KIDELTA LEARNING
Scalable AI for Automated Driving

Final Event | March 09, 2023

Towards Unsupervised Open World Semantic Segmentation

Discovery of Novelties and Class-Incremental Learning

How Can Neural Networks Discover the World?

- What is unknown?
- Have I seen something similar before?
- Give it a name and recognize it from now on.

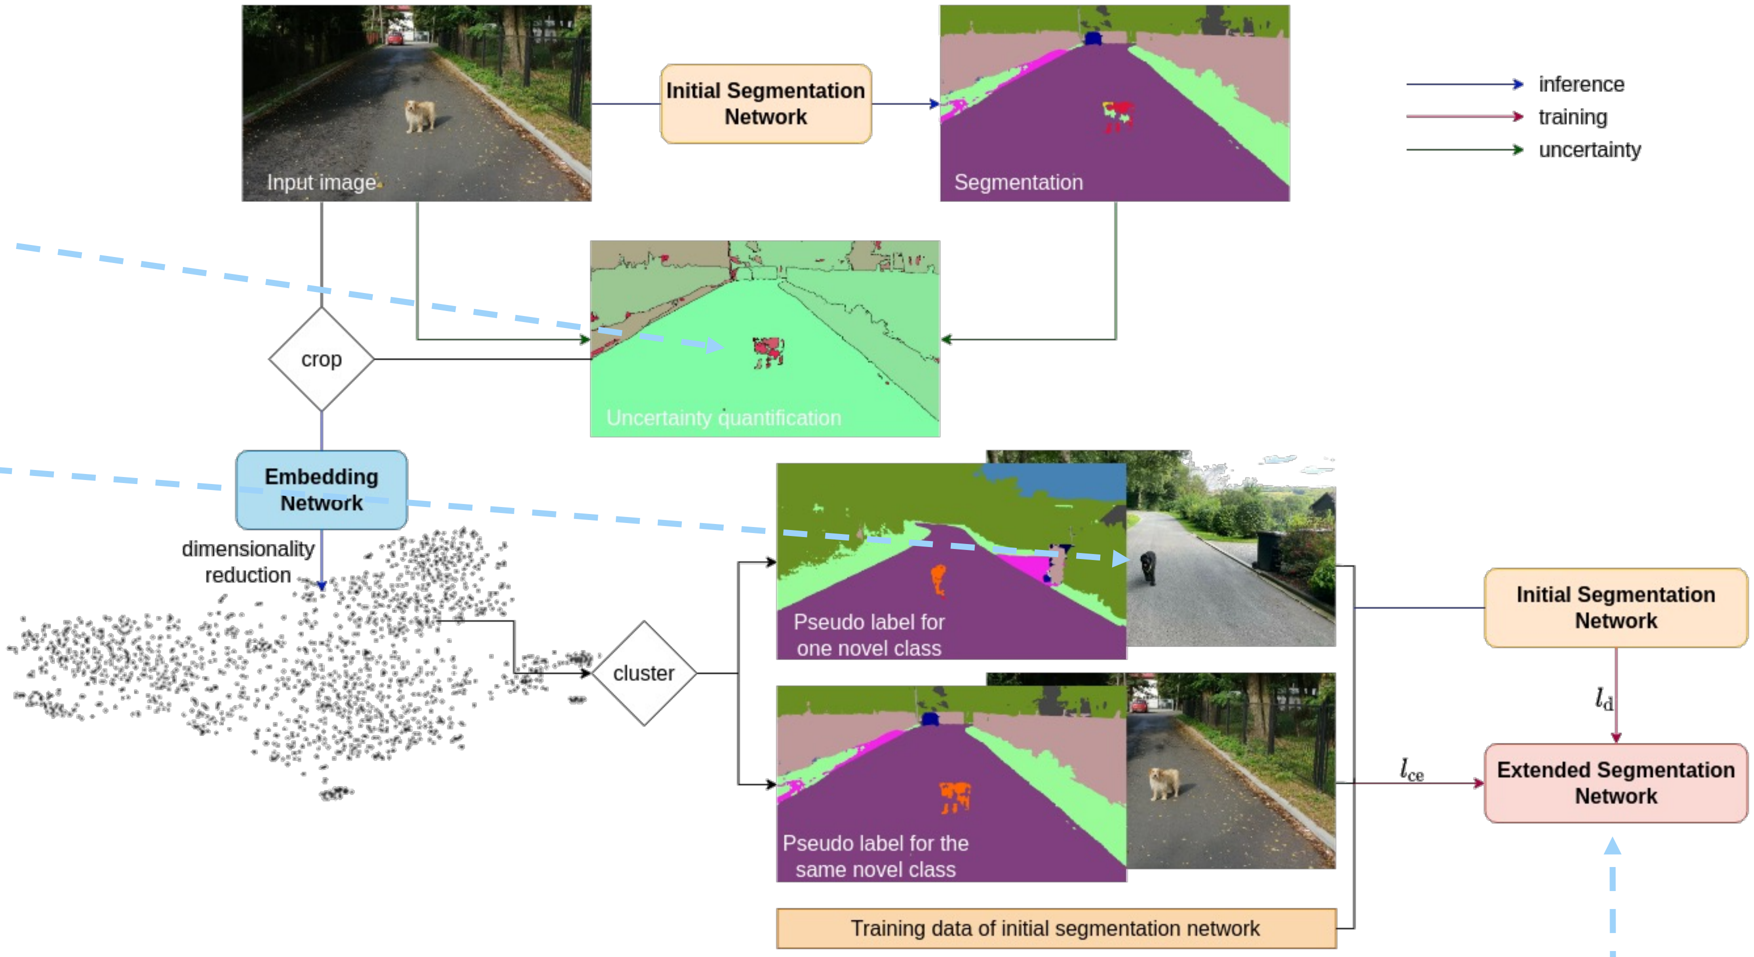

Motivation: Toy Example

- **Top left:** initial model trained on closed set of semantic classes, *e.g.* two moons
- **Top right:** OoD detection, *e.g.* by thresholding on maximized entropy
- **Bottom left:** clustering of OoD samples to create pseudo-labels, *e.g.* with k-means
- **Bottom right:** extending the model by fine-tuning it on the OoD enriched training data

