

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

Nils Y. Hammerla^{1,2}, Shane Halloran², Thomas Plötz²

¹babylon health, London, UK

²Open Lab, School of Computing Science, Newcastle University, UK
nils.hammerla@babylonhealth.com

Abstract

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 deep 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 activities 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 fANOVA framework, and provide guidelines for the practitioner who wants to apply deep learning in their problem setting.

1 Introduction

Deep learning represents the biggest trend in machine learning over the past decade. Since the inception of the umbrella term the diversity of methods it encompasses has increased rapidly, and will continue to do so driven by the resources of both academic and commercial interests. Deep learning has become accessible to everyone via machine learning frameworks like Torch7 [Collobert *et al.*, 2011], and has had significant impact on a variety of application domains [LeCun *et al.*, 2015].

One field that has yet to benefit of deep learning is Human Activity Recognition (HAR) in Ubiquitous Computing (ubicomp). The dominant technical approach in HAR includes sliding window segmentation of time-series data captured with body-worn sensors, manually designed feature extraction procedures, and a wide variety of (supervised) classifica-

tion methods [Bulling *et al.*, 2014]. In many cases, these relatively simple means suffice to obtain impressive recognition accuracies. However, more elaborate behaviours which are, for example, of interest in medical applications, pose a significant challenge to this manually tuned approach [Hammerla *et al.*, 2015]. Some work has furthermore suggested that the dominant technical approach in HAR benefits from biased evaluation settings [Hammerla and Plötz, 2015], which may explain some of the apparent inertia in the field of the adoption of deep learning techniques.

Deep learning has the potential to have significant impact on HAR in ubicomp. It can substitute for manually designed feature extraction procedures which lack the robust physiological basis that benefits other fields such as speech recognition. However, for the practitioner it is difficult to select the most suitable deep learning method for their application. Work that promotes deep learning almost always provides only the performance of the best system, and rarely includes details on how its seemingly optimal parameters were discovered. As only a single score is reported it remains unclear how this peak performance compares with the average during parameter exploration.

In this paper we provide the first unbiased and systematic exploration of the performance of state-of-the-art deep learning approaches on three different recognition problems typical for HAR in ubicomp. The training process for deep, convolutional, and recurrent models is described in detail, and we introduce a novel approach to regularisation for recurrent networks. In more than 4,000 experiments we investigate the suitability of each model for different tasks in HAR, explore the impact each model's hyper-parameters have on performance, and conclude guidelines for the practitioner who wants to apply deep learning to their application scenario. Over the course of these experiments we discover that recurrent networks outperform the state-of-the-art and that they allow novel types of real-time application of HAR through sample-by-sample prediction of physical activities.

2 Deep Learning in Ubiquitous Computing

Movement data collected with body-worn sensors are multivariate time-series data with relatively high spatial and temporal resolution (e.g. 20Hz - 200Hz). Analysis of this data in ubicomp is typically following a pipeline-based approach [Bulling *et al.*, 2014]. The first step is to segment the time-

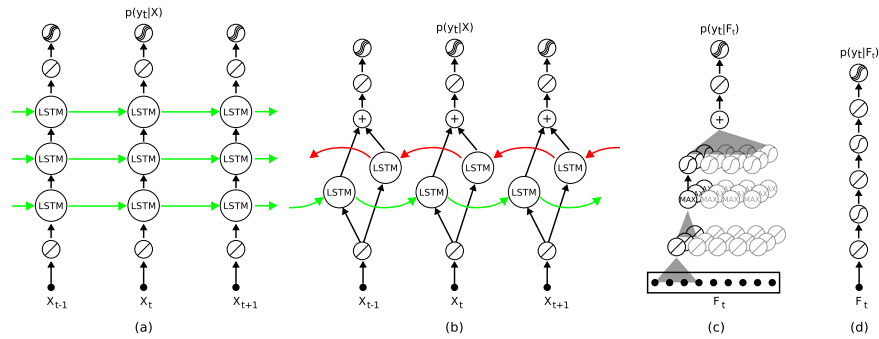


Figure 1: Models used in this work. From left to right: (a) LSTM network hidden layers containing LSTM cells and a final softmax layer at the top. (b) bi-directional LSTM network with two parallel tracks in both future direction (green) and to the past (red). (c) Convolutional networks that contain layers of convolutions and max-pooling, followed by fully-connected layers and a softmax group. (d) Fully connected feed-forward network with hidden (ReLU) layers.

