Genie of the Net: A New Approach for a Context-Aware Health Club

Susanna Pirttikangas[1], Jukka Riekki[1], Jouni Kaartinen[2], Jari Miettinen[3], Seppo Nissilä[3], and Juha Röning[1]

[1] P.O.Box 4500, FIN-90014 University of Oulu
{Susanna.Pirttikangas, Jukka.Riekki, Juha.Roning}@oulu.fi
[2] VTT Electronics, P.O.Box 1100, FIN-90571 OULU
Jouni.Kaartinen@vtt.fi
[3] Polar Electro Oy, Professorintie 5, FIN-90440 KEMPELE
{Jari.Miettinen, Seppo.Nissila}@Polar.fi

Abstract. This paper describes a context-aware Health Club for cyclists. The application is being built on Genie of the Net, a general software platform for context-aware applications. Genie of the Net is being developed by building systems for selected applications and defining a general architecture on the basis of the gained experience. Work on context recognition started with collecting data from cycling exercises. Data mining was used to find interesting and useful contexts from the collected data. The contexts were utilized in analyzing the exercises. Different techniques for visualizing the exercise data, recognized contexts and the results of the analysis were studied as well. A user interface was also built for testing visualization techniques.

1 Introduction

As technology is developed and various information services are becoming common, a large amount of data from everyday situations is becoming available. However, utilizing this technology and the data it produces in an efficient manner is an open research problem. There are several active fields of research solving this problem, such as ubiquitous computing [42], [43], context-aware computing [16], [36], [33], wearable computing [3], [29] and pervasive computing [11], [28]. Researchers in these overlapping fields try to develop services that can be used mostly without guidance or input by the user herself, and that are discreet and easy to use. Recognizing the context of the user is an essential part of these services.

Context recognition offers interesting challenges for data miners, since they are already familiar with suitable data processing methods. According to Mannila [26], data mining is about statistics, knowledge engineering, databases and visualization. These are also keywords in context-aware computing, and hence, the main focus of this paper.

Developing a context-aware application starts by specifying the application (equipment, services to be offered, etc.). Then, data is collected and analyzed to find out what kinds of contexts can be recognized, and which of them are useful in the application. Finally, methods for recognizing the selected contexts are implemented. This is clearly an iterative process, since every step consolidates the expertise of the field of research, and one might want to return to an earlier step as knowledge increases. Actually, this is the basic approach to data mining problems [31]. Hence, this approach for developing context-aware applications might be called context-aware data mining, or, more generally, ubiquitous data mining.

This paper describes a context-aware agent-based system, Genie of the Net, with an application scenario. In the next section, reported data mining methods for context recognition are discussed. The third section presents the Genie system. Genie collects information from sensors and databases, recognizes context based on this information, chooses relevant actions to serve the user on the basis of the recognized context, and performs the chosen actions. Genie is a general software platform for context-aware applications serving a mobile user.

Section 3.1 presents a Health Club application based on the Genie architecture. In this application, the Genie system offers the user services related to cycling exercises; services for planning an exercise schedule, creating instructions for exercising, instructing him or her during exercising, analyzing data collected during exercising, and visualizing the data and the results of the analysis.

Various traditional data mining problems and context-awareness have been studied in earlier projects [38], [39], [23], [37], [40], [21], [25], [24], [20], [14], [17], [19], [30], [27], [13], [6], [15]. There are a lot of challenges in developing the Genie architecture, but in this paper, the research approach is to develop context recognition methods for selected applications.

Section 4 describes the data that was collected for the development of context-aware services. Heart rate, cadence (cycling frequency), and the bike's speed were measured during cycling exercises. The height profile of the route was also available. The work involved in analyzing the collected data is presented. The data was visualized to get familiar with it. Clearly erroneous values were removed from the data in the preprocessing stage (section 4.2). The data was classified on the basis of the route's profile declinations (section 4.3). Different analyzing services and a user interface for evaluating cycling exercises are presented in section 4.4. Conclusions are drawn in the final section.

2 Related Work

The Context Toolkit [4], [33] developed at the Georgia Institute of Technology, is the most closely related to Genie. The Context Toolkit uses context widgets to communicate between the environment and the applications, while Genie uses agent technology for the communication. At Georgia Institute of Technology they have implemented a Meeting widget to identify several persons within the neighborhood of the user. The information for the toolkit is produced using, for example, speech recognition software, machine vision machinery, microfones, Active Badges [41], etc. The sensor data is filtered and interpreted, but for more sophisticated applications, machine learning and context history have to be catered for.

Microsoft research uses artificial intelligence to aid computer users (Tip Wizard in Excel), and they try to find new ways to improve the usability of their operating system and web-based services [8]. The Decision Theory and Adaptive Systems group (DTAS) is developing a Bayesian-based system to assist the user without interfering. Information about the users is gathered, and similar users are clustered to suggest personalized services. The group mentions that adaptive classification methods should be developed, so that the whole database would not have to be loaded whenever new users are introduced.

