Modeling Affective State using Learning Vector Quantization

de Vries, Jan

Publication date:
2014

Material based on:
de Vries, J. J. G., Pauws, S. C. and Biehl, M.: in press, Insightful Stress Detection from Physiology Modalities using Learning Vector Quantization, Neurocomputation

Hoofdstuk 4
EMOTION FROM A BODILY PERSPECTIVE

Abstract

Stress in daily life can lead to severe conditions as burn-out and depression and has a major impact on society. Being able to measure mental stress reliably opens up the ability to intervene in an early stage. We performed a large-scale study in which skin conductance, respiration and electrocardiogram were measured in semi-controlled conditions. Using Learning Vector Quantization techniques, we obtained up to 88% accuracy in the classification task to separate stress from relaxation. Relevance learning was used to identify the most informative features, indicating that most information is embedded in the cardiac signals. In addition to commonly used features, we also explored various novel features, of which the very-high frequency band of the power spectrum was found to be a very relevant addition.

4.1 Introduction

The harsh reality of daily life is that it becomes increasingly stressful. While certain levels of psychological stress help us perform optimally, prolonged exposure to stressors can have severe effects on wellbeing. Chronic stress is known to contribute to the development of, among others, cardiovascular diseases (Backé et al. 2012, Kivimäki et al. 2006) and has been found to contribute to high societal costs. In the US, for example, it has been estimated that job stress costs ”over $300 billion annually due to increased absenteeism, employee turnover, diminished productivity, medical, legal and insurance expenses, and workers’ compensation payments” (Rosch 2001).

The fine balance between the positive effects of short term stress and the detrimental effects of chronic stress on the one hand, and an increasingly demanding society on the other hand, indicate the need for assistance in balancing workload. Various products are available to help regulate mental stress (Heber et al. 2013, Westerink et al. 2014) including various biofeedback systems. One such biofeedback method is the stimulation of alpha-frequency brain waves, i.e., alpha neurofeedback (Dempster and Vernon 2009). Alpha brain waves are related to relaxation during wake, and stimulation of these waves are known to increase relaxation levels (Gruzelier 2002). The effects have been studied in the lab quite extensively, but only limitedly in circumstances that better reflect daily life. The application of neuro-feedback in a consumer device using the paradigm of music listening was researched (van Boxtel et al. 2012) in a double-blinded experiment with two types of control, as one aim of a comprehensive study. The effectiveness of such methods can, however, be further improved by providing them at the right moment to the right people. To that end, an objective method of measuring stress using easily and unobtrusively measurable physiological parameters is needed. This lead to the second aim of the aforementioned study: the development of such a method; which is subject of the present manuscript.

Several studies have attempted to classify stress from physiological measurements (Healey and Picard 2005, Zhai et al. 2005, Zhai and Barreto 2006, Choi and Gutierrez-Osuna 2009, Wijsman et al. 2011, Giakoumis et al. 2013) using various classification techniques. Among the more popular are Support Vector Machine (SVM) and Artificial Neural Network (ANN) (Sharma and Gedeon 2012). LVQ is a relatively novel technique that has been applied successfully to a wide range of classification challenges (Neural Networks Research Centre, Helsinki 2002), but rarely to classification of affect, and to the best of our knowledge, not yet to stress classification. The family of LVQ classification techniques use prototypes that are defined in the same mathematical space as the input data. The intuitive nature and ease of inspection give LVQ an advantage over less open-box methods such as SVM and ANN. We exploit this property of LVQ to gain new insights in the field of mental stress detection, where further understanding of the domain can help improve descriptive models (Sharma and Gedeon 2012).

In the present study, we set out to build classifiers to distinguish stress from relaxation using the three modalities of Electrocardiogram (ECG), Galvanic Skin Response (GSR), and Respiration (RSP). To that end we employ LVQ methods as well as SVM. We will use these methods to explore performance of uni-modal and multimodal classifiers in order to find out which signal is most rich in information to distinguish stressful reactions and investigate how individual features contribute. In the following, we will first create an overview of published affective and stress classifiers, then we describe the methods used, followed by results, discussion and conclusion.

4.2 Affect and Stress Classification

Whereas there are multiple definitions of stress that differ in various subtleties, an often used definition is that of Lazarus & Folkman: "Psychological stress is a relationship between the person and the environment that is appraised by the person as taxing or exceeding his or her resources and endangering his or her well-being" (Lazarus and Folkman 1984). Stress can be measured through a variety of physiological signals, among which Skin Conductance (SC), Skin Temperature (ST), Electrocardiogram (ECG), Blood Volume Pulse (BVP), Blood Pressure (BP), Electroencephalogram (EEG) and Electromyogram (EMG) (Sharma and Gedeon 2012). Because emotions and other affective states can also be measured using these signals, it is worth positioning our work in the light of other affective classifications as well.

Table 4.1 shows shows a snapshot of ten affect classification studies from physiology. It can be seen that a variety of physiological modalities is used as input, various techniques are applied and a variety of target classes are used. Because the number of classes, number of participants, prior probability of classes and methods used for validation vary between these studies, their performance cannot be compared directly. Nevertheless one can observe that there is room for improvement in terms of performance, which ranges between 61% and 86%, with the majority of performances between 70 and 80%.

Table 4.2 shows a detailed overview of studies that specifically classify mental stress. We observe that the performances reported are slightly higher than those reported for other affective states (Table 4.1). We observe that most studies report 'ordinary' cross validation in which data of participants is shared over training and test set, only a limited number of studies report participant-wise cross validation results in which participants are strictly separated over training and test set (i.e., no data of test-participants is used for training). The latter is generally more difficult than the former, which becomes also apparent in the performances in Table 4.2, but does better reflect the generalization performance (i.e., performance of the method for unseen users).

Table 4.1: *Review of ten machine learning studies employing different physiological signals to recognize various affective states.*

Reference	Modalities[1]	Ss[2]	Feat.[3]	Technique	Targets	Perf[4]
Sinha and Parsons (1996)	M	27	18	LDA	2 emotions	86%
Picard et al. (2001)	C,E,R,M	1	40	LDA	8 emotions	81%
Kim et al. (2004)	C,E,S	175		SVM	3 emotions	73%
Lisetti and Nasoz (2004)	C,E,S	29		kNN, LDA, ANN	6 emotions	86%
Rani et al. (2006)	C,E,S,M,P	15	46	kNN, SVM, RT, BN	3 emotions	86%
Kim and André (2008)	C,E,M,R	3	110	LDA, EMDC[5]	4 emotions	79%
Chanel et al. (2009)	C,E,R	11	18		3 emotions	66%
	B		18720		3 emotions	73%
	B,C,E,R		18738		3 emotions	70%
Hosseini et al. (2010)	C,E,R	15	38	SVM	2 arousal levels	77%
	B	15	21	LDA, SVM	2 arousal levels	85%
van den Broek, Lisý, Janssen, Westerink, Schut and Tuinenbreijer (2010)	E,M	21	10	kNN, SVM, ANN	4 emotions	61%
Katsis et al. (2008)	C,E,M,R	10	15	SVM, ANFIS	4 affect states	79%

[1] Abbreviations used: B Brain activity (EEG); C Cardiovascular activity (e.g., ECG and BVP);E Electrodermal activity (EDA); M Electromyogram (EMG); P Blood pressure; R Respiration; S Skin temperature

[2] Number of subjects

[3] Number of features

[4] Performance (accuracy)

[5] A tailored ensemble of binary classifiers

Table 4.2: *Review of machine learning studies employing different physiological signals to recognize stress.*

