot+20], an approach which has also demonstrated to indicate anomalous regions within an image [ORF20].

Recent works have proposed more dedicated solutions to anomaly segmentation. Among the resulting methods, many originate from uncertainty quantification. The intuition is that anomalous regions in an image correlate with high uncertainty. In this regard, early approaches estimate uncertainty using Bayesian deep learning, treating model parameters as distributions instead of point estimates [Mac92; Nea96]. Due to the computational complexity, approximations are mostly preferred in practice, which comprise, e.g., Monte-Carlo dropout [GG16], stochastic batch normalization [Ata+19], or an ensemble of neural networks [LPB17; GDS20]; with some of them also being extended to semantic segmentation in [BKC17; KG17; MG19]. Even when using approximations, Bayesian models still tend to be computationally expensive. Thus, they are not well suited to real time semantic segmentation which is required for safe automated driving.

This is why tackling anomaly segmentation with non-Bayesian methods are more favorable from a practitioner's point of view. Some approaches therefore include tuning a previously trained model to the task of anomaly detection, by either modifying its architecture or exploiting additional data. In [DT18], anomaly scores are learned by adding a separate branch to the neural network. In [HMD19; MH20] the network architecture is not changed but auxiliary outlier data, which is disjoint from the actual training data, is induced into the training process to learn anomalies. The latter idea motivated several works in anomaly segmentation [Blu+19; Bev+19; JRF20; CRG21]. Nonetheless, such models have to cope with multiple tasks, hence possibly leading to a performance loss with respect to the original semantic segmentation task [Van+21]. Moreover, when including outlier datasets in the training process, it cannot be guaranteed that the chosen outlier data is a good proxy for all possible anomalies.

Another recent line of works performs anomaly segmentation via generative models that reconstruct original input images. These methods assume that reconstructed images will better preserve the appearance of known image regions than that of unknown ones. Anomalous regions are then identified by means of pixel-wise discrepancies between the original and reconstructed image. Thus, such an approach is specifically designed to anomaly segmentation and has been extensively studied in [CM15; MVD17; Lis+19; Xia+20; Lis+21; Bia+21]. The main benefit of these approaches is that they do not require any OoD training data, allowing them to generalize to all possible anomalous objects. However, all these methods are limited by the integrated discrepancy module, i.e., the module that identifies relevant differences between the original and reconstructed image. In complex scenes, such as street scene images for automated driving, this might be a challenging task due to the open world setting.

Regarding the dataset landscape, only few anomaly segmentation datasets exist. The LostAndFound dataset [Pin+16] is a prominent example which contains anomalous objects in various streets in Germany while sharing the same setup as Cityscapes [Cor+16]. LostAndFound, however, considers children and bicycles as anomalies, even though they are part of the Cityscapes training set. This was filtered and refined in Fishyscapes [Blu+19]. Another anomaly segmentation dataset accompanies

the CAOS benchmark [Hen+20], which considers three object classes from BDD100k [Yu+20] as anomalies. Both, Fishyscapes and CAOS, try to mitigate low diversity by complementing their real images with synthetic data.

Efforts to provide anomalies in real images have been made in [Lis+19] by sourcing and annotating street scene images from the web and in [Lis+21; Sin+20] by capturing and annotating images with small objects placed on the road. Just recently, the datasets published alongside the SegmentMeIfY-ouCan benchmark [Cha+21] build upon those works, particularly contributing to broad diversity of anomalous street scenes as well as objects.

3.2 Incremental Learning in Semantic Segmentation

Building upon the detection of anomalies, training data can be enriched in order to learn novel classes. To avoid training from scratch, several approaches tackle the task of *incremental* or even *continuous learning*, which can be understood as adapting to continuously evolving environments. Besides learning novel classes, incremental learning also encompasses adapting to alternative tasks or other domains. A comprehensive framework to compare these different learning scenarios is provided in [VT19].

When learning novel classes, the primary issue incremental learning approaches face is a loss of the original performance on previously learned classes, that is commonly known as *catastrophic forget-ting* [MC89]. To overcome this problem, a model needs to be both, "stable" and "plastic", i.e., the model needs to retain its original knowledge while being able to adapt to new environments. The complexity of meeting these requirements at the same time is called the stability-plasticity-dilemma [AR05]. In this regard, proposed solution strategies can be separated into three categories, which are either based on architecture, regularization, or rehearsal. Most of these methods have been applied to image classification first.

Architecture strategies employ separate models for each sequential incremental learning task, combined with a selector to determine which model will be used for inference [Pol+01; Che+19; ARC20]. However, these approaches suffer from data imbalances, consequently standard classification algorithms tend to favor the majority class. Approaches to mitigate skewed data distributions are usually based on over- or undersampling. Another line of works, such as [Rus+16; RPR19], employ "growing" models, i.e., enlarging the model capacity by increasing the number of model parameters for more complex tasks. In [ACT17], the authors propose an automated approach to select the proper task-specific model at test time. More efficient approaches were introduced in [GK16; Yoo+18], that restrict the adaptation of parameters to relevant parts of the model in terms of the new task. The Self-Net [CME19] is made up of an autoencoder that learns low-dimensional representations of the models belonging to previously learned tasks. By that, retaining existing knowledge via approximating the old weights instead of saving them directly is accompanied with an implicit storage compression. The incremental adaptive deep model developed in [Yan+19] enables capacity scalability and sustainability by exploiting the fast convergence of shallow models at the initial stage and afterwards utilizing the power of deep representations gradually. Other procedures perform continuous learning, e.g., using a random-forest [Hu+19], an incrementally growing DNN, retaining a basic backbone [SAR20], or nerve pruning and synapse consolidation [Pen+21].

Regularization strategies can be further distinguished between weight regularization, i.e., measuring the importance of weights, and distillation, i.e., transferring a model's knowledge into another. The former identifies parameters with great impact on the original tasks that are suppressed to be updated. Elastic weight consolidation (EWC) [Kir+17] is one representative method, evaluating weight importance based on the Fisher information matrix, while the synaptic intelligence (SI) method [ZPG17] calculates the cumulative change of Euclidean distance after retraining the model. Both regularization methods were further enhanced, e.g., by combining them [Cha+18; AM19], or by including unlabeled data [Alj+18]. Another idea to maintain model stability was adapted in [Zen+21; Far+19], updating gradients based on orthogonal constraints. Bayesian neural networks are applied in [Lee+18b] to approximate a Gaussian distribution of the parameters from a single to a combined task.

Distillation is a regularization method, where the knowledge of an old model can be drawn into a new model to partly overcome catastrophic forgetting. Knowledge distillation, proposed in [HVD14], was originally invented to transfer knowledge from a complex into a simple model. The earliest approach, which applies knowledge distillation to incremental learning, is called learning without forgetting (LwF) [LH18]. A combination of knowledge distillation and EWC was proposed in [Sch+18]. Further approaches based on distillation loss are, e.g., [Jun+18; Yao+19; Kim+19; Lee+19].

Rehearsal or pseudo-rehearsal-based methods, which were already proposed in [Rob95], mitigate catastrophic forgetting by allowing the model to review old knowledge whenever it learns new tasks. While rehearsal-based methods retain a subset of the old training data, pseudo-rehearsal strategies construct a generator during retraining, which learns to produce pseudo-data as similar to the old training data as possible. Hence, they provide the advantages of rehearsal even if the previously learned information is unavailable. Methods reusing old data are, e.g., incremental classifier and representation learning (iCaRL) [Reb+17], which simultaneously learns classifiers and feature representation, or the method presented in [Cas+18], which proposes a representative memory. The bias correction (BiC) method [Wu+19] keeps old data in a similar manner, but handles the data imbalance differently. Most pseudo-rehearsal approaches include generative adversarial networks (GANs) [OOS17; Wu+18; Mel+19; Ost+19] or a variational autoencoder (VAE) [Shi+17]. The method presented in [Hou+18] combines distillation and retrospective (DR), whereby baseline approaches such as LwF are outperformed by a large margin.

Only few works exist, such as [TTA19; Kli+20; MZ21; URG22], that adapt incremental learning techniques to semantic segmentation. Most of them adjust knowledge distillation using only a small portion or even none of old data. One challenge of continuous learning for semantic segmentation is that images may contain unseen as well as known classes. Hence, annotations that are restricted to some task assign a great amount of pixels to a background class, exhibiting a semantic distribution shift. The authors of [Cer+20] provide a framework that mitigates biased predictions towards this background class. In general, existing works rely on supervision for incremental learning.

4 Anomaly Segmentation

The task of anomaly detection in the context of semantic segmentation, i.e., identifying anomalies at pixel-level, is commonly known as *anomaly segmentation*. For this task several approaches have been proposed that are either based on uncertainty quantification, generative models, or training strategies specifically tailored to anomaly detection. In this work, we will first review some of those wellestablished methods and, subsequently, report a performance comparison with respect to their capability of identifying anomalies. In particular, we will demonstrate empirically that entropy maximization yields great performance on this segmentation task, which is in accordance to the statement of the entropy's importance from the information-based perspective as presented in Section 2.

4.1 Methods

Let $\mathbf{x} \in \mathbb{I}^{H \times W \times 3}$, $\mathbb{I} = [0, 1]$, denote (normalized) color images of resolution $H \times W$. Feeding those images to a semantic segmentation network $\mathbf{F} : \mathbb{I}^{H \times W \times 3} \to \mathbb{R}^{H \times W \times S}$, the model produces pixel-wise class scores $\mathbf{y} = (y_{i,s})_{i \in \mathcal{I}, s \in \mathcal{S}} = \mathbf{F}(\mathbf{x}) \in \mathbb{R}^{H \times W \times S}$, with the set of pixel locations denoted by $\mathcal{I} = \{1, \ldots, H\} \times \{1, \ldots, W\}$ and the set of trained (hence known) classes denoted by $\mathcal{S} = \{1, \ldots, S\}$. The corresponding predicted segmentation mask is given by $\mathbf{m} = (m_i)_{i \in \mathcal{I}} \in \{1, \ldots, S\}^{H \times W}$, where for $m_i = \arg \max_{s \in \mathcal{S}} y_{i,s} \forall i \in \mathcal{I}$ the maximum a-posteriori probability principle is applied. Regarding the task of anomaly segmentation, the ultimate goal is then to obtain a score map $\mathbf{a} = (a_i)_{i \in \mathcal{I}} \in \mathbb{R}^{H \times W}$ that indicates the presence of an anomaly at each pixel location $i \in \mathcal{I}$ within image \mathbf{x} , i.e., the higher the score the more likely there should be an anomaly.

