A Case-Based Classification of Respiratory Sinus Arrhythmia

Markus Nilsson, Peter Funk

Mälardalen University
Department of Computer Science and Engineering
S:ta Ursulas väg 2A
P.O. Box 883, SE-721 23 Västerås, Sweden
{markus.nilsson;Peter.Funk}@mdh.se

Abstract. Respiratory Sinus Arrhythmia has until now been analysed manually by reviewing long time series of heart rate measurements. Patterns are identified in the analysis of the measurements. We propose a design for a classification system of Respiratory Sinus Arrhythmia by time series analysis of heart and respiration measurements. The classification uses Case-Based Reasoning and Rule-Based Reasoning in a Multi-Modal architecture. The system is in use as a research tool in psychophysiological medicine, and will be available as a decision support system for treatment personnel.

1 Introduction

This paper describes a system for pattern classification of Respiratory Sinus Arrhythmia (RSA). The patterns are classified with Case-Based Reasoning (CBR) and Rule-Based Reasoning (RBR) using physiological time series measurements. The system is developed to be a decision support system for treatment personnel, as well as a research tool in psychophysiological medicine. The next paragraph defines RSA and put it into clinical context.

Respiratory Sinus Arrhythmia is described as centrally modulated cardiac vagal and sympathetic efferent activities associated with respiration [3]. RSA occurs because the heart rate, i.e. the number of beats per minute, is variable. This Heart Rate Variability (HRV) is an effect of inhibitions on the sympathetic and parasympathetic systems while breathing. The sympathetic and the parasympathetic systems, which are a part of the autonomous nervous system, have different activity levels during different stages of the respiration cycle [9]. That is, RSA is a sinusoid pattern of the heart rate associated with the breathing. The pattern is directly connected to the state of the respiration of a normal healthy person, i.e. without cardiac and pulmonary dysfunctions. The pulse increases dramatically on an inhalation and decreases in the same fashion on an exhalation. This is illustrated in figure 3 in section 2.2.

Physicians detect irregular heart rate patterns by analysing the RSA. Some of the irregularities are dysfunctions caused by physiological and/or psychological stress. A common diagnostic method for detecting dysfunctions in RSA is to manually analyse sampled heart rate measurements together with an analysis of the measurements' frequency spectrum [4, 3, 9]. The dysfunctions are treated with cognitive behavioural sessions with psychologists, and with biofeedback training [8].

Time-series analysis in medical Case-Based Reasoning has previously been studied by Montani *et al.* [11, 12], where they integrated CBR, RBR and Model-Based Reasoning (MBR) in a Multi-Modal Reasoning (MMR) platform for managing, i.e. suggesting insulin therapy, for type 1 diabetic patients. Another CBR system which analyses time series is ICONS [17]. ICONS forecasts kidney functions by an extended CBR cycle which abstracts states from measurements and predicts trends from the states. Other related medical CBR systems are The Auguste project [10], Bichindaritz's CARE-PARTNER [2] and MNAOMIA [1], and Perner *et al.*'s airborne fungi detection system [16]. Further information of these systems can be found in Nilsson and Sollenborn's survey on medical CBR [14]. CBR in the medical domain was first pioneered in the late 1980's. Some of the early systems are PROTOS [5] and CASEY [6].

We propose a MMR system design for the classification of RSA, where CBR matches physiological parameters and RBR reduces the domain of cases. A system design for the classification of RSA is introduced in the next section. We evaluate the proposed system in section 3, and conclude the paper in section 4.
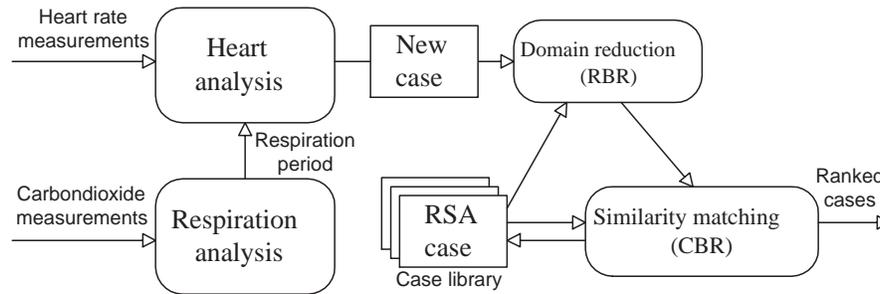
2 System Architecture

Fig. 1. A design for a classification system of Respiratory Sinus Arrhythmia.