series data into contiguous segments (or *frames*), either driven by some signal characteristics such as signal energy [Plötz *et al.*, 2012], or through a sliding-window segmentation approach. From each of the frames a set of features is extracted, which most commonly include statistical features or stem from the frequency domain.

The first deep learning approach explored to substitute for this manual feature selection corresponds to deep belief networks as auto-encoders, trained generatively with Restricted Boltzmann Machines (RBMs) [Plötz *et al.*, 2011]. Results were mixed, as in most cases the deep model was outperformed by a combination of principal component analysis and a statistical feature extraction process. Subsequently a variety of projects have explored the use of pre-trained, fully-connected networks for automatic assessment in Parkinson’s Disease [Hammerla *et al.*, 2015], as emission model in Hidden Markov Models (HMMs) [Zhang *et al.*, 2015; Alsheikh *et al.*, 2015], and to model audio data for ubicomp applications [Lane *et al.*, 2015].

The most popular approach so far in ubicomp relies on convolutional networks. Their performance for a variety of activity recognition tasks was explored by a number of authors [Zeng *et al.*, 2014; Ronaoo and Cho, 2015; Yang *et al.*, 2015; Ronaoo and Cho, 2015]. Furthermore, convolutional networks have been applied to specific problem domains, such as the detection of stereotypical movements in Autism [Rad *et al.*, 2015], where they significantly improved upon the state-of-the-art.

Individual frames of movement data in ubicomp are usually treated as statistically independent, and applications of sequential modelling techniques like HMMs are rare [Bulling *et al.*, 2014]. However, approaches that are able to exploit the temporal dependencies in time-series data appear as the natural choice for modelling human movement captured with sensor data. Deep recurrent networks, most notably those that rely on Long Short-Term Memory cells (LSTMs) [Hochreiter and Schmidhuber, 1997], have recently achieved impressive performance across a variety of scenarios (e.g. [Gregor *et al.*, 2015]). Their application to HAR has been explored in two settings. First, [Neverova *et al.*, 2015] investigated a variety of recurrent approaches to identify individuals based on

movement data recorded from their mobile phone in a large-scale dataset. Secondly, [Ordóñez and Roggen, 2016] compared the performance of a recurrent approach to CNNs on two HAR datasets, representing the current state-of-the-art performance on Opportunity, one of the datasets utilised in this work. In both cases, the recurrent network was paired with a convolutional network that encoded the movement data, effectively employing the recurrent network to only model temporal dependencies on a more abstract level. Recurrent networks have so far not been applied to model movement data at the lowest possible level, which is the sequence of individual samples recorded by the sensor(s).

3 Comparing deep learning for HAR

While there has been some exploration of deep models for a variety of application scenarios in HAR there is still a lack of a systematic exploration of deep learning capabilities. Authors report to explore the parameter space in preliminary experiments, but usually omit the details. The overall process remains unclear and difficult to replicate. Instead, single instantiations of e.g. CNNs are presented that show good performance in an application scenario. Solely reporting peak performance figures does, however, not reflect the overall suitability of a method for HAR in ubicomp, as it remains unclear how much effort went into tuning the proposed approach, and how much effort went into tuning other approaches it was compared to. How likely is a practitioner to find a parameter configuration that works similarly well for their application? How representative is the reported performance across the models compared during parameter exploration? Which parameters have the largest impact on performance? These questions are important for the practitioner, but so far remain unanswered in related work.

In this paper we provide the first unbiased comparison of the most popular deep learning approaches on three representative datasets for HAR in ubicomp. They include typical application scenarios like manipulative gestures, repetitive physical activities, and a medical application of HAR in Parkinson’s disease. We compare three types of models that are described below. To explore the suitability of each method

we chose reasonable ranges for each of their hyperparameters and randomly sample model configurations. We report on their performance across thousands of experiments and analyse the impact of hyperparameters for each approach¹.