Schilit et al [34], [1] have studied context-aware applications, and they divide these applications into four different categories: proximate selection concerning user interface issues, reconfiguration for the components and connections of the system, contextual information and commands, and context-triggered actions in a simple rule database. In the data mining point of view, the most interesting part is context-triggered actions that consist of information about the user, that is, identification, event type and location. When an event is detected, the application carries out a predefined action. The problem is how to create useful actions, maintain what is important, and delete futile and old information.

Classification of data into different contexts has been studied thoroughly by Schmidt et al [35], [22]. They had special equipment, a sensor box, for data gathering and they followed a typical data mining framework in context recognition. Several simple actions were performed and the data was stored and analyzed off-line. They calculated different statistics from the data and tried to find features that discriminate different situations. In their analysis, they used self-organizing maps and, based on this analysis, they formed a rule database for context recognition.

In wearable computing, the user carries all the sensor equipment on her person instead of using environmental sensor information, but this does not exclude the utilization of both the affective state of the user and the context of the environment [32]. Picard and Healey speak about the importance of labeling the situations accurately for analysis, and they have tried to create physiological fingerprints for the user's affective states [29], [10]. They noticed differences in physiological signals on different days and used Fisher linear discriminant projection [5] and leave one out test for discriminating between the different states.

In sports, data mining has been utilized at National Basketball Association (NBA). Advanced Scout is a PC-based data mining application developed at IBM [12]. The preprocessing of raw basketball data is automatically done by the program. Advanced Scout interprets patterns from the data, and allows the user to relate patterns to video tape.

Most of these context recognition tasks are very simple, including only the location context of the user, and they do not utilize the strength of data mining methods. The research by Schmidt and Microsoft offer an exception, maybe because they have enough data available. All in all, there are a lot of untapped possibilities for taking advantage of data mining methods in context recognition.

3 Genie

Genie collects information from sensors and databases, recognizes context based on this information, chooses relevant actions to serve the user on the basis of the recognized context, and performs the chosen actions. The Genie of the Net architecture is presented in Figure 1.

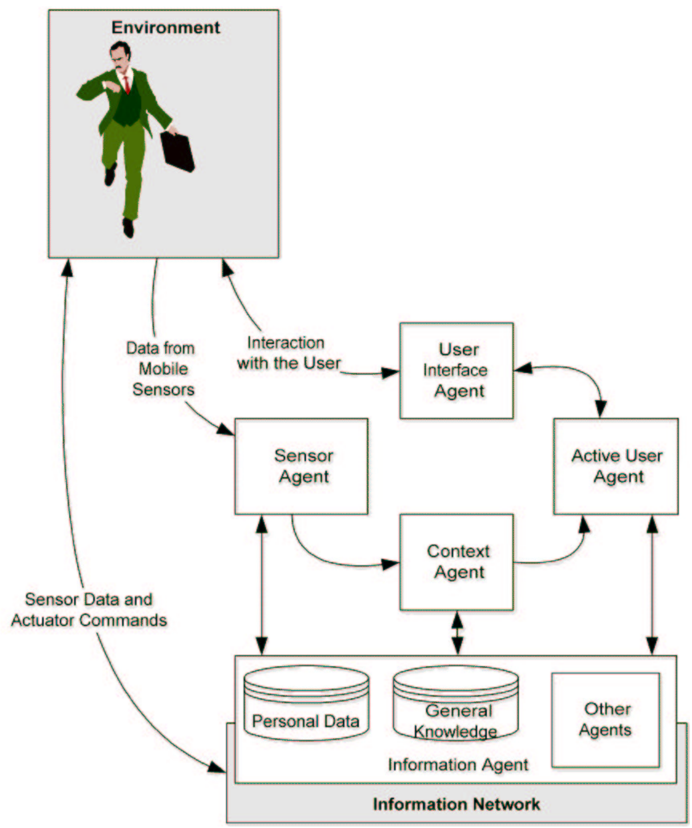

Fig. 1. Genie of the Net architecture [17, Fig. 1].

The Sensor Agent collects and preprocesses data from sensors. The Context Agent recognizes the user's context from this data and the other available information. The Information Agent provides information to the other agents and stores the sensor data and a description of the recognized context for later use. The Active User Agent exploits the recognized context in reasoning actions serving the user. The agent can, for example, request the User Interface Agent to present information to the user. In addition to scanning the context, the User Agent bases its reasoning on a user model containing different kinds of information about the user: name and other basic data, access rights, habits, plans, etc. Furthermore, the User Agent cooperates with the other agents. Although actuators are included in the architecture, the emphasis is on information management.

Although knowledge representation is an important part of the Genie architecture, it is not discussed here in more detail. The software components, that is, the agents communicate using a common language that contains a representation for all concepts (and their relations) shared by them, including those describing context.