Each of the methods employed in this section provides such score maps. Their underlying segmentation networks (DeepLabV3+, [Che+18]) are all trained on Cityscapes [Cor+16], i.e., objects not included in the set of Cityscapes object classes are considered as anomalies since they have not been seen during training and thus are unknown. The anomaly detection methods, however, differ in the way how the scores are obtained, which is why we briefly introduce the different techniques in the following.

Maximum softmax probability: The most commonly-used baseline for anomaly detection at image level is thresholding at the maximum softmax probability (MSP) [HG17]. Therefore, this method assumes that anomalies are attached a low confidence or, equivalently, high uncertainty. Using MSP in anomaly segmentation, the score map is computed via

$$a_i = 1 - \max_{s \in \mathcal{S}} \operatorname{softmax}(\mathbf{y}_i) \quad \forall \ i \in \mathcal{I} .$$
(17)

ODIN: A simple extension to improve MSP is applying temperature scaling as well as adding perturbations, which is known as out-of-distribution detector for Neural networks (ODIN) [LLS18]. In more

detail, let $t \in \mathbb{R}_{>0}$ be a hyperparameter for temperature scaling and let $\varepsilon \in \mathbb{R}_{\geq 0}$ be a hyperparameter for the perturbation magnitude. Then, the input x is modified as

$$\tilde{\mathbf{x}} = (\tilde{x}_i)_{i \in \mathcal{I}} \quad \text{with} \quad \tilde{x}_i = x_i - \varepsilon \operatorname{sign} \left(-\frac{\partial}{\partial x_i} \log \max_{s \in \mathcal{S}} \operatorname{softmax} \left(\frac{\mathbf{y}_i}{t} \right) \right) \quad \forall i \in \mathcal{I} ,$$
 (18)

yielding prediction $\tilde{\mathbf{y}} = F(\tilde{\mathbf{x}})$ for which thresholding is applied at the MSP, i.e., the anomaly score map is given by

$$a_i = 1 - \max_{s \in \mathcal{S}} \operatorname{softmax}(\tilde{\mathbf{y}}_i) \quad \forall i \in \mathcal{I} .$$
(19)

Mahalanobis distance: This anomaly detection approach estimates how well latent features fit to those observed in the training data. Let (L - 1) denote the penultimate layer of a network F with L layers. In [Lee+18a] the authors have shown that training a softmax classifier fits a class-conditional Gaussian distribution for the output features f_{L-1} . Hence, under that assumption

$$P\left(\mathbf{y}_{i}^{(L-1)} \mid \overline{y}_{i,s} = 1\right) = N\left(\mathbf{y}_{i}^{(L-1)} \mid \boldsymbol{\mu}_{s}, \boldsymbol{\Sigma}_{s}\right) \quad \forall i \in \mathcal{I} ,$$

$$(20)$$

where $\mathbf{y}^{(L-1)} = \mathbf{f}_{L-1}(\mathbf{x}) \in \mathbb{R}^{H \times W \times C_{L-1}}$ denotes the feature map of the penultimate layer given input \mathbf{x} , and $\overline{\mathbf{y}}$ the corresponding one-hot encoded final target. The minimal Mahalanobis distance $d_{\Sigma_s}(\mathbf{x}, \boldsymbol{\mu}_s)$ is then an obvious choice for an anomaly score map

$$a_i = \min_{s \in \mathcal{S}} d_{\boldsymbol{\Sigma}_s}(\mathbf{x}, \boldsymbol{\mu}_s) = \min_{s \in \mathcal{S}} (\mathbf{y}_i^{(L-1)} - \boldsymbol{\mu}_s)^\mathsf{T} \boldsymbol{\Sigma}_s^{-1} (\mathbf{y}_i^{(L-1)} - \boldsymbol{\mu}_s) \quad \forall \ i \in \mathcal{I} ,$$
(21)

cf. Equation (2). Note that the class means $\mu_s \in \mathbb{R}^{C_{L-1}}$ and class covariances $\Sigma_s \in \mathbb{R}^{C_{L-1} \times C_{L-1}}$ are generally unknown, but can be estimated by means of the training dataset.

Monte-Carlo dropout: In semantic segmentation, Monte Carlo dropout represents the most prominent technique to approximate Bayesian neural networks. According to [MG19], (epistemic) uncertainty is measured as the mutual information which might serve as anomaly score map, i.e.,

$$a_{i} = -\sum_{s \in \mathcal{S}} \left(\frac{1}{R} \sum_{r \in \mathcal{R}} \mathbf{p}_{i,s}^{(r)} \right) \log \left(\frac{1}{R} \sum_{r \in \mathcal{R}} \mathbf{p}_{i,s}^{(r)} \right) - \frac{1}{R} \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}} \mathbf{p}_{i,s}^{(r)} \log \mathbf{p}_{i,s}^{(r)} \quad \forall i \in \mathcal{I},$$
(22)

with $\mathbf{p}_i^{(r)} = (p_{i,s}^{(r)})_{s \in S} = \mathbf{softmax}(\mathbf{y}_i^{(r)})$ in the sampling round $r \in \mathcal{R} = \{1, \dots R\}$. Typically, $8 \leq R \leq 12$.

Void classifier: Neural networks can be trained to output confidences for the presence of anomalies [DT18]. One approach in this context is adding an extra class to the set S of previously trained classes of a semantic segmentation network, which then also requires annotated anomaly data to learn from. To this end, the void class in Cityscapes is a popular choice as proxy for all possible anomaly data

[Blu+19], in particular if the segmentation model was originally trained on Cityscapes. Thus, the softmax output of the additional class s = S + 1 represents the anomaly score map, i.e.,

$$a_{i} = \operatorname{softmax}_{s=S+1}(\mathbf{y}_{i}') \quad \forall i \in \mathcal{I} ,$$

$$\mathbf{p}'(-) \quad \mathbf{p}' = \mathbf{p}'(S) \quad \mathbf{p}'(S) \quad$$

where $\mathbf{y}' = (\mathbf{y}'_i)_{i \in \mathcal{I}} = (y'_{i,s})_{i \in \mathcal{I}, s \in \{1, \dots, S+1\}} = \mathbf{F}'(\mathbf{x}), \mathbf{F}' : \mathbb{I}^{H \times W \times 3} \to \mathbb{R}^{H \times W \times (S+1)}$

Learned embedding density: Let $\mathbf{f}_{\ell}(\mathbf{x}) \in \mathbb{R}^{H_{\ell} \times W_{\ell} \times C_{\ell}}$ denote the feature map, or equivalently feature embedding, at layer $\ell \in \mathcal{L} = \{1, \ldots, L\}$ of a semantic segmentation network. By employing normalizing flows, the true distribution of features $\mathbf{p}(\mathbf{f}_{\ell}(\mathbf{x})) \in \mathbb{I}^{H_{\ell} \times W_{\ell}}$, where $\mathbf{x} \in \mathcal{X}^{\text{train}}$ is drawn from the training dataset, can be trained via maximum likelihood, i.e., normalizing flows learn to produce the approximation $\hat{\mathbf{p}}(\mathbf{f}_{\ell}(\mathbf{x})) \approx \mathbf{p}(\mathbf{f}_{\ell}(\mathbf{x}))$ [Blu+19]. At test time, the negative log-likelihood measures how well features of a test sample fit to the feature distribution observed in the training data, yielding the anomaly score map

$$\mathbf{a} = \mathbf{u}\mathbf{p}^{\text{lin}}\left(-\log \hat{\mathbf{p}}(\mathbf{f}_{\ell}(\mathbf{x}))\right) \quad (\log \text{ applies } \log \text{ element-wise}) \tag{24}$$

with $\mathbf{up}^{\text{lin}} : \mathbb{R}^{H_{\ell} \times W_{\ell}} \to \mathbb{R}^{H \times W}$ denoting (bi-)linear upsampling.

Image resynthesis: After obtaining the predicted segmentation mask $\mathbf{m} \in \{1, \ldots, S\}^{H \times W}$, $\mathbf{m} = (m_i)_{i \in \mathcal{I}}$, this output can be further processed by a generative model $\mathbf{G} : \{1, \ldots, S\}^{H \times W} \to \mathbb{I}^{H \times W \times 3}$ aiming to reconstruct the original input image, i.e., $\mathbf{x}' = \mathbf{G}(\mathbf{m}) \approx \mathbf{x}$. This process is also called image resynthesis, and the intuition is that reconstruction quality for anomalous objects is worse than for those on which the generative model is trained on. To determine pixel-wise anomalies, a discrepancy network [Lis+19] $\mathbf{D} : \{1, \ldots, S\}^{H \times W} \times \mathbb{I}^{H \times W \times 3} \times \mathbb{I}^{H \times W \times 3} \to \mathbb{R}^{H \times W}$ can then be employed, which classifies whether one pixel is anomalous or not, based on information provided by \mathbf{m}, \mathbf{x}' , and \mathbf{x} . Here, \mathbf{D} is trained on intentionally triggered classification mistakes that are produced by flipping classes on predicted segmentation masks. The anomaly score map is given by the output of the discrepancy network, i.e.,

$$\mathbf{a} = \mathbf{D}(\mathbf{m}, \mathbf{x}', \mathbf{x}) = \mathbf{D}(\mathbf{m}, \mathbf{G}(\mathbf{m}), \mathbf{x}) .$$
(25)

SynBoost: The image resynthesis approach is limited by the employed discrepancy module **D**. In [Bia+21], the authors proposed to extend the discrepancy network by incorporating further inputs based on uncertainty, such as the pixel-wise softmax entropy

$$H_i(\mathbf{x}) = -\sum_{s \in \mathcal{S}} \operatorname{softmax}_s(\mathbf{y}_i) \log(\operatorname{softmax}_s(\mathbf{y}_i)) \quad \forall i \in \mathcal{I},$$
(26)

and the pixel-wise softmax probability margin

$$M_{i}(\mathbf{x}) = 1 - \max_{s \in \mathcal{S}} \left(\operatorname{softmax}(\mathbf{y}_{i}) \right) + \max_{s \in \mathcal{S} \setminus \{m_{i}\}} \left(\operatorname{softmax}(\mathbf{y}_{i}) \right) \quad \forall i \in \mathcal{I} .$$
(27)

Furthermore, **D** is trained on anomaly data provided by the Cityscapes void class. Thus, the anomaly score map is given by

$$\mathbf{a} = \mathbf{D}(\mathbf{m}, \mathbf{x}', \mathbf{x}, \mathbf{H}(\mathbf{x}), \mathbf{M}(\mathbf{x})) = \mathbf{D}(\mathbf{m}, \mathbf{G}(\mathbf{m}), \mathbf{x}, \mathbf{H}(\mathbf{x}), \mathbf{M}(\mathbf{x})) .$$
(28)

with $\mathbf{H}(\mathbf{x}) = (H_i(\mathbf{x}))_{i \in \mathcal{I}}$ and $\mathbf{M}(\mathbf{x}) = (M_i(\mathbf{x}))_{i \in \mathcal{I}}$.