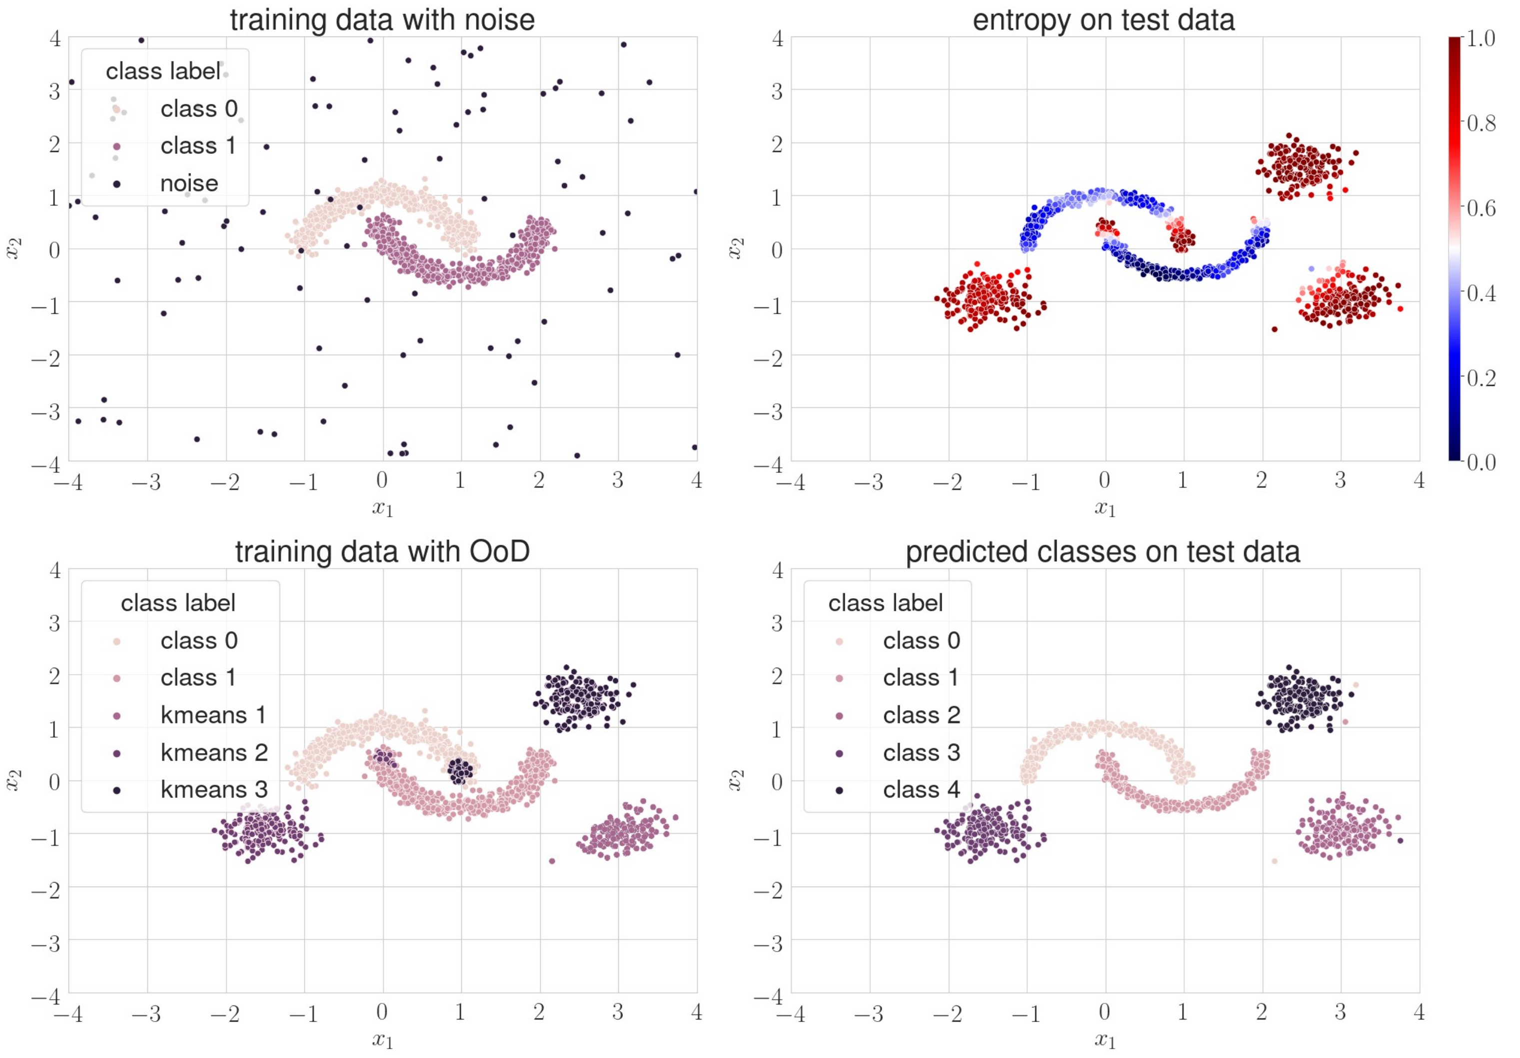

Introduction into Open World Recognition

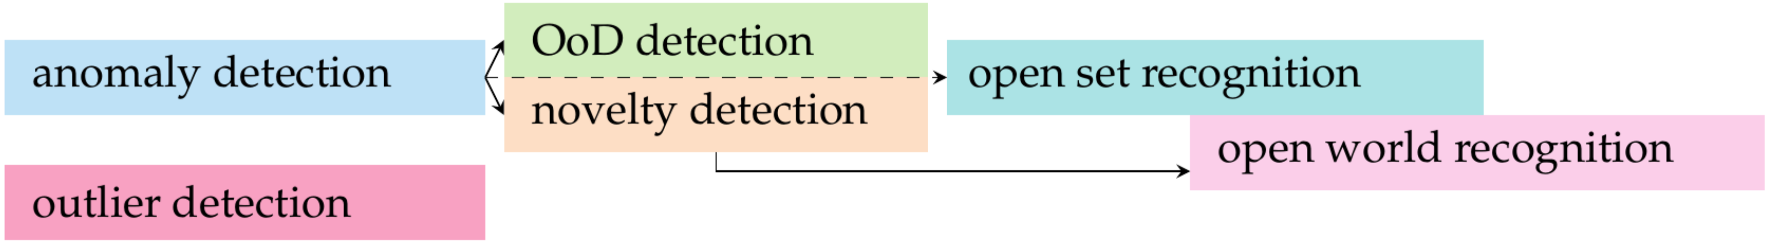

1. **Anomaly detection:** Out-of-distribution (OoD) detection (binary) vs. novelty detection (multi-class)

2. **Open set recognition:** semantic segmentation + OoD detection

3. **Open world recognition:** semantic segmentation + novelty detection + class-incremental learning

What are Out-of-Distribution (OoD) Objects?