A classification system for RSA is naturally divided into two initial analytical stages. Each stage analyses time series measurements. The first stage analyses

the respiration and the second stage analyses the heart measurements. Cases are thereafter created based on the findings in the analysis processes. Rules limit the number of cases for the matching procedure to compare to, and the cases that pass the filter are matched and ranked. The design is illustrated in figure 1. The system is a revised version of the two later parts of the design described in [13], the first part is processed in the hardware. Each part of the figure is described in the remainder of this section. The respiration analysis is described in subsection 2.1, followed by heart analysis, domain reduction, case matching and finally the user interface.

As RSA is quantified during a breath (a respiration cycle), a respiration analysis precedes the heart analysis. The respiration analysis locates when the respiration occurs and passes that information to the heart analysis.

2.1 Respiration analysis

A breath begins, by definition, on an inhalation. Hence, the respiration cycle starts when an exhalation stops and inhalation begins. Capnograph [7] measurements are used to pinpoint the beginning and end of the respiration cycle. The capnograph is a non invasive method, and measures the contents of carbon dioxide (CO_2) in exhaled air. Capnograph measurements are depicted in figure 2.
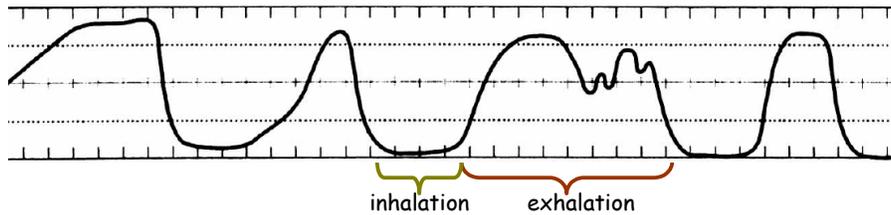

Fig. 2. Capnography measurements illustrating the respiration cycle, divided into inhalation and exhalation. The picture is adapted from [7].

Finding either the beginning or the end of the respiration cycle is actually sufficient to determine the entire respiration cycle, since the end of a respiration cycle marks the beginning of the next. A new breath starts, in the ideal case, when the level of CO_2 dramatically drops from circa 5% to just above 0%, followed by a steadily low level. This low level of CO_2 occurs during the entire inhalation. The level of CO_2 never reaches 0% because the surrounding air naturally contains CO_2, and it is also difficult to vacate the measuring sensor from all gases, even with a pump driven device.

A rough estimate of the respiration period is calculated by searching for a local maximum followed by a local minimum. The maximum represents the exhalation and the minimum the inhalation. A simulated annealing algorithm is then used on the first order derivates of the CO_2 measurements to find an approximate position between the maximum and minimum. The position is where the exhalation stops and the inhalation starts, i.e. where the respiration cycle begins.

Two points are identified, the first as the beginning and the second as the end of the respiration cycle. The samples in the respiration period are shifted in time due to lag in the sensor and additional delays associated with capnography measurements. A major delay is the transportation of CO_2 from the measuring point to the sensor. The CO_2 is sucked through a tube with a pump. The corrected measurements are then sent to the heart analysis as seen in figure 1.
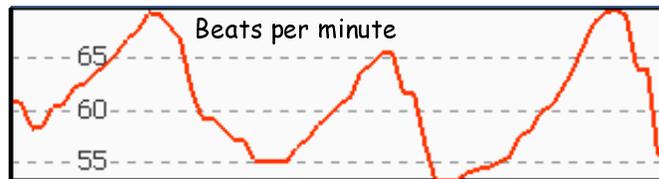
2.2 Heart analysis

Fig. 3. The heart rate variability, i.e. the oscillating effect of the heart can easily be seen in these heart rate measurements.

