Secure Linking of Patients and Tele-medical Sensors

Michael Højgaard Bojsen

Kongens Lyngby 2011
IMM-M.Sc-2011-77

Technical University of Denmark
Informatics and Mathematical Modelling
Building 321, DK-2800 Kongens Lyngby, Denmark
Phone +45 45253351, Fax +45 45882673
reception@imm.dtu.dk
www.imm.dtu.dk

Summary

The main task of this thesis is to identify if and how anomalies origin from a patient switch can be detected in medical data. By investigating various approaches and methods for pattern recognition a proof-of-concept model based on a heuristic approach is been analysed, designed and implemented.

The hypothesis is that newly arrived data from a fitness watch with heart rate monitoring capabilities can be validated against an already trained pattern signature is tested.

A state-of-the-art survey in telemedicine is performed from a technical perspective, where a case study from Frederiksberg Hospital is included, to expose the challenges and advantages in telemedicine the health care system can benefit from now and in the future.

All material found in this thesis is available on the attached CD in Appendix F.

Resumé

Hovedformålet med dette speciale er at identificere hvis og hvordan anomalier, stammende fra et patientskifte, kan blive opdaget i medicinsk data. Ved at undersøge forskellige tilgange og metoder for mønstergenkendelse er en prototype, baseret på en heuristisk tilgang, blevet analyseret, designet og implementeret.

Den testede hypotese er, at nyindsamlet data fra et sportsur med tilkoblet pulsmåler kan blive valideret mod en allerede trænet mønstersignatur.

En state-of-the-art undersøgelse i telemedicin er gennemført fra et teknisk perspektiv, hvor også et eksempel på et medicinsk forsøg fra Frederiksberg Hospital er inkluderet. Dette er gjort for at afdække, hvilke fordele og ulemper telemedicin kan bidrage med i sundhedssektoren nu og i fremtiden.

Alt materiale i dette speciale er tilgængeligt på den medfølgende CD i Appendix F.

Preface

This thesis was developed at Informatics and Mathematical Modelling at the Technical University of Denmark under the supervision of Christian Damsgaard Jensen in partial fulfilment of the requirements for acquiring the Ms.C. degree in engineering.

This thesis deals with aspects in matching patients with their medical sensors and detecting if a sensor, or patient, has been swapped.

The main focus is on developing a system in which a number of ideas will be tested and evaluated.

The thesis was developed during the period March 2011 - September 2011

Lyngby, September 2011

Michael Højgaard Bojsen

Acknowledgements

First and foremost I would like to thank Martin Bornhøft for proof reading and the many productive and funny brainstorm sessions we had during the process.

Secondly I would like to thank my supervisor Christian Damsgaard Jensen for his support throughout the project.

Jan Larsen, Intelligent Signal Processing Group, IMM, DTU should also be mentioned as he was able to guide me in the right direction after I reached a dead-end.

Contents

Chapter 1

Introduction

With the progress done in the field of telemedicine where more and more patients are "wired up" to some kind of monitoring device a crucial problem arises: How do we make sure that the worn device, and thereby the data collected from the device, is associated with the corresponding patient? A mix-up could result in misdiagnosing and following wrong or potential dangerous medication of the patient. E.g if a patient's blood pressure is monitored for a period of time but the medical sensors are coupled with a different patient. The other patient could have high blood pressure, while the first patient could have low blood pressure. Both patients would be misdiagnosed and mismedicated to treat the other patient's symptoms.

To distinguish patients from each other they have to be associated with some kind of profile or signature describing their course of illness. These signatures can be used to detect if a medical sensor or device has been or is being used by another person than the original patient.

One approach will be to use machine learning and/or neural networks to analyse the measured data for anomalies. This will be achieved by having a set learning period where the risk of a patient switch will be assumed minimal. This will only validate a patient against an assumption that the patient is the patient to whom the medical equipment is handed out to. This could be done by matching the pattern of the medical data against the previously learnt pattern and if possible

against the diagnosis given to the patient.

The alternative is to authenticate the patient each time a measurement has to be made. This would be clumsy and result in the whole concept of telemedine and remote monitoring to fall apart as the patient would be interrupted throughout the day making telemedicine very impractical.

The developed model evolves around heart rate monitor data obtained from simple sports watches. This enables us to get free real data from multiple volunteers with a reasonably quality. The collected data is used for tests and proof of concept setups.

1.1 Intended Use

The developed model is intended for use by a General Practitioner (GP) to detect whether or not telemetrically collected data is accurate and valid for a given patient.

It is meant as a help for the GP to be alerted if the collected data is polluted with another person's data (e.g. the patients next of kin), as such an event could potentially result in a wrong diagnosis and thereby following wrong medication.

This scenario can occur if a medical sensor is taken off e.g. in preparation for a shower. Here lies the possibility of another person taking it, either by mistake or with foul play in mind.

Another scenario where it is important to detect a change of patient evolves around the work flow taking place at the GPs office during a patient switch. Most medical devices have some kind of memory of the recorded data. If these devices are not reset or in some other way cleared of the recorded data, the possibility of importing data from a wrong patient into a journal could occur and have severe repercussions.

1.2 Thesis outline

The structure of this thesis follows a standard academic software approach with minor deviations. The thesis initially presents a description of critical terms and knowledge used throughout the project.

A state-of-the-art survey in telemedicine is performed in chapter 2 to get an overview of where telemedicine is today and how it is used. This is done by a case presentation of a current telemedical trail. The advantages and challenges in telemedicine are discussed. Furthermore a description of what telemedicine will be like in the future is included.

The system analysis in chapter 3 gives a tour through a proposed architecture for a telemedical system. A risk overview will also be presented for the system.

In the analysis in chapter 4 a number of use cases for the system are defined.

The data collection discussion in chapter 5 discusses the medical background for the data collected which are explained along with the data collection method; including devices and test subject selection. Furthermore a discussion of the collected data is performed.

In the theory chapter (chapter 6) the concept of signatures or profiles is touched upon and two main approaches, heuristic and machine learning, are discussed.

The design (chapter 7) of the proposed system is decided before the implementation (chapter 8) is discussed and explained.

An evaluation of extracted features from the collected data is performed in chapter 9.

Finally the future work of the system are discussed in chapter 10, followed by the conclusion of the thesis in chapter 11.

Chapter 2

Telemedicine

The following state of the art survey extends and clarifies the topics examined in the earlier project "NaPiLink"[4].

2.1 Background

The demographic of the world is changing these years as medical science can treat and cure more and more diseases. This results in more people living longer and thus the cost of medical systems increases.

To combat this increased cost of health care, telemedicine has been foreseen as a solution. This has resulted in a brand new marked in medical equipment, which will enable patients to be monitored by medical sensors without being hospitalised where patients can be exposed to other infections.

Remote monitoring of non-critical patients will enable hospitals to allocate less man power and thereby save money in salaries.

When a number of hospital beds can be virtually "moved" into the patient's own home it leaves more space in the hospitals for critically ill patients requiring more

than simple monitoring. This will give the critical patients better opportunity to be treated without waiting for extended periods of time for a treatment program in the hospital.

In remote areas like the Australian outback a telemedical infrastructure could enable local communities to get highly specialised medical advise without having to travel to the nearest hospital which can be hundreds of kilometres away.

The implementation of telemedicine is not without problems as it is crucial for patients trust in the system that every aspect is covered in risk assessments done by e.g. the Federal Drug Administration (FDA). Their approval is needed as a telemedical system will be considered as a medical device in itself.

Furthermore the success of telemedicine will depend heavily on how the patients perceive the system. The patients need to trust that the quality of their treatment in their own home, under no circumstances, will be any less than the quality of the treatment if it were performed at a traditional hospital.

2.2 State of the art

A number of telemedicine projects are currently being developed in Denmark. In this section we have selected one case of a telemedical project in process at Frederiksberg Hospital. This represents the current state-of-the-art in danish telemedical treatment and includes a number of key figures in cost of traditional medicine compared to telemedicine. Furthermore the complexity of such a system is described in a very practical manner while keeping the patients' safety and treatment on a high level.

This section also contains a table with the most common and available medical sensors for home use.

2.2.1 Case - Chronic obstructive Pulmonary Disease

Chronic Obstructive Pulmonary Disease (COPD) causes the patients lung function to decrease significantly. It is mostly caused by smoking. Typical symptoms are shortness of breath and chronic coughing. The Danish National Board of Health estimates that more than 300.000 danes suffer from COPD today [8].

To help patients with COPD Frederiksberg Hospital is doing a medical trial

to investigate whether or not telemedicine can improve the life quality for the patients and reduce the cost for the public health care system[5][7].

The normal course for a patient with COPD is that in every day life not much is required besides remembering to take the prescribed medication. However if the illness develops, changes or complications like influenza develops, the need for hospitalisation might rise.

The traditional procedure has been to treat the new come infection, make some adjustments to the patients medication and hospitalise the patient for observation for four to eight days. A cost estimate of one patient hospitalised for one day lies between DKK4.000 and DKK8.000[7]. The doctor responsible for the trial at Frederiksberg Hospital estimates the cost of a day of home hospitalisation to DKK1.400. In most cases a great improvement is seen within the first few hours after the adjusted medication is given. The next days are only used for observation to see if the medication should be adjusted further. When the medical staff is absolutely sure that the new medication has the wanted effect the patient can be sent home.

In the period where the patient is hospitalised a number of factors are less than perfect; the patient is away from home, meaning less or no visits from friends and family. During the hospitalisation the patients are normally housed in two, four or six person rooms with minimal privacy and comfort.

In the medical trial a selection of COPD patients are given the option to be hospitalised in their own home. This is done by giving the patient a computer with an attached webcamera and pulse oximeter. After having their medication adjusted as usual, the patient receives a thorough instruction in how to use the equipment and is thereafter sent home to recover after the incident that caused the patient to contact the hospital in the first place.

The medical staff in the trial will do their rounds in a more virtual manner; instead of walking from room to room to attend to the patients the medical staff will sit in front of a computer with camera and microphone attached and make video calls to their patients. These calls are made at a time agreed-upon with the patient.

The computer provided to the patient can connect the patient with both audio and visual communication to the staff at the hospital at the push of a button; an alarm will sound at the hospital in the same manner as if the patient had pulled the string above their bed in the hospital.

The patient can now, in the comfort of their own home, surrounded by family and the usual home care nurse, live a more or less normal life even when hospi-

talised. The patient can watch television, eat their usual food, go to sleep and wake up according to their own schedule. Their only commitment is to contact the hospital staff at the agreed time or in case of emergency.

Preliminary studies show that the patients get the same or better quality treatment while hospitalised in their own home. Patients report back that their life quality is significantly improved.

2.2.2 Available medical sensors

More and more manufacturers of medical equipment are targeting the consumer segment as microchip technology is becoming more available at prices suitable for the end-user marked.

The most common available medical sensors suitable for home use are listed in table 2.1

Sensor	Comment
Sphygmomanometer	Enables patients to monitor their own blood pressure
Blood glucose monitoring system	Essential for diabetes patients
Pulse oximeter	Inexpensive device using diffraction of light in two wave lengths to determine the oxygen saturation of blood
Heart rate monitor	Often found in fitness watches, utilising two electrodes in an elastic band to measure electrical potential between right and left side of the chest

Table 2.1: Table of a non-exhausting list of medical sensors available on the consumer market today.

