NEURAL INFORMATION PROCESSING SYSTEMS

CARLANE: A Lane Detection Benchmark for Unsupervised Domain Adaptation from Simulation to multiple Real-World Domains

Julian Gebele[1], Bonifaz Stuhr[1,2], Johann Haselberger[1,3]

[1] University of Applied Science Kempten
[2] Autonomous University of Barcelona
[3] Technische Universität Berlin

Introduction

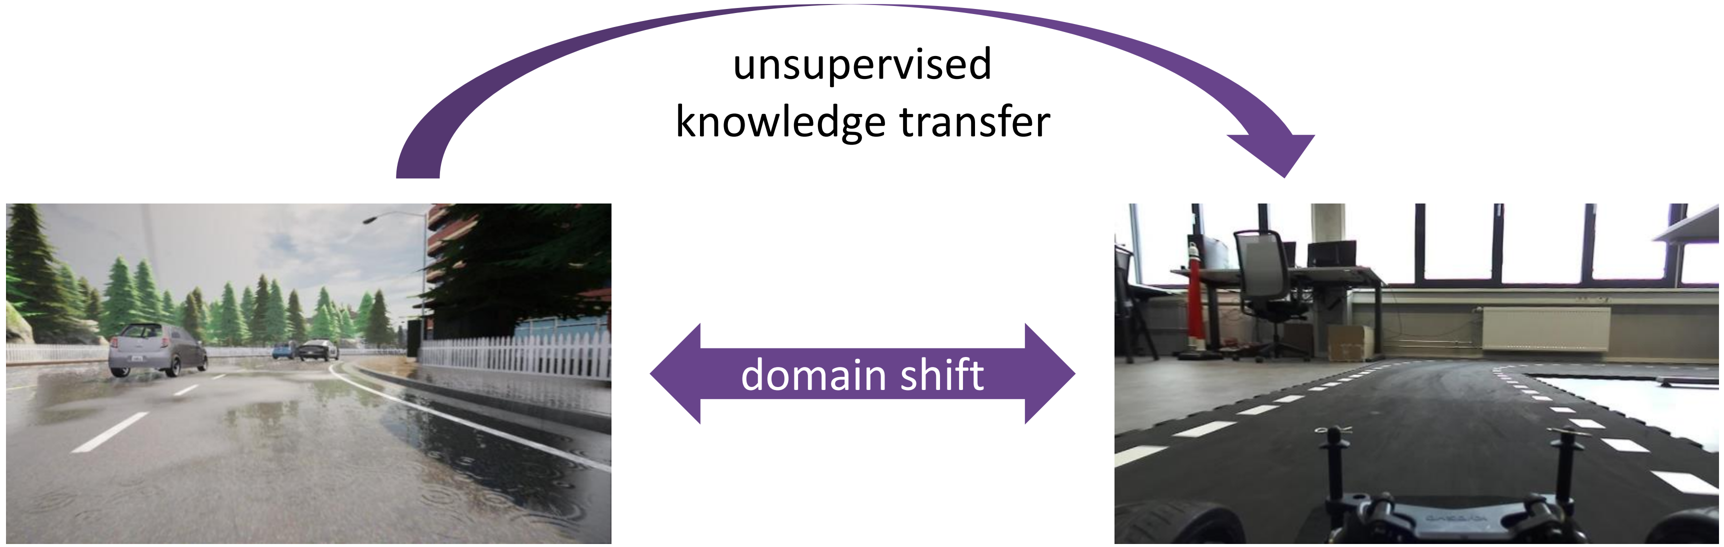

labeled source domain **unlabeled target domain**

- Unsupervised Domain Adaptation (UDA) demonstrates great potential to mitigate domain shifts by transferring models from labeled source domains to unlabeled target domains
- Only a few UDA works focus on lane detection for autonomous driving
- This can be attributed to the lack of publicly available datasets

Contributions

3-Way Sim-to-Real Benchmark

allowing single- and multi-target UDA

Establishing Baselines

evaluating well-known and own UDA methods

Several Dataset Tools

data collection agent labeling tool

» To the best of our knowledge, we are the first to adapt a lane detection model from simulation to multiple real-world domains

The CARLANE Benchmark

- The CARLANE Benchmark consists of three distinct sim-to-real datasets, which we build from our three different domains:

MoLane focuses on abstract lane markings in the domain of a 1/8th *Mo*del vehicle.

