

A Robotic Interface to a Medication Review Expert System

Claire D'Este, Daniel Reid and Byeong Ho Kang
School of Computing and Information Systems
University of Tasmania
Hobart, TAS 7001, Australia
{cedeste,bhkang}@utas.edu.au

Abstract

Medication problems are responsible for a significant percentage of health problems in the aged. A robot in the patient's home may assist in monitoring their condition and medication use. This application allows us to address issues with integrating the sensory and interactive capabilities of robots with the high-level problem solving of expert systems. The robot needs to provide and receive information from the patient with a simple interface, take sensor readings and use this data to make inferences using the expert system. This paper describes a system in development and gives a step-by-step example of an interaction between patient, robot and medication expert system.

1. Introduction

With an aging population and shortage of nursing staff, there will be increasing demand for technologies that assist the elderly and the disabled. A robot that lives with the patient in their home could monitor their condition and provide a variety of support services. The long-term aims for assistive robots include performing manual and cognitive tasks for the patient. These might range from bathing to scheduling medications. A major factor in this becoming a reality is the acceptance by the patients themselves. The interface between human and robot needs to be simple, reliable and unobtrusive. The robot also needs to interface between the real world and its high-level problem solving.

Rather than use a robot, it is possible to fit the patients home with inexpensive web cameras. However, if we want to perform sensor checks on the patients body, or use expensive sensors like thermal cameras, it would be more convenient and less expensive to have them onboard one device. A camera-laden home may also affect the patients sense of privacy, whereas a robot could be given commands from behind a closed door.

The robot should try and extract required information using means that are low impact on the patient. These patients are unlikely to have experience in programming robots and, due to the nature of the application, be unlikely to desire to learn. Elderly patients may also be uncomfortable with computer interfaces. We can explore these issues using simpler robots with reduced functionality. An important issue that we have chosen to focus on for this project is Medication Review (MR). This paper describes a system in development and gives a step-by-step example of an interaction between patient, robot and successfully evaluated Medication Review expert system [1].

1.1 Medication Review

Elderly patients without daily medical care often have issues with their medication. Drug usage problems result in 12% of all hospital admissions and costs \$400 million annually in Australia [7]. One initiative is the implementation of Medication Review systems. MRs are designed to assess potential problems by tracking the patients medical history, pathology results and their past and present drug regimen. This information is used to determine such issues as:

- medical conditions that are not being treated
- drugs they are on that may be affecting a condition
- combinations of drugs with similar ingredients that result in too high dosages
- drugs that are not required or contraindicated in their condition

Pharmacists in Australia are financially encouraged to implement MR systems. However, few take up this offer because of fear of error and a lack of confidence [10]. The research described in this paper makes use of a medication review knowledge base already developed using MCRDR that has demonstrated more consistent conclusions than a human expert [1].

1.2 Related Work

At Carnegie Mellon, they have been working for several years on a robot that provides reminders for elderly people regarding medication, meals and toilet breaks [8]. The robot not only alerts the patient, but can observe if the patient complies. The robot can also decide to not give a reminder if the patient is observed performing the task before the reminder is given. The focus of the research is on the action planning in the temporal domain. They use their own reasoning formalism called Quantitative Temporal Bayes Net.

Sincere Kourien Nursing Home in Japan has installed teddy bear robots that observe the elderly residents [6]. The robots have a speech interface. They time the responses to questions and the residents performing various tasks. Nursing staff are alerted if the robot determines there are any unexpected changes.

2 A Medication Review Robot

We could use a body of expert knowledge to create a specific knowledge base for the patient, but to provide the highest level of care we need to consider the patients changing condition. A robot running a medication review system could continuously consult the MR using sensory data acquired from the patient. For example, if the robot detects the patient has low blood sugar, the MR may recommend the patient eat something. The robot could also download pathology results, draw recommendations from the MR, and change the patients drug regimen immediately. Both scenarios do not require the patient to wait to be seen by a doctor or pharmacist and negative effects can be counteracted straight away. The robot could also remind the patient to take their medications on schedule, and eventually dispense the medications, to help prevent under or over dosages. The robot could easily cope with complex medication regimen, including reducing dosages over time.

Doctors may also add new general rules as new drugs are introduced, or as more information is gained from clinical trials. The doctor may also personalise the knowledge base with rules that apply to the specific patients response to certain medications. For example; unacceptable side effects from certain ingredients or apparently ineffective ingredients.