Reference	Modalities[1]	Ss[2]	Technique	Targets	Val.[3]	Perf[4]
Healey and Picard (2005)	C,E,R,M	9	LDA	3-level	CV	97%
Zhai et al. (2005)	E,C,O	6	SVM (linear kernel)	2-class	CV	57%
			SVM (RBF kernel)			60%
			SVM (sigmoid kernel)			80%
Zhai and Barreto (2006)	E,C,O,S	32	SVM	2-class	CV	90%
	C,O,S					90%
	E,O,S					90%
	E,C,S					61%
	E,C,O					89%
Choi and Gutierrez-Osuna (2009)	C,R	3	unspecified	2-class	CV within pp	83%
					pp-wise CV	69%
Wijsman et al. (2011)	C,R,E,M	21	Linear Bayes	2-class	CV	78%
			Quadratic Bayes			78%
			kNN			76%
			Fisher's Least Square			79%
Giakoumis et al. (2013)	E	24	LDA	2-class	CV	83%
	C					74%
	E,C					95%
	E,C				pp-wise CV	86%

[1] Abbreviations used: B Brain activity (EEG); C Cardiovascular activity (e.g., ECG and BVP);E Electrodermal activity (EDA); M Electromyogram (EMG); O Ocular Response (e.g., Pupil diameter); P Blood pressure; R Respiration; S Skin temperature
[2] Number of subjects
[3] Type of validation
[4] Performance (accuracy)

The study of Healey and Picard (2005) provided an exceptionally high performance of 97%. It should, however, be noted that their study is limited in the number of participants used (13) as well as using only one task for each stress level. Therefore, the high performance they obtained is likely biased by the specific set of participants and might reflect distinctions between the tasks rather than the stress levels. In general, we observe that the number of participants used in the studies is relatively limited: the studies included data from 3 to 32 participants. In our study we gathered data from more participants to have a more representative set of participants. We repeated measurements in 15 sessions to introduce temporal effects and environmental changes in the dataset that happen in daily life and influence the physiological measurements. Furthermore, we use multiple stressful tasks to induce more variety to better represent stressful situations in daily life.

Sharma and Gedeon (2012) made an extensive inventory of various aspects of stress detection. They conclude that ”Models developed to date that describe stress are quite simplistic. Generally, established techniques such as ANN and SVM have been used to model stress. Novel or more complex computational techniques are needed for stress models”. We believe that the application of LVQ classifiers can be such a novel computational technique and help gain more direct insight into the stress classification challenge and thereby provide valuable input to develop models that describe stress.

4.3 Method

The experiment performed to obtain the data that will be used in the analysis that is subject of this work is further described in van Boxtel et al. (2012). The following sections describe the most important details. The current problem is defined as a binary classification problem in discerning stressful from relaxation episodes from human physiological signals. Stressful episodes were operationalized as various mentally demanding tasks, relaxation episodes were operationalised listening to favourite music. Human physiological signals entail the following modalities: ECG, GSR and RSP.

4.3.1 Participants

Participants were recruited by means of a website that explained the procedures involved in the research in great detail. A total number of 171 persons indicated on the website that they wanted to participate in the research. 110 persons either did not follow up on our request, turned out to be unavailable at the time of the research, or decided to cancel their participation. The remaining 61 (20 male, 41 female) provided written informed consent. Their age ranged from 18 to 28 years (mean 21.2 years).

4.3.2 Design and procedure

Each participant returned 15 times within a period of 4 weeks for a session during which their physiology was measured. The sessions took place in a normal office room, in which each participant was seated in a comfortable reclining chair in front of a small table with a laptop on it. There were five such chairs and tables with laptops in the room, separated by wooden partitions, so that 5 participants could be trained at the same time by a single experimenter. The whole session was automated as much as possible. The experimenter supervised the sessions, and only took action in case something was wrong (usually bad electrode contacts, which were automatically signalled).

A training session on a particular day always consisted of the same sequence of tasks. After the signals were determined to be valid, a baseline measurement of five minutes rest with eyes opened was recorded, followed by 5 minutes with eyes closed. After that, 3 relaxation intervals of 8 minutes duration were interspersed by cognitive tasks lasting about 5 minutes each. The sequence of tasks are graphically represented in Figure 4.1. The (fixed) sequence was: Flanker task, relaxation 1, Stop-signal task, relaxation 2, Stroop task, relaxation 3, N-back task. During the relaxation intervals subconscious neuro-feedback was provided in three different ways, two of which are control conditions.

The interleaved task sequence was chosen for several reasons. First, it represents daily life stress, secondly it enhances changes in stress level which are particularly of interest for practical applications, and thirdly it provides a platform to test the relaxation effect of neuro-feedback.

Relaxation with Neuro-feedback

The participants were given a set of headphones that they used for listening to their favorite music. Participants could either bring their own music for that particular day on an MP3 player, or they could select that day's music from a playlist containing thousands of songs from various artists. There was no limitation to the kind of

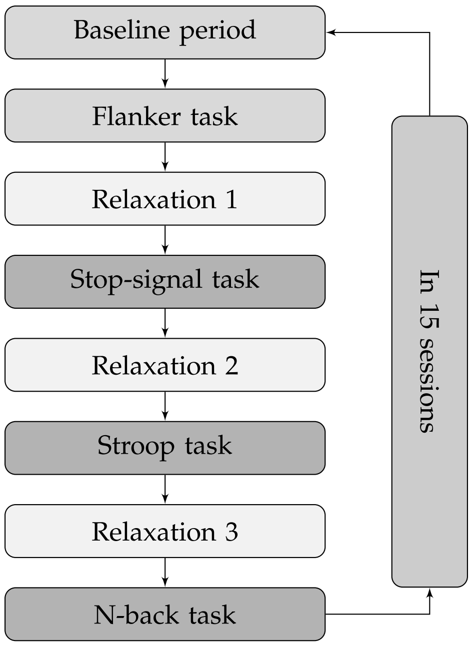

Figure 4.1: *Schematic outline of the experiment.*

music participants could listen to. Categories included genres like hard rock, easy listening and classical music.

As one part of this comprehensive study, the effects of neuro-feedback on relaxation were studied (van Boxtel et al. 2012). To that end, three conditions were used: alpha training and two types of controls, where one applies the same stimulation but at different (beta) frequencies that are not associated with relaxation and another control type where no stimulation is performed. Note that the stimulation was performed in a very subtle manner, as is described in the next paragraph, uses exactly the same setup over the three conditions, and has no direct effect on the peripheral physiological measurements taken (see Section 4.3.3).

The participants were randomly assigned to one of three groups: alpha training (A), random beta training (B), or control (C, music only); which was used for all sessions for this participant. Participants in group C listened to unaltered music, the music for the other two groups was altered by a high-pass filter of which the cut-off frequency was dynamically chosen. The cut-off frequency was adapted at real time based upon the frequency spectrum of the participants' EEG. To that end, the power in a target frequency range is calculated as relative to the total power (i.e., the power in the range 4-35Hz). The higher this relative power, the lower the cut-off frequency was chosen. The resulting effect is that lower (relative) power in the target frequency bands causes the low frequencies of the music to be filtered out, while high (relative) power in the target frequency will pass the music without much change (in the lower music frequencies). The target frequency ranges in the EEG spectrum were chosen as follows: The range for group A was based upon the power of alpha waves (8-12Hz), and for group B it was based upon beta waves (a randomized 4Hz bin in the range 16-36Hz). The alpha training was expected to increase relaxation while the other two types were not expected to have any effect on relaxation. These expectations were confirmed by van Boxtel et al. (2012).

