

KARDIO - A simulation of a cardiac care unit intended for the study of the ethical components of medical decision-making

Søren Holm¹, Peter Øhrstrøm², Claus Donner³

1. Centre for Social Ethics and Policy, University of Manchester, Oxford Road, Manchester M13 9PL, UK (soren.holm@man.ac.uk)

2. Department of Communication, Aalborg University, Langagervej 8, 9220 Aalborg øst, DK (poe@hum.auc.dk)

3. Department of Communication, Aalborg University, Langagervej 8, 9220 Aalborg øst, DK (cd@hum.auc.dk)

Abstract

The paper describes the KARDIO system - a computer simulation of a cardiac care unit (CCU) intended for use in the study of the ethical aspects of the decision-making of medical doctors. The rationale behind the system is described, as well as the trade-offs in the design phase. One of the underlying mathematical models is described in detail. The last part of the paper deals with the response of doctors towards the simulation.

Introduction

There is a growing interest in medical decision-making, as well as in medical ethics. This interest has also led to a growth in empirical studies looking at the ethical values of physicians and how these values influence the decisions they make.

This report describes the KARDIO simulation that simulates a cardiac intensive care unit (CCU). KARDIO was developed as a tool for the collection of data on medical decision-making, and the considerations behind its design are outlined below.

Methodological problems in the study of ethical decision-making

Studies of medical decisions in real life settings raise a number of problems of a both practical and research ethical nature. (See Holm et al 1993, 1996, 1997)

There are often practical problems in being allowed access to the medical settings, and data collection may also be complicated.

To limit the time needed for data collection many studies are done on the basis of medical records. This raises methodological problems, since the information in medical records is an edited/reconstructed version of what really happened. If the researcher is present in the medical setting it does, however, also create methodological problems because the mere fact of his or her presence may change the decision-making process in unpredictable ways.

The research ethical problems are connected to the need to obtain informed consent from doctors and patients. Even though these studies are primarily interested in the thought processes and decisions made by the doctors, the patients are inevitably drawn in, since what makes these decisions interesting is mainly, that they are made about patients and based on information about patients. It is impossible to study the decision-making of doctors without making the patients subsidiary study objects.

These problems have led a number of researchers in the field of empirical ethics to pursue a methodology utilising case-based questionnaires, i.e. questionnaires where a short clinical case is presented, and where the respondents are asked to answer a number of questions about what they would do in such a situation.

Some studies attempt to build in a temporal progression in the cases. This creates technical difficulties in a paper-and-pen questionnaire, because it is difficult to prevent the respondent from looking ahead to see what comes next in the case. It is also difficult to provide branching opportunities depending on the choices the respondent makes.

The case methodology is usually chosen in order to try to get behind the values doctors claim to have to the values they really have, and which are demonstrated by their responses to

cases. It would be possible to try to ask doctors directly what their values are, but such a study would encounter serious problems with respondents knowing the socially acceptable answer. Very few physicians would ever admit that socially unacceptable values influenced their decisions (e.g. racist or ageist values), even if they did.

A computer simulation seems the obvious solution to the problems of case based questionnaires in the study of the ethical aspects of medical decision-making. It does not create the same practical and ethical problems as a study of a real life setting, and it can, if the simulation is suitably constructed, present the respondent with a "naturally" developing scenario, where there is no possibility of looking ahead, and at the same time responsiveness to the decisions the respondent makes.

Features of the ideal setting to be simulated

The ideal setting for a study of decisions concerning ethical problems is a setting where such problems occur with a relatively high frequency, and where they are not "solved" in advance by legislation or a traditionally accepted solution.

Given the large human and economical investment in developing a simulation there are benefits in choosing as setting where a variety of ethical problems occur, since this will make it possible to study a number of these problems using the same basic simulation.

For pragmatic reasons it is also reasonable to choose a setting where common ethical problems occur so that the results may be extrapolated to other health care settings.

A setting where important ethical problems occur must also be preferred to one where only trivial problems occur. Finally the setting should not be so specialised that only very few physicians will possess the necessary factual knowledge to run the simulation.