2.1 Robots and Expert Systems

Current robot research concentrates almost entirely on the low level functions, such as navigating and recognising objects. Although there is no doubt these functions are necessary, this work addresses the general problem that users will soon have expectations of high level intelligence from their robots. The robots will need to solve problems at an

expert-level. The problem also remains for how these high and low levels will be integrated. This project aims towards a generic system that links lower level functions to an expert system. The specific knowledge domain and lower level functions should be easily customisable to the problem domain and robot platform.

Although expert systems are considered to be one of Artificial Intelligences success stories, the technology is still not being employed to its potential. Most expert system techniques are not reliable or maintainable enough to inspire the necessary confidence. Ripple Down Rules are one expert system method that overcomes these issues.

2.1.1 Ripple Down Rules

Ripple Down Rules are a method for adding rules incrementally to form a knowledge base. The rules can not only be created from batch learning, but also from on-line learning. RDRs lend themselves to robotic applications, because they can be constantly updated as the robot has new experiences in the environment, but are also simple enough to be hand-crafted by experts. RDRs allow for recovery by adding exception rules and new rules are added when a classification error is made.

Expert systems based on Ripple Down Rules are in production in various areas such as a major teaching hospital's chemical pathology laboratory, providing clinical interpretations of data for diagnostic reports [4]. Ripple Down Rules have also been used for online learning in robots [2]. Multiple Classification Ripple Down Rules (MCRDR) is an extension that allows the system to suggest several classifications [5]. The MR robot will use a variety of interfaces to satisfy conditions of rules in the MCRDR rule structure of Bindoff [1]. Once those conditions are met we can provide the user with conclusions or recommendations regarding their medication.

2.2 Platform

The robot in use for this project is a MOSTiTECH MIR mobile robot pictured in Figure 1. The MIR robots have a camera, touch screen, LED face, motion sensor and touch sensors. It has a videophone built onboard and may contact the doctor directly if it senses a potentially dangerous situation.

2.3 Conversational Agent

The process of performing a medication review can be defined by a series of goals; a series of steps required to accomplish those goals can then be planned. This system uses a step-based system, asking questions and receiving input from the user in a largely pre-planned order, allowing the

system to control the input that the user provides rather than asking open-ended questions that would allow the user to talk about anything - and to use context awareness to assist in the interpretation of the input.

When the user inputs a sentence, that sentence is broken up into individual words. Words that would commonly appear in sentences and are not medication-related are stripped out and ignored, as is punctuation. The system then searches the remaining words to find and extract relevant keywords that are related to the question that the users in answering.

Context awareness, in this case, will be used to refer to the method of interpreting information based on the current context that is, interpreting the information input by the user in regards to the context of the question. So, for example, if the system has asked the user what drugs they are taking, the system will then parse the input and extract all drug names it can then assume that the user is taking all of the drugs that they have named. Similarly, if the system asks what drugs the user has stopped taking, it will again find all drug names, but will assume that the user is not taking them. A question about medical conditions and observations will not look for drug names and vice versa by only searching for words that are relevant to the current question, the system can process input and produce output faster and with less computational cost.

While many conversational agents have been quite successful at conversing with the users and providing the appropriate feedback, they are primarily based on knowledge stored in non-dynamic predefined databases [5]. Another conversational agent of interest is the Medication Advisor project. Rather than performing medication reviews it helps users to manage their medication-taking by providing them with information and advice [3], and, again, it is based on a predefined and stable knowledge base. The conversational agent for this study, however, being based on an MCRDR knowledge base can be dynamically updated and expanded, and cope with the changing rule trees.

2.3.1 Exploiting the MCRDR for the conversational agent

As previously mentioned, we would like to reduce the load on the patient as much as possible. Therefore the conversational agent exploits the MCRDR rule structure in order to ask fewer questions. We ran a series of 15 test cases using real patient data. The number of questions aimed at clarifying the information entered by the patient was recorded. The number of questions asked was reduced by 79.3% without any reduction in accuracy [9] as compared with the original system of [1].