Entropy maximization: A desirable property of semantic segmentation networks is that they attach high prediction uncertainty to novel objects. To this end, the softmax entropy, see Equation (26), is one intuitive uncertainty measure. The segmentation network can be trained for high entropy on anomalous inputs via the multi-criteria training objective [CRG21]

$$J^{\text{total}} = (1 - \lambda) \mathbb{E}_{(\overline{\mathbf{y}}, \mathbf{x}) \sim \mathcal{D}} \left[J^{\text{CE}}(\mathbf{F}(\mathbf{x}), \overline{\mathbf{y}}) \right] + \lambda \mathbb{E}_{\mathbf{x} \sim \mathcal{D}^{\text{anom}}} \left[J^{\text{anom}}(\mathbf{F}(\mathbf{x})) \right],$$
(29)

where \mathcal{D} denotes non-anomaly training data (labels available) and $\mathcal{D}^{\text{anom}}$ denotes anomaly training data (no labels available). In this approach, the COCO dataset [Lin+14] represents a set of so-called *known unknowns*, which is used as proxy for $\mathcal{D}^{\text{anom}}$ with the aim to represent all possible anomaly data. Moreover, $\lambda \in \mathbb{I}$ is a hyperparameter controlling the impact of each single loss function on the overall loss J^{total} . For non-anomaly data, the loss function is chosen to be the commonly-used cross-entropy J^{CE} , while for anomaly data, i.e., for known unknowns, we have

$$J^{\text{anom}}(\mathbf{F}(\mathbf{x})) = -\frac{1}{H \cdot W} \sum_{i \in \mathcal{I}} \frac{1}{S} \sum_{s \in \mathcal{S}} \log \operatorname{softmax}_{s}(\mathbf{y}_{i}) , \quad \mathbf{x} \sim \mathcal{D}^{\text{anom}} .$$
(30)

Therefore, minimizing J^{anom} is equivalent to maximizing the softmax entropy since both reach their optimum when the softmax probabilities are uniformly distributed, i.e., $\operatorname{softmax}_s(\mathbf{y}_i) = \frac{1}{S} \forall s \in S, i \in \mathcal{I}$. After training, the anomaly score map is then given by the (normalized) softmax entropy

$$a_i = \frac{1}{\log S} H_i(\mathbf{x}) = -\frac{1}{\log S} \sum_{s \in S} \operatorname{softmax}_s(\mathbf{y}_i) \log(\operatorname{softmax}_s(\mathbf{y}_i)) \quad \forall i \in \mathcal{I}.$$
(31)

From an information-based point of view, the entropy contains significant contribution to the expected information. This particularly applies for instance predictions in semantic segmentation, which motivates the entropy maximization approach for the detection of unknown objects, cf. Section 2.

4.2 Evaluation and Comparison of Anomaly Segmentation Methods

Discriminating between anomaly and non-anomaly is essentially a binary classification problem. In order to evaluate the pixel-wise anomaly detection capability, we use the receiver operating characteristic (ROC) curve as well as the precision recall (PR) curve. While for the ROC curve the true positive rate is plotted against the false positive rate at varying thresholds, in the PR curve precision is plotted against recall at varying thresholds. Note that we consider anomalies as the positive class, i.e.,



Fig. 1. Qualitative comparison of anomaly score maps for one example image of RoadAnomaly21. Here, red indicates high anomaly scores while blue indicates low ones. The ground truth anomaly instance is highlighted by green contours. Note that the region of interest is restricted to the road, highlighted by red contours in the annotation.

correctly identified anomaly pixels are considered as true positive. In both curves, the degree of separability is then measured by the area under the curve (AUC), where better separability corresponds to a higher AUC.

The main difference between these two performance metrics is how they cope with class imbalance. While the ROC curve incorporates the number of true negatives (for the computation of the false positive rate), in PR curves true negatives are ignored and, consequently, more emphasis is put on finding the positive class. With the anomaly score maps as defined in Section 4.1, in our case, finding the positive class corresponds to identifying anomalies.

As evaluation datasets, we use LostAndFoundNoKnown [Blu+19] and RoadObstacle21 [Cha+21], which are both part of the public SegmentMeIfYouCan¹ anomaly segmentation benchmark. LostAnd-FoundNoKnown consists of 1043 road scene images where obstacles are placed on the road. This dataset is a subset of the prominent LostAndFound dataset [Pin+16] but considers only obstacles from object classes which are disjoint to those in the Cityscapes labels [Cor+16]. More precisely, images with humans and bicycles are removed such that the remaining obstacles in the dataset also represent anomalies to models trained on Cityscapes. Similar scenes can be found in RoadObstacle21. That dataset was published alongside the SegmentMeIfYouCan benchmark and contains 327 road obstacle scene images with diverse road surfaces as well as diverse types of anomalous objects. Both datasets restrict the region of interest to the road where anomalies appear. This task is extremely safety-critical as it is mandatory in automated driving to make sure that the drivable area is free of any hazard. All anomaly segmentation methods introduced in the preceding Section 4.1 are suited to be evaluated on these datasets. We provide a visual comparison of anomaly scores produced by the tested methods in Figure 1. We report numerical results in Table 1 and in the corresponding Figure 2.

¹www.segmentmeifyoucan.com

Detecting	and I	arnina	the	Unknown	in	Somantia	Sagmantation
Delecting	unu L	Jeurning	ine	Onknown	in	Semanne	Segmentation

	LostA	ndFoundNo	Known	RoadObstacle21				
Method	AuPRC ↑	AuROC \uparrow	$FPR_{95TPR}\downarrow$	AuPRC ↑	AuROC \uparrow	$FPR_{95TPR}\downarrow$		
Maximum Softmax	30.1	93.0	33.2	10.0	95.5	17.9		
ODIN	52.9	95.1	30.0	11.9	96.0	16.4		
Mahalanobis	55.0	97.5	12.9	19.5	95.1	21.7		
Monte Carlo Dropout	36.8	92.2	35.5	4.9	83.5	50.3		
Void classifier	4.8	79.5	47.0	10.4	89.7	41.5		
Embedding density	61.7	98.0	10.4	0.8	81.0	46.4		
Image resynthesis	42.7	96.4	17.4	37.5	98.6	4.7		
SynBoost	81.7	98.3	4.6	71.3	99.4	3.2		
Entropy maximization	77.9	98.0	9.7	76.0	99.7	1.3		

Table 1. Pixel-wise anomaly detection performance on the datasets LostAndFoundNoKnown and RoadObstacle21, respectively. The main evaluation metric represents the area under precision-recall curve (AuPRC). Moreover, the area under receiver operating characteristic (AuROC) and the false positive rate at a true positive rate of 95% (FPR_{95TPR}) are reported for further insights.

In general, we observe that anomaly detection methods originally designed for image classification, including MSP, ODIN and Mahalanobis, do not generalize well to anomaly segmentation. As the Mahalanobis distance is based on statistics of the Cityscapes dataset, the anomaly detection is likely to suffer from performance loss under domain shift. The same holds for Monte Carlo dropout and learned embedding density, particularly resulting in poor performance in RoadObstacle21, where various road surfaces are available. Therefore, those methods potentially act as domain shift classifier rather than as detector of unknown objects.

The detection methods based on autoencoders, namely image resynthesis and SynBoost, show to be better suited for the task of anomaly segmentation, clearly being superior to all the approaches that already have been discussed. Autoencoders are limited by their discrepancy module, and we observe that anomaly detection performance significantly benefits from incorporating uncertainty measures, as done by SynBoost. Only entropy maximization reaches similar anomaly segmentation performance, even outperforming SynBoost in RoadObstacle21. This again can be explained by the diversity of road surfaces, which detrimentally affects the discrepancy module.

As a final remark, we draw attention to the use of anomaly data. The void classifier follows the same intuition as entropy maximization by including known unknowns, but cannot reach nearly as good anomaly segmentation performance. We conclude that the COCO dataset is better suited as proxy for anomalous objects than the Cityscapes unlabeled objects. Moreover, the results of that method empirically demonstrate the impact of the entropy in anomaly segmentation, which is in accordance to the statement of the entropy's importance from the information perceptive described in Section 2.





Fig. 2. Receiver operating characteristic (left column) and precision recall (right column) curves for LostAndFoundNoKnowns (top row) and RoadObstacle21 (bottom row), respectively. Dashed red lines indicate the performance of random guessing, i.e., the "no-skill" baseline. The degree of separability between anomaly and non-anomaly is measured by the area under the curve.

Anomaly score / entropy	En	tropy m	aximiz	ation	Entropy maximization					
threshold		+ thre	sholdir	ıg	+ thre	esholdin	ıg + me	eta classifier		
$a_i \ge \tau, i \in \mathcal{I}$	FP↓	$\mathrm{FN}\downarrow$	$F_1\uparrow$	δ in % \downarrow	$FP\downarrow$	$\mathrm{FN}\downarrow$	$F_1\uparrow$	δ in % \downarrow		
$\tau = 0.30$	8,068	191	0.26	0.30	290	308	0.82	0.06		
$\tau = 0.40$	4,035	289	0.39	0.11	251	359	0.81	0.03		
$\tau = 0.50$	1,215	415	0.60	0.04	145	447	0.80	0.02		
$\tau = 0.60$	327	613	0.69	0.02	49	619	0.76	0.02		
$\tau = 0.70$	135	879	0.61	0.01	21	881	0.63	0.01		

Table 2. Detection errors at object level for LostAndFound anomalies at different anomaly score / entropy thresholds τ (to generate anomaly segmentation masks). The quantities false-positives (FP) and false-negatives (FN) are reported at segment level, with anomalies as positive class. The F₁ summarizes these quantities into an overall measure. By δ we denote the performance loss on the original task, which is the semantic segmentation of Cityscapes. In this context, we consider a performance loss of 1% as acceptable, particularly in regard of a significantly improved anomaly detection performance.

4.3 Combining Entropy Maximization and Meta Classification

Meta classification is the task of discriminating between a true positive prediction and a false positive prediction. For semantic segmentation, this idea was originally proposed in [Rot+20]. By means of hand-crafted metrics, which are based on dispersion measures, geometry features, or location information, all derived from softmax probabilities, meta classifiers have shown to reliably identify incorrect predictions at segment level. More precisely, connected components of pixels sharing the same class label are considered as segments in this context, and a false positive segment then corresponds to a segment-wise intersection-over-union (IoU) of zero.

