Unsupervised Learning for Mental Stress Detection
Exploration of Self-organizing Maps

Dorien Huysmans[1,3], Elena Smets[2,3], Walter De Raedt[3], Chris Van Hoof[2,3,4], Katleen Bogaerts[5,6], Ilse Van Diest[6] and Denis Helic[7]

[1]*KU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Leuven, Belgium*

[2]*KU Leuven, Department of Electrical Engineering (ESAT), Leuven, Belgium*

[3]*imec, Leuven, Belgium*

[4]*imec, Holst Centre, Eindhoven, The Netherlands*

[5]*REVAL - Rehabilitation Research Center, Faculty of Medicine and Life Sciences, Hasselt University, Diepenbeek, Belgium*

[6]*Research Group on Health Psychology, Department of Psychology, KU Leuven, Leuven, Belgium*

[7]*Knowledge Technologies Institute, Graz University of Technology, Graz, Austria*

Keywords: Mental Stress Detection, Skin Conductance, Electrocardiogram, Unsupervised Learning, SOM.

Abstract: One of the major challenges in the field of ambulant stress detection lies in the model validation. Commonly, different types of questionnaires are used to record perceived stress levels. These only capture stress levels at discrete moments in time and are prone to subjective inaccuracies. Although, many studies have already reported such issues, a solution for these difficulties is still lacking. This paper explores the potential of unsupervised learning with Self-Organizing Maps (SOM) for stress detection. In unsupervised learning settings, the labels from perceived stress levels are not needed anymore. First, a controlled stress experiment was conducted during which relax and stress phases were alternated. The skin conductance (SC) and electrocardiogram (ECG) of test subjects were recorded. Then, the structure of the SOM was built based on a training set of SC and ECG features. A Gaussian Mixture Model was used to cluster regions of the SOM with similar characteristics. Finally, by comparison of features values within each cluster, two clusters could be associated to either relax phases or stress phases. A classification performance of 79.0% (±5.16) was reached with a sensitivity of 75.6% (±11.2). In the future, the goal is to transfer these first initial results from a controlled laboratory setting to an ambulant environment.

1 INTRODUCTION

The concept of stress is difficult to capture because stress has both psychological as well as physiological aspects. Moreover, both of these aspects are complex and are typically caused by multiple factors. The psychological part has been described by multiple models such as the Demand-Control Model (Karasek and Theorell, 1992) and the Effort-Reward Imbalance Model (Siegrist, 2010). Physiologically, stress can be described by the activity and balance of the autonomic nervous system.

The interest in stress detection shifted from laboratory conditions to ambulatory, enabled by the growth of wearable sensor technology (Fahrenberg et al., 2007). Recently, wearables have been slowly introduced into daily-life studies of stress. This unveiled a major problem concerning validation. Validating ambulant stress detections is not a clearly defined process as there is no precise recording of what the participant's activities are. The only apparent way for validation are different types of questionnaires and diaries, filled in multiple times a day.

Kusserow et al. (2013) monitored the participants by a diary of daily activities (e.g. working, transport, conversation) and mood-state questionnaires which had to be filled in as soon as possible after perceiving stress arousal. However, most questionnaires were completed randomly and could not be related to the estimated stress-arousal phases.

Adams et al. (2014) performed an Experience Sampling Method study, enabled by their specifically

Huysmans, D., Smets, E., Raedt, W., Hoof, C., Bogaerts, K., Diest, I. and Helic, D.
Unsupervised Learning for Mental Stress Detection - Exploration of Self-organizing Maps.
DOI: 10.5220/0006541100260035
In *Proceedings of the 11th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2018) - Volume 4: BIOSIGNALS*, pages 26-35
ISBN: 978-989-758-279-0

designed smartphone app *SESAME*. Participants received approximately every half hour a notification to fill in the self-report and were free to fill in additional self-reports. In practice, many experience-sampling responses were delayed due to practical reasons of the application or occupation of the participants, or participants did not respond at all to notifications. Moreover, periods of time associated with very high levels of stress were under-reported.

Hovsepian et al. (2015) made an attempt to provide a gold standard for continuous stress measurements from wearable sensors and presented a stress model *cstress*. They prompted participants at random 15 times a day to fill in an Ecological Momentary Assessment (EMA). This EMA self-report served as the ground truth for field validation. The *cstress* model compensated for the arbitrary lag between the occurrence of a stressor and its self-report logging.

Validation of the ambulant data requires a different approach compared to lab studies. Labelling of the physiological data is often nonexisting or inaccurate . Although many studies have already reported these issues regarding stress level labelling in an ambulant environment (Adams et al., 2014; Hovsepian et al., 2015), these problems are rarely addressed in the analyses. This encourages the exploration of unsupervised stress detection algorithms.

Medina (2009) identified stress states from ECG signals using several unsupervised learning methods. These are clustering algorithms (including K-means and Spectral Clustering) and clustering ensemble methods as well as dimensionality reduction techniques (Principal Component Analysis and Forward Sequential Search) and evolutionary algorithms.