The various devices can either be handed out by a physician or can be bought by the patients themselves to monitor for values to check the state of ones body.

The most commonly available medical sensor in the consumer market must be the automatic sphygmomanometer for measuring blood pressure. The device works by having an inflatable cuff around the patient's arm to restrict blood flow and uses a manometer to measure the pressure. Utilising a microphone the amount of pressure in the cuff allowing the blood to just start flowing and the pressure where the flow is unimpeded are recorded. These measurements are recorded as systolic and diastolic blood pressures.

A blood glucose monitoring system is essential for diabetic patients to monitor their blood sugar level as both low and high levels can have serious effects on the health of the patient. By extracting a small amount of blood from the patient and applying glucose specific enzymes to it the amount of glucose can be estimated. The results of such measurements are important to record in some kind of formal way to enable the patient's GP to monitor the development of the condition and thereby help the patient regulate the medication.

A pulse oximeter is a small clip-like device designed for either a finger or an earlobe. The device functions by emitting two different wavelengths of light from two LEDs[1] through the skin of the patient. By observing the difference in the absorbance of the two wavelengths the device can calculate the amount of oxygenated haemoglobin in the patient's bloodstream. When the heart pumps fresh blood into the finger or earlobe this level rises and on that basis the patient's pulse can be calculated. It is often used to monitor patients suffering from COPD as the patient will often suffer from reduced blood oxygenation due to the ineffectiveness of their lungs.

A heart rate monitor is most commonly known either from a simple fitness watch or a much more complicated ECG/EKG[2]. Both are based on using electrodes on the skin of the patient's chest and monitors the electrical potential across two or more measuring points. From these electrodes the patient's heart rate can be extracted.

2.3 Challenges

The greatest technical challenge in telemedicine is to prove that the system ensures confidentiality, integrity, availability and authenticity of data. These requirements must be met before telemedicine can be used in other contexts than tests and trials. The data collected should be accurate and be of at least the same use as measurements done at hospitals and GP offices.

A number of risk assessments must be made to ensure that if anything fails there will be a fallback procedure that ensures that no treatment of a patient is compromised. A quick overview of the risks in a telemedical system follows in section 3.3 on page 15.

To be able to monitor patients in their own homes some kind of connectivity will be required. Depending on what or how much data needs to be transferred

[1] Light Emitting Diodes
[2] Electrocardiography

as little as a 56k modem or GPRS [3] can be sufficient. This will however not give the possibility of transferring live video in a resolution high enough for it to be used in diagnosis. If the required amount of data that needs to be transferred is large or needs to be real time, the requirements to the local connectivity will be strict. This can cause patients in some locations to be excluded from being hospitalised in their own home and will therefore require normal hospitalisation. These include, but is not limited to, locations without xDSL[4], high speed cellular service or technologies like WiMAX[5].

Another aspect to investigate would be the willingness of patients to be hospitalised in their own homes. Patients might have concerns about the quality of treatment and responsibility amongst the medical staff.

To ensure the correct treatment of more serious illnesses it would not be beneficial for either the health care system or the patient to be treated in their own home. Mostly treatment involving a large amount of monitoring would be able to benefit from telemedicine.

2.4 Advantages

Telemedicine in private and especially public health care is foreseen a great future as it can save many man hours of monitoring as well as many costly hospitalisations[7].

When the telemedical system is approved by the FDA a manufacturer of medical devices will fairly easy add telemedical capabilities to their product as the infrastructure for the different sensors are the same.

An added bonus of telemedicine is that the patients will have a higher quality of life while being hospitalised in their own home.

If we take a trip approximately 10 years into the future, telemedicine will be much more accepted and used. Both to replace the monitoring done in hospitals and even to add much more monitoring. This will result in general better health in the population as more illnesses can be diagnosed earlier than they can today.

We will see a more widespread use of telemedicine as technologies like 3G, LTE[6]

[3] General Packet Radio Service, a packet oriented mobile data service, based upon GSM.

[4] Digital Subscriber Line, a family of technologies that allow data transmission typically to the internet

[5] Worldwide Interoperability for Microwave Access

[6] Long Term Evolution, a standard build upon GSM/GPRS to provide a fast connection

and WiFi hotspots become more and more available to the general public. This will enable the wast majority of the population to stream live medical data to a service provider.

With the use of pattern recognition in telemedically obtained data the diagnosis process could be more or less automatic. Either to detect the very first telltale signs of new illnesses with patients suffering from chronic diseases or to detect changes or worsening of already known illnesses.

The whole health care system will be digitalised to allow patients to read their own journal from their own computer and thereby follow their treatment more closely.

via the cellular network

Chapter 3

System analysis

From a technical perspective telemedicine is well suited for cloud computing, as the amount of data will generally be large and complex while the answers after processing are wanted to be simple and clear. This enables the system designer to choose a cheap and low powered computer or microchip to collect and dispatch the raw medical data to "the medical cloud". At least three approaches can be chosen amongst;

1. either will the cloud service only sanitise the data and present them to the appropriate medical person (e.g. "Patient X's heart rate was 29 BPM [1] ")
2. the cloud service will be able to make some calculations and comparisons on the received data and present the medical person with a result (e.g. instead of presenting a result as in example 1 the service could provide a result such as "Patient X's heart rate was significantly lower compared to other measurements and to other patients data with the same medical history")
3. or the cloud service can do option 2 and furthermore give the patient some direct response (e.g. "Your blood pressure is unusually high, your GP has been notified and will see you next Tuesday")

[1] Beats per minute

In all cases the raw collected values should be included in the patients Electronic Patient Journal (EPJ) for review by the GP or other medical staff. An infrastructure like the described however rises a need for the EPJ and other medical records to be located in "the cloud" here meaning that all data should be available from all internet enabled computers. This will rise a number of risk assessments which we will not touch upon in this thesis.

No matter to what degree one wants the system to be intelligent a basic infrastructure must be set up.

3.1 General Architecture

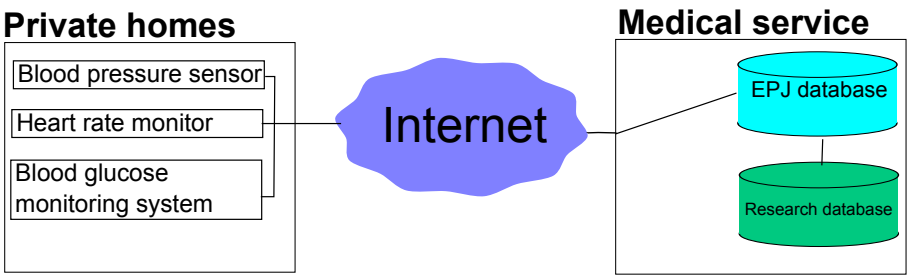

Figure 3.1: Raw data is collected by the patients at their homes, sent via the internet to a medical network

Figure 3.1 shows an overview of a general system. Patients use the medical sensors to collect data which is sent over the internet to a medical network where it is analysed. It should be noted that the above system is the absolute minimum architecture required. If deploying a medical cloud service it would be advantageous to have a specific service for technical maintenance of the medical sensors. This service could be provided by the manufacturer of the sensors.

If allowed by the patient, the data can be shared anonymously to a common database to help GPs to diagnose and treat patients with similar data or data patterns. These data can be used in research and to detect problems with the population as a whole. The common anonymous database will (ideally) be available to every health care professional and research team.

3.2 Binding of patient and medical sensor

A mismatch between a patient and a medical sensor not capable of transmitting live data is relatively rare as the patient will have to bring the sensor back to the GPs office. In this context, as the patient is already "validated", the mere possession of the medical device is sufficient. If the data is downloaded from the device as the patient is still with the GP, the GP will have the opportunity to ask the patient about any data looking odd and immediately get response about any unusual activity.

However a more strict protocol must be observed when a medical sensor *with* real time capabilities is handed out as a mismatch could have fatal consequences for the involved patients. As an absolute minimum some kind of double-check must be implemented. This could be having the patient writing down the serial number of the device on a receipt or a nurse or receptionist checking the device when the patient leaves the GPs office.

In the patients electronic journal the hand out should be noted.

3.3 Risk analysis

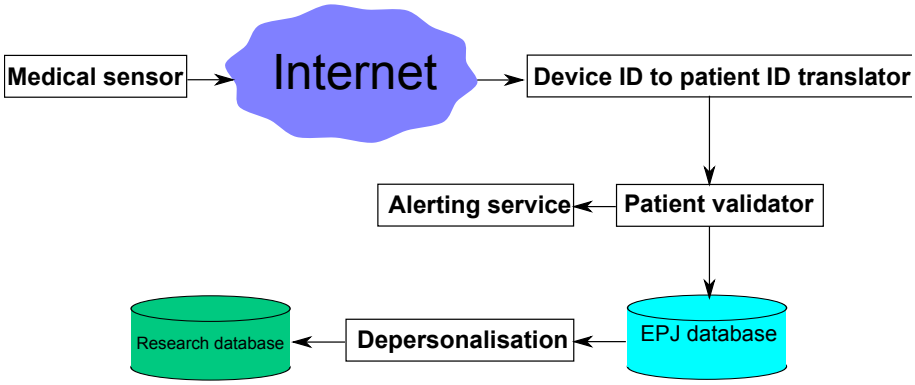

Figure 3.2: Component diagram of the data's path from sensor to EPJ

Figure 3.2 shows the different components in a telemedical setup.

Different challenges arises in the different components in the system. To trust that the system handles data in a secure and confidential manner a risk assess-

ment will give an overview of what measures must be taken to ensure that the risk of a patient journal containing wrong information is minimal.

When performing a risk analysis a tool is often needed to help the risk analyser to make sure the required steps are performed. In this case a framework like CORAS[6] could be relevant.

In stead of doing a full risk analysis on the whole system a quick overview of the components and their potential threats are discussed in the following.

It is expected that the transport of data between the components except for the *medical sensor* and *internet* will be inherently secured as it is performed in the same secure network.

3.3.1 Medical sensor

When data arrives to the system it is essential to be sure that the sender of the data is the device it claims to be, so some kind of mechanism for validating the authenticity (originator authentication) of the data transmitted from the device to our system is needed. Further the confidentiality of the data must not be compromised as medical data inherently contains sensitive information. The integrity of the data must be ensured from the medical sensor all the way to the EPJ as any changes of the data will result in corrupt data in the patients journal.

These requirements can be met by using some kind of certificate and key exchange and using an encrypted connection between the two end-points.

3.3.2 Device ID to patient ID translator

Under the assumption that the medical device is not wanted to contain information about the patient it is currently handed out to, some kind of translation from an unique device identification number to a unique patient identification is needed.

This step relates to the work flow at the GPs office where the device is handed out to the patient. In this process the *Device ID to patient ID translator* service must be given information about the time period in which a given translation is valid and most important what device is handed out to who.

If this service fail to give an answer corresponding with the reality the consequences will be significant as treatment (e.g. medication) depends on correct information in the patients journal.

The assumption mentioned above can be justified by looking at the alternative which will cause the process of handing out a device more cumbersome as the patients identification number (e.g. the social and security number (SSN)) must be written to the physical device. Also the patients privacy must be considered. If the patient looses the devices and it is possible for other people to extract the patients SSN and can relate the device to a specific condition, then the patients journal are exposed.