TuLane incorporates balanced and domain-randomized images from simulation as the source domain and the well-known *Tu*Simple dataset.

MuLane is a balanced combination of *M*oLane and T*u*Lane with two target domains.

Images sampled from our CARLANE Benchmark.

Dataset Collection - Simulation

- We use the open-source CARLA simulator
- Randomize multiple aspects of the agent and environment:
  - weather and daytime by adapting parameters such as cloud density, rain intensity, …
  - up to five neighbor vehicles are spawned randomly in the vicinity of the agent
  - ego vehicle position is varied by the data agent using a triangle wave function

Dataset	Ego Vehicle	Camera Position	Lane Deviation	Traffic	Pedestrians	World Objects	Daytime	Weather	City	Rural	Highway	Terrain	Lane Topology	Road Appearance
[4]	✗	✓	✓	✓	✗	✓	✓	✗	✗	✓	✗	✓	✓	✓
[19]	✗	✓	✓	✓	✗	✓	✓	✗	✗	✓	✗	✓	✓	✓
[5]	✗	✗	✓	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓
ours	✓	✓	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓

At this point, none of the other datasets is publicly available

[4] N. Garnett, R. Cohen, T. Pe'er, R. Lahav, and D. Levi, "3D-LaneNet: End-to-End 3D Multiple Lane Detection," in ICCV, pp. 1013 – 1021, 2019.
[5] C. Hu, S. Hudson, M. Ethier, M. Al-Sharman, D. Rayside, and W. Melek, "Sim-to-Real Domain Adaptation for Lane Detection and Classification in Autonomous Driving," 2022.
[19] N. Garnett, R. Uziel, N. Efrat, and D. Levi, "Synthetic-to-Real Domain Adaptation for Lane Detection," in ACCV, 2020.

Dataset Collection - Model Vehicle

- We capture data with our 1/8th model vehicle
- Tracks roughly contain the same proportion of straight and curved segments
- Randomizations:
  - alternating backgrounds, lighting conditions, surface materials, and lane topology
  - four different locations

