Monitoring Daily Activities of Multiple Sclerosis Patients with Connected Health Devices

Sourav Bhattacharya
Nokia Bell Labs, Cambridge, UK
sourav.bhattacharya@nokia-bell-labs.com

Alberto Gil C. P. Ramos
Nokia Bell Labs, Cambridge, UK
alberto.gil_ramos@nokia-bell-labs.com

Fahim Kawsar
Nokia Bell Labs, Cambridge, UK
TU Delft, Netherlands
fahim.kawsar@nokia-bell-labs.com

Nicholas D. Lane
Nokia Bell Labs, Cambridge, UK
University of Oxford, UK
nic.lane@nokia-bell-labs.com

Lynn M. Gionta
Teva Branded Pharmaceutical Products R&D, Inc.
Lynn.Gionta@tevapharm.com

Joanne Manidis
Teva Branded Pharmaceutical Products R&D, Inc.

Greg Silvesti
Teva Branded Pharmaceutical Products R&D, Inc.

Mathieu Vegreville
Nokia, Paris, France
matthieu.vegreville@nokia.com

UbiComp/ISWC'18 Adjunct,, October 8–12, 2018, Singapore, Singapore
ACM 978-1-4503-5966-5/18/10.
https://doi.org/10.1145/3267305.3267682

Abstract
We report results from a pilot study that focuses mainly on understanding the everyday life quality of patients suffering from multiple sclerosis through the lens of connected Nokia Health devices. Our dataset comprises of 198 individuals (184 females and 14 males) and the study lasted over six months. By analyzing carefully crafted user-studies and correlating with personal sensor data collected with Nokia devices, we found that the level of fatigue is one of the main sources of discomfort across the majority of the patients. We further perform an exploratory analysis, which provides an early indication that by actively monitoring and perturbing users' daily activity levels, such as increasing daily step-counts, sleep duration and decreasing body weight, patients can potentially reduce their daily fatigue level.

Author Keywords
Health, Connected Devices, Statistical Analysis, Machine Learning

ACM Classification Keywords
H.4.m [Information Systems Applications]: Miscellaneous; I.2.6 [Artificial Intelligence]: Learning

Introduction
Multiple Sclerosis (MS) is a disease of the nervous system, where the protecting layer of the nerve cells in the brain

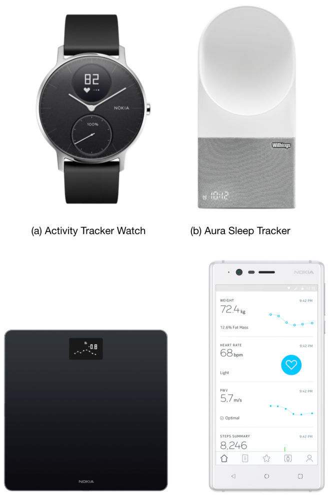

(a) Activity Tracker Watch (b) Aura Sleep Tracker

(c) Scale (d) Health Mate App

Figure 1: Nokia devices used: (a) Activity Tracker Watch collects number of steps, walking speed and distance, number of calories, sleep duration, light and deep sleep, (b) Aura Sleep Tracker records sleep duration, rapid eye movement, light and deep sleep, time awake at night, time to fall asleep, number of times awaken, time to be awake, bed in time, bed out time, nightly heart rate and respiratory rate, (c) Scale measures weight, body fat and water percentage, muscle mass and bone mass, (d) Health Mate App provides users' access to data and periodic questionnaires.