Objects from classes not included in the learnable semantic space of a deep neural network (DNN)

→ "none-of-the-known" objects

- **Example:** dogs are OoD objects for models trained on the Cityscapes dataset since there are no animals in the Cityscapes data

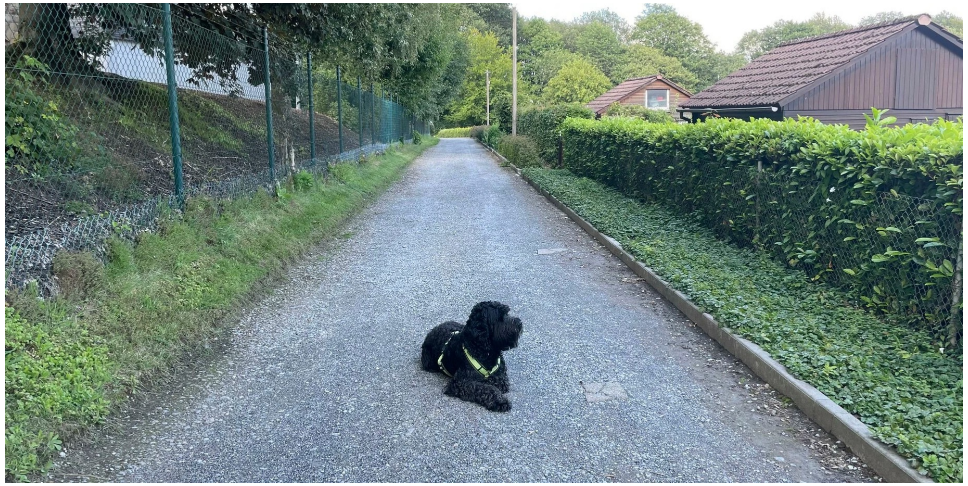

Open World Recognition

I. Data

in the open world, the DNN is confronted with OoD data.

II. Novelty Detection

novel classes must be recognized and labeled.

III. Incremental Learning

the DNN must be extended and fine-tuned on the annotated novel data.

Problems:

- lots of unlabeled data where novel classes might appear
- pixel-wise annotations are expensive

Our Contribution

- **Discovery:** detection and clustering of OoD instances to discover novel classes

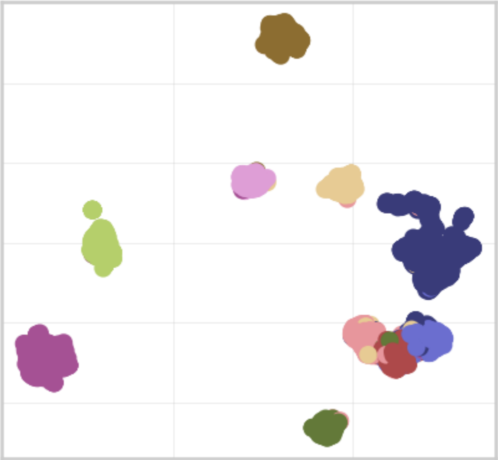

Idea: discover novel classes by clustering OoD instances in a low dimensional embedding space

- **Pseudo-Labeling:** creation of pseudo-labels based on the output of the OoD detector and on clusters

Example: pseudo labels for the classes human (left) and car (right).

Meta Regression as OoD Detector

- Construct metrics based on dispersion measures derived from the softmax probabilities of the underlying semantic segmentation DNN
- Apply a gradient boosting regressor to quantify the prediction quality
- Estimate the segment-wise IoU (from 0 (red) to 1 (green)) without having a ground truth
- Detect OoD segments by thresholding on the estimated IoU