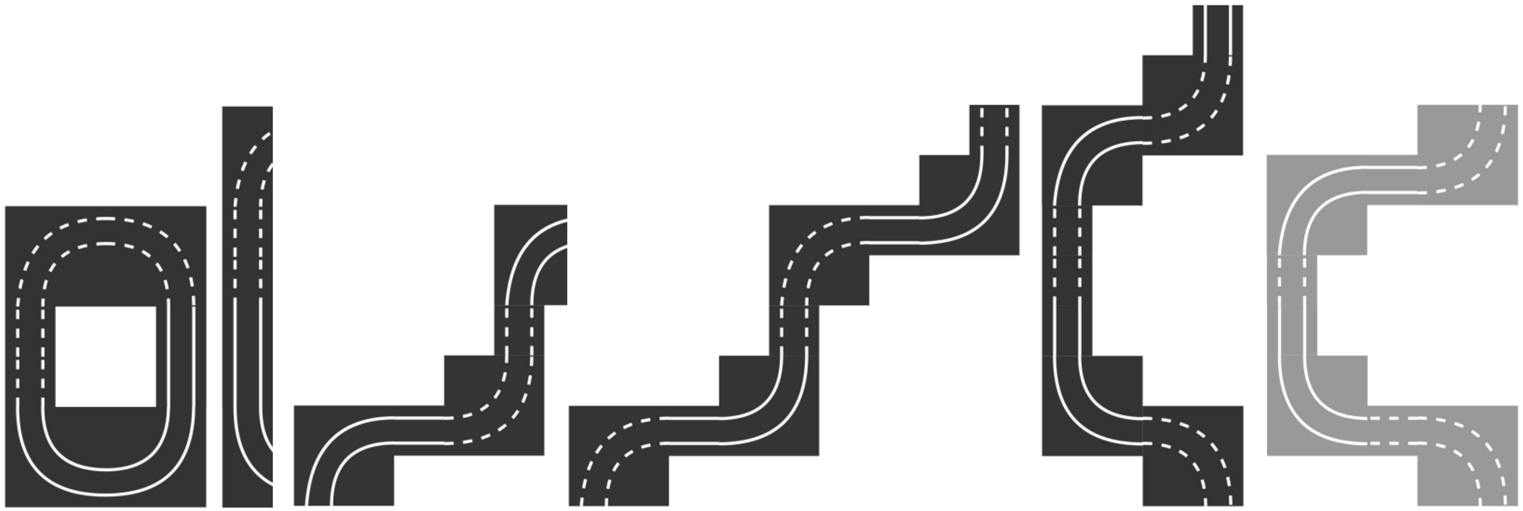

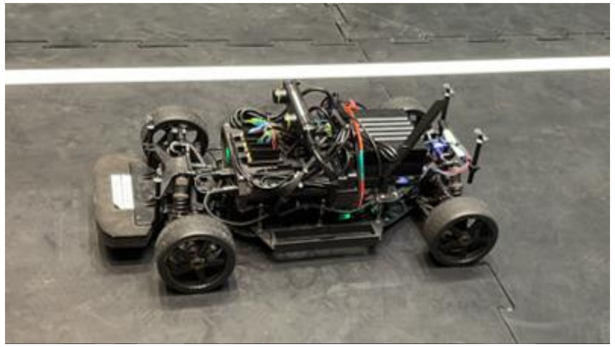

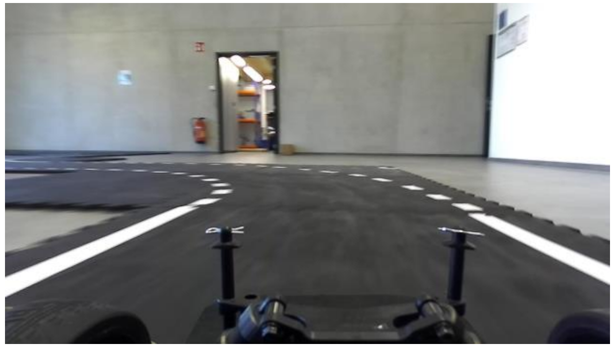

Dataset Collection - TuSimple

- We create a cleaned version of TuSimple[1] for the real-world target domain of TuLane
  - ensure that up to four lanes closest to the car are correctly labeled
- To clean the data, we utilize our publicly available labeling tool

[1] TuSimple, “TuSimple-benchmark.” https://github.com/TuSimple/tusimple-benchmark/tree/master/doc/lane_detection.

Dataset Statistics

- CARLANE contains a total of 163K unique images
- 118K images are annotated

Dataset overview. Unlabeled images denoted by *, partially labeled images denoted by **

Dataset	domain	total images	train	validation	test	lanes
MoLane	CARLA simulation	84,000	80,000	4,000	-	≤ 2
	model vehicle	46,843	43,843*	2,000	1,000	≤ 2
TuLane	CARLA simulation	26,400	24,000	2,400	-	≤ 4
	TuSimple	6,408	3,268	358	2,782	≤ 4
MuLane	CARLA simulation	52,800	48,000	4,800	-	≤ 4
	model vehicle + TuSimple	12,536	6,536**	4,000	2,000	≤ 4