An ideal setting would thus exhibit the following five features:

1. Frequent ethical problems
2. A variety of ethical problems
3. Ethical problems which are common and not "exotic"
4. Ethical problems which are important
5. Not too specialised a setting with regard to factual knowledge

Features of the ideal simulation

The features of an ideal simulation cannot be described in abstraction but depend on what purpose the simulation is being designed for. (See Øhrstrøm et al. 1991, 1992, 1996) If the intended use of the simulation is to train a person to perform a certain complex set of cognitive and motor skills (e.g. a cockpit simulation intended for pilots) a certain type of very life-like simulation is needed (real-time, full scale etc.). If the purpose of the simulation is more modest it may be neither necessary nor prudent to aim for absolute similarity to real life.

In the present context it is, for instance, possible to study more decisions within a given time-span if the simulation is not running in real time.

Given that doctors are usually not a captive audience who can be forced to run the simulation, it is also important that the simulation is perceived as interesting and worthwhile by physicians.

Finally there are always pragmatic considerations pointing in the direction of a simple simulation. An ideal simulation for the present purpose would thus exhibit the following four features:

1. As close to the real setting as possible
2. Interesting for physicians
3. Not too time consuming to run
4. Not too complicated to develop

Why simulate a cardiac care unit?

A number of medical settings can fulfil the criteria mentioned above. We have chosen to simulate a cardiac care unit (CCU) because we believe it fulfils some of the criteria for an ideal setting. It is, at the same time, also a setting, which it is possible to simulate in a way, which does not leave the simulation to far from the ideal simulation.

The patients in a CCU are critically ill, and the treatment of individual patients raise ethical problems concerning initiation and withdrawal of treatment, and resuscitation. There are at the same time also resource allocation problems since the number of beds in the CCU is limited, and it must be decided which patients should fill these beds. These problems are not unique to the CCU but occur many other places in the health care setting, although seldom with the same frequency. The problems do also fulfil the importance criterion since they concern life and death.

The CCU is a fairly specialised treatment setting, and a not insubstantial amount of factual knowledge is necessary.

The range of diseases that a patient in a CCU may present with is large, but there are few main groups which makes up the bulk of the patients. These are primarily myocardial infarction, disturbances in heart rhythm, and congestive heart failure. A simulation, which can handle these three main groups, will therefore be able to present a realistic picture to the respondent, and still be fairly simple.

The decision to simulate a whole CCU and not just individual patients was made in order to be able to study prioritisation problems. These problems are very common in health care, and they make it possible directly to study the influence of various patient characteristics (age, social status, gender, ethnicity etc.) on treatment and admission decisions.

There are a number of single patient simulations of patients with cardiac problems available, but pharmaceutical companies almost invariably produce them in order to show the (beneficial) effects of their products.

How does one read of ethical values from decisions in a simulation?

Decisions concerning clinical situations where there are ethical problems seldom come labelled as specifically ethical decisions.

One possible solution to this problem is by presenting the physician with two patients who are exactly similar in a medical sense (same diagnosis and prognosis), but who differ in one additional factor like age, ethnicity, or social status. If physicians consistently choose one of the patients for admission, or consistently treats one patient more aggressively than the other, there is strong evidence for some (ethical) value playing a part in this decision. In a simulation one can control the patient flow, and thus deliberately create situations where the physician has to choose between medically equivalent, but otherwise different, patients. In terms of creating situations to test ethical values related to age and ethnicity, making these variables explicit may, however, conceivably assign more pressure on the doctor to pick the "politically correct" answer. In real life, these differences are more implicit. In a simulation it is also possible to control other important factors like the stress level, and the pressure from outside authorities.

The KARDIO simulation

Based on the considerations above it was decided to develop a simulation of a 6-bed cardiac care unit, with a limited number of possible diseases, and a standard screen interface. We feel that this represents a reasonable compromise between the ideal demands, and the practical constraints in a project with limited resources. It gives us the possibility to study decisions concerning individual patients, as well as decisions concerning prioritisation.

In order to be certain to obtain data relevant to the analysis of ethical values, we decided to make certain decisions obligatory in the simulation. From our information gathering (see below) we know that it is routine in many CCUs to decide whether or not a given patient should be resuscitated, and whether or not the patient can be moved to an "ordinary" medical bed, following each ward round. We therefore made these two decisions obligatory, since they provide a pointer to the priorities of the physician running the simulation.

The user interface

The user interface follows standard Windows conventions, and is designed to require a minimum of keyboard input. This means that the basic user interaction is carried out by means of buttons and mouse operations. The interface is made as graphical as possible although there is at present no real-time video or speech involved.

