



An autonomous mobile system for the management of COPD

Maarten van der Heijden^{a,b,*}, Peter J.F. Lucas^a, Bas Lijnse^a, Yvonne F. Heijdra^c, Tjard R.J. Schermer^b

^a Institute for Computing and Information Sciences, Radboud University Nijmegen, The Netherlands

^b Department of Primary and Community Care, Radboud University Nijmegen Medical Centre, The Netherlands

^c Department of Pulmonary Diseases, Radboud University Nijmegen Medical Centre, The Netherlands

ARTICLE INFO

Article history:

Received 22 November 2012

Accepted 4 March 2013

Available online 15 March 2013

Keywords:

Chronic disease management

Bayesian networks

Decision support systems

eHealth

ABSTRACT

Introduction: Managing chronic disease through automated systems has the potential to both benefit the patient and reduce health-care costs. We have developed and evaluated a disease management system for patients with chronic obstructive pulmonary disease (COPD). Its aim is to predict and detect exacerbations and, through this, help patients self-manage their disease to prevent hospitalisation.

Materials: The carefully crafted intelligent system consists of a mobile device that is able to collect case-specific, subjective and objective, physiological data, and to alert the patient by a patient-specific interpretation of the data by means of probabilistic reasoning. Collected data are also sent to a central server for inspection by health-care professionals.

Methods: We evaluated the probabilistic model using cross-validation and ROC analyses on data from an earlier study and by an independent data set. Furthermore a pilot with actual COPD patients has been conducted to test technical feasibility and to obtain user feedback.

Results: Model evaluation results show that we can reliably detect exacerbations. Pilot study results suggest that an intervention based on this system could be successful.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

Increasing demands on health-care and continuous pressure from health-care authorities and insurance companies to reduce costs while maintaining quality of care has created a situation in which automation of particular parts of the patient's care process has attracted attention. Especially the provision of computer-aided assistance in the management of the patient's diseases is an attractive option. In the context of chronic diseases, patients are continuously at risk of deterioration of health, requiring regular medical checkups and monitoring of their health status by the treating medical doctor. Providing computer-aided support to the patient can relieve work-load of health-care workers, while helping patients self-manage their disease. However, the provision of computer-aided support to patients poses questions with respect to whether or not patients are able to profit from the support, how the patient data needed for that purpose can be collected and interpreted, and which technical infrastructure is most effective.

The idea to offer computer-aided support to patients at a distance from the treating hospital or practice is not new. Remote care facilitated by telecommunication technology has existed for some time already under different names such as: 'telehealth', 'telemedicine'

* Corresponding author at: Institute for Computing and Information Sciences, Radboud University Nijmegen, The Netherlands.

E-mail addresses: m.vanderheijden@cs.ru.nl (M. van der Heijden), peterl@cs.ru.nl (P.J.F. Lucas).

and more recently 'eHealth' (electronic health) and 'mHealth' (mobile health). The exact scope and definition of all these terms varies, as exemplified by the review of Oh et al. [1] that identifies 51 definitions of the term 'eHealth'. Although different definitions will place different emphases we can summarise the definitions as *facilitating health care irrespective of location by means of technology*. Practically speaking this means that information and communication technology assists in disease management, patient-doctor communication, patient education or any other application that promotes health.

The research described in this paper aimed at developing methods for computer-aided assistance, including event detection, alerting, monitoring and treatment advice, as part of chronic disease management at a distance from the hospital. Patients with diseases such as diabetes mellitus, chronic obstructive pulmonary disease (COPD), asthma and heart failure can benefit from assistance. We have applied it here to assist COPD patients, but the general framework is applicable to many other health-care situations. For COPD we collect data on respiratory symptoms, measure lung function and interpret these data by a probabilistic model to assess the risk of a clinically relevant worsening of symptoms due to an exacerbation.

The main challenge addressed by the research was to develop a computer-aided disease management framework that allows finding a proper balance between self-management by the patient and support in various forms by health-care workers. This depends on the patient's demands and wishes, the disease being managed and

the requirements of health-care workers. One of the implications of the resulting requirements was that disease-related patient data had to be interpreted automatically as part of the patient's self-management support. The system had to be capable of delivering autonomous assistance in disease management without being intrusive. Hence we decided for data monitoring and interpretation directly on a smartphone. Thus, the smartphone became the device of choice, enabling instant feedback to the patient. To cover the whole spectrum from complete self-management to distance support by health-care workers, collected patient data can also be relayed to health-care providers through a central server for both disease management configuration and inspection.

Use of a smartphone makes it possible to forego the need for a personal computer (PC) with internet connectivity. This has the advantage that whereas most people are used to responding to phone alerts, sending reminders via email or a website may have little effect on the patient's behaviour. This is also very much in the spirit of health care that is no longer tied to a particular location.

If we compare this to usual care we see that patients have regular but fairly sporadic contact with their physician unless there is an acute reason for an unscheduled visit. Our system is capable of advising patients to take measures at an earlier stage than usual and enables easier communication between patients and health-care providers. Although parts of what the system provides has been used before in the context of eHealth or decision-support systems meant for health-care workers as user, the idea to place both healthcare users and workers on an equal footing by offering patients sophisticated, mobile decision support is new. The more common eHealth systems are mainly focussed on sending small sets of measurements from the patient to health-care providers, who still have to interpret these findings [2]. In the system described below, clinical knowledge is incorporated by means of a probabilistic graphical model in the smartphone, making it possible to provide relevant clinical advice automatically, to assist in patient self-management.

The following questions are explored in the rest of the paper. It is investigated in what way patients can be empowered with disease management assistance to prevent exacerbations of COPD. Another relevant question is whether it is feasible to automatically interpret monitoring data by probabilistic models to detect clinically relevant events. This question is explored by means of statistical model validation methods. As a prerequisite we examine choices with respect to the important features of the disease management system in terms of hardware and software. Finally, the usability of the system was investigated with the help of COPD patients. Summarising, we report on a system for COPD exacerbation management, that has the novel feature of including automatic data interpretation by a probabilistic risk model, enabling autonomous operation to support patient self-management.

In the next section, we start with some background on the clinical problem addressed in the research. Then the architecture of the disease management system is described in Section 3, focusing on its design and technical capabilities. Furthermore we report on a pilot study that investigates technical and clinical feasibility with a number of patients. The probabilistic model that is used for data interpretation is explained and evaluated in Section 5. Then in Section 6 we compare our system with telemonitoring requirements laid out by Peirce et al. [3] and with existing work on COPD telehealth, followed by a general discussion in Section 7.

2. Chronic obstructive pulmonary disease

Chronic obstructive pulmonary disease, or COPD for short, is a chronic lung disease with high impact on patient well-being and with considerable health-care related costs [4]. Exacerbations – acute events of worsening of COPD-related health status – are important events in the progression of COPD, such that monitoring

patients in a home setting to detect exacerbation onset may be warranted [5]. In this paper we focus on disease management and specifically on detecting and managing the occurrence of exacerbations of COPD at an early stage. We aim to decrease the impact of COPD on the patient's quality of life, and prevent unscheduled doctor visits and hospitalisation due to exacerbations.

Chronic obstructive pulmonary disease is estimated to affect some 64 million people worldwide¹ and is one of the major chronic diseases in terms of both morbidity and mortality. COPD affects the respiratory system, decreasing lung capacity and obstructing airways, thus interfering with normal breathing. Patients often suffer from a combination of emphysema and chronic bronchitis, causing shortness of breath and therefore reducing their capability of performing day-to-day activities. The main cause of COPD is prolonged exposure to tobacco smoke; other causes include severe air pollution. COPD is currently not curable, but treatment does reduce the burden considerably. For further information on COPD see e.g. the Global Initiative for Chronic Obstructive Lung Disease (GOLD).²

An important aspect of COPD which is particularly relevant in the present context is the progressive nature of the disease. Specifically exacerbations have a profound impact on the patient's well-being and on health-care costs [4]. These exacerbations are mainly caused by airway infections resulting in symptom worsening [6]. Important to note is also that patients with frequent exacerbations usually have faster disease progression, which makes exacerbation prevention a particularly relevant goal. Additionally, a faster treatment response to exacerbations leads to better recovery [7]. We can distinguish different clinical approaches to defining exacerbations. Due to limited observability it is not feasible to give a practical definition in terms of pathophysiology, hence exacerbations are usually defined as an increase in symptoms; in terms of use of medication; or in terms of unscheduled health care use. We will return to this point in Section 5.

