Accepted at the IEEE 20th International Workshop on Multimedia Signal Processing (MMSP 2018).

A Recurrent Neural Network for Multisensory Non-Intrusive Load Monitoring on a Raspberry Pi

Alon Harell, Stephen Makonin, and Ivan V. Bajić
School of Engineering Science, Simon Fraser University, Burnaby, Canada
Email: aharell@sfu.ca, smakonin@sfu.ca, ibajic@ensc.sfu.ca

***Abstract*—Understanding how appliances consume power is important for energy conservation. Non-intrusive load monitoring (NILM) helps meeting energy conservation goals by inferring individual appliance power usage from a single measurement. By using additional, readily available, sensor information such as weather data, it is possible to improve the accuracy of NILM. In this paper we present an example case of disaggregating two appliances using a recurrent neural network, and show how the use of multisensory data (weather and power) as an input improves performance. We demonstrate the system on a Raspberry Pi – a cost-efficient platform with limited computational capabilities.**

***Index Terms*—NILM, disaggregation, RNN, LSTM, weather**

I. Prototype Description

Non-intrusive load monitoring (NILM) [?] traditionally uses the aggregate power meter as the only sensor. Multisensory NILM can help improve the accuracy of appliance usage detection. In order to demonstrate the usefulness of multisensory data for NILM, we compare the performance of two neural networks on readings from two sub-meters taken from the AMPds2 [?] dataset. Both networks are hybrid recurrent neural networks, which output both an estimated total power and a simplified ON/OFF state classification, for each appliance. A detailed schematic of the architecture is presented in Fig. ??.

One network will use only the total power as an input, while the other will also incorporate temperature readings. Training, validation, and testing were all performed on a subset of the data from the AMPds2 [?] dataset. In AMPds2, power readings are taken from a central meter as well as 20 sub-meters in one house over two years. In order to simplify the problem, we take only two of the sub-meters (HPE and OFE) and manually add them to create the aggregated power signal. These two sub-meters were selected because they have opposite correlation with the temperature measurement while almost no correlation with each other. Training was done on the first 600 days of measurement (500 for training, 100 for validation) and testing was performed on the following 100 days. Training was limited to 100 epochs or until validation loss plateaued. After training, the weights of both networks were loaded on to a Raspberry Pi 3 computer (model B V1.2), which runs each network and performs power disaggregation faster than real-time, presenting dynamically updating graphs of both appliances and their comparison to ground truth data.

This work was funded in part by IC-IMPACTS.

Having both the classification and total power estimate at the output allows us to compare the two networks using both the F-score and Estimation Accuracy [?]. As seen in Table ??, the inclusion of temperature in the input data improves the performance in both metrics, especially for OFE.

TABLE I
F1-Score and Estimation Accuracy Results

Input	Sub-meter	F1-Score	Est Acc
Power	HPE	0.9997	0.966
	OFE	0.790	0.535
	Overall	0.865	0.912
Power + Temperature	HPE	0.9996	0.976
	OFE	0.847	0.688
	Overall	0.903	0.939

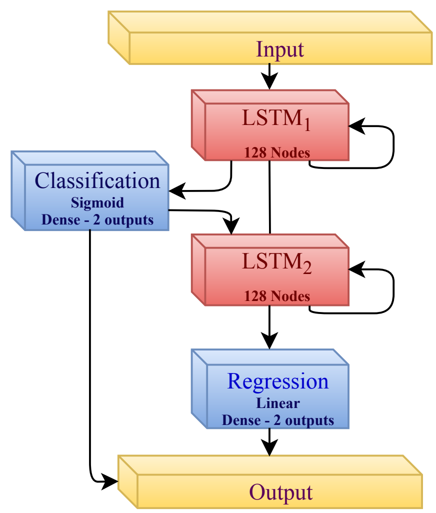

Fig. 1. Network architecture