Lizenz: CC BY-SA

netidee STIPENDIEN

Abstract der fertigen Arbeit

Object detection is nowadays a highly demanded computer vision task, being applied also in safety-critical domains, such as autonomous driving, pedestrian detection, and medical image analysis. Yet, models very often remain ignorant of the uncertainty in their predictions, i.e., they may produce overconfident false predictions. To build trustworthy models, this uncertainty must be captured and handled carefully. No model is perfect, and there may always be situations in which it is better to make no prediction at all, rather than give a best-effort guess.

Therefore, in this thesis, we will explore what approaches enable object detection networks to generate estimates of their uncertainty. As these estimation methods typically do not come for free, their computational overhead has to be considered carefully. We focus on Edge AI, a domain where real-time predictions are crucial, such as for automated traffic management in a smart city. Here, information needs to be processed directly at the source with the limited hardware available.

In our work, we searched and compared the literature for promising approaches and made necessary modifications to them, enabling state-of-the-art object detection with uncertainty estimation on the Edge. In total, we implemented and evaluated four different approaches on uncertainty estimation, covering a simple baseline, ensemble networks, probabilistic approximations, and learning based on the theory of evidence. Training and evaluation are conducted across six different autonomous driving datasets, creating a domain-shift scenario for evaluating the quality of uncertainty estimation. We selected eight different metrics to evaluate the performance of our models in both object detection and uncertainty estimation. Our results show significant differences in performance between the approaches and that complex uncertainty estimation approaches may, in many cases, not outperform a simple baseline.