From the 61 participants, 50 completed all training sessions (and without technical problems). The participants were distributed over the three groups as follows: A (alpha training): $N = 18$ (12 female; mean age 20.7 ± 1.8 years); B (random beta training): $N = 12$ (9 female; mean age 20.6 ± 1.5 years); C (control, music only): $N = 20$ (15 female; mean age 21.0 ± 2.1 years). Further details on the neuro-feedback training can be found in van Boxtel et al. (2012), the present study focusses at the stress and relaxation aspects of this study.

The mentally demanding tasks are further detailed in the following, taken from the study protocol (Sitskoorn et al. 2009).

Stop-signal task

> " The stop-signal task basic choice reaction time task. A green triangle (0.050 of screen width) on a black background is presented on the computer screen. Subjects have to indicate as a fast as possible the direction of the triangle. For a triangle to the left, subjects press the most left button of a button box and when it points to the right, the most right button has to be pressed. In one third of the trials the green arrow becomes red for 100ms and no answer has to be given, as depicted in Figure 4.2. When subjects are able to stop their response, the next time the stop signal will be given 50ms later to make it more difficult. When subjects give a response despite the presence of a stop signal, the signal appears the next trial 50 ms earlier to make it easier for the subject to stop the response. The task starts with a stop signal delay time of 250 ms and depending on the reaction of the subject, the stop signal delay time changes. Logan and Cowan (1984) fitted performance on this task in a formal model. The present task will use staircase tracking of response rate to arrive at

50% of successfully stopped trials, which is an optimal value for estimating inhibitory efficiency (Stop Signal Reaction Time (SSRT)). After one trial was finished, a fixation cross of 0.004 of the screen width appeared between 1 and 2 seconds on the screen before the next trial started. " (Sitskoorn et al. 2009)

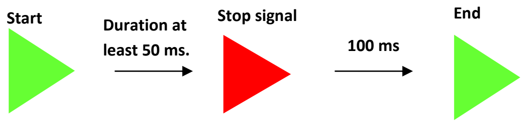

Figure 4.2: *Example of a Stop-signal task. The trial starts with a green arrow that depending on the subject's performance is green for a certain amount of time (at least 50 ms). After this time, the stop signal is initiated and the arrow becomes red for 100ms. This is followed by a green arrow that marks the end of the trial. It is the aim of the task that subjects do not give a response when the arrow becomes red.*

Stroop task

" The computerized version of the Stroop Color Word Test (SCWT) (Zysset et al. 2001) is used as a measure of executive functioning. In the Stroop task, subjects have to indicate whether the meaning of a word is the same as the color of which another word is printed in. Both words are not presented at exactly the same time to make it more difficult for the subject. In our version of the Stroop task, both words are printed above each other and the first word is presented 150 ms before the other word. In the case of Figure 4.3, the word "geelïs presented 150 ms before the word "rood". Both words are visible during 500 ms. In this period, subjects have to indicate whether the color of the upper word is the same as the meaning of the lower word. This requires inhibition of the automatic response to read the color word (Hammes 1971). Hence, this test is considered a measure of 'disinhibition' and it generally has high reliability (Bouma et al. 1996).
In the example, the color of the word "geelïs red and the meaning of the lower word is red, so the trial is correct. When the color and color name correspond subjects press with their index finger the 'yes'-button, if not, they press the 'no'-button. Whether the right or left index finger will be used for the yes and no response is counterbalanced between sessions.

Congruent trials are trials in which the color of the upper word is the same as the meaning of this word. For example, the word ”Blueïs written in blue ink and it means blue. When the color of the upper word is not the same as its meaning, the trial is called incongruent. In our experiment, four colors and the corresponding color names are used, namely red, yellow, blue and green. However, also the sign ”XXXXïs used as an upper word. The expectation is that subjects will make fewer mistakes when ”XXXXïs used as the upper word, because this word has no meaning and therefore subjects only have to deal with the color and not the meaning of the word. To keep between trials the attention of the subjects, a fixation cross with a variable duration between 1 and 2 seconds is presented on the computer screen. The duration of the fixation cross is variable to prevent a fixed rhythm of predicting and answering to the stimulus. ” (Sitskoorn et al. 2009)

Figure 4.3: *An example of an incongruent matching trial in the Stroop task. The word ”geel”means yellow but is written in red ink, so it is incongruent. The upper word is presented 150 ms before the lower word. Both words are visible during 500 ms. Between trials, a fixation cross with a duration between 1 and 2 seconds is presented on the computer screen.*

N-back task

” The N-Back task is a working memory task, introduced by Kirchner (1958), and requires subjects to decide whether each stimulus in a sequence matches the one that appeared N items previously. For example in a 3-back task subjects have to decide whether a letter currently presented on the screen is the same as three letters earlier. Our version of the N-back task is a 2-back task, meaning that subjects should decide whether the letter on the screen was the same as two letters ago. The used test set consists of 8 letters, namely B, F, K, H, M, Q, R and X. We decided not to use vowels to prevent the formation of words, which are more easily remembered than single letters. Furthermore, we use letters who are spatially different to be sure that when subjects make an error, it is caused by the difficulty of the task and not by confusion whether a letter was a V or W for example. Each letter will be presented for half a second on the computer screen. Between the end of a stimulus and the beginning of the next one, a fixation cross appeared on the screen. Except the 2-back trials, also lure trials were included in the task, such as depicted in Figure 4.4. These were trials in which the trial was 1-back or 3-back. When a trial was 2-back, subjects respond to the target by pressing the 'yes' button with their index finger. Whether the right or left index finger has to be used for the yes response is counterbalanced between sessions. " (Sitskoorn et al. 2009)

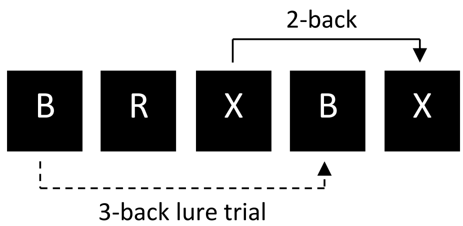

Figure 4.4: *An example of a sequence of letters in the N-back task. The last X matches the letter that was presented two items ago (X) and is therefore a 2-back. In this case subjects have to press the ijes-button on the button box, indicating that it was a 2-back. The letter B on the fourth position of the sequence is the same as the first one and is an example of a 3-back lure trial.*

Flanker task

" The Eriksen Flanker task (Eriksen and Eriksen 1974) is a basic choice reaction time task. In the task, five horizontally aligned arrows (size arrows: 0.050 of the screen width, space between arrows: 0.050 of the screen width) are presented on a 15 inch computer screen (resolution 1440 x 900 pixels, refresh rate 60 Hz) and subjects have to indicate the direction of the middle arrow (see Figure 4.5). Subjects can indicate this direction with a button box of which the most left button is pressed for an arrow pointing to the left and the most right button for an arrow pointing to the right. The two arrows on the left and right side of the middle arrow are flanker arrows and presented 150 ms before the middle arrow. These arrows are meant to distract the subject. The four flanker arrows always point in the same direction to the left or right. In this way, two situations can occur, namely that the flankers point in the same direction as the middle arrow (congruent) or that the flankers point in the opposite direction of the middle arrow (incongruent). The middle arrow with flankers will be present for 500 ms. After this period, a fixation cross (0.004 of screen width) appears at the same position as the middle arrow, namely in the center of the screen. To prevent that subjects learn when the next trial will start, the duration of the fixation cross will vary between 1 and 2 seconds. " (Sitskoorn et al. 2009)

Figure 4.5: *Stimuli used in the Flanker task. The two arrows on the left and right side of the middle arrow are the flanker arrows and presented 150 ms before the middle arrow appears. The arrows are white and presented on a black background. Subjects have to indicate the direction of the middle arrow. a) Congruent situation. The flanker arrows point in the same direction as the middle arrow. b) Incongruent situation. The flanker arrows point into the opposite direction of the middle arrow.*

For the classification analysis described in the Section 4.3.4, we selected the data gathered during the three relaxation tasks and the Stop-signal, Stroop and N-back tasks (as mentally stressful tasks). We did not include the Flanker task in the analysis as it turned out that participants were able to master the Flanker task very well after only a few attempts, thereby strongly reducing the mental stressfulness of the task in subsequent sessions. After each task the participants were asked to rate their level of stress vs relaxation on a visual analogue scale. Effectiveness of the induction of stress vs relaxation was tested by applying an ANOVA with repeated measures to these reported levels of stress.