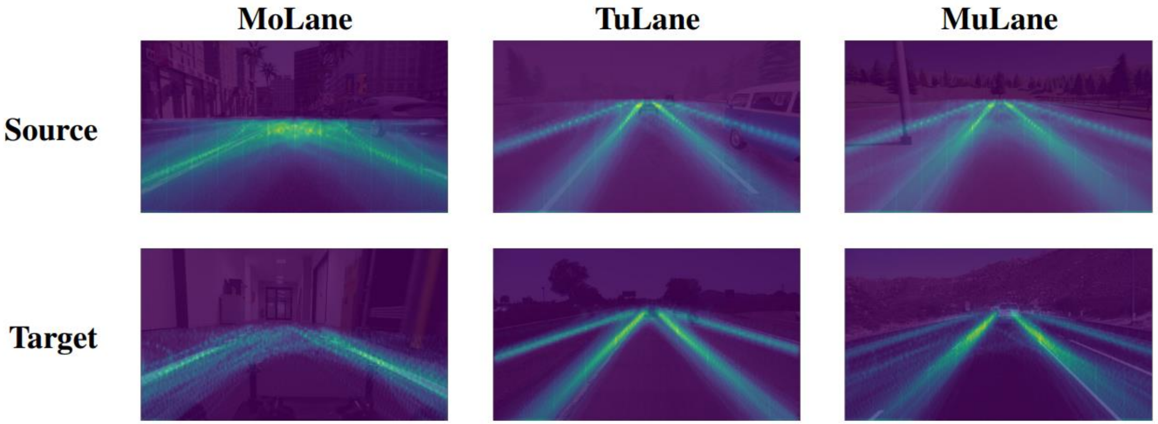

Lane annotation distributions of the three subsets of CARLANE.

Benchmark Experiments

- We use **UFLD**[2] as lane detection backbone
- The following UDA methods are evaluated: **DANN**[3], **ADDA**[6], **SGADA**[7], and **SGPCS** (ours)
- To work with UFLD the UDA methods had to be adopted for grid-based lane detection
- Results are given as an average over five runs
- Baselines, implementations, and weights are publicly available

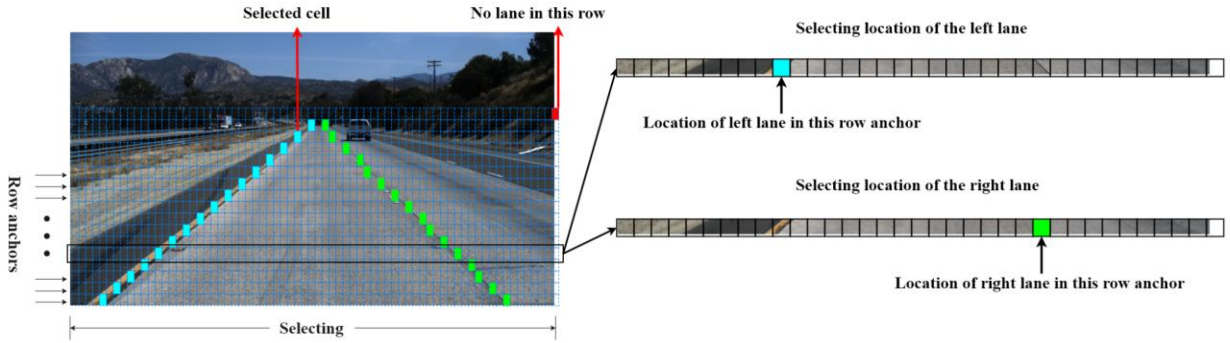

UFLD's grid-based lane detection.

[2] Qin, Zequn, Huanyu Wang, and Xi Li. "Ultra fast structure-aware deep lane detection." European Conference on Computer Vision. Springer, Cham, 2020.
[3] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. March, and V. Lempitsky, "Domain-Adversarial Training of Neural Networks," Journal ofMachine Learning Research, vol. 17, no. 59, pp. 1–35, 2016.
[6] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial Discriminative Domain Adaptation," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), July 2017.
[7] I. B. Akkaya, F. Altinel, and U. Halici, "Self-Training Guided Adversarial Domain Adaptation for Thermal Imagery," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, pp. 4322–4331, June 2021.

