Multiple Alarm Management with Self-Organizing Maps

Satu Tamminen[1], Susanna Pirttikangas[1], Seppo Nissilä[2], Vesa Pentikäinen[2], Kauko Väinämö[1], and Juha Röning[1]

[1] Machine Vision and Media Processing Group, Infotech Oulu
University of Oulu, Computer Engineering Laboratory, FIN-90401 OULU
[2] Polar Electro Oy, Professorintie 5, FIN-90440 KEMPELE

Abstract. A method for process control is considered. Abnormal states of the process have been identified with dynamic control limits separately for each measured signal. A self-organizing map (SOM) has been used to combine the information.
In this research, the methods were applied to the physical system of humans, which cannot be regulated like industrial processes, except in clinical conditions. The data were collected during the spring 1996 and consist of over eight weeks of physical measurements and diaries recorded in a home environment by four test subjects.
The research shows that this method can be used to monitor the system of a human being. Physical activities seem to dominate the results. This is understandable, because some of the chosen variables reflect immediately the effects of physical stress.
Keywords: process control, dynamic control limits, self-organizing maps

1 Introduction

The purpose of this research was to develop a simple and descriptive device for people interested in their own well-being. It helps to recognize the changes in physical parameters occurring during normal life. There is no ambition to substitute the doctor in diagnosing, but to help people to maintain good health.

The first step in process modelling was to define the normal and abnormal states of the system. The measured signals of the system can be analyzed with dynamic control limits in order to monitor their variability. When abnormalities are detected, a new binary variable can be composed to identify hazardous observations. In this application, dynamic control limits were used as indicators of fast changes in independent signals.

The calculation of dynamic control limits can be done on-line. For data reduction, old data are deleted when new data are entered in the system. A history of three days seems to be adequate for setting up the limits.

The monitoring of a process usually requires an analytical system model. In industrial applications this is often difficult to accomplish, and for a human being it is impossible, because the environment cannot be controlled and the person's routines frequently change in every-day life. Neural networks provide a nonlinear solution where no analytical process model is needed. The self-organizing map is an unsupervised method, needs no target values and provides a visualization tool for temporal process control. The areas of the map can be identified based on the content of each area.

For eight selected variables, dynamic control limits were computed and (0,1) indicators for normal and abnormal observations were obtained. The most descriptive indicators were chosen among a total of 16, and they were entered into the self-organizing map with the original signals.

The research shows that this method can be used to monitor the system of a human being. Since, however, everyday activities are not repetitive, it is too ambiguous to try to predict the future states. Physical activities seem to dominate the results. This is understandable, because some of the chosen variables reflect immediately the effects of physical stress. No general implication for mental stress was found, however. This was due to the fact that all the test subjects lived quite a steady life during the measurement period.

2 Data Description

The data for this research were obtained during the spring of 1996. Fourteen healthy middle-aged male volunteers collected their physiological data daily for eight to ten weeks. Four of them were selected for this first phase of the research. Their R-R interval and activity were measured continuously with an R-R interval recorder and an activity monitor during the daytime. A R-R series consists of the time spans measured between two R peaks in an ECG signal. Diastolic and systolic blood pressure along with body temperature were recorded three times a day by the subjects. The first measurements were made in the morning after waking up. The second measurements were made between 2:00 and 8:00 PM and the third in the evening before going to bed. The quality of sleep was evaluated at night. During the measurement period, all the subjects were living normal life. The variation in measurement times was due to unsupervised self-measurements.

Furthermore, the subjects filled in a diary, indicating their daily emotional states, such as fatigue, happiness, pain etc. The amounts of coffee, tea, cigarettes and alcohol consumed were also reported. The notes on the emotions during the daytime were made at 2:00-8:00 PM and those on the rest of the day before going to bed. The physical exercise and meal times were also recorded.

The measurement time of over eight weeks produced a very large data set. To make the signals compatible, the continuously measured variables were discretized by taking averages for one-hour spans simultaneous to the discretely measured variables. Thus, three values for each day resulted in data vectors of 170 observations or longer.

The data were analyzed with different statistical methods. This analysis and discussions with experts led to the selection of eight variables with high quality, stability and descriptivity. The others were discarded because of the long missing periods (several days or even weeks) and noisy or clearly erroneous measurements. The variables were diastolic and systolic blood pressure, the mean and the standard deviation of R-R intervals, activity, body temperature, weight and quality of sleep.