The study by Grigore and Bornoiu (2014) investigated stress detection by using electrodermal features. Their evaluation method relied on observations of the SC signal by an expert observer, combined with questions to the test subjects about their state, to mark a recorded signal as stress or relax. An unsupervised method was preferred and they proposed to use a Kohonen Neural Network, also known as Self-Organizing Map (SOM). Training the SOM was unsupervised, though to point out regions of the SOM related to stress or relax phases, the authors compared the neural activation patterns with the signal from the expert observer. As such, quite elaborate expert information was essential for marking of the SOM.

The application of unsupervised techniques are relatively unexplored within the field of mental stress detection. The SOM in particular has been successfully applied in many other fields such as brain computer interfaces (Han and Kim, 2016) and geophysics (Aguado et al., 2008; de Matos et al., 2007; Bauer et al., 2012). In the current research an algorithmic pipeline is explored based on the unsupervised learning algorithm SOM. The purpose is to rule out the use of labels. The pipeline's input are heart rate variability (HRV) and skin conductance response features. These are derived from lab recorded ECG and SC signals during a series of stress-inducing tasks. The algorithmic pipeline only relies on the recorded physiological signals and no expert observations are required for marking stress and relax states within the SOM. The findings will enhance our understanding of the link between physiological signals and stressors and may advise further strategies for stress detection in ambulatory settings.

2 MATERIALS AND METHODS

2.1 Experimental Setup and Data Set

The goal is to define a psychophysiological stress profile of test persons. During laboratory experiments, participants have to complete three different stress-inducing tasks. During these tasks, the participants were monitored with the NeXus-10KMII (MindMedia, Herten, The Netherlands) and Health Patch (imec, Leuven, Belgium) to measure skin conductance (SC) and the electrocardiogram (ECG).

2.1.1 Test Subjects

The data set consisted of a group of 12 test subjects (age 37.3 $\pm$8.8). Within this group, there were five male participants and seven female participants, recruited at Tumi Therapeutics, a multidisciplinary ambulatory treatment center specialized in the treatment of stress-related symptoms and syndromes. They all reported stress-related complaints, suffering from chronic stress, but were not diagnosed with any clinical disorder (e.g. depression or burnout). This assessment was performed by a therapist. One test subject was not included for further analysis as the data recorded by Health Patch was too noisy.

2.1.2 Experimental Protocol

The laboratory experiment lasted for 14 minutes and was set up as seen in Figure 1. The subjects had to complete three stress tasks of two minutes. All three tests are commonly used to induce stress in laboratory settings (Liao and Carey, 2015). The tasks were each separated by a two minutes resting phase. Before the

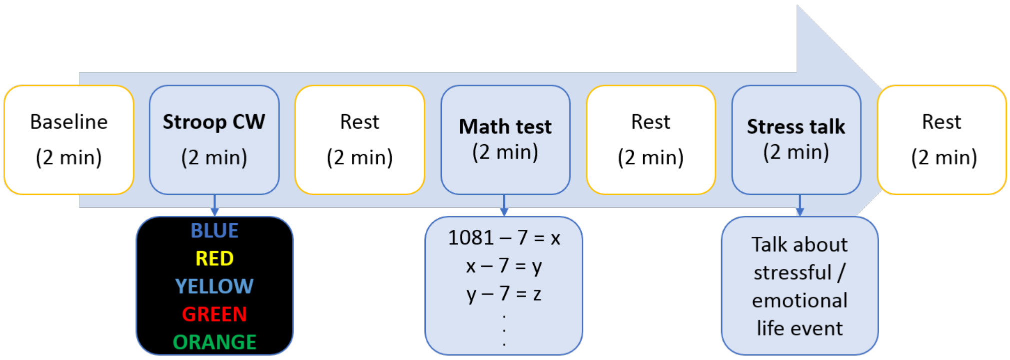

Figure 1: Experimental protocol.

first task a two minutes baseline was recorded. The first stress task was the *Stroop Color Word Test* (Van der Elst et al., 2006), words of colours were written in a different colour as the colour the word represent, e.g. the word *blue* is printed in red ink. The test subject had to say the colour of the ink as correct and as fast as possible. The challenge is to suppress the instinctive response of saying the colour the word represents. The correct answer is *red* in this example. This is a commonly used stress task. Additional stress could be added when the test supervisor urges the test subject to be faster or to say *wrong* when a mistake has been made.

The second test was a calculation test in which the participant continuously had to subtract the number 7 from the number 1081. In the same manner as the Stroop test, additional stress could be added by the test supervisor.

The final stress task was a stress talk, the participant had to talk about a very stressful or emotionally negative event in his life and to recall his feelings related to this event. The test supervisor could ask questions such as *How did you feel?*.

2.1.3 Sensors and Signals

Two sensors were applied: the NeXus-10MKII and the Health Patch. The NeXus-10MKII (Mind Media BV, Herten, The Netherlands), referred to as Nexus, is not a wearable sensor, though highly accurate. Therefore this sensor could serve as a gold standard to compare measurements of other sensors. The following signals were measured: blood volume pulse (BVP) and skin conductance. BVP was sampled at 128Hz and SC at 32Hz. The Health Patch is a wearable monitoring system developed by imec. The sensor is a patch consisting of a sensor node and an electronic module to record the electrocardiogram (ECG). Sampling frequency was 256Hz.

Based on BVP of Nexus and the extracted heart rate from the Health Patch signal, both sensors were visually synchronised in time. The signals of 4 test subjects could not be synchronised, as the Health Patch data lacked quality. Therefore the data set consisted of 7 test subjects.