3.1 Health Club

In the Health Club, the Genie system guides the user in a cycling exercise. A typical scenario is as follows: Before exercising, the user plans an exercise schedule at her terminal and possibly also outlines more detailed instructions for each exercise. Genie automatically presents a calendar containing the exercise schedule to the user in the specified context and reminds about the forthcoming exercise. Before the exercise, the user checks from her terminal what kind of exercise she has planned to perform. During the exercise, a heart rate monitor records her heart rate, cadence, and the bike's speed. After the exercise, she goes to her terminal and loads the data into the system. The system recognizes the context history of the exercise based on the collected data and other available information (e.g., the height profile of the route). The system analyzes the exercise and presents the collected data, the context history and the results of the analysis to the user. The user might be advised during the exercise as well, although the current hardware does not support such a service.

The Health Club system actually contains several subsystems, the user's Genie being one of them. The user (i.e., Genie) does not need to interact with other systems, but the information provided by them enables a larger set of services. Additional information about, for example, the route cycled, enables different methods of analyzing the exercise. Sharing exercise information with other users allows comparisons. Furthermore, sharing calendars makes it possible to plan group exercises.

4 Methods

4.1 Data

The development of the Health Club application was initiated with the collection of the data. A strict route was formed where one cyclist exercised. The route was 11.5 km long and consisted of about four kilometers of uphill, about four kilometers of downhill and the rest was classified according to the profile as level ground. The data were collected in September 2000 and they contain 9 identical exercises in one direction along the route and 4 exercises in the opposite direction.

The measurement device was a heart rate monitor, Polar XTrainer Plus[1]. The Heart rate of the cyclist, the cadence and the speed were measured during a training period with a sample frequency of 5 seconds. The cyclist pushed an interval time button at certain points of the route. An example of an exercise is presented as a time series in Figure 2.

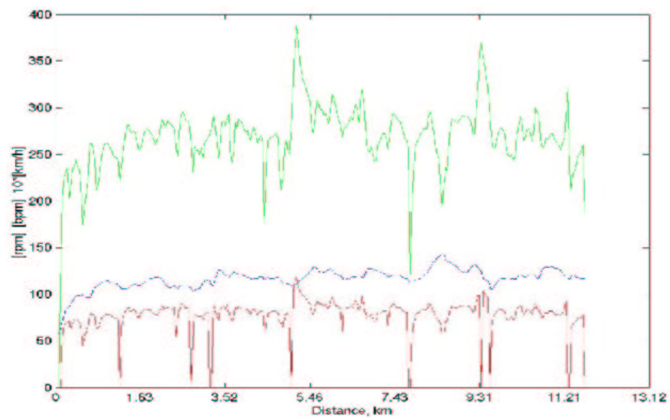

Fig. 2. Data collected during a cycling exercise. The curves show, from top down: velocity, heart rate and cadence, and the corresponding units at the left are 10*km/h, beats per minute and rounds per minute.

One goal was to resolve whether it is possible to identify changes in the declination of the ground using these three measurements. The system would then be able to automatically instruct the user, for example, allowing a higher heart rate for the cyclist when going uphill. The contexts *cyclist going uphill*, *cyclist going downhill* and *cyclist on level ground* could be used in off-line analysis if an accurate route profile is not available. They would also help in possible online guidance for the cyclist or as additional telemetry information for coaches during a competition. Of course, the measurement system could be augmented by extra sensors, but the cost of the system was to be low.

It is also essential for the cyclist to study the data collected from the exercises herself. If Genie offers analyzing services, it is important to find suitable ways of presenting the results, hence, visual data mining has to be considered.

Another goal was to allow several users to share each other's exercise data. If an exercise is done on the same route, two exercises can be compared. The exercise data needs to be analyzed and synchronized for comparison. One could also compare her own exercises. There are several simple ways to analyze different exercises: minimum values and maximum values for each variable, mean values, values between different thresholds, differences in distances at certain time points, etc. Trend information for each variable would also be interesting.

[1] This product is a trademark of Polar Electro Oy.

For example, if the cyclist's speed in an exercise has a downward trend, her technique might be wrong.

4.2 Preprocessing

The preprocessing stage was performed to get familiar with the data and the special area of cycling. Clearly erroneous values were removed and the data was prepared for further analysis.

First a thorough statistical analysis of the data was performed. All the exercises were done with a light moderate intensity zone, meaning that the heart rate of the cyclist was between 60-70 % percent of the person's maximum heart rate most of the time. Since there are only three variables, it was easy to visualize the data in three-dimensional space, but as the correlations between heart rate and speed are not very high due to the cyclist's interactions with other people or due to traffic lights, altitudes or wind, etc, in the route, the result was a large cluster that was rather challenging to interpret.

Differences were calculated from these three variables. This led to a situation where a scatter-plot was not adequate for visualization, and self-organizing maps were used to visualize this higher dimensional data.

There were some erroneous heart rate values, because if the cyclist had to stop for some reason and move away from the bicycle, the sensor signal did not reach the measurement device. In all such cases, the device stored zero values for heart rate, and linear interpolation was performed for these erroneous values. If speed was zero and cadence was not, the cadence was set to zero, too.