or in the spinal cord are damaged, resulting in a varying degree of physical, mental and psychological problems in patients. MS is a lifelong condition and can potentially be very serious, causing severe disability. The disease is found to be two to three times more common in women than men and currently over hundred thousand people are estimated to be affected by it in the UK alone. Patients suffering from severe MS often require long term health monitoring and support in their daily lives. Constant medical monitoring adds a significant cost to the healthcare system. Health clinic monitoring also adds difficulties in patient's personal life by requiring frequent visits to the health centre. One potential way to better understand patients' symptoms is to use wearable devices to monitor aspects of patient's life, such as quality of sleep, amount of daily locomotion, and variation of Body Mass Index (BMI) over time. Patient's history would provide additional insights to the medical professionals to tailor a better treatment to the individual patient. Towards this goal, we perform a pilot study of a number of patients suffering from various degrees of MS. Specifically, we monitor patients' everyday lives with Nokia devices and periodically run carefully crafted user studies to understand the sources of discomfort in their lives. In addition to performing basic correlation analysis between the survey dataset and the sensor data collected using the connected devices, we tried to understand how subtle everyday behavior changes could result in a better life quality, e.g., decreasing the fatigue level of a patient. In our analysis, we first perform a statistical analysis to understand dependencies across sensor measurements, then we model patients' fatigue levels using Deep Neural Networks (DNNs), and finally we study how to recommend statistically meaningful actionable feedback to improve patient wellbeing through input exploration on the learned DNN.

Related Work

Personal wearables have augmented clinical studies by allowing real-time and unobtrusive sensing of medical attributes. Patients are becoming comfortable with personal wearables, allowing not only the measurement of physical attributes, but also capturing individual and social activities [1]. DNNs are becoming popular within health analytics research and have been successfully used in predicting the chance of achieving personal weight objectives [2], and aggregate user statistics, e.g., ranking across population, in a privacy sensitive manner [3]. Wearable sensors have also been used to remotely assess the level of mobility impairment, common in patients suffering from MS, by continuously monitoring their gait outside health clinics [4], and to enable the study of the effects of stress on MS patients [5], enabling the creation of remote tools to decrease patients' stress levels to improve their quality of life.

Data Collection

Our pilot study comprises of 198 patients (184 female and 14 male) and spans a period of six months. At the beginning of the study we provided all patients with the Nokia Health devices to monitor their health attributes (c.f., Fig. 1's caption for a comprehensive list) via the Nokia Activity Tracker Watch, Aura Sleep Tracker and Scale, as well as to provide them with a set of daily, weekly and monthly questionnaires, with filling time of 10sec, 2min and 1min, regarding their wellbeing through the Nokia mobile app. The questionnaires were used to track a number of life qualities including fatigue level, ability to sleep well, walking difficulties, capacity to perform self-care, ability to perform usual activities, levels of pain or discomfort, anxiety or depression, general quality of life and treatment adherence. These short questionaries, filled in by each patient, provide qualitative subjective information about their perceived well-being. For each user we also collected a number of static informa-

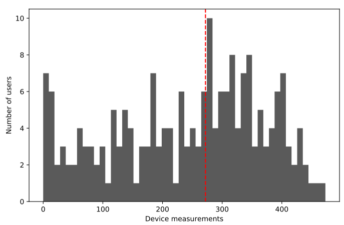

Figure 2: Histogram of users contributing data to the pilot study via the Nokia Activity Tracker Watch, Aura Sleep Tracker and Scale, illustrating a condensed view of the variability of device usage across users, plotting the number of users versus the number of device measurements (specifically, each one device aggregates its daily measurements into one record per day which we henceforth refer to one device measurement), where the red line in the figure denotes the median of the users across the number of device measurements, which thereby separates high frequency users from low frequency ones. In particular, it can be seen that the majority of the users used the Nokia devices frequently throughout the study.

tion such as gender, age category, the type of MS, the level of mobility impairment and comorbidities across patients, which serve to augment the feature space of the dataset. Timely responses to the questionnaires and aggregated sensor measurements over each day, form the core of our study dataset. As can be seen from Fig. 2, most users use their health devices regularly over the period of six months, contributing hundreds of daily measurements, which translate to a good daily coverage across users providing a rich raw dataset to perform our statistical analysis.

Analysis and Results

We now briefly present the key findings from our pilot study and then provide details of our analysis and evaluation results. Our main findings can be summarised as follows:

- We focussed on understanding the impact of daily activities on the fatigue level of patients suffering from MS, as it is one of their main sources of discomfort.
- The level of fatigue shows a significantly positive correlation with body weight and BMI, and negative correlation with daily step counts and walking speed.
- Via DNNs we show how to perturb users recent daily activity patterns, to recommend subtle lifestyle changes, in order to potentially reduce their fatigue levels.