The state of the respiratory system is observable via symptoms including dyspnea (breathlessness), productive cough, wheezing breath and decreased activity due to breathlessness. Besides these symptoms a number of physiological signs are relevant, in particular the forced expiratory volume in 1 s (FEV₁) and blood oxygen saturation. FEV₁ measures airway obstruction by testing to what extent the patient can overcome obstructive and restrictive resistance during forced exhalation. A number of other indicators of deterioration exist, like blood oxygen pressure, inflammatory proteins and white blood-cell counts. However, measuring these factors requires hospital-grade equipment and incurs considerable inconvenience for the patient. Blood oxygen pressure can be observed by proxy with a pulse-oximeter that measures blood oxygen saturation.

3. Remote disease management

The long term nature of COPD and associated exacerbation risk require that any system, deployed in a home-care setting, takes into account not only efficacy, but also usability as important factors in the design. This section describes the current system design, the choices we made and some of the issues that arose during the development and implementation.

3.1. System overview

3.1.1. General architecture

In Fig. 1 a graphical representation of the general idea behind our disease management setup is shown. The system consists of

¹ World Health Organization <http://www.who.int/mediacentre/factsheets/fs315/en/index.html>. Accessed: January 2013.

² www.goldcopd.com.

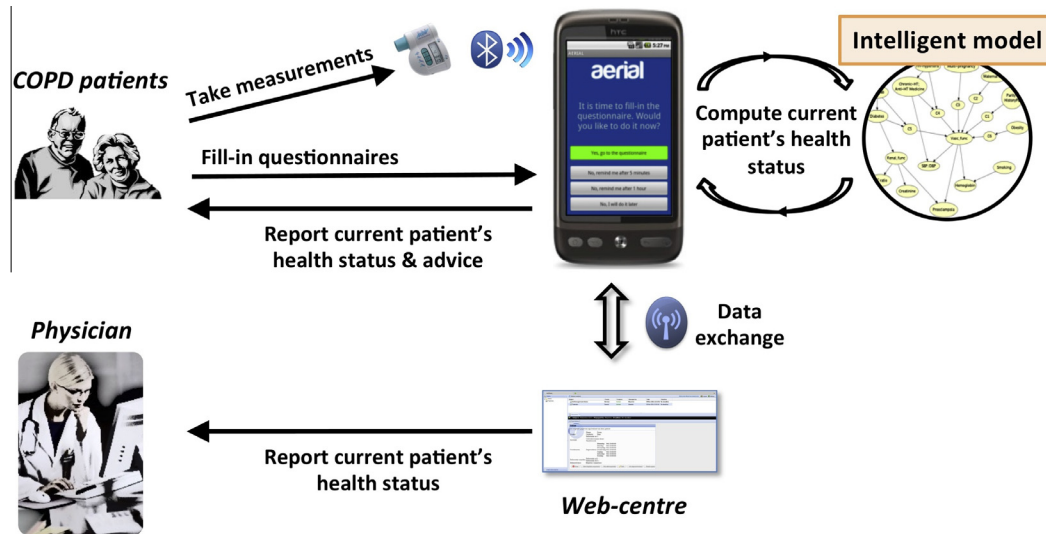


Fig. 1. Schematic of the system setup.

a smartphone as the main component taking care of communication and computation. Sensors are used to obtain objective information on the patient's health status, transmitted wirelessly to the smartphone. Questionnaire data is collected from the patient on the smartphone. A web-based system allows scheduling tasks and collecting patient data centrally. The web-centre receives the data from the smartphone and provides data access for health-care workers. Patient data are interpreted in the smartphone by means of a disease-specific probabilistic model that incorporates clinical variables considered relevant by a clinical specialist.

3.1.2. The COPD-specific system

The system we used for COPD management is a specific instantiation of the general system just described. For COPD we were interested in measuring lung function and blood-oxygen levels to complement the recorded clinical symptoms. To do so we employed a sensor interface that took care of the Bluetooth communication and to which a micro-spirometer and pulse-oximeter were connected.

Before going into more detail on the various components let us first describe the monitoring process. At regular intervals, adjustable in frequency and in time of the day, the patient gets an automatic reminder for data entry from the smartphone. The patient is presented with a simple touch-interface to answer a set of questions about COPD symptoms and is subsequently asked to perform a spirometry test and pulse-oximeter measurement. The results of the measurements are transmitted to the phone, and entered in a Bayesian network model to determine the probability of an exacerbation. In addition, the data is synchronised with the web-centre, which allows the responsible health-care workers to examine the patient data; depending on the situation this may be a nurse specialised in lung diseases, general practitioner or pulmonologist. If necessary, the patient can be advised to take action, based on the model's prediction.

Monitoring patients for relapse or a new occurrence of exacerbation provides valuable additional information compared to usual care, as exacerbations are often underreported, with patients only contacting their physician when the exacerbation is already quite severe. For the population of patients with regular exacerbations, this intervention provides a means for the patients themselves and for health-care professionals to stay better informed on the patient's health status and to act accordingly.

3.2. Design considerations

We now discuss a number of design issues that are relevant for automatic disease management. Although we focus on the requirements for our COPD system, these points also have more general value for similar systems.

3.2.1. Population

Chronic disease management is by its very nature a long term effort, which makes it particularly important to find a careful balance between costs and benefits. The intrusiveness of monitoring systems and costs, both monetary and in terms of patient time investment, require a precise definition of the target population. We argue that those that stand to gain the most from our COPD system are the patients with moderate to severe COPD and frequent exacerbations (more than 2 per year). These patients suffer greatly from the consequences of exacerbations, hence providing regular data to detect exacerbations in an early stage will in general be more acceptable. The most appropriate time to start the intervention would be directly after emergency treatment, because at that time the goal of preventing hospitalisation is clearly relevant for the patient. To increase the chance of success in general, psychological factors associated with intervention acceptance and technology acceptance in particular should be taken into account when introducing a disease management system.

3.2.2. Security

The system can operate autonomously in which case no data is transmitted. If data is communicated to health-care providers, all data transmission from the smartphone to the server is encrypted (HTTPS) to protect the privacy sensitive nature of the data. Also access-rights to the data in the web-centre are controlled and patients should give prior consent. Since these are general issues when working with patient data, we will not focus on them here, but they remain important.

3.2.3. Usability and instruction

Ease of use is a critical requirement for any system that has to be used on a regular basis for a prolonged period of time. Since the interval between exacerbations is usually in the order of months, one should take care to reduce patient effort to a minimum, lest patients would stop entering data due to it being inconvenient. The patient population is relatively old on average –

possibly not very experienced with the kind of technology offered – hence to facilitate understanding the system provides the ability to do practice runs. Nurses will have a supportive role in training patients. Specifically, at the start of the monitoring process, the nurse will enter the relevant information in the web-centre and instruct the patient on the use of the smartphone and sensors. A practice run serves to check whether the process has been understood and also provides baseline measurement values. In addition to the verbal instructions, the patient receives a written step-by-step instruction for reference. Our initial tests showed that a written instruction may be insufficient and therefore visual instructions, for example with screen captures of the phone, should be provided as well. As nurses are often just as pressed for time as physicians and also cannot be expected to be experts on the technicalities of the system, we provided a support phone-number for both patients and the nurses to assist in resolving any problems they might have. When introducing technology, especially to older patients, instruction by example appears to be preferable to only written instructions.

