



The Fishyscapes Benchmark: Measuring Blind Spots in Semantic Segmentation

Hermann Blum¹ · Paul-Edouard Sarlin² · Juan Nieto³ · Roland Siegwart¹ · Cesar Cadena¹

Received: 25 January 2021 / Accepted: 3 August 2021 / Published online: 14 September 2021
© The Author(s) 2021

Abstract

Deep learning has enabled impressive progress in the accuracy of semantic segmentation. Yet, the ability to estimate uncertainty and detect failure is key for safety-critical applications like autonomous driving. Existing uncertainty estimates have mostly been evaluated on simple tasks, and it is unclear whether these methods generalize to more complex scenarios. We present Fishyscapes, the first public benchmark for anomaly detection in a real-world task of semantic segmentation for urban driving. It evaluates pixel-wise uncertainty estimates towards the detection of anomalous objects. We adapt state-of-the-art methods to recent semantic segmentation models and compare uncertainty estimation approaches based on softmax confidence, Bayesian learning, density estimation, image resynthesis, as well as supervised anomaly detection methods. Our results show that anomaly detection is far from solved even for ordinary situations, while our benchmark allows measuring advancements beyond the state-of-the-art. Results, data and submission information can be found at <https://fishyscapes.com/>.

Keywords Semantic segmentation · Anomaly detection · Uncertainty estimation · Out-of-distribution detection · Autonomous driving

1 Introduction

Deep learning has had a high impact on the precision of computer vision methods (Chen et al., 2018; He et al., 2017; Fu et al., 2018; Sun et al., 2018) and enabled semantic understanding in robotic applications (McCormac et al., 2018; Florence et al., 2018; Liang et al., 2018). However, while these algorithms are usually compared on closed-world datasets with a fixed set of classes (Geiger et al., 2012; Cordts et al., 2016), the real-world is uncontrollable, and an incorrect reaction by an autonomous agent to an unexpected input can have disastrous consequences (Bozhinoski et al., 2019).

As such, to reach full autonomy while ensuring safety and reliability, decision-making systems need information

about outliers and uncertain or ambiguous cases that might affect the quality of the perception output. As illustrated in Fig. 1, deep convolutional neural networks (CNNs) react unpredictably for inputs that deviate from their training distribution. In the presence of outlier objects, this is interpolated with the available classes at high confidence. Existing research to detect such behaviour is often labeled as out-of-distribution (OoD), anomaly, or novelty detection, and has so far focused on developing methods for image classification, evaluated on simple datasets like MNIST or CIFAR-10 (Malinin and Gales, 2018; Papernot & McDaniel, 2018; Hendrycks & Gimpel, 2017; Lee et al., 2018; Ruff et al., 2018; Golan & El-Yaniv, 2018; Choi et al., 2018; Sabokrou et al., 2018; Pidhorskyi et al., 2018). How these methods generalize to more elaborate network architectures and pixel-wise uncertainty estimation has not been assessed in prior work.

Motivated by these practical needs, we introduce ‘Fishyscapes’, a benchmark that evaluates uncertainty estimates for semantic segmentation. The benchmark measures how well methods detect potentially hazardous anomalies in driving scenes. Fishyscapes is based on data from Cityscapes (Cordts et al., 2016), a popular benchmark for semantic segmentation in urban driving. Our benchmark consists of (i) Fishyscapes

Communicated by O. Veksler.

This work was supported by the Hilti Group.

✉ Hermann Blum
blumh@ethz.ch

¹ Autonomous Systems Lab, ETH Zürich, Zürich, Switzerland

² Computer Vision and Geometry Group, ETH Zürich, Zürich, Switzerland

³ Microsoft Research, Zürich, Switzerland

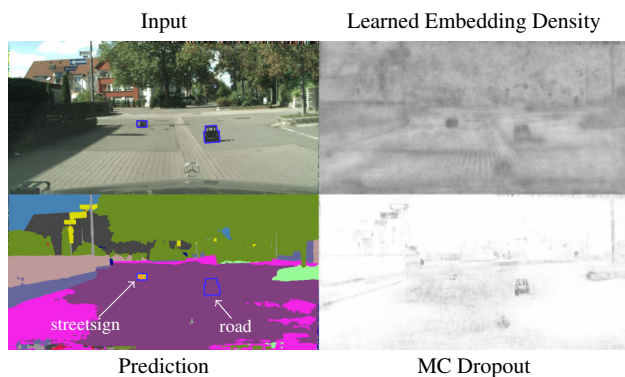


Fig. 1 When exposed to an object type unseen during training, a state-of-the-art semantic segmentation model (Chen et al., 2018) predicts familiar labels (*streetsign*, *road*) with high confidence. To detect such failures, we evaluate various methods that assign a pixel-wise out-of-distribution score, where higher values are darker. The blue outline is added for illustration.

Web, where images from Cityscapes are overlaid with objects that are regularly crawled from the web in an open-world setup, and (ii) Fishyscapes Lost and Found, that builds up on a road hazard dataset collected with the same setup as Cityscapes (Pinggera et al., 2016) and that we supplemented with labels.

To provide a broad overview, we adapt a variety of methods to semantic segmentation that were originally designed for image classification. Because segmentation networks are much more complex and have high computational costs, this adaptation is not trivial, and we suggest different approximations to overcome these challenges.

Our experiments show that the embeddings of intermediate layers hold important information for anomaly detection. Based on recent work on generative models, we develop a novel method using density estimation in the embedding space. However, we also show that varying visual appearance can mislead feature-based and other methods. None of the evaluated methods achieves the accuracy required for safety-critical applications. We conclude that these remain open problems, with our benchmark enabling the community to measure progress and build upon the best performing methods so far.

To summarize, our contributions are the following:

- We introduce the first public benchmark evaluating pixel-wise uncertainty estimates in semantic segmentation, with a dynamic, self-updating dataset for anomaly detection.
- We report an extensive evaluation with diverse state-of-the-art approaches to uncertainty estimation, adapted to the semantic segmentation task, and present a novel method for anomaly detection.

- We show a clear gap between the alleged capabilities of established methods and their performance on this real-world task, thereby confirming the necessity of our benchmark to support further research in this direction.

2 Related Work

Here we review the most relevant works in semantic segmentation and their benchmarks, and methods that aim at providing a confidence estimate of the output of deep networks.

2.1 Semantic Segmentation

State-of-the-art models are fully-convolutional deep networks trained with pixel-wise supervision. Most works (Ronneberger et al., 2015; Badrinarayanan et al. 2017; Chen et al., 2016; Chen et al., 2018) adopt an encoder-decoder architecture that initially reduces the spatial resolution of the feature maps, and subsequently upsamples them with learned transposed convolution, fixed bilinear interpolation, or unpooling. Additionally, dilated convolutions or spatial pyramid pooling enlarge the receptive field and improve the accuracy.

Popular benchmarks compare methods on the segmentation of objects (Everingham et al., 2010) and urban scenes. In the latter case, Cityscapes (Cordts et al., 2016) is a well-established dataset depicting street scenes in European cities with dense annotations for a limited set of classes. Efforts have been made to provide datasets with increased diversity, either in terms of environments, with WildDash (Zendel et al. 2018), which incorporates data from numerous parts of the world, or with Mapillary (Neuhold et al., 2017), which adds many more classes. Recent data releases add multi-sensor and multi-modality recordings on top of that (Sun et al., 2020; Geyer et al., 2020; Caesar et al., 2020). Like ours, some datasets are explicitly derived from Cityscapes, the most relevant being Foggy Cityscapes (Sakaridis et al., 2018), which overlays synthetic fog onto the original dataset to evaluate more difficult driving conditions. The Robust Vision Challenge¹ also assesses generalization of learned models across different datasets.

Robustness and reliability are only evaluated by these benchmarks through ranking methods according to their accuracy, without taking into accounts the uncertainty of their predictions. Additionally, despite the fact that one cannot assume that models trained with closed-world data will only encounter known classes, these scenarios are rarely quantitatively evaluated. To our knowledge, WildDash (Zendel et al. 2018) is the only public benchmark that explicitly reports uncertainty w.r.t. OoD examples. These are however

¹ <http://www.robustvision.net/>.

drawn from a very limited set of full-image outliers, while we introduce a diverse set of objects, as WildDash mainly focuses on accuracy. Complementarily, the Dark Zurich dataset (Sakaridis et al., 2020) allows for uncertainty-aware evaluation of semantic segmentation models with regard to deprived sensor inputs, i.e. evaluating aleatoric uncertainty.