Physicians observe both the HRV and the frequency spectrum of the HRV when they classify RSA. The beginning and the end of a HRV period is based on the respiration analysis. The HRV period span over the same time period as the respiration period, and is calculated from heart rate measurements. The heart rate measurements are mean-valued electrocardiogram (ECG) measurements. The conversion from ECG to heart rate measurements are automatically computed in the hardware [1]. HRV measurements are depicted in figure 3.

$$\sum_{i=1}^{n} \left(HR(i) - \frac{\sum_{j=1}^{n} HR(j)}{n} \right) = 0 \tag{1}$$

The frequency spectrum is calculated when the samples for the HRV have been collected. However, some pre-processing is required before a frequency spectrum can be calculated. The physicians are only interested in the oscillation of

---

[1] The AirPas and cStress hardware environments from PBM StressMedicine are used to measure physiological parameters.

the sequence of samples, HR, that make up the HRV, when they observe the frequency spectrum. The sample sequence has to be shifted to oscillate around its own mean value, as seen in equation 1. If not, a large portion of the lower end of the frequency spectrum is mixed with non relevant oscillations due to the nature of the heart rate samples. The heart rate samples are always positive numbers with a range of about 50-90 beats per minute, which unintentionally create large sine waves, or low frequencies within the measurement sequence.
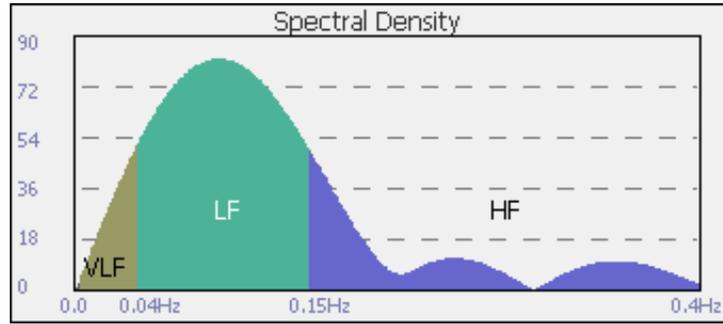

Fig. 4. A frequency spectrum of a typical RSA. Physicians are only interested in the range from 0 to $0.4Hz$. The spectrum is divided into three major frequency bands. Very Low Frequencies (VLF), Low Frequencies (LF) and High Frequencies (HF), as various physiological variables appear within these individual bands.

The output sample rate from the hardware sensors is 2 Hz; and a normal breath is in the range of 6-12 seconds. Hence, there are usually too few samples in the HRV to make any useful frequency transformation. The solution is to pad, or to add, the sample sequence with zeroes. Padding with zeroes does not affect the frequency distribution in the spectrum. The sample sequence is padded to 2048 samples. The samples are then transformed to the frequency spectrum using a Fast Fourier Transformation (FFT). The length, or power value, of each frequency is calculated from the FFT's output of complex numbers, see equation 2, and figure 4.

$$Power(f) = \sqrt[2]{FFTreal(f)^2 + FFTimg(f)^2} \qquad (2)$$

Physicians study additional parameters in their classification of RSA. The additional parameters are notch patterns and peak-to-valley differences in the heart rate measurements. The peak-to-valley value is the ΔY difference of the maximum and minimum heart rate sample values. Notches are irregular dips in the otherwise smooth heart rate oscillation. The notches have different significances depending on where they occur. Both peak-to-valley and notches are calculated.

2.3 Cases and domain reduction

Cases contain all above described parameters and measurements, with one addition, first order derivates of the heart rate measurements are also included. A case is represented by 17 features, 10 of which are actively used in the matching process. Five features contain time-series or other sequences, they are the CO_2 measurements, heart rate measurements, first order derivates of the heart rate measurements, the power and frequency table of the FFT. The remaining 7 structure the Case, such as the RSA class, a couple of variables to speed up the matching process. Sample frequencies for both heart rate and CO_2 is also included, and finally a second classification of RSA. The second classification is based on another quantification approach for RSA [4] and is not of interest for this paper.