4.3.3 Measurements

GSR was recorded from the left index finger, ECG was recorded from an electrode placed on the left wrist, and RSP was measured using a chest belt with stretch sensor. The signals were sampled at a rate of 1024 Hz (ECG), and 256 Hz (GSR, and RSP) by a 24 bit A/D converter on a Nexus-10 portable device (MindMedia B.V.,

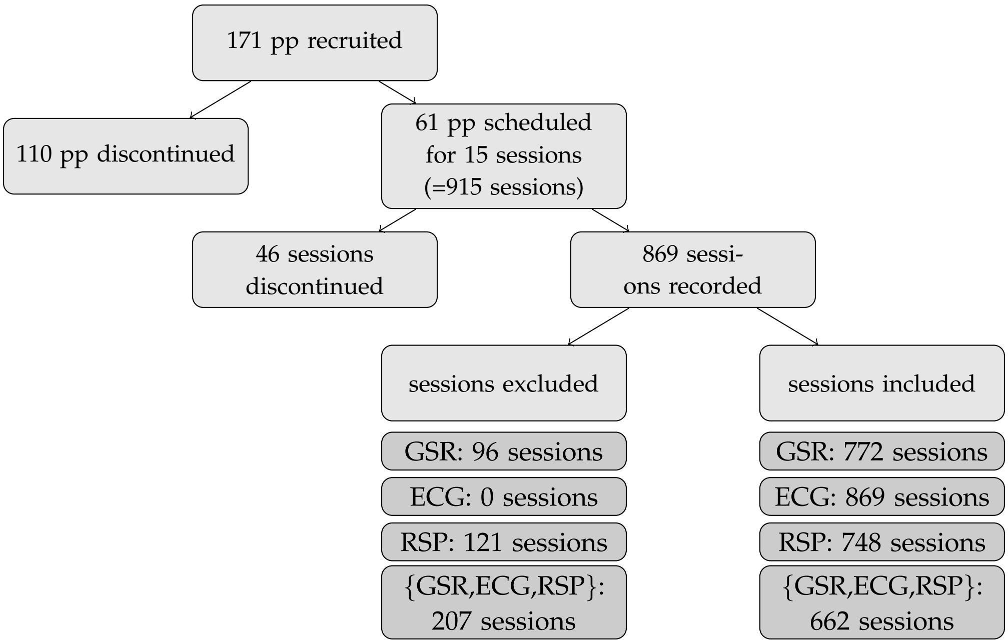

Figure 4.6: *Schematic outline of the data selection/exclusion process. Left branches show exclusion, and right branches inclusion of sessions.*

The Netherlands).

For each participant there were 15 sessions scheduled totalling 915 ($= 61 * 15$) sessions, of which 46 were discontinued due to technical problems or unconformity of participants, yielding 869 sessions. Due to bad signal quality (e.g., signals out of range of the measuring equipment) we further excluded, 96, and 121 sessions for GSR and RSP respectively, resulting in 772, and 748 sessions from analyses for these signals. No ECG sessions needed to be excluded. In total 662 sessions contained valid signals for all modalities. Figure 4.6 depicts the data selection (or exclusion) schematically.

Preprocessing & Feature extraction

The steps taken during preprocessing and feature extraction are schematically depicted in Figure 4.7. As a first step of preprocessing, the signals were downsampled to 512 Hz (ECG), and 128 Hz (GSR, and RSP). Subsequently, signals were analyzed

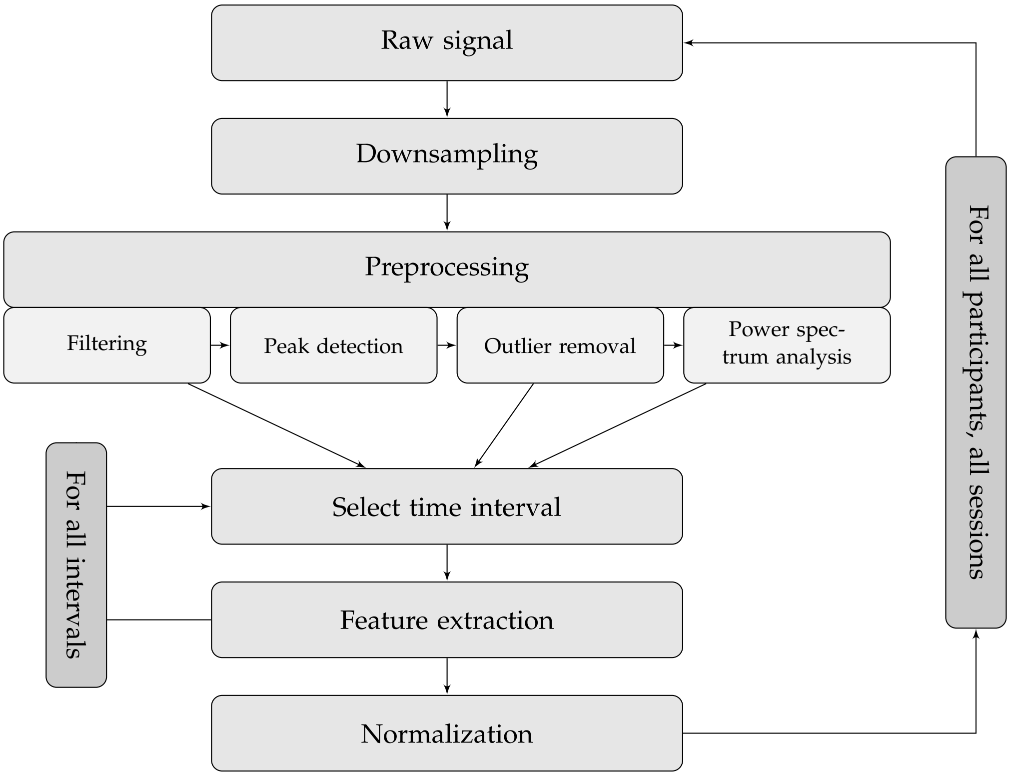

Figure 4.7: *Schematic overview of the feature extraction process.*

through the following dedicated preprocessing methods:

ECG preprocessing consisted of the following steps (as outlined in de Waele et al. (2009)): R-peak detection, IBI outlier removal, and Heart Rate Variability (HRV) analysis. R-peak detection was performed using a pattern matching technique (Poor 1994). The resulting intervals between the R-peaks, called Inter-Beat Interval (IBI), are filtered for outliers by using a sliding window histogram. In order to estimate frequency domain HRV features, an Autoregressive-Moving Average (ARMA) time series model was used to derive power in the frequency bands defined in the HRV guidelines paper (Malik et al. 1996), ranging from 0.04 to 0.15 Hz, which is known to vary with parasympathetic nervous system activity (Grossman and Taylor 2007). Next to the frequency domain HRV features, a variety of time domain HRV features is calculated, given that no 'golden standard' for HRV has been defined (Allen et al. 2007), as well as several features based on plain IBIs.

