Abstract

This master's thesis investigates the application of transfer learning for object classification using ultrasound sensors. The background of this work lies in ultrasound-based driver assistance systems intended to detect objects in the vehicle environment during parking or maneuvering situations.

The methods studied encompass fine-tuning, freezing, and pretraining on audio data outside the ultrasound range. A domain adaptation of sensor configurations, considering the speed of airborne sound and energy losses due to distances traveled, was explored. In addition to an existing dataset, two additional datasets of different sensor configurations were created.

The results demonstrate the effectiveness of fine-tuning in adapting models to new sensor configurations. Pretraining with audio data outside the ultrasound range enhances performance, with freezing of this pretraining leading to losses. Domain adaptation shows performance improvements in studies without fine-tuning to the target sensor configuration. The height offset of a sensor does not exhibit significant effects on the results. The progression of horizontal sensor distances does not determine the optimal order for fine-tuning. The performance gains achieved through fine-tuning are therefore in big parts attributed to increased data volume.