A case belongs to one of the stereotypical classes of RSA identified in [18]. A class can contain an arbitrary number of example cases. The classes are clustered into larger groups, the clustering criterion is based on the number of notches the heart rate measurement contains[2]. A notch is a smaller or larger dip in the heart rate pattern, as illustrated in figure 5.
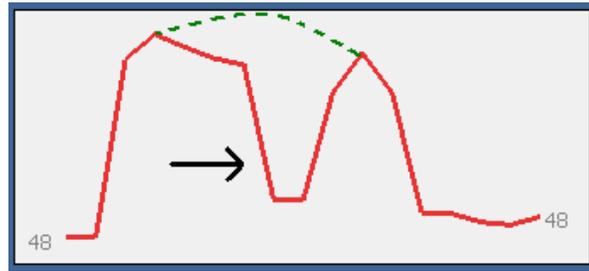

Fig. 5. The arrow marks a notch in the heart rate pattern. The rules concluded this is a notch due to the steep fall of the heart rate in the middle of the sequence. The dotted line marks a, would be normal, i.e. non-dysfunctional, RSA pattern.

A class is not limited to one group. An RSA class may end up in several clusters. Rules trigger new cases for notches. One rule indicates whether there is a change in the direction of the heart rate signal, i.e. starts falling after an inclination and vice versa. Another rule monitors the acceleration of any changes and a third decides if the change is big enough to be considered a notch or if it merely is a part of the naturally oscillation. The rules do also consider where the dips occur before they are concluded as notches. The number of notches for the entire sequence of heart rate measurements is calculated. This determines

---

[2] Stereotypical classes and their clusters may change whenever new knowledge from psychophysiological research is available.

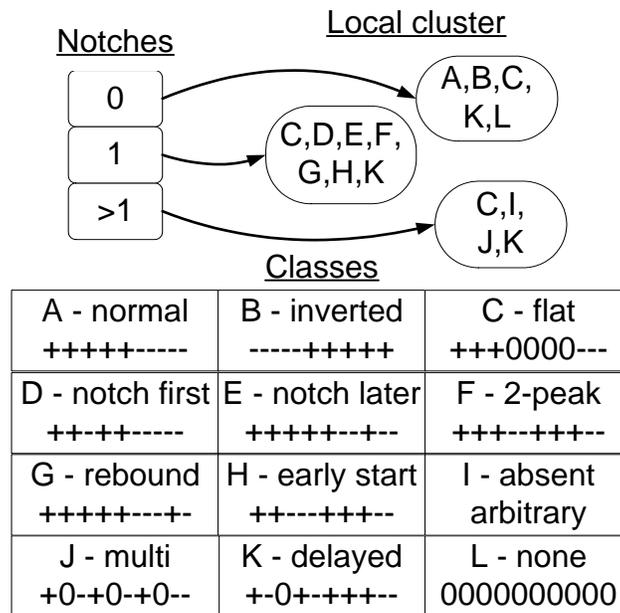which cluster of classes the matching procedure is to use. The current clusters are illustrated in figure 6.

Notches	Local cluster
0	A,B,C, K,L
1	C,D,E,F, G,H,K
>1	C,I, J,K

Classes

A - normal	B - inverted	C - flat
+++++-----	-----+++++	+++0000---
D - notch first	E - notch later	F - 2-peak
++-++-----	+++++--+--	+++--+++--
G - rebound	H - early start	I - absent
+++++---+-	++---+++--	arbitrary
J - multi	K - delayed	L - none
+0-+0-+0--	+-0+-+++--	0000000000

Fig. 6. The local clusters and the corresponding stereotypical classes of RSA. Each RSA class is described with its typical structure in the form of first order derivates. + is an increase in the heart rate, - a decrease and 0 is status quo. The RSA patterns are visualised with examples of sample sequences and corresponding frequencies in [18].

2.4 Case similarities

A new case is matched with stored cases by calculating the similarity of the heart rate measurements and the heart rate frequency spectrum. The new case is matched with all cases in all the classes of the local cluster, as can be seen in figure 6.

The frequency match calculates the distance between two frequency vectors by comparing the spectral density of each individual frequency. This is calculated throughout the entire length of the vector, i.e. all the elements in the vector. The power, i.e. amplitude, corresponding to each frequency is normalized against the maximum power (largest power value in the vector) in order to be in the range of $0 - 1$. The difference between the stored and the new case's frequency powers are weighted, and the difference is accumulated. The total difference between