Bevandic et al. (2019) experiment with OoD objects for semantic segmentation by overlaying objects on Cityscapes images in a manner similar to ours. They however assume the availability of a large OoD dataset, which is not realistic in an open-world context, and thus mostly evaluate supervised methods. In contrast, we assess a wide range of methods that do not require OoD data. Mukhoti and Gal (2018) introduce a new metric for uncertainty evaluation and are the first to quantitatively assess misclassification for segmentation. Yet they only compare few methods on normal ID data. The MVTec benchmark (Bergmann et al., 2019) compares a range of anomaly segmentation methods on images of single objects to find industrial production anomalies. It mostly compare methods that focus on low-power computing. Following our work, the CAOS benchmark (Hendrycks et al., 2019) also compares anomaly segmentation methods in simulated and real-world driving scenes. While their results confirm our finding that most established methods scale poorly to semantic segmentation, their methodology lacks open-world testing, which we argue later is important for true anomaly detection.

2.2 Uncertainty Estimation

There is a large body of work that aims at detecting OoD data or misclassification by defining uncertainty or confidence estimates.

Probabilistic modeling of a neural network's output is a straightforward approach in uncertainty estimation. The softmax score, i.e. the classification probability of the predicted class, was shown to be a first baseline (Hendrycks & Gimpel, 2017), although sensitive to adversarial examples (Goodfellow et al., 2015). Its performance was improved by ODIN (Liang et al., 2018), which applies noise to the input with the Fast Gradient Sign Method (FGSM) (Goodfellow et al., 2015) and calibrates the score with temperature scaling (Guo et al., 2017). Probabilistic modelling has been extended further in Deep Belief Networks that propagate activation distributions throughout the network (Frey and Hinton, 1999; Loquercio et al., 2020).

Bayesian deep learning (Gal, 2016; Kendall and Gal, 2017) adopts a probabilistic view by designing deep models whose outputs and weights are probability distributions instead of point estimates. Uncertainties are then defined as dispersions of such distributions, and can be of several types. *Epistemic* uncertainty, or model uncertainty, corresponds to the uncertainty over the model parameters that best

fit the training data for a given model architecture. As evaluating the posterior over the weights is intractable in deep non-linear networks, recent works perform (MC) sampling with dropout (Gal & Ghahramani, 2016) or ensembles (Lakshminarayanan et al., 2017). *Aleatoric* uncertainty, or data uncertainty, arises from the noise in the input data, such as sensor noise. Both have been applied to semantic segmentation (Kendall and Gal, 2017), and successively evaluated for misclassification detection (Mukhoti & Gal, 2018), but only on ID data and not for OoD detection. Malinin and Gales (2018) later single out *distributional* uncertainty to represent model misspecification with respect to OoD inputs. Their approach however was only applied to image classifications on toy datasets, and requires OoD data during the training stage. To address the latter constraint, Lee et al. (2018) earlier proposed a Generative Adversarial Network (GAN) that generates OoD data as boundary samples. This is however very challenging to scale to complex and high-dimensional data like high-resolution images of urban scenes. Recently, Bayesian methods investigated the inductive bias of network structures beyond weights (Wilson & Izmailov, 2020). For example, Antoran et al. (2020) extracts meaningful uncertainties from an 'ensemble' of network activations at varying depth, and Yehezkel Rohekar et al. (2019) employs a sampling scheme for architectures.

OoD and novelty detection is often tackled by non-Bayesian approaches. As such, feature introspection amounts to measuring discrepancies between distributions of deep features of training data and OoD samples, using either (NN) statistics (Papernot & McDaniel, 2018; Mandelbaum & Weinshall, 2017) or Gaussian approximations (Lee et al., 2018; Amersfoort et al., 2020). These methods have the benefit of working on any classification model without requiring specific training. Recently, connections between feature density and Bayesian uncertainties have been investigated (Postels et al., 2020). On the other hand, approaches specifically tailored to perform OoD detection include one-class classification (Ruff et al., 2018; Golan & El-Yaniv, 2018), which aim at creating discriminative embeddings, density estimation (Choi et al., 2018; Nalisnick et al., 2019), which estimate the likelihood of samples w.r.t to the true data distribution, and generative reconstruction (Sabokrou et al., 2018; Pidhorskyi et al., 2018; Gong et al., 2019), which use the quality of auto-encoder reconstructions to discriminate OoD samples. Richter and Roy (2017) apply the latter to simple real images recorded by a robotic car and successfully detect new environments.

3 Benchmark Design

Because it is not possible to produce ground truth for uncertainty values, evaluating estimators is not a straightforward

task. We thus compare them on the proxy classification task (Hendrycks & Gimpel, 2017) of detecting anomalous inputs. The uncertainty estimates are seen as scores of a binary classifier that compares the score against a threshold and whose performance reflects the suitability of the estimated uncertainty for anomaly detection.

Such an approach however introduces a major issue for the design of a public OoD detection benchmark. With publicly available ID training data A and OoD inputs B , it is not possible to distinguish between an uncertainty method that informs a classifier to discriminate A from any other input, and a classifier trained to discriminate A from B . The latter option clearly does not represent progress towards the goal of general uncertainty estimation, but rather overfitting.

To this end, we (i) only release a small validation set with associated ground truth masks, while keeping larger test sets hidden, (ii) continuously evaluate submitted methods against a dynamically changing, synthetic dataset, and (iii) compare the performance on the dynamic dataset with evaluations on real-world data. Additionally, all submissions to the benchmark must indicate whether any OoD data was used during training, which is cross-checked with linked publications.

Examples from all benchmark datasets are shown in Fig. 2.

3.1 Does the Method Work in an Open World?

The open world scenario describes the problem that an autonomous agent who is freely interacting with the world has to be able to deal with the unexpected at all times. To test perception methods in an open world scenario, a benchmark therefore needs to present truly unexpected inputs. We argue that this is never truly possible with a fixed dataset that by design has limited diversity, and over time may simply identify those methods that deal best with the kind of objects included in the dataset. Instead, we propose a dynamically changing dataset that samples diverse objects at every iteration.

In general, there are three options to generate such dynamic datasets: At every iteration, one may (i) capture new data in the wild and annotate, (ii) render new objects in simulation, or (iii) capture new objects in the wild, but blend them into already annotated scenes. While data from the wild is essential to test methods in realistic settings, annotation for semantic segmentation is very expensive and not a sustainable way to generate new datasets multiple times per year. Between (ii) and (iii) there is an essential trade-off. Rendering in 3D ensures physically viable object placement and consistent lighting. Images of diverse objects in the wild are much better available than textured 3D models and can be blended into real-world scenes. We acknowledge that there is an ongoing debate whether photorealistic rendering engines or modern blending techniques achieve more realistic images, which was touched upon by a response-work to

this benchmark (Hendrycks et al., 2019). In this work, we decided to base our dataset FS Web on approach (iii). In the following, we describe a blending-based reference dataset FS Static and the dynamically changing dataset FS Web.

FS Static is based on the validation set of Cityscapes (Cordts et al., 2016). It has a limited visual diversity, which is important to make sure that it contains none of the overlaid objects. In addition, background pixels originally belonging to the void class² are excluded from the evaluation, as they may be borderline OoD. Anomalous objects are extracted from the generic Pascal VOC (Everingham et al., 2010) dataset using the associated segmentation masks. We only overlay objects from classes that cannot be found in Cityscapes: *aeroplane, bird, boat, bottle, cat, chair, cow, dog, horse, sheep, sofa, tvmonitor*. Objects cropped by the image borders or objects that are too small to be seen are filtered out. We randomly size and position the objects on the underlying image, making sure that none of the objects appear on the ego-vehicle. Objects from mammal classes have a higher probability of appearing on the lower-half of the screen, while classes like birds or airplanes have a higher probability for the upper half. The placing is not further limited to ensure each pixel in the image, apart from the ego-vehicle, is comparably likely to be anomalous. To match the image characteristics of cityscapes, we employ a series of postprocessing steps similar to those described in Abu Alhaija et al. (2018), without those steps that require 3D models of the objects to e.g. adapt shadows and lighting.

To make the task of anomaly detection harder, we add synthetic fog (Sakaridis et al., 2018; Dai et al., 2020) on the in-distribution pixels with a per-image probability. This prevents fraudulent methods to compare the input against a fixed set of Cityscapes images. The dataset is split into a minimal public validation set of 30 images and a hidden test set of 1000 images. It contains in total around $4.5e7$ OoD and $1.8e9$ ID pixels. The validation set only contains a small disjoint set of pascal objects to prevent few-shot learning on our data creation method.

FS Web is built similarly to FS Static, but with overlay objects crawled from the internet using a changing list of keywords. Our script searches for images with transparent background, uploaded in a recent timeframe, and filters out images that are too small. The only manual process is filtering out images that are not suitable, e.g. with decorative borders or watermarks. The dataset for March 2019 contains $4.9e7$ OoD and $1.8e9$ ID pixels. As the diversity of images and color distributions for the images from the web is much greater than

² *void* in cityscapes is defined as: forms of horizontal ground-level structures that do not match any class, things that might not be there anymore the next day/hour/minute (e.g. movable trash bin, buggy, bag, wheelchair, animal), clutter in the background that is not distinguishable, or any objects that do not match a class (e.g. visible parts of the ego vehicle, mountains, street lights, back side of signs).