3 Dynamic Control Limits

The idea of using confidence limits of the expected mean in process control is not new. Control charts are an essential part of statistical process control, as they distinguish the normal variability of the process from abnormal, presuming the distribution of the signal to be normal [1]. Control charts have a middle line and two straight control lines defined on the basis of the confidence limits computed from at least four samples of the process, or else the result is not reliable. The process is under control if all measurements remain between the control lines.

If the process is measured continuously, a sliding technique can be used. If there are l variables, let x_{ij}, where $i = 1, 2, \ldots$ and $j = 1, 2, \ldots, l$, stand for the measured value at time i for the j:th variable. The limits for the value entering the system at time k ($k > 4$) are calculated from $x_{1j}, x_{2j}, \ldots, x_{(k-1)j}$. The second limits for the value at time $k+1$ are calculated from $x_{2j}, x_{3j}, \ldots, x_{kj}$ and so on. In this way, the limits are computed every time a new value enters the system and they evolve along with the signal, in other words they are dynamic.

In cases that are not predictable and smooth, e.g. certain industrial processes, the present value of the signal might be strongly dependent on previous values. Furthermore, if certain values of the process show better correlation than the others, a weighting procedure for the signal history is a way to approach the problem.

If w_{ij} stand for the weights, the expected mean $(\bar{x}_w)_j$ for the weighted signal is

$$(\bar{x}_w)_j = \sum_{i=1}^{k} w_{ij} x_{ij} / \sum_{i=1}^{k} w_{ij} \tag{1}$$

and the standard deviation $\sigma_{(x_w)_j}$

$$\sigma_{(x_w)_j} = \sqrt{\sum_{i=1}^{k} w_{ij} (x_{ij} - (\bar{x}_w)_j)^2 / \sum_{i=1}^{k} w_{ij}}. \tag{2}$$

The confidence limits can be formed as follows

$$\begin{aligned} UL_{(x_w)_j} &= (\bar{x}_w)_j + T\sigma_{(x_w)_j} / \sqrt{\textstyle\sum_{i=1}^{k} w_{ij}} \\ LL_{(x_w)_j} &= (\bar{x}_w)_j - T\sigma_{(x_w)_j} / \sqrt{\textstyle\sum_{i=1}^{k} w_{ij}}, \end{aligned} \tag{3}$$

where T is a constant defining the width of the limits, $UL_{(x_w)_j}$ is the upper limit, and $LL_{(x_w)_j}$ is the lower limit.

If the signals contain a lot of irregularities, a suitable filtering structure is needed. The dynamic control limits are not defined based on the original signal, but on a particular flattened signal. The present value of the flattened signal is obtained from the history of the dynamic control limits by defining the percentage between the original signal and the previous upper or lower limit. If the original signal exceeds these limits by this percentage, the signal is flattened. In this way the erroneous and irregular values of signals do not affect the adaptation of the dynamic limits too much. On the other hand, if there are known limits the signal has to conform to, this *a priori* information can be noted in defining the dynamic limits.

When the dynamic control limits are ready, we have an upper and a lower boundary value for each signal. Every time a signal is beyond these values, there might be something wrong with the system. From now on, this overdrafting will be referred to as an alarm.

With this approach, two (0,1) indicators are established; one for a lower alarm and one for an upper alarm of the signal. There are thus two more variables for each signal.

4 Self-Organizing Maps

The self-organizing map (SOM) provides a data-driven approach to process monitoring and modeling. The method has the advantage that little or no *a priori* information is needed about the system domain and it is not necessary to define the process model analytically [2].

The monitored process should be static or else the map is not able to visualize the data correctly. A problem will arise especially if the future measurements have some kind of trend or gradual development.

The theory of SOM is utilized here only briefly. According to Kohonen [3], a SOM converts nonlinear statistical dependencies between high-dimensional data into simple geometric relationships, usually on a two-dimensional grid. It can therefore serve as a clustering tool for problems where traditional methods are not efficient enough.

Each neuron of the SOM is represented by an l-dimensional weight vector. The neurons of the map are connected to adjacent neurons by neighborhood relations, which assign the topology of the map. Hexagonal or rectangular topology is usually used. In the basic SOM algorithm, the topology and the number of neurons are fixed at the beginning.

The weight vectors of the map can be initialized randomly by setting the vector components to random values that are evenly distributed in the area of corresponding data vector components. Linear initialization can also be used. In that case the weight vectors are initialized in an orderly fashion along the two-dimensional subspace spanned by the eigenvectors of the input data.