3.3.3 Patient validator

Even if the *device ID to patient ID translator* service do contain correct information no guarantee can be made that only the patient the equipment is handed out to is using it.

It would be of benefit if a service could validate the received data against previously recorded data to validate that the patient using the equipment does not change during a monitoring period.

Such a service would have to be trained with the original patient's medical data pattern or in some other way be provided with a pattern to match the incoming data against. The service should be able to detect if the person using the device is not the original patient. The service must be adaptive as patients patterns may slowly change, e.g. if the illness develops or improves.

The correct detection of a patient change is critical as a patient will be diagnosed wrong as the data basis would be incorrect.

3.3.4 Alerting service

With the propagation of medical data the need for privacy emerges. The *alerting service* should inform relevant and only relevant medical personal about the detected patient change. This information would have to travel along secure and trusted channels to the medical professional.

As this transport must be expected to leave the secure network again some kind of certificate and key exchange must be in place.

Dependent on how confident the *patient validator* is that a patient different alerting measures can be taken.

3.3.5 EPJ database

When the data firstly arrives at the *EPJ database* it is crucial that the data is associated with the correct patient. This requirement is fulfilled by the *device ID to patient ID translator* and *patient validator* services. The main risks in the *EPJ database* related to who can access what data when. For example should a GP be able to see all patients' journals or just the patients associated with his office? What if a number of GPs is co-located at the same clinic, it would be of convenience to the patients that a number of doctors can attend to their medical needs. A lot of questions could be asked and the answers must be sought in the law governing medical privacy in the applicable country.

3.3.6 Depersonalisation

In the *depersonalisation* step where the highly private data is transferred to the *research database* there is a risk if although some information, like social security number, name and address, is removed one can still associate the recorded data to a specific patient. This could be if a patient suffers from a very rare disease.

3.4 Summary

A risk assessment is paramount to uncover the risks inherently connected with telemedicine. By having to deal with sensitive, personal and life dependant data the greatest number of precautionary measures must be taken. This will include using state-of-the-art encryption and certificate exchange in all parts of systems handling telemedical data. To allow so many sensors and patients to use telemetrically collected medical data the transfer must be assumed to use the internet.

In this thesis we will concern ourselves with the challenges in the *patient validator* as this concept is very interesting in regards both to patient privacy and patient safety.

The system will be able to extract signatures from medical sensor data and

use these signatures to validate if newly arrived data is associated with the same patient. The validation should present the result in a percentage of the probability that the arrived data is from the original patient.

The system will not be able to *identify* patients solely based on the arrived data as this is assumed to be handled in a *device ID to patient ID translator* service.

If the system does not contain any signatures for a patient the received data will serve as a training set for the system.

In short the system requirements will be:

1. Be able to extract signatures from medical sensor data
2. Be able to validate if the extracted signature origins from the proposed patient
3. Store extracted signatures for different patients
4. Adapt stored signatures to allow for slow changes in data signatures

Chapter 4

Analysis

In this chapter an analysis of the model to be developed is performed. In chapter 3 we narrowed the prototype down to the *patient validator* component. Assuming the other components exist and functions in the manner described an analysis of the *patient validator* can be undertaken. This analysis will serve as a guide for how the system should behave.

A critical aspect is that the system only will alert the medical staff when a patient has been swapped even when the patient's medical data pattern evolves e.g. due to changes in the patient's physical fitness or the patients body adapts to the medication. This will require some kind of adaptive pattern recognition mechanism.

4.1 Overview

The system will receive telemedical data along with a patient identification to which the data is validated against using a heuristic approach with signature extraction from slices to fulfil the requirements stated in section 3.4.

The incoming data will be considered authenticated, kept confidential and being

unmodified. These mechanisms are assumed to be implemented in the transport from the *medical sensor* through the internet to the *device ID to patient ID translator*, all the way to the *EPJ database* and will therefore not be included in this chapter.

4.2 Signatures

To validate the incoming data against an existing patient some kind of signature generation/extraction must be done. A number of approaches can be considered, e.g. extracting statistical features from the data or using neural networks to generate weight functions.

4.2.1 Generation

The incoming data will originate from humans who we will expect to have some kind of habits during every day life which will be expressed as patterns in the medical data. This can be the patient's breakfast causing the blood glucose level to rise or the bike ride to work causing the patient's heart rate to increase. Most people will have fairly the same pattern for at least workdays and another more diffuse pattern for weekends and vacations. The patterns will remain fairly the same as long as the patient's day consists of the same actions or events at roughly the same time. In a perfect world this would be true and a pattern could easily be extracted as only measurement error and external noise should be handled. However the reality is different and will require a mechanism for handling offsets both in the point of time for events and even the skipping of events or actions during a day.

4.2.2 Storage

To store and retrieve the generated signatures some kind of database-like structure is needed. The datastore will have to be able to quickly and efficiently retrieve the signature for a given patient when requested. This will enable the system to validate the incoming data against a patient's signature in a timely fashion and thereby trigger an alert to the medical staff in near real time.

4.2.3 Matching

Assuming we now can extract signatures for the incoming data we will need to match it against the signatures already obtained for the patient.

To be able to match the incoming data we will need to previously have received enough data to be able to generate a signature/profile for the patient. If that is not the case the incoming data should be regarded as belonging to the proposed patient and a signature/profile should be generated from it.

When a patient is correctly validated against the signature obtained from the datastore the newly arrived data's signature should be handed to the datastore to enhance the signature available for future validations.

4.3 Alerting service

The *alerting service* will get a notification from the *patient validator* component if the received data does not match with the proposed patient. The *alerting service* should then contact appropriate medical staff to alert them to the indication of a patient switch. The alerting mechanism will not be covered in the implemented prototype.

4.4 Use cases

In this section a number of use cases are presented. Use cases in the sense that with a given set of inputs an specific behaviour is expected. The use cases are condensed from the requirements in section 3.4.

In each use case the *system user* is defined as the *actor* as no physical person will utilise the system directly. All interactions from patients with the prototype is through the medical device handed out by their GP. The GP will receive the results directly from the *alerting service* and do therefore not have any interaction with the *patient validator*.

Besides the *actor* the *precondition* for the use case is defined. This describes what is assumed to be in the system already. Furthermore the *main path* of the use case

The selected use cases are chosen as they represent the basic scenarios suitable for the proposed system. Each use case related to one or more requirements from section 3.4.

4.4.1 Learning patterns from a patient

Actor	System user
Precondition	The system contains no data about the patient.
Postcondition	The system has learnt the patients telemetrical pattern.
Requirements	1, 2, 3
Main path (M)	1. Patient wears/uses the medical equipment as per the GPs instructions.

4.4.2 Matching new data with existing patterns from the same patient

Actor	System user
Precondition	The system contains already previously learnt data about the patient.
Postcondition	The system matches the patient's telemetrical pattern with the already learnt pattern and improves the stored signature/profile for better future matching.
Requirements	1, 2, 3, 4
Main path (M)	1. Patient wears/uses the medical equipment as per the GPs instructions.

4.4.3 Matching new data with existing patterns from a different patient

Actor	System user
Precondition	The system contains already previously learnt data about patient A.
Postcondition	The system can not match patient B's telemetrical pattern with the already learnt pattern from patient A.
Requirements	1, 2, 3
Main path (M)	1. Patient B wears/uses the medical equipment handed out to patient A.

4.5 Summary

In this chapter the concept of signatures was discussed along with some reflections on how to match a signature with a number of previously obtained signatures. Furthermore the scope of this thesis was clarified by opting not to include other components but the *patient validator*.

A number of use cases and more specific requirements was presented to be more specific of what requirements falls into the scope of the *patient validator* component.

These defined requirements enables us to continue our investigation with data collection and experimentation with techniques of signature extraction.

Chapter 5

Data Collection

To investigate if it is possible to extract a signature from medical data distinctly enough to determine if a patient swap has occurred some data must be acquired. This data will enable us to test various data extraction methods for this validation of a given patient stays the same.

As medical/biometric data can cover many aspects, the goal was to select data satisfying the following criteria:

- Only contain one direct parameter as this simplifies the requirements and number of parameters in the system
  - Criteria excluding ECG/EKG[1]
- Easy and non-invasive to acquire
  - Criteria excluding blood glucose
- Cheap or easy accessible measuring devices
  - Criteria excluding treadmill max load test

[1] Electrocardiography

The selected parameter was heart rate in beats per minute (BPM) as it is continuous, non-invasive, easy and cheap to obtain and thereby satisfy the above criteria.

A theory is that the pattern in our heart rate is distinct enough to validate newly arrived data against a stored signature/profile to conclude if a patient swap has taken place.

A GP would typically want to monitor a patient's heart rate to examine the patient for anxiousness, arrhythmia or other heart related symptoms.

5.1 Medical background

A number of diseases causes patients to have distinct heart rate patterns.

Some patients suffering from diabetes can have symptoms of autonomic dysfunction. Autonomic dysfunction is a condition where the body's autonomous systems do not work properly. This can be expressed as constant tachycardia where the patient will have a heart rate >120 beats per minute even in a resting state[1].

Anxiety attacks will influence the patient's heart rate by raising it significantly while the patient is currently experiencing the attack. These attacks typically last up to 10 minutes[3]. Dependent on the severity, nature and on how well the condition is handled by medication the number of attacks can vary from only a couple of times a year up to multiple times during a day.

A more common condition is atrial flutter which is a heart disease due to a reentry circuit in the right atrium with secondary activation of the left atrium[9] causing a ventricular rate of 150 beats per minute.

Patients suffering from fever caused by infections often show tendency of tachycardia.

The above mentioned conditions will have significantly different heart rate patterns compared to healthy people.

5.2 Measurement devices

With the support from Pallas Informatics a number of identical fitness watches was obtained. Furthermore the CEO gave consent to that we could use his employees as test subjects, off course with their own consent. This allowed us to collect data from people from a somewhat diverse group of people.

The fitness watches was Garmin's Forerunner 60 for men with an optional digital heart rate monitor. This model was chosen based on its relatively low price, log capacity, battery capacity, weight and previously obtained knowledge about how to programmatically interact with Garmin products. The specifications for the watch can be found in Appendix B.

5.3 Test subject selection

We wanted to have a somewhat representative selection of test subjects. To achieve this we tried to get persons with different gender, age, height, weight and activity level. Not all combinations could be found but some. See Appendix A for full information about the test subjects and the amount of data collected from each.

It was not possible to recruit any test subjects with any known medical conditions concerning their heart.

5.3.1 Test subject briefing

Before obtaining any information or data from the test subjects they were given an oral briefing about the project in general and how the collected data would be used. The briefing contained at least the following:

- A promise of anonymity
- Information about what data is collected
- Information about what the data is used for

5.4 Observations

During the data collection period the test subjects reported back two main observations as follows.

5.4.1 Missing data points

A number of test subjects reported that the watches had a tendency to "fall out", meaning that no measurements from the heart rate monitor were recorded. With a little "fiddling around" with the heart rate monitor the link was restored. However the watch had a tendency to make unrealistic measurements a few seconds after the connection was restored.