A profile and coordinates of the route were acquired from the map service of the city of Oulu, and this information was used in the classification of the measurements to different geographical attributes. It was noticed that when the distance was calculated from the speed, which had been measured only once every 5 seconds, there was a cumulative error in the location of the cyclist. Therefore, the measurements had to be synchronized with the map coordinates according to the cyclist's own interval markings and the known positions on the map. This adjustment is not very reliable either, because the cyclist was sometimes not able to push the interval button at the specific location (other objects on the route), and pushed the button too late.

The route was classified as downhill, uphill and level ground according to the angles of the z-axes on the map in monotonous acclivity and declination. To reduce the error in the distance calculation, the most frequent class within 20 meters was assigned to each observation. The turns on the route, where it was assumed the cyclist had to use brakes, were also calculated and marked. All in all, six different classes were formed for the route: downhill, level, uphill, downhill and turning, level and turning, and uphill and turning.

4.3 Classification

Classification aimed at distinguishing different ground declinations for each observation. The target classes were obtained in the preprocessing stage (downhill, uphill, level ground, etc.). This way, without any additional measurement device, the cyclist would know if the reason for a rising heart rate would be an uphill and not the consequence of wrong tactics, for example. Self-Organizing Maps (SOM) [18] were used for unsupervised classification. Multilayer perceptrons [9] were used for supervised classification.

Unsupervised Classification There were several different approaches for the derivation of suitable features for the SOM. The large scatter-plot cluster encouraged us to perform histogram equalization for the variables (speed, cadence, heart rate and the differences). Histogram equalization is an image processing technique used to improve the contrast of an image [7].

In another experiment, the variables and the differences were fuzzified according to the histograms and second-order statistics. The cyclist is probably not so interested in the exact value of the heart rate but, depending on the type and the goals of the exercise, but in whether it is high or low enough. We used three different kinds of membership functions for fuzzification: triangular, z-shaped and s-shaped [2] for average, low and high values, respectively.

In all the experiments, the inputs were labeled according to the route profile classes and fed into the map. It was assumed that the information in the inputs would be enough for profile classification. For example, if heart rate rises and velocity goes down while pedaling, the SOM would cluster these observations to uphill.

The self-organizing maps formed with these different kinds of input data did not show any significant clusters for the observations and profile classes. Figure 3 shows a SOM which was taught with inputs normalized using histogram equalization. The classes are 1, 2, 3, 4, 9 and 12, meaning downhill, level, uphill, downhill and turning, level and turning, and uphill and turning, respectively.

Supervised Classification A multilayer perceptron was used for supervised classification. The inputs for the perceptron were the current values of the variables (speed, cadence and heart rate) and the observations within a certain history window. Several different lengths of history windows were tested.

The results with the self-organizing maps led to a different choice for the target classes. First the downhills and uphills were divided into steep and mild downhills and uphills, and multilayer perceptrons with different sizes were taught with Levenberg-Marquardt optimization [9].

Finally, only three target classes, uphill, downhill and flat ground were chosen as targets. The variables were normalized to the scale of [-1,1]. It was found that the uphill observations are most distinguishable. With a three-layer perceptron and a history window of three observations, 42 percent of the downhill observations, 24 percent of the level ground observations and 82 percent of the uphill observations were classified correctly in the training set. These results indicate that there is potential in recognizing the different ground declinations, but the classification requires more work.

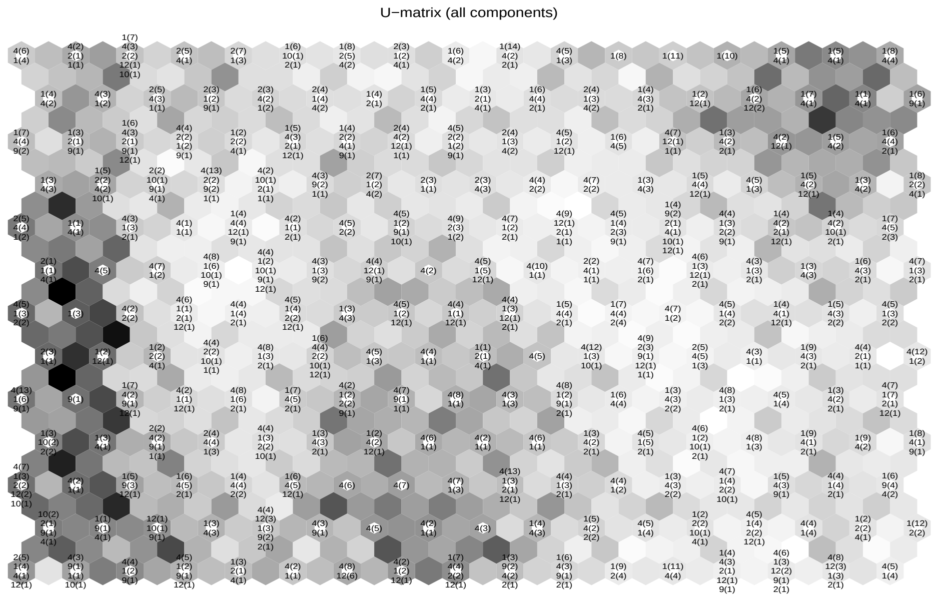

Fig. 3. A self-organizing map for route declination clustering. The labels are profile declinations and the number of similar observations in the same neuron are in parentheses.