3.2.4. Data collection rate

Depending on the health status of the individual patient the rate of data acquisition can be varied, which can be automated based on the acquired data and the model. As long as a patient has low risk of an exacerbation, monitoring can take place on a weekly basis, keeping the time investment at a minimum. If a patient is at risk according to model predictions, the system check-in can be scheduled daily to ensure the possible exacerbation is detected and acted upon appropriately. A week between registrations may however be too long as an infection can develop in a matter of days, which means that unscheduled registrations are also an important part of the monitoring process. To facilitate self-management, it is important that the patient is in control of the registration schedule and can initiate registrations in addition to the predefined schedule. In this way patients can check their probability of exacerbation (appropriately visualised) when they feel they might be at risk. Eventually, this would hopefully lead to patients who no longer need our supporting technology, instead being able to decide by themselves when to contact their physician. The system, thus, has a dual role, providing information to the patient and health-care workers and helping the patient assess their COPD related health status, for better control over their disease.

3.2.5. Sensor data reliability

One of the critical points in the current system is the lung function measurement, which provides objective data to augment the subjective symptom data. Whereas oximetry is a simple manoeuvre clipping the sensor on your finger and waiting for a few seconds, spirometry is harder to perform well. In a home-care setting there is a risk of obtaining little useful data. Although patients that would use our system are diagnosed with COPD and therefore have to perform regular spirometry tests, it remains of paramount importance to properly instruct the patient at the start of the intervention. To adhere to the spirometry standards the FEV₁ measurement is performed thrice and the best value is used. For the sake of usability we do not continue until three successful measurements have been made, as this would raise the bar unrealistically high. A reduction of the problem comes from the fact that we only take FEV₁ into account which is easier and far less strenuous for the patient than measuring forced vital capacity (FVC), as is customary in lung function testing, and requires complete exhalation. In general, the reliability of all data that is gathered should be considered. Sensors can provide objective data only when they are sufficiently reliable.

3.2.6. Feedback

A further consideration is the kind of advice that the system should provide based on the model prediction of exacerbation risk. Typical advice would be to contact the physician, but a useful next step would be to advise, for example, to increase the dosage of bronchodilator drugs, corticosteroids or to start with antibiotics. The exact implementation of different kinds of advice should be decided in cooperation with clinicians, in accordance with clinical guidelines. We return to this point in Section 7.

3.3. Hardware components of the disease management system

3.3.1. Smartphone

Currently our system runs on an HTC Desire smartphone as an application in the Android OS. In principle any Android phone with Bluetooth capability should suffice, which makes the platform fairly general. If available the fast 3.5G HSDPA connection was used, but due to the use of asynchronous communication and the relatively limited amounts of data being transmitted, data can also be transmitted to the central server when only a slower 2G GPRS connection can be realised. Recall that the self-management part of system also functions without any server connectivity at all.

3.3.2. Sensor interface

The phone communicates with the sensors via a Mobi, a Bluetooth-capable multichannel sensor-interface, from Twente Medical Systems International. In our case a Nonin pulse-oximeter and a custom micro-spirometer were connected to the Mobi. An important advantage of using the Mobi sensor interface is the availability of the communication-protocol specification, enabling us to integrate the sensor readings seamlessly into the Android application. Most of the other micro-spirometers on the market do not allow this, which makes them unsuitable for easy deployment in a home setting. The monitoring kit is shown in Fig. 2.

3.3.3. Pulse-oximeter

The pulse-oximeter used in this study was a Nonin Medical 8000AA, which is an industry standard pulse-oximeter. SpO₂ accuracy is 70–100% ± 2 digits.

3.3.4. Spirometer

We used custom-made pneumotachograph micro-spirometer prototypes by Twente Medical Systems International. These spirometers were newly developed to interact with the Mobi sensor interface and have the advantage of providing raw data such that analysing the spirometer readings is possible without requiring external software. This enables tight integration with our applica-



Fig. 2. The monitoring kit consisting of a smartphone, sensor interface and sensors.




Fig. 3. Schematic depiction of a spirometric flow-curve. Shaded area indicates FEV₁.

tion, which is difficult or impossible with most commercial spirometers on the market.

Measuring raw spirometric data requires that we do some signal processing. The raw data signal is processed on the smartphone to obtain FEV₁ values, i.e. forced expiratory volume in 1 s. In Fig. 3 a schematic spirometric flow-curve is shown. The shaded area is the FEV₁, which is computed from the raw data by first determining the peak of the curve, then finding the lowest point in the half second prior to the peak and integrating the flow during one second from there.

3.4. The Aerial application

As the project in which the intervention was developed was called *Aerial*, the smartphone application inherited the same name. Android applications are written in Dalvik, a programming language that uses Java syntax, but is compiled to run on the Android OS without the Java Virtual Machine. The application provides the following functionality: a timed alarm to signal the registration, with on-screen buttons to delay the start of the registration; a touch-screen interface for the questionnaire, which consists of eight yes/no questions (see Table 1); on-screen instructions for performing the measurements and Bluetooth communication with the Mobi sensor-interface to receive the measured data; processing of the spirometry data to compute the forced expiratory volume in 1 s (FEV₁); computation of the probability of an exacerbation based on the observed data (see also the section on the Bayesian network and model implementation below); asynchronously transmission of the observed data to the server over a secured data connection. In Fig. 4 screen captures are shown of the alarm screen and of the question about cough.

Table 1

Monitoring questionnaire. All on-screen text was in Dutch and has been translated to English for ease of exposition.

Since the last registration, did you have dyspnea worse than usual?
Since the last registration, did you have more sputum than usual?
Since the last registration, did you have sputum of a different colour than usual?
Since the last registration, did you cough more than usual?
Since the last registration, did you have more wheezing breath than usual?
Since the last registration, were you limited in performing day-to-day activities due to COPD?
Since the last registration, did you have a feeling of malaise?
Since the last registration, did you have a fever?

3.5. Risk model

The main component of our system is the probabilistic risk model. Based on the data that is gathered, the model can compute the probability of an exacerbation. The model we use is a Bayesian network [8,9]; Section 5 is devoted to the description of the development and evaluation of this model. The Bayesian network has been implemented using the lightweight reasoning engine EBayes³. Since EBayes is written in Java the model inference could easily be integrated in the Android application. Due to the relatively small size of the Bayesian network and the processing power of modern smartphones, the inference does not have to be deferred to a server but can be performed on site. This has the advantage that even when mobile phone network coverage is suboptimal the application can still provide a probability estimate and accompanying advice. Recall that this capability is also in line with our aim of patient empowerment as mentioned in the introduction.

3.6. Web-centre

The web-centre is the administration web-application that was built using the workflow management system iTasks [10]. The workflow system implements advanced features to generate and coordinate tasks and provides a generic (web)interface. Since the data management involved with patient monitoring is suitable to be represented as a workflow, iTasks provides a simple and effective way to construct the web-centre. In Fig. 5 an impression of the web-centre is shown. Shielded by a secure login and different levels of access rights, the web-centre provides a workflow to enrol patients in the system as follows: first a user account is made with a registration number as identifier; the regular registration schedule is established in consultation with the patient; and a monitoring kit is assigned to the patient; then a practice registration can be performed (registrations are scheduled on the server; the smartphone regularly polls for new messages from the server but can also be forced to synchronise manually), of which the measurement results can be saved as baseline values; finally a start date and time for the registration schedule can be chosen. In addition, the web-centre has basic administration functionality to manage users and monitoring kit assignment as needed in research projects. Importantly, all the data received from the phone can be made available for inspection by the relevant health-care workers, thereby offering a means to stay apprised of the patient's health status that is not available in usual care.

4. Technical feasibility and usability pilot study

Normally, there is a long way to go between the development of a prototype system and its actual use within a health-care setting. The involvement of patients in the development of a system is of crucial importance to improve a system based on their feedback and to ensure that the system meets the requirements.

We performed a system evaluation with a limited number of stable COPD patients recruited from the University Centre for Chronic Disease Dekkerswald in Groesbeek, the Netherlands, mainly to ensure that they were able to use the system effectively. We tested the monitoring equipment and data entry procedures with help of the patients and also studied the usability of the system from a patient's point of view.

