KI DELTA LEARNING
Scalable AI for Automated Driving

A Low-Complexity Approach for Domain Adaptation

Joshua Niemeijer, Jörg P. Schäfer

Use-Case

Our research addresses unsupervised learning techniques for domain adaption to mitigate or even eliminate manual labeling efforts. Our primary focus lies on the task of semantic segmentation. We are interested in the domain changes from synthetic to real data and in real to real domain changes.

Figure 1: Left: Source domain performance; Right target domain performance

Technical Problem

The training dataset for DNNs defines a specific data domain. The trained DNNs perform well in this source domain. However, when they are applied to a different domain, the target domain, the performance of the DNN drops (see Figure 1). Unfortunately, supervised adaptation to the target domain is laborious since labeled data is expensive. The scientific problem resides in the different distribution of source and target domains. How to adapt a model trained on the source domain without annotations on the target domain data?

Technical Solution

We present a low-complexity approach to address the domain discrepancy. We aim to align both distributions through semantic Self-Supervision. To that, we compute the class prototypes that represent the classes in the feature space. We assume that target domain feature representations are closer to the correct class centroids than to the incorrect ones. We exploit this property to improve this clustering and increase the target domain's segmentation quality. Given the improvements on the target domain, we can generate high-quality pseudo labels for self-training on the target domain. Self-training on the target domain results in the feature space's alignment of target and source domain distributions. This aids semantic clustering.

Thus we end up with the synergistic effect seen in Figure 3 by iteratively improving the pseudo labels .

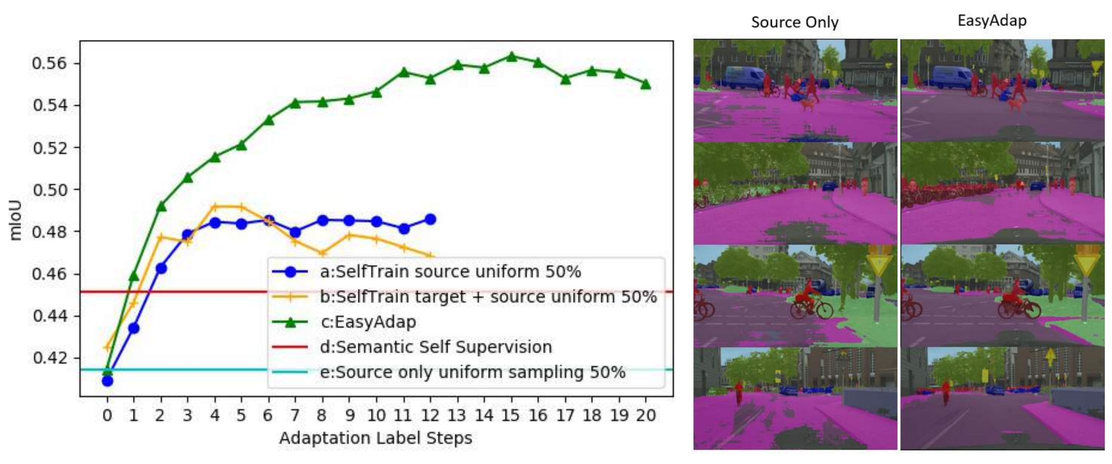

Figure 3: Left: MIoU values over iteration. Right: Qualitative examplesv

Evaluation

We evaluated our approach on different real-to-real and synthetic-to-real domain changes. Representative in Figure 3, the adaptation from the synthetic Gta5 to the real Cityscapes dataset is depicted. We improve the source-only training from 38% to 56% mIoU. Despite our low complexity approach, we achieved near SOTA performance, which makes our approach especially applicable to real-world scenarios. Since it is infeasible to adapt to each domain, we also explored how different approaches generalize to unseen real-world domains. Figure 4 shows that the adapted models struggle in some domains, indicating the necessity for further generalization research.

