Transfer Learning Across Human Activities Using a Cascade Neural Network Architecture

Xin Du, Katayoun Farrahi and Mahesan Niranjan

Cascade learning is a new adaptive approach to training deep neural networks. It is particularly suited to transfer learning, as learning is achieved in a layer-wise fashion, enabling the transfer of selected layers to optimize the quality of transferred features. In the domain of Human Activity Recognition (HAR), where the consideration of resource consumption is critical, cascade learning is of particular interest as it has demonstrated the ability to achieve significant reductions in computational and memory costs with negligible performance loss. In this paper, we evaluate the use of cascade learning and compare it to end to end learning in various transfer learning experiments, all applied to HAR. We consider transfer learning across objectives, for example, opening the door features transferred to opening the dishwasher. We additionally consider transfer across sensor locations on the body, as well as across datasets. Over all of our experiments, we find that cascade learning achieves state of the art performance (i.e. not too different from end to end learning) with the advantage of requiring fewer parameters (often a single transferred layer). Finally, the overall results we demonstrate for transfer learning using cascade learning and our deep learning architecture, while not directly comparable to previously published work due to varying experimental conditions, are significantly higher for the challenging task of transferring over datasets.