³ F.G. Cozman, <http://www.cs.cmu.edu/javabayes/EBayes/index.html/>. Accessed: January 2013.

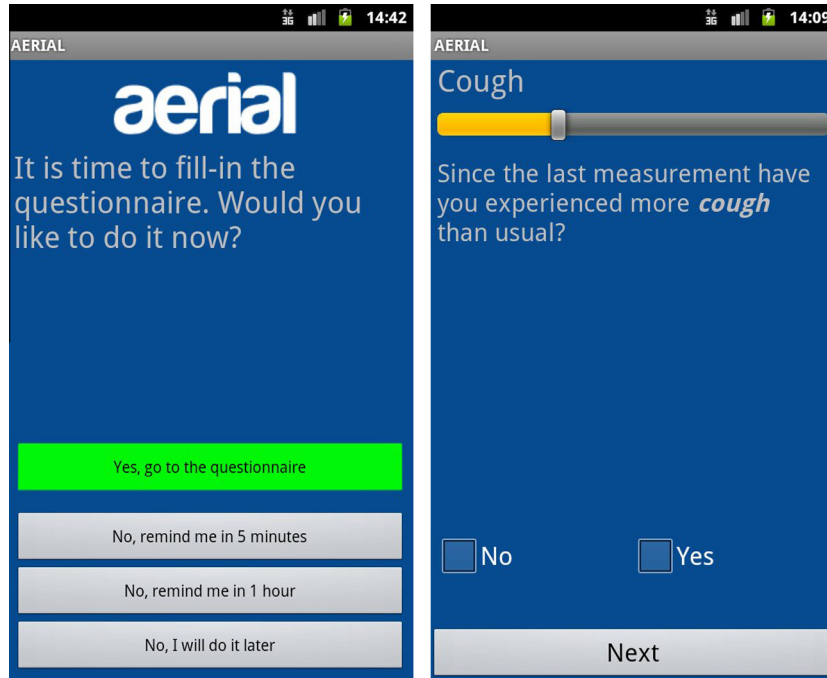


Fig. 4. Screen captures of the application start screen at registration time and the question about cough.

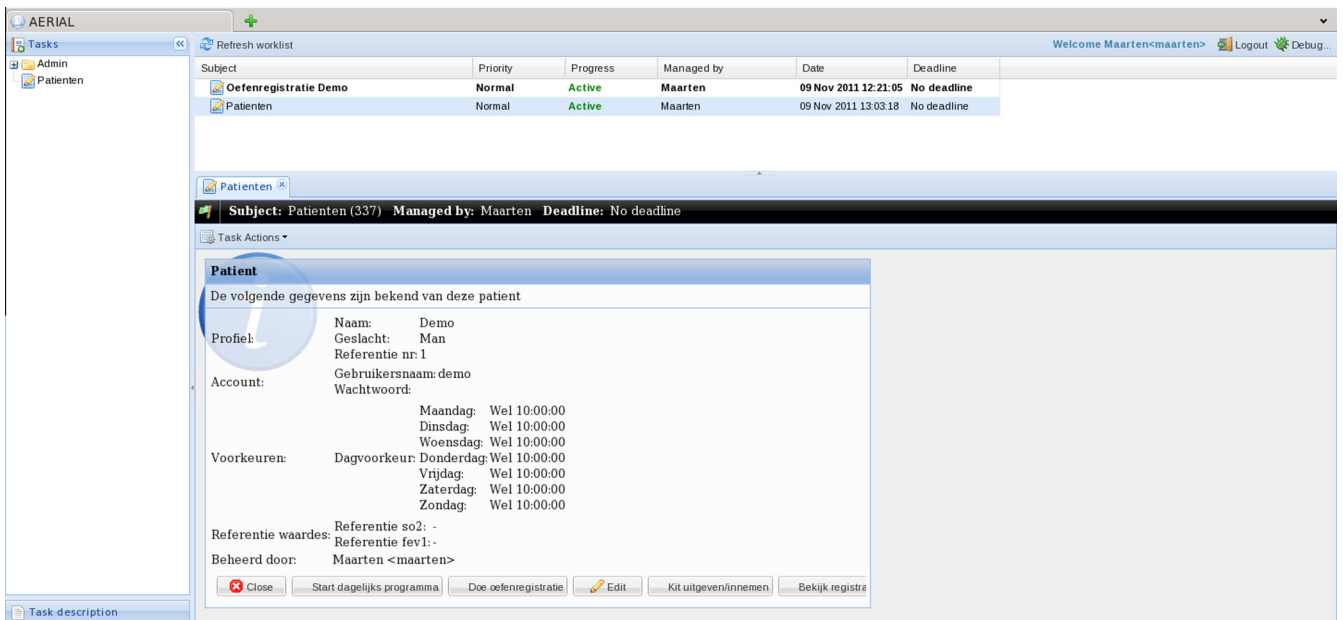


Fig. 5. A screen capture of the patient administration web-interface.

4.1. Methods

For this study five stable patients were recruited by a convenience sample from the lung rehabilitation program who gave informed consent to participate in the study. Using three exacerbation monitoring kits (as described in Section 3) patients were monitored for a duration of 9 days, starting in January 2011. Patients were contacted daily during this time to answer a set of questions and perform spirometry and oximetry measurement. The answers were entered as evidence into the Bayesian network model to determine the probability of occurrence of an exacerbation. These predictions have not been used for patient ad-

vice yet, as further model validation is required. Patients were asked to report malfunctioning which together with the received data and server-logs could be used to verify system performance. At the end of the monitoring period a semi-structured evaluation interview was held – both with the patient and with the health-care staff involved – to obtain qualitative feedback on the usability of the system. In the interview we established whether or not the patients understood the procedure and the questions and whether they found the phone-interface sufficiently usable. Also, we checked anomalous data we received (if any) and asked for suggestions for improvement. The evaluation results of the first two patients were used to improve the system before starting with the

second group. As this was a technical feasibility test, the clinical data obtained were only used to check for errors in the application or obvious model inaccuracies.

4.2. Results

The pilot study allowed us to identify some problems in the server-side software, for example with respect to adequately recovering from connection errors, which could be amended relatively easily. After these amendments, the system functioned adequately at the technical level. With respect to usability, some problems became apparent at an early stage, such as the response buttons being placed too close together for comfortable use and an unclearly worded question resulting in confusion. In the evaluation interviews that were carried out with the patients the consensus was that the system could be useful to gain insight in the disease, was easy to use and not found to be intrusive. In particular patients indicated that they would be willing to use such a system in a home-care setting, which will need to be verified more rigorously in the next research stage. Thus, the patients' impression of the system after using it was quite positive and made us confident that the current system offers a suitable basis for moving toward a commercial product.

5. A probabilistic model for the detection of COPD exacerbation

The conclusion reached above, that the system was suitably functioning from a technical point of view, leaves the question unanswered whether it is also sufficiently accurate for predicting exacerbations, the ultimate aim of the development of the system. This question is addressed below.

5.1. Bayesian networks

For the data interpretation to detect exacerbations we use a *Bayesian network* [8], i.e. a probabilistic graphical model represented as a pair $BN = (G, P)$. Here, $G = (V, A)$ is a directed acyclic graph consisting of *vertices* V , corresponding one-to-one to random variables of interest, and *arcs* $A \subseteq V \times V$, representing dependencies between variables. Furthermore, P is the quantitative part, denoting a joint probability distribution, specified in terms of a family of conditional probability distributions of the form $P(V_i | pa(V_i))$, that is the probability that V takes on a specific value given the values of its parent variables. The network represents the joint distribution over the random variables, which can be factored according to the independences represented in the graph, resulting in:

$$P(V_1, V_2, \dots, V_n) = \prod_{i=1}^n P(V_i | pa(V_i)),$$

where $V_i \in V$ is the representation of a random variable in the graph G . Probabilities of interest can be computed from this joint probability. In this case we are mainly interested in the probability of an exacerbation given the evidence obtained from monitoring.