SGPCS (ours)

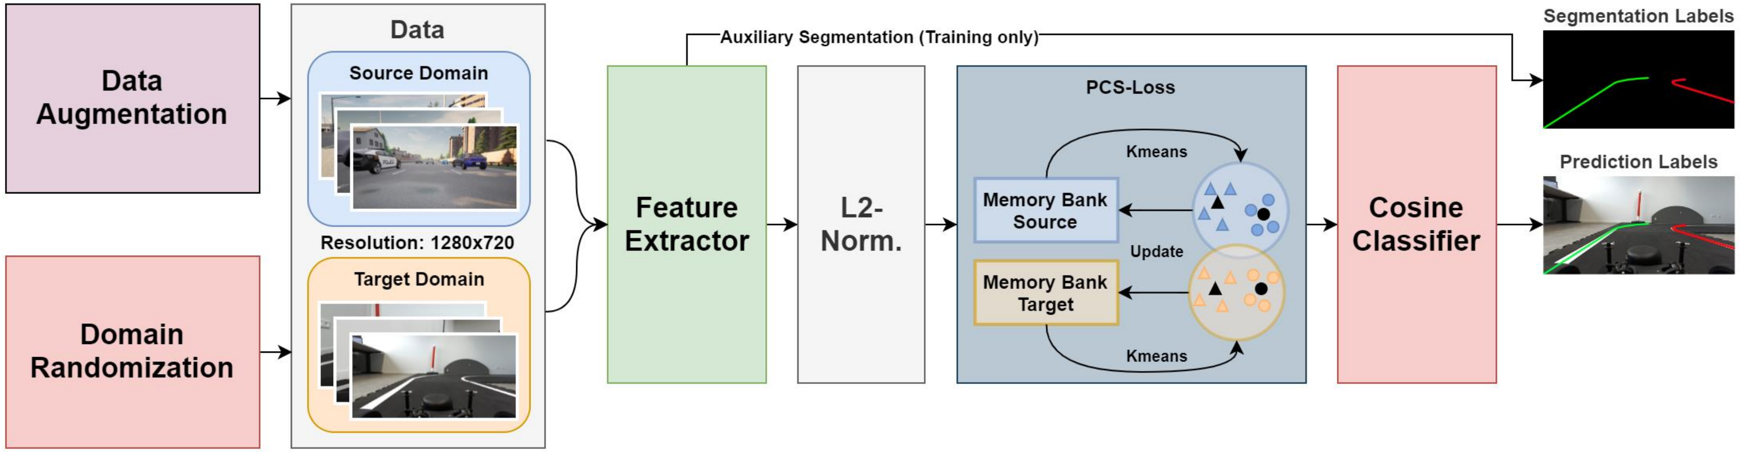

- We build upon PCS [8] and perform in-domain contrastive learning and cross-domain self-supervised learning via cluster prototypes
  - Memory bank features are updated with a momentum of 0.5
  - Spherical K-means [9] (K=2500) is used to cluster memory bank features into prototypes
- Our objective function comprises the in-domain and cross-domain losses from PCS, all losses from UFLD, and our pseudo loss for grid-based lane detection

[8] X. Yue, Z. Zheng, S. Zhang, Y. Gao, T. Darrell, K. Keutzer, and A. Sangiovanni-Vincentelli, “Prototypical Cross-domain Self-supervised Learning for Few-shot Unsupervised Domain Adaptation,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2021.
[9] J. Johnson, M. Douze, and H. Jégou, “Billion-scale similarity search with GPUs,” IEEE Transactions on Big Data, vol. 7, no. 3, pp. 535–547, 2019.

Benchmark Experiments

- Adaptation methods are not able to achieve comparable results to the supervised baselines (UFLD-TO):
  - Maximum accuracy gain of 4.55%
  - High false positive (FP) and false negative (FN) rates
- FP and FN rates increase significantly on the multi-target task of MuLane

ResNet-18	MoLane LA	MoLane FP & FN	TuLane LA	TuLane FP	TuLane FN	MuLane LA	MuLane FP	MuLane FN
UFLD-SO	89.39	25.25	87.43	34.21	23.48	88.02	50.24	26.08
DANN [12]	87.65±0.48	29.97±1.21	88.74±0.32	32.71±0.52	21.64±0.65	86.01±0.67	55.33±1.22	36.30±1.90
ADDA [13]	92.85±0.17	10.61±0.77	90.72±0.15	29.73±0.36	17.67±0.42	89.83±0.33	46.79±0.43	20.57±0.63
SGADA [21]	93.82±0.10	**7.13**±0.22	**91.70**±0.13	**28.42**±0.34	**16.10**±0.43	90.71±0.10	**45.13**±0.32	**17.26**±0.36
SGPCS (ours)	**93.94**±0.04	7.16±0.16	91.55±0.13	28.52±0.21	16.16±0.26	**91.57**±0.22	45.49±0.63	17.39±0.88
UFLD-TO	97.35	0.50	94.97	18.05	3.84	96.57	34.06	2.49