4.4 Analyzing Services and Visualization

In the case of a specific application, as in data mining problems in general, help from specialists is essential. In the development of the Health Club application, discussions were held with coaches from different sports, particularly those who work with cyclists.

The season for a cyclist consists of four different phases; improving and maintaining basic condition, building up for competition, the competition period and a transitional recovery phase. Within all these periods, there are certain types of exercises to do; those that improve explosive quickness, that is, interval and uphill exercises, those that improve aerobic fitness, those that help the body to recover, etc. Clearly, the Health Club could offer different analyzing services for different types of exercise.

The study of different analyzing services and visualization techniques was initiated by building a graphical user interface. The interface was made in a Matlab environment, because it is easy to try different kinds of appearances, and to test and modify the interface. The interface and services will be implemented into Genie of the Net later.

In the user interface, the information from data mining could be used in making a plan for an upcoming exercise or there could be a predefined plan depending on the time of the season. After the exercise, the user would gain explanations about what went right and what went wrong in the exercise. The cyclist would also get a plan to improve from weaknesses and avoid observed failure. If the cyclist breaks down at the end of an exercise, she would have instructions for a slower speed in the beginning of the next exercise in the same route, for example. This work is under development.

Later on, when the classification for the measurements is available, this information can replace the route profile in the user interface. The preprocessed measurements (fuzzy variables) can also be shown in the user interface.

The exercises are ordered according to the date and the type of exercise in the database. The user can pick up an exercise and view it in the interface (figure 4), where the statistics of the exercises are shown in addition to the time series plots of the measured variables. If they are available, the route coordinates and profile are shown, also.

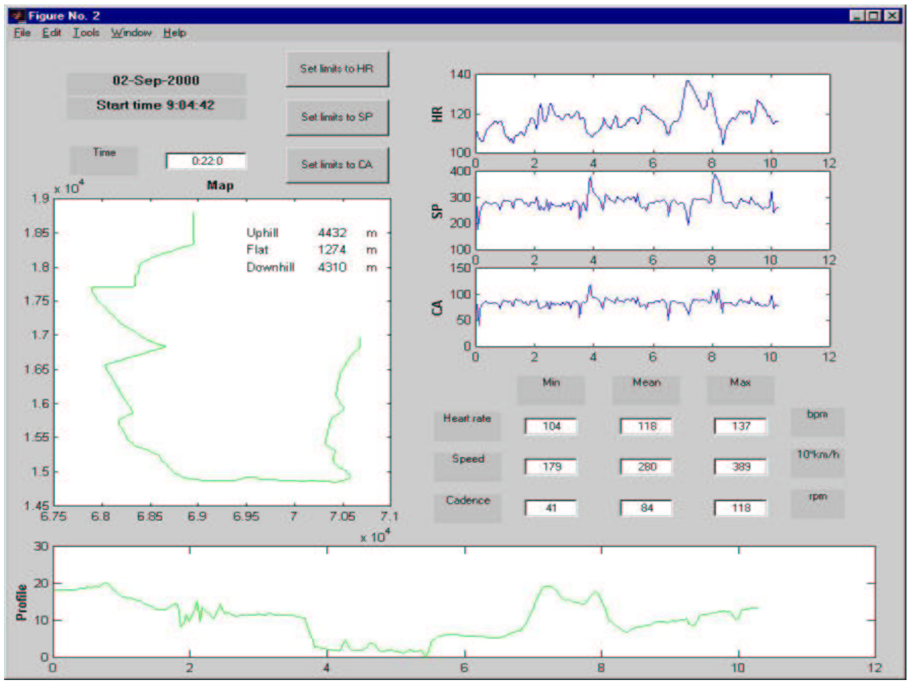

Fig. 4. User interface for evaluating a cycling exercise. Heart rate, speed, and cadence are shown at the top right. A 2-dimensional map of the exercise route is shown at the top left and the route's height profile at the bottom.

The user can select thresholds for each variable to observe how other variables and the possible route profile are related within these thresholds. It is also possible to compare two exercises completed on the same route, as in figure 5.

5 Conclusions

Context-aware computing can utilize the methods of data mining in several different ways. In this paper, data mining was performed for cycling data to find