3.1 Deep feed-forward networks (DNN)

We implemented deep feed-forward networks, which correspond to a neural network with up to five hidden layers followed by a softmax-group. The DNN represents a sequence of non-linear transformations to the input data of the network. We follow convention and refer to a network with N hidden layers as N -layer network. Each hidden layer contains the same number of units, and corresponds to a linear transformation and a rectified-linear (ReLU) activation function. We use two different regularisation techniques: i) Dropout: during training each unit in each hidden layer is set to zero with a probability p_{drop} , and during inference the output of each unit is scaled by $1/p_{drop}$ [Srivastava *et al.*, 2014] (dropout-rate is fixed to 0.5 for all experiments); ii) Max-in norm: After each mini-batch the incoming weights of each unit in the network are scaled to have a maximum euclidean length of d_{in} . To limit the number of hyperparameters of the approach we chose not to perform any generative pre-training and to solely rely on a supervised learning approach. The input data fed into the network corresponds to frames of movement data. Each frame consists of a varying number of s samples from \mathbb{R}^d , which are simply concatenated into a single vector $F_t \in \mathbb{R}^{s*d}$. The model is illustrated in figure 1(d).

The DNN is trained in a mini-batch approach, where each mini-batch contains 64 frames and is stratified with respect to the class distribution in the training-set. We minimise the negative log likelihood using stochastic gradient descent.

3.2 Convolutional networks (CNN)

CNNs aim to introduce a degree of locality in the patterns matched in the input data and to enable translational invariance with respect to the precise location (i.e. time of occurrence) of each pattern within a frame of movement data. We explore the performance of convolutional networks and follow suggestions by [Srivastava *et al.*, 2014] in architecture and regularisation techniques. The overall CNN architecture is illustrated in figure 1(c). Each CNN contains at least one temporal convolution layer, one pooling layer and at least one fully connected layer prior to a top-level softmax-group. The temporal convolution layer corresponds to a convolution of the input sequence with n_f different *kernels* (feature maps) of width k_w . Subsequent max-pooling is looking for the maximum within a region of width m_w and corresponds to a sub-sampling, introducing translational invariance to the system. Width of the max-pooling was fixed to 2 throughout all experiments. The output of each max-pooling layer is transformed using a ReLU activation function. The subsequent fully connected part effectively corresponds to a DNN and follows the same architecture outlined above.

For regularisation we apply dropout after each max-pooling or fully-connected layer, where the dropout-probability p_{drop}^i in layer i is fixed for all experiments

($p_{drop}^1 = 0.1$, $p_{drop}^2 = 0.25$, $p_{drop}^{i>2} = 0.5$). Similar to the DNN we also perform max-in norm regularisation as suggested in [Srivastava *et al.*, 2014]. The input data fed into the CNN corresponds to frames of movement data as in the DNN. However, instead of concatenating the different input dimensions the matrix-structure is retained ($F_t \in \mathbb{R}^s \times \mathbb{R}^d$). The CNN is trained using stratified mini-batches (64 frames) and stochastic gradient descent to minimise negative log likelihood.

3.3 Recurrent networks

In order to exploit the temporal dependencies within the movement data we implemented recurrent neural networks based on LSTM cells in their *vanilla* variant that does not contain *peephole* connections [Greff *et al.*, 2015]. This architecture is recurrent as some of the connections within the network form a directed cycle, where the current timestep t considers the states of the network in the previous timestep $t - 1$. LSTM cells are designed to counter the effect of diminishing gradients if error derivatives are backpropagated through many layers “through time” in recurrent networks [Hochreiter *et al.*, 2001]. Each LSTM cell (unit) keeps track of an internal state (the *constant error carousel*) that represents its “memory”. Over time the cells learn to output, overwrite, or null their internal memory based on their current input and the history of past internal states, leading to a system capable of retaining information across hundreds of time-steps [Hochreiter and Schmidhuber, 1997].