Fig. 2 Qualitative examples of Fishyscapes Static (rows 1–2) and Fishyscapes Web (rows 3–5) and Fishyscapes Lost and Found (rows 6–8). The ground truth contains labels for ID (blue) and OoD (red) pixels, as well as ignored void pixels (black). We additionally show the

output of the best method per dataset in column 4 and the best method without OoD training in the last column. We report the AP of each method output in its top right corner (Color figure online).

those from Pascal VOC, we also adapt our overlay procedure. In total, we follow these steps, some of which were added from June 2019 onwards (marked with *):

- in case the image does not already have a smooth alpha channel, smooth the mask of the objects around the borders for a small transparency gradient
- adapt the brightness of the object towards the mean brightness of the overlaid pixels
- apply the inverse color histogram of the Cityscapes image to shift the color distribution towards the one found on the underlying image*
- radial motion blur*
- depth blur based on the position in the image*
- color noise

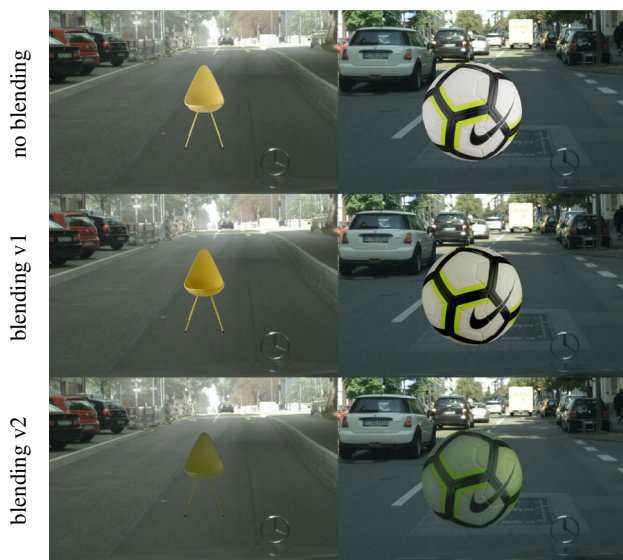


Fig. 3 Illustration of the blending process and improvements (v2) applied in June 2019. While color adaptation to the predominantly gray Cityscapes images is visually most obvious, important improvements in v2 include depth and motion blur, as well as glow effects.

- glow effects to simulate overexposure*

Figure 3 shows an illustration of the blending results.

As discussed, the blending process is part of a trade-off to make an open-world dataset feasible. To further ensure that methods do not overfit to any artifacts created by the blending process, but detect anomalies based on their semantics and appearance, we include a sample of ID objects in the blending dataset. For this, we create a database from objects in the Cityscapes training dataset (*car, person, truck, bus, train, bike*) where we manually filter out any occluded instances. We then decide at random for every image whether to blend an anomalous object or a Cityscapes object, where we skip random placement and histogram adaptation for the latter. This addition was introduced in FS Web Jan 2020. An example can be seen in Fig. 2.

As indicated, the postprocessing was improved between iterations of the dataset. Because the purpose of the FS Web dataset is to measure any possible overfitting of the methods through a dynamically changing dataset, we will continue to refine also this image overlay procedure, updating our method with recent research results. Any update to the blending is also applied to the FS Static validation set, allowing submissions to validate the effect of blending improvements.

3.2 Does the Method Work on Real Images?

As discussed in Sect. 3.1, capturing and annotating driving scenes multiple times per year is not sustainable, which made it necessary to use synthetic data generation for the

dynamic dataset. However, for safe deployment it is equally important to test methods under real-world conditions. This is the purpose of the FS Lost and Found dataset in our benchmark.

FS Lost and Found is based on the original Lost and Found dataset (Pinggera et al., 2016). However, the original dataset only includes annotations for the anomalous objects and a coarse annotation of the road. It does not allow for appropriate evaluation of anomaly detection, as objects and road are very distinct in texture and it is more challenging to evaluate the anomaly score of the objects compared to eg. building structures. In order to make use of the full image, we add pixel-wise annotations that distinguish between *objects* (the anomalies), *background* (classes contained in Cityscapes) and *void* (anything not contained in Cityscapes classes that still appears in the training images). Additionally, we filter out those sequences where the ‘road hazards’ are children or bikes, because these are part of regular Cityscapes data and not anomalies. We subsample the repetitive sequences, labelling at least every sixth image, and remove images that do not contain objects. In total, we present a public validation set of 100 images and a testset of 275 images, based on disjoint sets of locations.

While the Lost and Found images were captured with the same setup as Cityscapes, the distribution of street scenery is very different. The images were captured in small streets of housing areas, industrial areas, or on big parking lots. The anomalous objects are usually very small and are not equally distributed on the image. Nevertheless, the dataset allows to test for real images as opposed to synthetic data, therefore preventing any overfitting on synthetic image processing. This is especially important for parameter tuning on the validation set.

3.3 Metrics

We consider metrics associated with a binary classification task. Since the ID and OoD data is unbalanced, metrics based on the (ROC) are not suitable (Saito & Rehmsmeier, 2015). We therefore base the ranking and primary evaluation on the (AP). However, as the number of false positives in high-recall areas is particularly relevant for safety-critical applications, we additionally report the false positive rate at 95% recall (FPR_{95}). This metric was also used in Hendrycks and Gimpel (2017) and emphasizes safety.

Semantic classification is not the goal of our benchmark, but uncertainty estimation and outlier detection should not come at high cost of segmentation accuracy. We therefore additionally report the mean (IoU) of the semantic segmentation on the Cityscapes validation set.

For safety-critical systems, it is not only important to detect anomalies, but also to be fast enough to allow for a reaction. We therefore report the inference time of joint seg-

mentation and anomaly detection per single frame. Times are measured over 500 images of the Cityscapes validation set on a GeForce 1080 Ti GPU.

4 Evaluated Methods

We now present the methods that are evaluated in Fishyscapes. In a first part, we describe the existing baselines and how we adapted them to the task of semantic segmentation. We then propose a novel method based on learned embedding density. Finally, we list those methods that were submitted to the public benchmark so far.

All approaches are applied to the state-of-the-art semantic segmentation model DeepLab-v3+ (Chen et al., 2018). Further implementation details are listed in the supplementary material.

4.1 Baselines

Softmax The maximum softmax probability is a commonly used baseline and was evaluated in Hendrycks and Gimpel (2017) for OoD detection. We apply the metric pixel-wise and additionally measure the softmax entropy, as proposed by Lee et al. (2018), which captures more information from the softmax.

OoD training While we generally strive for methods that are not biased by data, learning confidence from data is an obvious baseline and was explored in DeVries and Taylor (2018). As we are not supposed to know the true OoD distribution, we do not use Pascal VOC, but rather approximate unknown pixels with the Cityscapes *void* class. In our evaluation, we (i) train a model to maximise the softmax entropy for OoD pixels, or (ii) introduce void as an additional output class and train with it. The uncertainty is then measured as (i) the softmax entropy, or (ii) the score of the void class.

Bayesian DeepLab was introduced by Mukhoti and Gal (2018), following Kendall and Gal (2017), and is the only uncertainty estimate already applied to semantic segmentation in the literature. The epistemic uncertainty is modeled by adding Dropout layers to the encoder, and approximated by T (MC) samples, while the aleatoric uncertainty corresponds to the spread of the categorical distribution. The total uncertainty is the predictive entropy of the distribution \mathbf{y} ,

$$\hat{\mathbb{H}}[\mathbf{y}|\mathbf{x}] = - \sum_c \left(\frac{1}{T} \sum_t y_c^t \right) \log \left(\frac{1}{T} \sum_t y_c^t \right), \tag{1}$$

where y_c^t is the probability of class c for sample t . The epistemic uncertainty is measured as the mutual information (MI) between \mathbf{y} and the weights \mathbf{w} ,

$$\hat{\mathbb{H}}[\mathbf{y}, \mathbf{w}|\mathbf{x}] = \hat{\mathbb{H}}[\mathbf{y}|\mathbf{x}] - \frac{1}{T} \sum_{c,t} y_c^t \log y_c^t. \tag{2}$$

Dirichlet DeepLab Prior networks (Malinin and Gales, 2018) extend the framework of Gal (2016) by considering the predicted logits \mathbf{z} as log concentration parameters $\boldsymbol{\alpha}$ of a Dirichlet distribution, which is a prior of the predictive categorical distribution \mathbf{y} . Intuitively, the spread of the Dirichlet prior should model the distributional uncertainty, and remain separate from the data uncertainty modelled by the spread of the categorical distribution. To this end, Malinin and Gales (2018) advocate to train the network with the objective:

$$\begin{aligned} \mathcal{L}(\theta) = & \mathbb{E}_{p_{\text{in}}} [\text{KL} [\text{Dir}(\boldsymbol{\mu}|\boldsymbol{\alpha}_{\text{in}})||p(\boldsymbol{\mu}|\mathbf{x}; \boldsymbol{\theta})]] \\ & + \mathbb{E}_{p_{\text{out}}} [\text{KL} [\text{Dir}(\boldsymbol{\mu}|\boldsymbol{\alpha}_{\text{out}})||p(\boldsymbol{\mu}|\mathbf{x}; \boldsymbol{\theta})]] \\ & + \text{CrossEntropy}(\mathbf{y}, \mathbf{z}). \end{aligned} \tag{3}$$