Opening window

The opening window consists of a graphical representation of the 6 beds in the CCU. Under each bed there is a field displaying the name, age, pulse, blood pressure, and pain level of the patient in the bed.

The opening window also contains a clock displaying the remaining time for the present wardround and buttons for ending a ward round, or ending the program.

Patient record window

The patient record window contains data about the individual patient. A picture of the patient is presented in the upper left-hand corner, and just below this is a field with the patients name, occupation, and age. The scrollable patient record occupies most of the right side of the window. Below this there is a continuous one-lead monitoring ECG. The patient record window is designed to give a quick overview of the state of the patient, and enable the physician to make initial decisions about further information gathering or initial Treatment decisions.

Pulse and blood pressure graph

This graph displays the patients pulse (red markings), blood pressure (blue markings), and temperature (green markings) against time. It is identical with the graph often found in patient records, or attached to the end of the patients bed.

Doctors orders

This window is the gateway to further information or to the implementing of treatment.

On the left hand side is a column of buttons where various information can be ordered. From top to bottom the buttons are "blood biochemistry", "X-ray", "12-lead ECG", "echo-cardiography" (not yet implemented). The last button in this column moves the physician to the treatment ordering windows (See below). The buttons for diagnostic tests are duplicated on the right hand side of the window, where the doctor gets access to the most recent results as well as to any previous results.

Between these two columns of buttons is a list with four options for the discharging of the patient: "no discharge", "discharge to ordinary medical ward", "discharge to home", "dead". When a patient is discharged his or her bed becomes empty.

Empty beds are a precondition for new admissions. During the running of the simulation the doctor will at various intervals be presented with new patients in the emergency department. In each case the doctor will have to decide whether or not the patient should be admitted to the CCU. If the doctor decides to admit, but there are no empty beds, a bed must be made available by the discharge of one of the patients already in the CCU.

Blood biochemistry

The blood biochemistry module consists of two windows an ordering window and a result window. The normal range is listed.

Chest X-rays

The X-ray module also consists of an ordering window where chest X-rays (urgent or non-urgent) can be ordered and a result window. The results are presented as text.

12-lead ECG

This button gives immediate access to a 12-lead ECG.

Treatment ordering windows

A very large number of drugs are relevant in the treatment of patients in a CCU, and some of these drugs are available in a range of different formulations and strengths, and under different proprietary names. The windows of the treatment ordering module are therefore designed hierarchically.

Cardiac arrest

During a ward round an alarm may sound and indicate that one of the patients is in cardiac arrest. The window changes to the patient record window for the patient in question where the ECG shows a pattern compatible with ventricular fibrillation. The physician is asked whether or not resuscitation should be attempted. If resuscitation is not attempted or if it is unsuccessful the patient dies, and the physician is asked to discharge the patient.

If resuscitation is successful the patient is resuscitated to the worst of the haemodynamic states in the model, and the physician will have to take steps to treat this situation.

New patient

When a new patient presents him- or herself in the emergency department a window is shown which is basically identical to the patient record window, except that no ECG is shown. Based on the information in the patient record the doctor will have to decide whether to admit the patient to the CCU, to an ordinary medical ward, or not to admit at all.

DNR-status and transferability

Between two ward rounds the physician is asked to make decisions for each patient concerning two items: a) should the patient be attempted resuscitated if he or she goes into cardiac arrest (i.e. Do-Not-Resuscitate status (DNR-status)), and b) can the patient be transferred to an ordinary medical ward if a new patient presents in the emergency department (transferability).

The decisions do not influence the simulation in any way, and do not bind the doctor if the relevant situation occurs. Similar decisions are made and recorded in real CCUs.

Logging of data

All user decisions and actions not included in the patient record are continuously logged in a separate log-file. This file is not accessible by the user. All events are time stamped. This means that it is possible to describe the users interactions with the system in great detail, including the users information seeking strategy. The decisions taken between two ward rounds

concerning DNR-status and transferability for each individual patient are also logged and provide data for simple analysis.

Collection of domain specific knowledge

KARDIO does, in a certain sense, work as an inverted expert system. It does not produce a diagnosis based on information about the patient, but instead information about the patient based on a diagnosis and a mathematical model. The development of KARDIO therefore required collection of both explicit and implicit domain knowledge. This knowledge was collected from the medical literature and from discussions with doctors and a visit to a CCU.