The BVP signal of Nexus was not further processed as more detailed HRV information can be extracted from the ECG signal.

2.2 Self-organizing Map

SOMs represent higher-dimensional data as a globally ordered two-dimensional map. The SOM can be seen as an elastic grid of nodes fitted to the input signal space, while preserving the topological relationships of the signal space (Kohonen et al., 2001) .

Here, the input signal space is an n-dimensional feature space. Every node i is associated with a weight vector $w_i = [\mu_{i1}, \mu_{i2}, \ldots, \mu_{in}]^T \in \mathbb{R}^n$. The input feature vector is $x_{stim} = [\xi_1, \xi_2, \ldots, \xi_n]^T \in \mathbb{R}^n$ (Arnrich et al., 2010). The feature vector x_{stim} is mapped to the *best-matching node* c by comparing it with all weight vectors w_i. As a metric of similarity, the smallest Euclidean distance is searched:

$$c = \arg\min_i ||x_{stim} - w_i||. \quad (1)$$

During training, topological relationships of the input feature space are projected onto the two-dimensional SOM by adapting the weight vectors w_c. Nodes that are topographically close to the *best-maching node c* are also activated to learn from the same input x_{stim}. This results in a smoothing effect on the weight vectors of nodes in the neighbourhood and eventually leads to *global ordering* of the map. Input vectors are presented to the map in a random order. Given x_{stim} at time t, the update of the weight vector w_i of node i is as follows:

$$w_i(t+1) = w_i(t) + h_{ci}(t)[x_{stim}(t) - w_i(t)]. \quad (2)$$

The initial values of the $w_i(0)$ can be arbitrary.

Table 1: Feature set.

Physiological signal	Feature	Description
SC	SC PH	signal power in a phasic SC signal
SC	SC RR	SC responses rate
SC	SC DIFF2	signal power in second difference from SC signal
SC	SC MAG	sum of the magnitudes of SC responses
SC	SC DUR	sum of the duration of SC responses
SC	Slope	slope of the regression line of the signal
SC	PH %50	percentile 50 of peak height
SC	PH %85	percentile 85 of peak height
ECG	mean HR	mean heart rate
ECG	SDNN	standard deviation of all normal RR intervals (i.e. NN intervals)
ECG	RMSSD	root-mean-square successive difference of all normal RR intervals
ECG	LF HRV	low frequency HRV (power in the 0.04-0.15 Hz band)
ECG	HF HRV	high frequency HRV (power in the 0.15-0.4 Hz band)
ECG	LFHF HRV	ratio (LF HRV) / (HF HRV)

The neighbourhood function $h_{ci}(t)$ can be defined in terms of the Gaussian function:

$$h_{ci}(t) = \alpha(t) \cdot \exp\left(-\frac{||r_c - r_i||^2}{2\sigma^2(t)}\right), \quad (3)$$

with $0 < \alpha(t) < 1$ the *learning-rate factor* and σ the width of the kernel, both decreasing monotonically in time. $r_c \in \mathbb{R}^2$ and $r_i \in \mathbb{R}^2$ are the location vectors in the SOM of nodes c and i, and with increasing $||r_c - r_i||$, $h_{ci} \to 0$.

After training, a test set can be mapped onto the SOM to determine its set of best-matching nodes.

2.3 Training of SOM

The first stage was mapping the higher-dimensional feature space onto a two-dimensional grid, while preserving the topological relationships within the data. Measurement data was mapped by the SOM algorithm to different areas on this grid. Component planes visualised the relationship between variables.

The SOM was trained using SC and ECG features. In total 14 features were derived from SC and ECG signals (Table 1) of the laboratory dataset. This is a set of frequently applied features for stress analysis found in literature. The studies of Boucsein (2012), Kappeler-Setz et al. (2010) and Wijsman et al. (2011) focus on SC analysis. Other research focuses on stress reactions using HRV: Vrijkotte et al. (2000), Hjortskov et al. (2004), Melillo et al. (2011) and Taelman et al. (2009). Additionally, several studies apply a combination of physiological signals (Zhai et al., 2005; Healey and Picard, 2005). The calculation of features was performed with a window size of 50s and a step size of 20s. These parameters were found to be optimal for performance after cross validation. Additionally, features were normalised within every test subject by Z-score standardization to account for inter-subject variation.

The topology of the SOM lattice was chosen according to Vesanto and Alhoniemi (2000) and Aguado et al. (2008). The number of nodes M was heuristically determined as $M = 5\sqrt{N}$ with N the number of samples in the data set. The aspect ratio of the lattice is the square root of the ratio of the two largest eigenvalues of the data set. As the average training data set during cross validation consisted of 267 feature vectors, the SOM topology contained 80 nodes [10×8].

Training of the SOM was performed by the SOM functions of the PyMVPA package for multivariate pattern analysis in Python (Hanke et al., 2008). The learning rate α in Eq. 3 was by default set to 0.05. The maximum number of iterations was set to 400, at which the SOM should be converged. The nodes of the SOM are settled at a location in feature space. The value of each feature at every node can be visualised as component planes. The advantage is that relationships between features can be interpreted graphically.