Although other probabilistic models, such as based on logistic regression or artificial neural networks, could have been used for detecting exacerbations in patients, Bayesian networks appeared to be the most suitable method for our research. The main reason is their inherent flexibility, as it is possible to develop Bayesian networks from expert knowledge only, whereas learning of parameters and structure from data is also an option [11]. This is a necessary feature because of the ironic situation that although COPD is a very common disorder, there is currently not much patient data available from which probabilistic models can be learned. This is caused by the early stage in which much of the eHealth research still is; few researchers have tried to monitor sufficient number of

parameters of a sufficient number of patients at home in a sufficiently frequent way, e.g. more than once every week.

Bayesian networks also support reasoning about individual patients: although the model describes general relations between the variables of interest, all predictions are personalised by entering patient-specific data. The model is thus capable of making predictions for individual situations, and can provide 'what-if-predictions by entering virtual evidence. Furthermore Bayesian networks are able to provide probability estimates given partial patient data.

5.2. A Bayesian network for COPD

For our COPD models, the main outcome variable is *exacerbation*, focussing on a symptom based definition. But the nature of a Bayesian network allows us to easily inspect probabilities for any variable. The network contains two hidden variables, namely *infection* and *lung function* which cannot be observed directly, but whose values can be derived based on indirect measures, such as body temperature for infection and the forced expiratory volume in 1 s (FEV₁) for lung function. Other important variables are the symptoms that one might expect a patient to report, such as dyspnea (breathlessness), sputum volume and purulence, cough, wheeze, general malaise, fever, and whether performing daily activities is difficult due to COPD. Additionally, the clinical signs FEV₁ and blood oxygen saturation (SpO₂) are included. Other signs and lab results are not included as these are generally not available in a home-care setting. All variables except for FEV₁ are binary and either encode presence or absence, or normal or abnormal states, which should be interpreted in comparison to the individual patient's baseline values. As FEV₁ is an important, objective measure it has five states: 'Normal'; -10%; -20%; -30%; <-30%; this division is based on expert opinion but would perhaps be represented more naturally as a continuous variable (as volume is clearly continuous), which requires appropriate continuous data.

5.2.1. Expert knowledge model

The goal of using expert opinion to construct the model was to obtain an understandable clinically justifiable model, in the absence of sufficient data that would allow learning the structure from the data. The contrast with a purely data driven approach lies in the fact that we can ask the expert for causal influences, instead of correlations obtained from data via statistical methods. We can therefore argue for the validity of the structure of the Bayesian network in causal terms, of course, without supposing that a Bayesian network is a causal network. We were interested primarily in the probability of an exacerbation, which is a clinically relevant event that is mediated by a number of physiological processes and influenced by environmental factors. We cannot hope to take into account every relevant factor, but it is clear that the symptoms of an exacerbation are a consequence of reduced lung function. Unobserved processes that constitute either normal or abnormal lung functioning can be represented by hidden variables. The measurements SpO₂ and FEV₁ are objective, but not sufficient to fully characterise lung function. The symptoms *dyspnea* and *wheeze* are direct consequences of reduced lung function. Whether patients are capable of performing daily activities also depends on whether the lungs are capable of providing sufficient oxygen. Next in the causal view are possible causes of reduced lung function. The most prevalent cause of an exacerbation is a lower airway infection, which is in itself a combination of unobservable processes, modelled by the variable *infection*. The presence of an infection is then the cause of further symptoms like *fever*, general *malaise* and *sputum* production, the latter in turn causing cough to clear the sputum accumulation in the airways.

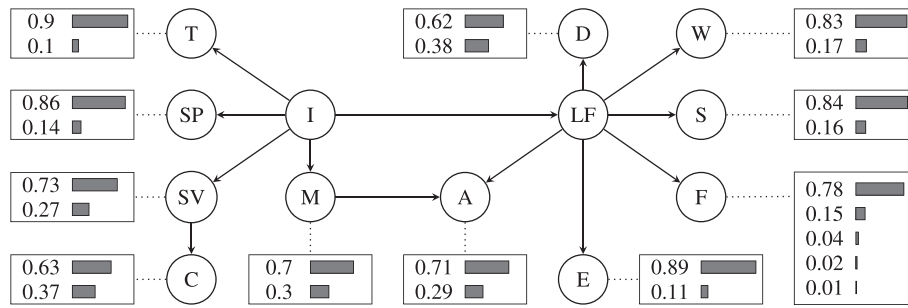


Fig. 6. Expert opinion based Bayesian network with prior, marginal probabilities shown (top probability is Normal state). A = activity, C = cough, D = dyspnea, E = exacerbation, F = FEV₁, I = infection, LF = lung function, M = malaise, S = SpO₂, SP = sputum purulence, SV = sputum volume, T = temperature and W = Wheeze.

The resulting expert-opinion based COPD-exacerbation prediction model is depicted in Fig. 6. The Bayesian network has been constructed in close cooperation with two pulmonologists (lung specialists) of the Radboud University Nijmegen Medical Centre. The relevant variables that are related to the occurrence of exacerbations were identified by the domain experts and supported by recent clinical literature (e.g. [6]). As a second step the dependence relations between the variables were elicited from the domain experts, resulting in the qualitative part of the model. Subsequently, probabilities were estimated. As this is a fairly difficult task for clinicians not used to the abstract representation in terms of conditional probabilities, we started off by asking them to provide qualitative constraints of the type ‘Is $P(X = \text{true} | \text{pa}(X) = x) <, = \text{ or } > P(X = \text{false} | \text{pa}(X) = x)$ ’ (expressed in words). Based on these constraints an initial network was constructed to give some insight in how, approximately, the resulting Bayesian networks would look. We then went over different test cases to adjust the probabilities until the expert thought the resulting predictions were acceptable. The entire process has – in a different context – been previously described in [12].

5.2.2. Parameters from data

We have used two data sets to learn parameters and to validate the model. We started with a data set from an earlier project on detecting exacerbations at the Radboud University Medical Centre [13]. We will refer to this data set with A. This set consists of questionnaire answers of 86 Dutch COPD patients, 54 of whom had an exacerbation during the study. The data has been acquired bi-weekly during 2006–2007 via automatic telephone interviews, resulting in time series data with a total of 1922 data entries, of which 162 entries provided exacerbation data according to a symptom based definition of exacerbation. Later we will study the influence on the results by different definitions of an exacerbation. Data was only available for a subset of the variables in the model, specifically: exacerbation, dyspnea, sputum volume and purulence, wheeze, cough, and activity.

As a comparison we also validated the predictions on an unrelated data set [14], that provides time series of COPD-exacerbation related variables. We refer to this data set with B. This data set is used solely for external validation, so no model parameters were learned from this set. As the data is used for testing our model retrospectively, the set of variables is incomplete. However, it should give some indication of the generalisability of the results. The data consists of time series from 13 patients of the London COPD cohort who had an exacerbation, with a total of 2849 data entries, of which 406 were during an exacerbation. In contrast to the Dutch data set, these data were collected on a daily basis. The data contains values for the variables dyspnea, sputum volume and purulence, wheeze, cough, temperature, and SpO₂.

Parameter estimation proceeded as follows. Since not all data sets include all variables of interest, parameters were learned for those variables for which data was available. The usual approach

to learning parameters from data entails computing a maximum likelihood estimate of the parameters, let us call them θ , for the given data D :

$$\arg \max_{\theta} P(D|\theta),$$

or from a Bayesian perspective, the maximum a posteriori (MAP) estimate:

$$\arg \max_{\theta} P(\theta|D) = \arg \max_{\theta} P(D|\theta)P(\theta)/P(D),$$

by Bayes’ rule. Note that the denominator $P(D)$ is constant when maximising the expression for the parameter θ and therefore irrelevant. For complete data this can be computed easily when assuming that the data is independent and identically distributed (i.i.d. – assumption), but when data is incomplete – as it almost always will be in a clinical context – this should somehow be taken into account. We used the expectation–maximisation algorithm, popularised by Dempster, Laird and Rubin [15], which iteratively adjusts the parameters by alternating between estimating values for the missing data given the current parameters and using the completed dataset to compute new maximum likelihood parameters. Although this does not guarantee that the optimal parameters are found, Dempster et al. proved that the procedure improves the likelihood at each iteration, resulting in a local maximum of the likelihood function.

