German Medical Data Sciences: A Learning Healthcare System U. Hübner et al. (Eds.) © 2018 The authors and IOS Press. This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0). doi:10.3233/978-1-61499-896-9-155

Modeling Activity Tracker Data Using Deep Boltzmann Machines

Martin TREPPNER^{a,1}, Stefan LENZ^a, Harald BINDER^a and Daniela ZÖLLER^a ^aInstitute of Medical Biometry and Statistics, Faculty of Medicine and Medical Center - University of Freiburg

Abstract. Commercial activity trackers are set to become an essential tool in health research, due to increasing availability in the general population. The corresponding vast amounts of mostly unlabeled data pose a challenge to statistical modeling approaches. To investigate the feasibility of deep learning approaches for unsupervised learning with such data, we examine weekly usage patterns of Fitbit activity trackers with deep Boltzmann machines (DBMs). This method is particularly suitable for modeling complex joint distributions via latent variables. We also chose this specific procedure because it is a generative approach, i.e., artificial samples can be generated to explore the learned structure. We describe how the data can be preprocessed to be compatible with binary DBMs. The results reveal two distinct usage patterns in which one group frequently uses trackers on Mondays and Tuesdays, whereas the other uses trackers during the entire week. This exemplary result shows that DBMs are feasible and can be useful for modeling activity tracker data.

Keywords. Deep Learning, Deep Boltzmann Machines, Activity Trackers

1. Introduction

Wearable devices, more specifically activity trackers, are attracting considerable interest due to their ability to measure physical activity continuously. The identification of patterns in the corresponding vast amounts of data generated in these settings poses a challenge to conventional linear modeling approaches.

Few researchers have addressed the problem of modeling the joint distribution of large quantities of activity tracker data using deep learning techniques. This paper seeks to address this issue by employing a deep Boltzmann machine (DBM) which has been shown to be a promising method in many applications like single-cell genomics, object recognition, and information retrieval [1-3]. Additionally, it provides a generative model, i.e., artificial samples can be generated from a trained model for exploring the structure learned by the approach.

Our analysis aims to develop methods that can detect structures in large unlabeled datasets and broaden current knowledge of weekly activity tracker usage patterns. Furthermore, the results could subsequently provide information on associations of latent activity patterns with health outcomes.

¹ Corresponding Author, Martin Treppner, Institute of Medical Biometry and Statistics (IMBI), Stefan-Meier-Str. 26, 79104 Freiburg, Germany; E-mail: treppner@imbi.uni-freiburg.de.

Due to the complex structure of activity tracker data, the modeling task is often subdivided. Bai et al., for example, employed a two-stage model in which the presence of activity and the activity intensity given any activity at all are modeled separately using linear models [4]. In [5], the authors utilized random forests to model physical activity but were only able to train their model on data from a strictly controlled setting, which is not directly comparable to activity data from the use of activity trackers by the general population. In contrast to the approach considered in our study, [6] use supervised learning methods in a controlled surrounding, i.e., a gold standard is needed in the learning procedure.

2. Methods

The data for this study was obtained from *openhumans.org* where users of activity trackers can donate their activity records for the purpose of scientific studies. We had access to publicly shared Fitbit data from 29 individuals and extracted their daily step counts (median number of steps = 7181). Since our goal was to explore weekly usage patterns data was prepared such that each week represents one sample. Accordingly, step counts were available for 3577 weeks (median number of steps in a given week i at day j, j = 1 corresponding to Monday, and so on.

After deleting all weeks in which no data was recorded, we dichotomized the step counts for activity/inactivity of the tracker. This approach was adapted from [4] using an indicator function:

$$\mathbf{1}_{A}(\mathbf{x}_{ij}) \coloneqq \begin{cases} 1, & \mathbf{x}_{ij} > 0\\ 0, & \text{else} \end{cases}$$
(1)

The following subsections give a brief overview of deep Boltzmann machines.

2.1. Deep Boltzmann Machines (DBMs)

In order to identify the usage patterns mentioned above, we employ deep Boltzmann machines. This method has the potential to outperform previous approaches for wearable device data because it can learn a hierarchy of abstract features due to its deep architecture [7]. In addition, DBMs can process vast quantities of unlabeled data which is inevitable in settings where data is obtained from fitness trackers under real-life conditions.