The first term forces ID samples to produce sharp priors with a high concentration $\boldsymbol{\alpha}_{\text{in}}$, computed as the product of smoothed labels and a fixed scale α_0 . The second term forces OoD samples to produce a flat prior with $\boldsymbol{\alpha}_{\text{out}} = \mathbf{1}$, effectively maximizing the Dirichlet entropy, while the last one helps the convergence of the predictive distribution to the ground truth. We model pixel-wise Dirichlet distributions, approximate OoD samples with void pixels, and measure the Dirichlet differential entropy.

kNN Embedding. Different works (Papernot & McDaniel, 2018; Mandelbaum & Weinshall, 2017) estimate uncertainty using kNN statistics between inferred embedding vectors and their neighbors in the training set. They then compare the classes of the neighbors to the prediction, where discrepancies indicate uncertainty. In more details, a given trained encoder maps a test image \mathbf{x}' to an embedding $\mathbf{z}'_l = \mathbf{f}_l(\mathbf{x}')$ at layer l , and the training set \mathbf{X} to a set of neighbors $\mathbf{Z}_l := \mathbf{f}_l(\mathbf{X})$. Intuitively, if \mathbf{x}' is OoD, then \mathbf{z}' is also differently distributed and has e.g. neighbors with different classes. Adapting these methods to semantic segmentation faces two issues: (i) The embedding of an intermediate layer of DeepLab is actually a map of embeddings, resulting in more than 10,000 kNN queries for each layer, which is computationally infeasible. We follow Mandelbaum and Weinshall (2017) and pick only one layer, selected using the FS Lost and Found validation set. (ii) The embedding map has a lower resolution than the input and a given training embedding $\mathbf{z}_l^{(i)}$ is therefore not associated with one, but with multiple output labels. As a baseline approximation, we link $\mathbf{z}_l^{(i)}$ to all classes in the associated image patch. The relative density (Mandelbaum & Weinshall, 2017) is then:

$$D(\mathbf{z}') = \frac{\sum_{i \in K, c'=c_i} \exp \left(- \frac{\mathbf{z}' \mathbf{z}^{(i)}}{|\mathbf{z}'| |\mathbf{z}^{(i)}|} \right)}{\sum_{i \in K} \exp \left(- \frac{\mathbf{z}' \mathbf{z}^{(i)}}{|\mathbf{z}'| |\mathbf{z}^{(i)}|} \right)}. \tag{4}$$

Here, c_i is the class of $\mathbf{z}^{(i)}$ and c' is the class of \mathbf{z}' in the downsampled prediction. In contrast to Mandelbaum and Weinshall (2017), we found that the cosine similarity from Papernot and McDaniel (2018) works well without additional losses. Finally, we upsample the density of the feature map to the input size, assigning each pixel a density value.

As the class association is unclear for encoder-decoder architectures, we also evaluate the density estimation with k neighbors independent of the class:

$$D(\mathbf{z}') = \sum_{i \in K} \exp\left(-\frac{\mathbf{z}' \cdot \mathbf{z}^{(i)}}{|\mathbf{z}'| |\mathbf{z}^{(i)}|}\right). \quad (5)$$

This assumes that an OoD sample \mathbf{x}' , with a low density w.r.t \mathbf{X} , should translate into \mathbf{z}' with a low density w.r.t \mathbf{Z}_l .

4.2 Learned Embedding Density

We now introduce a novel approach that takes inspiration from density estimation methods while greatly improving their scalability and flexibility.

Density estimation using k NN has two weaknesses. First, the estimation is a very coarse isotropic approximation, while the distribution in feature space might be significantly more complex. Second, it requires to store the embeddings of the entire training set and to run a large number of NN searches, both of which are costly, especially for large input images. On the other hand, recent works (Choi et al., 2018; Nalisnick et al., 2019) on OoD detection leverage more complex generative models, such as normalizing flows (Dinh et al., 2017; Kingma & Dhariwal, 2018; Dinh et al., 2014), to directly estimate the density of the input sample \mathbf{x} . This is however not directly applicable to our problem, as (i) learning generative models of images that can capture the entire complexity of e.g. urban scenes is still an open problem; and (ii) the pixel-wise density required here should be conditioned on a very (ideally infinitely) large context, which is computationally intractable.

Our approach mitigates these issues by learning the density of \mathbf{z} . We start with a training set \mathbf{X} drawn from the unknown true distribution $\mathbf{x} \sim p^*(\mathbf{x})$, and corresponding embeddings \mathbf{Z}_l . A normalizing flow with parameters θ is trained to approximate $p^*(\mathbf{z}_l)$ by minimizing the negative log-likelihood (NLL) over all training embeddings in \mathbf{Z}_l :

$$\mathcal{L}(\mathbf{Z}_l) = -\frac{1}{|\mathbf{Z}_l|} \sum_i \log p_\theta(\mathbf{z}_l^{(i)}). \quad (6)$$

The flow is composed of a bijective function \mathbf{g}_θ that maps an embedding \mathbf{z}_l to a latent vector η of identical dimensionality and with Gaussian prior $p(\eta) = \mathcal{N}(\eta; 0, \mathbf{I})$. Its loglikelihood

is then expressed as

$$\log p_\theta(\mathbf{z}_l) = \log p(\eta) + \log \left| \det \left(\frac{d\mathbf{g}_\theta}{d\mathbf{z}} \right) \right|, \quad (7)$$

and can be efficiently evaluated for some constrained \mathbf{g}_θ . At test time, we compute the embedding map of an input image, and estimate the NLL of each of its embeddings. In our experiments, we use the Real-NVP bijector (Dinh et al., 2017), composed of a succession of affine coupling layers, batch normalizations, and random permutations.

The benefits of this method are the following: (i) A normalizing flow can learn more complex distributions than the simple k NN kernel or mixture of Gaussians used by Lee et al. (2018), where each embedding requires a class label, which is not available here; (ii) Features follow a simpler distribution than the input images, and can thus be correctly fit with simpler flows and shorter training times; (iii) The only hyperparameters are related to the architecture and the training of the flow, and can be cross-validated with the NLL of ID data without any OoD data; (iv) The training embeddings are efficiently summarized in the weights of the generative model with a very low memory footprint.

Input preprocessing (Liang et al., 2018) can be trivially applied to our approach. Since the NLL estimator is an end-to-end network, we can compute the gradients of the average NLL w.r.t. the input image by backpropagating through the flow and the encoder.

A flow ensemble can be built by training separate density estimators over different layers of the segmentation model, similar to Lee et al. (2018). However, the resulting NLL estimates cannot be directly aggregated as is, because the different embedding distributions have varying dispersions and dimensions, and thus densities with very different scales. We propose to normalize the NLL $N(\mathbf{z}_l)$ of a given embedding by the average NLL of the training features for that layer:

$$\bar{N}(\mathbf{z}_l) = N(\mathbf{z}_l) - \mathcal{L}(\mathbf{Z}_l). \quad (8)$$

This is in fact a (MC) approximation of the differential entropy of the flow, which is intractable. In the ideal case of a multivariate Gaussian, \bar{N} corresponds to the Mahalanobis distance used by Lee et al. (2018). We can then aggregate the normalized, resized scores over different layers. We experiment with two strategies: (i) Using the minimum detects a pixel as OoD only if it has low likelihood through all layers, thus accounting for areas in the feature space that are in-distribution but contain only few training points; (ii) Following Lee et al. (2018), taking a weighted average, with weights given by a logistic regression fit on the FS Lost and Found validation set, captures the interaction between the layers.