5.4.2 Irritation and discomfort

The test subjects reported back that the heart monitor belt was not that pleasant to wear for extended periods due to the fact that it is made of plastic and synthetic material which will absorb moisture from the skin and cause irritation.

The irritation caused most of the test subjects to refuse to participate in further data collection or them to only be willing to provide data a few hours per day.

5.5 Data pre-processing

After collecting the raw data three filters was applied:

- Each test subject data was separated from each other
- Data points not containing a pulse value or a value of zero was discarded. This do not corrupt the data set as all data points contain a timestamp. The data points with a zero or empty value are a sign of the watch loosing connection with the heart rate monitor belt or the belt loosing contact with the test subjects skin.
- A simple low pass filter was applied, smoothing the data to correct for sudden spikes in the data values. A low pass filter was chosen as it uses the time factor to smooth the values.

After this the data points were sorted by timestamps and divided into days (i.e. [00:00-00:00[)

5.5.1 Data degradation

Data set	Average smoothing correction
A	0.3341
B	0.3492
C	0.3766
D	0.3676
F	0.3846
G	0.3262
H	N/A
I	0.3943
J	0.3346
K	0.3951

Table 5.1: Table showing corrected amount pr. data point in each data set

To investigate if the applied filter corrupted or in other ways manipulated the data the difference between the original data and the filtered data is averaged out on all data points. The results from this test can be seen in table 5.1

5.6 Data collection procedure

A total of five identical Garmin Forerunner 60 fitness watches was obtained. There were all configured to log and save heart rate data every fifth second from the paired heart rate monitor belt. Calculation of spend energy and other features was not relevant to enable in this context.

The data collection depended on people voluntary offered to wear a heart rate monitor and a watch with no other reward than a print out of the data collected from them.

Before a watch would be handed out to a test subject the logged data in the watch would be transferred to a PC using a wireless Garmin ANT dongle[2]. The watch was then reset to factory defaults to eliminate any settings compromising

[2]ANT is a proprietary wireless sensor communication protocol

the collected data. After the reset the wireless link between the heart rate monitor belt and the watch was verified.

When the watch was handed to the test subject a quick introduction of the usage of the watch was given; including how to start and stop logging and how to get the watch to show the current heart rate and the current time of day.

During the monitoring the test subjects could see their current heart rate on the watches display to verify that the link between watch, heart rate monitor and skin was established. The test subjects were instructed to check that the watch logged data during the course of their day especially in the morning. Further they were instructed to moisturise their skin and the electrodes in the heart rate monitor belt using saliva to enhance the connection.

5.7 Data quality

In the collected data the watches tendency to lose connection with either the test subjects skin or the data link between the heart rate monitor belt and the watch were distinct. Especially during the night where the test subjects must be presumed to sleep a loss of data spanning multiple hours were observed.

After the connection was restored typically there was a number of implausible measurements either values considered to be too high or low.

Besides the above mentioned issues the data quality was generally of high quality and by doing manual counting during 15-second periods the measured values could be validated as being true and accurate.

5.8 Data analysis

In figure 5.1 on page 33 two different test subjects data for the same day (2011-06-21) was plotted on the same graph to visually verify that the two patterns were different. As it can be seen in the graph test subject A generally had a lower heart rate opposed to test subject C.

In test subject A's data a couple of spikes around 16:48 and 19:12 is observed. The second spike is presumably from some kind of sports activity.

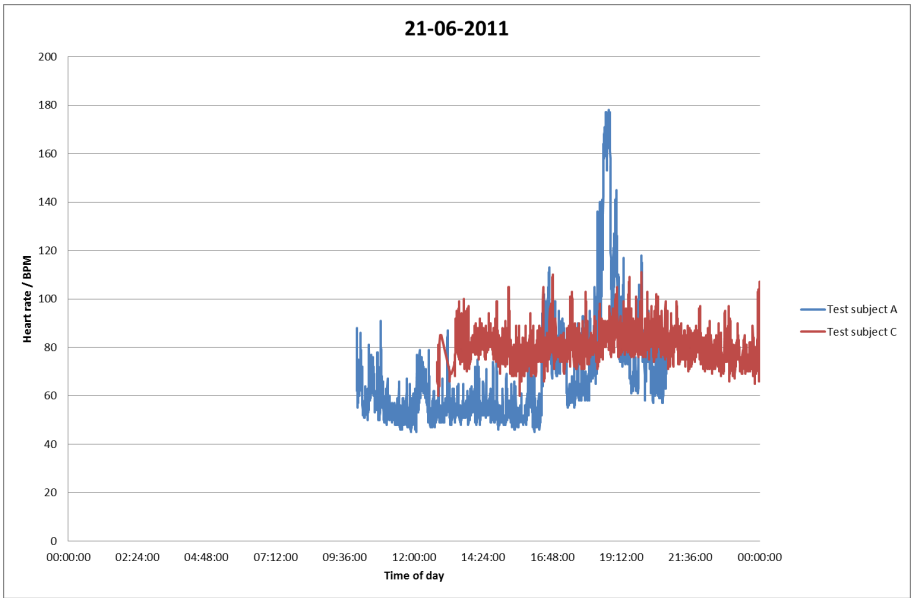

Figure 5.1: Plot with test subject A and C's data for 2011-06-21

In figure 5.2 on page 34 two days, Tuesday and Thursday, in the same week for test subject A are plotted. The baseline for the data is visually the same with a visually estimated average in the period 09:36 to around 16:00 on 62-64 BPM. The noticeable spike mentioned before is present in the two data samples but with a slight difference in amplitude and a marked difference in which time period the spike occurred.

In figure 5.3 on page 34 the two same days are graphed for test subject C. As noted previous on the 2011-06-21 were the data all in a confined interval centred around 80 BPM. In the data for 2011-06-23 a number of spikes can be observed. These spikes are short bursts of activity perhaps from climbing a flight of stairs. The spikes were not observed on the first day. On both days a slight rise in heart rate was observed in the afternoon compared to before lunch. This can be related to a desk job during the day and more activity in the test subjects spare time.

Also a couple of spikes with low values can be observed. These values seem abrupt and unlikely as the value changes quickly from around 75 BPM to 50 BPM.

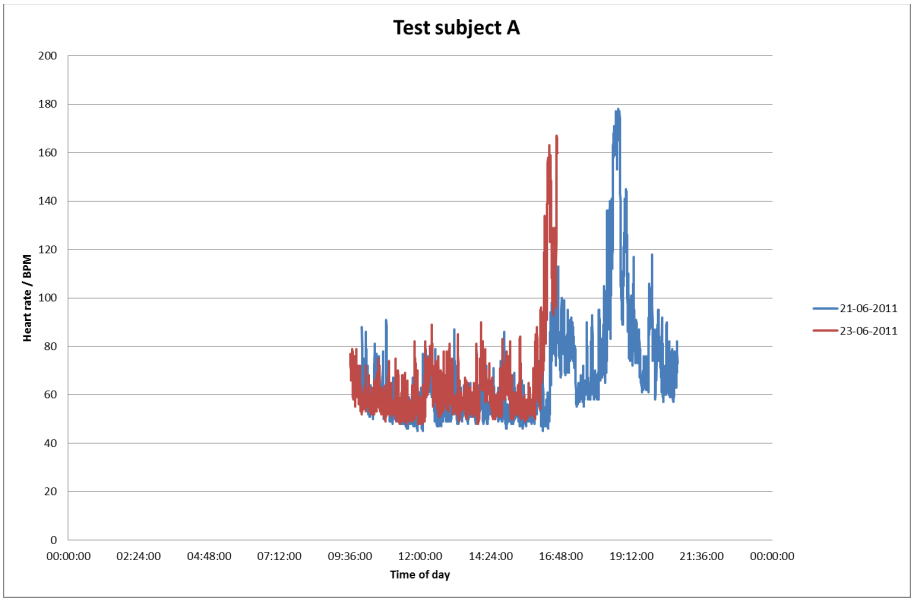

Figure 5.2: Plot with test subject A's data for 2011-06-21 and 2011-06-23

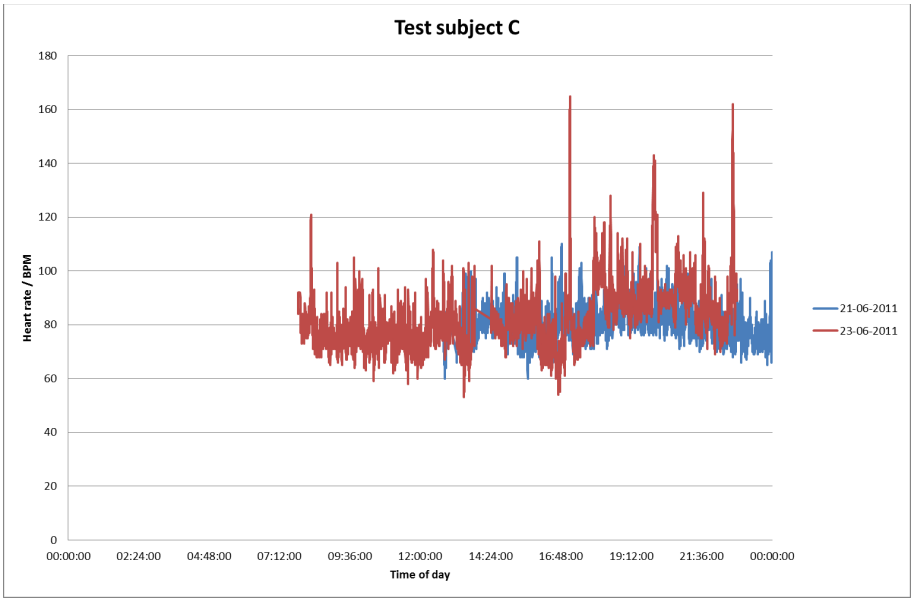

Figure 5.3: Plot with test subject C's data for 2011-06-21 and 2011-06-23

5.9 Exploratory Data Analysis

To investigate if there is a coherence between the different features extracted for the signatures we scatter plot two features against each other to see if there is a pattern.

In figure 5.4 two scatter plots are seen. First the absolute minimum value plotted against the absolute maximum value for each person. Some unrealistic values are seen in this plot (a minimum heart rate of 16 BPM is implausible low for a healthy person).

In the plot it is noted that the *min vs. max* plot spans more values both along the x- and the y-axis. This is expected as the plot is based on absolute values and is therefore a lot more sensitive to changes as the plotted values essentially only depends on one value per axis.

The *10th percentile vs. 90th percentile* plot is less scatted along the x-axis as the plotted values depends on the data set as a whole, not only single measurements. Along the y-axis the range is about the same as for the absolute value plots.

