

## Master's Thesis for Computer Science / Engineering\*



## Domain Adaptation via Bayesian Active Learning by Disagreement

Deep neural networks for autonomous driving are usually trained on data from specific geographic regions and often suffer performance drops when deployed in new environments due to domain shift, which can be interpreted as a mismatch between the learned posterior over model parameters and the target-domain data distribution. One example is transferring traffic-light classifiers trained on Japanese data to European or German road scenes, where signal layouts, colors, and pictograms differ significantly.

In this thesis, a pre-trained traffic-light classification model from the open-source Autoware stack will be adapted to European traffic-light data using Bayesian Active Learning by Disagreement (BALD). By exploiting epistemic uncertainty, BALD selects the most informative target-domain samples, those that reveal which source-domain assumptions are wrong in the new environment, for annotation and fine-tuning.

The goal is to demonstrate robust and label-efficient domain adaptation for safety-critical traffic-light perception in autonomous driving.

The proposed thesis entails the following tasks:

- + Literature research on Bayesian deep learning, active learning, and domain adaptation
- + Integration of Bayesian inference into a traffic light classifier
- + Implementing BALD-based sample selection in Python
- + Fine-tuning the Autoware pre-trained model on European traffic light data
- + Visualizing uncertainty, disagreement, and performance improvements

## Relevant Literature

- Gal, Y., Islam, R., & Ghahramani, Z. (2017). Deep Bayesian Active Learning with Image Data. Proceedings of the 34th International Conference on Machine Learning.
- Kendall, A., & Gal, Y. (2017). What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?
- Fregin, A., Muller, J., Krebel, U., & Dietmayer, K. (2018). The DriveU Traffic Light Dataset: Introduction and Comparison with Existing Datasets.
- Khorasani, A., Chen, N., Oswal, Y. D., Gopalan, A. S., Kolemen, E., & Shwartz-Ziv, R. (2026). Beyond the Loss Curve: Scaling Laws, Active Learning, and the Limits of Learning from Exact Posteriors (preprint).
- Autoware Foundation. Autoware Perception Stack Documentation.
- Models: EfficientNet-b1 and MobileNet-v2 (YOLOX for detection)

This thesis is conducted as a full Master's thesis at KIT.

\* For interested students of the math faculty, there is also a possibility.

**Institute of Measurement and Control Systems (MRT)**  
Prof. Dr.-Ing. Christoph Stiller

**Advisor:**  
Dominik Strutz, M.Sc.

**Programming language(s):**  
Python (proficient)

**System, Framework(s):**  
Linux, PyTorch, Lightning

**Required skills:**

- Familiarity with academic research papers
- Prior knowledge in probability and statistics
- Experience in deep learning (and computer vision)
- Ability to work independently

**Language(s):**  
English

For more information please contact:

**Dominik Strutz**

Room: 039  
Phone: +49 721 608-46769  
Email: Dominik.Strutz@kit.edu

To directly apply, please send your CV and transcript of records (ToR; for first term graduate students also the ToR of the previous program/degree) to the e-mail address provided above.