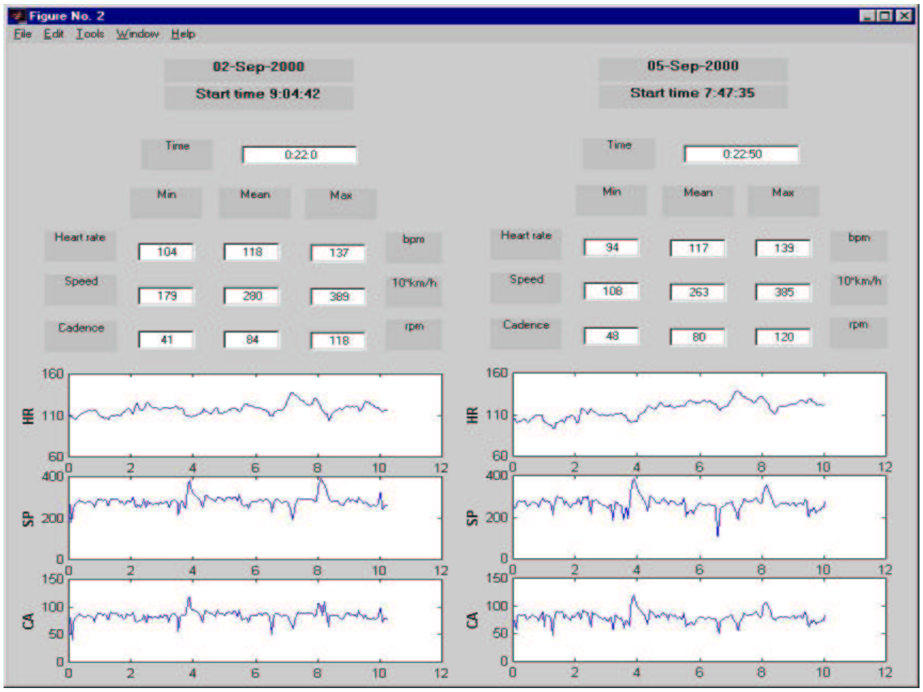

Fig. 5. User interface for comparing two cycling exercises.

interesting contexts that can be useful for a cyclist planning an exercise schedule. An analysis of the context history results in comments on the exercise, pointing out, for example, that the user started at too high a speed. Furthermore, the success in keeping the heart rate, cadence and speed within the intended intervals can be judged. Numerous simple tools can be offered to the user for studying the minimum and maximum values of the collected data, comparing two exercises, matching the collected data against the height profile, calculating the time that elapsed between two locations, etc.

A comparison of two exercises becomes easier if the profile information of the route is available, but analyzing services can also be offered without this profile. Comparing two exercises completed on different routes is a challenging problem. The route could be divided into periods of certain lengths and these periods could be compared more closely.

The role of the user is very important, as nobody, at least not yet, wants to be completely dependent on the artificial intelligence of the equipment. Thus, the data analysis has to be made understandable for the user, so that she can select which analyzing services to use, and know how to use them. The user interface developed in this research is clear and illustrative, and it is easy to add interactivity and simple analyzing tools, such as estimating fitness or long-term progress, to the interface. When several users can access each other's exercise data, they can compare their tactics and exercise data, as long as privacy issues are handled.

In the SOM experiments, the variables did not contain enough information about the dynamics of the situation. For example, when the cyclist is going uphill, the history information should be considered from further back than just one observation. If the earlier context was the cyclist in a steep downhill, the speed does not slow down as rapidly as if the earlier context would have been level ground.

The classification of the observations into different route declinations needs additional work, so that the contexts: *cyclist going uphill*, *cyclist going downhill* and *cyclist on level ground* can be recognized. Other resistant forces, like wind, the material of the ground, rain, the brakes, etc, confuse the classification results. The low sample frequency was also a problem in the classification, since the location of the cyclist had to be approximated. Another reason for these results is that too small declinations in the target classes were interpreted as uphills and downhills. In this work, the tested person is in good shape, and his heart does not probably react to small changes in declination.

The ground declination contexts could be recognized straightforwardly using additional measurements, for example, a tri-axial acceleration sensor or an altimeter (by Polar). Sideways swinging of the bicycle, indicating bad technique, could also be recognized with this kind of posture measurement. The classification would be easier if the sample frequency for the measurements was higher, but this requires modification of the existing hardware. Another future challenge is to classify the measurements of several cyclists, since there are differences in heart rate, cycling tactics and equipment.

Cyclists use several different kinds of tactics, depending on their individual strengths, but the use of the gears separates a beginner and a professional. If it is possible to detect the moment at which a gear has been shifted, the effect of different shifting tactics on the overall performance could be analyzed. This would also be valuable information for the recognition of ground declinations.

The visualization and analysis of the exercises is perhaps the most important part for a cyclist in analyzing her own progress and tactics. The interface will be evaluated by cyclists to find out their real needs. Future work includes implementing the Matlab version of the data processing and the user interface as components of the Genie system.

Furthermore, a wider set of contexts than those experienced during the exercises could be recognized. For example, a calendar containing the planned exercises could be presented to the user in certain contexts. The Health Club is an example of a more general application area, in which the observed activities are compared with planned activities. For example, an office worker might plan a time schedule, and the system could evaluate how accurately the plan was followed.

The methods developed for the Health Club context recognition can be used in different applications, also. A data set in an office environment, consisting of illumination, humidity, temperature, touch, and acceleration measurements while the user is in various contexts, was collected. The main difficulty is that the points in time, where the context has changed are not accurately known.

Now that off-line analysis is done for the data, the lack of actual information from these situations complicates context recognition. It is very important for the analyzer to get the right reference for every situation that occurs during different activities.

