KI DELTA LEARNING
Scalable AI for Automated Driving

Towards Unsupervised Open World Semantic Segmentation

Svenja Uhlemeyer, Matthias Rottmann, Hanno Gottschalk

Semi- and Unsupervised Learning

Use-Case in Delta Learning

Open world semantic segmentation involves identifying pixels which belong to unknown objects and incrementally learning novel classes. It is desirable to perform such an incremental learning task in an unsupervised fashion, as labeling is expensive and time-consuming.

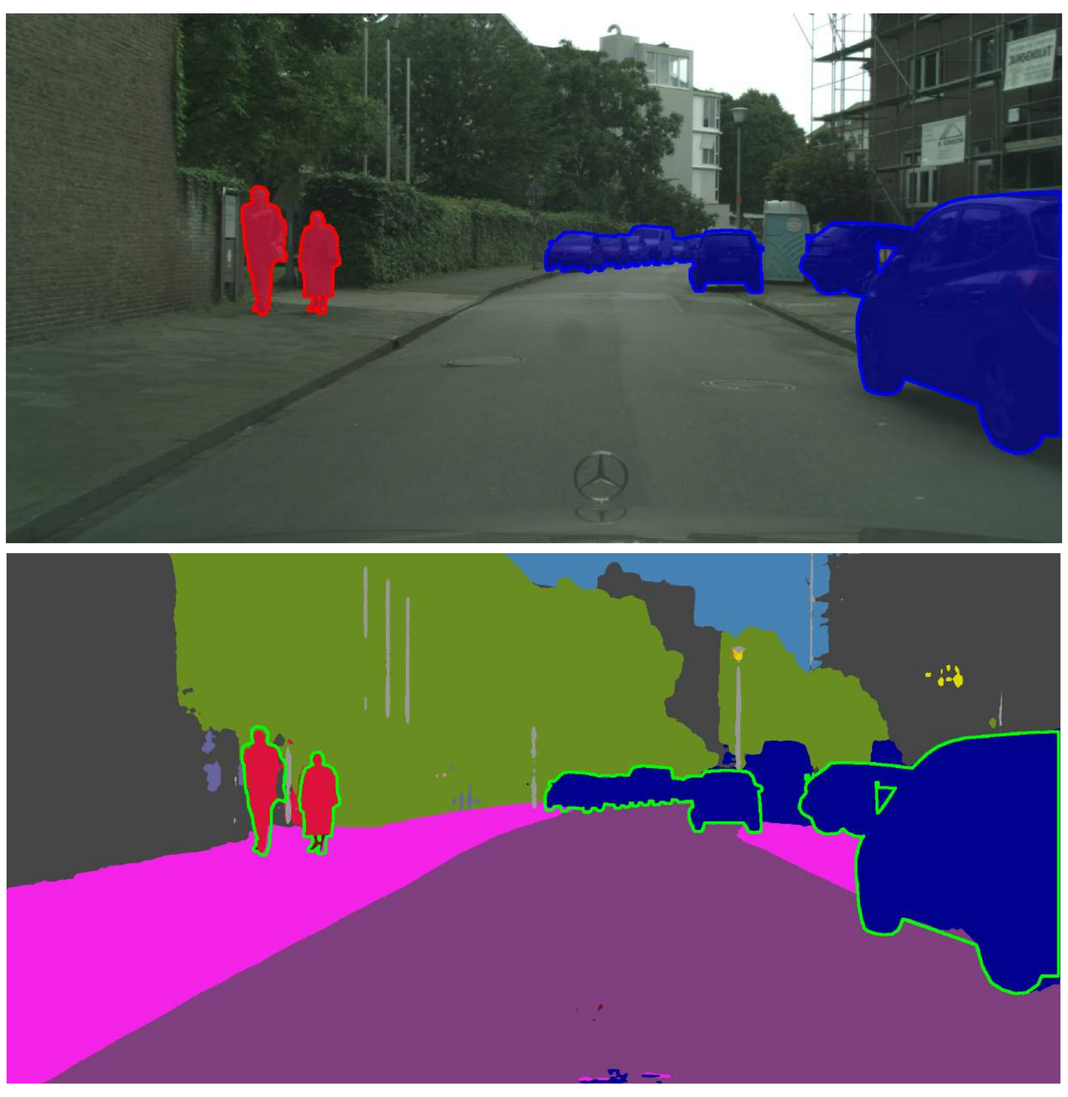

Figure 1: Incremental learning of humans and cars as novel classes without any ground truth annotations (© Cityscapes dataset)

Approach

Our proposed method consists of four steps:
(1) DNN are usually trained on a closed set of semantic classes and thus unable to recognize anomalous instances. We employ a post-processing meta regression model [1] to estimate the prediction quality of semantic segmentation masks. By thresholding on the predicted quality score, we obtain binary anomaly masks, wherein each connected component constitutes an anomaly.