the vectors is then normalized to become the similarity distance. A similarity distance of 0 represents two identical frequency vectors. A pseudo code follows:

```
similarity_distance = 0

/* for all frequencies */
FOR (frequency = 1 TO number_of_frequencies)
  /* Normalize the powers (range 0-1) for the stored case */
  IF (maximum_power > 0)
    power_storedcase = frequency / maximum_power_storedvector
  ELSE power_storedcase = 0
  ENDIF

  /* Normalize the powers (range 0-1) for the new case */
  IF (maximum_power > 0)
    power_newcase = frequency / maximum_power_newvector
  ELSE power_newcase = 0
  ENDIF

  /* calculate the weighted difference between the two
     frequencies and accumulate the differences */
  IF (power_storedcase > power_newcase)
    similarity_distance += (power_storedcase-power_newcase)*weight
  ELSE
    similarity_distance += (power_newcase-power_storedcase)*weight
  ENDIF
ENDFOR

/* normalize the similarity for the vector (range 0-1) */
similarity_distance /= frequencies
```

A sequence of heart rate measurements is not constant in length unlike the frequency vector. This is due to the variable time of an individual breath, and the heart rate measurements are coupled with the respiration cycle. The heart rate measurements are matched by their first order derivates. Changes within the heart rate pattern are easier to find if two vectors with first order derivates are used instead of the original measurements.

The derivates of the new case are interpolated to match the number of derivates in the stored case, if they do not contain equal number of samples. The distance is calculated for every pair of derivates and compiled to a normalised similarity distance for the entire heart rate sequence, much like the frequency similarity measurement.

The similarities of the measurements and the spectrum are merged to one similarity distance for the entire case, $similarity\ distance = frequency\ dist. \times derivate\ dist$. The cases are ranked based on the similarity distance, lower value equals closer match. The list is then purged from duplicate cases of the same

class, as there is often several example cases of the same class appearing in a sequence. As an example, a list of cases indicating the following RSA classes $3, 3, 3, 3, 9, 9, 1, 1, 1, 1, 2$ is transformed to $3, 9, 1, 2$.

The relevant cases, with the closest similarity, are finally presented to the user. The frequency and derivate similarity distances are also available to the user.

2.5 Case library

The case library consists of approximately 50 cases. These 50 cases represents the existing 12 stereotypical classes of RSA [18]. Each class is represented with several example cases. The examples are distributed that the more common classes has more examples than the rare occurring classes.

There is no adaptation of the cases in the system, since the solution domain consists (currently) of 12 stereotypical classes of RSA. The RSA classification tries to label the physiological measurements with a single integer, i.e. this solution domain is not suited for adaptation, and is typical for problem domains that are not well understood [19]. The system has several operational modes: classification, evaluation and learning. The system classifies RSA patterns in the classification mode. An evaluation of the system performance compared to a user is added to the classification in the evaluation mode. The learning mode accepts new cases to the case base, as well as deletion and modification of existing cases.

2.6 User interface

As mentioned in the introduction, one of the systems tasks is to serve as a research tool for researchers in psychophysiology. Hence, a windowed environment was chosen to display the measurements. The user can freely choose what measurements or parameters he/she wants to work with, as well as enabling different operational modes. A screenshot of the system is displayed in figure 7. The screenshot illustrates the complexity of classifying RSA.

The system is currently implemented in C++ as an application for the Windows platform. The application uses OpenGL to display graphics, as by doing this it is easier to port the application to other platforms in the future, due to OpenGL's OS independent interface.

3 Evaluation

This section contains a first evaluation of the RSA classification. The first evaluation was also the first time leading experts in the field of psychophysiology came in contact with the system.

