DeepVentilation: Learning to Predict Physical Effort from Breathing

Sagar Sen[1*], **Pierre Bernabé**[1] and **Erik Johannes B. L. G. Husom**[2]
[1]Simula Research Laboratory, Pb 135, 1325 Lysaker, Norway
[2]Department of Physics, University of Oslo, Norway
{sagar, pierbernabe}@simula.no, ejhusom@uio.no

Abstract

Tracking physical effort from physiological signals has enabled people to manage required activity levels in our increasingly sedentary and automated world. Breathing is a physiological process that is a reactive and realistic representation of physical effort. In this demo, we present **DeepVentilation**, a deep learning system to predict *minute ventilation* in litres of air a person moves in one minute uniquely from real-time measurement of rib-cage movement due to breathing. DeepVentilation has been trained on input signals of expansion and contraction of the rib-cage obtained using a non-invasive respiratory inductance plethysmography sensor to predict minute ventilation as observed from a face/head mounted exercise spirometer. The system is used to track physical effort closely matching our perception of actual exercise intensity. The source code for the demo is available here: ***https://github.com/simula-vias/DeepVentilation***

1 Introduction

Global physical activity levels have declined substantially over the last five decades [Ozemek *et al.*, 2019]. Therefore, a market for consumer wearable devices to track number of steps [Bassett *et al.*, 2017], and heart rate [Thomson *et al.*, 2019] has grown enormously in our so-called *health society* [Adams, 2019]. Step counting is tremendously popular with devices such as the FitBit [Diaz *et al.*, 2015] and the Apple Watch [Veerabhadrappa *et al.*, 2018] where people aim to reach the quintessential 10,000 steps per day [Schneider *et al.*, 2006]. Heart rate monitors have been popular with athletes for several years since their invention in 1977 by Polar Elektro. The integration of *near infrared spectroscopy* (NIRS) on smartwatches has made heart rate measurement ubiquitous. Heart rate measures *intensity of work* and goes beyond step counting to provide a more fine-grained feedback on physical effort. However, the heart rate exhibits *cardiovascular drift* [Coyle and Gonzalez-Alonso, 2001] which refers to the increase in heart rate that occurs during prolonged endurance exercise with little or no change in workload. In addition, heart rate is slow to react to the real physical effort which sometimes varies quickly such as in high-intensity interval training (HIIT). Therefore, we ask, can our breathing help us predict physical effort in a more reactive and representative manner?

DeepVentilation is a deep learning system that has been trained to predict *minute ventilation* in litres per minute directly from the expansion and contraction of the breathing muscles around the ribcage. The output minute ventilation is the amount of air a person moves in one minute which is typically measured using a face/head mounted exercise spirometer. While, the input is the measurement of breathing forces (in *millivolts* across a strain gauge) due to ribcage expansion and contraction. It is measured from a *respiratory inductance plethysmography* (RIP) sensor called Flow [1] [Laugstøl, 2018]. The raw ribcage movement signals contain information about change in lung volume which DeepVentilation leverages to predict minute ventilation. Minute ventilation as predicted by DeepVentilation instantly follows exercise intensity (in comparison to standard heart rate) matching the user's perception of physical effort. The measurement of ground truth minute ventilation using a face/head mask and spirometer does not limit breathing to the best of our knowledge of scientific literature. However, the face/head mask is not always portable and convenient in many sports activities including swimming and contact sports.

The rest of the article is organized as follows. In Section 2, we describe the data set used to train DeepVentilation. In Section 3, we present the complete architecture of DeepVentilation and it evaluate with respect to ground truth data from a spirometer. We conclude in Section 4.

2 Training Data

Predicting minute ventilation required us to collect ground truth data from an exercise spirometer at the same time as obtaining data from a RIP sensor. The data collection has been performed by exercise physiologists at the Norwegian School of Sports Science[2] as part of a joint project.