(2) For comparability reasons, we feed image patches with anomalies into an image classification network which is trained on ImageNet and extract the features in the penultimate layer. To enable visualizability and transparency, we employ dimensionality reduction techniques to obtain a 2D embedding space as proposed in [2].

(3) We utilize clustering methods like DBSCAN for the discovery of novel classes in the 2D embedding space. Each cluster (of sufficient size) gets a unique class ID, and for each anomaly which belongs to this cluster, we generate pseudo labels as demonstrated in Figure 2.

(4) We extend the DNN by the number of recognized novel classes and fine-tune it on a mix of pseudo labeled test data and replayed training data to mitigate catastrophic forgetting. Therefore, we also apply knowledge distillation.

Evaluation

We performed six different experiments, where we demonstrate that, without access to ground truth and even with few data, a DNN's class space can be extended by novel classes, achieving considerable segmentation accuracy. E.g., in the experiment visualized in Figure 1, we obtain IoU scores of 40.22% and 81.27% for the novel classes human and car, respectively (see Figure 3).

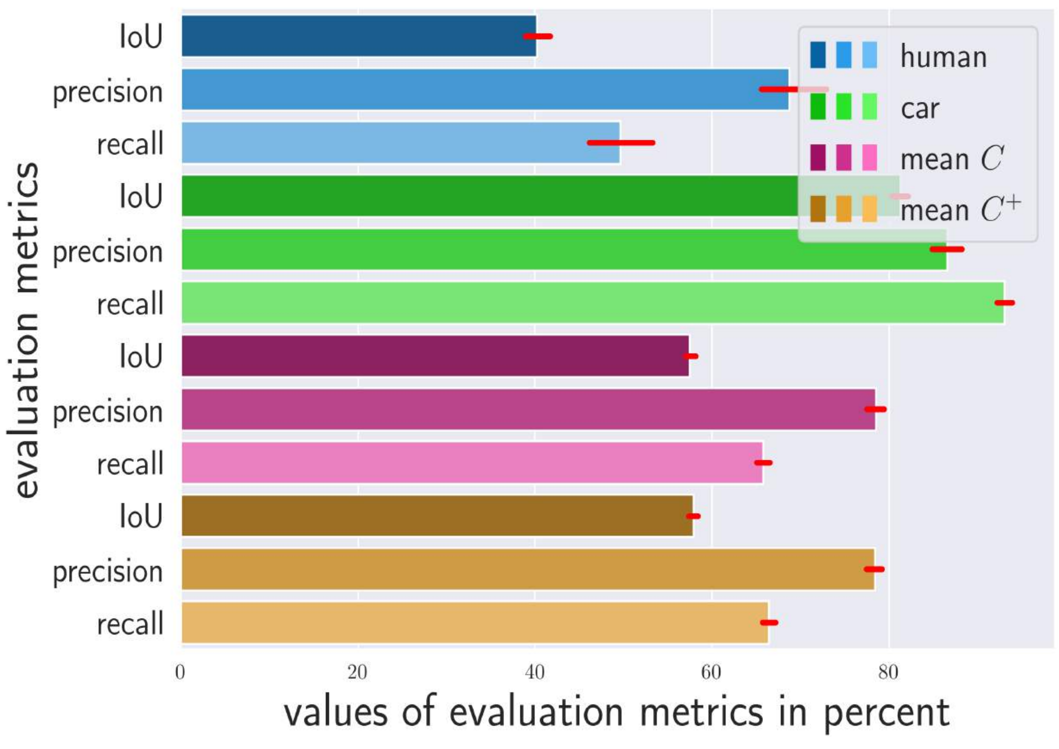

Figure 3: IoU, precision and recall values for incrementally earned classes human and car, as well as averaged over the old and over all classes

References:

[1] Matthias Rottmann et al., Prediction error meta classification in semantic segmentation: Detection via aggregated dispersion measures of softmax probabilities.
[2] Philipp Oberdiek et al., Detection and retrieval of out-of-distribution objects in semantic segmentation.