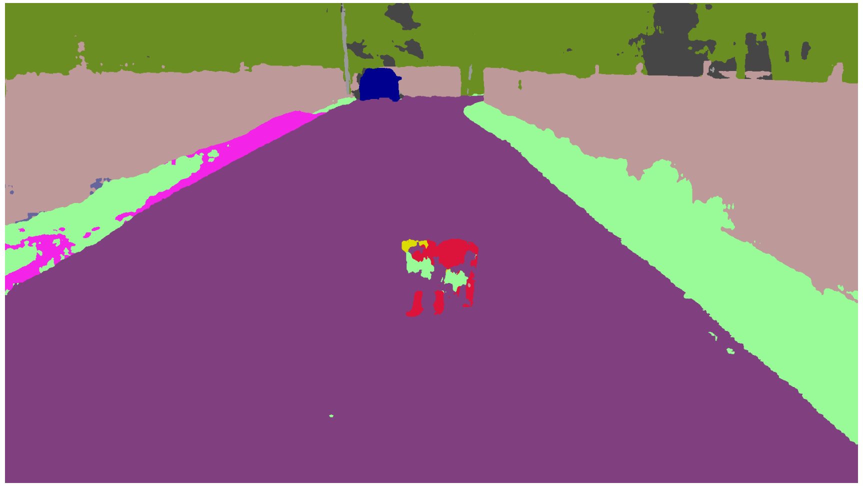

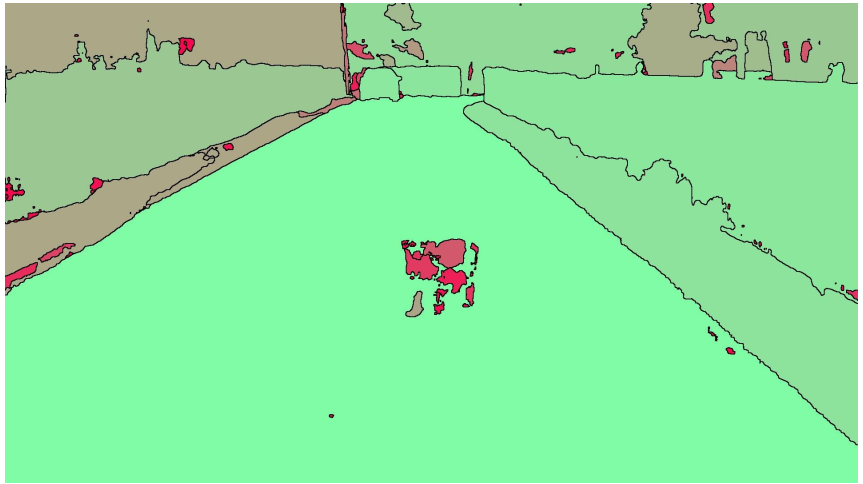

Embedding Network

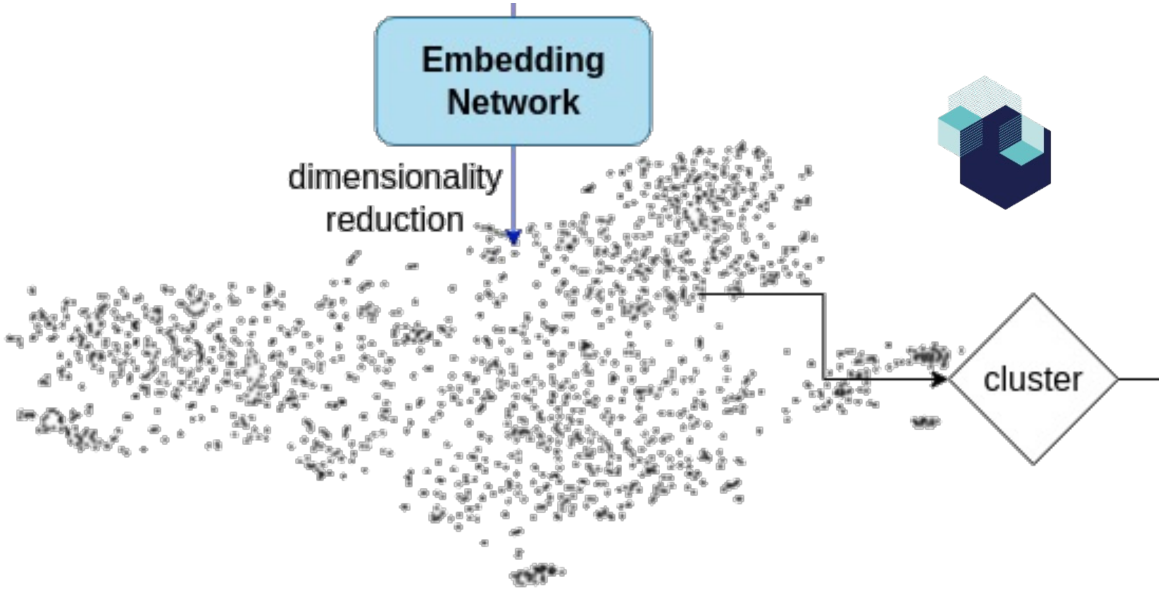

- Employ an image classification CNN which is trained on ImageNet 1000 classes as embedding network
- Feed images patches of OoD objects into the embedding network
- Extract the features of the penultimate layer
- Reduce the dimension of these features, e.g. with PCA, t-SNE or UMAP
- Build clusters in the low-dimensional space, e.g. with k-means or DBSCAN
- Each cluster constitutes a novel class

Pseudo-Labeling

- Pseudo labels are generated for all images/OoD instances per cluster