GSR was preprocessed using the SCRGauge method described in Kohlish (1992) which first subsamples the GSR signal to 1 Hz, uses cubic splines interpolation followed by a dedicated local maximum detection which is triggered by exceeding a certain gradient. Backward and forward searches are subsequently applied to detect the onset of Skin Conductances Responses (SCRs), and half recovery times. The raw GSR signal was used to derive several Skin Conductance Level (SCL) features from, the extracted SCRs to derive SCR features from, and from the residual signal that resides after subtracting SCRs from the raw signal, using the technique described in de Vries and van der Zwaag (2010) we derived features that represent purely the tonic part of GSR.

RSP signals were first lowpass filtered (cut-off 0.5Hz) and then analyzed for individual breaths. Using a localized min/max filter (Lemire 2006), local minima and maxima are detected. When found in the right order, they characterize a single breath. Based upon the distribution of identified breath amplitudes in a signal, too small or too large breaths (outliers) are removed. After this preprocessing the RSP signal is characterized by a sequence of breaths similar to the IBI signal for ECG.

All features have been calculated over equal length time intervals in order to avoid bias in duration dependent features (such as standard deviations) towards certain tasks. To this end, the first 5 minutes (which is the minimal duration of tasks) of measured signals from each task was taken to derived the feature values. A complete overview of extracted features can be found in Table 4.3. The specific features have been chosen such that they express the dynamics known to be relevant (Dawson et al. 2000, Stern et al. 2001, Task Force of the European Society of Cardiology the North American Society of Pacing Electrophysiology 1996) as they are modulated by the Autonomous Nervous System (ANS) that responds to stress. From this large set of features, we compiled a subset of features representing the most often used features in literature (inspired by the list in van den Broek et al. (2009)). They are marked with an asterisk in Table 4.3. In order to combine the data gathered from the different physiological signals to be used by a single classifier, we applied feature level fusion.

As highlighted in van den Broek, van der Zwaag, Healey, Janssen and Westerink (2010), there are many different techniques for normalization. In addition to the choice of which correction formula to use, the choice in defining the baseline period, there is also the choice of correcting on signal or feature level. The aim of normalization is to reduce the variance that occurs due to differences in physiology between participants, but also the long-term changes in physiology over time within participants, e.g., due to differences in physical fitness or the environment (such as temperature and humidity) (Boucsein 1992, Boucsein 2012). We have chosen for z-correction, a technique that compensates both for baseline level and variation, and is not too sensitive to outliers. Rather than applying the correction to the raw signals (which would only make sense for the skin conductance level), we apply it to the features derived, and we use the entire recording (all tasks) as reference signal, as suggested for e.g., SCR amplitude by Boucsein (Boucsein 1992, Boucsein 2012). Hence, after computing the features per task, we applied z-correction ($x_{corr} = \frac{x-\mu}{\sigma}$) to compensate for differences in physiological baselines between people and sessions. In this formula, μ represents the mean of a feature's values over all tasks within a single session (for a single participant), and σ the respective standard deviation.

Table 4.3: *Features extracted from the raw and preprocessed signals.*

ECG	IBI min	minimal IBI
	IBI max	maximal IBI
	IBI mean *	mean IBI
	IBI std *	standard deviation of IBIs, also referred to as SDNN
	IBI amp	amplitude of IBIs (max-min)
	IBI power VLF *	power of IBIs in very low frequency band ($0 - 0.04$ Hz)
	IBI power LF *	power of IBIs in low frequency band ($0.04 - 0.15$ Hz)
	IBI power HF *	power of IBIs in high frequency band ($0.15 - 0.4$ Hz)
	IBI power VHF	power of IBIs in very high frequency band ($0.4 - 1$ Hz)
	IBI power LH *	ratio between IBI power LF and HF
	IBI RMSSD *	root mean square of successive differences of IBIs
	IBI PNN50	proportion of IBIs > 50 ms
	IBI SDSD	standard deviation of successive differences of IBIs
GSR	SCL mean *	mean SCL
	SCL std *	standard deviation of SCL
	SCL grad	gradient of SCL (estimated by best linear fit)
	SCL min	minimal SCL
	SCL max	maximal SCL
	SCR freq *	number of SCRs per second
	SCR max amp	maximal amplitude of SCRs
	SCR mean amp *	mean amplitude of SCRs
	SCR sum amp	sum of amplitudes of SCRs
	SCR mean rise time *	mean rise time of SCRs
	SCR mean rec time *	mean half recovery time of SCRs
	SCR mean rise rec	mean ratio of rise time and half recovery time of SCRs
	SCR mean rise amp	mean ratio of rise time and amplitude of SCRs
	SCR mean rec amp	mean ratio of half recovery time and amplitude of SCRs
	SCRc SCL mean	mean SCL after correcting for SCRs
	SCRc SCL std	standard deviation of SCL after correcting for SCRs
	SCRc SCL grad	gradient of SCL after correcting for SCRs (estimated by best linear fit
	SCRc SCL min	minimal SCL after correcting for SCRs
	SCRc SCL max	maximal SCL after correcting for SCRs
RSP	mean rate *	mean respiration rate
	median rate	median respiration rate
	mean amp *	mean amplitude of respirations
	mean inh time	mean inhalation time
	mean exh time	mean exhalation time
	mean cycle	mean respiration time
	mean duty cycle	ratio between mean inhalation time and cycle
	mean inh exh	ratio between mean inhalation and exhalation time

* This feature is included in the commonly used set of features representing all modalities.

4.3.4 Classification analysis

In order to answer our research question of discerning stress from relaxation using physiology as input, we applied a selection of classifiers and further optimized their parameter settings using data from the individual physiological modalities (GSR, ECG and RSP) as well as the combined multi-modal dataset. Finally, we used the trained LVQ classifiers to derive which features were most influential in distinguishing stress from relaxation.

Learning Vector Quantization (LVQ) comprises a family of classifiers that is of open box nature, that is, they provide direct insight into the information learned by the classifier. LVQ, initially proposed by Kohonen (1990), defines prototypes $\mathbf{w}_T$ in the same (mathematical) space as the data (samples ξ) to represent the classes. These prototypes are directly interpretable as they show characteristics of classes in terms of the features chosen. During training, samples are presented sequentially, and for each sample the closest prototype(s) are updated by moving them towards or away from the presented sample. Several variants have been proposed, amongst which Robust Soft Learning Vector Quantization (RSLVQ) (Seo and Obermayer 2003), which introduces soft prototype assignments which act similarly to a soft window around the decision boundary (Witoelar et al. 2011), and Generalized Matrix Learning Vector Quantization (GMLVQ) (Schneider et al. 2009a), which introduces a relevance matrix Λ that is trained along with the prototypes. $\Lambda = \Omega^\top \Omega$, with Ω of size $M \times N$, is used to adapt the distance measure used by LVQ according to Equation 2.17. We have trained GMLVQ both with $2 \times N$ and $5 \times N$ sized matrices Ω, but since we observed identical performances, we will only report results of $2 \times N$ sized Ω. We will present results for RSLVQ and GMLVQ using one prototype per class as using more prototypes per class did not improve the results. In addition, we apply SVM (Vapnik 1998), a very popular technique in this domain of biomedical engineering. Next to linear SVM, which will be reported in the results, we also applied SVM with an RBF kernel, which however, did not improve upon the results.

Cross validation

In order to estimate generalization performance, we employed a cross validation scheme. Because physiological data shows large variation between participants (Gale and Edwards 1983, Boucsein 1992, Boucsein 2012), the most applicable, but also most challenging classification task is to separate training and test data not only per sample, but per subject. Hence we used 10 fold participant-wise cross validation, which divides the set of participants in tenths and repeatedly uses data from

90% of participants for training and the rest for testing. The results reported are means and standard deviations over 10×10-fold participant-wise cross validations. In addition to the participant-wise cross validation, we also performed 'ordinary' cross validation in which data of single participants can be split over both training and test set, thereby leaking some information from 'training participants' to the test set.