We implement two flavours of LSTM recurrent networks: i) deep forward LSTMs contain multiple layers of recurrent units and are connected “forward” in time (see figure 1(a)); and ii) bi-directional LSTMs which contain two parallel recurrent layers that stretch both into the “future” and into the “past” of the current time-step, followed by a layer that concatenates their internal states for timestep t (see figure 1(b)).

Practically these two flavours differ significantly in their application requirements. A forward LSTM contextualises the current time-step based on those it has seen previously, and is inherently suitable for real-time applications where, at inference time, the “future” is not yet known. Bi-directional LSTMs on the other hand use both the future and past context to interpret the input at timestep t , which makes them suitable for offline analysis scenarios.

In this work we apply recurrent networks in three different settings, each of which is trained to minimise the negative log likelihood using *adagrad* [Duchi *et al.*, 2011] and subject to max-in norm regularisation.

In the first case the input data fed into the network at any given time t corresponds to the current frame of movement data, which stretches a certain length of time and whose dimensions have been concatenated (as in the DNN above). We denote this model as *LSTM-F*. The second application case of forward LSTMs represents a real-time application, where each sample of movement data is presented to the network in the sequence they were recorded, denoted *LSTM-S*. The final scenario sees the application of bi-directional LSTMs to the same sample-by-sample prediction problem, denoted *b-LSTM-S*.

¹Source-code will be made publicly available

3.4 Training RNNs for HAR

Common applications for RNNs include speech recognition and natural language processing. In these settings the context for an input (e.g. a word) is limited to its surrounding entities (e.g. a sentence, paragraph). Training of RNNs usually treats these contextualised entities as a whole, for example by training an RNN on complete sentences.

In HAR the context of an individual sample of movement data is not well defined, at least beyond immediate correlations between neighbouring samples, and likely depends on the type of movement and its wider behavioural context. This is a problem well known in the field, and affects the choice of window length for sliding window segmentation [Bulling *et al.*, 2014].

In order to construct b mini-batches that are used to train the RNN we initialise a number of positions $(p_i)_b$ between the start and end of the training-set. To construct a mini-batch we extract the L samples that follow each position in $(p_i)_b$, and increase $(p_i)_b$ by L steps, possibly wrapping around the end of the sequence. We found that it is important to initialise the positions randomly to avoid oscillations in the gradients. While this approach retains the ordering of the samples presented to the RNN it does not allow for stratification of each mini-batch w.r.t. class-distribution.

Training on long sequences has a further issue that is addressed in this work. If we use the approach outlined above to train a sufficiently large RNN it may “memorise” the entire input-output sequence implicitly, leading to poor generalisation performance. In order to avoid this memorisation we need to introduce “breaks” where the internal states of the RNN are reset to zero: after each mini-batch we decide to retain the internal state of the RNN with a carry-over probability p_{carry} , and reset it to zero otherwise. This is a novel form of regularisation of RNNs, which should be useful for similar applications of RNNs.

4 Experiments

The different hyper-parameters explored in this work are listed in table 1. The last column indicates the number of parameter configurations sampled for each dataset, selected to represent an equal amount of computation time. We conduct experiments on three benchmark datasets representative of the problems typical for HAR (described below). Experiments were run on a machine with three GPUs (NVIDIA GTX 980 Ti), where two model configurations are run on each GPU except for the largest networks.

After each epoch of training we evaluate the performance of the model on the validation set. Each model is trained for at least 30 epochs and for a maximum of 300 epochs. After 30 epochs, training stops if there is no increase in validation performance for 10 subsequent epochs. We select the epoch that showed the best validation-set performance and apply the corresponding model to the test-set.