[*]Contact Author

[1]http://www.sweetzpot.com/flow
[2]http://www.nih.no

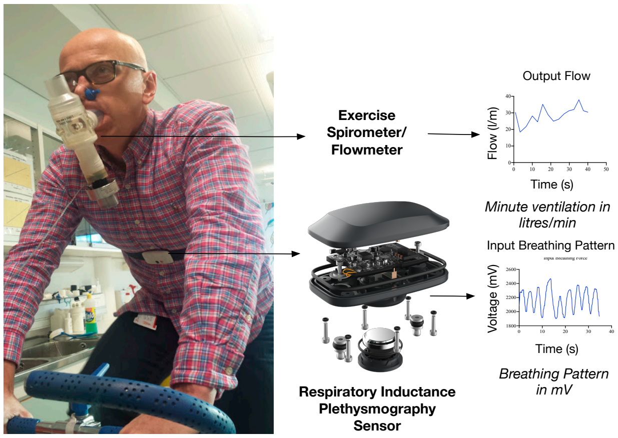

Figure 1: Training Data from a Spirometer and a RIP sensor

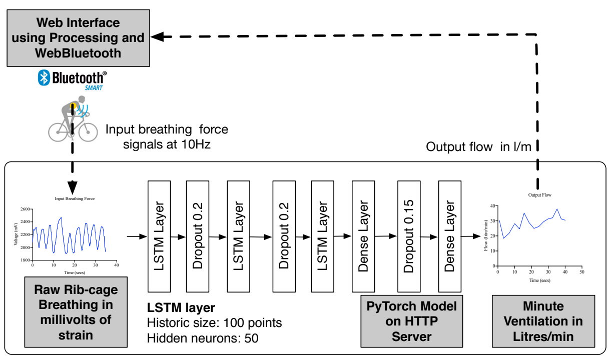

Figure 2: DeepVentilation's Architecture

We carried out measurements from five subjects (all male, aged 26±1 years), on a cycle ergometer (17980 Lode Excalibur Sport, Lode BV, Groningen, Netherlands) during a sub-maximal exercise test (at three different power levels) as well as an incremental exercise/ramp test. These measurements were repeated for each subject on two separate days. We obtained input rib-cage movement data using the Flow RIP sensor containing a *semi-conductor strain gauge* measuring forces through the click on button attached to a chest strap as shown in Figure 1. The output tidal volume and minute ventilation was measured simultaneously with a Douglas Bag (components from Harvard Apparatus, Kent, UK) also shown in Figure 1. The data was synchronized by means of *three deep breaths*. The measurements were carried out over a period of seven weeks. The data is made available in four columns: *timestamp (s), minute ventilation (l/m), ribcage movement (mV), heart rate (bpm)* to a deep learning model.

3 Architecture

DeepVentilation's architecture is illustrated in Figure 2. Ribcage breathing data (strain in the range 0 to 4096mV) is transmitted via Bluetooth Low Energy protocol [Nikodem and Bawiec, 2020] to a Web Application running on Google Chrome. DeepVentilation transforms a sliding window of received values to one value in litres/min (l/m) using several layers of *long short term memory networks* (LSTMs)[Greff *et al.*, 2016] as shown in Figure 2. The LSTM network model is implemented in PyTorch[Ketkar, 2017] and is available as a running web service through a RESTful web API [Richardson *et al.*, 2013]. We use an LSTM model because its recurrent neural network architecture is capable of learning long-term dependencies in a sequence of raw breathing data. The dropout layers in the architecture are used to regularize each LSTM layer by dropping neurons with a probability of 0.2. It is an effective method to remove large weights that may cause over-fitting of the data. The last layer in the network is a dense fully connected layer that transforms 10 values from the last LSTM node to a single value in l/m. A sequence of 100 points of raw breathing data sampled at 10Hz is sent in a request to the neural network model. The web application receives a response from the API and renders the predicted minute ventilation in real-time. The real-time feedback is developed using pure Javascript and sensor connectivity is handled via Web Bluetooth [3].

It is important to note that the neural network model is only 50 Kb in size and will eventually be embedded directly into a web page on the client side using the TensorFlow.js[4] library. This will be achieved by converting PyTorch to the ONNX standard[5], and subsequently to TensorFlow and TensorFlow.js. If the client side device is GPU accelerated then the speed of prediction will be comparable to running on the model on a server. This makes it possible to embed the model directly into the browser that can also be accessed on a smartphone as Web Bluetooth is available on most Android phones.

3.1 Evaluation

We compared the output of DeepVentilation's LSTM model to the ground truth data obtained from the exercise spirometer as shown in Figure 3. The evaluation shown in the figure corresponds to a ramp test where a user starts at low power output and gradually increases power output until he in this case could no longer increase power output. In 90 % of the breathing data we observed a maximum deviation of 20 % from the ground truth. The LSTM model exhibits lower fluctuation which gives a perception of stability to a user. We believe that the model will improve in accuracy if a second sensor is used to also measure abdominal breathing. Nevertheless, for an easier user experience our aim is to achieve reasonable accuracy primarily from ribcage breathing.