5.2.3. Comparison between expert and data models

We developed exacerbation predictions models both on the basis of expert opinion and from data. One would hope and expect that these models at least have some similarity. Although the parameters will usually not be the same we can make a qualitative comparison using the constraints we elicited from the expert. As described above we asked the expert to order the values of the random variables based on which would be more likely given the values of the parent variables. We would expect that the parameters learned from data would follow the same patterns.

For most of the variables for which we had data, the ordering constraints were met, except for the symptom variables *sputum purulence* and *activity* and the outcome variable *exacerbation*, which warranted more detailed inspection. For *sputum purulence* (SP) the qualitative constraints were that the conditional probabilities given *infection* were ordered as follows: $P(\text{SP} = \text{normal} | \text{I} = \text{false}) > P(\text{SP} = \text{abnormal} | \text{I} = \text{false})$ and $P(\text{SP} = \text{abnormal} | \text{I} = \text{true}) > P(\text{SP} = \text{normal} | \text{I} = \text{true})$. Analogous constraints were established for *activity* given *lung function*. When we look at the parameters learned from the data we can see that in both these cases the direction of the effect is correct and the probabilities are close to meeting the constraints. For instance the parameters for *sputum purulence* are $P(\text{SP} = \text{abnormal} | \text{I} = \text{false}) = 0.02$ and $P(\text{SP} = \text{abnormal} | \text{I} = \text{true}) = 0.43$, indicating that the probability increases considerably with the presence of an infection, but not quite enough to reach the limit of 50% which would reverse the

ordering. The threshold of 50% follows from the qualitative constraints, because if for a binary variable X the probability $P(X = 0) < P(X = 1)$ then by definition $P(X = 1) > 0.5$. Something similar holds for *activity*. For the *exacerbation* variable the situation is a bit different in the sense that the distance to the threshold of 50% is larger, but the difference between the probability of an exacerbation given normal lung function or abnormal lung function is large. It appears that the qualitative constraint was too strict and there are more situations with decreased lung function that are not an exacerbation. However as the evaluation that follows will show, the difference between the probabilities for normal and abnormal lung function is sufficiently large to result in enough discriminative power.

5.3. Model evaluation

To gain some insight in model accuracy we can study the behaviour personalised with patient data. In the network of Fig. 6 prior probabilities were shown, but by entering observations we can personalise the probability estimates to be consistent with a particular patient status. Table 2 contains the probabilities of some key variables, given different evidence situations. The causal nature of the model ensures that the probability of an exacerbation increases when more symptoms are reported or signs measured.

ROC-analysis To understand how the model behaves in general we performed an ROC-analysis, using 10-fold cross-validation with data set A , on the sub-model for which data was available to compare the expert opinion probabilities with learned parameters. That is, the variables in the set {exacerbation, dyspnea, sputum volume, sputum purulence, cough, wheeze, activity}. Prediction are obtained for each data point by inserting evidence into the network and inferring the resulting probability of an exacerbation. Concretely this first means that the observed variables are clamped to the observed value (and thus have probability 1). Second, the probability of an exacerbation is computed given the observations (using standard Bayesian network inference algorithms, see e.g. [8]). In Fig. 8 the resulting ROC-curves are shown for both the data and expert model, conveying that the predictions can indeed distinguish the exacerbation cases. For the data model we find a mean area under the curve (AUC) of 0.93; and for expert model a mean AUC of 0.97.

Since the construction of our model we have obtained the data set (B) from Hurst et al. [14] described above. It is worthwhile to compare the results on this data set with the data set we used originally. To do so we looked at the subset of variables in B that intersect with the original data, that is {exacerbation, dyspnea, sputum volume, sputum purulence, cough, wheeze}. We then computed the model predictions for *exacerbation* given the observations on these variables and again performed an ROC-analysis. The area under the curve for these predictions was AUC = 0.87. As is common in ROC-analysis, we computed the optimal cut-off probability p^* by means of the point on the ROC curve closest to the (0,1) coordinates:

Table 2

Model behaviour for different evidence situations. First column indicates which variables are observed to be abnormal. Other columns show probability of being abnormal (e.g. symptom is present).

	Dyspnea	Sputum purulence	Cough	SpO ₂	Exacerbation
No evidence	0.38	0.14	0.37	0.16	0.11
Cough	0.43	0.20	1	0.20	0.17
Cough, Wheeze	0.68	0.45	1	0.41	0.46
Dyspnea, SpO ₂	1	0.51	0.53	1	0.55
Dyspnea, Sputum volume, FEV ₁ (−20%)	1	0.67	0.80	0.60	0.71

Table 3

Exacerbation detection performance for different data sets and definitions of exacerbation (see text for details).

Data set	AUC	TPR	FPR	ACC
A	0.93	1	0.16	0.85
A_{Drug}	0.81	0.79	0.3	0.72
$A_{Contact}$	0.85	0.87	0.3	0.72
B	0.87	0.88	0.20	0.81
$B p_{mean}^A$	–	1	1	0.14
$B p_{max}^A$	–	0.88	0.20	0.81

$$p^* = \arg \min_p \left((1 - \text{tpr}_p)^2 + \text{fpr}_p^2 \right),$$

where ‘tpr_p’ and ‘fpr_p’ are the true (false) positive rate that depend of the chosen threshold value p . The optimal cut-off value gives a general, problem-independent balance between true positive and false positive rates. Table 3 shows a comparison of the prediction results of the cross-validation model on the original and the new data in terms of AUC and the true (false) positive rate and accuracy at cut-off p^* (computed separately for each data set). As a further interesting evaluation we took the model learned from A and decided upon a cut-off probability which we then applied to the data from B . Hence we used the new data as if we did not have the correct labels, in the sense that we did not compute the optimal cut-off point. Since the A model was learned via cross-validation we first computed p^* for each of the cross-validation runs and compared two options, either taking the mean cut-off of the cross-validation runs p_{mean}^A ; or the max cut-off probability p_{max}^A (row 5 and 6 in Table 3). Note that using the mean is too aggressive a cut-off as it results in all data being classified as an exacerbation. This can be explained by the class imbalance in the data, as misclassifying only a few data points has a relatively large influence on the performance measures. The optimal cut-off point is in some sense overfitted to the training data. Using the max cut-off however, we see that the performance is the same as for the optimal value for data set B . This sensitivity to the precise cut-off point will likely decrease when we obtain more data (especially positive examples).

Definitions of exacerbation When using data to learn parameters it is important to note that clinically there are different definitions for when a worsening of COPD-related health status should be considered an exacerbation. The article by Bischoff et al. [13] that describes the study data we use here, also acknowledges this point by distinguishing exacerbations defined in terms of symptoms; drug prescriptions; and health care contact. In Fig. 7 the counts are shown for data points according to the following definitions:

Symptom-definition	Presence for at least 2 days of at least the major symptoms dyspnea, sputum volume increase and sputum purulence.
Drug-definition	Start of a course of corticosteroids or antibiotics.
Contact-definition	Unscheduled contact with health care workers due to worsening of COPD.

Interestingly enough there are many data points for which these definitions do not agree, which will of course influence model accuracy. Hence it is worthwhile to study the effect of using other definitions on predictions. Row 2 and 3 of Table 3 show the performance statistics of using the *drug* or *contact* definition with the model learned with the *symptom* definition. As expected this results in a considerable drop in performance, as the definitions capture different aspects of exacerbations.