ResNet-34	LA	FP & FN	LA	FP	FN	LA	FP	FN
UFLD-SO	90.35	22.25	89.42	32.35	21.19	89.17	48.86	23.67
DANN [12]	90.91±0.42	19.73±1.51	91.06±0.14	30.17±0.20	18.54±0.25	88.76±0.22	48.93±0.47	24.16±0.89
ADDA [13]	92.39±0.26	12.17±0.84	91.39±0.16	28.76±0.30	16.63±0.36	90.22±0.39	45.84±0.54	19.49±0.90
SGADA [21]	93.31±0.10	9.41±0.16	92.04±0.09	28.18±0.20	15.99±0.24	**91.63**±0.03	**44.18**±0.12	**16.23**±0.16
SGPCS (ours)	**93.53**±0.25	**8.24**±0.91	**93.29**±0.18	**25.68**±0.48	**12.73**±0.59	91.55±0.17	44.75±0.28	16.41±0.44
UFLD-TO	97.21	0.30	94.43	20.74	7.20	96.54	33.76	2.03

» FP and FN represent wrongly detected and missing lanes, which can lead to crucial impacts on autonomous driving functions

Benchmark Experiments

- In accordance with the quantitative results, we observe only a slight adaptation of the source and target domain features for ADDA, SGADA, and SGPCS compared to the supervised baseline UFLD-SO

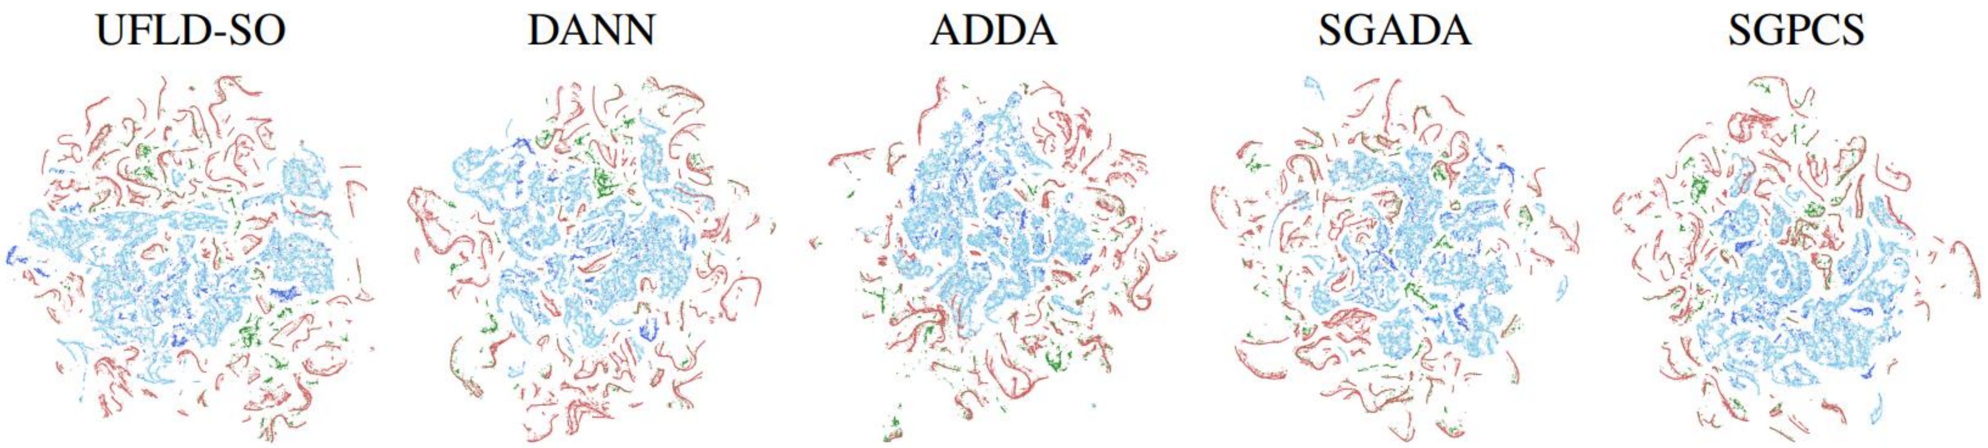

t-SNE visualization of MuLane dataset. The source domain is marked in blue, the real-world model vehicle target domain in red, and the TuSimple domain in green.