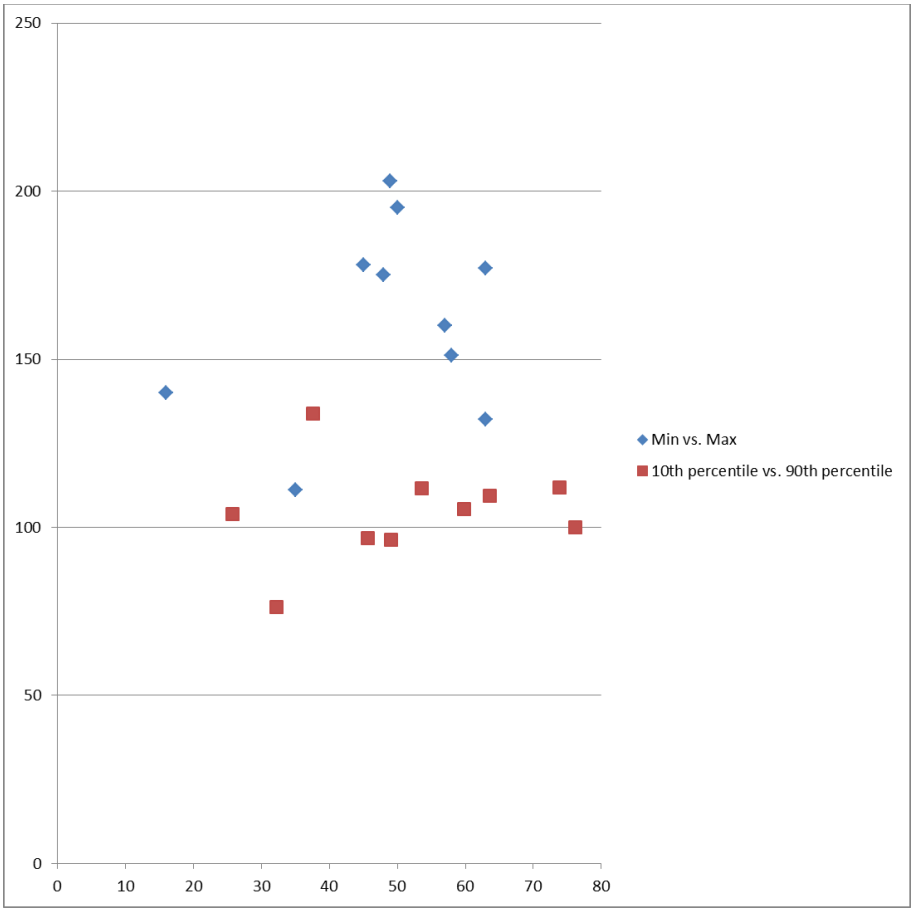

Figure 5.4: Scatter plot with both max vs. min and 10th percentile vs. 90th percentile. The x-axis represents the lower value and the y-axis the higher. Both axis is in beats per minute (BPM).

Chapter 6

Theory

The main goal is to be able to validate weather or not newly arrived data belongs to the same person. Two approaches will be investigated in this chapter: a heuristic and a machine learning.

The advantage of the heuristics approach compared to the machine learning is that the calculations will typically be less demanding and therefore easier to implement into the medical sensors itself.

6.1 Heuristic approach

There are a number of different heuristic approaches that are interesting, which all depends on how the data is perceived. A common feature of the reviewed heuristic approaches are some kind of slicing.

Slicing divides a data set into smaller bands of data to enable us to generate more than a single signature for one test subject for each recorded day.

The slices enables us to validate a patient's model against a fully trained signature set with partial data set recorded. A disadvantage however is that the

training data has to be of better quality and quantity.

The number of slices in a data set will be dependent on what type of data is needed to be extracted. Consider heart rate data of a patient with an illness causing abnormally low resting heart rate. In that case it would make sense to regard a days measurements as a whole to search for the minimum value. Then consider a patient suffering from anxiousness. In that case one would expect a number of heart rate peaks during a day, coherent with the anxiousness attacks. In this context the patients minimum and maximum values for the day are not relevant. Instead the difference in maximum heart rate from hour to hour is of much more value to extract a pattern from.

An example of collected data can be found in Appendix C

6.1.1 Signatures

To reduce the amount of data that needs to be stored for extended periods of time, it is useful to generate some kind of signature for a given data set.

These signatures have to describe as much as possible of the data set and still be compact and lightweight.

A simple approach for heart rate data would be to use the maximum and minimum values of BPM for a given period of time. This can be extended with information about the distribution of the contained values, e.g. the median or average of the values. The standard deviation of the values would also be of interest as it describes the density of the spectrum.

As the minimum and maximum values are very sensitive to outliers one would/could choose to use the 10 percentiles instead.

6.1.2 Signature extraction - Slicing

A number of different approaches have been investigated to find an appropriate method for signature extraction. The following lists the most promising approaches.

6.1.2.1 Value slicing

An example of value slicing can be seen in figure 6.1 on page 40

Value slicing is a horizontal slicing method where the measurements are divided into slices based upon their value. A strategy could be to base a signature for the slice on the number (or period) of measurements in the different ranges of heart rate.

6.1.2.2 Time slicing

An example of time slicing can be seen in figure 6.2 on page 40

Time slicing is a vertical slicing method where e.g. a twenty-four hour period is divided into a number of smaller periods.

6.1.2.3 Trigger slicing

Like time slicing, trigger slicing is a vertical slicing method. The first slice starts at a specific event, e.g. first time of a day where the heart rate of a patient rises above a set value.

An example of trigger slicing can be seen in figure 6.3 on page 40. In this example the first slice should be disregarded.

6.1.3 Comparison of slicing methods

Each of the described slicing methods have their forces and weaknesses. The most advanced method is *trigger slicing* as it introduces a number of parameters; threshold value for when to start slicing. Introducing as little new parameters as possible will be of benefit as the tuning of the system will be simpler.

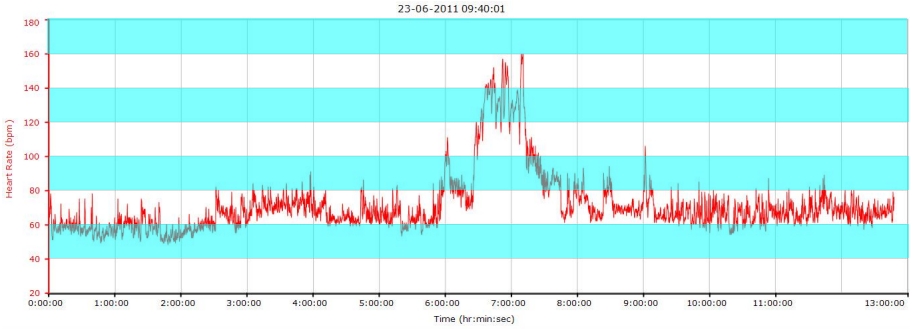

Figure 6.1: An example of value slicing

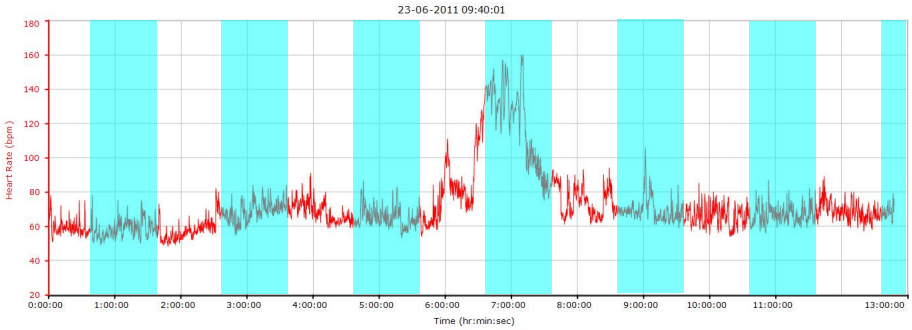

Figure 6.2: An example of time slicing. Note the offset due to the starting time of the collected data.

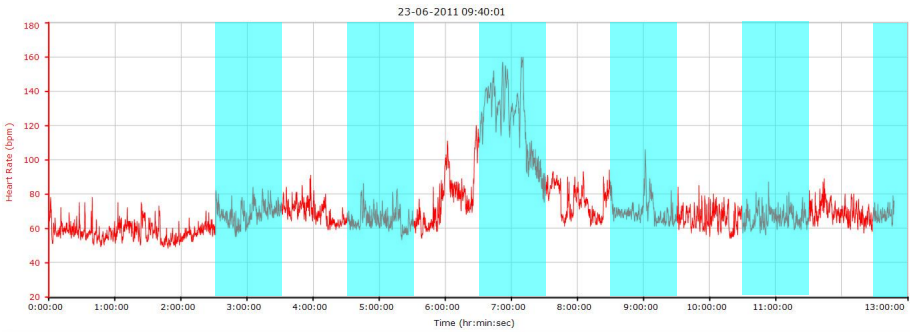

Figure 6.3: An example of trigger slicing

6.2 Machine learning

Machine learning is known for its ability to learn and recognise patterns in different contexts. For pattern recognition a multilayer feedforward network would typically be used as these can be trained to fit nonlinear functions and can handle noisy data. The resulting weight functions can however be hard to understand by the fact that they are initialised randomly and as an effect of this the machine learning approach might not be a suitable solution in medical equipment as requirements typically are to have deterministic programs.

6.3 Summary

To sum the different approaches up the proposed system should have some kind of signatures defined for each slice in the data sets. The easiest concept to grasp and comprehend is the time slicing as it is easily configurable with only one parameter. Value slicing requires more fine-tuning perhaps even patient individual. The most complex and customisable is trigger slicing as the system designer would have most parameters to utilise to get the pattern signatures to match the patient.

Machine learning is not investigated further in this thesis as it is more difficult to understand the inner workings due to the complex nature of e.g. a multilayer feedforward network.

If the simple approach of time slicing can prove to be useful it would be the most interesting because of the simplicity.

Chapter 7

Design

In this design chapter a description of a prototype solution which can show how the different strategies mentioned in chapter 6 perform on real data. The implemented solution will fulfil the role of *patient validator* from the system proposed in chapter 3. Furthermore it will obey the contents of the requirements stated in section 3.4 and the use cases in chapter 4.

7.1 Overall system design

The overall design of the system can be seen in figure 7.1. The implemented system will receive the raw medical data from the *device ID to patient ID translator* along with the patient who is matched to the devices' ID.

The implemented design will have to know the concept of a *signature* to be able to validate newly arrived data. The prototype will evolve around this concept as it is the most prominent feature of the prototype. A superficial description of the data structure is discussed.

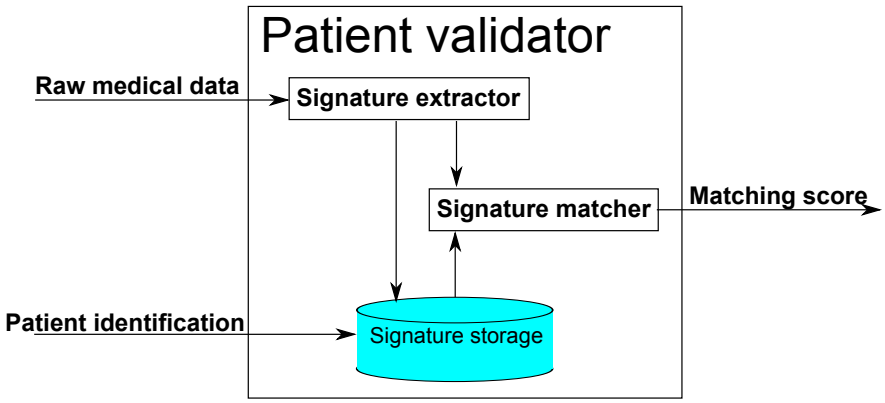

Figure 7.1: A detailed view of how the interaction between the different components is defined.