The work reported in this paper gave us a valuable insight into ubiquitous data mining and its differences to traditional data mining. The main difference between ubiquitous and traditional data mining frameworks is that traditionally data has not been gathered for a specific purpose, making the analysis secondary, whereas in ubiquitous environments the collection of data has to be planned carefully to fit the application. Furthermore, context-aware applications place new challenges on data mining, as the same methods have to work for different users and in different environments.

Acknowledgments

This work was funded by TEKES, Nokia, Polar Electro, CIM-Tech, Solid, Tauno Tönning Foundation and Infotech Oulu graduate School. VTT Electronics participated in developing the Genie system. The authors would like to thank Anu Himola for her work. Petteri Ellilä and Ulla Elsilä are acknowledged for their advice in coaching.

References

[1] N. Adams, R. Gold, B.N. Schilit, M.M. Tso, and Roy Want. An infrared network for mobile computers. In *USENIX Symposium on Mobile & Location-Independent Computing*, pages 41–52, Cambridge, Massachusetts, August 2-3 1993. USENIX Association.

[2] C.H. Chen, editor. *Fuzzy Logic and Neural Network Handbook*. McGraw-Hill Series on Computer Engineering. McGraw-Hill Inc., 1996.

[3] B. Clarkson, K. Mase, and A. Pentland. Recognizing user context via wearable sensors. In *Proc. Fourth International Symposium on Wearable Computers, ISWC2000*, pages 69–75, Sheraton Colony Square, Atlanta, GA, USA, Oct. 18-19 2000.

[4] A.K. Dey and G.D. Abowd. The context toolkit: Aiding the development of context-aware applications. In *Workshop on Software Engineering for Wearable and Pervasive Computing*, June 2000.

[5] R.O. Duda and P.E. Hart. *Pattern Classification and Scene Analysis*. Wiley-Interscience, 1973.

[6] E.Tuulari. Context aware hand-held devices. VTT publications:412, VTT Electronics, Espoo, 2000.

[7] R.C. Gonzalez and R.E. Woods. *Digital Image Processing*. Addison-Wesley Publishing Company, Inc., 1993.

[8] S. Hamilton. Inside microsoft research. *Computer*, 31(1):51–58, January 1998.

[9] S. Haykin. *Neural Networks, A Comprehensive Foundation*. Macmillan Publishing Company, 1994.

[10] J. Healey and R. Picard. Digital processing of affective signals. In *Proc. ICASSP*, Seattle, Washington, May 12-15 1998.

[11] A.C. Huang, B.C. Ling, and S. Ponnekanti. Pervasive computing: What is it good for? In *Proc. ACM International Workshop on Data Engineering for Wireless Mobile Access*, pages 84–91, 1999.

[12] I.Bhandari, E.Colet, J.Parker, Z.Pines, R.Pratap, and K.Ramanujam. Advanced scout: Data mining and knowledge discovery in NBA data. *Data Mining and Knowledge Discovery Journal*, 1(1):121–125, 1997.

[13] J.Himberg, J.Mäntyjärvi, and P.Korpipää. Using PCA and ICA for exploratory data analysis in situation awareness. In *IEEE Conference on multisensor fusion and integration in intelligent systems*, 2001. accepted.

[14] J.Kaartinen, Y.Hiltunen, P.Kovanen, and P.T.Ala-Korpela. Application of self-organizing maps for the detection and classification of human blood plasma lipoprotein lipid profiles on the basis of 1h nmr spectroscopy data. *NMR in Biomedicine*, 11:168–176, 1998.

[15] J.Mäntyjärvi, J.Himberg, P.Korpipää, and H.Mannila. Extracting the context of a mobile device user. In *IFAC Symposium on analysis, design, and evaluation of human-machine systems*, 2001. accepted.

[16] J.Pascoe, N.Ryan, and D.Morse. Issues in developing context-aware computing. In *Proc. HUC 1999*, Karsruhe, Germany, September 27-29 1999.

[17] J. Kaartinen, E. Mierisová, J.M.E. Oja, J-P. Usenius, R.A. Kauppinen, and Y. Hiltunen. Automated quantification of human brain metabolites by artificial neural network analysis from in vivo single-voxel 1h NMR spectra. *Journal of Magnetic Resonance*, 134:176–179, 1998.

[18] T. Kohonen. *Self-Organizing Maps.* Springer-Verlag Berlin Heidelberg New York, 1997.

[19] P. Korpipää. Visualizing constraint-based temporal association rules. *AIEDAM: Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 2001. Accepted.

[20] K.Väinämö, J.Röning, S.Nissilä, T.Mäkikallio, and M.Tulppo. Artificial neural networks for aerobic fitness approximation. In *Proc. International Conference on Neural Networks (ICNN '96)*, pages 1939–1944, Washington, D.C., USA, June, 3-6 1996.

[21] K.Väinämö, T.Mäkikallio, M.Tulppo, and J.Röning. A neuro-fuzzy approach to aerobic fitness classification: a multistructure solution to the context-sensitive feature selection problem. In *Proc. WCCI '98*, pages 797–802, Anchorage, Alaska, USA, May, 4-9 1998.

