Deep and Dilated Residual Networks for Human Activity Recognition

Enrique Marquez, Katayoun Farrahi, Jonathon Hare, and Mahesan Niranjan

Electronics and Computer Science, University of Southampton

The standard approach to apply deep learning to sequence modeling applications is to use recurrent neural networks to capture the history of patterns over time. Current state of the art deep learning architectures for human activity recognition (HAR) include recurrent neural networks (or variations of these) as the model core, either as ensembles [1], or in combination with convolutional layers [2] to classify sequences of activities over time. Very recently, researchers in machine learning have shown that generic convolutional neural networks without recurrent structures can outperform recurrent architectures for a number of applications related to sequence modeling [3]. In this work, we further validate that recurrent neural networks are not necessary on HAR applications and demonstrate that generic convolutional architectures are state-of-the-art for HAR outperforming previous architectures on three datasets. We consider deep generic convolutional neural networks, with residual connections and discuss the optimal architectures and performance results for the Opportunity, PAMAP2, and DaphneGait datasets available on the UCI repository. We further evaluate the use of dilated convolutions to exponentially increase the size of the receptive field. For the datasets evaluated we find that dilations can improve performance for certain activities where a larger history of data can improve the predictive performance. Overall, we find that our deep residual neural networks, without any recurrent units, outperform all previously reported results for HAR. Our methods perform consistently best on all three datasets, showing that our approach also generalises well to various types of human activities and sensor placements. We show that hierarchies of convolutions are sufficient to capture the temporal dynamics of sensor data for the problem of activity classification.

References

[1] Guan, Yu, and Thomas Plötz. "Ensembles of deep lstm learners for activity recognition using wearables." Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 1.2 (2017): 11.

[2] Ordóñez, Francisco Javier, and Daniel Roggen. "Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition." Sensors 16.1 (2016): 115.

[3] Bai, Shaojie, J. Zico Kolter, and Vladlen Koltun. "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling." arXiv preprint arXiv: 1803.01271 (2018).