7.2 Expected input data

After the initial pre-processing of the input data it should be ready for further processing in the signature extractor and signature matcher.

The assumed input is the raw TCX files already separated into batches from each test subject. Both the data and the patient's ID are given to the system.

The data is simple in structure as it is in principle lists of data points composed of a time stamp and a single value.

7.3 Structure

Herein the structure of the system is explained in more detail. We assume that the surrounding system contains the components declared in chapter 3.

7.3.1 Signature extraction

The signatures must have a relative small footprint compared to storing the whole data set. This is not to limit the amount of storage as the complete data

set must be stored in the *EPJ database* but to ease the calculations needed to be done to validate the arrived data.

When the raw data arrives a signature must be extracted. This is done by dividing the data into *slices* as per chapter 6. This will enable us to match the patient with smaller portions of data as we get a more graduated picture of the data's pattern.

By extracting multiple features per slice we will get a better chance of validating the incoming data against the patient by using a simple average of the different scores of the features.

The extracted features must relate directly to the obtained data and improve as more data for a patient are obtained. This can be accomplished by storing the features average and the standard deviation value for each slice in the data. By using the average of these a good signature for the slice should be obtained. When enough data is collected one could use a sliding window average instead to disregard older data. The average could also be weighted so the older data is slowly phased out.

7.3.2 Signature storage

The *signature storage* should save the collected signatures extracted from the incoming data. This will enable us to match the different features against a larger and larger data basis with more confidence in the validation as a consequence.

A *dictionary type* with an underlying *database* would be the obvious choice for storing the extracted signatures. The *database* could be skipped when working with smaller amounts of data.

7.3.3 Signature matching

The *signature matching* component

Grouping the collected slice signatures from the same patient into lists of such signatures allows us to extract features on how the underlying spectrum of individual features is distributed. By looking at the average and standard deviation for both the average and standard distribution of the extracted features a signature of the patient's data emerges.

This *training set signature* relates directly to the arrived data's slice signatures meaning that they can be compared directly. By using this approach a slices signature feature can be validated using statistics and thereby extract a probability for each feature. The average of the probabilities can be considered to be the validation score or the confidence the system have that a given patient matches with the provided data.

7.4 Expected output data

The output of the program should be a score of how good a match a newly arrived data set is to the expected patient. A score is useful to pass on to another system which will be able to see if a poor score is due to changes in medication or environmental changes. If a true/false result was given it would be hard to distinguish between normal changes (a patient can have a very irregular pattern) and real patients switches. In a final implementation it should be possible for the GP to set some parameters on how much fluctuation would be considered normal.

Furthermore it would be useful if the signatures can be saved in an easy to use format.

Chapter 8

Implementation

The chosen implementation strategy is to have a very simple program capable of giving a quick overview of the calculated values for the given input data.

This is accomplished by having a very flat data structure exploiting the built-in features like LINQ in the .NET framework.

8.1 General Implementation Architecture

The prototype is implemented using Microsoft C# 4.0. The main argument for choosing this is previous experience with this framework and the capability to use Microsoft's LINQ framework [1]. Further C# is the fifth most used programming language according to [10]. All programming, structuring and debugging are done in Microsoft Visual Studio 2010 Ultimate.

A statistics library developed by Martin Bornhøft has been used and extended for calculating normal distribution, standard deviation and percentiles.

[1] Language-INtegrated Query extensions

8.2 Acquisition of data

Import of the data from the fitness watches are done manual as these use a proprietary USB protocol. To facilitate this import from the watches a combination of Garmin ANT Agent and Garmin Training Center was used.

From Garmin Training Center the workouts can be exported contained in XML files with the TCX extension.

8.2.1 TCX-format

The TCX format is an XML format containing information about physical workouts. The format designed for fitness software as it contains information like calories burnt during the workout, duration, distance, start time and heart rate information.

The root element of the XML file is *TrainingCenterDatabase*, named after the program use to generate it, Garmin Training Center. The root element contains two sub-elements; *Activities* and *Author*.

The *Author* element contains information about the program that generated the file, identifying various version and build numbers.

The *Activities* element contains *Activity* sub-elements .

Each *Activity*, as a workout is called in the TCX context, contains a number of *Tracks* containing *Trackpoints*. Each *Trackpoint* contains an UTC timestamp for the given measurement and the measurements itself.

8.3 Internal Data representation

8.3.1 PulseModel

A class called *PulseModel* gets initialised with a TCX file as an argument. The file is then parsed and the timestamps and heart rate data are stored internally in a *Dictionary*[2] type with the timestamp as the key and heart rate as the value.

[2] From the .NET framework's System.Collections.Generic namespace

The *PulseModel* class is used for parsing, containing the data in the program and doing very simple operations on the given data set like offsetting, slicing and smoothing the data. These features can be structured in a *SliceSignature* instance.

8.3.2 SliceSignature

The *SliceSignature* class is merely a structure for holding a number of feature parameters.

8.3.3 HeuristicSlice

The *HeuristicSlice* class is used like *PulseModel* to contain the raw data. From a *HeuristicSlice* instance one can get various features return based on the raw data.

The two most complex features of the *HeuristicSlice* are *Completeness* and *Error*.

Completeness calculates how much of the expected data the *HeuristicSlice* contains.

Error takes another *HeuristicSlice* as an argument. The absolute difference between each data point is summed up and divided with the number of data points contained. This is used for calculating how much a smoothing operation on a data set impacts the data.

8.3.4 TrainingSetSignature

The *TrainingSetSignature* class contains a Dictionary of normal distributions (*NDist* for each feature contained in the *HeuristicSlice* class with the start time for the *Heuristicslice* for which the data is extracted as the key.

This is the main signature to describe a test subject's data.

8.3.5 TrainingSet

The *TrainingSet* class is initialised with an instance of the *PulseModel* class. The *TrainingSet* constructor divides the data into smaller *PulseModel* instances each containing data for one day.

The constructor then runs each day's data through except for the last day contained in the data set, as this is used for validation of the trained model.

For each day used to train the model a smoothing using a low pass filter is performed. This results in a new smoother *PulseModel* which is sliced into 10 even sized slices using Time Slicing (see figure 6.2). For comparison the original non-smoothed *PulseModel* is sliced in the same manner.

For each slice its *SliceSignature* is extracted and added to a Dictionary containing a list of *SliceSignatures* for each start time of the slices. However if the completeness of the *HeuristicSlice* is below a set threshold the data is not added to the dictionary.

After the training signatures have been added to the *TrainingSet* the last day in the given data set is used to validate how good the *TrainingSet* has been taught. This *ValidationResult* is stored in the *TrainingSet* instance.

Furthermore the average amount of data correction per data point the applied filter inflicted is held in the *AvgSmoothingError* member.

From the *TrainingSet* a *TrainingSetSignature* can be calculated.

8.4 Test program flow

In the developed prototype and test program it is desirable to have a simple linear flow.

The program expects the collected TCX data to be located in a specific folder and that the data files are prefixed with a single letter indicating the test subject.

Firstly all the raw TCX-files for each test subject is loaded into a *PulseModel*. After all the test subjects data has been loaded a *TrainingSet* is generated from this *PulseModel*. The *TrainingSet* is then stored in a Dictionary with the test subject's letter as the key.

The test program then tries to match each *SliceSignature* in each *TrainingSet* with all other *TrainingSet*s. These results are save into a number of CSV[3] files named after the test subject.

A number of parameters are written to another CVS file to be used in a exploratory data analysis in section 5.9.

Each test subjects signature is saved to an appropriate file for later usage.

[3] Comma-separated values

Chapter 9

Evaluation

With the developed prototype at hand it is very easy to test a number of hypotheses to evaluate the proposed methods for validating medical data against a trained model.

9.1 Size of data basis

By training the system with the available data except the last day an independent check of the data could be performed.

The results from this can be seen in table 9.1

The expected validation results would show a high score to validate that the implemented system could positively validate patients with their own data.

With a validation score threshold of 75% for a confirmed match, the data shows that a validation could not be performed with confidence except for test subject C which is the data set containing most data points.

In figure 9.1 a plot of the correlation between the number of collected data points per test subject and the resulted validation score. This suggests that

Data set	Match score
A	43.01%
B	38.70%
C	87.29%
D	42.62%
F	51.94%
G	24.77%
H	N/A
I	63.46%
J	20.67%
K	69.02%

Table 9.1: Table showing the match between data from the same patient.

the amount of collected data in these experiments had not been sufficient to validate a patient with enough confidence for practical purposes. From the plot it seems that a validation score of 75% would be achieved with around 50.000 data points in the training set for each test subject.

Assuming a completeness of data of 70% and sampling every 5 seconds an estimated minimum of 4.1 complete days should be collected and used for the learning set for a patient as seen in equation 9.1.

$$\frac{50000}{70\%} * \frac{5}{60 * 60 * 24} = 4.1336 \tag{9.1}$$

9.2 Validation of use cases

To validate that the prototype can honour the set use cases in section 4.4 a number of tests are performed in the following.

9.2.1 Use case: Learning patterns from a patient

The system was trained with all available data for a given test subject except the last day period and hereafter was the last days data validated against the trained system.

This resulted in the generation of a signature for the given test subject. An

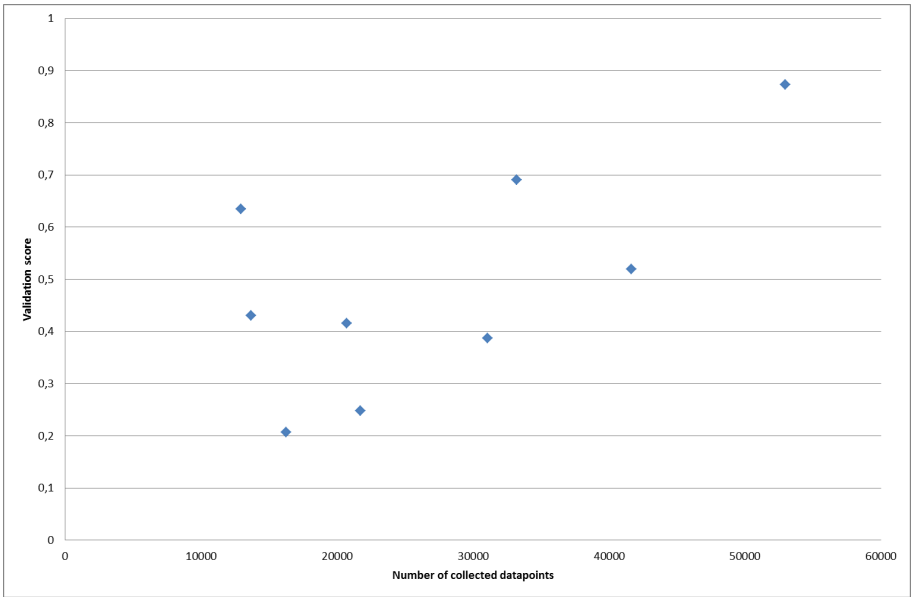

Figure 9.1: Number of data points per test subject plotted against the validation score in table 9.1

example of such a signature can be seen in appendix D.