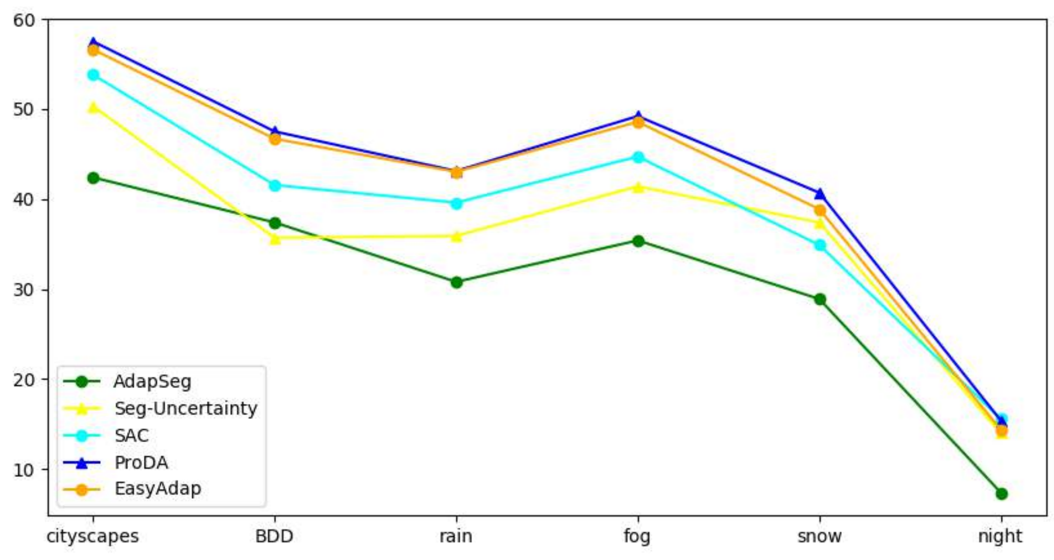

Figure 4: Different approaches for adaptation from GTA5 to Cityscapes yield models that are evaluated on third real world domains.

References:

1. „The cityscapes dataset for semantic urban scene understanding“ Cordts 2016
2. “Playing for data: Ground truth from computer games” Richter 2016

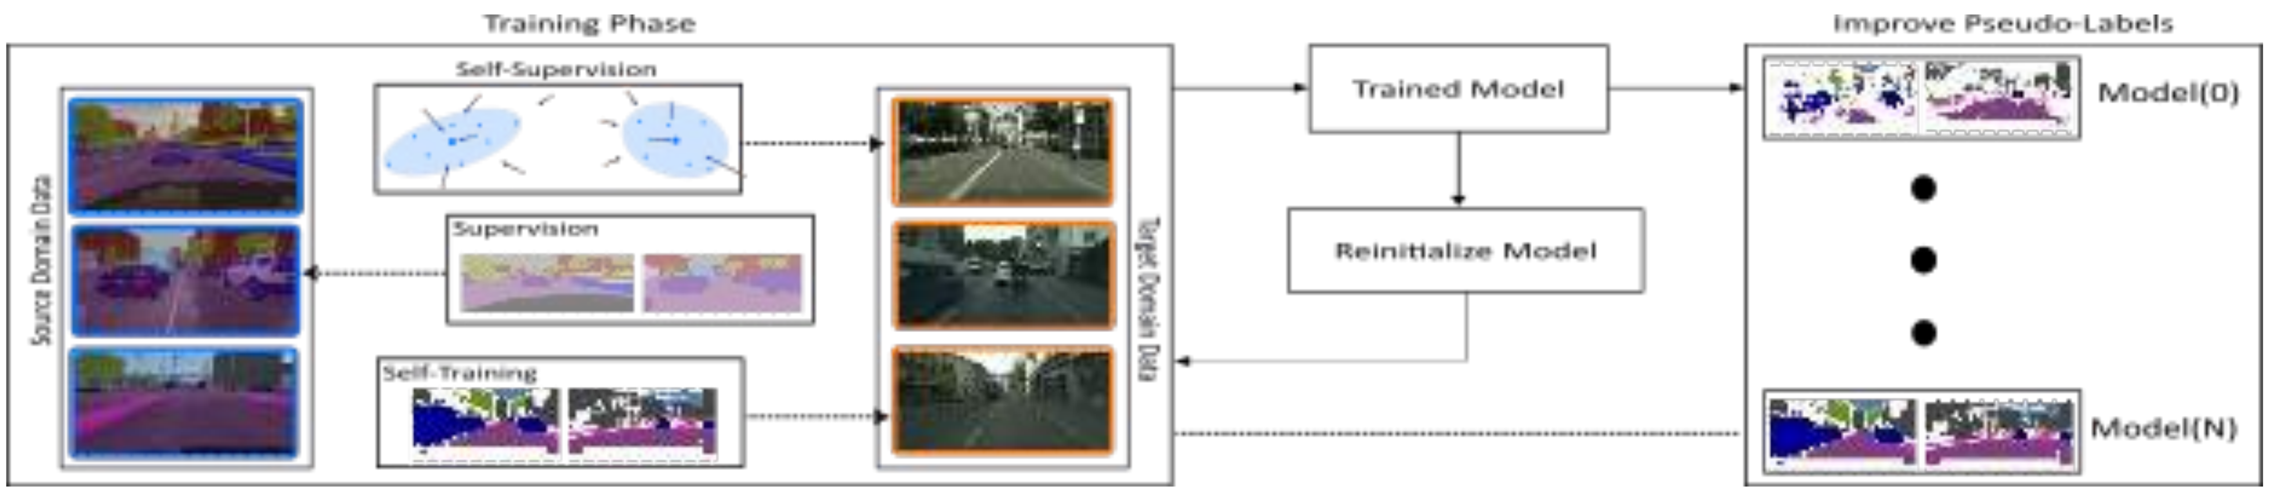

Figure 2: An overview of the internal dependencies (dashed lines): The source-domain training depends on the source-domain data and labels; the self-training of the target-domain training depends on the target-domain data and its previously created pseudo-labels; the self-supervision depends on the target-domain data and the feature clusters of the source-domain data.

Environment

Partners

External partners

For more information contact:

Joshua.Niemeijer@dlr.de
Joerg.Schaefer@dlr.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.

www.ki-deltalearning.de @KI_Familie KI Familie

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