The case-base was initialized with stereotypical cases produced by domain experts. The cases are described in [18]. The cases are supposed to cover all known
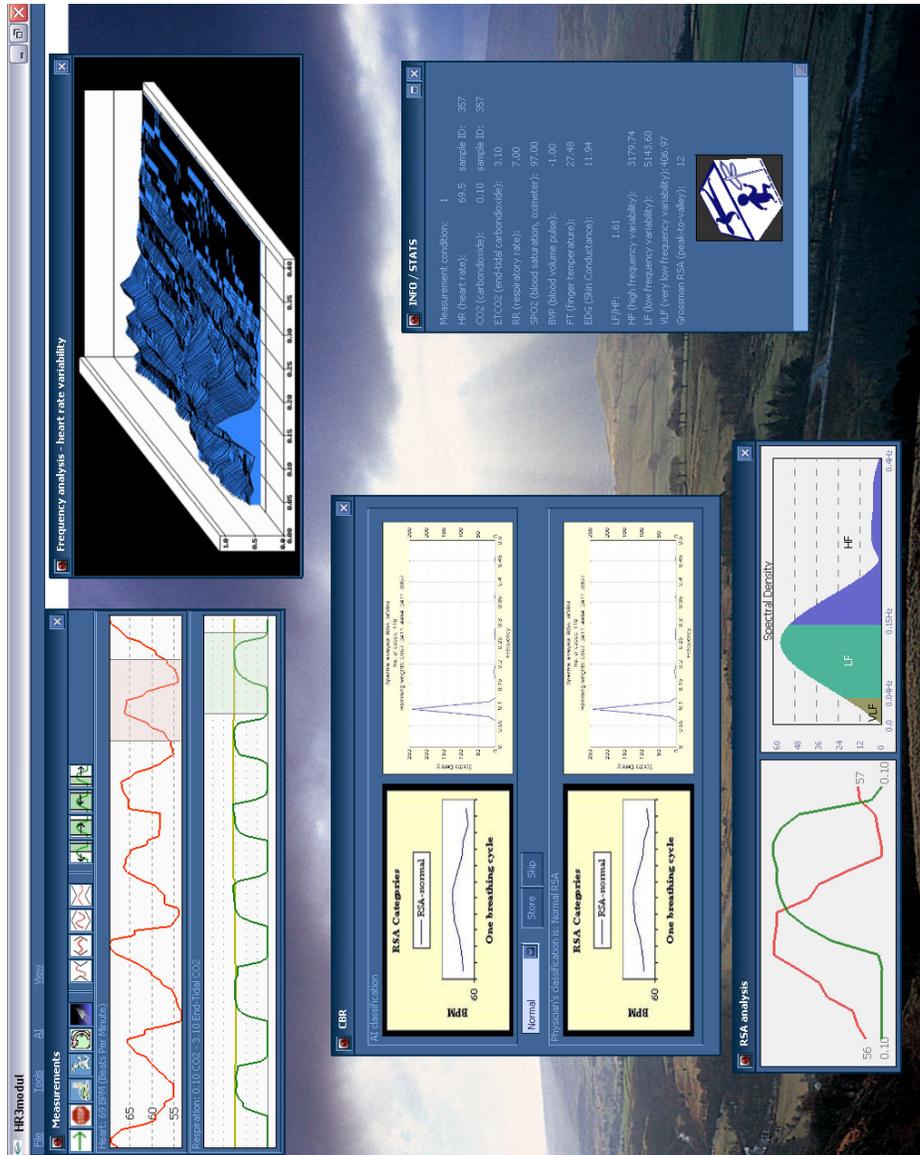
Fig. 7. A screenshot of the application christened HR3modul. HR3modul is a tool for classification of Respiratory Sinus Arrhythmia. The top left shows heart rate, CO_2, and end-tidal CO_2 measurements. Below the measurements is a classification, and at the bottom an RSA period that is being classified. The top right shows a frequency spectrum for each analysed RSA and below it shows additional information physicians/clinicians find useful.

classifications of RSA, i.e. cover the entire domain. Additional cases were also added to the case library. The additional cases belong to one of the stereotypical classes, and were added to facilitate an easier matching process. An example of an additional case is where the heart rate is constant during the entire respiration cycle, i.e., $\forall_i(s_i \in S : s_i = 0)$ after the conversion in equation 1. There exist no frequencies in a straight line.