The meta classification approach can straightforwardly be adapted to post-process anomaly segmentation masks. This seems particularly reasonable in combination with entropy maximization. Since entropy maximization generally increases the sensitivity towards predicting anomalies, it is possible that the entropy is also increased at pixels belonging to non-anomalous objects. In the latter case, this would yield false positive anomaly instance predictions, which, however, can be identified and discarded afterwards by meta classification. The concept of trading false-positive detection for anomaly detection performance is motivated by [Cha+20]. Moreover, meta classifiers are expected to considerably benefit from entropy maximization, since in the original work [Rot+20] the entropy as metric has already been observed to be well correlated to the segment-wise IoU.

In our experiments on LostAndFound [Pin+16], we employ a logistic regression as meta classifier that is applied as a post-processing step on top of softmax probabilities. We observe that the meta classifier is capable of reliably removing false-positive anomaly instance predictions, which in turn significantly improves detection performance of anomalous objects. The meta classification performance is reported in Table 2, a visual example is given in Figure 3. We note that meta classification



Fig. 3. Meta classification as quality rating of anomaly instance predictions. Before applying the meta classifier (left), the anomaly segmentation mask contains anomaly instance predictions (orange segments), with some false-positives on the road. Based on softmax probabilities, the meta classifier performs a prediction quality rating (middle, red corresponds to poor quality), which is then used to remove false positive anomaly instance predictions (right). Note that the region of interest is restricted to the road, where ground truth anomalous objects (or obstacles) are indicated by green contours.

is applied to segmentation masks as input. Therefore, the output of the combination of entropy maximization and meta classification does not yield pixel-wise anomaly scores to compare against the methods presented in Section 4.1.

The idea of meta classification can even be used to directly identify potential anomalous objects in the semantic segmentation mask, see [ORF20], which will be subject to discussion in the following section about unsupervised learning of unknown objects.

5 Discovering and Learning Novel Classes

If certain types of anomalies appear frequently, it might be reasonable to include them as additional learnable classes of the segmentation model [ORF20]. In this section, we examine an unsupervised learning method in order to further process anomaly predictions with the goal to produce labels corresponding to novel classes. Afterwards, we will perform an incremental learning approach to train a model on novel classes by means of the retrieved unsupervised labels. This overall procedure has originally been proposed in [URG22].

5.1 Unsupervised Identification and Segmentation of a Novel Class

Consider the dataset $\mathcal{D}^{\text{test}} \subseteq \mathcal{X}$ of unlabeled images $\mathbf{x} = (x_i)_{i \in \mathcal{I}} \in \mathbb{I}^{H \times W \times 3}$, along with a semantic segmentation network $\mathbf{F} : \mathbb{I}^{H \times W \times 3} \to \mathbb{R}^{H \times W \times S}$ trained on the set of classes $\mathcal{S} = \{1, \ldots, S\}$. Moreover, let $\mathbf{a} = (a_i)_{i \in \mathcal{I}} \in \mathbb{R}^{H \times W}$ denote a score map, as introduced in Section 4, which assigns the degree of anomaly to each pixel $i \in \mathcal{I}$ in image \mathbf{x} . In this section, we employ the unsupervised anomaly segmentation technique described as a three-step procedure in what follows. 1. Image embedding: Image retrieval methods are commonly applied to construct a database of images that are visually related to a given image. On that account, such methods must quantify visual similarities, i.e., measuring the discrepancy or "distance" between images. A simple idea is averaging over the pixel-wise differences. However, this approach is extremely sensitive towards data transformation such as rotation, variation in light, or different resolutions. More advanced approaches make use of visual descriptors that extract the elementary characteristics of the visual contents, e.g., color, shape, or texture. These methods are invariant to data transformation, i.e., they perform well in identifying images representing the same item. If we want to detect different instances of the same category, deep learning methods represent the state-of-the-art. In this regard, convolutional neural networks (CNNs) achieve very high accuracy in image classification tasks. These networks extract features of the same category result in similar feature vectors. We now adapt this idea to identify anomalies that belong to the same class.

Let $\mathcal{K}_{\mathbf{a}|\mathbf{x}}$ denote the set of connected components within $(a_i^{(\tau)})_{i\in\mathcal{I}}, a_i^{(\tau)} := \mathbf{1}_{\{a_i\geq\tau\}} \forall i \in \mathcal{I}$ for a given threshold $\tau \in \mathbb{R}$, after processing image \mathbf{x} . Furthermore, let $\mathcal{K} := \bigcup_{x\in\mathcal{X}} \mathcal{K}_{\mathbf{a}|\mathbf{x}}$ denote the set of all predicted anomaly components in $\mathcal{D}^{\text{test}}$. For each component $k \in \mathcal{K}_{\mathbf{a}|\mathbf{x}}$, we tailor the input \mathbf{x} to the image crop $\mathbf{x}^{(k)} = (x_i)_{i\in\mathcal{I}'}, \mathcal{I}' \subseteq \mathcal{I}$ by means of the bounding box around $k \in \mathcal{K}_{\mathbf{a}|\mathbf{x}}$. By feeding the crop $\mathbf{x}^{(k)}$ to an image classification network \mathbf{G} , we map $\mathbf{x}^{(k)}$ onto its feature vector $\mathbf{g}^{(k)} := \mathbf{G}_{L-1}(\mathbf{x}^{(k)}) \in \mathbb{R}^n, n \in \mathbb{N}$ for all $k \in \mathcal{K}$. Here, \mathbf{G}_{L-1} denotes the output of the penultimate layer of \mathbf{G} .

2. Dimensionality reduction: Feature vectors extracted by CNNs are usually very high-dimensional. This evokes several problems regarding the clustering of such data. The first issue is known as *curse of dimensionality*, i.e., the amount of required data explodes with increasing dimensionality. Furthermore, distance metrics become less precise. Dimensionality reduction approaches project the feature vectors onto a low-dimensional representation, either by feature elimination, selection, or extraction. The latter creates new independent features as a combination of the original vectors and can be further distinguished between linear and non-linear techniques. A linear feature extraction approach, named principal component analysis (PCA) [Pea01], aims at decorrelating the components of the vectors by a change of basis, such that they are mostly aligned along the first axes. Thereby, little information is lost if we drop the last components. A more recent non-linear method is t-distributed stochastic neighbor embedding (t-SNE) [MH08], which uses conditional probabilities representing pairwise similarities. Let us consider two feature vectors $g^{(k)}$, $g^{(k')}$ with $k, k' \in \mathcal{K}$ and let $p_{k|k'} \in \mathbb{I}$ denote their similarity under a Gaussian distribution. Employing a Student t-distribution with one degree of freedom in the low-dimensional space then provides a second probability $q_{k|k'} \in \mathbb{I}$. Hence, t-SNE aims at minimizing the following sum (or Kullback-Leibler divergence) [MH08]

$$\sum_{k \in \mathcal{K}} \sum_{k' \in \mathcal{K}} \mathbf{p}_{k|k'} \log\left(\frac{\mathbf{p}_{k|k'}}{\mathbf{q}_{k|k'}}\right)$$
(32)

using gradient descent. We first perform dimensionality reduction via PCA, which is then followed by t-SNE. In our experiments, we observe that this combination of methods improves the effectiveness of

mapping anomaly predictions onto a two-dimensional embedding space. Here, the embedding ideally creates neighborhoods of visually related anomalies.

3. Novelty segmentation: If anomalies of the same category are detected more frequently, they are expected to form a bigger cluster in the embedding space. Those clusters can be identified by employing algorithms such as density-based spatial clustering of applications with noise (DBSCAN) [Est+96]. This algorithm supports the idea of non-supervision since it does not require any information of the potential anomaly data, such as e.g., the number of clusters. To this end, DBSCAN divides data points into core points, border points, and noise, depending on the size of the neighborhood $\varepsilon \in [0, \infty)$ and the minimal number of a core point's neighbors $\delta \in \mathbb{N}$.

More precisely, let $\tilde{\mathbf{g}}^{(k)} \in \mathbb{R}^2$ denote the two-dimensional representation of $\mathbf{x}^{(k)}$. Then, $\tilde{\mathbf{g}}^{(k)}$ is considered as a core point, if the corresponding point-wise density $\rho(\tilde{\mathbf{g}}^{(k)}) := |\{\tilde{\mathbf{g}}^{(k')} : \|\tilde{\mathbf{g}}^{(k)} - \tilde{\mathbf{g}}^{(k')}\| < \varepsilon, k' \in \mathcal{K}\}| \geq \delta$, i.e., the ε -neighborhood of $\tilde{\mathbf{g}}^{(k)}$ contains at least δ points including itself. We denote the neighborhood of a core point by $\tilde{\mathbf{g}}^{(k)}$, which corresponds to a component $\hat{k} \in \mathcal{K}$, as $B_{\hat{k}} := \{\tilde{\mathbf{g}}^{(k')} : \|\tilde{\mathbf{g}}^{(\hat{k})} - \tilde{\mathbf{g}}^{(k')}\| < \varepsilon, k' \in \mathcal{K}\}$. If $\tilde{\mathbf{g}}^{(k)}$ is not a core point but belongs to a core point's neighborhood, we call it a border point. Otherwise, i.e., if $\tilde{\mathbf{g}}^{(k)}$ is neither a core point nor within a core point's neighborhood, we call it noise.

Finally, a cluster $C_j \subset \mathcal{K}, j \in \mathcal{J} := \{1, \dots, J\}$ of components is formed by merging overlapping neighborhoods B_k , yielding $J \in \mathbb{N}$ clusters in total. In other words, clusters are formed from connected core points and their neighborhoods' border points. Given $\rho(\tilde{\mathbf{g}}^{(k)})$, we can determine the cluster density of C_j , e.g.,

as the maximum
$$\max_{k \in C_j} \rho(\tilde{\mathbf{g}}^{(k)})$$
, or as the average $\frac{1}{|\mathcal{C}_j|} \sum_{k \in C_j} \rho(\tilde{\mathbf{g}}^{(k)})$

The cluster $C^* \subset K$, which is the cluster of highest density given a sufficient cluster size, is then selected to be further processed. To this end, let us consider the predicted segmentation mask $\mathbf{F}(\mathbf{x}) =$ $\mathbf{m} = (m_i)_{i \in \mathcal{I}}$, where $m_i = \arg \max_{s \in S} y_{i,s}, i \in \mathcal{I}$. The pseudo labels $\tilde{\mathbf{y}} = (\tilde{y}_i)_{i \in \mathcal{I}}$ for the originally unlabeled \mathbf{x} are then obtained by setting $\tilde{y}_i = S + 1$ if pixel location *i* belongs to a component $k \in C^*$, and $\tilde{y}_i = m_i$ otherwise.

