Making Neural Networks Adaptive to Changes in the Dimensions of Sensor Data

Mohammad Malekzadeh[1] Richard G. Clegg[2] Andrea Cavallaro[2] Hamed Haddadi[1]
[1]Imperial College London [2]Queen Mary University of London

Overview

Motion sensors embedded in wearable and mobile devices allow for dynamic selection of sensor streams and sampling rates, enabling useful applications, *e.g.* for power management or control of data sharing. While deep neural networks (DNNs) achieve competitive accuracy in sensor data classification, current DNN architectures only process data coming from a fixed set of sensors with a fixed sampling rate, and changes in the dimensions of their inputs cause considerable accuracy loss, unnecessary computations, or failure in operation.

To address this problem, we introduce a *dimension-adaptive pooling* (DAP) layer that makes DNNs robust to temporal changes in sampling rate and in sensor availability (see Figure 1). DAP operates on convolutional filter maps of variable dimensions and produces an input of fixed dimensions suitable for feedforward and recurrent layers. Building on this architectural improvement, we propose a *dimension-adaptive training* (DAT) procedure to generalize over the entire space of feasible data dimensions at the inference time. DAT comprises the random selection of dimensions during the forward passes and optimization with accumulated gradients of several backward passes. We then combine DAP and DAT to transform existing non-adaptive DNNs into a *Dimension-Adaptive Neural Architecture* (DANA) without altering other architectural aspects. Our solution does not need up-sampling or imputation, thus reduces unnecessary computations at inference time. Experimental results, on four benchmark datasets of human activity recognition, show that DANA prevents losses in classification accuracy of the state-of-the-art DNNs, under dynamic sensor availability and varying sampling rates.

We show that DANA outperforms the standard DNNs over a range of sampling rates and retains accuracy when some sensors are unavailable at inference time. For instance, on a dataset of 3 sensors and 13 activities, DANA keeps classification accuracy similar to the original DNN in a range of 6Hz to 50Hz and its accuracy only falls from 95% to around 90% and 85% in case of missing one or two of the three sensors, respectively, while the original DNN cannot handle these changes, or achieve at most 75% and 55% accuracy with resampling and imputation prepossessing.

Evaluation

We evaluate DANA on four public datasets of human activity recognition: *UCI-HAR*, *UTwente*, *MobiAct*, *MotionSense*, and show how to transform three benchmark DNN architectures for sensor-based human activity recognition into an adaptive neural network. We show how DAP works

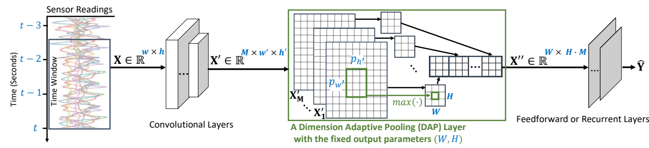

Figure 1: Overview of DAP Layer, with parameters (W, H), that enables a DNN performing classification on input data of variable dimensions.

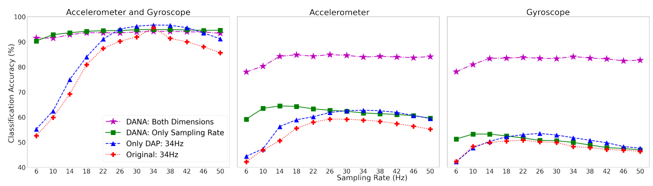

Figure 2: Classification accuracy of a benchmark DNN on UCI-HAR dataset for different sensors and sampling rate.

with and without using DAT. We also evaluate the advantages of DANA at training time, compared with three other training procedures: *standard*, *weight averaging*, and *meta-learning*. Moreover, for the inference time, we compare DANA with two alternatives baselines: *imputation* and *re-sampling* with and without *data augmentation*. Finally, we perform across datasets experiment to show the generalization of DANA. The complete version of this abstract, code, and data to reproduce results are publicly available at `https://github.com/mmalekzadeh/dana`.

As a representative example of our results, Figure 2 shows the performance of DANA in different sensor selection and sampling rate scenarios. Figure 2 (left) shows that the adaptivity to the sensor selection is not associated with a large penalty when all sensors are present. Figure 2 (middle) shows that when the gyroscope is deselected, DANA maintains its accuracy around 85% while the accuracy of other DNNs falls rapidly to around 60%. Similarly, Figure 2 (right) shows that when the accelerometer is deselected the accuracy for DANA remains around 85% while other DNNs fall to 50% or less. It is interesting that while for other DNNs deselecting accelerometer data causes more accuracy loss than deselecting gyroscope, the type of the deselected sensor has a reduced effect on DANA.

Conclusion

We presented DANA, a solution to make deep neural networks adaptive to temporal changes in the dimensions of the input data to cope with adaptive sampling and sensor selection. DANA provides a single trained model that retains high classification accuracy across a range of settings, thus avoiding the need for a separate classifier for each setting at inference time. DANA imposes no limitations on the type of DNN and is flexible in shaping the DNNs without adding or removing trainable parameters.