4.1 Datasets

We select three datasets typical for HAR in ubicomp for the exploration in this work. Each dataset corresponds to

an application of HAR. The first dataset, Opportunity, contains manipulative gestures like opening and closing doors, which are short in duration and non-repetitive. The second, PAMAP2, contains prolonged and repetitive physical activities typical for systems aiming to characterise energy expenditure. The last, Daphnet Gait, corresponds to a medical application where participants exhibit a typical motor complication in Parkinson’s disease that is known to have a large inter-subject variability. Below we detail each dataset:

The **Opportunity dataset (Opp)** [Chavarriaga *et al.*, 2013] consists of annotated recordings from on-body sensors from 4 participants instructed to carry out common kitchen activities. Data is recorded at a frequency of 30Hz from 12 locations on the body, and annotated with 18 mid-level gesture annotations (e.g. Open Door / Close Door). For each subject, data from 5 different runs is recorded. We used the subset of sensors that did not show any packet-loss, which included accelerometer recordings from the upper limbs, the back, and complete IMU data from both feet. The resulting dataset had 79 dimensions. We use run 2 from subject 1 as our validation set, and replicate the most popular recognition challenge by using runs 4 and 5 from subject 2 and 3 in our test set. The remaining data is used for training. For frame-by-frame analysis, we created sliding windows of duration 1 second and 50% overlap. The resulting training-set contains approx. 650k samples (43k frames).

The **PAMAP2 dataset** [Reiss and Stricker, 2012] consists of recordings from 9 participants instructed to carry out 12 lifestyle activities, including household activities and a variety of exercise activities (Nordic walking, playing soccer, etc). Accelerometer, gyroscope, magnetometer, temperature and heart rate data are recorded from inertial measurement units located on the hand, chest and ankle over 10 hours (in total). The resulting dataset has 52 dimensions. We used runs 1 and 2 for subject 5 in our validation set and runs 1 and 2 for subject 6 in our test set. The remaining data is used for training. In our analysis, we downsampled the accelerometer data to 33.3Hz in order to have a temporal resolution comparable to the Opportunity dataset. For frame-by-frame analysis, we replicate previous work with non-overlapping sliding windows of 5.12 seconds duration with one second stepping between adjacent windows (78% overlap) [Reiss and Stricker, 2012]. The training-set contains approx. 473k samples (14k frames).

The **Daphnet Gait dataset (DG)** [Bachlin *et al.*, 2009] consists of recordings from 10 participants affected with Parkinson’s Disease (PD), instructed to carry out activities which are likely to induce freezing of gait. Freezing is a common motor complication in PD, where affected individuals struggle to initiate movements such as walking. The objective is to detect these freezing incidents, with the goal to inform a future situated prompting system. This represents a two-class recognition problem. Accelerometer data was recorded from above the ankle, above the knee and on the trunk. The resulting dataset has 9 dimensions. We used run 1 from subject 9 in our validation set, runs 1 and 2 from subject 2 in our test set, and used the rest for training. In our analysis, we downsampled the accelerometer data to 32Hz. For frame-by-frame analysis, we created sliding windows of 1 second

Category	log-uniform?	LR	LR decay	L	momentum	max-in norm	p_{carry}	#layers	#units	#conv.-layers	KW 1	KW 2	KW 3	nF 1	nF 2	nF 3	#experiments												
		Learning																Regularisation		Architecture									
		y	y															-	-	-	-	-	-	-	-	-	-	-	-
DNN	max	10^{-1}	10^{-3}	-	0.99	4.0	-	5	2048	-	-	-	-	-	-	-	-	1000											
	min	10^{-4}	10^{-5}	-	0.0	0.5	-	1	64	-	-	-	-	-	-	-	-												
CNN	max	10^{-1}	10^{-3}	-	0.99	4.0	-	3	2048	3	9	5	3	128	128	128	256												
	min	10^{-4}	10^{-5}	-	0.0	0.5	-	1	64	1	3	3	3	16	16	16													
LSTM-F	max	10^{-1}	-	64	-	4.0	1.0	3	384	-	-	-	-	-	-	-	128												
	min	10^{-3}	-	8	-	0.5	0.0	1	64	-	-	-	-	-	-	-													
LSTM-S	max	10^{-1}	-	196	-	4.0	1.0	3	384	-	-	-	-	-	-	-	128												
	min	10^{-3}	-	32	-	0.5	0.0	1	64	-	-	-	-	-	-	-													
b-LSTM-S	max	10^{-1}	-	196	-	4.0	1.0	1	384	-	-	-	-	-	-	-	128												
	min	10^{-3}	-	32	-	0.5	0.0	1	64	-	-	-	-	-	-	-													

Table 1: Hyper-parameters of the models and the ranges of values explored in experiments.

duration and 50% overlap. The training-set contains approx. 470k samples (30k frames).