2.4 Analysis of Trained SOM

The second stage consisted of exploring the patterns emerged in the SOM during training. These patterns were outlined by clustering. Every node of the SOM was assigned to a cluster by the clustering method Gaussian Mixture Models. The implementation was based on the scikit-learn package of Gaussian Mixture Models (Pedregosa et al., 2011). The main interest was finding a *stress* and a *relax* cluster. Therefore, the number of components was set to two.

To determine the feasibility of the algorithmic pipeline for stress detection, it was validated by a leave-one-participant-out (LOO) cross validation scheme. This means the dataset was split in n folds, with n being the number of participants. One fold contains the data of one test person. With each iteration, the SOM was trained by *n-1* folds, leaving out the data of one test person. The resulting trained SOM and its clustering will differ slightly with every iteration.

The quality of the different clusters can be expressed in terms of cohesion and separation of the clusters. These factors are merged in the silhouette coefficient (Pedregosa et al., 2011; Rousseeuw, 1987). The advantage of this performance characteristic is that it does not rely on labels. The silhouette coefficient s of sample i is defined as:

$$s(i) = \frac{b(i) - a(i)}{max(a,b)}, \tag{4}$$

with a the mean intra-cluster distance and b the mean distance to all samples of the nearest cluster. As only two clusters are considered, b is simply the other cluster than to which i is assigned to. The range of the silhouette coefficient is between -1 and 1. If $s(i)$ approaches 1, it implies that $a(i) \ll b(i)$ and the mean intra-cluster distance is much smaller than the mean distance to samples of the other cluster. Therefore, sample i is well-clustered and assigned to the right cluster. When $s(i)$ is about zero, the sample i lies equally far from both clusters and it is not clearly defined which is the right cluster. If $s(i)$ is close to -1, $a(i) \gg b(i)$, this means that the sample is misclassified. The final silhouette score of the LOO cross validation is the average of n testing procedures.

To further analyse the clusters, the statistics of the clusters were derived. Feature values of all nodes within one cluster were gathered and represented in a boxplot.

Additionally, labels were introduced exclusively for performance calculations. Thereby, the procedure presented in this paper can be compared to other supervised methods. During the stress experiment, two possible phases are alternated, a *relax* phase and a *stress* phase. We assume that the stress tasks effectively induced stress as has been shown in earlier research (Smets et al., 2016). The references therefore, do not rely on subjectively reported stress levels. In order to classify unseen data in these two phases, the training data is split and labelled as data recorded during *relax* (label 0) or *stress* (label 1) phase. The latency of stress onset after the start of a stress-inducing task or the fading of a stress response after ending the task are not taken into account. The classification performance is the average of sensitivity and specificity. The final classification performance of the LOO cross validation is the average of n testing procedures.

3 RESULTS AND DISCUSSION

3.1 SOM Structure

Component planes present graphically the value of each feature at every node after training. For illustrative purpose, the SOM was trained by the whole data set instead of partially by $n-1$ folds. The corresponding components planes are depicted in Figure 2. Colour bars next to every component indicate the value at the node, for which red means high and blue low values.

Variables with similar colours in corresponding regions are positively correlated. These correlations are confirmed in the correlation matrix (Figure 3). It can be seen that components *SCPh, SCRR, SCmag, SCdur, PH %50, PH %85* and *mean heart rate* exhibit low values at the lower central region of the SOM. The upper left and right corners are dominated by high values. These characteristics are related to an elevated level of sympathetic activity of the autonomous nervous system (McCorry, 2007; Task Force of the European Society of Cardiology the North American Society of Pacing Electrophysiology, 1996). *SCdiff2* and *slope* show correlation as well, for which *SCdiff2* contains more extreme positive values. This is explained by the fact that *slope* is based on the first derivative of the SC signal and *SCdiff2* on the second derivative. Only the strongest variation in signal is retained. *Mean heart rate* and *RMSSD* appear strongly negatively correlated. Mathematically, these features are calculated with a similar formula. Moreover, literature confirms that an increase in heart rate and a lower vagal tone (*RMSSD*) are signs of stress (Vrijkotte et al., 2000). *LFHF* has a positive correlation with *mean HR* as seen centrally in the component plane, with reduced values and in the upper left corner for increased values. Furthermore, *HF* exhibits large regions with a value below its average, especially in the corners where other features have increased values. The study of Hjortskov et al. (2004) confirms these observations as they found a reduction in the high-frequency component of HRV and an increase in low-to-high frequency ratio during a stress situation. Additionally, a stable low-frequency component was reported, which is not clear from the component plane in this paper.