As with other neural networks, the SOM will also yield poor results if erroneous data are used. Therefore, the input data must be pre-processed carefully. The data variables must be quantitative, e.g. symbolic data should be transformed into a suitable form. If the scales of the input variables are very different, the variables should be normalized, which gives all the variables an equal influence in the training phase of the SOM. The SOM is able to handle missing values, but if a large number of components is missing, it will affect the reliability of the map.

The SOM is iteratively trained, and it tends to approximate the probability density of the data. Let $X_{n \times l}$, where n is the number of measured values and l is the number of variables, denote the whole data set. One sample vector is the i:th row of the matrix X. At each training step, one sample vector is randomly drawn from the input data set and the similarities between this vector and all the weight vectors are computed. The best-matching unit (BMU) is the map unit whose weight vector is closest to the sample vector. When the BMU is found, the values of BMU and its topological neighbors are updated.

The quality of mapping can be measured using the average quantization error, which is the average distance between the input vectors and the corresponding BMUs. The accuracy of the results can be visualized with independent quantization errors for each observation. For better understanding of the results, the map can be characterized by labeling the data units. If no labels are available, inspection of the weight vectors and the clusters of the map may help in characterization [4].

The unified distance matrix can be used for structure visualization of the SOM. The matrix consists of the distances between the map units on a two-dimensional map, and the matrix can be represented by a grey-level image. The lighter the color between any two map units is, the smaller is the relative distance between them. Component plane representation visualizes relative values of the weight vectors separately for each input variable.

For more information about self-organizing maps, the comprehensive book of Kohonen is recommended [3].

5 Application Results

5.1 Health Indicators

There are many events that influence the behavior of the selected signals of the test subjects. The calculation of dynamic control limits is simple, and the limits give us guidelines for research on the system. In this alarm system, the critical points of the data are not as important as the drastic change. But if the signal remains longer in a hazardous area identified based on *a priori* knowledge, a continuous alarm will be given. For some variables these hazardous areas are easy to define (e.g. temperature), but for most of the variables the definition is highly dependent on the subject.

As mentioned above, eight variables were chosen for the study of the dynamic behavior of human physical structure. After filtering the artifacts, the idea of dynamic control limits was applied to the data. The limits were calculated using the Matlab version 5.3. and a Sun Ultra 10 workstation.

The variables modeled were diastolic (DBP) and systolic (SBP) blood pressure, the mean and the standard deviation of R-R intervals (RR and RRSTD respectively), activity (ACT), weight (WEIGHT), body temperature (TEMP) and quality of sleep (QS%). They were assumed to describe the person's physical condition in the best possible way. Both SBP and DBP were chosen, because they are affected by different physical states [5]. Figure 1 shows the dynamic control limits for the diastolic blood pressure of one person. It can be seen that both the extreme values of DBP and the rapid changes are recognized. Later on the alarms are marked as _ah if lower and _yh if upper alarm is considered.

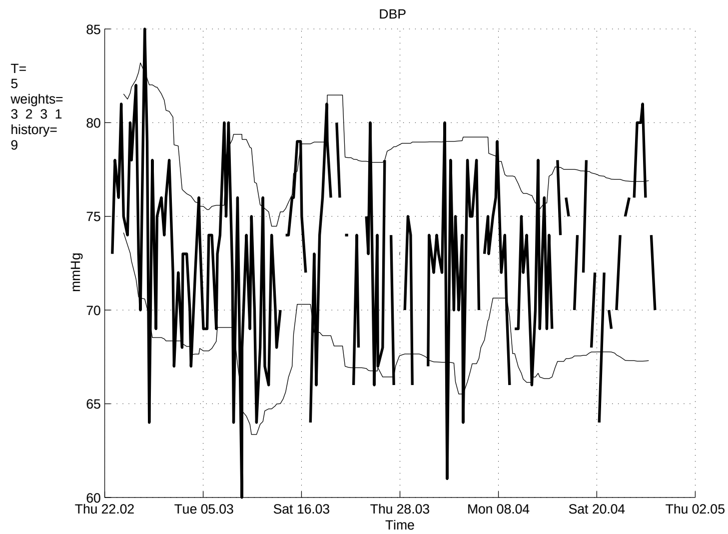

Fig. 1. The dynamic control limits for diastolic blood pressure and the original signal. Time is shown on the x-axis and mmHg on the y-axis.

The sliding history (earlier referred to k) was chosen to be nine for all variables. This means that three days of measurements were taken into account in building the present value of the dynamic control limit and all the previous values were ignored.