4.2 Influence of hyper-parameters

In order to estimate the impact of each hyperparameter on the performance observed across all experiments we apply the fANOVA analysis framework. fANOVA [Hutter *et al.*, 2014] determines the extent to which each hyperparameter contributes to a network’s performance. It builds a predictive model (random forest) of the model performance as a function of the model’s hyperparameters. This non-linear model is then decomposed into marginal and joint interaction functions of the hyperparameters, from which the percentage contribution to overall variability of network performance is obtained. fANOVA has been used previously to explore the hyperparameters in recurrent networks by [Greff *et al.*, 2015].

For the practitioner it is important to know which aspect of the model is the most crucial for performance. We grouped the hyperparameters of each model into one of three categories (see table 1): i) *learning*, parameters that control the learning process; ii) *regularisation*, parameters that limit the modelling capabilities of the model to avoid overfitting; and iii) *architecture*, parameters that affect the structure of the model. Based on the variability observed for each hyperparameter we estimate the variability that can be attributed to each parameter category, and to higher order interactions between categories.

4.3 Performance metrics

As the datasets utilised in this work are highly biased we require performance metrics that are independent of the class distribution. We opted to estimate the mean f1-score:

$$F_m = \frac{2}{|c|} \sum_c \frac{\text{prec}_c \times \text{recall}_c}{\text{prec}_c + \text{recall}_c} \quad (1)$$

Related work has previously used the weighted f1-score as primary performance metric (for Opportunity). In order to compare our results to the state-of-the-art we estimate the weighted f1-score:

$$F_w = 2 \sum_c \frac{N_c}{N_{\text{total}}} \frac{\text{prec}_c \times \text{recall}_c}{\text{prec}_c + \text{recall}_c}, \quad (2)$$

Performance	PAMAP2	DG	OPP	
	F_m	F_1	F_m	F_w
DNN	0.904	0.633	0.575	0.888
CNN	0.937	0.684	0.591	0.894
LSTM-F	0.929	0.673	0.672	0.908
LSTM-S	0.882	0.760	0.698	0.912
b-LSTM-S	0.868	0.741	0.745	0.927
CNN	[Yang <i>et al.</i> , 2015]		-	0.851
CNN	[Ordóñez and Roggen, 2016]		0.535	0.883
DeepConvLSTM	[Ordóñez and Roggen, 2016]		0.704	0.917
Delta from median	ΔF_m	ΔF_1	ΔF_m	mean
DNN	0.129	0.149	0.357	0.221
CNN	0.071	0.122	0.120	0.104
LSTM-F	0.10	0.281	0.085	0.156
LSTM-S	0.128	0.297	0.079	0.168
b-LSTM-S	0.087	0.221	0.205	0.172

Table 2: Best results obtained for each model and dataset, along with some baselines for comparison. Delta from median (lower part of table) refers to the absolute difference between peak and median performance across all experiments.

where N_c is the number of samples in class c , and N_{total} is the total number of samples.

5 Results

Results are illustrated in figure 2. Graphs (a-c) show the cumulative distribution of the main performance metric on each dataset. Graph (d) illustrates the effect of each category of hyper-parameter estimated using fANOVA.

Overall we observe a large spread of peak performances between models on OPP and DG, with more than 15% mean f1-score between the best performing approach (b-LSTM-S) and the worst (DNN) on OPP (12% on DG) (see table 2). On PAMAP2 this difference is smaller, but still considerable at 7%. The best performing approach on OPP (b-LSTM-S) outperforms the current state-of-the-art by a considerable margin of 4% mean f1-score (1% weighted f1-score). The best CNN discovered in this work further outperforms previous results reported in the literature for this type of model by more than 5% mean f1-score and weighted f1-score (see table 2). The good performance of recurrent approaches, which model movement at the sample level, holds the potential for novel (real-time) applications in HAR, as they alleviate the need for segmentation of the time-series data.




