UQnet: quantifying uncertainty in trajectory prediction by a non-parametric and generalizable approach

Abstract:

Predicting the trajectories of road agents is fundamental for self-driving cars. Trajectory prediction contains many sources of uncertainty in data and modelling. A thorough understanding of this uncertainty is crucial in a safety-critical task like auto-piloting a vehicle. We need to distinguish between the uncertainty caused by partial observability (of all the possible factors that may affect a drivers' near-future decisions), the so-called aleatoric uncertainty, and the uncertainty of using a model (an abstraction!) in new scenarios (possibly not present in the data with which it was calibrated), the so-called epistemic uncertainty. In this paper we propose a new framework to systematically quantify both sources of uncertainty. Specifically, to approximate the spatial distribution of an agent's future position we propose a new 2D histogram-based deep learning model combined with deep ensemble techniques which measure both aleatoric and epistemic uncertainty by entropy-based quantities in trajectory forecasting. The proposed Uncertainty Quantification Network (UQnet) employs a causal part to enhance its generalizability. Experiments on the INTERACTION dataset show that UQnet significantly improves the generalizability of the missing rate compared to the state-of-the-art. Further analysis shows that high aleatoric uncertainty cases are mainly caused by heterogeneous driving behaviours and unknown intended directions. Based on this aleatoric uncertainty component, we estimate the lower bound of the mean-square-error and final-displacement-error as an indicator for the predictability of trajectories. Furthermore, we use the epistemic uncertainty to identify rare cases in the test set. Our results illustrate that domain knowledge on speed-dependent driving behaviours is essential for adapting a model from low-speed to high-speed situations. Our paper contributes to motion forecasting with a new framework, that recasts this problem in terms of model generalisation, and puts forward methods to quantify the resulting uncertainty.