KI ABSICHERUNG
Safe AI for Automated Driving

Out-of-Distribution Detection in Semantic Segmentation

Robin Chan, Svenja Uhlemeyer,
Matthias Rottmann, Hanno Gottschalk

DNN specific Safety Concern: Unreliable confidence information (SC-1.1)

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

- Objects not fitting the definition of classes used in the training data of a model:

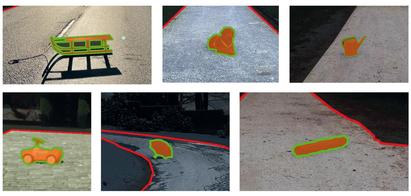

Figure 1: Different OoD objects in street scenes

Entropy Maximization

- Deep neural networks (DNNs) for semantic segmentation are usually trained to operate on a closed set of object classes
- Retrain for high softmax entropy on OoD objects while retaining original semantic segmentation performance:

$$(1-\lambda)\ \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}_{in}}\left[\mathcal{L}_{in}(f(\mathbf{x}),y)\right]\ +\ \lambda\ \mathbb{E}_{\mathbf{x}'\sim\mathcal{D}_{out}}\left[\mathcal{L}_{out}(f(\mathbf{x}'))\right]$$

- Cityscapes as in-distribution dataset, COCO as out-distribution dataset:

Figure 2: Random sample from the COCO dataset

- Superior pixel level OoD detection performance over several established baselines:

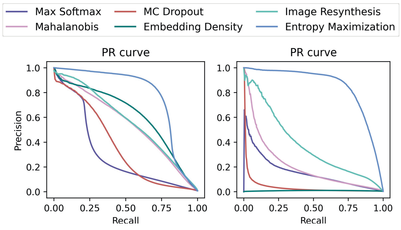

Figure 3: Lost&Found (left), RoadObstacle (right)

OoD Prediction after Entropy Maximization

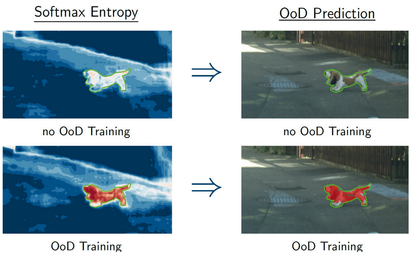

Figure 4: OoD object prediction via thresholding

Safety Hypothesis:

The method addresses the safety concern Unreliable Confidence Information of DNNs. It enhances the performance of detecting semantically unknown objects in the semantic segmentation of street scenes by means of the prediction uncertainty of DNNs.

Meta Classification

- Construct metrics from the softmax probabilities of the semantic segmentation DNN
- Apply logistic regression to remove false positive OoD object predictions

Figure 5: Meta Classification for OoD detection

- Improved error rates wrt. OoD predictions
- Marginal sacrifice of Cityscapes semantic segmentation performance of at most 1%

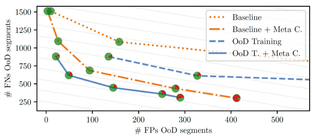

Figure 6: Error rates wrt. OoD object predictions

Generalization Capabilities

- Also performing well on truly unseen OoD objects in the newly introduced SegmentMeIfYouCan benchmark

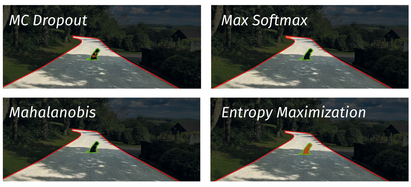

Figure 7: Final OoD object segmentation masks

References:

[1] R. Chan et al. „Entropy Maximization and Meta Classification for Out-of-Distribution Detection in Semantic Segmentation“, ICCV 2021.
[2] R. Chan et al. „SegmentMeIfYouCan: A Benchmark for Anomaly Segmentation“, NeuRIPS 2021 Track on Datasets and Benchmarks.

BERGISCHE UNIVERSITÄT WUPPERTAL

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

Image sources:
Fig 1 und 7: Segment Me If You Can
Fig 2: Microsoft COCO / Flickr
Fig 4: BUW (Heatmaps, links), Daimler AG (Raw Bilder, rechts)
Andere: BUW

KI Absicherung 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

https://www.ki-absicherung-projekt.de @KI_Familie KI Familie