Benchmark Experiments

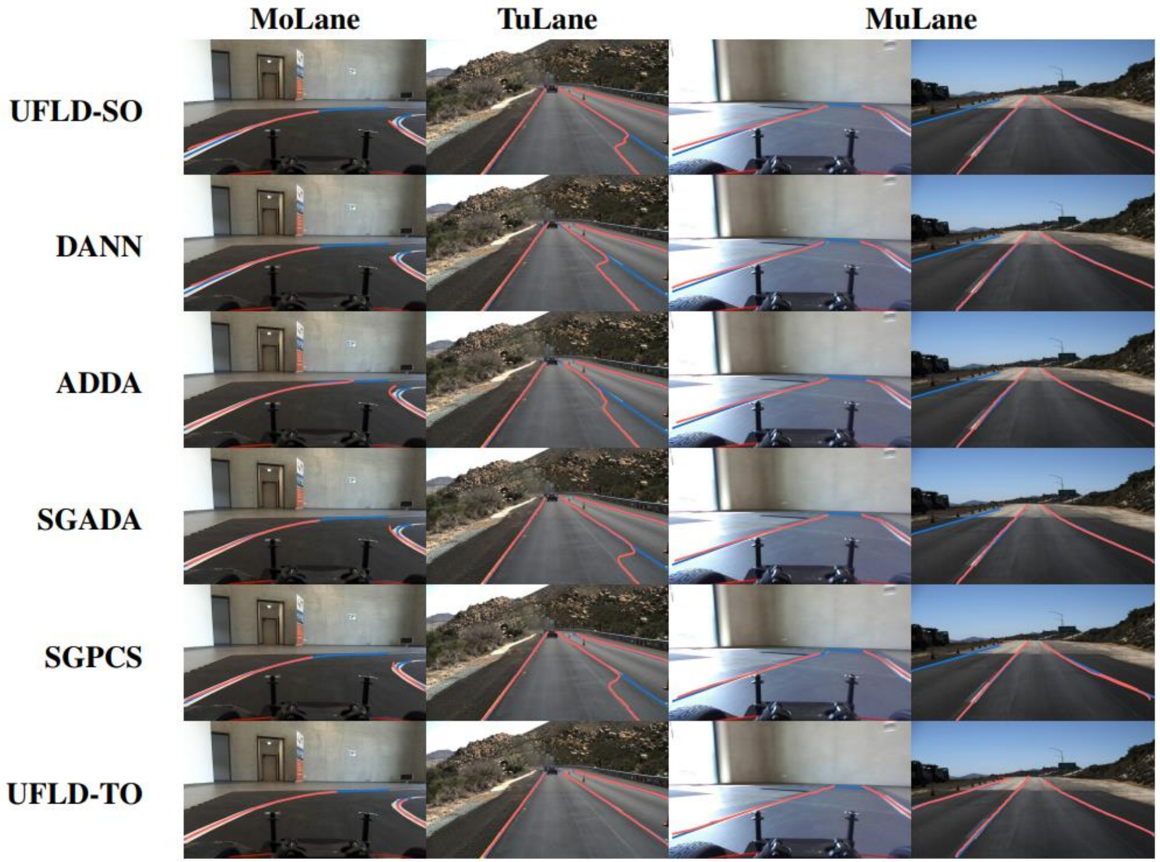

Qualitative results of target domain predictions. Ground truth lane annotations are marked in blue, predictions in red.