Table 1 Benchmark results

method	score	FS Lost & Found		FS Web Oct 20		FS Web Jan 20		FS Static		re-training	OoD data	Cityscapes mIoU	time [s]
		AP \uparrow	FPR ₉₅ \downarrow	AP \uparrow	FPR ₉₅ \downarrow	AP \uparrow	FPR ₉₅ \downarrow	AP \uparrow	FPR ₉₅ \downarrow				
Random	random uncertainty	00.3	95.0	01.6	95.0	01.5	95.0	02.5	95.0	✗	✗	80.3	-
Softmax	max-probability	01.8	44.8	11.8	40.1	11.1	37.8	12.9	39.8	✗	✗	80.3	0.05
	entropy	02.9	44.8	16.6	39.8	15.6	37.5	15.4	39.8				0.29
kNN Embedding	density	03.5	30.0	24.9	35.2	22.3	31.7	44.0	20.2	✗	✗	80.3	7.89
	relative class density	00.8	100.0	09.6	100.0	8.3	100.0	15.8	100.0				6.45
Image Resynthesis	resynthesis difference	05.7	48.1	12.5	51.3			29.6	27.1	✗	✗	81.4	2.39
Learned Embedding Density (ours)	single-layer NLL	03.0	32.9	21.4	44.0	18.8	41.1	40.9	21.3		✗		0.29
	minimum NLL	04.3	47.2	32.6	55.8	30.2	49.5	62.1	17.4	✗	✗	80.3	1.53
	logistic regression	04.7	24.4	29.2	38.8	26.0	33.8	57.2	13.4		✓		1.53
SynBoost	dissimilarity score	43.2	15.8	61.3	18.9			72.6	18.8	✗	✓	81.4	2.32
Bayesian DeepLab	mutual information	<i>09.8</i>	38.5	<i>35.8</i>	25.7	<i>33.8</i>	<i>19.7</i>	48.7	15.5	✓	✗	73.8	3.62
	fixed patches	15.7	76.9	65.3	19.6	61.4	28.2	82.9	05.1			77.7	0.12
	random patches	21.2	36.9	67.2	10.3			86.2	02.4	✓	✓	77.3	0.12
Outlier Head	combined probability	30.9	22.2	64.0	18.8			84.0	10.3			77.3	0.12
	max-entropy	01.7	30.6	28.2	22.6	28.7	18.4	27.5	23.6	✓	✓	79.0	0.29
OoD training	void classifier	10.3	22.1	43.0	14.0	42.8	12.0	45.0	19.4		✓	70.4	0.05
	prior entropy	34.3	47.4	30.0	76.6	33.9	75.5	31.3	84.6	✓	✓	70.5	0.05

Gray columns represent the primary metric of the benchmark. Methods are only evaluated on those FS Web datasets with object images appearing on the web after their submission date. For every metric and dataset, the best performance is marked bold and the best performance without OoD training is marked italic

4.3 Submitted Methods

The following methods were submitted to our benchmark since it went online in August 2019. They were not implemented or trained by us, but we include an overview since they are part of the benchmark results.

An *outlier head* can be added in a multi-task fashion to many semantic segmentation architectures. Bevandic et al. (2019) trains the head in a supervised fashion on both ID and OoD data samples. The training is executed simultaneously with the segmentation training. The outlier detection head then returns a pixel-wise anomaly score. Submitted were three variants of this method where the exact descriptions are in submission for publication.

Image resynthesis uses reconstruction to estimate the fit of an input to the training data distribution of a generative model. While auto-encoders such as described in Sect. 2 scale poorly to the level of detail in urban driving, good results have been achieved with generative adversarial networks (Wang et al., 2018; Isola et al., 2017) that synthesize driving scenes from semantic segmentation. Lis et al. (2019) uses such a method to find outliers by comparing the original and resynthesized image, where they train the comparison on flipped semantic labels in the ID data and therefore do not require outliers in training. While the original work (Lis et al., 2019) experimented with lower resolution segmentation data, Di Biase et al. (2021) submitted an adapted, scaled-up model.

Synboost is a modular approach that combines introspective uncertainties and input reconstruction into a pixel-wise dissimilarity score. Further details are described in Di Biase et al. (2021).

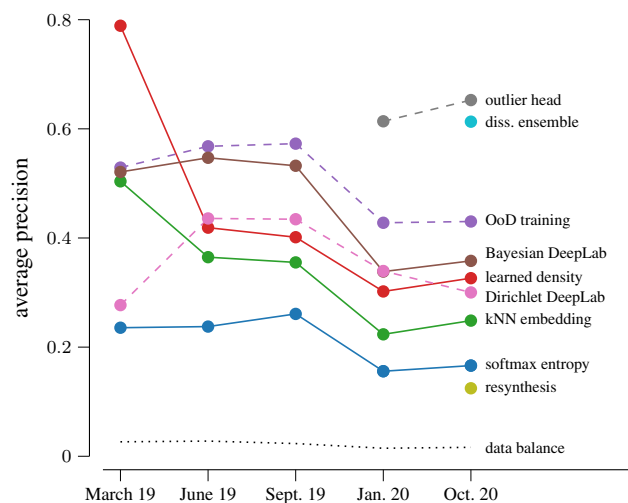


Fig. 4 Performance evolution over the different iterations of the FS Web dataset. We only plot the best-performing variant of each method. Methods that train on OoD data are plotted with dashed lines. Notable changes are the better blending method in June 19 and the inclusion of blended ID objects in January 20, which changed the data-balance.

5 Discussion of Results

We show in Table 1 the results of our benchmark as of December 2020 for the aforementioned datasets and methods. Qualitative examples of all methods are shown in Fig. 5.

Softmax confidence Confirming findings on simpler tasks (Lee et al., 2018), the softmax confidence is not a reliable score for anomaly detection. While training with OoD data clearly improves the softmax-based detection, it is not much




Fig. 5 Successful and failed examples for all methods on the Fishyscapes Lost and Found dataset. Input images overlaid with the evaluation labels are on the left, predicted anomaly scores on the right of each exam-

ple pair. For every method, we show the best variant. The red circles highlight anomalies that are missed by the method or indistinguishable from noise.

better than Bayesian DeepLab, that does not require such data.

Difference between datasets For most methods, there is a clear performance gap between the data from Lost and Found and the other datasets. We attribute this to two factors. First, the dataset contains a lot of images with only very small objects. This is indicated by the AP of the random classifier, which equals to the fraction of anomalous pixels. Second, the qualitative examples show the more challenging nature of the Lost and Found dataset with e.g. false positives for the void classifier or outlier head, and cases where small anomalous objects are not detected at all e.g. for the Bayesian DeepLab or Softmax Entropy.

We further investigate the results on FS Web over time in Fig. 4. While most methods follow overall trends that can be attributed to the difficulty of the individual objects or differences in data balance, it becomes clear that (i) embedding based methods were picking up blending artifacts in FS Web March 2019, and (ii) Dirichlet DeepLab is performing very inconsistently. (i) appears to be fixed with the advanced blending from June 2019, since the introduction of blended ID objects did not have any effect on embedding based methods. (ii) could indicate a degree of overfitting to specific object types, because Dirichlet DeepLab is trained on OoD data.

Semantic segmentation accuracy The data in table 1 illustrates a tradeoff between anomaly detection and segmentation performance. Methods like Bayesian DeepLab or Outlier Head are consistently among the best methods on all datasets, but need to train with special losses that reduce the segmentation accuracy by up to 10%. If segmentation accuracy is important, methods that do not require any retraining are particularly interesting.

Supervision with OoD data appears to be important for good anomaly detection. On every dataset, the best method required OoD data and is at least 38% better than any ‘unsupervised’ method. While training with OoD data can in principle lead to overfitting to specific objects, the results on FS Web, which was designed specifically to resemble open-world settings, show that the Outlier Head or Dissimilarity Ensemble are very robust to diverse anomalies.

We however want to emphasize that anomaly detection and uncertainty estimation are very different principles. Our benchmark therefore serves the dual purpose of finding either the best anomaly segmentation method or well-scalable uncertainty estimates, that are simply tested on the proxy task of anomaly detection. Comparing Bayesian DeepLab and the void classifier shows that good uncertainty estimation methods can even compete with some supervised methods, but so far not with specifically designed anomaly segmentation methods.

Inference time differs significantly between methods. Methods can be broadly sorted into two categories, where the first do a single pass through a (sometimes modified) DeepLabv3+ architecture and the second category applies additional processing on top of this forward pass. Our measurements show that methods in the second category have up to two orders of magnitude higher inference time. The only exception marks the single-layer embedding density, where inference time is comparable to single pass methods. While nearly all methods³ were executed as optimised tensorflow graphs, measurements are still dependent on the implementation details and possible parallelization is limited by GPU memory constraints. For example, the difference between softmax max-prob, softmax entropy, and dirichlet entropy can only be explained with inefficiencies in the softmax entropy implementation that cause a difference of more than 0.2 s.

Challenges in method adaptation The results reveal that some methods cannot be easily adapted to semantic segmentation. For example, retraining required by special losses can impair the segmentation performance, and we found that these losses (e.g. for Dirichlet DeepLab) were often unstable during training or did not converge. Other challenges rise from the complex network structures which complicate the translation of class-based embedding methods such as deep

k-nearest neighbor (Papernot & McDaniel, 2018) to segmentation. This is illustrated by the performance of our simple implementation.

6 Conclusion

In this work, we introduced *Fishyscapes*, a benchmark for anomaly detection in semantic segmentation for urban driving. Comparing state-of-the-art methods on this complex task for the first time, we draw multiple conclusions:

- The softmax output from a standard classifier is a bad indicator for anomaly detection.
- Most of the better performing methods required special losses that reduce the semantic segmentation accuracy.
- Supervision of anomaly segmentation methods with OoD data consistently outperformed unsupervised methods even in open-world scenarios.

Overall, the methods compared in our benchmark so far leave a lot of room for improvement. To safely deploy semantic segmentation methods in autonomous cars, further research is required. As a public benchmark, *Fishyscapes* supports the evaluation of new methods on urban driving scenarios.