- Combination of the predictions of the semantic segmentation DNN and the OoD detector
- Class ID of OoD objects depends on their corresponding cluster (same ID for the whole cluster)

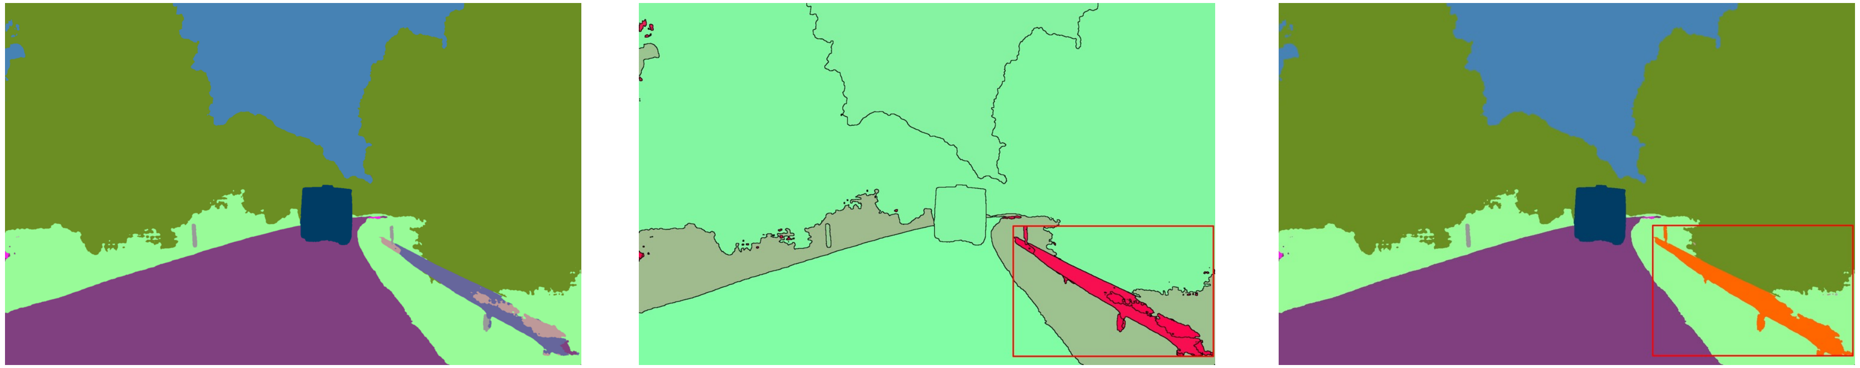

Class-Incremental Learning

- Class-incremental learning causes catastrophic forgetting
- 1. strategy: replay of training data
- 2. strategy: knowledge distillation
- The DNN is extended by output neurons for novel classes
- We only train parameters of the decoder
- The output of the initial DNN is computed for knowledge distillation