5.2 Class-Incremental Learning

Let $\tilde{\mathcal{Y}}$ denote the set of pseudo labels, then the training data for some novel class S + 1 can be represented by $\mathcal{D}^{S+1} \subseteq \mathcal{D}^{\text{novel}} \times \tilde{\mathcal{Y}}$, where $\mathcal{D}^{\text{novel}}$ denotes the set of previously-unseen images containing novel classes. By extending the semantic segmentation network \mathbf{F} to $\mathbf{F}^+ : \mathbb{I}^{H \times W \times 3} \to \mathbb{R}^{H \times W \times (S+1)}$ and retraining \mathbf{F}^+ on \mathcal{D}^{S+1} , we perform incremental learning to add a novel and previously unknown class to the semantic space of \mathbf{F} .



Segmentation prediction of the initial model



Fig. 4. Predicted semantic segmentation mask of the initial model's prediction (left) and corresponding segment-wise quality estimation (right) for one example from the Cityscapes test split. Green color indicates a high segment-wise IoU, red color indicates a low one.

Regularization: Knowledge distillation is a subcategory of regularization strategies aiming to mitigate a catastrophic forgetting, i.e., these strategies try to mitigate performance loss on the previouslylearned classes $S = \{1, ..., S\}$ while learning the additional class S + 1. In [MZ19], the authors adapted incremental learning techniques to the task of semantic segmentation. Among others, they introduced the overall objective

$$J^{\text{total}}(\mathbf{x}, \tilde{\mathbf{y}}) = (1 - \lambda) J^{\text{CE}}(\mathbf{F}^+(\mathbf{x}), \tilde{\mathbf{y}}) + \lambda J^{\text{D}}(\mathbf{F}^+(\mathbf{x}), \mathbf{F}(\mathbf{x})), \quad \lambda \in \mathbb{I} ,$$
(33)

where $(\mathbf{x}, \tilde{\mathbf{y}}) \in \mathcal{D}^{S+1}$. Here, J^{CE} denotes the common cross-entropy loss over the enlarged set of class indices $\mathcal{S}^+ := \{1, \ldots, S+1\}$ and J^{D} the distillation loss. The latter loss is defined as

$$J^{\mathrm{D}}(\mathbf{F}^{+}(\mathbf{x}), \mathbf{F}(\mathbf{x})) := -\frac{1}{H \cdot W} \sum_{i \in \mathcal{I}} \sum_{s \in \mathcal{S}} \operatorname{softmax}_{s}(\mathbf{y}_{i}) \log(\operatorname{softmax}_{s}(\mathbf{y}_{i}^{+}))$$
(34)

with $\mathbf{y} = \mathbf{F}(\mathbf{x})$ and $\mathbf{y}^+ = \mathbf{F}^+(\mathbf{x})$. Knowledge distillation can be further improved by freezing the weights of the encoder part of \mathbf{F}^+ during the training procedure [MZ19].

Rehearsal: If the original training data $\mathcal{D}^{\text{train}} \subseteq \mathcal{X} \times \mathcal{Y}$ of network **F** is available, in incremental learning such data is usually re-integrated into the training set of the extended network \mathbf{F}^+ , i.e., the training samples are drawn from $\mathcal{D}^{\text{train}} \cup \mathcal{D}^{S+1}$. To save computational costs of training and to balance the amount of old and new training data, established methods, e.g., [Rob95], only use a subset of $\mathcal{D}^{\text{train}}$. This subset is typically obtained by randomly sampling a set from $\mathcal{D}^{\text{train}}$ that matches the size of $|\mathcal{D}^{S+1}|$.

In combination with knowledge distillation, rehearsal strategies can be employed to mitigate a loss of performance on classes that are related to the novel class. This issue may arise e.g., through visual similarity such as between classes like *bus* and *train*, or due to class affiliation as in the case of *bicycle* and *rider*. Relevant classes can be identified by their frequency of being predicted on the relabeled





Fig. 5. Relative frequency of old classes being predicted by the initial model on pixels that are assigned to the novel class. Thus, the subset of \mathcal{D}_{CS}^{train} included in the retraining should mainly involve bicycles, motorcycles, and cars.

pixels, i.e.,

$$\nu_s^{\text{tot}} := \sum_{(\mathbf{x}, \tilde{\mathbf{y}}) \in \mathcal{D}^{S+1}} \left| \{ i \in \mathcal{I} \mid m_i = s \land \tilde{y}_i = S+1 \} \right| \ \forall \ s \in \mathcal{S} ,$$
(35)

and hence

$$\nu_s^{\text{rel}} := \frac{\nu_s^{\text{tot}}}{\sum_{s' \in \mathcal{S}} \nu_{s'}^{\text{tot}}} \quad \forall \ s \in \mathcal{S} .$$
(36)

The subset of $\mathcal{D}^{\text{train}}$ is then randomly sampled under the constraint that there are at least $\nu_s^{\text{rel}} |\mathcal{D}^{S+1}|$ images containing the class s for all $s \in S$.

5.3 Experiments and Evaluation

In the following experiments, we will employ a DeepLabV3+ [Che+18] model with an underlying WiderResNet38 [Zhu+19] backbone for semantic segmentation. This network is initially trained on a set of 17 classes, which we will extend by a novel class. The already trained classes are the Cityscapes training classes except *pedestrian* and *rider*, i.e., we exclude any *human* in the training process of our initial semantic segmentation network \mathbf{F} .

The initial model was trained on the Cityscapes [Cor+16] training data. For the incremental learning process, we use a portion of those data and combine them with the generated disjoint training set \mathcal{D}^{S+1} containing previously unseen images and pseudo labels on novel objects. Here, the images from



Prediction of initial model

Novel class in validation image

Prediction of extended model

Fig. 6. Comparison of the predicted semantic segmentation masks before (left) and after (right) adapting the model to the novel *human* class (orange) for one example of the Cityscapes validation split (middle). Here, the novel components are highlighted in orange, green contours indicate the ground truth annotation of the novelty.

	metric	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	human	car	truck	snq	train	motorcycle	bicycle	mean excl. human	mean incl. human
e	IoU	97.34	80.63	88.91	47.24	51.03	52.90	55.44	66.66	89.95	56.29	93.76	00.00	90.61	69.66	76.90	70.35	24.45	54.57	68.63	64.82
for	Precision	98.35	89.39	92.80	74.57	66.76	72.68	75.04	86.22	93.60	77.66	96.38	00.00	92.97	80.23	88.59	83.33	28.57	59.30	79.79	75.36
مّ	Recall	98.96	89.16	95.50	56.32	68.41	66.02	67.98	74.61	95.85	67.17	97.18	00.00	97.27	84.09	85.35	81.87	62.92	87.24	80.94	76.44
	IoU	97.46	80.78	89.30	47.48	49.31	53.25	55.28	65.72	90.17	54.53	93.47	41.42	91.21	69.30	72.52	62.06	30.45	57.72	68.24	66.75
fteı	Precision	98.68	89.31	93.11	78.33	69.48	73.74	76.02	88.99	94.21	75.66	95.69	59.73	95.26	84.88	87.26	91.68	64.38	76.01	84.28	82.91
а –	Recall	98.75	89.43	95.62	54.67	62.95	65.70	66.96	71.54	95.46	66.13	97.57	57.48	95.45	79.06	81.11	65.76	36.61	70.57	76.08	75.05

Table 3. Evaluation of the Cityscapes validation split before and after incremental learning the novelty *human* (highlighted in gray) with knowledge distillation and rehearsal. The classes *pedestrian* and *rider* are aggregated to the novel class *human*. All other classes are treated as background, i.e., they are ignored during training, regarding the data from \mathcal{D}^{S+1} .

 \mathcal{D}^{S+1} are drawn from the Cityscapes test data. For evaluation purposes, we use the Cityscapes validation data. Hence, during the incremental learning process only known objects are presented to the model except humans and a few instances, such as the ego-car or mountains in an image background, belonging to the Cityscapes void category.

We use the idea of meta classification, similarly as introduced in Section 4.3, to rate the prediction quality of predicted semantic segmentation masks. Here, the meta task is used to estimate the segment-wise IoU first, see Figure 4, on which we apply thresholding (at $\tau = 0.5$) to determine potential anomalies, cf. [ORF20]. We employ gradient boosting as meta model, which achieves a coefficient of determination of $R^2 = 82.51\%$ in estimating the segment-wise IoU on the Cityscapes validation split.

In accordance to Section 2 and as already observed in Section 4.3, the softmax entropy is again one of the main metrics included in the meta model to identify anomalous predictions. Thus, the entropy shows to have great impact on meta classification performance, which, similarly, has also been observed in [Cha+20; CRG21].



Fig. 7. A comparison of anomaly scores obtained by meta classification (left) and entropy (right) on an image from RoadObstacle21. The dog on the image is the anomaly of interest (indicated by green contours), which would have been overlooked by meta classification but entirely detected by the entropy.

Given anomaly segmentation masks, we perform image embedding using the encoder of the image classification network DenseNet201 [Hua+17], that is pretrained on ImageNet [Den+09]. Next, we reduce the dimensionality of the resulting feature vectors to 50 via PCA and further to 2 by applying t-SNE. In [ORF20], a qualitative and quantitative evaluation of different embedding approaches is provided. Note that t-SNE is non-deterministic, i.e., we obtain slightly different embedding spaces for different runs. In our experiments, employing DBSCAN with parameters $\varepsilon = 2.5$ and $\delta = 15$ produces a *human*-cluster including 91 components from 76 different images. The most frequently predicted class of these components are *car*, *motorcycle*, and *bicycle* with $\nu_{11}^{rel} = 24.84\%$, $\nu_{15}^{rel} = 26.69\%$ and $\nu_{16}^{rel} = 33.53\%$, respectively, see Figure 5.

We train the extended model \mathbf{F}^+ as described in Section 5.2 for 70 epochs, weighting the loss functions in Equation (33) equally, i.e., $\lambda = 0.5$. The extended model is capable of retaining its initial knowledge by achieving an mIoU score of 68.24% on the old classes when evaluating on the Cityscapes validation data. This yields a marginal loss of only 0.39% compared to the initial model \mathbf{F} . At the same time, \mathbf{F}^+ predicts the novel human class with a class IoU of 41.42%, without a single annotated human instance in the training data \mathcal{D}^{S+1} . A visual example of the applied unsupervised novelty segmentation approach is provided in Figure 6, more details on the numerical evaluation is given in Table 3.

5.4 Outlook on Improving Unsupervised Learning of Novel Classes

In the preliminary experiments presented in this section, we have demonstrated that a semantic segmentation network can be extended by a novel class in an unsupervised fashion. As a basis to start, the investigated unsupervised learning approach requires anomaly segmentation masks. Currently, these are obtained by meta classification [Rot+20], which is, however, not a method specifically tailored for the task of anomaly segmentation. In other words, the obtained masks are possibly inaccurate. To be even more precise on the limitation of plain meta classification, this method is only able to find anomalies when the segmentation model produces a (false positive) object prediction on those anomalies. By design, meta classifiers cannot find overlooked instances, e.g., obstacles on the road that also have been classified as road. As an illustration of this issue, we refer to Figure 7.