We identified the following pieces of information as important for the clinical assessment of an individual patient:

- a. History
- b. Current state
 - Clinical presentation
 - Pulse
 - Blood pressure
 - Temperature
 - ECG (single lead continuous and twelve lead)
 - Stethoscopic findings
- c. Paraclinical findings
 - Chest X-ray
 - Blood biochemistry
 - Echo-cardiography

We decided to implement all these in the simulation, but due to time constraints echo-cardiography has not yet been implemented.

The most difficult design decision related to the specific domain was the decision concerning how often the simulation should be updated, and what the connection should be between real time, and the time scale of the simulation. In the end we decided to build the simulation on a model containing two "ward rounds", a morning and an evening "ward round". Ward rounds do occur in real life, and they are a major feature of the social organisation of hospital life. By choosing the ward round as the unit of time, we were able to compress the time scale considerably, and thereby to study more decisions, before the respondent becomes tired of running the simulation.

The mathematical model of disease development in myocardial infarction

An underlying mathematical model governs the development of the clinical condition of each patient. The following stochastic model has been devised for myocardial infarction.

The model deals with the temporal development of a number of patients. For each patient a risk parameter R (between 0 and 20) as well as an infarct size IS (between 0 and 20) are given. For each patient there is also given a period of time TT (in hours) from the first occurrence of the heart disease to the arrival at the cardiac care unit.

In the model the states of the patients is basically described in terms of a variable corresponding to the haemodynamical state (HD). This variable can range over the values n (normal), lvf (left ventricular failure), sh (shock), and d (dead). We shall use n=4, lvf=3, sh=2, and d=1.

The temporal model is discrete. The time parameter is denoted T. The states of a patient are only updated every 12 hours and when certain drug injections are given to the patient. At every instant of updating a random number RN, is picked and a new HD value is calculated as a function of the RN and the old HD value. It should be noted that without medical treatment the HD value cannot increase.

if $1 - (PP(1) + PP(2) + PP(3))$ is negative then

$$PP(4) = 0$$

else

$$PP(4) = 1 - (PP(1) + PP(2) + PP(3)).$$

$$PP(3) = R/40$$

$$PP(2) = R/60$$

$$PP(1) = 0,05 + R/200$$

These PP values hold for T=1. For the following instants of updating the PP values are reduced according to the following recursive formulae:

$$PP(3)_{[T=i+1]} = PP(3)_{[T=i]} \exp(-\ln 2 * 0,33)$$

$$PP(2)_{[T=i+1]} = PP(2)_{[T=i]} \exp(-\ln 2 * 0,5)$$

$$PP(1)_{[T=i+1]} = PP(1)_{[T=i]} \exp(-\ln 2 * 0,5)$$

The probabilities P(1), P(2), P(3), and P(4) are found by normalization of PP(1), PP(2), PP(3), and PP(4), such that the sum of the probabilities in the outcome space is 100%.

At every moment when the state of a patient is updated the pain level can be calculated as a function of the IS number, the value of T+TT, and the medical treatment. Also the temperature (Temp), the heart rate (PL), the blood pressure (P1 and P2), the ECG and the x-ray descriptions for chest X-rays are calculated as a function of HD, the previous HD (i.e. the HD at T-1), T+TT, IS and the medical treatment.

Temp is calculated as a variation from normal temperature i.e. 36,8 deg. Celsius in the morning and 37,3 deg. Celsius in the evening. If the patient has a myocardial infarction a small rise in temperature is modelled.

Calculation of pain as a function of IS:

$$PAIN(0) = 0$$

$$PAIN(X) = 1, \text{ for } 0 < X < 5$$

$$PAIN(X) = 2, \text{ for } 5 < X < 10$$

$$PAIN(X) = 3, \text{ for } 10 < X < 15$$

$$PAIN(X) = 4, \text{ for } 15 < X < 20$$

A time dependent coefficient is introduced in the following way, where T is the number of ward rounds since admission:

$$K(0) = K(1) = K(2) = 1$$

$$K(T) = 1 - (T - 2)/6, \text{ for } T = 3, 4, 5, 6, 7, 8$$

$$K(T) = 0, \text{ for } T > 8$$

Now a modified pain level can be calculated as follows:

$$\text{pain}(T) = PAIN(K(T) * IS) + 1$$

The modified pain level is then adjusted according to the pain relieving treatment given and transformed into a verbal description that is presented in the patient records.