4.4 Results

Per participant, we have 15 sessions comprising 3 repetitions of 2 tasks, one operating in a stress condition and one operating in a relaxation condition. As dependent variable, we asked participants to report stress level using a visual analogue scale from zero to one (mean for relax condition: 0.29; for stressful condition: 0.38). It is evident that mean reported stress levels are derived from same participants measured in different sessions, repetitions and tasks, not from different participants. This refers to a within-subject or repeated measure design for statistical analysis. To demonstrate the induction of stress by means of mental and relaxation tasks, the aim of the analysis is to reject the null hypothesis of no difference in mean 'reported stress level' between stress conditioned tasks and relaxation conditioned tasks. We conducted an ANOVA with repeated measures with 'reported stress level' as dependent variable, and session (15), repetitions (3) and tasks (2) as within-subject independent variables. Missing values in reported stress levels were dealt with by means of case-based exclusion. We found a significant main effect for tasks ($F(1,36) = 38.3$, $p < 0.001$) allowing us to reject the null hypothesis of no difference in reported stress levels.

Tables 4.4 and 4.5 show the means and standard deviations per feature per task. They indicate, e.g., that the number of SCRs observed in the relaxation tasks is generally lower than in the mentally stressful tasks, as is the average SCL, respiration rate and amplitude. The heart rate variability, measured e.g. through IBI RMSSD, is generally higher in the stressful tasks. Within the mentally stressful or relaxation tasks the three subtasks show similar feature values.

The classification results of the participant-wise cross validation are shown in Table 4.6 when using data from single modalities, and Table 4.7 when combining data from multiple modalities. It can be observed that the performances are very similar over different classification techniques. Comparing the single modalities, the classifiers perform best on ECG, followed by GSR and RSP. Combining features from

Table 4.4: *Statistics (means and standard deviations) per feature per task of non-normalized data. See Table 4.3 for explanation of the feature names used.*

Feature name | Task	Relax 1	Relax 2	Relax 3	N-back	Stop-signal	Stroop
SCR freq	0.029 ± 0.034	0.029 ± 0.036	0.031 ± 0.036	0.057 ± 0.057	0.059 ± 0.059	0.062 ± 0.063
SCL mean	2.043 ± 1.398	2.122 ± 1.454	2.184 ± 1.460	2.447 ± 1.537	2.304 ± 1.518	2.416 ± 1.554
SCL std	0.248 ± 0.201	0.244 ± 0.197	0.242 ± 0.193	0.215 ± 0.184	0.240 ± 0.209	0.237 ± 0.195
RSP rate	15.432 ± 3.827	15.080 ± 3.867	14.951 ± 3.877	16.303 ± 3.700	16.488 ± 3.391	16.399 ± 3.580
RSP mean amp	6.367 ± 4.755	6.452 ± 5.016	6.419 ± 4.978	6.947 ± 5.124	6.781 ± 5.114	7.048 ± 5.766
IBI mean	0.825 ± 0.112	0.839 ± 0.113	0.853 ± 0.116	0.842 ± 0.115	0.835 ± 0.113	0.836 ± 0.114
IBI RMSSD	0.053 ± 0.034	0.057 ± 0.039	0.060 ± 0.038	0.073 ± 0.056	0.071 ± 0.052	0.071 ± 0.050

Table 4.5: *Statistics (means and standard deviations) per feature per task of normalized data. See Table 4.3 for explanation of the feature names used.*

Feature name | Task	Relax 1	Relax 2	Relax 3	N-back	Stop-signal	Stroop
SCR freq	-0.343 ± 0.712	-0.269 ± 0.845	-0.199 ± 0.843	0.433 ± 0.977	0.380 ± 0.848	0.441 ± 0.882
SCL mean	-0.202 ± 0.645	-0.048 ± 0.727	0.096 ± 0.803	0.647 ± 0.889	0.329 ± 0.650	0.563 ± 0.705
SCL std	0.069 ± 0.887	0.035 ± 0.891	0.031 ± 0.920	-0.205 ± 0.876	0.009 ± 0.854	0.021 ± 0.830
RSP rate	-0.279 ± 0.813	-0.449 ± 0.825	-0.522 ± 0.912	0.065 ± 1.050	0.197 ± 0.823	0.164 ± 0.860
RSP mean amp	-0.149 ± 0.874	-0.164 ± 0.851	-0.123 ± 1.007	0.139 ± 1.035	0.059 ± 0.849	0.131 ± 0.923
IBI mean	-0.065 ± 0.745	0.237 ± 0.793	0.627 ± 0.874	0.348 ± 1.016	0.251 ± 0.716	0.201 ± 0.809
IBI RMSSD	-0.238 ± 0.825	-0.114 ± 0.814	0.187 ± 0.917	0.585 ± 0.912	0.463 ± 0.751	0.455 ± 0.805

the three modalities improves performance, and adds an additional 3.5 percentage points on top of the performance obtained using only ECG features. This, however, only when including more than just the set of commonly-used features. For the richest data set covering all features, GMLVQ performs best with just under 87% accuracy. The corresponding Area Under the Curve (AUC) of the Receiver Operating Characteristic (ROC) is 0.95. Using 'ordinary' cross validation, performance is slightly higher with accuracies up to 88%.

The training performance is presented in Tables 4.8 and 4.9, again for uni- and multimodal input, respectively. It can be observed that the training performances for LVQ and linear SVM are only slightly higher than their respective test performances (Tables 4.6 and 4.7), while for the ANN and RBF SVM there is a larger difference in performance. We also verified whether parameter settings could be optimized, w.r.t. generalization performance, using only training data (and training performance) and found that this is possible on the present dataset, within two percentage points of the overall optimum, for all classifiers except kNN and ANN.

Table 4.6: *Test performance of 10x10 fold participant-wise cross validation on normalized data using single-modalities. Accuracy and AUC are listed as mean and standard deviation over the 10 repetitions of the 10-fold cross validation.*

	RSP			GSR			ECG		
Method	Hyper-parameter	Accuracy	AUC	Hyper-parameter	Accuracy	AUC	Hyper-parameter	Accuracy	AUC
Baseline		50.0%	0.5		51.4%	0.5		50.0%	0.5
kNN	$k = 11$	$69.0\% \pm 0.3\%$	$NaN \pm NaN$	$k = 11$	$71.9\% \pm 0.2\%$	$NaN \pm NaN$	$k = 11$	$81.7\% \pm 0.2\%$	$NaN \pm NaN$
ANN	$N_{hidden} = 3$	$70.4\% \pm 0.5\%$	0.748 ± 0.091	$N_{hidden} = 3$	$73.6\% \pm 0.6\%$	0.813 ± 0.070	$N_{hidden} = 5$	$84.1\% \pm 0.2\%$	0.913 ± 0.094
SVM - linear	$C = 1$	$71.0\% \pm 0.2\%$	0.748 ± 0.089	$C = 0.1$	$74.8\% \pm 0.3\%$	0.814 ± 0.062	$C = 0.1$	$83.2\% \pm 0.1\%$	0.897 ± 0.116
SVM - RBF	$C = 1$	$69.7\% \pm 0.5\%$	0.745 ± 0.068	$C = 10$	$69.7\% \pm 0.3\%$	0.775 ± 0.075	$C = 1$	$81.6\% \pm 0.3\%$	0.882 ± 0.106
Means		$65.7\% \pm 0.1\%$	0.684 ± 0.117		$71.3\% \pm 0.1\%$	0.778 ± 0.064		$76.8\% \pm 0.2\%$	0.846 ± 0.130
GLVQ		$67.7\% \pm 0.2\%$	0.709 ± 0.115		$71.6\% \pm 0.3\%$	0.783 ± 0.064		$77.2\% \pm 0.2\%$	0.848 ± 0.136
RSLVQ	$v_{soft} = 0.5$	$70.9\% \pm 0.3\%$	0.746 ± 0.091	$v_{soft} = 1$	$74.8\% \pm 0.2\%$	0.812 ± 0.062	$v_{soft} = 1$	$83.1\% \pm 0.1\%$	0.895 ± 0.109
GRLVQ		$68.0\% \pm 0.4\%$	0.707 ± 0.121		$66.6\% \pm 1.9\%$	0.715 ± 0.107		$70.2\% \pm 1.8\%$	0.721 ± 0.151
GMLVQ		$71.0\% \pm 0.3\%$	0.744 ± 0.091		$74.4\% \pm 0.3\%$	0.810 ± 0.058		$83.4\% \pm 0.2\%$	0.895 ± 0.115