The weighting procedure for DBP was defined to follow the person's diurnal blood pressure rhythm [6]. The weights were chosen to strengthen the previous value and the measurement made at the same time on the previous day. The value between these two was also weighted. A similar weighting procedure was used for the other variables, too.

Because of the long measurement period, the subjects were not able to carry out the measurements continuously. The data therefore contain a great deal of missing values. This is also due to problems with the measurement devices. The missing values were replaced by the average of the earlier values of the signals within the sliding history.

5.2 Combining Alarms

The combination of alarms is important, because it helps to attribute the interactions between the variables to the subjective records and to control the whole system simultaneously.

Self-organizing maps have been used to monitor many different industrial processes [2]. As far as the human being is concerned some substantial differences can be found compared to industrial processes, however. First of all, the conditions in normal life are not homogeneous and not always predictable. The physical system of a human being is not stable, and many normal every day activities may affect the measured signals considerably. SOM visualizes this complex system into a simple and descriptive picture. A Sun Ultra 10 workstation and SOM_TOOLBOX [7] with Matlab 5.1 was used for the training and visualization of the self-organizing map.

For the four subjects, the data were divided into a training period and a test period. The test period covered the last four weeks from the data set, and it was used to confirm that the contents of the areas remain similar when new data is fed into the map.

The size of the map was 30x21 units and linear initialization was used. After the training, the map was compared with the diary and other records in order to find a correlation between the results and daily activities or subjective feelings.

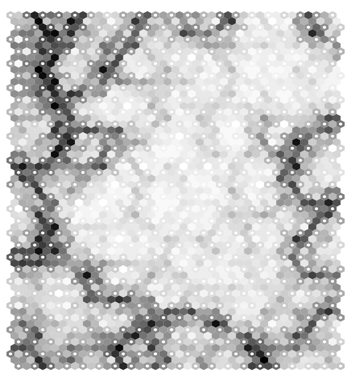

Fig. 2. Self-organizing map for an alarm combination.

Figure 2 contains one large light-colored area as well as some smaller areas. The large area stands for the observations with no alarms and the small areas for different alarm combinations. The content of the areas can be seen in Figure 3, where the component planes of each variable are shown.

A small R-R average and low heart rate variability overlapped for several measurements, most of which had been recorded during or right after physical exercise.

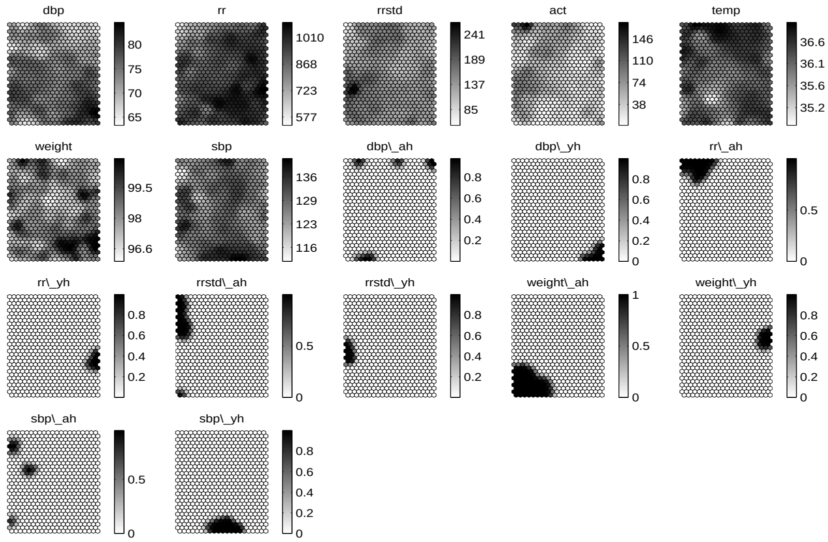

Fig. 3. Component planes of SOM.

Body temperature may have biased the results, because it can be easily measured wrongly, and actual temperatures can therefore be higher than in the data. This would explain some of the odd results.

Figure 4 visualizes the quantization errors for the training set. The bigger the circle, the poorer is the mapping result. Poor results in the alarm areas are due to the high variability between the other measurements in that area. The normal area has a large amount of training data points and therefore distinct separation of the data points.

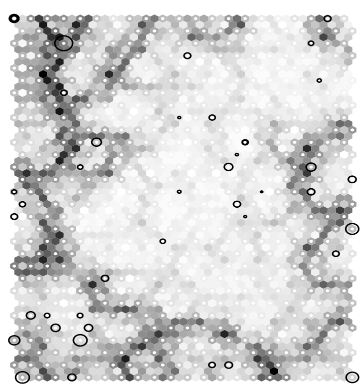

Fig. 4. Quantization errors for the map.

