

# Deep, Convolutional, and Recurrent Models for Human Activity Recognition using Wearables

## Abstract

Human activity recognition (HAR) in ubiquitous computing is beginning to adopt deep learning to substitute

## 1. Introduction

Human activity recognition (HAR) in ubiquitous computing is beginning to adopt deep learning to substitute for well-established analysis techniques that rely on hand-crafted feature extraction and classification techniques.

From these isolated applications of custom depp architectures it is, however, difficult to gain an overview of their suitability for problems ranging from the recognition of manipulative gestures to the segmentation and identification of physical activites like running or ascending stairs.

In this paper we rigorously explore deep, convolutional, and recurrent approaches across three representative datasets that contain movement data captured with wearable sensors.

We describe how to train recurrent approaches in this setting, introduce a novel [regularisation](#) approach, and illustrate how they outperform the state-of-the-art on a large benchmark dataset.

Across thousands of recognition experiments with randomly sampled model configurations we investigate the suitability of each model for different tasks in HAR, explore the impact of hyperparameters using the fANVOA framework, and provide guidelines for the practitioner who wants to apply deep learning in their problem setting.

## 6. Discussion

In this work we explored the performance of state-of-the-art deep learning approaches for Human Activity Recognition using wearable sensors.

We described how to train recurrent approaches in this setting and intoduced a novel regularisation approach.

In thousands of experiements we evaluated the performance of the models with randomly sampled hyperparameters.

We found that bi-directional **LSTMs** outperform the current state-of-the-art on Opportunity, a large benchmark dataset, by a considerable margin.

However, interesting from a practitioner's point of view is not the peak performance for each mode, but the process of parameter exploration and insights into their suitability for different tasks in HAR.

Recurrent networks outperform convolutional networks significantly on activities that are short in duration but have a natural ordering, where a recurrent approach benefits from the ability to contextualise observations across long periods of time.

For bi-directional **RNNs** we found that the number of units per layer has the largest effect on performance across all datasets.

For prolonged and repetitive activities like walking or running we recommend to use **CNNs**.

Their average performance in this setting makes it more likely that the practitioner discovers a suitable configuration, even though we found some RNNs that work similarly well or even outperform CNNs in this setting.

We further recommend to start exploring learning-rates, before optimising the architecture of the network, as the learning-parameters had the largest effect on performance in our experiments.

We found that models differ in the spread of recognition performance for different parameter settings.

Regular **DNNs**, a model that is probably the most approachable for a practitioner, requires a significant investment in parameter exploration and shows a substantial spread between the peak and median performance.


Practitioners should therefore not discard the model even if a preliminary exploration leads to poor recognition performance.

More sophisticated approaches like CNNs or RNNs show a much smaller spread of performance, and it is more likely to find a configuration that works well with only a few iterations.

- <https://arxiv.org/pdf/1604.08880.pdf>
- <http://gradschoolstory.net/terry/readingpapers/>

---

## Plugin Backlinks:



From:  
<https://jace.link/> - **Various Ways**

Permanent link:  
<https://jace.link/open/deep-convolutional-and-recurrent-models-for-human-activity-recognition-using-wearables>

Last update: **2020/06/02 09:25**