Figure 2: (a)-(c): Cumulative distribution of recognition performance for each dataset. (d): results from fANOVA analysis, illustrating impact of hyperparameter-categories on recognition performance (see table 1).

The distributions of performance scores differ between the models investigated in this work. CNNs show the most characteristic behaviour: a fraction of model configurations do not work at all (e.g. 20% on PAMAP2), while the remaining configurations show little variance in their performance. On PAMAP2, for example, the difference between the peak and median performance is only 7% mean f1-score (see table 2). The DNNs show the largest spread between peak and median performance of all approaches of up to 35.7% on OPP. Both forward RNNs (LSTM-F, LSTM-S) show similar behaviour across the different datasets. Practically all of their configurations explored on PAMAP2 and OPP have non-trivial recognition performance.

The effect of each category of hyperparameter on the recognition performance is illustrated in figure 2(d). Interestingly, we observe the most consistent effect of the parameters in the CNN. In contrast to our expectation it is the parameters surrounding the learning process (see table 1) that have the largest main effect on performance. We expected that for this model the rich choice of architectural variants should have a larger effect. For DNNs we do not observe a systematic effect of any category of hyperparameter. On PAMAP2, the correct learning parameters appear to be the most crucial. On OPP it is the architecture of the model. Interestingly we observed that relatively shallow networks outperform deeper variants. There is a drop in performance for networks with more than 3 hidden layers. This may be related to our choice to solely rely on supervised training, where a generative pre-training may improve the performance of deeper networks.

The performance of the frame-based RNN (LSTM-F) on OPP depends critically on the carry-over probability introduced in this work. Both always retaining the internal state and always forgetting the internal state lead to the low performance. We found that p_{carry} of 0.5 works well for most settings. Our findings merit further investigation, for example into a carry-over *schedule*, which may further improve LSTM performance.

Results for sample-based forward LSTMs (LSTM-S) mostly confirm earlier findings for this type of model that found learning-rate to be the most crucial parameter [Greff *et al.*, 2015]. However, for bi-directional LSTMs (b-LSTM-S) we observe that the number of units in each layer has a

surprisingly large effect on performance, which should motivate practitioners to first focus on tuning this parameter.

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 introduced a novel regularisation approach. In thousands of experiments 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 model, 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.