In our approach, we consider a two-layer Boltzmann machine where we denote the visible layer as **v** and use $\mathbf{h}^1, \mathbf{h}^2$ for the first and second hidden layer, respectively. Furthermore, we restrict ourselves to a DBM with no within layer connections. The DBM enables modeling of the joint distribution of a large number of Bernoulli variables. In this context, these variables represent whether an individual made use of an activity tracker or not at a specific day. Broadly speaking, DBMs consist of stacked sets of visible and hidden nodes in which each layer captures complex, higher-order correlations between the activities of hidden features in the layer below [2].

Following the definitions and outline of [2] we define the energy of the state $\{v,h^1,h^2\}$ as:

$$\mathbf{E}(\mathbf{v}, \mathbf{h}^1, \mathbf{h}^2; \mathbf{\theta}) = -\mathbf{v}^{\mathrm{T}} \mathbf{W}^1 \mathbf{h}^1 - \mathbf{h}^{1\mathrm{T}} \mathbf{W}^2 \mathbf{h}^2$$
(2)

where $\theta = {\mathbf{W}^1, \mathbf{W}^2}$ are the model parameters, representing the symmetric interactions between layers. In energy-based models, low energy corresponds to high probabilities whereas high energy represents a low probability.

Next, we define the probability of the visible vector **v**:

$$p(\mathbf{v}; \theta) = \frac{1}{Z(\theta)} \sum_{\mathbf{h}^1, \mathbf{h}^2} \exp\left(-E(\mathbf{v}, \mathbf{h}^1, \mathbf{h}^2; \theta)\right)$$
(3)

Where $Z(\theta)$ denotes the intractable normalizing constant, i.e., the partition function. Furthermore, the conditional distributions over the visible and the two sets of hidden units are given by logistic functions σ :

$$\mathbf{p}(\mathbf{h}_{j}^{1} = 1 | \mathbf{v}, \mathbf{h}^{2}) = \sigma(\sum_{i} W_{ij}^{1} \mathbf{v}_{i} + \sum_{m} W_{jm}^{2} \mathbf{h}_{j}^{1})$$

$$\tag{4}$$

$$p(\mathbf{h}_{m}^{2} = 1 | \mathbf{h}^{1}) = \sigma\left(\sum_{j} W_{im}^{2} \mathbf{h}_{i}^{1}\right)$$
(5)

$$\mathbf{p}(\mathbf{v}_{i} = 1 | \mathbf{h}^{1}) = \sigma\left(\sum_{j} W_{ij}^{1} \mathbf{h}_{j}\right)$$
(6)

2.2. Training Procedure

To carry out stochastic gradient ascent on the log-likelihood we make use of the following parameter update rule for the simplified case with one hidden layer **h**:

$$\Delta \mathbf{W} = \nu \left(\mathbb{E}_{\mathsf{P}_{\mathsf{data}}}[\mathbf{v}\mathbf{h}^{\mathsf{T}}] - \mathbb{E}_{\mathsf{P}_{\mathsf{model}}}[\mathbf{v}\mathbf{h}^{\mathsf{T}}] \right)$$
(7)

Here, $\mathbb{E}_{P_{data}}[\cdot]$ is referred to as the data-dependent expectation while we denote $\mathbb{E}_{P_{model}}[\cdot]$ as the data-independent expectation. Also, the learning rate ν determines the influence each individual training sample has on the updates of \mathbf{W} . To train the model's expectations we used stochastic approximation procedures which are outlined in [2]. The data-dependent expectation was approximated using variational learning, where we can characterize the true posterior distribution by a fully factorized distribution [8]. Besides, the data-independent expectation was approximated using Gibbs sampling [8]. By stacking multiple restricted Boltzmann machines (RBMs), where only two layers are considered simultaneously, the resulting DBM can learn internal representations which enable us to identify complex statistical structures within the hidden layers [8]. To this end, we adopted the greedy layerwise pre-training which is detailed in [2]. In this framework [2] introduce modifications to the first and the last RBM of the stack so that the parameters $\theta = \{\mathbf{W}^1, \mathbf{W}^2\}$ are initialized to reasonable values. On this basis, the parameters can be improved during the approximate likelihood learning of the entire DBM [7]. In our analyses, we set the number of visible nodes \mathbf{v} and the number of nodes in the first hidden layer \mathbf{h}^1 to seven in order to represent each weekday. We use one node in the terminal hidden layer \mathbf{h}^2 since we aimed to detect two groups of usage patterns. Hence, an active node in the terminal hidden layer represents one group while an inactive node represents another pattern. Furthermore, we set the learning rate ν to 0.007 during

pre-training and increased it to 0.008 for the training of the entire DBM. The number of epochs was held constant at 40 for pre-training as well as for training the entire DBM. Data analysis was performed using Julia Version 0.6.2 and the Julia package BoltzmannMachines.jl (*https://github.com/binderh/BoltzmannMachines.jl.git*).