The pulse is calculated from the HD value and the modified pain level in the following way:

$$PL(HD, T) = PL0(HD) * (1 + \log(\text{pain}(T) * 0,75) - \log 0,75)$$

where

$$PL0(4) = 72, \quad PL0(3) = 90, \quad PL0(2) = 60, \quad PL0(1) = 0$$

PL(2), PL(3), and PL(4) are changed randomly by maximally 10%.

The blood pressures are also calculated from the HD value in the following way:

$$BP(HD, T) = BP0(HD) * (1 + \log(\text{pain}(T) * 0,75) - \log 0,75)$$

where

$$BP0(4) = 140 \text{ (diastolic: 80)}$$

$$BP0(3) = 100 \text{ (diastolic: 60)}$$

$$BP0(2) = 70 \text{ (diastolic: 40)}$$

$$BP0(1) = 0$$

BP(2), BP(3), and BP(4) are changed randomly by maximally 10%.

If the systolic blood pressure BP (syst.) comes under 50, the patient dies (i.e. HD=1)

Drug effects can be of three different kinds in the model:

1. Initial reduction of IS
2. Change in transition probabilities
3. Symptomatic effects

1. One class of drugs (thrombolytic drugs) can reduce the basic risk parameters R and IS if given immediately upon arrival of the patient.

2. A number of drugs can change the transition probabilities between the 4 haemodynamic states, or can improve the patients haemodynamic state (the latter is impossible without treatment).

3. Drugs with symptomatic effects will change the patients pain-level, blood pressure etc. but will not change any of the parameters of the underlying model which determines the future course of the patients development.

For each drug, or group of drugs, a simple pharmacokinetic and pharmacodynamic model is implemented governing onset and decay of the pharmacological effect.

For many drugs there is a lethal dose such that the patient dies immediately if more than the lethal dose is given.

Initial studies

Initial validation studies have been carried out where a number of physicians with varying levels of experience have been asked to run the simulation, and comment on the design of the user-interface, the "behaviour" of the disease models, the response to treatment etc..

These initial studies showed that the user-interface is accepted by physicians, that they are willing to run the simulation for up to 2 hours, and that they can at times become engrossed in treating their "patients". The testing thus indicates that the simulation is sufficiently life-like to be valuable in the study of decision-making. We therefore decided to move on to research use of the program.

Our initial studies of medical reasoning using a scenario designed to investigate if and how the chronological age of the patient influences treatment decisions, show that doctors do not consciously pick up on the fact that their decisions have ethical implications, even in situations where they discriminate against older patients. They see the task as a purely "technical" task, to be solved at a technical level.

Acknowledgements

The research was supported by grants from the Danish Heart Association, the Danish Healthsciences Research Council, and Fonden af 1870.

References

Derczynski, W.; Holm, S., and Szawarski, Z. 1993. Attitudes to capital punishment: a comparison of Polish and Danish doctors. *Bulletin of Medical Ethics* 88: 17-24.

Holm, S. 1995. The medical hierarchy and perceived influence on technical and ethical decisions. *Journal of Internal Medicine* 237: 487-92.

Holm, S.; Gjersøe, P.; Grode, G.; Hartling, O.; Ibsen, K.E., and Marcussen, H. 1996. Ethical reasoning in mixed nurse-physician groups. *Journal of Medical Ethics* 22: 168-73.

Holm, S. 1997. Metodi standardizzati per la rilevazione dei dati in bioetica. In: Gindro, S.; Bracalenti, R.; and Mordini, E. (eds.). *La Ricerca in bioetica: politiche, metodi, strategie*, 41-46. Rome: CIC International.

Nielsen, F.; and Øhrstrøm, P. 1991. *Design and Implementation of a Symptom Language used for Diagnosing Component Faults in a Nuclear Power Plant, RISØ-I-582*. Risø: Risø National Laboratory.

Øhrstrøm, P. ; Nielsen, F.R.; and Pedersen, S.A. 1992. Diagnosing Component Faults in a Generic Nuclear Power Plant Using Counterfactual and Temporal Reasoning. In

Proceedings from 4th MOHAWC Workshop, 62-87. Risø: Risø National Laboratory.

Øhrstrøm, P., and Andersen, H.B. 1996. *SCENARIO ASSISTANT - Using a Tense Logical Language to Handle Scenarios Supporting Training of Coordination and Communication*. Risø: Risø National Laboratory.