More interesting is the performance of the expert opinion model as this is based on the experience and clinical gaze of the pulmonologist instead of a technical definition. The result, using the same

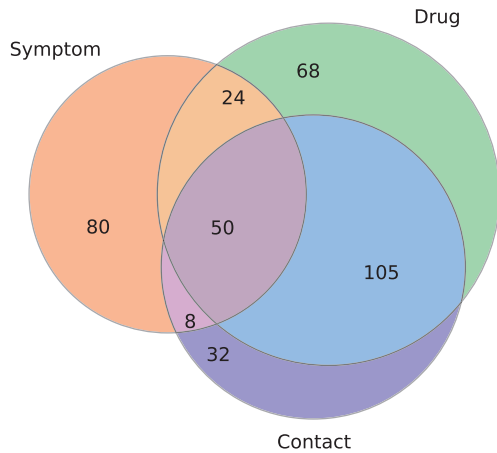


Fig. 7. Venn diagram of data-point counts for different definitions of exacerbation in data set A.

indicators as before, is shown in Table 4. Interestingly enough we observe the same pattern here as with the learned parameters, in the sense that the symptom based definition performs considerably better than the event based definitions. A possible explanation is that drug use and health care contact data are based on self reporting which may have introduced noise (both over-reporting and under-reporting could have occurred due to for example misremembering). Since the symptom-based definition is derived from more variables it is less likely that someone would mistakenly remember multiple symptoms. Another explanation might be that by explicitly constructing a model of symptoms we primed our expert to focus on the symptomatic part of exacerbations. Perhaps this balance would have been different had we included medication in the model. This is a topic for further investigation.

These results indicate that we can reliably detect exacerbations as they are happening, which is a practically useful result. The final trade-off between true and false positives should be based on more data, taking into account what definition of exacerbation is most useful for disease management. Compared to usual care it is an

Table 4 Performance of the expert model for different definitions of exacerbation.

Definition	AUC	TPR	FPR	ACC
Symptom	0.97	1	0.10	0.91
Drug	0.82	0.75	0.24	0.75
Contact	0.85	0.83	0.25	0.76

improvement even if we catch only some of the exacerbations that are now detected too late. Ultimately, the more difficult task of predicting exacerbations still lies ahead, as currently we lack sufficient data to model the temporal progression. In further testing of the system we also plan to gather the necessary data to construct a temporal model, which requires trend analysis of time courses of symptoms and signs leading up to an exacerbation.

6. Related work

6.1. Telemonitoring requirements

In a recent study by Peirce et al. [3] an extensive analysis of current telemonitoring practice was performed. Based on this analysis a set of requirement questions was identified that should be answered to obtain an effective telemonitoring system. Although we think that our system is more general than purely monitoring, it still appears useful to consider these questions and try to answer them for our system. *Why monitor?* Early detection of deterioration leads to improved outcomes for patients. *Who to monitor?* In general deciding on the target population is important for the success of an intervention. For our COPD system we target those patients that suffer from frequent exacerbations (more than 2 per year), which is a clearly defined group that stands to gain from timely intervention on exacerbations. *How to monitor?* Our system uses a combination of objective and subjective data, while trying to maintain a minimal effort requirement for patients. *What to monitor?* The variables in our model have been identified in cooperation with clinicians and the value of each variable is determined statistically from data to ensure we take into account the most relevant parameters. That is, we started out from clinical requirements, not from technological ones. However to obtain a usable

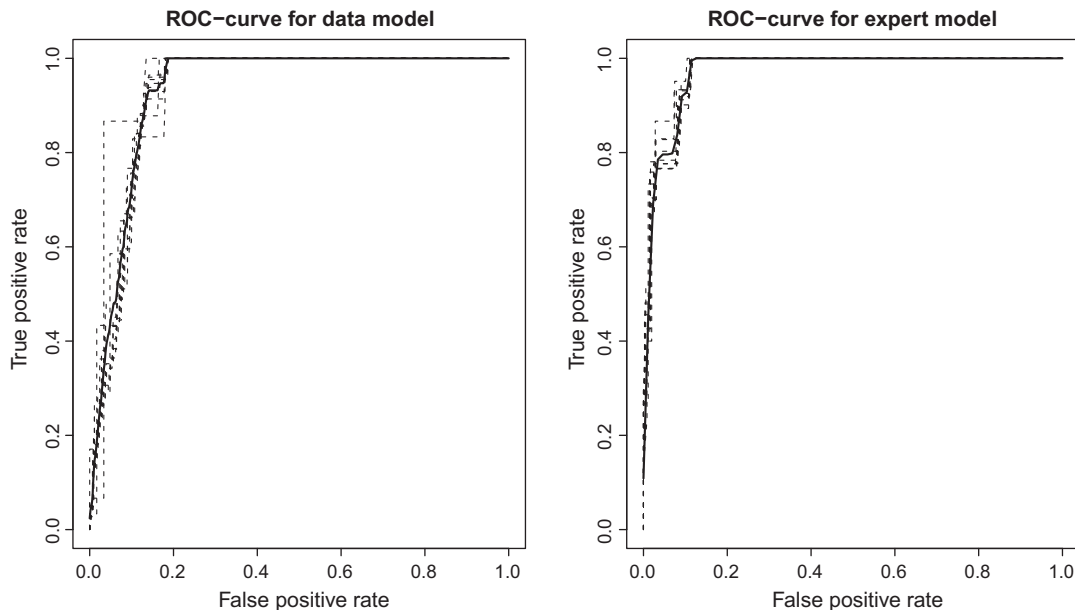


Fig. 8. ROC-curves, cross-validation results (dashed lines) and average (bold).

system, in practice some concessions had to be made to what is possible technically and cost-wise. *How to detect deterioration early?* This is an important point that forms the basis of our current research. Now that we have a system that can detect exacerbations, the next step is modelling temporal progression in order to observe trends that allow us to predict an approaching exacerbation. Although Peirce et al. labelled early detection as a shortcoming of many systems, it is one of the main reasons for the development of our system. *How to display information appropriately?* We have tried to keep the user interface as simple as possible to accommodate our users who may not be proficient with modern smartphones. The presentation of the feedback based on model predictions is still a topic of further research. However as was recognised by Peirce et al. a combination of a simple feedback mechanism (say a smiley face, traffic light or some kind of simple colour coding scheme) with a more in depth analysis of the data (e.g. trend graphs for symptoms and measurements) will likely turn out useful. Patients can then choose the level of detail with which they are comfortable. *Who to notify? How to respond?* Alert notifications are an important part of the practical embedding of a telemonitoring system in day-to-day care. We think that both the patient and their health care providers should stay part of the decision process, even if data interpretation is automated as it is here. Therefore, patients keep the initiative in contacting their physician, but the system can provide advice on when to do so. *How to store data?* It depends on the local health care structure whether integration with GP or hospital information systems is desirable. From the perspective of patient-empowerment/self-management, it is useful that our system can also be employed stand alone, leaving access rights with the patient. Summarising, our system is already well-equipped with respect to these telemonitoring requirements, yet refinements can also be identified by further attention to answering these questions.

6.2. Related telemonitoring systems

Monitoring COPD-patients has been a topic of interest for some time now, but the level of automation is usually limited. Basically we can distinguish two types of earlier work, telemonitoring where technology only provides a means of communication between patient and physician; and systems that have some kind of automated alerting or data interpretation. Systems of the first kind have for example been reported by [16], where pulse-oximetry data was sent to the hospital via a telephone line and analysed by hand; [17], where a nurse set up a teleconference with a physician; and [18], where patients with respiratory failure were monitored by telephone. The system reported in [19] is more similar to ours in the sense that both questionnaire answers and measurements are acquired and alerts are sent when readings fell outside the previously established normal range. Also in [20] a monitoring system is reported that checks symptoms by means of a questionnaire and automatically checks for deviant answers. These studies all focus on general COPD care, whereas we focus on early detection of exacerbations. Also, by employing a Bayesian network we are able to perform more automatic data interpretation. Recently, an exacerbation risk alert-call system has been tested, based on weather forecasts, which was reported to be effective in predicting risk periods [21]. Another recent study on predicting exacerbations found that breathing frequency is a good predictor [22]. Currently their approach is limited because breathing frequency was measured via a home oxygen-therapy system. If a convenient way to monitor breathing frequency could be found, this would be a possible extension of our current system. For further references, see also the review of home care for COPD by Polisena et al. [23] and the telehealthcare review by McLean et al. [2].