Having now several anomaly segmentation methods at hand, which we, e.g., introduced in Section 4.1, it seems obvious to replace the underlying anomaly segmentation method by a more sophisticated one as future work. In this light, given the decent performance of the examined unsupervised learning method relying only on meta classification, combining this method with entropy maximization represents a particularly promising approach to further improve the presented novelty training, cf. again Section 4.3.

Conclusions

Semantic segmentation as a supervised learning task is typically performed by models that operate on a given set containing a fixed number of classes. This is in clear contrast to the open world scenarios to which practitioners contemplate the usage of segmentation models. There are important capabilities that standard segmentation models do not exhibit. Among them is the capability to know when they face an object of a class they have not learned – i.e., to perform anomaly segmentation – as well as the capability to realize that similar objects, presumably of the same (yet unknown) class, appear frequently and should be learned either as a new class or be attributed to an existing one. In this work, we have seen first promising results for two tasks, for anomaly segmentation as well as for the detection and unsupervised learning of new classes.

For anomaly segmentation, we considered a number of generic baseline methods stemming from image classification as well as some recent anomaly segmentation methods. Since the latter clearly outperforms the former, this stresses the need for the development of methods specifically designed for anomaly segmentation. With the entropy maximization method, we have demonstrated empirically as well as theoretically that good proxies in combination with training on anomaly examples for high entropy could be key to anomaly segmentation capabilities. Particularly on the challenging RoadObstacle21 dataset with diverse street scenarios, entropy maximization yields great performance, outperforming many other established methods. While there exists a moderate number of datasets for anomaly segmentation, there is clearly still the need of additional datasets. The number of possible unknown object classes not covered by these datasets is evidently enormous. Furthermore, also the vast variety of possible environmental conditions and further domain shifts that may occur, possibly also in combination with unknown objects, continuously demand their exploration.

For detection and unsupervised learning of new classes, we have demonstrated in preliminary experiments that a combination of well-established dimensionality reduction and clustering methods along with the advanced uncertainty quantification method for semantic segmentation called MetaSeg is well able to detect unknown classes of which objects appear relatively frequently in a given test set. Indeed, MetaSeg can also be used to define segmentation proposals for pseudo ground-truths of new classes, which can also be learned incrementally by the segmentation model. For the considered scenario of subsequently learning humans within the Cityscapes dataset, this approach yields an reasonable segmentation performance on the novel class without significantly losing performance on the original classes. This examined methodology may help to incorporate new classes into existing models with low human labeling effort. The necessity for this will occur repeatedly in future. An example are the electric scooters that have recently arisen in several metropolitan areas across the globe. This is an example for a global phenomenon. However, also local phenomena, such as boat trailers at the coast, could be of interest. Such classes could be initially incorporated into an existing model using the presented methodology in this work. Afterwards, the initial performance could be further improved with active learning approaches, such as presented in [Col+21], still requiring only a small amount of human labeling effort. It is also an open question, to which extent the proposed method can be used iteratively to improve the performance on a new class. Also for this track of research, the lack of data for pursuing that task is a limiting factor as of now.

Acknowledgments

The research leading to these results is funded by the German Federal Ministry for Economic Affairs and Climate Action within the projects "KI Absicherung - Safe AI for Automated Driving", grant no. 19A19005R, and "KI Delta Learning - Scalable AI for Automated Driving", grant no. 19A19013Q, respectively. The authors would like to thank the consortia for the successful cooperation. Supported by:



The authors also acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC).

- [ACS19] Matt Angus, Krzysztof Czarnecki, and Rick Salay. "Efficacy of Pixel-Level OOD Detection for Semantic Segmentation". In: *arXiv* (Nov. 2019), pp. 1–13. arXiv: 1911.02897 (cit. on p. 7).
- [ACT17] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. "Expert Gate: Lifelong Learning With a Network of Experts". In: *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition (CVPR). Honolulu, HI, USA, July 2017, pp. 7120–7129 (cit. on p. 9).
- [AHK01] Charu Aggarwal, Alexander Hinneburg, and Daniel Keim. "On the Surprising Behavior of Distance Metrics in High Dimensional Spaces". In: *Proceedings of the International Conference on Database Theory*. London, UK, Jan. 2001, pp. 420–434 (cit. on p. 2).
- [Alj+18] Rahaf Aljundi et al. "Memory Aware Synapses: Learning What (Not) to Forget". In: *arXiv* (Oct. 2018), pp. 1–24. arXiv: 1711.09601 (cit. on p. 10).
- [AM19] Mohammed Amer and Tomás Maul. "Reducing Catastrophic Forgetting in Modular Neural Networks by Dynamic Information Balancing". In: *arXiv* (Dec. 2019), pp. 1–15. arXiv: 1912.04508 (cit. on p. 10).
- [AR05] Wickliffe Abraham and Anthony Robins. "Memory Retention The Synaptic Stability Versus Plasticity Dilemma". In: *Trends in Neurosciences* 28.2 (2005), pp. 73–78 (cit. on p. 9).
- [ARC20] Shivang Agarwal, Ajita Rattani, and C. Ravindranath Chowdary. "AILearn: An Adaptive Incremental Learning Model for Spoof Fingerprint Detection". In: *arXiv* (Dec. 2020), pp. 1–16. arXiv: 2012.14639 (cit. on p. 9).
- [Asa+21] Hayk Asatryan et al. "A Convenient Infinite Dimensional Framework for Generative Adversarial Learning". In: *arXiv* (Jan. 2021), pp. 1–29. arXiv: 2011.12087 (cit. on p. 4).
- [Ata+19] Andrei Atanov et al. "Uncertainty Estimation via Stochastic Batch Normalization". In: Proceedings of the International Symposium on Neural Networks (ISNN). Moscow, Russia, July 2019, pp. 261–269 (cit. on p. 8).
- [AY01] Charu C. Aggarwal and Philip S. Yu. "Outlier Detection for High Dimensional Data".
 In: Proceedings of the ACM SIGMOD International Conference on Management of Data (MOD). May 2001, pp. 37–46 (cit. on p. 2).

- [Bev+19] Petra Bevandić et al. "Simultaneous Semantic Segmentation and Outlier Detection in Presence of Domain Shift". In: *Proceedings of the German Conference on Pattern Recognition (GCPR)*. Dortmund, Germany, Oct. 2019, pp. 33–47 (cit. on p. 8).
- [Bia+21] Giancarlo Di Biase et al. "Pixel-Wise Anomaly Detection in Complex Driving Scenes". In: *arXiv* (Mar. 2021), pp. 1–16. arXiv: 2103.05445 (cit. on pp. 8, 13).
- [BKC17] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. "Bayesian SegNet: Model Uncertainty in Deep Convolutional Encoder-Decoder Architectures for Scene Understanding". In: *Proceedings of the British Machine Vision Conference (BMVC)*. Sept. 2017, pp. 1–12 (cit. on p. 8).
- [Blu+19] Hermann Blum et al. "Fishyscapes: A Benchmark for Safe Semantic Segmentation in Autonomous Driving". In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops. Seoul, Korea, Oct. 2019, pp. 2403–2412 (cit. on pp. 7, 8, 13, 15).
- [Bre+00] Markus M. Breunig et al. "LOF: Identifying Density-Based Local Outliers". In: *ACM SIGMOD Record* 29.2 (2000), pp. 93–104 (cit. on p. 2).
- [Cas+18] Francisco M. Castro et al. "End-to-End Incremental Learning". In: *arXiv* (Sept. 2018), pp. 1–21. arXiv: 1807.09536 (cit. on p. 10).
- [Cer+20] Fabio Cermelli et al. "Modeling the Background for Incremental Learning in Semantic Segmentation". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. virtual conference, June 2020, pp. 9233–9242 (cit. on pp. 2, 10).
- [Cha+18] Arslan Chaudhry et al. "Riemannian Walk for Incremental Learning: Understanding Forgetting and Intransigence". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. Munich, Germany, Sept. 2018, pp. 556–572 (cit. on p. 10).
- [Cha+20] Robin Chan et al. "Controlled False Negative Reduction of Minority Classes in Semantic Segmentation". In: *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*. virtual conference, July 2020, pp. 1–8 (cit. on pp. 18, 24).
- [Cha+21] Robin Chan et al. "SegmentMeIfYouCan: A Benchmark for Anomaly Segmentation". In: *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track.* virtual conference, Dec. 2021 (cit. on pp. 9, 15).
- [Che+18] Liang-Chieh Chen et al. "Encoder-Decoder With Atrous Separable Convolution for Semantic Image Segmentation". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. Munich, Germany, Sept. 2018, pp. 833–851 (cit. on pp. 11, 23).
- [Che+19] Aida Chefrour et al. "A Novel Incremental Learning Algorithm Based on Incremental Vector Support Machina and Incremental Neural Network Learn++". In: *Revue d'Intelligence Artificielle* 33.3 (2019), pp. 181–188 (cit. on p. 9).