9.2.2 Use case: Matching new data with existing patterns from the same patient

In table 9.1 the results from matching data which is not contained in the already learnt data is matched against the trained set. This use case did not pass as the data basis for the tests were insufficient as discussed in section 9.1

9.2.3 Use case: Matching new data with existing patterns from a different patient

In appendix E a table showing the validation scores for each training set against each other.

If the signature for test subject A is compared with test subject F (the two test subjects who has delivered the most data points) an average score of 16.9%

validation score which must be considered to be an indicator of a patient switch. The other way around is the average score 32.4%. This would result of a rejection or alert to be sent via the alerting service to the GP.

In table 9.3 the available data set is cross matched with each other. Note that the matches of the same test subjects data set is based on a signature containing the data being matched against and therefore biased! Also note that the validation score A to B is not the same as B to A. The difference in the two directions is because of one signature accepting a more broad spectrum of values than the other.

It can be seen in the resulting data that with only one exception a rejection of data would have taken place with a validation score threshold of 75% (marked with **bold** in the table).

The exception where test subject C validated against test subject J with a score of 82.69% must be attributes to poor quality of the data from test subject J. The data set is the third smallest collected and the data set with the lowest validation score against itself. This points to that test subject J had an irregular heart beat pattern on the observed days. After manual inspection of the recorded data a fairly low completeness is marked in the data.

9.3 Validation with tachycardic patients

Ideally a number of patients suffering from tachycardic, or in some other way had an unique heart rate pattern, had delivered heart rate measurements for use in this investigation. This would be nearly impossible to get such data unless one or more of the test subjects by coincidence suffered tachycardia.

To validate that the developed system will distinguish patients suffering from tachycardic caused by e.g. infection and fever a data set resembling the symptoms of such a patient was generated. This was done by adding a fixed offset to all data points contained in a data set obtained from a test subject. By adding different amounts various degrees of infection can be simulated.

Using the data set from test subject C three additional data sets were generated with an offset of 10, 20 and 30 BPM. These were then validated against the other collected data sets. The validation scores can be seen in table 9.2. Test subject C's original data is included in the test suite to validate that the signature will change.

Data set	Offset +10	Offset +20	Offset +30
A	26.78%	18.78%	16,40%
B	26.78%	22.52%	18.51%
C	46.47%	33.67%	28.27%
D	40.36%	30.05%	24.66%
F	38.90%	29.81%	24.94%
G	27.44%	22.58%	17.96%
I	51.74%	36.42%	27.38%
J	54.13%	32.14%	26.10%
K	59.06%	56.21%	37.78%

Table 9.2: Table showing the match between generated data and the different test subjects.

From the data in table 9.2 the validation score falls, as expected, with progressively larger deviation from the original.

	A	B	C	D	F	G	I	J	K
A	**100.00%**	19.00%	14.11%	21.30%	16.93%	24.20%	18.45%	18.64%	14.98%
B	43.85%	**89.87%**	42.85%	51.58%	40.96 %	43.62%	42.38%	46.28%	27.60%
C	30.86%	50.28%	**78.98%**	50.00%	62.26%	33.14%	59.78%	**82.69%**	52.46%
D	35.60%	37.51%	34.92%	**87.37%**	33.09%	27.18%	55.02%	52.30%	33.81%
F	42.72%	53.75%	60.35%	49.51%	74.01%	44.80%	50.08%	64.77%	38.43%
G	24.20%	19.22%	14.38%	11.86%	17.16%	**100.00%**	7.86%	12.94%	12.14%
I	18.45%	12.58%	29.02%	26.42%	23.18%	7.86%	**100.00%**	35.80%	20.40%
J	18.64%	30.61%	44.86%	31.75%	32.49%	12.94%	35.80%	**100.00%**	31.49%
K	15.05%	19.35%	31.76%	32.87%	25.23%	9.64%	35.34%	42.99%	**80.31%**

Table 9.3: Table showing the match between data from different test subjects.

Chapter 10

Future extension

If time allowed it would be obvious to collect more data to have a substantial data basis to work with. This would include a number of test subjects monitored continuously for preferably 5 days or more. The results derived from more data would be very interesting as the they would be more clear and conclusive than what was discovered in this thesis.

To improve the performance of the signature matching algorithm one could look into doing something to eliminate the constrains of using fixed lengths of slices or one could evaluate the score of a slice using the neighbouring slices as well.

A future extension to investigate would be to implement machine learning to enhance the pattern recognition. Neural networks or genetic programming would be two interesting approaches to an extension of this project.

Instead of using heuristic signatures another type of signature should be extracted from the raw data, perhaps by exploiting how fast a test subject's heart rate can rise and fall as this is an indicator of the test subject's physical condition.

Future work could include a solution with the algorithm analysing data live (telemetrically) or directly in the sampling device, as this would move the project from a theoretical project to a "proof of concept" directly testable with real

patients.

To enhance the patient validation another sensor than a heart rate monitor could be chosen as this parameter did not prove itself as a reliable validator. It would be interesting to see if blood pressure could be more distinct for a patient. This is however doubtful as both the systolic and diastolic pressure can and will span over a large range of values.

Another approach to combat the similarity of the inputted medical data would be to receive input from more than one sensor or sensor-type. This would enable us to create a more distinct signature/profile for a given patient and thereby increase the certainty that the correct patient is using the provided medical device.

Chapter 11

Conclusion

As mentioned in the introduction the main challenges with the increased use of telemedicine in the public health sector are the risk of misdiagnosis and mismedication due to failure to relate collected data with the correct patient.

The state of the art in telemedicine is presented with a case study of a trial current done at Frederiksberg Hospital. This survey shows that both the health care system as well as the patients can benefit from telemedicine; the health care system can save money and medical personnel on hospitals utilising telemedicine and the patients gain a higher quality of life during hospitalisations which is foremost an advantage for people suffering from chronic diseases.

To investigate how to minimise this risk of misdiagnosis and mismedication a proposed telemedical system is analysed from a bird's eye view. The proposed systems different components and their roles are analysed with emphasis on the receiving and first processing of medical data. From this analysis the requirements for a *patient validator* are extracted.

With a more in-depth analysis of the *patient validator* and its requirements the concept of *signatures* for a patient's data is clarified. This allows us to investigate and experiment with signature extraction from collected data from test subjects.

With the collected data it has not been possible to make any conclusions regarding the possibility of validating received measurements on the basis of previous data. The main reason for this is the lack of sufficient data. This can be attributed to the test subjects quitting during the test period due to discomfort of the provided heart rate monitor belts.

A prototype of a heuristic approach is developed to test and experiment with different types of signature extraction. During this process the best performance for the simplest solution is concluded to the time slicing.

This thesis analyses and investigates the urgent issue of validating data obtained telemedically from medical devices to prevent inappropriate medication being prescribed to a patient. A number of approaches are discussed and evaluated with the concept of signatures for smaller segments of medical data being validated against a model trained with previously obtained data. By using this approach the collected data from each test subject can be matched against each other and with other test subject's data. A coherence between the amount of data collected and the confidence of matching can be observed.

To enhance the confidence and reliability of the system more data is needed to reach an absolute conclusion whether or not the developed solution is based on a viable approach.

Appendix A

Test subjects

This appendix lists the various information on the test subjects.

ID	Gender	Age (years)	Height (cm)	Weight (kg)	BMI	Activity level	Data points
A	M	28	186	84.0	24.28	4.0	13694
B	F	45	157	50.0	20.28	3.5	31087
C	M	28	180	82.5	25.46	2.0	52972
D	M	30	172	73.0	24.68	3.0	20704
E	M	17	181	73.0	21.37	4.0	0
F	M	24	195	95	24.98	3.0	41624
G	M	45	180	77	23.77	4.5	21733
H	M	33	177	86	27.45	2.0	8960
I	M	45	180	77	23.77	2.0	12959
J	F	60	172	60	20.28	2.5	16242
K	F	36	167	60	21.51	2.5	33211

BMI = Body Mass Index

Appendix B

Garmin Forerunner 60 Specifications

This appendix lists the specifications as informed by the manufacturer.

Item	Value
Unit dimensions, WxHxD:	5.6 x 3.8 x 1.3 cm
Display size, WxH:	2.0 x 2.8 cm
Display resolution, WxH:	56 x 31 pixels
Touchscreen:	no
Weight:	44 g
Battery:	coin cell battery (CR2032)
Battery life:	1 year
Water resistant:	yes (50m)
GPS-enabled:	no
Heart rate monitor:	yes

Source: Garmins website [2]

Appendix C

Data sample

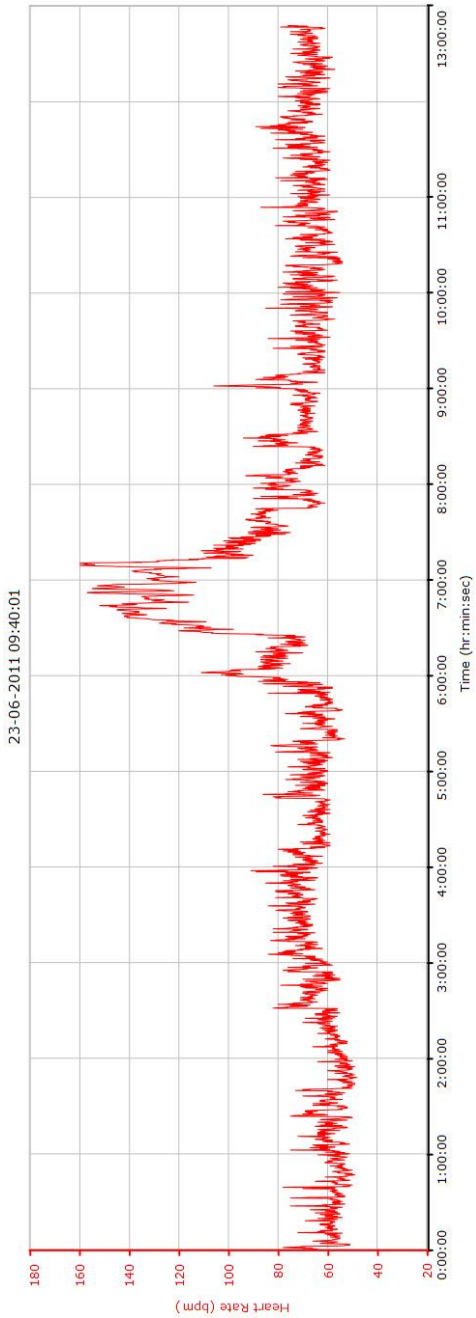

Figure C.1: An example of collected data.