Table 4.7: *Test performance of 10x10 fold participant-wise cross validation on normalized data combining multiple modalities. Accuracy and AUC are listed as mean and standard deviation over the 10 repetitions of the 10-fold cross validation.*

	Multi-selection			Multi-all		
Method	Hyper-parameter	Accuracy	AUC	Hyper-parameter	Accuracy	AUC
Baseline		51.3%	0.5		51.3%	0.5
kNN	$k = 11$	$79.1\% \pm 0.2\%$	-	$k = 11$	$81.7\% \pm 0.3\%$	-
ANN	$N_{hidden} = 3$	$81.6\% \pm 0.4\%$	0.848 ± 0.126	$N_{hidden} = 5$	$85.4\% \pm 0.5\%$	0.936 ± 0.017
SVM - linear	$C = 0.1$	$81.4\% \pm 0.3\%$	0.848 ± 0.138	$C = 0.1$	$86.6\% \pm 0.2\%$	0.954 ± 0.015
SVM - RBF	$C = 10$	$73.8\% \pm 0.3\%$	0.792 ± 0.138	$C = 0.001$	$51.3\% \pm 0.0\%$	0.849 ± 0.034
Means		$78.9\% \pm 0.2\%$	0.821 ± 0.132		$80.2\% \pm 0.2\%$	0.915 ± 0.029
GLVQ		$78.9\% \pm 0.2\%$	0.818 ± 0.129		$80.6\% \pm 0.2\%$	0.913 ± 0.029
RSLVQ	$v_{soft} = 1$	$81.6\% \pm 0.3\%$	0.847 ± 0.133	$v_{soft} = 1$	$86.6\% \pm 0.2\%$	0.950 ± 0.016
GRLVQ		$64.7\% \pm 2.0\%$	0.641 ± 0.130		$72.2\% \pm 1.1\%$	0.773 ± 0.080
GMLVQ		$81.6\% \pm 0.2\%$	0.845 ± 0.135		$86.7\% \pm 0.2\%$	0.951 ± 0.016

Table 4.8: *Training performance of 10x10 fold participant-wise cross validation on normalized data using single-modalities. Accuracy and AUC are listed as mean and standard deviation over the 10 repetitions of the 10-fold cross validation.*

	RSP			GSR			ECG		
Method	Hyper-parameter	Accuracy	AUC	Hyper-parameter	Accuracy	AUC	Hyper-parameter	Accuracy	AUC
Baseline		51.4%	0.5		51.4%	0.5		50.0%	0.5
kNN	$k = 11$	$75.3\% \pm 0.1\%$	-	$k = 11$	$77.1\% \pm 0.1\%$	-	$k = 11$	$84.8\% \pm 0.0\%$	-
ANN	$N_{hidden} = 3$	$72.7\% \pm 0.1\%$	0.798 ± 0.009	$N_{hidden} = 3$	$77.3\% \pm 0.2\%$	0.846 ± 0.007	$N_{hidden} = 5$	$85.7\% \pm 0.1\%$	0.929 ± 0.003
SVM - linear	$C = 1$	$71.9\% \pm 0.0\%$	0.783 ± 0.009	$C = 0.1$	$75.6\% \pm 0.0\%$	0.818 ± 0.006	$C = 0.1$	$83.7\% \pm 0.0\%$	0.903 ± 0.004
SVM - RBF	$C = 1$	$78.4\% \pm 0.1\%$	0.869 ± 0.005	$C = 10$	$100.0\% \pm 0.0\%$	1.000 ± 0.000	$C = 1$	$97.2\% \pm 0.0\%$	0.996 ± 0.000
Means		$66.1\% \pm 0.0\%$	0.712 ± 0.011		$71.7\% \pm 0.0\%$	0.777 ± 0.007		$77.3\% \pm 0.0\%$	0.851 ± 0.005
GLVQ		$68.3\% \pm 0.1\%$	0.738 ± 0.009		$71.9\% \pm 0.1\%$	0.780 ± 0.009		$77.5\% \pm 0.1\%$	0.850 ± 0.006
RSLVQ	$v_{soft} = 0.5$	$71.7\% \pm 0.1\%$	0.782 ± 0.009	$v_{soft} = 1$	$75.6\% \pm 0.0\%$	0.816 ± 0.006	$v_{soft} = 1$	$83.6\% \pm 0.0\%$	0.899 ± 0.004
GRLVQ		$67.9\% \pm 0.2\%$	0.717 ± 0.013		$67.1\% \pm 0.8\%$	0.689 ± 0.062		$70.3\% \pm 1.3\%$	0.720 ± 0.081
GMLVQ		$71.9\% \pm 0.1\%$	0.778 ± 0.011		$76.0\% \pm 0.1\%$	0.816 ± 0.006		$83.8\% \pm 0.0\%$	0.901 ± 0.004

Table 4.9: *Training performance of 10x10 fold participant-wise cross validation on normalized data combining multiple modalities. Accuracy and AUC are listed as mean and standard deviation over the 10 repetitions of the 10-fold cross validation.*

	Multi-selection			Multi-all		
Method	Hyper-parameter	Accuracy	AUC	Hyper-parameter	Accuracy	AUC
Baseline		51.3%	0.5		51.3%	0.5
kNN	$k = 11$	$83.4\% \pm 0.0\%$	-	$k = 11$	$85.8\% \pm 0.1\%$	-
ANN	$N_{hidden} = 3$	$84.0\% \pm 0.1\%$	0.918 ± 0.005	$N_{hidden} = 5$	$92.0\% \pm 0.3\%$	0.971 ± 0.003
SVM - linear	$C = 0.1$	$82.9\% \pm 0.0\%$	0.906 ± 0.006	$C = 0.1$	$88.2\% \pm 0.0\%$	0.947 ± 0.002
SVM - RBF	$C = 10$	$100.0\% \pm 0.0\%$	1.000 ± 0.000	$C = 0.001$	$51.3\% \pm 0.0\%$	1.000 ± 0.000
Means		$79.6\% \pm 0.0\%$	0.878 ± 0.006		$81.0\% \pm 0.0\%$	0.886 ± 0.003
GLVQ		$79.6\% \pm 0.0\%$	0.876 ± 0.007		$81.5\% \pm 0.0\%$	0.885 ± 0.003
RSLVQ	$v_{soft} = 1$	$82.7\% \pm 0.0\%$	0.904 ± 0.005	$v_{soft} = 1$	$88.2\% \pm 0.0\%$	0.943 ± 0.002
GRLVQ		$66.3\% \pm 1.1\%$	0.710 ± 0.075		$72.9\% \pm 0.5\%$	0.774 ± 0.021
GMLVQ		$83.0\% \pm 0.0\%$	0.903 ± 0.005		$88.4\% \pm 0.1\%$	0.943 ± 0.002