3.1 Evaluation data set

A data set of approximately 100 pre-recorded measurements was used in the evaluation. The measurements were recorded in a cStress system, and were measured from a normal population of 17 year olds. Pre-recorded measurements are parsed and simulated in the HR3modul system as if they were real-time measurements streamed directly from hardware.
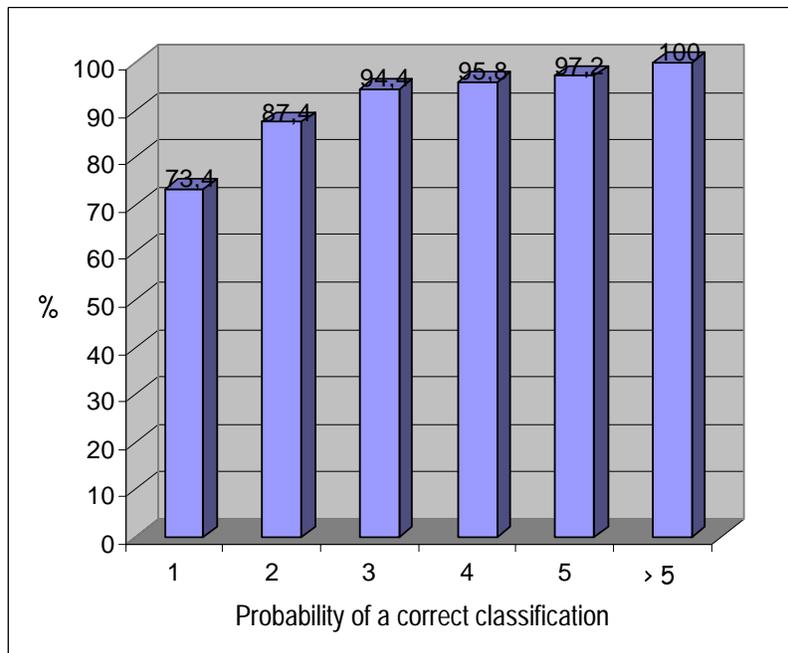

Fig. 8. Evaluation of the classification system. The columns represent the probability of an accurate classification, ranging from the same classification as an expert in the first attempt to the correct classification in the 5th attempt. All attempts beyond the 5th are summarised in the rightmost column.

3.2 Results

The evaluation was conducted with the help of the domain experts. Cases of special interest for an accurate classification were pushed to the case library. An evaluation mode was enabled when the case library contained enough example cases of RSA. The evaluation mode collects statistics of the accuracy of the classification system. The case library used in the evaluation consists of approximately 50 cases. The cases represent the existing stereotypical classes of RSA.

However, as this was the first time the physicians had an opportunity to view every individual RSA, i.e. the HRV per respiration cycle; new patterns of RSA were discovered. This invalidates the notion of total domain coverage by the cases in the case library, since the new RSA patterns do not fit into any of the stereotypical classes described in section 3. The new patterns are for the most parts deviations of similar existing patterns. A first hint from the physiologists points towards a reorganisation of the class structure; super classes with subclasses as an addition to totally new patterns.

Nevertheless, statistics were collected from the evaluation. A summary of the statistics are presented in figure 8. The figure represents the accuracy, i.e. similarity, of the classification system in a comparison with a domain expert. The leftmost column represents the probability that the first RSA class suggested by the system is the same as the expert would choose. The second column from the left represents the probability that the expert's choice of class is the same as either the first or the second RSA class suggested by the system. The rest of the columns proceed in the same manner, from left to right. All statistics of the similarities beyond the 5th suggested class have been summarised into the rightmost column.

4 Conclusions

We have presented a MMR design for the classification of RSA. The design uses two analytical stages of time series measurements from the heart and from exhaled air. The analytical stages process the time series measurements so they will conform into cases, for later similarity comparisons. A RBR stage limits the number of RSA classes that have to be considered in the matching, and a CBR stage makes a similarity match with the cases from the remaining RSA classes.

The MMR design for the classification of RSA seems to be reliable, as 19 out of 20 cases in the evaluation data set were among the three top most suggested classes. But the spread of the accuracy also suggests that it would be beneficial of using a differential diagnosis in the future, if the accuracy do not reach 100%. The evaluation also showed that even the experts benefit from the system, as they discovered new patterns of RSA while using the system.

The algorithm for calculating the similarity is sufficient for the time being. But some sort of improvement will probably be necessary as the case library grows. An interesting approach is to reduce the dimensions of the time-series features by for instance using the D-HS^T indexing [15], as it is specially suited for temporal time-series. Adaptation might also be a future issue as the case library is adapted for super classes and subclasses. A case could possible be adapted from a super class to fit a subclass or vice versa.