We aggregated the raw data into input and target pairs $\{(\boldsymbol{x}^{(i)}, y^{(i)})\}_{i=1}^{n}$, $\boldsymbol{x}^{(i)} \in \mathbb{R}^l$, $y^{(i)} \in \{0, 1, \ldots, m-1\}$ for supervised classification learning and stratified (based on user anonymised hash identification, gender and age category features as well as the target fatigue scores) into training (about 90%) and test (roughly 10%) sets. Given the small number of male users, namely 14, and that the test set contains roughly 10% of the total data, the generalisation capabilities of our analysis should be understood to apply mainly to female patients. Using this dataset we focus on understanding correlations between everyday activities of patients with their reported fatigue levels, and on exploring the predictive power of DNNs to understand potential effects of actionable health recommendations.

Correlation Analysis

The correlation analysis on the training set aims to identify pairs of variables that show positive or negative coevolution. We use Pearson coefficients with p-values for the analysis. As the target variable we use the level of fatigue reported by the patients. We found that fatigue level correlates positively with the body weight ($\rho = 0.24$, $p \ll 0.01$), positively with the BMI ($\rho = 0.2$, $p \ll 0.01$), negatively with the steps ($\rho = -0.14$, $p \ll 0.01$) and negatively with the walking speed ($\rho = -0.12$, $p \ll 0.01$). Further study is required to understand if causal relation exists or not.

Predictive Modeling

For predictive inference we model the relationship between the input features and the target scores with a DNN having a softmax non-linearity at the output layer $f_{\mathrm{DNN}}(\boldsymbol{x}; \boldsymbol{\theta}) \in \mathbb{R}^m$, $\boldsymbol{\theta} \in \mathbb{R}^k$ via $y_{\mathrm{pred}}(\boldsymbol{x}; \boldsymbol{\theta}) := \mathrm{argmax}_j\{[f_{\mathrm{DNN}}(\boldsymbol{x}; \boldsymbol{\theta})]_j\}$, and minimise the cross-entropy between a one-hot encoding of $y^{(i)}$ onto $\mathbb{R}^m$ and $f_{\mathrm{DNN}}(\boldsymbol{x}^{(i)}; \boldsymbol{\theta})$, via mini-batch gradient descent across the samples in the training folds generated in cross-validation (stratified via user anonymised hash identification). After the optimal model is obtained via cross-validation, resulting from $\boldsymbol{\theta}$ initialisation and hyperparameters optimization, then a new model instance is trained on the entire training set. Once this last training is concluded, the final $\boldsymbol{\theta}^*$ is stored. After $\boldsymbol{\theta}^*$ is selected, the test set is then used to assess the predictive quality of the final model. Having trained $f_{\mathrm{DNN}}(\boldsymbol{x}; \boldsymbol{\theta}^*)$ to predict the overall fatigue level $y_{\mathrm{pred}}(\boldsymbol{x}; \boldsymbol{\theta^*})$ of a user given activity data $\boldsymbol{x}$ collected using his/her Nokia devices, the trained DNN allows us to make perturbations in the input space and measure the overall fatigue prediction level that the user

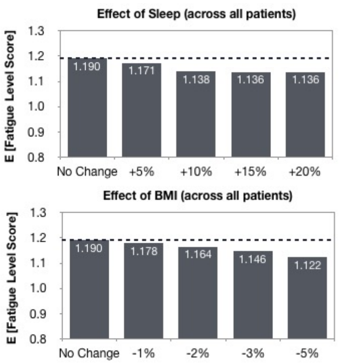

Figure 3: Illustration of the variation of the expected value of the predicted fatigue (1) (y-axis) across different local optimum input perturbations $\boldsymbol{\delta}^*_{\mathcal{T}_1}$ (x-axis) along different one-dimensional subspaces, namely along the sleep feature (up) and along the BMI feature (down). The results translate into insight for potential life-style recommendation effects of perturbations on the expected fatigue level across patients from $\mathcal{T}_2$.