References

- [Alsheikh *et al.*, 2015] Mohammad Abu Alsheikh, Ahmed Selim, Dusit Niyato, Linda Doyle, Shaowei Lin, and Hwee-Pink Tan. Deep activity recognition models with triaxial accelerometers. *arXiv:1511.04664*, 2015.
- [Bachlin *et al.*, 2009] Marc Bachlin, Daniel Roggen, Gerhard Troster, Meir Plotnik, Noit Inbar, Inbal Meidan, Talia Herman, Marina Brozgol, Eliya Shaviv, Nir Giladi, et al. Potentials of enhanced context awareness in wearable assistants for parkinson’s disease patients with the freezing of gait syndrome. In *ISWC*, 2009.
- [Bulling *et al.*, 2014] Andreas Bulling, Ulf Blanke, and Bernt Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys (CSUR)*, 46(3):33, 2014.
- [Chavarriaga *et al.*, 2013] Ricardo Chavarriaga, Hesam Saha, Alberto Calatroni, Sundara Tejaswi Digumarti, Gerhard Tröster, José del R Millán, and Daniel Roggen. The opportunity challenge: A benchmark database for on-body sensor-based activity recognition. *Pattern Recognition Letters*, 2013.
- [Collobert *et al.*, 2011] Ronan Collobert, Koray Kavukcuoglu, and Clément Farabet. Torch7: A matlab-like environment for machine learning. In *BigLearn, NIPS Workshop*, 2011.
- [Duchi *et al.*, 2011] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.
- [Greff *et al.*, 2015] Klaus Greff, Rupesh Kumar Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *arXiv preprint arXiv:1503.04069*, 2015.
- [Gregor *et al.*, 2015] Karol Gregor, Ivo Danihelka, Alex Graves, and Daan Wierstra. Draw: A recurrent neural network for image generation. *arXiv:1502.04623*, 2015.
- [Hammerla and Plötz, 2015] Nils Y Hammerla and Thomas Plötz. Let’s (not) stick together: pairwise similarity biases cross-validation in activity recognition. In *UbiComp*, 2015.
- [Hammerla *et al.*, 2015] Nils Y Hammerla, James M Fisher, Peter Andras, Lynn Rochester, Richard Walker, and Thomas Plötz. Pd disease state assessment in naturalistic environments using deep learning. In *AAAI*, 2015.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [Hochreiter *et al.*, 2001] Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, and Jürgen Schmidhuber. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001.
- [Hutter *et al.*, 2014] Frank Hutter, Holger Hoos, and Kevin Leyton-Brown. An efficient approach for assessing hyperparameter importance. In *ICML*, pages 754–762, 2014.
- [Lane *et al.*, 2015] Nicholas D Lane, Petko Georgiev, and Lorena Qendro. Deeppear: robust smartphone audio sensing in unconstrained acoustic environments using deep learning. In *UbiComp*, pages 283–294. ACM, 2015.
- [LeCun *et al.*, 2015] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 2015.
- [Neverova *et al.*, 2015] Natalia Neverova, Christian Wolf, Griffin Lacey, Lex Fridman, Deepak Chandra, Brandon Barbello, and Graham Taylor. Learning human identity from motion patterns. *arXiv:1511.03908*, 2015.
- [Ordóñez and Roggen, 2016] Francisco Javier Ordóñez and Daniel Roggen. Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1):115, 2016.
- [Plötz *et al.*, 2011] Thomas Plötz, Nils Y Hammerla, and Patrick Olivier. Feature learning for activity recognition in ubiquitous computing. In *IJCAI*, 2011.
- [Plötz *et al.*, 2012] Thomas Plötz, Nils Y Hammerla, Agata Rozga, Andrea Reavis, Nathan Call, and Gregory D Abowd. Automatic assessment of problem behavior in individuals with developmental disabilities. In *UbiComp*, 2012.
- [Rad *et al.*, 2015] Nastaran Mohammadian Rad, Andrea Bizzego, Seyed Mostafa Kia, Giuseppe Jurman, Paola Venuti, and Cesare Furlanello. Convolutional neural network for stereotypical motor movement detection in autism. *arXiv:1511.01865*, 2015.
- [Reiss and Stricker, 2012] Attila Reiss and Didier Stricker. Introducing a new benchmarked dataset for activity monitoring. In *ISWC*, 2012.
- [Ronaoo and Cho, 2015] Charissa Ann Ronaoo and Sung-Bae Cho. Deep convolutional neural networks for human activity recognition with smartphone sensors. In *Neural Information Processing*, pages 46–53. Springer, 2015.
- [Ronaoo and Cho, 2015] Charissa Ann Ronaoo and Sung-Bae Cho. Evaluation of deep convolutional neural network architectures for human activity recognition with smartphone sensors. In *Proc. of the KIISE Korea Computer Congress*, pages 858–860, 2015.
- [Srivastava *et al.*, 2014] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *JMLR*, 2014.
- [Yang *et al.*, 2015] Jian Bo Yang, Minh Nhut Nguyen, Phyo Phyo San, Xiao Li Li, and Shonali Krishnaswamy. Deep convolutional neural networks on multichannel time series for human activity recognition. In *IJCAI*, 2015.
- [Zeng *et al.*, 2014] Ming Zeng, Le T Nguyen, Bo Yu, Ole J Mengshoel, Jiang Zhu, Pang Wu, and Juyong Zhang. Convolutional neural networks for human activity recognition using mobile sensors. In *MobiCASE*, pages 197–205. IEEE, 2014.
- [Zhang *et al.*, 2015] Licheng Zhang, Xihong Wu, and Dingsheng Luo. Human activity recognition with hmm-dnn model. In *ICCI*, pages 192–197. IEEE, 2015.