Online Continual Learning for Human Activity Recognition

Martin Schiemer[1], Lei Fang[1], Simon Dobson[1], Juan Ye[1]

[1]School of Computer Science, University of St Andrews, UK.

Contributing authors: ms400@st-andrews.ac.uk; lf28@st-andrews.ac.uk; simon.dobson@st-andrews.ac.uk; juan.ye@st-andrews.ac.uk;

Abstract

Sensor-based human activity recognition (HAR), with the ability to recognise human activities from wearable or embedded sensors, has played a significant role in many applications, such as personal health monitoring, smart homes, and manufacturing. The real-world, long-term deployment of these HAR systems drives a crucial research question: *how to evolve the HAR model automatically over time to accommodate changes in an environment or activity patterns*. This paper presents an online continual learning (OCL) scenario for HAR, where sensor data arrives in a streaming manner which contains unlabelled samples from previously learnt or novel activities. We introduce a technique, called *OCL-HAR*, which automatically detects novel activities from streaming sensor data and extends the current HAR model with sparsely labelled data. Compared to the existing class-incremental continual learning setting, our scenario does neither assume explicit task boundaries nor abundant labelled training data for new tasks. To the best of our knowledge, this is the first online continual learning technique specifically designed for HAR, targeting realistic real-world, long-term, deployment challenges. OCL-HAR does not assume *a prior* knowledge about any activities of interest, and it can automatically extend the model to accommodate new activity classes.

We have empirically evaluated OCL-HAR on four third-party, publicly available HAR datasets. We analysed our technique under different conditions with varying degrees of data availability and training data percentage. Our findings indicate that this OCL scenario poses a challenge for state-of-the-art continual learning techniques that have substantially underperformed. Our technique OCL-HAR has consistently outperformed them in all experimental configurations, resulting in micro and macro F1 score improvements of up to 0.17 and 0.23.