2.4 Input and Output

Figure 1 illustrates the interfaces between patient, robot, conversational agent and expert system. The robot speaks information to the patient in the form of questions that might help make new conclusions, and the conclusions themselves. The robot achieves this using text-to-speech. The patient provides answers to the questions via the robot using a variety of interfaces including the touchscreen, touch sensors and speech. These unparsed sentences are sent to the conversational agent, which parses them to find conditions. The conditions are then fed to the expert system to match on rules. The robot also gains sensor readings from the patient either from processed images, or medical sensors such as temperature or pulse. These readings are directly used by the expert system as input conditions. If the expert system is able to infer any conclusions from the input conditions from conversational agent and robot sensors then the conclusions are sent to the conversational agent. The conversational agent parses the conclusions and converts it into an easily understandable sentence. This text sentence is then sent back to the robot which uses text-to-speech to alert the patient.

3 An Interaction Example

In this section, we will work step-by-step through a patient's interaction with the robot.

3.1 Ask questions

Before the robot asks questions it ensures that its body is facing the patient, is an appropriate distance from them and then pans and tilts its head to appear to be maintaining eye contact. The robot's sensors will then be within easy access for the patient, as well as enforcing in them that the robot is paying attention to them. The apparent eye contact is achieved by searching for a face in the space and continuing to track it during the interaction. The robot requests information from the patient in case the medical record needs to be updated:

Robot: "Have you stopped or started any medications?"

Patient: "I have stopped taking Lithium and started taking Metformin"

The conversational agent parses the sentence which then triggers an update of the medical record. The robot then asks:

Robot: "Have you been experiencing any new symptoms?"

Patient: "I have been feeling a bit nauseus."



Figure 1. Input and output between patient, robot, conversational agent and expert system

3.2 Medical Record

One of the main aims of our approach is to minimise the input requirements from the patient. Therefore we attempt to use any information we can gain that will not interfere with the patient. The robot has the patient's medical record, and will also store answers the patient has previously given. The next step is to load any information from the medical record that can be used as conditions for the expert system. In this example, the patient's medical record states that they suffer from Anaemia, Hypertension and Ischaemic heart disease (IHD) and are taking Iron supplements and a Tricyclic Antidepressant. As the patient informed us they have begun taking Metformin then we also add the drug family Biguanide.

3.3 Robot sensing

There are a variety of medical sensors that may be connected to the robot via serial or USB port. These include blood pressure, temperature, pulse and blood sugar. As this patient suffers from Anaemia we have decided to connect a haemoglobin meter. The patient is asked to prick their finger and drop a sample into the meter. The resulting reading is 89 and this is added as haemoglobin female once we have gained the patient's gender from the medical record.

The robot also attempts to estimate the patient's hip and waist ratio from vision, and from that determine their Body Mass Index (BMI). BMI can be used to evaluate obesity, which, if present, will have an impact on the effectiveness of the patient's medication. In this example, the patient was found to have a BMI of 26, which is considered to be overweight, but not obese. We may develop this subsystem further to make dietary recommendations based on the patient's BMI.

3.4 Check for conclusions

Once we have asked the questions and checked the medical record we have a list of conditions we can attempt to match with rules. The system traverses the MCRDR rule tree travelling down branches in order to find the deepest-level conclusions that have all their rule conditions satisfied. The conditions we have are:

General Information: Age = 62, Sex = Female, Marital status = Married.

Current Conditions: Nausea, Anaemia, IHD, Hypertension.

Sensor Readings: Haemoglobin female = 89, BMI = 26.

Current Drugs: Metformin, Iron, Antidepressant tricyclic.

The conditions for this patient match on the following rules:

if Nausea == Current \wedge Metformin == Y \wedge Age \geq 50
then Nausea while on Metformin

if Anaemia == Current
 \wedge Haemoglobin female \leq 110 \wedge Iron == N
then Ongoing Anaemia Despite Treatment

if Biguanide == Y \wedge Haemoglobin female \leq 110
then Anaemia with Predisposing Drug

if Antidepressant tricyclic == Y
 \wedge IHD == Current
then Multiple potential causes of arrhythmia

3.5 Ask questions based on partial matches

The following rules also match partially. We then use the non-matching conditions of these rules to ask further questions of the patient.

if IHD == Current \wedge Chest pains == Current
then IHD with ongoing chest pain

if Shortness of breath == Current
 \wedge Hypertension == Current \wedge CCF != Current
 \wedge IHD == Current
then Shortness of breath in patient at risk of Congestive Cardiac Failure (CCF)

The conversational agent takes the non-matching conditions and forms questions that the robot speaks aloud to the patient.

Robot: "Have you been experiencing chest pains?"

Patient: "No."

Robot: "Have you been experiencing shortness of breath?"

Patient: "Yes, sometimes."