Funding Open Access funding provided by ETH Zurich.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix

Here we provide additional experimental evaluations as well as details on the proposed datasets and the evaluated methods.

A Misclassification Detection

Additionally to anomaly detection, we test some methods on the detection of misclassifications from the semantic segmentation output. Misclassification detection is another proxy

³ Image Resynthesis and SynBoost were submitted as pytorch models.

classification task that correlates with uncertainty. However, misclassification mixes uncertainty from

- noise in the input (*aleatoric* uncertainty)
- model uncertainty
- shifts in data balance (softmax classification implicitly learns a prior distribution of the classes over the training set)

Nevertheless, failure detection is an important problem for deployment on autonomous agents, e.g. as part of sensor fusion mechanisms, and misclassification detection is used in different related work (Hendrycks & Gimpel, 2017; Lakshminarayanan et al., 2017; Guo et al., 2017; Jiang et al., 2018) to benchmark uncertainty estimates.

Dataset We test misclassification detection on a diverse mixture of different data sources that introduce sources of uncertainty in the input. From Sakaridis et al. (2018), we select all images. From Dai et al. (2020), we map classes *sky* and *fence* to *void*, as their labelling is not accurate and sometimes areas that are not visible due to fog are simply labelled *sky*. For WildDash (Zendel et al. 2018), we use all images. For Mapillary Vistas (Neuhold et al., 2017), we sample 50

Table 2 Mapping of mapillary classes onto our used set of classes for misclassification detection

Mapillary label	Used label
<i>construction–barrier–fence</i>	<i>fence</i>
<i>construction–barrier–wall</i>	<i>wall</i>
<i>construction–flat–road</i>	<i>road</i>
<i>construction–flat–sidewalk</i>	<i>sidewalk</i>
<i>construction–structure–building</i>	<i>building</i>
<i>human–person</i>	<i>person</i>
<i>human–rider–*</i>	<i>rider</i>
<i>nature–sky</i>	<i>sky</i>
<i>nature–terrain</i>	<i>terrain</i>
<i>nature–vegetation</i>	<i>vegetation</i>
<i>object–support–pole</i>	<i>pole</i>
<i>object–support–utility–pole</i>	<i>pole</i>
<i>object–traffic–light</i>	<i>traffic light</i>
<i>object–traffic–sign–front</i>	<i>traffic sign</i>
<i>object–vehicle–bicycle</i>	<i>bicycle</i>
<i>object–vehicle–bus</i>	<i>bus</i>
<i>object–vehicle–car</i>	<i>car</i>
<i>object–vehicle–motorcycle</i>	<i>motorcycle</i>
<i>object–vehicle–on–rails</i>	<i>train</i>
<i>object–vehicle–truck</i>	<i>truck</i>
<i>marking–*</i>	<i>road</i>
<i>anything else</i>	<i>void</i>

Table 3 Misclassification detection results

method	score	FS Misclassification		
		max J ↑	AP ↑	mIoU
Baseline	random uncertainty	00.0	38.9	45.5
Softmax	max-probability	43.6	67.4	45.5
	entropy	43.5	68.4	
OoD training	max-entropy	44.3	71.3	35.8
Bayesian DeepLab	mutual information	40.7	70.4	30.3
	predictive entropy	41.6	73.8	
Dirichlet DeepLab	prior entropy	29.7	65.0	37.5
kNN Embedding	density	40.7	68.0	45.5
	relative class density	31.7	58.0	

Gray column represents the primary metric

random images from the validation set and apply the label mapping described in Table 2.

During evaluation all pixels labelled as *void* are ignored.

Evaluated Methods From the methods evaluated on anomaly detection, we note that the void classifier produces meaningless results for misclassification detection since a high void output score produces the exact misclassification it is detecting. Furthermore, we did not evaluate the learned embedding density.

Results of our evaluation are presented in Table 3 and qualitative examples in Fig. 6. Differently from anomaly detection, the softmax score is expected to be a good indicator for classification uncertainty, and indeed shows competitive results. For Bayesian DeepLab, we find the predictive entropy to be a better indicator of misclassification, which was also observed by Kendall and Gal (2017). The kNN density shows results similar to the other methods, hinting that embedding-based methods cannot be entirely classified as OoD-specific, but may also be able to detect input noise that is very different from the training distribution. Overall, the experiments do not reveal a single method that performs significantly better than others.

B Details on the Methods

In this section we provide implementation details on the evaluated methods to ease the reproducibility of the results presented in this paper.

B.1 Semantic Segmentation Model

We use the state-of-the-art model DeepLabv3+ (Chen et al., 2018) with Xception-71 backbone, image-level features, and dense prediction cell. When no retraining is required, we use the original model trained on Cityscapes⁴.

⁴ <https://github.com/tensorflow/models/blob/master/research/deeplab>.

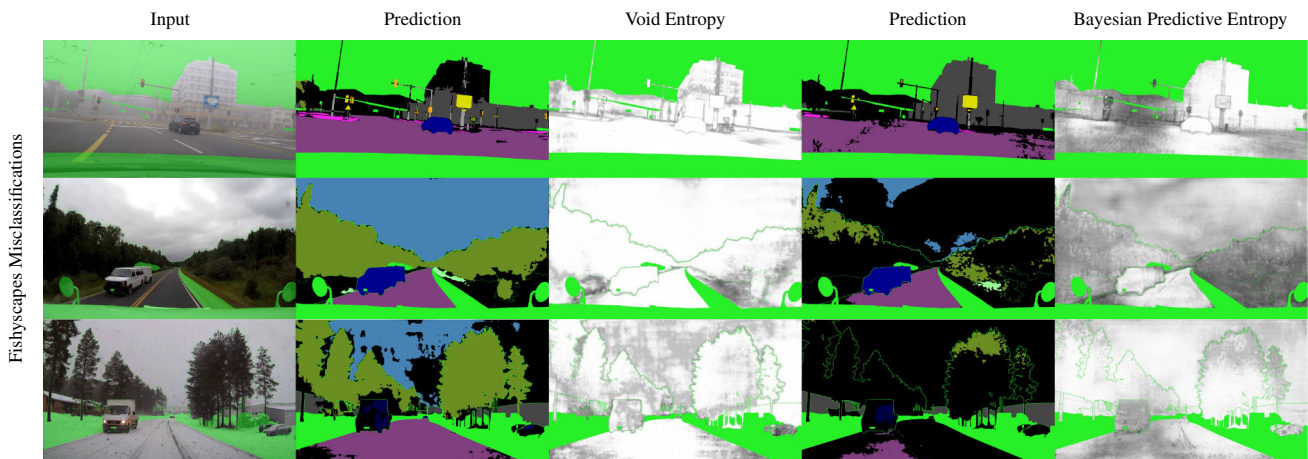


Fig. 6 Qualitative examples of misclassification detection. Predictions correspond to the uncertainty maps to their right. Misclassifications are marked in black, while ignored *void* pixels are marked in bright green. Better methods should assign a high score (dark) to misclassified

pixels. While the different trainings clearly lead to different classification performances, none of the methods captures all the misclassified pixels.

B.2 Softmax

ODIN (Liang et al., 2018) applies input preprocessing and temperature scaling to improve the OoD detection ability of the maximum softmax probability. Early experiments on Fishyscapes showed that (i) temperature scaling did not improve much the results of this baseline, and (ii) input preprocessing w.r.t. the softmax score is not possible due to the limited GPU memory and the large size of the DeepLab model. As the maximum probability is anyway not competitive with respect to the other methods, we decided to not further develop that baseline.

B.3 Bayesian DeepLab

We reproduce the setup described by Mukhoti and Gal (2018). As such, we use the Xception-65 backbone pretrained on ImageNet, and insert dropout layers in its middle flow. We train for 90k iterations, with a batch size of 16, a crop size of 513×513 , and a learning rate of $7 \cdot 10^{-3}$ with polynomial decay.

B.4 Dirichlet DeepLab

Following Malinin and Gales (2018), we interpret the output logits of DeepLab as log-concentration parameters α and train with the loss described by Eq. (3) and implemented with the TensorFlow Probability (Dillon et al., 2017) framework. For the first term, the target labels are smoothed with $\epsilon = 0.01$ and scaled by $\alpha_0 = 100$ to obtain target concentrations. To ensure convergence of the classifier, we found it necessary to downweight both the first and second terms by 0.1 and

to initialize all but the last layer with the original DeepLab weights.

We also tried to replace the first term by the negative log-likelihood of the Dirichlet distribution but were unable to make the training converge.

B.5 kNN Embedding

Layer of Embedding. As explained in Sect. 4.1, we had to restrict the kNN queries to one layer. A single layer of the network already has more than 10,000 embedding vectors and we need to find k nearest neighbors for all of them. Querying over multiple layers therefore becomes infeasible. To select a layer of the network, we test multiple candidates on the FS Lost and Found validation set. We experienced that our kNN fitting with *hnsplib*⁵ (Malkov & Yashunin, 2018) was not deterministic, therefore we provide the average performance on the validation set over 3 different experiments. Additionally, we had to reduce the complexity of kNN fitting by randomly sampling 1000 images from Cityscapes instead of the whole training set (2975 images).