Appendix D

Signature sample

	Minimum		Maximum		Mean		Std. dev.		10th percentile		90th percentile	
Slice	Mean	SD.	Mean	SD.	Mean	SD.	Mean	SD.	Mean	SD.	Mean	SD.
00:00:00	62,000	3,000	92,000	3,000	71,694	3,000	3,858	3,000	72,694	3,000	70,694	3,000
02:24:00	53,000	3,000	95,000	3,000	64,036	3,000	4,247	3,000	62,534	3,000	65,537	3,000
04:48:00	53,000	3,000	93,000	3,000	63,194	3,000	5,885	3,000	58,004	3,000	68,383	3,000
07:12:00	60,000	3,000	123,000	3,000	77,235	3,000	10,575	3,000	62,470	3,000	92,000	3,000
09:36:00	52,667	4,509	83,333	6,028	63,279	5,617	5,476	0,669	58,949	4,271	67,608	6,988
12:00:00	55,333	1,528	94,333	10,214	69,434	2,580	5,988	0,970	64,053	1,526	74,815	4,323
14:24:00	54,000	1,732	115,000	36,387	73,130	12,018	14,065	13,272	49,402	18,324	96,859	42,317
16:48:00	56,500	10,607	146,500	40,305	85,976	24,137	24,688	20,060	37,194	24,182	134,758	72,456
19:12:00	65,000	3,000	113,000	3,000	73,769	3,000	6,248	3,000	67,850	3,000	79,687	3,000
21:36:00	55,000	3,000	93,000	3,000	68,402	3,000	6,283	3,000	62,413	3,000	74,391	3,000

Table D.1: The extracted signature for test subject B *SD.* = *Standard deviation*

Appendix E

Results

Here follows all the calculated validation/matching scores for all available data. These data can also be found on the included DVD (Appendix F) in CSV format.

Matching for data set A using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	1,000	N/A	N/A	N/A	1,000	1,000	1,000	1,000	1,000	N/A
B	0,072	NaN	NaN	NaN	0,332	0,239	0,073	0,148	0,276	NaN
C	0,126	NaN	N/A	NaN	0,180	0,138	0,147	0,023	0,233	NaN
D	0,374	N/A	N/A	N/A	0,437	0,135	0,076	0,097	0,159	NaN
F	0,283	N/A	N/A	NaN	0,166	0,126	0,099	0,008	0,334	NaN
G	N/A	N/A	N/A	N/A	N/A	0,655	0,263	0,000	0,050	NaN
I	N/A	N/A	N/A	N/A	N/A	0,250	0,050	0,191	0,247	NaN
J	N/A	N/A	N/A	NaN	0,248	0,175	0,142	0,006	0,361	NaN
K	0,113	N/A	N/A	NaN	0,181	0,122	0,144	0,065	0,274	NaN

Matching for data set B using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	0,072	N/A	N/A	N/A	0,353	0,226	0,797	0,907	0,276	N/A
B	1,000	1,000	1,000	1,000	0,737	0,736	0,735	0,779	1,000	1,000
C	0,166	0,709	N/A	0,421	0,329	0,103	0,695	0,749	0,264	0,421
D	0,248	N/A	N/A	N/A	0,839	0,449	0,882	0,655	0,129	0,409
F	0,159	N/A	N/A	0,362	0,239	0,189	0,878	0,767	0,368	0,315
G	N/A	N/A	N/A	N/A	N/A	0,081	0,958	0,834	0,076	0,232
I	N/A	N/A	N/A	N/A	N/A	0,213	0,821	0,871	0,027	0,187
J	N/A	N/A	N/A	0,338	0,507	0,240	0,855	0,705	0,260	0,335
K	0,214	N/A	N/A	0,250	0,003	0,000	0,563	0,658	0,234	0,286

Matching for data set C using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	0,126	N/A	N/A	N/A	0,559	0,292	0,513	0,027	0,335	N/A
B	0,166	0,709	NaN	0,595	0,752	0,501	0,582	0,198	0,351	0,672
C	1,000	1,000	N/A	0,706	0,747	0,749	0,726	0,714	0,733	0,733
D	0,130	N/A	N/A	N/A	0,746	0,627	0,711	0,138	0,307	0,841
F	0,436	N/A	N/A	0,648	0,858	0,720	0,822	0,500	0,450	0,547
G	N/A	N/A	N/A	N/A	N/A	0,364	0,673	0,334	0,084	0,202
I	N/A	N/A	N/A	N/A	N/A	0,831	0,832	0,150	0,274	0,902
J	N/A	N/A	N/A	0,717	0,869	0,704	0,794	0,783	0,945	0,976
K	0,078	N/A	N/A	0,043	0,682	0,803	0,823	0,388	0,673	0,707

Matching for data set D using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	0,374	N/A	N/A	N/A	0,437	0,306	0,242	0,517	0,283	N/A
B	0,248	NaN	NaN	NaN	0,534	0,683	0,143	0,408	0,201	0,409
C	0,130	NaN	N/A	NaN	0,277	0,520	0,243	0,485	0,424	0,366
D	1,000	N/A	N/A	N/A	1,000	0,779	0,779	0,779	0,779	1,000
F	0,249	N/A	N/A	NaN	0,114	0,582	0,292	0,416	0,222	0,441
G	N/A	N/A	N/A	N/A	N/A	0,419	0,326	0,306	0,228	0,080
I	N/A	N/A	N/A	N/A	N/A	0,644	0,252	0,691	0,851	0,313
J	N/A	N/A	N/A	NaN	0,427	0,835	0,221	0,502	0,396	0,757
K	0,237	N/A	N/A	NaN	0,087	0,476	0,138	0,495	0,468	0,466

Matching for data set F using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	0,928	N/A	N/A	N/A	0,337	0,255	0,333	0,024	0,686	N/A
B	0,462	NaN	NaN	0,237	0,505	0,694	0,315	0,427	0,887	0,773
C	0,762	NaN	N/A	0,389	0,633	0,660	0,478	0,767	0,581	0,558
D	0,747	N/A	N/A	N/A	0,446	0,775	0,263	0,200	0,256	0,779
F	0,779	N/A	N/A	0,779	0,727	0,710	0,737	0,726	0,735	0,728
G	N/A	N/A	N/A	N/A	N/A	0,351	0,542	0,735	0,303	0,309
I	N/A	N/A	N/A	N/A	N/A	0,734	0,841	0,294	0,133	0,502
J	N/A	N/A	N/A	0,033	0,749	0,976	0,743	0,667	0,711	0,655
K	0,243	N/A	N/A	0,000	0,311	0,525	0,353	0,376	0,528	0,738

Matching for data set G using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	NaN	N/A	N/A	N/A	NaN	0,655	0,263	0,000	0,050	N/A
B	NaN	NaN	NaN	NaN	NaN	0,236	0,160	0,257	0,076	0,232
C	NaN	NaN	N/A	NaN	NaN	0,168	0,088	0,230	0,061	0,172
D	NaN	N/A	N/A	N/A	NaN	0,154	0,102	0,119	0,138	0,080
F	NaN	N/A	N/A	NaN	NaN	0,138	0,205	0,323	0,010	0,182
G	N/A	N/A	N/A	N/A	N/A	1,000	1,000	1,000	1,000	1,000
I	N/A	N/A	N/A	N/A	N/A	0,210	0,052	0,074	0,009	0,048
J	N/A	N/A	N/A	NaN	NaN	0,244	0,030	0,270	0,000	0,103
K	NaN	N/A	N/A	NaN	NaN	0,210	0,087	0,099	0,084	0,127

Matching for data set I using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	NaN	N/A	N/A	N/A	NaN	0,250	0,050	0,191	0,247	N/A
B	NaN	NaN	NaN	NaN	NaN	0,240	0,163	0,012	0,027	0,187
C	NaN	NaN	N/A	NaN	NaN	0,386	0,381	0,141	0,288	0,255
D	NaN	N/A	N/A	N/A	NaN	0,243	0,058	0,244	0,463	0,313
F	NaN	N/A	N/A	NaN	NaN	0,220	0,412	0,147	0,100	0,280
G	N/A	N/A	N/A	N/A	N/A	0,210	0,052	0,074	0,009	0,048
I	N/A	N/A	N/A	N/A	N/A	1,000	1,000	1,000	1,000	1,000
J	N/A	N/A	N/A	NaN	NaN	0,466	0,638	0,001	0,152	0,533
K	NaN	N/A	N/A	NaN	NaN	0,198	0,157	0,049	0,292	0,324

Matching for data set J using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	NaN	N/A	N/A	N/A	0,248	0,175	0,142	0,006	0,361	N/A
B	NaN	NaN	NaN	0,338	0,522	0,422	0,185	0,081	0,260	0,335
C	NaN	NaN	N/A	0,561	0,495	0,386	0,329	0,440	0,539	0,390
D	NaN	N/A	N/A	N/A	0,427	0,340	0,119	0,025	0,237	0,757
F	NaN	N/A	N/A	0,270	0,376	0,291	0,373	0,231	0,378	0,355
G	N/A	N/A	N/A	N/A	N/A	0,244	0,030	0,270	0,000	0,103
I	N/A	N/A	N/A	N/A	N/A	0,466	0,638	0,001	0,152	0,533
J	N/A	N/A	N/A	1,000	1,000	1,000	1,000	1,000	1,000	1,000
K	NaN	N/A	N/A	0,226	0,170	0,235	0,151	0,431	0,553	0,438

Matching for data set K using percentiles

Set	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10
A	0,113	N/A	N/A	N/A	0,190	0,000	0,133	0,230	0,237	N/A
B	0,214	NaN	NaN	0,250	0,102	0,017	0,099	0,101	0,196	0,569
C	0,078	NaN	N/A	0,177	0,222	0,284	0,240	0,298	0,642	0,569
D	0,237	N/A	N/A	N/A	0,069	0,118	0,140	0,197	0,567	0,973
F	0,114	N/A	N/A	0,113	0,271	0,160	0,223	0,255	0,311	0,572
G	N/A	N/A	N/A	N/A	N/A	0,004	0,054	0,109	0,166	0,149
I	N/A	N/A	N/A	N/A	N/A	0,217	0,228	0,276	0,411	0,635
J	N/A	N/A	N/A	0,226	0,203	0,071	0,190	0,577	0,792	0,950
K	1,000	N/A	N/A	1,000	0,737	0,737	0,736	0,716	0,720	0,779

Appendix F

Attached material

F.1 CD

All materials are attached on this CD.

The material consists of:

- Source code of prototype
- Obtained heart monitor data
- Generated signatures

Bibliography

[1] June Edhouse. Broad complex tachycardia-part i. *BMJ 324 : 719 doi: 10.1136/bmj.324.7339.719*, 2002.

[2] Garmin. Garmin fr60 specs. *https://buy.garmin.com/shop/shop.do?pID=27483*, 2011.

[3] Anthony S Hale. Abc of mental health: Anxiety. *BMJ 314 : 1886*, 1997.

[4] Michael Højgaard Bojsen. Napilink. 2011.

[5] Time Magazine. New lessons from the old world. *http://www.time.com/time/magazine/article/0,9171,1901483-1,00.html*, 2009.

[6] Ketil Stølen Mass Soldal Lund, Bjørnar Solhaug. *Model-Driven Risk Analysis*. Springer, 2011.

[7] Dagens medicin (Danish). Jeg vil kunne gå stuegang fra sommerhuset. *http://www.dagensmedicin.dk/nyheder/2010/05/21/jeg-vil-kunne-ga-stuegang-/*, 2010.

[8] Sundhedsstyrelsen / National Board of Health. Early detection, follow-up, treatment and rehabilitation. *http://www.sst.dk/English/Health_promotion/COPD.aspx*, 2011.

[9] Richard Irons Steve Goodacre. Atrial arrhythmias. *BMJ 324 : 594 doi: 10.1136/bmj.324.7337.594*, 2002.

[10] TIOBE. Tiobe programming community index 2011. *http://www.tiobe.com/index.php/content/paperinfo/tpci/index.html*, 2011.