Benchmark Experiments

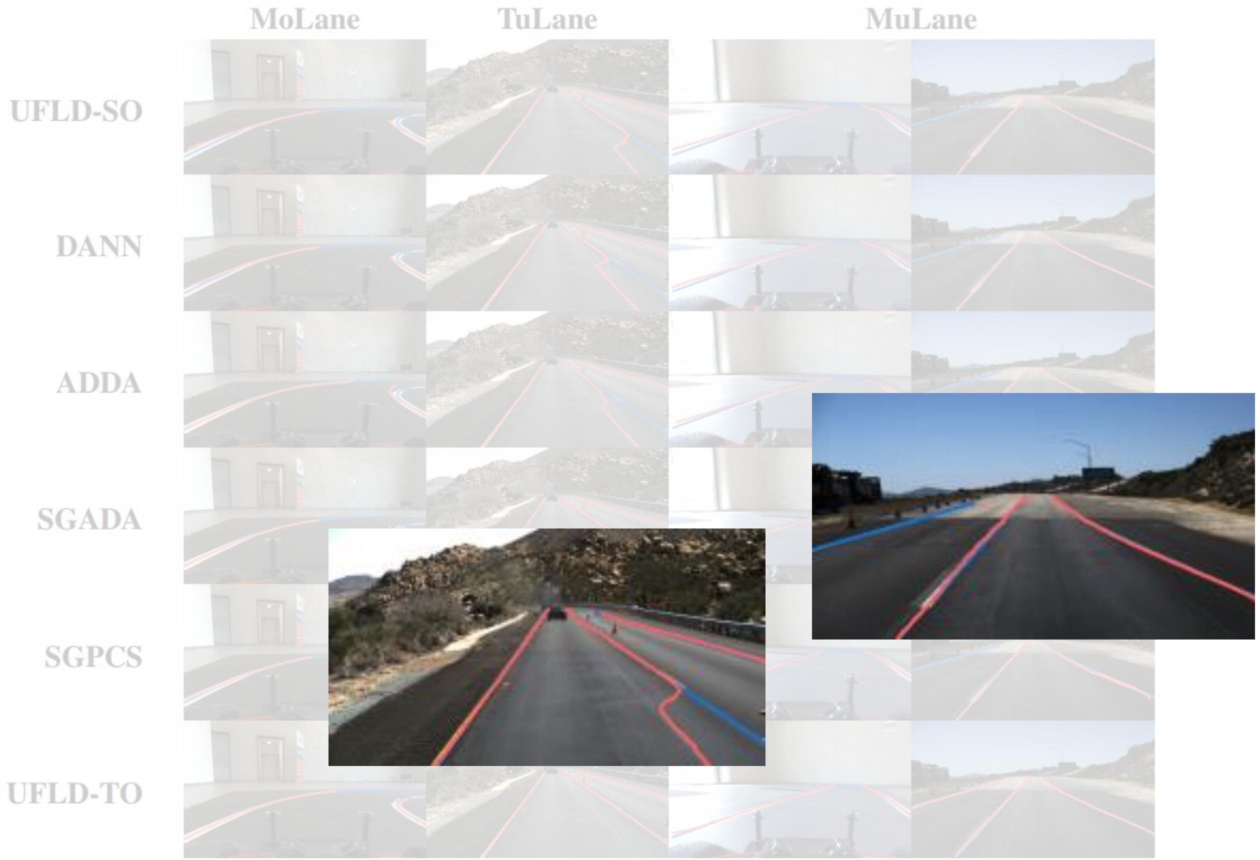

Qualitative results of target domain predictions. Ground truth lane annotations are marked in blue, predictions in red.

Conclusion

We proposed CARLANE, the first 3-way sim-to-real domain adaptation benchmark for 2D lane detection.

The current difficulties of the examined UDA methods to adequately align the source and target domains confirm the need for the proposed CARLANE benchmark.

UDA methods should be tested with care and under the right conditions on a full-scale car. However, real-world testing in the model vehicle domain can be carried out in a safe and controlled environment.

https://carlanebenchmark.github.io/

IFM
Institut für Fahrerassistenz
und vernetzte Mobilität

Hochschule
Kempten
University of Applied Sciences

Technische
Universität
Berlin

UAB
Universitat Autònoma
de Barcelona