For the kNN density, we provide the results for different layers in Table 4.

For class-based embedding, we perform a similar search for the choice of layer. The result can be found in Table 5.

Number of Neighbors We select k according to Tables 6 and 7. All values are measured with the same kNN fitting. As the computational time for each query grows with k , small values are preferable. Note that by definition, the relative class density needs a sufficiently high k such that not all neighbors are from the same class.

⁵ <https://github.com/nmslib/hnsplib>.

Table 4 Parameter search of the embedding layer for kNN density

DeepLab layer	(AP)
xception_71/middle_flow/block1/unit_8	1.00 ± .02
xception_71/exit_flow/block2	1.80 ± .01
aspp_features	2.97 ± .47
decoder_conv0_0	3.84 ± .19
decoder_conv1_0	2.46 ± .09

The AP is computed on the validation set of FS Lost and Found. Based on these results, we use the layer decoder_conv0_0 in all our experiments

Table 5 Parameter search of the embedding layer for class based relative kNN density

DeepLab layer	(AP)
xception_71/middle_flow/block1/unit_8	9.6 ± .0
xception_71/middle_flow/block1/unit_10	9.7 ± .0
xception_71/exit_flow/block2	9.7 ± .1
aspp_features	2.3 ± .7
decoder_conv0_0	2.8 ± .1
decoder_conv1_0	3.1 ± .2

The AP is computed on the validation set of FS Static. Based on these results, we use the layer xception_71/exit_flow/block2 in all our experiments

Table 6 Parameter search for the number of nearest neighbors for kNN embedding density

k	(AP)
1	42.3
2	44.6
5	47.7
10	50.9
20	52.2
50	52.7
100	52.5

As computing time increases with k , we select $k = 20$

Table 7 Parameter search for the number of nearest neighbors for the class based kNN relative density

k	(AP)
5	5.4
10	6.7
20	7.9
50	9.3
100	9.9
200	10.0

As computing time increases with k , we select $k = 100$

B.6 Learned Embedding Density

Flow architecture The normalizing flow follows the simple architecture of Real-NVP. We stack 32 steps, each one

composed of an affine coupling layer, a batch normalization layer, and a fixed random permutation. As recommended by Kingma and Dhariwal (2018), we initialize the weights of the coupling layers such that they initially perform identity transformations.

Flow training For a given DeepLab layer, we export the embeddings computed on all the images of the Cityscapes training set. The number of such datapoints depends on the stride of the layer, and amounts to 22M for a stride of 16. We keep 2000 of them for validation and testing, and train on the remaining embeddings for 200k iterations, with a learning rate of 10^{-4} , and the Adam optimizer. Note that we can compare flow models based on how well they fit the in-distribution embeddings, and thus do not require any OoD data for hyperparameter search.

Layer selection OoD data is only required to select the layer at which the embeddings are extracted. The corresponding feature space should best separate OoD and ID data, such that OoD embeddings are assigned low likelihood. We found that it is critical to extract embeddings before ReLU activations, as some dimensions might be negative for all training points, thus making the training highly unstable. We show in Table 8 the (AP) on the FS Lost and Found validation set for different layers. We first observe that we did not achieve training convergence for those layers that showed best results in the kNN method. This may be due to the high dimensionality of these layers, and/or because the flow is not well suited to approximate these distributions. We also notice that overall layers in the encoder middle flow work best, while Mukhoti and Gal (2018) insert dropout layers at this particular stage. While we do not know the reason behind their design decision, we hypothesize they found these layers to best model the epistemic uncertainty.

Table 8 Cross-validation of the embedding layer for the learned density

DeepLab layer	AP
xception_71/entry_flow/block5	1.27
xception_71/middle_flow/block1/unit_4	2.14
xception_71/middle_flow/block1/unit_6	2.38
xception_71/middle_flow/block1/unit_8	2.41
xception_71/middle_flow/block1/unit_10	2.52
xception_71/middle_flow/block1/unit_12	2.22
aspp_features	–
decoder_conv0_0	0.16
decoder_conv1_0	2.77

The AP is computed on the validation set of FS Lost and Found. Based on these results, we use the layer decoder_conv1_0 in all our experiments. We could not manage to make the training of the aspp_features layer converge, most likely due to a very peaky distribution that induces numerical instabilities

Table 9 Cross-validation of the input preprocessing for the learned density

Noise ϵ	AP on FS Static validation	Static test
None	36.0	52.5
0.1	38.4	–
0.2	39.1	–
0.25	39.2	55.4
0.3	39.2	–
0.35	39.2	–
0.4	39.1	–
0.5	39.0	–
1.0	36.6	–

Based on these results, we apply noise with magnitude $\epsilon = 0.25$ in all our experiments

Effect of input preprocessing As previously reported by Liang et al. (2018; 2018), we observe that this simple input preprocessing brings substantial improvements to the detection score on the test set. We show in Table 9 the AP for different noise magnitudes ϵ .

References

- Abu Alhaija, H., Mustikovela, S. K., Mescheder, L., Geiger, A., & Rother, C. (2018). Augmented reality meets computer vision: Efficient data generation for urban driving scenes. *International Journal of Computer Vision*, 126(9), 961–972. <https://doi.org/10.1007/s11263-018-1070-x>.
- Antorán, J., Allingham, J., & Hernández-Lobato, J. M. (2020). Depth uncertainty in neural networks. *Advances in neural information processing systems*, 33.
- Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). SegNet: A deep convolutional encoder–decoder architecture for image segmentation. *IEEE TPAMI*, 39(12), 2481–2495. <https://doi.org/10.1109/TPAMI.2016.2644615>.
- Bergmann, P., Fauser, M., Sattlegger, D., & Steger, C. (2019). MVTEC AD—A comprehensive Real-World dataset for unsupervised anomaly detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 9592–9600).
- Bevardić, P., Krešo, I., Oršić, M., & Šegvić, S. (2019). Simultaneous semantic segmentation and outlier detection in presence of domain shift. In *German conference on pattern recognition* (pp. 33–47).
- Bozhinoski, D., Di Ruscio, D., Malavolta, I., Pelliccione, P., & Crnkovic, I. (2019). Safety for mobile robotic systems: A systematic mapping study from a software engineering perspective. *Journal of System Software*, 151, 150–179. <https://doi.org/10.1016/j.jss.2019.02.021>.
- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., & Beijbom, O. (2020). Nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 11621–11631).
- Chen, L. C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2016). DeepLab. *IEEE TPAMI*, <https://doi.org/10.1109/TPAMI.2017.2699184>.
- Chen, L. C., Zhu, Y., Papandreou, G., Schroff, F., & Adam, H. (2018). Encoder–decoder with atrous separable convolution for semantic image segmentation. *ECCV*, <https://doi.org/10.1007/978-3-030-01234-249>.
- Choi, H., Jang, E., & Alemi, A. A. (2018). WAIC, but why? Generative ensembles for robust anomaly detection (Preprint). <https://arxiv.org/abs/1810.01392>.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., et al. (2016). The cityscapes dataset for semantic urban scene understanding. *CVPR*, <https://doi.org/10.1109/CVPR.2016.350>.
- Dai, D., Sakaridis, C., Hecker, S., & Van Gool, L. (2020). Curriculum model adaptation with synthetic and real data for semantic foggy scene understanding. *International Journal of Computer Vision*, 128(5), 1182–1204.
- DeVries, T., & Taylor, G. W. (2018). Learning confidence for out-of-distribution detection in neural networks.
- Di Biase, G., Blum, H., Siegwart, R., & Cadena, C. (2021). Pixel-wise anomaly detection in complex driving scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 16918–16927).
- Dillon, J. V., Langmore, I., Tran, D., Brevdo, E., Vasudevan, S., Moore, D., Patton, B., Alemi, A., Hoffman, M., & Saurous, R. A. (2017). Tensorflow distributions.
- Dinh, L., Krueger, D., Bengio, Y. (2014). NICE: Non-linear independent components estimation (Preprint). <https://arxiv.org/abs/1410.8516>.
- Dinh, L., Sohl-Dickstein, J., & Bengio, S. (2017). Density estimation using real NVP. In *ICLR*.
- Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., & Zisserman, A. (2010). The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2), 303–338. <https://doi.org/10.1007/s11263-009-0275-4>.
- Florence, P. R., Manuelli, L., & Tedrake, R. (2018). Dense object nets: Learning dense visual object descriptors by and for robotic manipulation. In: A. Billard, A. Dragan, J. Peters, J. Morimoto (eds.) *Conference on robot learning (CoRL), proceedings of machine learning research* (Vol. 87, pp. 373–385). PMLR.
- Frey, B. J., & Hinton, G. E. (1999). Variational learning in nonlinear gaussian belief networks. *Neural Computation*, <https://doi.org/10.1162/089976699300016872>.
- Fu, H., Gong, M., Wang, C., Batmanghelich, K., & Tao, D. (2018). Deep ordinal regression network for monocular depth estimation. In *CVPR* (pp. 2002–2011). IEEE. 10.1109/CVPR.2018.00214.
- Gal, Y. (2016). Uncertainty in deep learning. Ph.D. thesis, University of Cambridge.
- Gal, Y., Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: M. F. Balcan, K. Q. Weinberger (Eds.), *ICML, proceedings of machine learning research* (Vol. 48, pp. 1050–1059). New York: PMLR.
- Geiger, A., Lenz, P., & Urtasun, R. (2012). Are we ready for autonomous driving? The KITTI vision benchmark suite. In: *CVPR* (pp. 3354–3361). ieeexplore.ieee.org. 10.1109/CVPR.2012.6248074.
- Geyer, J., Kassahun, Y., Mahmudi, M., Ricou, X., Durgesh, R., Chung, A. S., Hauswald, L., Pham, V. H., Mühlegg, M., Dorn, S., Fernandez, T., Jänicke, M., Mirashi, S., Savani, C., Sturm, M., Vorobiov, O., Oelker, M., Garreis, S., Schuberth, P. (2020). A2D2: Audi autonomous driving dataset (Preprint). <https://arxiv.org/abs/2004.06320>.
- Golan, I., & El-Yaniv, R. (2018). Deep anomaly detection using geometric transformations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *NeurIPS* (pp. 9758–9769). Curran Associates Inc.
- Gong, D., Liu, L., Le, V., Saha, B., Mansour, M. R., Venkatesh, S., & van den Hengel, A. (2019). Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly

- detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 9592–9600).
- Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *ICLR*.
- Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In: D. Precup, Y. W. Teh (Eds.), *ICML, proceedings of machine learning research* (Vol. 70, pp. 1321–1330). PMLR, International Convention Centre.
- He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2017). Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision* (pp. 2961–2969).
- Hendrycks, D., Basart, S., Mazeika, M., Mostajabi, M., Steinhardt, J., & Song, D. (2019). A benchmark for anomaly segmentation (Preprint). <https://arxiv.org/abs/1911.11132>.
- Hendrycks, D., & Gimpel, K. (2017). A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*.
- Isola, P., Zhu, J. Y., Zhou, T., Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 1125–1134). IEEE. 10.1109/cvpr.2017.632.
- Jiang, H., Kim, B., Guan, M., & Gupta, M. (2018). To trust or not to trust a classifier. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 31, pp. 5545–5556). Curran Associates Inc.
- Kendall, A., Gal, Y. (2017). What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS* (pp. 5574–5584).
- Kingma, D. P., & Dhariwal, P. (2018). Glow: Generative flow with invertible 1x1 convolutions. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *NeurIPS* (pp. 10236–10245). Curran Associates Inc.
- Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *NIPS* (pp. 6402–6413).
- Lee, K., Lee, H., Lee, K., & Shin, J. (2018). Training confidence-calibrated classifiers for detecting Out-of-Distribution samples. In *ICLR*.
- Lee, K., Lee, K., Lee, H., & Shin, J. (2018). A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *NeurIPS*.
- Liang, M., Yang, B., Wang, S., & Urtasun, R. (2018). Deep continuous fusion for multi-sensor 3D object detection. In *ECCV* (pp. 663–678). Springer. 10.1007/978-3-030-01270-039.
- Liang, S., Li, Y., & Srikant, R. (2018). Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*.
- Lis, K., Nakka, K., Fua, P., & Salzmann, M. (2019). Detecting the unexpected via image resynthesis. In *Proceedings of the IEEE international conference on computer vision* (pp. 2152–2161).
- Loquercio, A., Segu, M., & Scaramuzza, D. (2020). A general framework for uncertainty estimation in deep learning. *IEEE Robotics and Automation Letters*, 5(2), 3153–3160. <https://doi.org/10.1109/LRA.2020.2974682>.
- Malinin, A., & Gales, M. (2018). Predictive uncertainty estimation via prior networks. In *NeurIPS*.
- Malkov, Y. A., & Yashunin, D. A. (2018). Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(4), 824–836.
- Mandelbaum, A., & Weinshall, D. (2017). Distance-based confidence score for neural network classifiers (Preprint). <https://arxiv.org/abs/1709.09844>.
- Mccormac, J., Clark, R., Bloesch, M., Davison, A., & Leutenegger, S. (2018). Fusion++: Volumetric Object-Level SLAM. In *International conference on 3D vision (3DV)* (pp. 32–41). IEEE. 10.1109/3DV.2018.00015.
- Mukhoti, J., & Gal, Y. (2018). Evaluating bayesian deep learning methods for semantic segmentation (Preprint). <https://arxiv.org/abs/1811.12709>.
- Nalisnick, E., Matsukawa, A., Teh, Y. W., Gorur, D., & Lakshminarayanan, B. (2019). Do deep generative models know what they don't know? In *ICLR*.
- Neuhold, G., Ollmann, T., Bulo, S. R., & Kotschieder, P. (2017). The papillary vistas dataset for semantic understanding of street scenes. In *ICCV* (pp. 5000–5009). IEEE. 10.1109/ICCV.2017.534.
- Papernot, N., & McDaniel, P. (2018). Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning (Preprint). <https://arxiv.org/abs/1803.04765>.
- Pidhorskyi, S., Almohsen, R., & Doretto, G. (2018). Generative probabilistic novelty detection with adversarial autoencoders. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), *NeurIPS* (pp. 6822–6833).
- Pinggera, P., Ramos, S., Gehrig, S., Franke, U., Rother, C., & Mester, R. (2016). Lost and found: Detecting small road hazards for self-driving vehicles. In *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 1099–1106). IEEE. 10.1109/IROS.2016.7759186.
- Postels, J., Blum, H., Cadena, C., Siegwart, R., Van Gool, L., & Tombari, F. (2020). Quantifying aleatoric and epistemic uncertainty using density estimation in latent space (Preprint). <https://arxiv.org/abs/2012.03082>.
- Richter, C., & Roy, N. (2017). Safe visual navigation via deep learning and novelty detection. In *Robotics: Science and systems (RSS)*. Robotics: Science and Systems Foundation. 10.15607/RSS.2017.XIII.064.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. https://doi.org/10.1007/978-3-319-24574-4_28.
- Ruff, L., Vandermeeulen, R., Goernitz, N., Deecke, L., Siddiqui, S. A., Binder, A., Müller, E., Kloft, M. (2018). Deep one-class classification. In J. Dy, A. Krause (Eds.), *ICML, proceedings of machine learning research* (Vol. 80, pp. 4393–4402). PMLR, Stockholm. <http://proceedings.mlr.press/v80/ruff18a.html>.
- Sabokrou, M., Khalooei, M., Fathy, M., & Adeli, E. (2018). Adversarially learned one-class classifier for novelty detection. *CVPR*. <https://doi.org/10.1109/cvpr.2018.00356>.
- Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLoS ONE*, 10(3), e0118432. <https://doi.org/10.1371/journal.pone.0118432>.
- Sakaridis, C., Dai, D., Hecker, S., & Van Gool, L. (2018). Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *ECCV* (pp. 707–724). Springer. 10.1007/978-3-030-01261-8_42.
- Sakaridis, C., Dai, D., & Van Gool, L. (2018). Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126(9), 973–992. <https://doi.org/10.1007/s11263-018-1072-8>.
- Sakaridis, C., Dai, D., & Van Gool, L. (2020). Map-guided curriculum domain adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. <https://doi.org/10.1109/TPAMI.2020.3045882>.
- Sun, D., Yang, X., Liu, M., & Kautz, J. (2018). PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. *CVPR*. <https://doi.org/10.1109/CVPR.2018.00931>.
- Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al. (2020). Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2446–2454).

- Van Amersfoort, J., Smith, L., Teh, Y. W., & Gal, Y. (2020). Uncertainty estimation using a single deep deterministic neural network. In *International conference on machine learning* (pp. 9690–9700).
- Wang, T. C., Liu, M. Y., Zhu, J. Y., Tao, A., Kautz, J., & Catanzaro, B. (2018). High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 8798–8807).
- Wilson, A. G., & Izmailov, P. (2020). Bayesian deep learning and a probabilistic perspective of generalization (Preprint). <https://arxiv.org/abs/2002.08791>.
- Yehezkel Rohekar, R., Gurwicz, Y., Nisimov, S., & Novik, G. (2019). Modeling uncertainty by learning a hierarchy of deep neural connections. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 32, pp. 4244–4254). Curran Associates, Inc.
- Zendel, O., Honauer, K., Murschitz, M., Steininger, D., & Fernandez Dominguez, G. (2018). Willdash-creating hazard-aware benchmarks. In *ECCV* (pp. 402–416). [openaccess.thecvf.com.10.1007/978-3-030-01231-125](https://openaccess.thecvf.com/10.1007/978-3-030-01231-125).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.