3. Results

After having obtained estimates of the parameters θ , we used the DBM to generate new observations for the visible layer to explore the learned structure. Specifically, we used the DBM to compute the deterministic potential for the activation of the hidden nodes \mathbf{h}^1 given that the nodes in the terminal layer \mathbf{h}^2 were active/inactive. We then propagated the deterministic potential through the network to obtain the visible potential. Subsequently, we generated 10,000 uniformly distributed random numbers between 0 and 1 and assigned the value 1 if the visible potential was higher than the random number. Next, we used the generated data to plot the learned patterns in a heat map displayed in Figure 1. From the graph, we can distinguish two clear usage patterns. The upper pattern denoted as "on" shows that there is the tendency to use activity trackers at the beginning of the week on Monday and Tuesday and slightly increased usage on the weekend, while the lower pattern indicates a high usage throughout the whole week. A robustness check of the results was done by dropping each case individually and calculating the visible potential again, resulting in similar patterns.





Figure 1. Heatmap of weekly activity tracker usage generated by the DBM. The upper pattern "on" shows a group that frequently uses activity trackers on Mondays and Tuesdays. The lower row indicates a group with regular usage throughout the whole week.

4. Discussion

One limitation is our sample size of 29 individuals. Being able to acquire more data and validate the model is an essential next step. Besides, integrating other measurements like heart rate and sleeping behavior could improve our model substantially. Based on this, we could adapt the structure of our DBM to find even more usage patterns or model activity intensity. Another downside regarding our methodology is the dichotomization of step counts which gives away valuable information in the data. This can be addressed by using different cutoffs to incorporate the number of steps per day, but other, direct modeling approaches may be more useful.

Nevertheless, we believe our work could be the basis for future studies which take the continuous measurements from activity trackers into account. To further our research, we plan to use the partitioning approach presented in [7]. Doing this, we could potentially combine a binary RBM, to model the presence/absence of tracker usage, with a Gaussian RBM, for learning the activity intensity patterns and for modeling the temporal structure over a longer time period.

5. Conclusion

We have presented a deep learning approach, more precisely a deep Boltzmann machine, to model the complex joint distribution of activity tracker data and showed that it can be used successfully to extract meaningful activity tracker usage patterns. Most importantly, we were able to reveal two distinct weekly usage patterns in which one group mostly uses trackers on Mondays and Tuesdays, whereas the other uses trackers during the entire week. Understanding these usage patterns, which might e.g. correspond to individuals with planned activity vs. individuals with continuous self-optimization, could be useful for subsequent analyses.

6. Conflict of Interest

The authors state that they have no conflict of interests.

References

- C. Angermueller, H.J. Lee, W. Reik et al., DeepCpG: Accurate prediction of single-cell DNA methylation states using deep learning, *Genome Biology*, 8(1) (2017), 67.
- [2] R. Salakhutdinov, G. Hinton, Deep Boltzmann Machines, Proceedings of the Twelth International Conference on Artificial Intelligence and Statistics, 5 (2009), 448-455.
- [3] N. Srivastava, R. Salakhutdinov, G. Hinton, Modeling documents with a Deep Boltzmann Machine, Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence, (2013), 616-624.
- [4] J. Bai, Y. Sun, J.A. Schrack et al., A two-stage model for wearable device data, *Biometrics*, 2017.
- [5] K. Ellis, J. Kerr, S. Godbole et al., A random forest classifier for the prediction of energy expenditure and type of physical activity from wrist and hip accelerometers, *Physiological measurement*, **35(11)** (2014), 2191–2203.
- [6] J. Staudenmayer, D. Pober, S. Crouter et al., An artificial neural network to estimate physical activity energy expenditure and identify physical activity type from an accelerometer, *Journal of applied physiology* (1985), 107(4) (2009), 1300–1307.
- [7] M. Hess, S. Lenz, T.J. Blätte et al., Partitioned learning of deep Boltzmann machines for SNP data, *Bioinformatics (Oxford, England)*, 33(20) (2017), 3173–3180.
- [8] R. Salakhutdinov, G. Hinton, An efficient learning procedure for deep Boltzmann machines, *Neural computation*, 24(8) (2012), 1967–2006.