[22] K. Van Laerhoven. *On-line Adaptive Context Awareness Starting From Low-level Sensors.* PhD thesis, University of Brussels, 1999.

[23] P. Laurinen, J. Röning, and H. Tuomela. Steel slab temperature modelling using neural- and bayesian networks. In *Soft Computing and Intelligent Systems for Industry (SOCO/ISFI 2001)*, June, 26-29 2001. accepted.

[24] S. Lipponen, K.Väinämö, T.Mäkikallio, M.Tulppo, and J.Röning. Approximating aerobic fitness with neural networks applying PCA. In *Proc. 1997 Finnish Signal Processing Symposium (FINSIG'97)*, pages 159–162, Pori, Finland, May 22 1997.

[25] S. Lipponen, T.Mäkikallio, M.Tulppo, and J.Röning. Finding structure in fitness data. In *Conference on The Practical Application of Knowledge Discovery and Data Mining*, pages 101–109, London, UK, March, 25-27 1998.

[26] H. Mannila. Methods and problems in data mining. In Foto N. Afrati and Phokion Kolaitis, editors, *Database Theory - ICDT '97, 6th International Conference, Delphi, Greece,1997, Proceedings*, volume 1186 of *Lecture Notes in Computer Science*, pages 41–55. Springer, January 8-10 1997.

[27] J. Mäntyjärvi, T.Seppänen, and E.Tuulari. Hand gesture recognition of a mobile device user. In *Conference on Multimedia and Expo, ICME 2000*, New York, USA, 30 July - 2 August 2000. IEEE.

[28] W. Mark. Turning pervasive computing into mediated spaces. *IBM Systems Journal*, 38(4):677–692, 1999.

[29] R.W. Picard and J.Healey. Affective wearables. In *IEEE ISWC Proceedings*, volume 1, October 1997.

[30] P.Poptani, J.Kaartinen, R.K.Gupta, M.Niemitz, Y.Hiltunen, and R.A Kauppinen. Diagnostic assessment of brain tumours and non-neoplastic brain disorders in vivo using proton nuclear magnetic resonance spectroscopy and artificial neural networks. *Journal of Cancer Research and Clinical Oncology*, 125:343–349, 1999.

[31] D. Pyle. *Data Preparation for Data Mining*. Morgan Kaufmann Publishers, 1999.

[32] B.J. Rhodes, N. Minar, and J. Weaver. Wearable computing meets ubiquitous computing: reaping the best of both worlds. In *The Third International Symposium on Wearable Computers. Digest of Papers.*, pages 141–149, October 18-19 1999.

[33] D. Salber, A.K. Dey, and G.D. Abowd. The context toolkit: Aiding the development of context-enabled applications. In *Proc. Conference on Human Factors in Computing Systems, CHI'99*, pages 434–441, 1999.

[34] B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *Workshop on Mobile Computing Systems and Applications*, pages 85–90, December 8-9 1994.

[35] A. Schmidt, K.A. Aidoo, A. Takaluoma, U. Tuomela, K. Van Laerhoven, and W. Van de Velde. Advanced interaction in context. In Hans-W. Gellersen, editor, *Handheld and Ubiquitous Computing, First International Symposium*, pages 89–101, Karlsruhe, Germany, September 1999. Springer.

[36] A. Schmidt, A. Takaluoma, and J. Mäntyjärvi. Context-aware telephony over WAP. In *The Second International Symposium on Handheld and Ubiguitous Computing, HUC2k*, 24-27 September 2000. Short paper.

[37] S.Tamminen, P.Laurinen, and J.Röning. Comparing regression trees with neural networks in aerobic fitness approximation. In *Proc. of International ICSC Symposium on Advances in Intelligent Data Analysis (AIDA'99)*, pages 414–419, Rochester, NY, USA, June, 22-25 1999.

[38] S. Tamminen, S. Pirttikangas, S. Nissilä, V. Pentikäinen, K. Väinämö, and J. Röning. Multiple alarm management with self-organizing maps. In *Proc. Intelligent Data Analysis in Medicine and Pharmacology*, pages 135–144, Washington, DC, USA, November 6–10 1999.

[39] S. Tamminen, S. Pirttikangas, and J. Röning. Self-organizing maps in adaptive health monitoring. In *Proc. International Joint Conference of Neural Networks (IJCNN2000)*, pages 259–266, July, 24-27 2000.

[40] K. Väinämö, J. Röning, and P. Laurinen. Smart archive for online learning systems. In *TOOLMET 2000 - 3rd International Symbosium on Tool Environments and Development Methods for Intelligent Systems*, Oulu, Finland, April, 13-14 2000.

[41] R. Want and A. Hopper. Active badges and personal interactive computing objects. *IEEE Trans. Consumer Electronics*, 38(1):10–20, February 1992.

[42] M. Weiser. The computer for the 21st century. *Scientific American*, pages 66–75, September 1991.

[43] M. Weiser. Some computer science issues in ubiquitous computing. *Communications of the ACM*, 36(7):75–84, July 1993.