[CM15]	Clement Creusot and Asim Munawar. "Real-Time Small Obstacle Detection on High- ways Using Compressive RBM Road Reconstruction". In: <i>Proceedings of the IEEE In-</i> <i>telligent Vehicles Symposium (IV)</i> . Seoul, Korea, June 2015, pp. 162–167 (cit. on p. 8).
[CME19]	Blake Camp, Jaya Krishna Mandivarapu, and Rolando Estrada. "Self-Net: Lifelong Learn- ing via Continual Self-Modeling". In: <i>arXiv</i> (July 2019), pp. 1–16. arXiv: 1805.10354 (cit. on p. 9).
[Col+21]	Pascal Colling et al. "MetaBox+: A New Region Based Active Learning Method for Se- mantic Segmentation Using Priority Maps". In: <i>Proceedings of the International Confer-</i> <i>ence on Pattern Recognition Applications and Methods (ICPRAM)</i> . virtual conference, Feb. 2021, pp. 51–62 (cit. on p. 27).
[Cor+16]	Marius Cordts et al. "The Cityscapes Dataset for Semantic Urban Scene Understanding". In: <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-</i> <i>nition (CVPR)</i> . Las Vegas, NV, USA, June 2016, pp. 3213–3223 (cit. on pp. 8, 11, 15, 23).
[CRG21]	Robin Chan, Matthias Rottmann, and Hanno Gottschalk. "Entropy Maximization and Meta Classification for Out-of-Distribution Detection in Semantic Segmentation". In: <i>Proceedings of the IEEE International Conference on Computer Vision (ICCV)</i> . virtual conference, Oct. 2021, pp. 5128–5137 (cit. on pp. 8, 14, 24).
[Cyb89]	George Cybenko. "Approximation by Superpositions of a Sigmoidal Function". In: <i>Mathematics of Control, Signals and Systems</i> 2.4 (1989), pp. 303–314 (cit. on p. 4).
[Den+09]	Jia Deng et al. "ImageNet: A Large-Scale Hierarchical Image Database". In: <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)</i> . Miami, FL, USA, June 2009, pp. 248–255 (cit. on p. 25).
[DF07]	Laurens De Haan and Ana Ferreira. <i>Extreme Value Theory: An Introduction</i> . Springer, 2007 (cit. on p. 6).
[DT18]	Terrance DeVries and Graham W. Taylor. "Learning Confidence for Out-of-Distribution Detection in Neural Networks". In: <i>arXiv</i> (Feb. 2018), pp. 1–12. arXiv: 1802.04865 (cit. on pp. 8, 12).
[Est+96]	Martin Ester et al. "A Density-Based Algorithm for Discovering Clusters in Large Spa- tial Databases With Noise". In: <i>Proceedings of the ACM SIGKDD International Confer-</i> <i>ence on Knowledge Discovery and Data Mining (KDD)</i> . Portland, OR, USA, Aug. 1996, pp. 226–231 (cit. on p. 21).
[Far+19]	Mehrdad Farajtabar et al. "Orthogonal Gradient Descent for Continual Learning". In: <i>arXiv</i> (Oct. 2019), pp. 1–13. arXiv: 1910.07104 (cit. on p. 10).
[GDS20]	Fredrik K. Gustafsson, Martin Danelljan, and Thomas Bo Schön. "Evaluating Scalable Bayesian Deep Learning Methods for Robust Computer Vision". In: <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops</i> . virtual conference, June 2020, pp. 1289–1298 (cit. on p. 8).

Preprint – Preprint – Preprint – Preprint – Preprint – Preprin

- [GG16] Yarin Gal and Zoubin Ghahramani. "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning". In: *Proceedings of the International Conference on Machine Learning (ICML)*. New York, NY, USA, June 2016, pp. 1050–1059 (cit. on p. 8).
- [GK16] Alexander Gepperth and Cem Karaoguz. "A Bio-Inspired Incremental Learning Architecture for Applied Perceptual Problems". In: *Cognitive Computation* 8.5 (2016), pp. 924–934 (cit. on p. 9).
- [HAB19] Matthias Hein, Maksym Andriushchenko, and Julian Bitterwolf. "Why ReLU Networks Yield High-Confidence Predictions Far Away from the Training Data and How to Mitigate the Problem". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Long Beach, CA, USA, June 2019, pp. 41–50 (cit. on p. 7).
- [Hen+19] Dan Hendrycks et al. "Using Self-Supervised Learning Can Improve Model Robustness and Uncertainty". In: *Proceedings of the Conference on Neural Information Processing Systems (NIPS/NeurIPS)*. Vancouver, BC, Canada, Dec. 2019, pp. 15637–15648 (cit. on p. 7).
- [Hen+20] Dan Hendrycks et al. "Scaling Out-of-Distribution Detection for Real-World Settings". In: *arXiv* (Dec. 2020), pp. 1–12. arXiv: 1911.11132 (cit. on p. 9).
- [HG17] Dan Hendrycks and Kevin Gimpel. "A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. Toulon, France, Apr. 2017, pp. 1–12 (cit. on pp. 7, 11).
- [HMD19] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. "Deep Anomaly Detection With Outlier Exposure". In: Proceedings of the International Conference on Learning Representations (ICLR). New Orleans, LA, USA, May 2019, pp. 1–18 (cit. on p. 8).
- [Hou+18] Saihui Hou et al. "Lifelong Learning via Progressive Distillation and Retrospection". In: Proceedings of the European Conference on Computer Vision (ECCV). Munich, Germany, Aug. 2018, pp. 452–467 (cit. on p. 10).
- [HTF07] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer, 2007 (cit. on p. 4).
- [Hu+19] Chunyu Hu et al. "A Novel Feature Incremental Learning Method for Sensor-Based Activity Recognition". In: *IEEE Transactions on Knowledge and Data Engineering* 31.6 (2019), pp. 1038–1050 (cit. on p. 10).
- [Hua+17] Gao Huang et al. "Densely Connected Convolutional Networks". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Honolulu, HI, USA, July 2017, pp. 4700–4708 (cit. on p. 25).

[HVD14]	Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. "Distilling the Knowledge in a Neural Network". In: <i>Proceedings of the Conference on Neural Information Processing Systems (NIPS/NeurIPS) Workshops</i> . Montréal, QC, Canada, Dec. 2014, pp. 1–9 (cit. on p. 10).
[HZ21]	Jiangpeng He and Fengqing Zhu. "Unsupervised Continual Learning via Pseudo Labels". In: <i>arXiv</i> (July 2021), pp. 1–9. arXiv: 2104.07164 (cit. on p. 2).
[JRF20]	Nicolas Jourdan, Eike Rehder, and Uwe Franke. "Identification of Uncertainty in Artificial Neural Networks". In: <i>Proceedings of the Uni-DAS e.V. Fahrerassistenz und automatisiertes Fahren (FAS) Workshop</i> . virtual conference, July 2020, pp. 1–11 (cit. on p. 8).
[Jun+18]	Heechul Jung et al. "Less-Forgetful Learning for Domain Expansion in Deep Neural Net- works". In: <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> . New Orleans, LA, USA, Feb. 2018, pp. 3358–3365 (cit. on p. 10).
[KG17]	Alex Kendall and Yarin Gal. "What Uncertainties Do We Need in Bayesian Deep Learn- ing for Computer Vision?" In: <i>Proceedings of the Conference on Neural Information</i> <i>Processing Systems (NIPS/NeurIPS)</i> . Long Beach, CA, USA, Dec. 2017, pp. 5574–5584 (cit. on p. 8).
[Kim+19]	Dahyun Kim et al. "Incremental Learning With Maximum Entropy Regularization: Re- thinking Forgetting and Intransigence". In: <i>arXiv</i> (Feb. 2019), pp. 1–10. arXiv: 1902. 00829 (cit. on p. 10).
[Kir+17]	James Kirkpatrick et al. "Overcoming Catastrophic Forgetting in Neural Networks". In: <i>arXiv</i> (Jan. 2017), pp. 1–13. arXiv: 1612.00796 (cit. on p. 10).
[Kle09]	Jussi Sakari Klemelä. Smoothing of Multivariate Data: Density Estimation and Visual- ization. John Wiley & Sons, 2009 (cit. on p. 4).
[Kli+20]	Marvin Klingner et al. "Class-Incremental Learning for Semantic Segmentation Re- Using Neither Old Data Nor Old Labels". In: <i>Proceedings of the IEEE Intelligent Trans-</i> <i>portation Systems Conference (ITSC)</i> . virtual conference, Sept. 2020, pp. 1–8 (cit. on p. 10).
[KNT00]	Edwin M. Knorr, Raymond T. Ng, and Vladimir Tucakov. "Distance-Based Outliers: Algorithms and Applications". In: <i>The International Journal on Very Large Data Bases</i> 8.3-4 (2000), pp. 237–253 (cit. on p. 2).
[Kri+09]	Hans-Peter Kriegel et al. "Outlier Detection in Axis-Parallel Subspaces of High Dimensional Data". In: <i>Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)</i> . Bangkok, Thailand, Apr. 2009, pp. 831–838 (cit. on p. 2).
[KSZ08]	Hans-Peter Kriegel, Matthias Schubert, and Arthur Zimek. "Angle-Based Outlier De- tection in High-Dimensional Data". In: <i>Proceedings of the ACM SIGKDD International</i> <i>Conference on Knowledge Discovery and Data Mining (KDD)</i> . Las Vegas, NV, USA, Aug. 2008, pp. 444–452 (cit. on p. 2).

- [Lee+18a] Kimin Lee et al. "A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks". In: *Proceedings of the Conference on Neural Information Processing Systems (NIPS/NeurIPS)*. Montréal, QC, Canada, Dec. 2018, pp. 7167–7177 (cit. on pp. 7, 12).
- [Lee+18b] Sang-Woo Lee et al. "Overcoming Catastrophic Forgetting by Incremental Moment Matching". In: *arXiv* (Jan. 2018), pp. 1–16. arXiv: 1703.08475 (cit. on p. 10).
- [Lee+19] Kibok Lee et al. "Overcoming Catastrophic Forgetting With Unlabeled Data in the Wild".
 In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). Seoul, Korea, Oct. 2019, pp. 312–321 (cit. on p. 10).
- [LH18] Zhizhong Li and Derek Hoiem. "Learning Without Forgetting". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 40 (2018), pp. 2935–2947 (cit. on p. 10).
- [Lin+14] Tsung-Yi Lin et al. "Microsoft COCO: Common Objects in Context". In: Proceedings of the European Conference on Computer Vision (ECCV). Zurich, Switzerland, Sept. 2014, pp. 740–755 (cit. on p. 14).
- [Lis+19] Krzysztof Lis et al. "Detecting the Unexpected via Image Resynthesis". In: *Proceedings* of the IEEE International Conference on Computer Vision (ICCV). Seoul, Korea, Oct. 2019, pp. 2152–2161 (cit. on pp. 8, 9, 13).
- [Lis+21] Krzysztof Lis et al. "Detecting Road Obstacles by Erasing Them". In: *arXiv* (Apr. 2021), pp. 1–18. arXiv: 2012.13633 (cit. on pp. 8, 9).
- [LLS18] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. "Enhancing the Reliability of Out-ofdistribution Image Detection in Neural Networks". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. Vancouver, BC, Canada, Apr. 2018, pp. 1–15 (cit. on pp. 7, 11).
- [LPB17] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. "Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles". In: *Proceedings of the Conference on Neural Information Processing Systems (NIPS/NeurIPS)*. Long Beach, CA, USA, Dec. 2017, pp. 6402–6413 (cit. on p. 8).
- [Mac92] David J. C. MacKay. "A Practical Bayesian Framework for Backpropagation Networks". In: *Neural Computation* 4.3 (1992), pp. 448–472 (cit. on p. 8).
- [Mah36] Prasanta Chandra Mahalanobis. "On the Generalized Distance in Statistics". In: *Proceedings of the National Institute of India* 2.1 (1936), pp. 49–55 (cit. on p. 3).
- [MC89] Michael McCloskey and Neal Cohen. "Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem". In: *Psychology of Learning and Motivation* 24 (1989), pp. 109–165 (cit. on p. 9).
- [Mel+19] Diego Mellado et al. "Self-Improving Generative Artificial Neural Network for Pseudorehearsal Incremental Class Learning". In: *Algorithms* 12.10 (2019), pp. 1–17 (cit. on p. 10).