References

1. Isabelle Bichindaritz. Mnaomia: Improvingcase-based reasoning for an application in psychiatry. In *Artificial Intelligence in Medicine: Applications of Current Technologies*, pages 14–20. AAAI, 1996.
2. Isabelle Bichindaritz, Emin Kansu, and Keith M. Sullivan. Case-based reasoning in care-partner: Gathering evidence for evidence-based medical practice. In *Advances in CBR: 4th European Workshop*, pages 334–345. ECCBR'98, September 1998.
3. Tuan Pham Dinh, Helene Perrault, Pascale Calabrese, Andre Eberhard, and Gila Benchetrit. New statical method for detection and quantification of respiratory sinus arrhytmia. *IEEE Transactions on Biomedical Engineering*, 46(9):1161–1165, September 1999.
4. P. Grossman and J. Van Beek et al. A comparison of three quantification methods for estimation of respiratory sinus arrhytmia. *Psychophysiology*, 27(6):702–714, November 1990.
5. Ellis Raymond Bareiss jr. *PROTOS: A unified approach to concept representation, classification and, learning*. UMI Dissertation Services, 300 North Zeeb Road, P.O. Box 1346, Ann Arbor, Michigan 48106-1346, USA, 1988.
6. Phyllis Koton. *Using Experience in Learning and Problem Solving*. MIT Press, 1988.
7. Bernard Landis and Patricia M.Romano. A scoring system for capnogram biofeedback: Preliminary findings. *Applied Psychophysiology and Biofeedback*, 23(2):75–91, 1998.
8. Paul M. Lehrer, Evgeny Vaschillo, and Bronya Vaschillo. Resonant frequency biofeedback training to increase cardiac variability: Rationale and manual for training. *Applied Psychophysiology and Biofeedback*, 25(3):177–191, 2000.
9. Marek Malik and John Camm et al. Heart rate variability - standards of measurement, physiological interpretation, and clinical use. *European Heart Journal*, 17:354–381, March 1996.
10. Cindy Marling and Peter Whitehouse. Case-based reasoning in the care of alzheimer's disease patients. In *Case-Based Research and Development*, pages 702–715. ICCBR'01, July/August 2001.
11. Stefania Montani, Paolo Magni, Riccardo Bellazzi, Cristiana Larizza, Abdul V. Roudsari, and Ewart R. Carson. Integrating model-based decision support in a multi-modal reasoning system for managing type 1 diabetic patients. *Artificial Intelligence in Medicine*, 29:131–151, 2003.
12. Stefania Montani, Paolo Magni, Abdul V. Roudsari, Ewart R. Carson, and Riccardo Bellazzi. Integrating different methodologies for insulin therapy support in type 1 diabetic patients. In *Artificial Intelligence in Medicine*, pages 121–130. AIME'01, July 2001.

13. Markus Nilsson, Peter Funk, and Mikael Sollenborn. Complex measurement classification in medical applications using a case-based approach. In *Workshop proceedings: Case-Based Reasoning in the Health Sciences*, pages 63–72. International Conference on Case-Based Reasoning, June 2003.

14. Markus Nilsson and Mikael Sollenborn. Advancements and trends in medical case-based reasoning: An overview of systems and system development. In *Proceedings of the 17th International FLAIRS Conference*, pages 178–183, May 2004.

15. David Patterson, Mykola Galushka, and Niall Rooney. An effective indexing and retrieval approach for temporal cases. In *Proceedings of the 17th International FLAIRS Conference*, pages 190–195, May 2004.

16. P. Perner, T. Günther, H. Perner, G. Fiss, and R. Ernst. Health monitoring by an image interpretation system - a system for airborne fungi identification. In *Medical Data Analysis*, volume 2868 of *Lecture Notes in Computer Science*, pages 62–74, 2003.

17. Rainer Schmidt and Lothar Gierl. Temporal abstractions and case-based reasoning for medical course data: Two prognostic applications. In *Machine Learning and Data Mining in Pattern Recognition*, volume 2123 of *Lecture Notes in Computer Science*, pages 23–34, 2001.

18. Bo von Schéele. *Classification Systems for RSA, ETCO2 and other physiological parameters*. PBM Stressmedicine, Heden 110, 821 31 Bollnäs, Sweden, 1999.

19. Ian Watson. *Applying Case-Based Reasoning: Techniques for Enterprise Systems*. Morgan Kaufmann Publishers Inc, 340 Pine St, 6th floor, San Fransisco, CA 94104, USA, 1997.