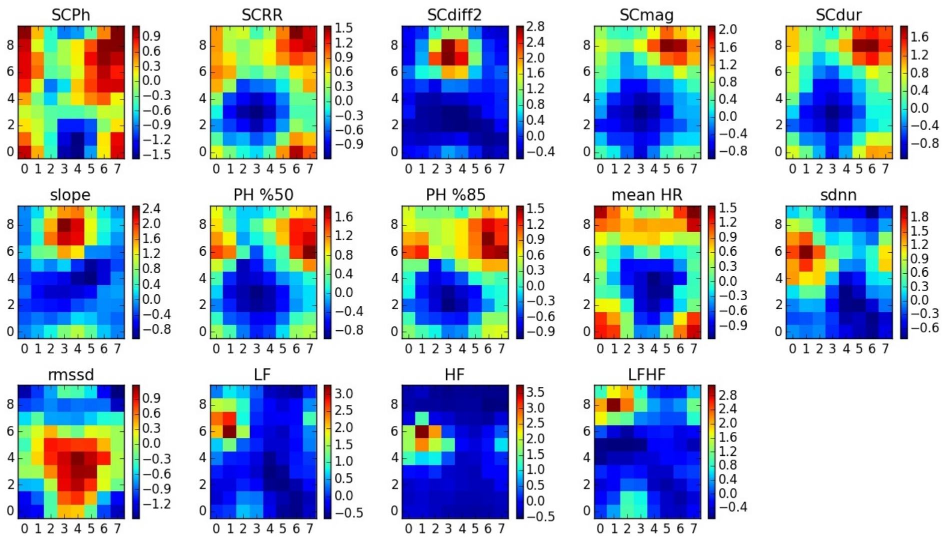

Figure 2: Component planes of trained SOM. Training was performed with whole data set for illustrative purpose.

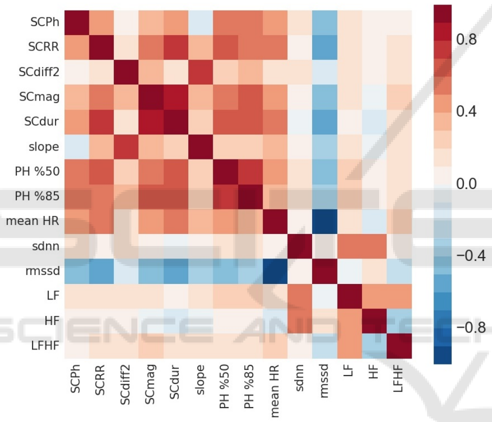

Figure 3: Correlation values between variables.

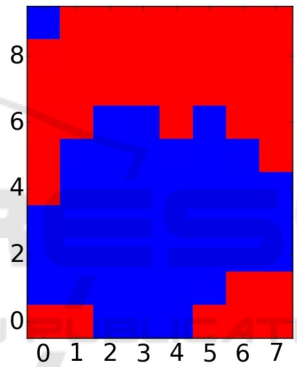

Figure 4: Clustering of trained SOM.

3.2 Clusters

The outcome of clustering of the SOM trained on the whole data set can be seen in Figure 4. A clear correlation can be seen between the clusters and the component planes of the SOM. The center of the SOM is a cluster coloured in blue, which corresponds to regions with low feature values observed in the component planes. The second cluster is the surrounding area, coloured in red. This corresponds to the components planes of *SCPh, SCRR* and *mean HR*. This visual observation is a first quality check of the resulting clusters.

The silhouette coefficient after ten runs of LOO validation was 0.301 ($\pm$0.0152). The standard deviation indicates that all silhouette coefficients were approximately equivalent. This outcome is acceptable as it is sufficiently larger than zero, meaning most samples were assigned to the right cluster.

3.3 Cluster Identification

Boxplots of both clusters over different training folds were compared. It was seen that these clusters do not capture random patterns every training fold, yet exhibit repeating patterns. This indicates that clusters have similar characteristics.

Furthermore, as the training of the SOM was unsupervised, it was not known to which state a cluster belongs, *stress* or *relax*. From literature (Boucsein, 2012; Kappeler-Setz et al., 2010; Vrijkotte et al., 2000; Hjortskov et al., 2004) it is known that during stress both *SCph* values and *mean HR* are high. This prior knowledge was applied onto the first training fold. The boxplots of leaving out test subject 1 were taken as a reference and classified. Interestingly, the pattern of all boxplot values corresponded to characteristics that can be assigned to either *stress* or *relax*. Cluster A was characterised by boxplots with negative feature averages, while these in cluster B were positive or close to zero. Therefore cluster A could be associated to *relax* and cluster B to *stress*.

Cluster boxplots of other iterations were compared to these classified boxplots. To determine corresponding boxplots, the Root Mean Squared Error (RMSE) between the average of boxplots was compared. Corresponding boxplots have a minimum RMSE between them. Subsequently, corresponding boxplots define corresponding states (*stress* or *relax*) of the clusters. The results are depicted in Figure 5, 6, 7, and 8, for SC features in cluster *relax* and *stress* and for ECG features in cluster *relax* and *stress* respectively. Iteration 1 to 4 are depicted, while omitting iteration 5 to 7 for improved readability of the figure.

It can be clearly observed that similar boxplot patterns exist between training iterations. Boxplots of a particular feature are within the same range over different training iterations. Moreover, within every training iteration (compare Figure 5 with Figure 6 and Figure 7 with Figure 8), it can be seen that the feature values are low, i.e. under the zero mean, within the *relax* cluster, and high, i.e. higher than the zero mean, within the *stress* cluster. Exceptions are *RMSSD* and *HF*, as expected from literature. The decrease of *HF* in stress situations is less explicit compared to *RMSSD*. These results were consistent with literature and observations made for the component planes.

3.4 Performance

Test data points were mapped to the SOM to determine their best-matching nodes and predict the state of *stress* or *relax*. The points were assigned to a state corresponding to the cluster the node belongs too. The clustering of the trained SOM with mapped test data and their actual labels is depicted in Figure 9. Label 1 corresponds to *stress*, label 0 to *relax*.