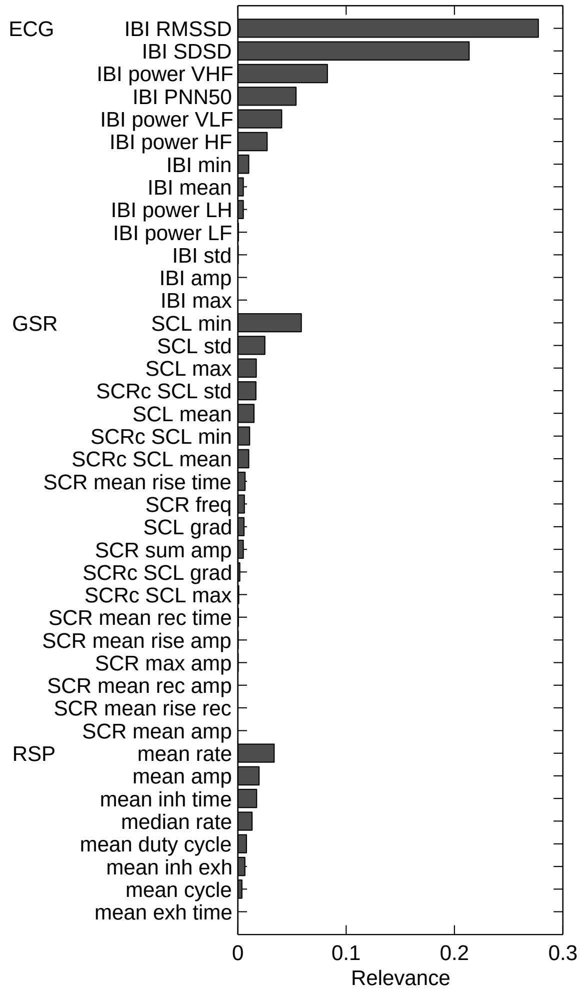

Figure 4.8: *Relevances trained by GMLVQ. For explanation of the feature names used, see Table 4.3.*

The diagonal elements of the relevance matrix trained by GMLVQ indicate relevance of individual features to the classifier's decisions. Figure 4.8 shows these relevances and indicates that most informative features come from the ECG modality, followed by GSR and RSP. The most important individual contributions come from the time domain HRV measures SDSD and RMSSD, followed by the frequency domain HRV measured by the power in the VHF range. The most influential GSR feature is the minimum SCL and for RSP the mean rate.

4.5 Discussion

The classification performances for uni-modal stress classification range from 71% for RSP, to 83.4% for ECG indicating that most information can be found in the ECG signal. By combining modalities the performance is further increased to 86.7% in participant-wise cross-validation. Using 'ordinary' cross-validation we obtained an accuracy of 87.7%. Both LVQ methods we employed (RSLVQ and GMLVQ) performed well and slightly outperformed the popular SVM, that we included for reference.

From the training performances we observe that ANN and RBF SVM have the tendency to overtrain and thereby have a poor generalization performance. The parameter settings of LVQ variants as well as linear SVM could be optimized using training performances, thereby enabling to build a classifier solely on training data without the need for a test or validation set while yielding good generalization performance.

With these accuracies, our methods outperform the affective classifiers that are listed in Table 4.1. They also outperform the participant-wise validated stress-classifiers in Table 4.2 as well as most others using other cross validation schemes. In comparison to other studies, we used data obtained from a larger number of participants, and also repeated measurements for every participant in 15 sessions spread over several weeks. Thereby, we used a more representative sample of participants and obtained reliable estimates of generalization performance.

The most important feature was found to be the RMSSD. Although it reflects high frequency modulation of heart rate that, in general, is affected by Respiratory Sinus Arrhythmia (RSA), RMSSD has been shown to be unaffected of breathing (Penttilä et al. 2001). The second most influential measure was the related SDSD. The PNN50 that correlates with RMSSD (Task Force of the European Society of Cardiology the North American Society of Pacing Electrophysiology 1996), was also identified as quite relevant. It is worth noting that the two most influential features (RMSSD and SDSD) have been found earlier to be most reliable measures for short term intervals (McNames and Aboy 2006), i.e., in the order of 5 minutes, which reflects the measurement time in our experiments. The importance of various HRV measures for distinguishing can be explained by the fact that they reflect parasympathetic (HF power, RMSSD) and sympathetic (LF power, SDNN) nervous system responses which are known to relate to the fight-or-flight response and dampening response, respectively (Task Force of the European Society of Cardiology the North American Society of Pacing Electrophysiology 1996).

Out of the frequency domain HRV measures VHF power we found most influential. Many studies that include cardiac activity as measure for stress only use features from the lower frequency (up to HF) ranges and do not usually consider frequencies above 0.4Hz. This might be due to various reasons. First, the mechanisms that affect VHF power are not well understood. Second, higher frequency ranges cannot be measured reliably through the most commonly used modality BVP which has less sharp peaks, thereby not allowing for a very accurate detection of heart beats which is reflect particularly in inaccuracy in the higher frequencies of HRV. Our use of ECG enabled the reliable use of VHF power as a measure of stress.

We also inspected the prototypes trained by GMLVQ to verify that the stress prototype, as compared to the relax prototype, is characterized by higher heart rate and generally higher HRV values with the exception of RMSSD. Especially for HRV there are varying results published (Mathewson et al. 2010, Wright et al. 2007, Vuksanović and Gal 2007), which is confirmed by Berntson and Caciopo, who observed that "it is clear that no single pattern of autonomic adjustments and associated changes in heart rate variability will apply universally across different stressors" (Berntson and Cacioppo 2004). In our study we included three different stressors to induce stressful situations, thereby creating more robustness against this effect. Further, the stress prototype is characterized by more SCRs, higher maximum SCL and faster, though deeper, breathing. These findings are in line with observations made by others (Boucsein 2012, Houtveen et al. 2002).

4.6 Conclusion

We have successfully built classifiers of stress from three physiological modalities and observed that the cardiac activity made the strongest uni-modal classifier with over 83% accuracy. Combining the three modalities into a multi-modal classifier improved performance further up to 88% accuracy. By using data from a large sample of participants and repeated sessions we ensured good generalizability to unseen users. The LVQ techniques slightly outperformed well-known techniques such as SVM. These open-box methods allowed us to observe the most important features for stress detection. These were the time domain HRV measures RMSSD and SDSD. The third most important features was found to be very high-frequency HRV from ECG. Most other studies use BVP to measure cardiac activity, however that does not allow for accurate VHF HRV measurements. Therefore it might be advisable for methods that aim at stress detection to use ECG rather than BVP as measurement modality of cardiac activity.

The classifiers built and the knowledge gained on important features for the distinction between stress and relaxation using physiological parameters have brought us one step closer to the realization of a system that can monitor physiology during the day and help its users to monitor their stressful moments during the day. In case a certain quota has been reached or a stressful period reaches a certain duration such a system could trigger the user and offer a means to relief the stress, e.g., a paced breathing exercise (Westerink et al. 2014).

While we have setup our experiments such that they represent daily life as well as possible, the measurements were taken during semi-lab circumstances. Future work should look into the application of the developed classifiers in daily-life measurements and observe their performance. This brings the challenge of reliable ground truth measurements, however this might become more and more feasible with the rapid development of various technologies such as Google Glass (Google 2014) that can capture context. It would be interesting to expand the classifier to also be able to classify other affective phenomena such as emotions and moods.