7. Discussion

We aimed to develop a clinically useful system to assist in patient management and in this paper have described the progress towards that goal. In this section we discuss limitations of the current implementation, lessons learned from the development, and we look at future work, extensions and possible generalisations.

The current version of our system has some limitations in the implementation that are a consequence of an iterative design process. We envision an autonomous system capable of disease management, but it is necessary to first construct and test the various components. Currently our probabilistic model functions as a classifier labelling the monitoring data as an exacerbation with some probability. We have shown that we can do this reliably, also on new data. To further assist in disease management we would like to infer the probability of health states ahead of time. In general chronic disease management can benefit from focussing on the analysis of trends in time series. We are currently working on a temporal Bayesian network that extends the current model to take time into account.

Currently, we provide the patient with information about the likelihood of an exacerbation as a minimal kind of support. A logical next step is to analyse in more detail what kind of advice would be beneficial. In order to do so we will also take into account more background information about the patient, most notably current therapy plans. Patient feedback can then be extended to include adjusting therapy. So while the system we have now is mainly an alerting system – which is already an improvement compared to usual care – it will be extended to more specific advice. Clinical guidelines and input from clinicians is indispensable in formulating the right advice. Also more in general, for a decision support system to become successful it is important to have the cooperation of both patients and clinicians. As users of the system they will be able to comment on usability and suggest improvements.

These extensions are part of the route to a system that is capable of both autonomous functioning and more traditional alerting. Besides the improvements mentioned already, future work will include a larger evaluation on prospective data of model performance leading up to a randomised trial to evaluate the efficacy of the intervention. The current work however provided us with valuable insights about usability, by asking patients for feedback; allowed us to test the technical aspects of monitoring; and ensured that detecting exacerbations was possible. These findings form a good basis for the extended evaluation that is now being planned.

8. Conclusion

We have described the results of the development of a novel autonomous mobile assistance system for disease management, in particular focussed on COPD. The system described in this paper uses probabilistic reasoning to automatically interpret patient specific data to assist patients in managing COPD exacerbations. Initial testing shows that applying the system is technically feasible and patients are capable and willing to use the system. The model is well-founded on expert knowledge, literature and data, providing effective exacerbation detection. We have thus produced and evaluated on a pilot scale an advanced system architecture for home management of COPD exacerbations, with promising results. Future work will involve a more extensive test in a home-care setting, finally leading to a system capable of exacerbation detection in an early stage such that COPD exacerbation impact can be reduced and patients can self-manage their disease.

Acknowledgments

This work is supported by the Netherlands Organisation for Health Research and Development, ZonMw Project 300050003.

References

- [1] Oh H, Rizo C, Enkin M, Jadad A. What is eHealth? A systematic review of published definitions. *J Med Internet Res* 2005;7(1):e1.
- [2] McLean S, Nurmatov U, Liu J, Pagliari C, Car J, Sheikh A. Telehealthcare for chronic obstructive pulmonary disease. *Cochrane Database Syst Rev* 2011;7.
- [3] Peirce S, Hardisty A, Preece A, Elwyn G. Designing and implementing telemonitoring for early detection of deterioration in chronic disease: defining the requirements. *Health Inform J* 2011;17(3):173–90.
- [4] Viegi G, Pistelli F, Sherrill D, Maio S, Baldacci S, Carrozzi L. Definition, epidemiology and natural history of COPD. *Eur Respir J* 2007;30(5):993–1013.
- [5] McKinstry B, Pinnock H, Sheikh A. Telemedicine for management of patients with COPD? *The Lancet* 2009;374(9691):672–3.
- [6] Wedzicha J, Seemungal T. COPD exacerbations: defining their cause and prevention. *The Lancet* 2007;370(9589):786–96.
- [7] Wilkinson T, Donaldson G, Hurst J, Seemungal T, Wedzicha J. Early therapy improves outcomes of exacerbations of chronic obstructive pulmonary disease. *Am J Respir Crit Care Med* 2004;169(12):1298–303.
- [8] Pearl J. Probabilistic reasoning in intelligent systems: networks of plausible inference. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.; 1988.
- [9] Koller D, Friedman N. Probabilistic graphical models: principles and techniques. The MIT Press; 2009.
- [10] Lijnse B, Plasmeijer R. iTasks 2: iTasks for end-users. In: Proceedings of the 21st symposium on implementation and application of functional languages; 2009.
- [11] Lucas P, van der Gaag L, Abu-Hanna A. Bayesian networks in biomedicine and health-care. *Artif Intell Med* 2004;30:201–14.
- [12] Lucas P, Boot H, Taal B. Computer-based decision-support in the management of primary gastric non-hodgkin lymphoma. *Methods Inform Med* 1998;37:206–19.
- [13] Bischoff E, Boer L, Molema J, Akkermans R, van Weel C, Vercoulen J, Schermer T. Validity of an automated telephonic system to assess COPD exacerbation rates. *Eur Respir J* 2012;39(5):1090–6.
- [14] Hurst J, Donaldson G, Quint J, Goldring J, Patel A, Wedzicha J. Domiciliary pulse-oximetry at exacerbation of chronic obstructive pulmonary disease: prospective pilot study. *BMC Pulm Med* 2010;10(1):52.
- [15] Dempster A, Laird N, Rubin D. Maximum likelihood from incomplete data via the EM algorithm. *J Roy Statist Soc Ser B (Methodol)* 1977;39(1):1–38.
- [16] Maiolo C, Mohamed E, Fiorani C, Lorenzo AD. Home telemonitoring for patients with severe respiratory illness: the Italian experience. *J Telemed Telecare* 2003;9(2):67–71.
- [17] Vontetsianos Th, Giovanis P, Katsaras Th, Rifopoulou A, Mpirmpa G, Giaboudakis P, et al. Telemedicine-assisted home support for patients with advanced chronic obstructive pulmonary disease: preliminary results after nine-month follow-up. *J Telemed Telecare* 2005;11(S1):86–8.
- [18] Vitacca M, Assoni G, Pizzocaro P, Guerra A, Marchina L, Scalvini S, Glisenti F, Spanevello A, Bianchi L, Barbano L, Giordano A, Balbi B. A pilot study of nurse-led, home monitoring for patients with chronic respiratory failure and with mechanical ventilation assistance. *J Telemed Telecare* 2006;12(7):337–42.
- [19] Paré G, Scotte C, St Jules D, Gauthier R. Cost-minimization analysis of a telehomecare program for patients with chronic obstructive pulmonary disease. *Telemed J e-Health* 2006;12(2):114–21.
- [20] Trappenburg J, Niesink A, de Weert-van Oene G, van der Zeijden H, van Snippenburg R, Peters A, Lammers J, Schrijvers A. Effects of telemonitoring in patients with chronic obstructive pulmonary disease. *Telemed J e-Health* 2008;14(2):138–46.
- [21] Halpin D, Laing-Morton T, Spedding S, Levy M, Coyle P, Lewis J, Newbold P, Marno P. A randomised controlled trial of the effect of automated interactive calling combined with a health risk forecast on frequency and severity of exacerbations of COPD assessed clinically and using EXACT PRO. *Prim Care Respir J* 2011;20(3):324–31.
- [22] Yañez A, Guerrero D, Pérez de Alejo R, Garcia-Rio F, Alvarez-Sala J, Calle-Rubio M, et al. Monitoring breathing rate at home allows early identification of COPD exacerbations. *CHEST J* 2012 [Epub].
- [23] Polisen J, Tran K, Cimon K, Hutton B, McGill S, Palmer K, Scott R. Home telehealth for chronic obstructive pulmonary disease: a systematic review and meta-analysis. *J Telemed Telecare* 2010;16(3):120–7.