Now the second rule matches, but not the first. We add the warnings regarding CCF to the list of conclusions.

3.6 Make recommendations

Once all information has been and clarified, and the applicable conclusions have been found, the user is informed of the results.

Robot: "There are some problems with your medication, please contact your GP or pharmacist as soon as possible. Would you like more information?"

Patient: "Yes."

Initially they are simply told that there could be some potential medication-related problems and they should see their medical professional. They are then asked if they want more information; if they agree they are presented with information and each specific conclusion that is found. The system itself does not advise the user to take any action, but to see a human expert to get advice on what they should do, as incorrect or misunderstood information from the system could have dire medical consequences. The conclusions we found were:

- Nausea while on Metformin
- Multiple potential causes of arrhythmia
- Anaemia with Predisposing Drug
- Ongoing Anaemia Despite Treatment

- Shortness of breath in patient at risk of Congestive Cardiac Failure (CCF)

Unfortunately the conclusion names have been entered by pharmacy rather than software experts. Therefore there is not a standard format for conclusion naming. However, the conversational agent attempts to phrase the conclusions in a natural way and use information from the conditions used to fire the rule. For example:

Robot: "There is concern with your Metformin medication contributing to your Anaemia and your Antidepressant tricyclic contributing to your arrhythmia."

We may choose to reduce the information provided to the patient. Rather than alarming them with discussion of Congestive Cardiac Failure we can say:

Robot: "Please inform your doctor of your shortness of breath in relation to your IHD."

The robot is equipped with a video phone and wireless internet can hold contact information for each patients GP and/or pharmacist. Upon reaching the end of the review an e-mail can be compiled and sent to them containing the medication review and the precise details of all problems found and the causes of those problems. They could then use this information to develop a care plan for the user. The video phone could be used in emergency situations, such as when the robot determines the patient has collapsed.

3.7. Conclusions

The possible applications for robots greatly increases when we bridge the gap between low-level robotic sensing and high-level expert knowledge. However, as more people are able to use robots for their needs, the role of the interface becomes increasingly important. High-level problem solving robots will only be as successful as the users satisfaction. Therefore ensuring that the interface is efficient and effective will have significant impact on the uptake of RDR-based robotic systems. The Medical Review system provides us with many types of interaction between human and robot and give us the opportunity to address problems that will be common to many applications.

References

- [1] I. Bindoff, P. Tenni, B. Kang, and G. Peterson. Intelligent decision support for medication review. *Advances in Knowledge Acquisition and Management, Lecture Notes in Artificial Intelligence*, 2006.
- [2] C. D'Este and C. Sammut. Active object discovery for communicating with humans. In *Australasian Conference on Robotics and Automation (ACRA 2005)*, Sydney, Australia, 5-7 December 2005.

- [3] G. Ferguson, J. Allen, N. Blaylock, D. Byron, N. Chambers, M. Dzikovska, L. Galescu, X. Shen, R. Swier, and M. Swift. The medication advisor project: Preliminary report. Technical Report Technical Report 776, CS Dept., U. Rochester, 2002.
- [4] B. Kang and P. Compton. Maintenance approach to case based reasoning. In *Advances in Case-Based Reasoning, EWCBR-94*, 1994.
- [5] B. Kang, P. Compton, and P. Preston. Multiple classification ripple down rules: Evaluation and possibilities. In *The 9th Knowledge Acquisition for Knowledge Based Systems Workshop*.
- [6] J. Lytle. Robot bears care for the elderly. <http://news.bbc.co.uk/1/hi/sci/tech/1829021.stm>, February 2002.
- [7] G. Peterson. The future is now: the importance of medication review. *Australian Pharmacist*, pages 268–75, 2002.
- [8] M. Pollack, S. Engberg, J. Matthews, S. Thrun, L. Brown, D. Colbry, C. Orosz, B. Peintner, S. Ramakrishnan, J. Dunbar-Jacob, C. McCarthy, M. Montemerlo, J. Pineau, and N. Roy. Pearl: A mobile robotic assistant for the elderly. In *Workshop on Automation as Caregiver: the Role of Intelligent Technology in Elder Care (AAAI)*.
- [9] D. Reid. A conversational medication review interface. Master's thesis, School of Computing, University of Tasmania, 2008.
- [10] D. Rigby. The challenge of change establishing an hmr service in the pharmacy. *Australian Pharmacist*, pages 214–217, 2004.