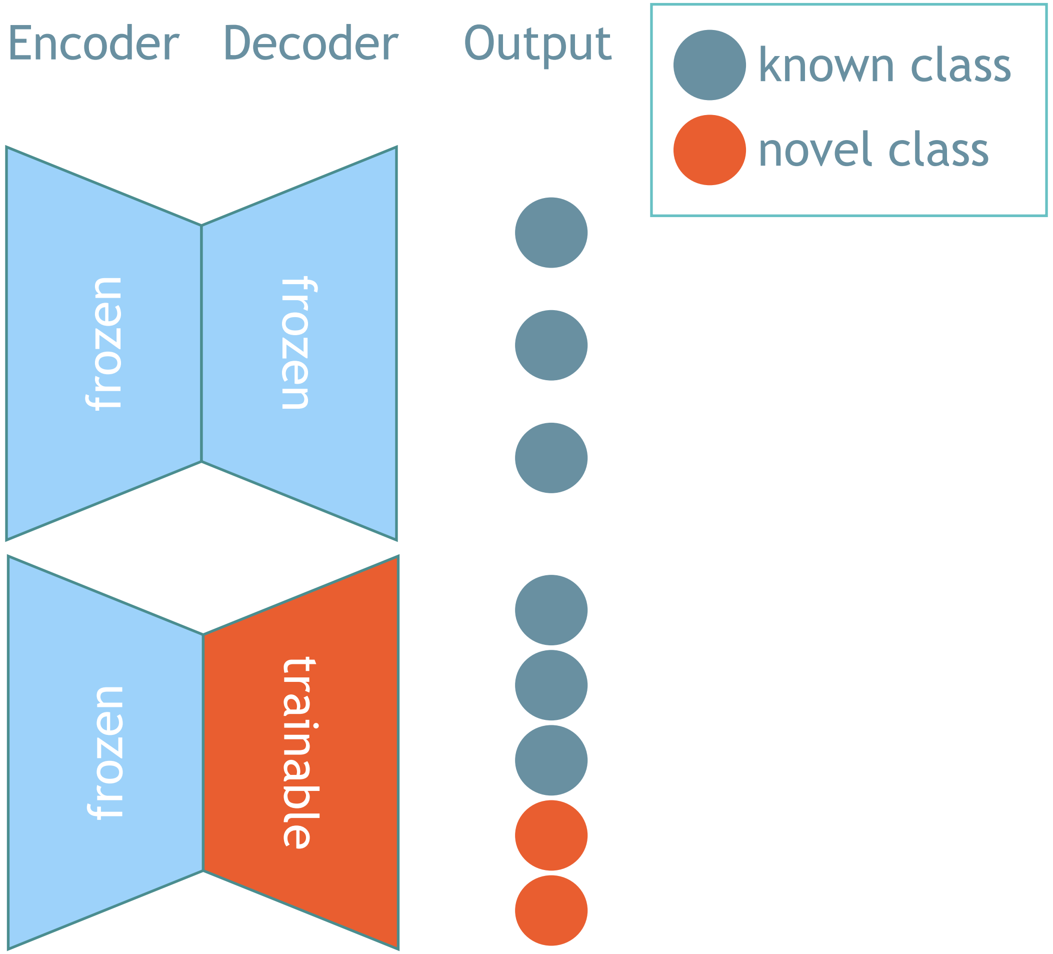

Results

image & novelty annotation ground truth

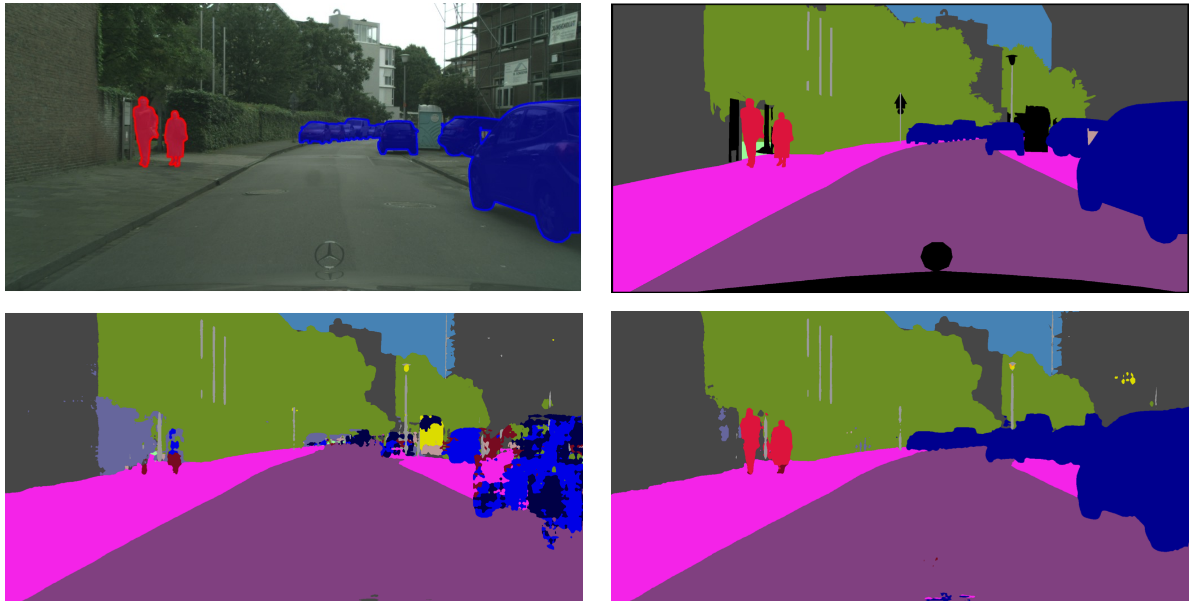

prediction of initial DNN prediction of extended DNN

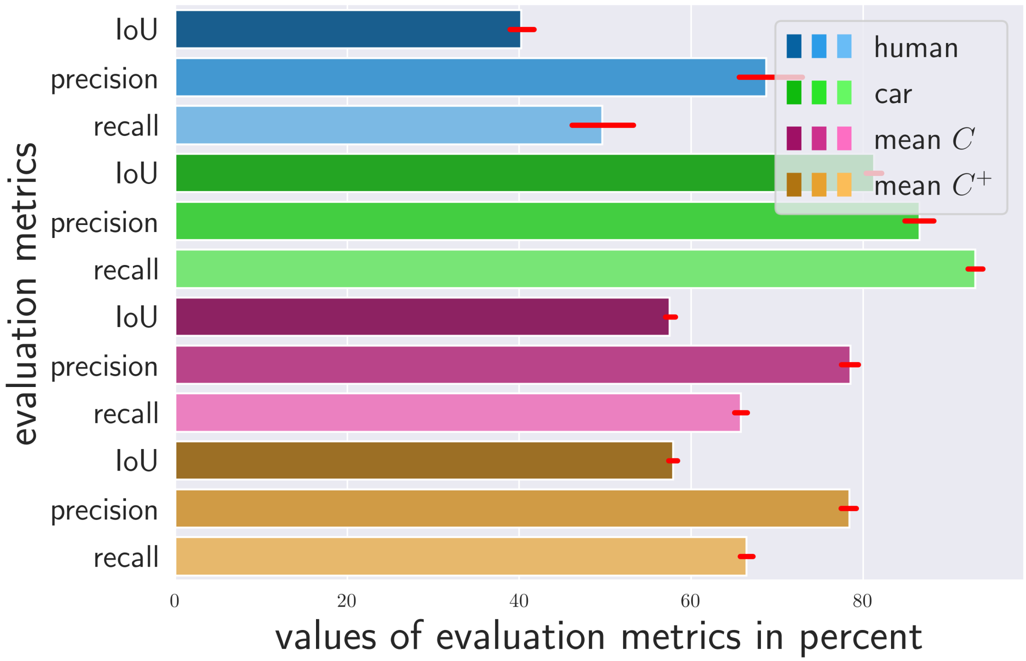

- Novel classes: human (40.22 ± 1.77 % IoU) and car (81.27 ± 1.16 % IoU)
- Performance on old classes: initial DNN achieves 56.99 %, extended DNN 57.52 ± 0.80 % mean IoU

Conclusion

- we can incrementally learn novel classes without any ground truth of the novel class
- our method benefits significantly from high performance networks and anomaly detectors
- shortcomings: clustering algorithm is highly sensitive towards chosen hyperparameters & it is not guaranteed that there are no clusters of known classes

- our method can be used as a baseline for future approaches
- paper and code are publicly available

Thank you!

paper

code

KIDELTA LEARNING

Scalable AI for Automated Driving

Svenja Uhlemeyer | University of Wuppertal | uhlemeyer@math.uni-wuppertal.de

Matthias Rottmann | University of Wuppertal | rottmann@uni-wuppertal.de

Hanno Gottschalk | University of Wuppertal | hgottsch@uni-wuppertal.de

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.

KI FAMILIE

Supported by:

Federal Ministry for Economic Affairs and Climate Action

on the basis of a decision by the German Bundestag

www.ki-deltalearning.de @KI_Familie KI Familie