The SOM_TOOLBOX provided a tool for tracking the system's state in time with the trajectory. If abnormal states do not occur too often, the tool is very useful, but for human beings it only jumps around the map. It is impossible to predict where the system is going to shift next.

The other test subjects showed some results similar to those described above. There were some areas where the measurement time dominated the map, and exercise also had an impact on the results.

6 Discussion

A sliding history of three days seems to be a good choice for the calculation of dynamic control limits. If the history is longer, the borders are smoother but the confidence limits wider. If the history is shorter, the confidence limits are too narrow. The weights chosen reflect the local variability of the signals. If there were no weights, the control limits would be far too wide.

Specification of the suitable percentage of smoothing the signal in order to identify the erroneous and extreme measurements has to be done individually for each person and variable, because the ranges are dissimilar for distinct persons and different variables.

The dynamic limits are fairly adaptive, and they therefore provide no alarms for trends. In this way, important aspects of the person's general health may go unnoticed. For the subjects at hand, this is not a problem, because the measuring periods were quite stable, which was investigated carefully.

The method has only been applied to four subjects so far, and the results cannot be generalized to larger populations. But the idea can be utilized straightforwardly for new subjects. If a totally different application is concerned, some effort is required for variable selection and parameter adjustment.

The role of physical exercise is remarkable in the self-organizing map. It is difficult to extract other abnormalities from the data when everyday activities dominate the measurements.

It was an interesting finding that, for the investigated subjects, physical exercise affected the measured signals quite similarly. Furthermore, the influence remained for a while after the exercise.

Based on the findings on these four subjects, it is clear that no map can adequately apply to all the subjects, because the reactions caused by different stress states occur individually.

There are some differences between morning and evening measurements. Afternoon and evening measurements tend to appear in the same areas, however. The reasons why the measurement times vary are different for the subjects. One subject has a map that differentiates the mornings from the other measurements by DBP and SBP. The subject had no high values for blood pressure in the mornings. The maps of the other subjects did not show similar incidents, partly because their overall variability in blood pressure was larger. The other subject had a map which was differed because of activity: he walked almost every morning during the calculation of heart rate and activity. His heart rate was also high. Another explanation might be that his alcohol consumption was really high.

One subject had a low percentage of quiet sleep during the test period. The cause for this was not found, but it affected the daytime coffee consumption. The others did not have similar incidents, but the whole measurement periods were more or less stable.

An analysis of the quantization error in the map for different measurements showed that the biggest errors were in the alarm areas. This was due to the fact that the measurements with alarms were not so homogeneous as the normal measurements, and small alarm areas showed quite large variability. The same tendency was also found during the test periods. When visualizing the response surfaces of the map, notable quantization errors also occurred whenever the winning neuron of the map was not clear.

The high number of missing values was a problem for most of the subjects, and extrapolation of the signals was performed. This procedure does not give right values for the data and may bias the results, but at this point of research it helped to develop the tools for analysing the results.

References

[1] H. Kume. *Laadun parantamisen tilastolliset menetelmät.* Metalliteollisuuden keskusliitto, Helsinki, 1989. in Finnish.

[2] Esa Alhoniemi, Jaakko Hollmén, Olli Simula, and Juha Vesanto. Process Monitoring and Modeling Using the Self-Organizing Map. *Integrated Computer Aided Engineering*, 6(1):3–14, 1999.

[3] Teuvo Kohonen. *Self-Organizing Maps.* Springer-Verlag, Germany, 1997. (Second Extended Edition).

[4] Juha Vesanto. SOM-Based Data Visualization Methods. *Intelligent Data Analysis*, 3(2), 1999.

[5] Väinö Turjanmaa, Martti Tuomisto, Mats Fredrikson, Seppo Kalli, and Arto Uusitalo. Blood Pressure and Heart Rate Variability and Reactivity as Related to Daily Activities in Normotensive Men Measured with 24-h Intra-Arterial Recording. *Journal of Hypertension*, 9(7):665–673, 1991.

[6] J. Staessen, C. J. Bulpitt, E. O'Brien, J. Cox, R. Fagard, A. Stanton, L. Thijs, S. Van Hulle, G. Vyncke, and A. Am. The Diurnal Blood Pressure Profile. A Population Study. *American Journal of Hypertension*, 5(6):386–392, 1992.

[7] http://www.cis.hut.fi/projects/somtoolbox/.