The model can be fine tuned for one person if the person is willing to provide data for ventilation and ribcage movement. This is would be an interesting N-of-1 trial [McDonald *et al.*, 2017]. However, the goal of the presented model is be as generic as possible such that it is robust to the breathing patterns of a large number of people. At the time of writing, the model is trained on 5 users already and does not need to be trained for each user before use. The test user was not part of the training data as the data was collected in 2017

[3] https://webbluetoothcg.github.io/
[4] https://www.tensorflow.org/js
[5] https://onnx.ai/

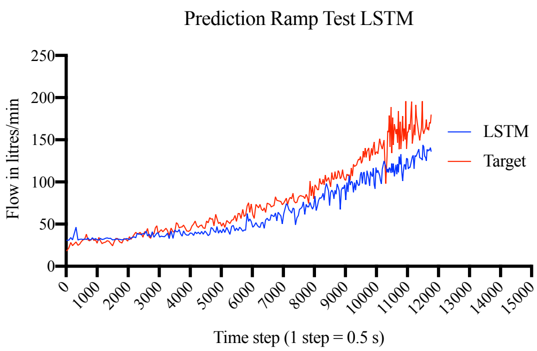

Figure 3: Evaluation of LSTM Model in DeepVentilation

independently for 5 users. The idea is a proof of concept with promising results. However, we need more data which we will be collecting from people with different attributes : ethnicity, age, weight, height, trained/untrained, temperate/tropical/cold climate and with the sensor positioned with certain variability in tension and position.

4 Conclusion

DeepVentilation is a system that predict physical effort based on minute ventilation. It can continually improve with additional data from different endurance sports. DeepVentilation can handle uncertainty due to muscular artifacts from movements other than breathing, sensor position on the body and strap tightness. It can also be trained to distinguish between genders, age, weight, height, and fitness level for diverse user groups. Novel neural network models such as attention-based transformer models[Vaswani *et al.*, 2017] can improve prediction accuracy as it is independent of the sequence of input data.

References

[Adams, 2019] Mary Louise Adams. Step-counting in the "health-society": Phenomenological reflections on walking in the era of the fitbit. *Social Theory & Health*, 17(1):109–124, 2019.

[Bassett *et al.*, 2017] David R Bassett, Lindsay P Toth, Samuel R LaMunion, and Scott E Crouter. Step counting: a review of measurement considerations and health-related applications. *Sports Medicine*, 47(7):1303–1315, 2017.

[Coyle and Gonzalez-Alonso, 2001] Edward F Coyle and J Gonzalez-Alonso. Cardiovascular drift during prolonged exercise: new perspectives. *Exercise and sport sciences reviews*, 29(2):88–92, 2001.

[Diaz *et al.*, 2015] Keith M Diaz, David J Krupka, Melinda J Chang, James Peacock, Yao Ma, Jeff Goldsmith, Joseph E Schwartz, and Karina W Davidson. Fitbit®: An accurate and reliable device for wireless physical activity tracking. *International journal of cardiology*, 185:138–140, 2015.

[Greff *et al.*, 2016] Klaus Greff, Rupesh K Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *IEEE transactions on neural networks and learning systems*, 28(10):2222–2232, 2016.

[Ketkar, 2017] Nikhil Ketkar. Introduction to pytorch. In *Deep learning with python*, pages 195–208. Springer, 2017.

[Laugstøl, 2018] Arne Laugstøl. Ventilation measurement devices, methods and computer program product, March 22 2018. US Patent App. 15/560,073.

[McDonald *et al.*, 2017] Suzanne McDonald, Francis Quinn, Rute Vieira, Nicola O'Brien, Martin White, Derek W Johnston, and Falko F Sniehotta. The state of the art and future opportunities for using longitudinal n-of-1 methods in health behaviour research: a systematic literature overview. *Health Psychology Review*, 11(4):307–323, 2017.

[Nikodem and Bawiec, 2020] Maciej Nikodem and Marek Bawiec. Experimental evaluation of advertisement-based bluetooth low energy communication. *Sensors*, 20(1):107, 2020.

[Ozemek *et al.*, 2019] Cemal Ozemek, Carl J Lavie, and Øivind Rognmo. Global physical activity levels-need for intervention. *Progress in cardiovascular diseases*, 2019.

[Richardson *et al.*, 2013] Leonard Richardson, Mike Amundsen, Michael Amundsen, and Sam Ruby. *RESTful Web APIs: Services for a Changing World.* " O'Reilly Media, Inc.", 2013.

[Schneider *et al.*, 2006] Patrick L Schneider, David R Bassett Jr, Dixie L Thompson, Nicolaas P Pronk, and Kenneth M Bielak. Effects of a 10,000 steps per day goal in overweight adults. *American Journal of Health Promotion*, 21(2):85–89, 2006.

[Thomson *et al.*, 2019] Elizabeth A Thomson, Kayla Nuss, Ashley Comstock, Steven Reinwald, Sophie Blake, Richard E Pimentel, Brian L Tracy, and Kaigang Li. Heart rate measures from the apple watch, fitbit charge hr 2, and electrocardiogram across different exercise intensities. *Journal of sports sciences*, 37(12):1411–1419, 2019.

[Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

[Veerabhadrappa *et al.*, 2018] Praveen Veerabhadrappa, Matthew Duffy Moran, Mitchell D Renninger, Matthew B Rhudy, Scott B Dreisbach, and Kristin M Gift. Tracking steps on apple watch at different walking speeds. *Journal of general internal medicine*, 33(6):795–796, 2018.