4

- [MG19] Jishnu Mukhoti and Yarin Gal. "Evaluating Bayesian Deep Learning Methods for Semantic Segmentation". In: *arXiv* (Mar. 2019), pp. 1–13. arXiv: 1811.12709 (cit. on pp. 8, 12).
- [MH08] Laurens van der Maaten and Geoffrey Hinton. "Visualizing Data Using t-SNE". In: *Journal of Machine Learning Research* 9.86 (2008), pp. 2579–2605 (cit. on p. 20).
- [MH20] Alexander Meinke and Matthias Hein. "Towards Neural Networks that Provably Know When They Don't Know". In: *Proceedings of the International Conference on Learning Representations (ICLR)*. virtual conference, Apr. 2020, pp. 1–18 (cit. on pp. 7, 8).
- [MVD17] Asim Munawar, Phongtharin Vinayavekhin, and Giovanni De Magistris. "Limiting the Reconstruction Capability of Generative Neural Network Using Negative Learning". In: *Proceedings of the IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*. Tokyo, Japan, Sept. 2017, pp. 1–6 (cit. on p. 8).
- [MZ19] Umberto Michieli and P. Zanuttigh. "Incremental Learning Techniques for Semantic Segmentation". In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*. Seoul, Korea, Oct. 2019, pp. 3205–3212 (cit. on pp. 2, 22).
- [MZ21] Umberto Michieli and Pietro Zanuttigh. "Knowledge Distillation for Incremental Learning in Semantic Segmentation". In: *Computer Vision and Image Understanding* 205 (2021), p. 103167 (cit. on p. 10).
- [Nea96] Radford M. Neal. *Bayesian Learning for Neural Networks*. Springer, 1996 (cit. on p. 8).
- [OOS17] Augustus Odena, Christopher Olah, and Jonathon Shlens. "Conditional Image Synthesis With Auxiliary Classifier GANs". In: *Proceedings of the International Conference on Machine Learning (ICML)*. Sydney, NSW, Australia, Aug. 2017, pp. 2642–2651 (cit. on p. 10).
- [ORF20] Philipp Oberdiek, Matthias Rottmann, and Gernot A. Fink. "Detection and Retrieval of Out-of-Distribution Objects in Semantic Segmentation". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. virtual conference, June 2020, pp. 1331–1340 (cit. on pp. 8, 19, 24, 25).
- [Ost+19] Oleksiy Ostapenko et al. "Learning to Remember: A Synaptic Plasticity Driven Framework for Continual Learning". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Long Beach, CA, USA, June 2019, pp. 11313– 11321 (cit. on p. 10).
- [Pap+03] Sprios Papadimitriou et al. "LOCI: Fast Outlier Detection Using the Local Correlation Integral". In: *Proceedings of the International Conference on Data Engineering*. Bangalore, India, Mar. 2003, pp. 315–326 (cit. on p. 2).
- [Pea01] Karl Pearson. "On Lines and Planes of Closest Fit to Systems of Points in Space". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (1901), pp. 559–572 (cit. on p. 20).

Preprint – Preprint – Preprint – Preprint – Preprint

[Pen+21]	Jian Peng et al. "Overcoming Long-Term Catastrophic Forgetting Through Adversarial Neural Pruning and Synaptic Consolidation". In: <i>IEEE Transactions on Neural Networks and Learning Systems (TNNLS)</i> (Feb. 2021). early access, pp. 1–14 (cit. on p. 10).
[Pin+16]	Peter Pinggera et al. "Lost and Found: Detecting Small Road Hazards for Self-Driving Vehicles". In: <i>Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)</i> . Daejeon, Korea, Oct. 2016, pp. 1099–1106 (cit. on pp. 8, 15, 18).
[Pol+01]	Robi Polikar et al. "Learn++: An Incremental Learning Algorithm for Supervised Neural Networks". In: <i>IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)</i> 31.4 (2001), pp. 497–508 (cit. on p. 9).
[Reb+17]	Sylvestre-Alvise Rebuffi et al. "iCaRL: Incremental Classifier and Representation Learn- ing". In: <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern</i> <i>Recognition (CVPR)</i> . Honolulu, HI, USA, July 2017, pp. 5533–5542 (cit. on p. 10).
[Rob95]	Anthony V. Robins. "Catastrophic Forgetting, Rehearsal and Pseudorehearsal". In: <i>Connect. Sci.</i> 7.2 (1995), pp. 123–146 (cit. on pp. 10, 22).
[Ros52]	Murray Rosenblatt. "Remarks on a Multivariate Transformation". In: <i>The Annals of Mathematical Statistics</i> 23.3 (1952), pp. 470–472 (cit. on p. 4).
[Rot+20]	Matthias Rottmann et al. "Prediction Error Meta Classification in Semantic Segmen- tation: Detection via Aggregated Dispersion Measures of Softmax Probabilities". In: <i>Proceedings of the International Joint Conference on Neural Networks (IJCNN)</i> . virtual conference, July 2020, pp. 1–9 (cit. on pp. 8, 18, 25).
[RPR19]	Deboleena Roy, Priyadarshini Panda, and Kaushik Roy. "Tree-CNN: A Hierarchical Deep Convolutional Neural Network for Incremental Learning". In: <i>arXiv</i> (Sept. 2019), pp. 1–30. arXiv: 1802.05800 (cit. on p. 9).
[RRS00]	Sridhar Ramaswamy, Rajeev Rastogi, and Kyuseok Shim. "Efficient Algorithms for Min- ing Outliers from Large Data Sets". In: <i>Proceedings of the ACM SIGMOD International</i> <i>Conference on Management of Data (MOD)</i> . May 2000, pp. 427–438 (cit. on p. 2).
[Rud87]	Walter Rudin. Real and Complex Analysis. McGraw-Hill, Inc., 1987 (cit. on p. 4).
[Rus+16]	Andrei A. Rusu et al. "Progressive Neural Networks". In: <i>arXiv</i> (Sept. 2016), pp. 1–14. arXiv: 1606.04671 (cit. on p. 9).
[SAR20]	Syed Shakib Sarwar, Aayush Ankit, and Kaushik Roy. "Incremental Learning in Deep Convolutional Neural Networks Using Partial Network Sharing". In: <i>IEEE Access</i> 8 (2020), pp. 4615–4628 (cit. on p. 10).
[Sch+18]	Jonathan Schwarz et al. "Progress & Compress: A Scalable Framework for Continual Learning". In: <i>arXiv</i> (July 2018), pp. 1–13. arXiv: 1805.06370 (cit. on p. 10).
[Shi+17]	Hanul Shin et al. "Continual Learning With Deep Generative Replay". In: <i>Proceedings of the Conference on Neural Information Processing Systems (NIPS/NeurIPS)</i> . Long Beach, CA, USA, Dec. 2017, pp. 2990–2999 (cit. on p. 10).

[Sin+20] Aasheesh Singh et al. "LiDAR Guided Small Obstacle Segmentation". In: arXiv (Mar. 2020), pp. 1–8. arXiv: 2003.05970 (cit. on p. 9). Mayu Sakurada and Takehisa Yairi. "Anomaly Detection Using Autoencoders With Non-[SY14] linear Dimensionality Reduction". In: Proceedings of the Workshop on Machine Learning for Sensory Data Analysis (MLSDA). Gold Coast, QLD, Australia, Dec. 2014, pp. 4-11 (cit. on p. 4). [TTA19] Onur Tasar, Yuliya Tarabalka, and Pierre Alliez. "Incremental Learning for Semantic Segmentation of Large-Scale Remote Sensing Data". In: IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 12.9 (2019), pp. 3524–3537 (cit. on p. 10). [URG22] Svenja Uhlemeyer, Matthias Rottmann, and Hanno Gottschalk. Towards Unsupervised Open World Semantic Segmentation. Jan. 2022. arXiv: 1709.02435 (cit. on pp. 10, 19). [Van+21] Simon Vandenhende et al. "Multi-Task Learning for Dense Prediction Tasks: A Survey". In: IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI). early access. Jan. 2021, pp. 1–20 (cit. on p. 8). [VT19] Gido M. van de Ven and Andreas S. Tolias. "Three Scenarios for Continual Learning". In: arXiv (Apr. 2019), pp. 1–18. arXiv: 1904.07734 (cit. on p. 9). [Wu+18] Yue Wu et al. "Incremental Classifier Learning With Generative Adversarial Networks". In: arXiv (Feb. 2018), pp. 1–10. arXiv: 1802.00853 (cit. on p. 10). [Wu+19] Yue Wu et al. "Large Scale Incremental Learning". In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Long Beach, CA, USA, June 2019, pp. 374–382 (cit. on p. 10). [Xia+20] Yingda Xia et al. "Synthesize Then Compare: Detecting Failures and Anomalies for Semantic Segmentation". In: Proceedings of the European Conference on Computer Vision (*ECCV*). virtual conference, Aug. 2020, pp. 145–161 (cit. on p. 8). [Yan+19] Yang Yang et al. "Adaptive Deep Models for Incremental Learning: Considering Capacity Scalability and Sustainability". In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD). Anchorage, AK, USA, July 2019, pp. 74–82 (cit. on p. 9). [Yao+19] Xin Yao et al. "Adversarial Feature Alignment: Avoid Catastrophic Forgetting in Incremental Task Lifelong Learning". In: Neural Computation 31.11 (2019), pp. 2266–2291 (cit. on p. 10). [Yoo+18] Jaehong Yoon et al. "Lifelong Learning With Dynamically Expandable Networks". In: arXiv (June 2018), pp. 1–11. arXiv: 1708.01547 (cit. on p. 9). [Yu+20] Fisher Yu et al. "BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning". In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). virtual conference, June 2020, pp. 2636–2645 (cit. on p. 9).

Preprint – Preprint – Preprint – Preprint – Preprint

Bibliography

[Zen+21]	Guanxiong Zeng et al. "Continuous Learning of Context-dependent Processing in Neural Networks". In: <i>arXiv</i> (June 2021), pp. 1–18. arXiv: 1810.01256 (cit. on p. 10).
[Zhu+19]	Yi Zhu et al. "Improving Semantic Segmentation via Video Propagation and Label Re- laxation". In: <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern</i> <i>Recognition (CVPR)</i> . Long Beach, CA, USA, June 2019, pp. 8856–8865 (cit. on p. 23).
[ZPG17]	Friedemann Zenke, Ben Poole, and Surya Ganguli. "Continual Learning Through Synap- tic Intelligence". In: <i>Proceedings of the International Conference on Machine Learning</i> (<i>ICML</i>). Sydney, NSW, Australia, Aug. 2017, pp. 3987–3995 (cit. on p. 10).

37