After ten runs of LOO validation, the average testing performance is 79.0% (±5.16). The testing performance is the average of sensitivity and specificity. The sensitivity, indicating the ability to recognize stress phases, has a value 75.6% (±11.2).

Grigore and Bornoiu (2014) applied SOM as well for stress detection, using similar SC features and reported an average recognition rate of 86.25%. Different was the their labelling system for validation of their outcomes. An expert observer evaluated the SC signal in combination with participant questionnaires to manually label the input signal. Recognition rates will be higher as labelling of the data is based on a priori evaluation of the physiological signals. Furthermore, it is not clear how their *average recognition rate* is computed. Moreover, their number of participants is not reported for comparison.

Smets et al. (2016) had a similar experimental setup and reported a maximum performance rate for non-personalized models of 82.7% using SVM. Similar features for ECG and SC were applied, with additional temperature and respiration features. This is comparable to the performance in this paper. As unsupervised techniques are generally harder to apply, the potential of SOM for stress detection exists.

4 FUTURE WORK

This paper only covers data derived from laboratory experiments, however the SOM technique has been introduced to tackle problems in modelling and validation of ambulant data. Therefore, it would be of great interest to apply this algorithmic pipeline onto ambulant data and to further determine its feasibility.

Several papers apply an intermediate step between training and clustering the SOM, namely the creation of a U-matrix followed by Kmeans clustering (Aguado et al., 2008) or a gradient function of the SOM followed by the watershed segmentation algorithm (Bauer et al., 2012). The U-matrix or gradient function displays the distance between neigbouring nodes and allows visual delineation of the clusters. The subsequent clustering or segmentation step would only be performed in a two-dimensional space. The approach with building a U-matrix has been performed, though no clear conclusions could be drawn from visual inspection of the U-matrix. Most probably the border between the *stress* and *relax* cluster cannot be clearly drawn here. One of the reasons is that physiological responses of the stress tests slowly vanish into periods defined as *relax*. Future models would benefit from taking into account response and recovery periods.

Advances can be made in more detailed evaluation of the clustering step. Other parameters or other cluster algorithms could be explored to better separate *stress* from *relax* clusters. An interesting aspect would be to add more clusters to capture stress levels directly from the SOM or determine the confidence of a predicted stress level. A first step would be to add a third cluster outlining intermediate values and focus on the extreme values of *stress* and *relax*.

A limitation of the study was the size of the data set which was rather small. For the current research, this was not a major problem as it focussed on the exploration of the algorithmic pipeline. For future validation, larger data sets are required.

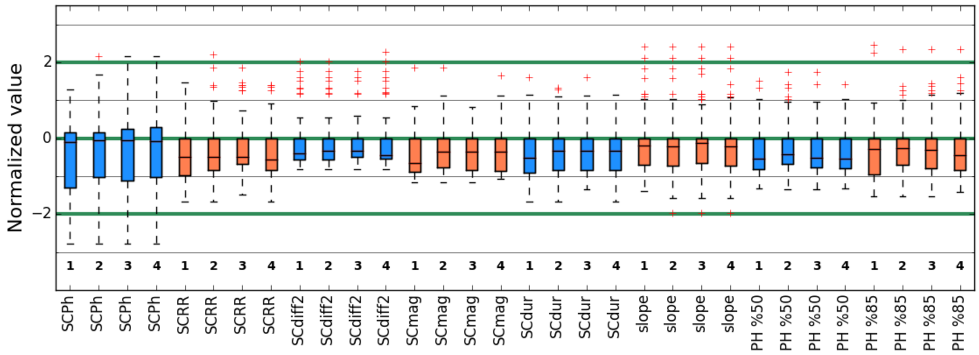

Figure 5: Boxplots of SC feature values in cluster *relax*, with number referring to the training iteration.

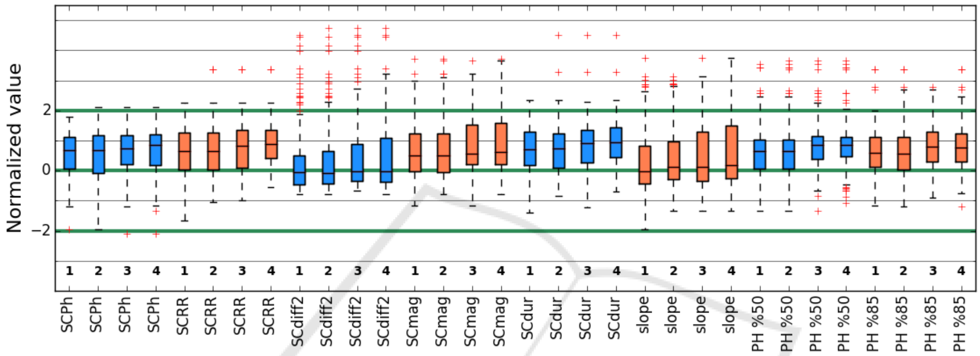

Figure 6: Boxplots of SC feature values in cluster *stress*, with number referring to the training iteration.