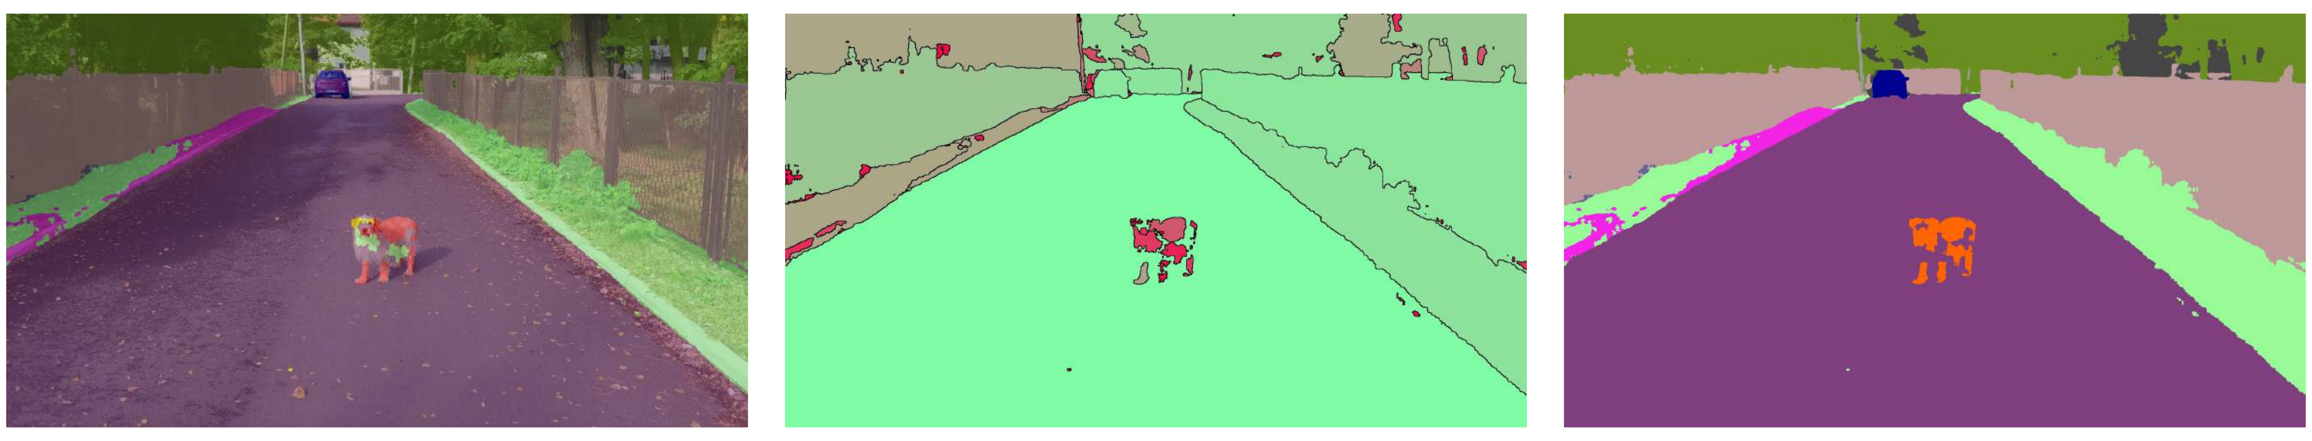

Figure 2: Pseudo labeling (right) as a combination of semantic segmentation (left) and anomaly segmentation (middle) (© SegmentMeIfYouCan, RoadAnomaly21 dataset)

Partners

External partners

For more information contact:
uhlemeyer@math.uni-wuppertal.de

code paper

KI Delta Learning is a project of the KI Familie. It was initiated and developed by the VDA Leitinitiative autonomous and connected driving and is funded by the Federal Ministry for Economic Affairs and Climate Action.

Supported by:
Federal Ministry for Economic Affairs and Climate Action
on the basis of a decision by the German Bundestag

KI FAMILIE

www.ki-deltalearning.de @KI_Familie KI Familie