

Praxis der Forschung

Bayesian Knowledge Distillation for Motion Prediction Models

Methods for Big Data, Scientific Computing Center

Advisor

MSc Christian Schlauch

Contact

christian.schlauch@partner.kit.edu

Examiner

Prof. Dr. Nadja Klein

Possible Start

As soon as possible

Background

Motion prediction is a key component of autonomous driving systems, enabling the anticipation of future behavior of surrounding traffic participants. Recent advances have led to highly capable multi-agent motion prediction models based on transformers and graph neural networks that achieve state-of-the-art performance on large-scale benchmarks. However, these models are often computationally expensive and difficult to deploy in latency-critical applications.

Knowledge Distillation (KD) aims to transfer knowledge from a large, high-performing teacher model to a smaller and more efficient student model. Classical distillation methods primarily focus on matching predictions and retaining the accuracy of the teacher model. However, a major open research question is whether distilled models can also retain the uncertainty characteristics and robustness properties of their teachers. In particular, it remains unclear how uncertainty-aware distillation affects performance in safety-critical **long-tail scenarios** and under **out-of-distribution (OOD)** conditions, which are often underrepresented in standard benchmark evaluations but highly relevant for real-world deployment. **Bayesian distillation** methods are particularly attractive in this context, because understanding prediction confidence is crucial for safe decision-making in autonomous driving.

The goal of this thesis is to investigate, implement, and evaluate Bayesian distillation methods for autonomous driving motion prediction. The work can be conducted using an **existing research software stack** in **Python, Pytorch** containing state-of-the-art motion prediction models, large-scale datasets, and evaluation pipelines for benchmark, long-tail, and OOD performance assessment.

Task Details

- Conduct a literature review on Bayesian distillation for autonomous driving and related domains.
- Implement and analyze state-of-the-art distillation approaches, including, for example, output-level, distribution-level, and cross-architecture distillation methods.
- Select and extend a distillation approach with a focus on uncertainty calibration and robustness in long-tail scenarios of motion prediction.

- Systematically evaluate the strength and limitations of implemented methods using available motion prediction benchmarks.
- Analyze the adaption of the distillation approach to related research domains and derive future research directions.