would have if he/she would have had slightly different activities. This minor variation in routine would then be quantified and returned to the user as a form of actionable feedback. Note that we partition the test set $\mathcal{T}$ into two disjoint halves $\mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2$ (stratified) for proper assessment of the generalisation capabilities of the input perturbation analysis. Namely, we explore various perturbations in only the first half $\mathcal{T}_1$ and (having chosen the final input perturbations) assess the impact of the perturbation only on the second half $\mathcal{T}_2$. This methodology lends itself to consider perturbations $\boldsymbol{\delta} \in \mathbb{R}^l$, and study those perturbation which achieve a decrease in the expected value of the predicted fatigue, namely $\mathbb{E}_{\boldsymbol{x} \in \mathcal{T}_1}\{y_{\text{pred}}(\boldsymbol{x} + \boldsymbol{\delta}; \boldsymbol{\theta}^*)\} \leq \mathbb{E}_{\boldsymbol{x} \in \mathcal{T}_1}\{y_{\text{pred}}(\boldsymbol{x}; \boldsymbol{\theta}^*)\}$. Once different final perturbations are chosen, we check how those small perturbations, which correspond to minor changes in the users lives affect their fatigue levels. Note also that $\boldsymbol{\delta} \mapsto \mathbb{E}_{\boldsymbol{x} \in \mathcal{T}_1}\{y_{\text{pred}}(\boldsymbol{x} + \boldsymbol{\delta}; \boldsymbol{\theta}^*)\}$ may have very many local minima, and that the minimisation of this objective function becomes increasingly difficult as the number of features becomes larger, i.e., $l \gg 1$. For this reason, we manually select only a few relevant features, and consider only the relatively small subspaces of $\mathbb{R}^l$ spanned by those features. Note that since the correlation analysis done was based only on the training set, we can, without any bias risk, use the correlation results to have a heuristic sense of what could be good choices for the selected features, which limits the search space and makes the minimisation efficient. Once the minimisation is complete based on the test subset $\mathcal{T}_1$, locally optimal perturbations $\boldsymbol{\delta}^*_{\mathcal{T}_1}$ are kept. For generalisation assessment, we consider the separate test subset $\mathcal{T}_2$, and evaluate the expected fatigue level

$$\mathbb{E}_{\boldsymbol{x} \in \mathcal{T}_2}\{y_{\text{pred}}(\boldsymbol{x} + \boldsymbol{\delta}^*_{\mathcal{T}_1}; \boldsymbol{\theta}^*)\}. \tag{1}$$

Examples of such changes in activities and their effects on the expected fatigue levels are shown in Fig. 3, for two basic forms of perturbations, in which $\boldsymbol{\delta}^*_{\mathcal{T}_1}$ is non-zero only for the feature which corresponds to the total sleep time series, and for the feature that corresponds to the BMI time series. Interestingly, subtle changes in the patient's activities along these features indicate a drop in the overall fatigue level.

Conclusion

We conducted a pilot study to gain insights into the effects of everyday life choices on the fatigue levels of patients suffering from various degrees of MS, by monitoring their daily activities through connected Nokia Health devices. Our analysis focussed on a correlation analysis to find out main factors significantly impacting the fatigue level across patients. Additionally, we conducted a predictive modeling analysis using DNNs to show potential directions of providing actionable feedback in the form of subtle changes in lifestyle to potentially decrease the level of fatigue.

REFERENCES

1. Ziemssen, Kern & Thomas. Multiple sclerosis: clinical profiling and data collection as prerequisite for personalized medicine approach. BMC Neurol, 2016.
2. Velickovic, Karazija, Lane, Bhattacharya, Liberis, Lio, Chieh, Bellahsen & Vegreville. Cross-modal Recurrent Models for Weight Objective Prediction from Multimodal Time-series Data. arXiv:1709.08073, 2017.
3. Velickovic, Lane, Bhattacharya, Chieh, Bellahsen & Vegreville. Scaling Health Analytics to Millions without Compromising Privacy using Deep Distributed Behavior Models. EAI PervasiveHealth, 2017.
4. Moon, McGinnis, Seagers, Motl, Sheth, Wright, Ghaffari & Sosnoff. Monitoring gait in Multiple Sclerosis with novel wearable motion sensors. PLOS ONE, 2017.
5. Martinez & Picard, Wearable Technologies for Multiple Sclerosis. International Conference on SPWID, 2016.