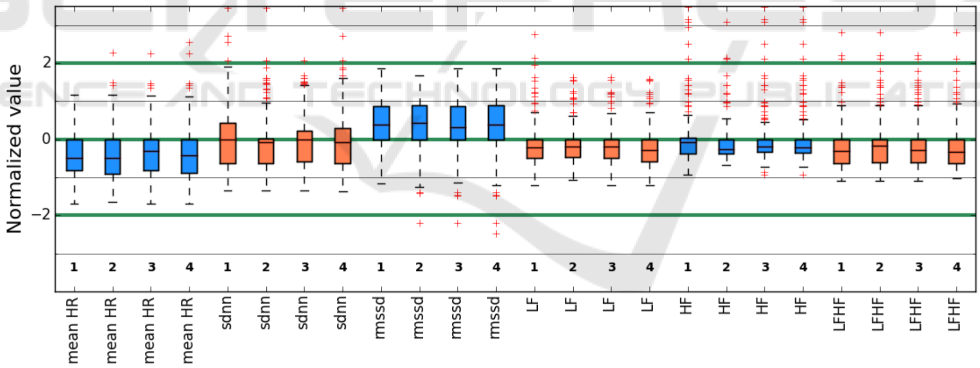

Figure 7: Boxplots of ECG feature values in cluster *relax*, with number referring to the training iteration.

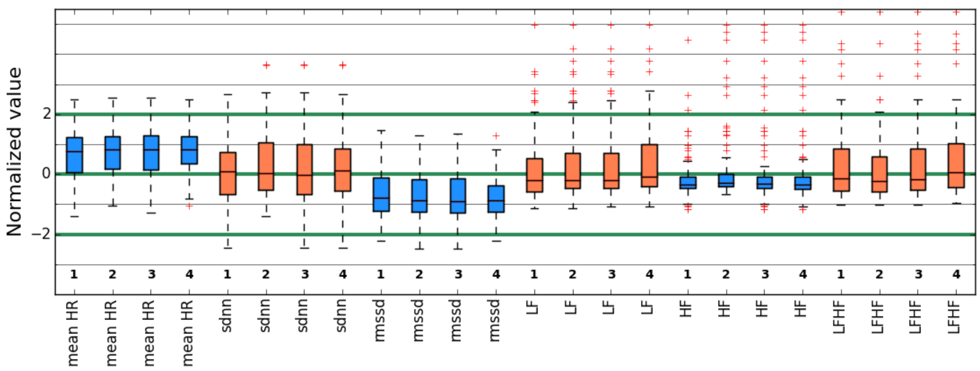

Figure 8: Boxplots of ECG feature values in cluster *stress*, with number referring to the training iteration.

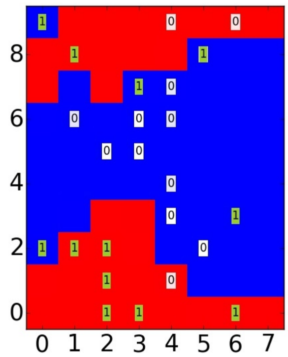

Figure 9: Clustered SOM with labels of projected testing data. Red cluster indicates a region classified as *stress* and blue for *relax*. Labels 1 indicate *stress* data points and labels 0 *relax* data points.

It is suggested to validate the developed model of unsupervised learning with *Self-Organizing Maps* against the field-data and field self-reports of ambulatory models such as the *cStress* model (Hovsepian et al., 2015). This model aims to provide a gold standard for continuous stress assessment of ambulant data. Validation against this data set would contribute to building a gold standard as more methods can be compared in the future.

5 CONCLUSIONS

An unsupervised algorithmic pipeline based on SOM and clustering has been introduced to explore the feasibility of SOM for unsupervised stress detection. It was tested on laboratory data. After ten runs of LOO validation, the testing performance is 79.0% ($\pm$5.16). As this was comparable to the performance of a supervised algorithm on a very similar test setup, the technique based on SOM is considered to be suitable for mental stress detection. Future research should investigate if the results obtained here in a controlled laboratory setting can be transferred to an ambulant environment.

ACKNOWLEDGEMENTS

We thank the therapists of Tumi Therapeutics for their help in patient recruitment and data collection. The authors report no conflict of interest with the current manuscript. The research was partly funded by a Ph.D. grant of the Flanders Innovation & Entrepreneurship agency (VLAIO) and by imec funds 2017.

REFERENCES

Adams, P., Rabbi, M., Rahman, T., Matthews, M., Voida, A., Gay, G., Choudhury, T., and Voida, S. (2014). Towards personal stress informatics: Comparing minimally invasive techniques for measuring daily stress in the wild. In *Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare*, PervasiveHealth '14, pages 72–79. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

Aguado, D., Montoya, T., Borras, L., Seco, A., and Ferrer, J. (2008). Using som and pca for analysing and interpreting data from a p-removal sbr. *Engineering Applications of Artificial Intelligence*, 21(6):919 – 930.

Arnrich, B., Kappeler-Setz, C., La Marca, R., Tröster, G., and Ehlert, U. (2010). Self organizing maps for affective state detection. In *Machine Learning for Assistive Technologies*.

Bauer, K., Muoz, G., and Moeck, I. (2012). Pattern recognition and lithological interpretation of collocated seismic and magnetotelluric models using self-organizing maps. *Geophysical Journal International*, 189(2):984–998.

Boucsein, W. (2012). *Electrodermal Activity*. Springer.

de Matos, M. C., Osorio, P. L. M., and Johann, P. R. S. (2007). Unsupervised seismic facies analysis using wavelet transform and self-organizing maps. *Geophysics*, 72(1):9–21.

Fahrenberg, J., Myrtek, M., Related, K. P., and Perrez, M. (2007). Ambulatory assessment - monitoring behavior in daily life settings. *European Journal of Psychological Assessment*, 23(4):206–213.

Grigore, O. and Bornoiu, I.-V. (2014). Kohonen neural network stress detection using only electrodermal activity features. *Advances in Electrical and Computer Engineering*, 14(3):71–78.

Han, J.-S. and Kim, G.-J. (2016). A method of unsupervised machine learning based on self-organizing map for bci. *Cluster Computing*, 19(2):979–985.

Hanke, M., Halchenko, Y. O., Sederberg, P. B., J.Hanson, S., Haxby, J. V., and Pollmann, S. (2008). Pymvpa: A python toolbox for machine-learning based data analysis. In *annual meeting of the Society for Neuroscience*.

Healey, J. A. and Picard, R. W. (2005). Detecting stress during real-world driving tasks using physiological sensors. *IEEE Transactions on Intelligent Transportation Systems*, 6:156 – 166.

Hjortskov, N., Rissén, D., Blangsted, A. K., Fallentin, N., Lundberg, U., and Søgaard, K. (2004). The effect of mental stress on heart rate variability and blood pressure during computer work. *European Journal of Applied Physiology*, 92(1):84–89.

Hovsepian, K., al'Absi, M., Ertin, E., Kamarck, T., Nakajima, M., and Kumar, S. (2015). cstress: Towards a gold standard for continuous stress assessment in the mobile environment. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubi-*

quitous Computing, UbiComp '15, pages 493–504, New York, NY, USA. ACM.

Kappeler-Setz, C., Arnrich, B., Schumm, J., La Marca, R., Tröster, G., and Ehlert, U. (2010). Discriminating stress from cognitive load using a wearable eda device. *IEEE Transactions on Information Technology in Biomedicine*, 14(2):410–417.

Karasek, R. and Theorell, T. (1992). *Healthy Work: Stress, Productivity, and the Reconstruction of Working Life*. Basic Books.

Kohonen, T., Schroeder, M. R., and Huang, T. S. (2001). *Self-Organizing Maps*. Springer-Verlag New York, Inc, 3rd edition.

Kusserow, M., Amft, O., and Tröster, G. (2013). Monitoring stress arousal in the wild. *IEEE Pervasive Computing Magazine*, 12(2):28–37.

Liao, L. and Carey, M. (2015). Laboratory-induced mental stress, cardiovascular response and psychological characteristics. *Reviews in Cardiovascular Medicine*, 16(1):28–35.

McCorry, L. K. (2007). Physiology of the autonomic nervous system. *American Journal of Pharmaceutical Education*, 71.

Medina, L. (2009). Identification of stress states from ecg signals using unsupervised learning methods. In *Portuguese Conf. on Pattern Recognition - RecPad*.

Melillo, P., Bracale, M., and Pecchia, L. (2011). Nonlinear heart rate variability features for real-life stress detection. case study: students under stress due to university examination. *BioMedical Engineering OnLine*, 10(1):96.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Rousseeuw, P. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.*, 20(1):53–65.

Siegrist, J. (2010). Effort-reward imbalance at work and cardiovascular diseases. *International Journal of Occupational Medicine and Environmental Health*, 23(3):279–285.

Smets, E., Casale, P., Großekathöfer, U., Lamichhane, B., De Raedt, W., Bogaerts, K., Van Diest, I., and Van Hoof, C. (2016). *Comparison of Machine Learning Techniques for Psychophysiological Stress Detection*, pages 13–22. Springer International Publishing.

Taelman, J., Vandeput, S., Spaepen, A., and Van Huffel, S. (2009). *Influence of Mental Stress on Heart Rate and Heart Rate Variability*, pages 1366–1369. Springer Berlin Heidelberg, Berlin, Heidelberg.

Task Force of the European Society of Cardiology the North American Society of Pacing Electrophysiology (1996). Heart rate variability standards of measurement, physiological interpretation, and clinical use.

Van der Elst, W., Van Boxtel, M., Van Breukelen, G., and Jolles, J. (2006). The stroop color-word test. *Assessment*, 13(1):62–79. PMID: 16443719.

Vesanto, J. and Alhoniemi, E. (2000). Clustering of the self-organizing map. *IEEE Transactions on neural networks*, 11(3):586–600.

Vrijkotte, T., Van Doornen, L., and De Geus, E. (2000). Effects of work stress on ambulatory blood pressure, heart rate, and heart rate variability. *Hypertension*, page 886.

Wijsman, J., Grundlehner, B., Liu, H., Hermens, H., and Penders, J. (2011). Towards mental stress detection using wearable physiological sensors. In *33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC 2011*, pages 1798–1801. IEEE Engineering in Medicine & Biology Society.

Zhai, J., Barreto, A., and Chin, C. (2005). Realization of stress detection using psychophysiological signals for improvement of human-computer interactions. In *Proceedings. IEEE SoutheastCon, 2005.*, pages 415–420.