Consequence Logics for Cancer Risk Counselling Systems: A Specification Framework

Peter McBurney and Simon Parsons
Department of Electronic Engineering
Queen Mary and Westfield College
University of London
London E1 4NS
United Kingdom

John Fox and David Glasspool
Advanced Computation Laboratory
Imperial Cancer Research Fund
P. O. Box 123
Lincoln's Inn Fields
London WC2A 3PX
United Kingdom

November 5, 1999

Abstract

We outline a preliminary specification framework for a prognostic expert system to advise General Practitioners and their patients on the individual risks of breast cancer. The system is prognostic in the sense that it considers and presents the likely consequences of alternative monitoring, prophylactic, treatment and lifestyle options available to the patient. In addition to predicting cancer risk under various genetic, lifestyle and treatment scenarios, the system is intended to assist patients to make informed decisions regarding their treatment options. To do so requires consideration of the likely consequences of actions, the impacts of these as perceived by the patient and the actual decision-making style(s) of the patient concerned. Such information is generally only qualitative in nature or, if quantitative, is often not understood well by patients. Consequently, a qualitative formalism is required and we propose a Logic of Consequences using argumentation. We argue that it is important that the system adopt a neutral (non-normative) approach to the decision-making style of patients.

Key words: Argumentation, cancer, consequence, intelligent systems, risk counselling.

1 Introduction

The Imperial Cancer Research Fund's (ICRF's) Risk Assessment in Genetics (RAGs) project aims to develop theoretically-grounded decision support systems for prediction of health risks and their communication to patients. Work on the project is being carried out in two major streams. The first is concerned with developing a general theoretical approach to the representation of uncertainty, while the second applies this theoretical background to the development of decision support software for a specific medical application involving the communication of risk: Genetic counselling of women at risk from breast cancer. The methods developed, however, are intended to be applicable to any kind of genetic counselling and the management of personal risk generally

One of the results of the revolution in molecular genetics is the increasing availability of tests for the presence of genetic abnormalities which predispose certain individuals to contracting particular diseases at some point in their lives. For example, the location of the BRCA1 gene for familial breast cancer has recently been determined on chromosome 17, and it is confidently expected that a reliable test for the presence of the gene will be available within one or two years. In the UK alone this will impact of the order of 30,000 women who are estimated to carry the gene. Given a positive test (which implies an 80% lifetime risk of contracting the disease) such women will require supportive counselling to help them understand their individual circumstances and their options for monitoring, prophylaxis and therapy.

The aim of genetic counselling is to provide individuals with information about their personal risk of developing a disease in such a way as to encourage appropriate health care decisions (e.g. to attend for regular screening and increase the chances of early detection). Unfortunately counsellors are not well served with techniques which are proven to facilitate understanding of personal risks. To remedy this, and as part of the RAGs project, the ICRF has developed a decision support system for breast and ovarian cancer [3, 4]. The system is designed to assist the General Practitioner (GP) in assessing and communicating genetic risk information to women concerned about their personal risk. The tool consists of a user interface which allows a family tree to be constructed for the presenting patient, incorporating information on the incidence of cancers. This interfaces with a risk calculation protocol which generates an assessment of risk level and recommendations for patient management, along with an explanation for the conclusions reached. The user interface component is designed to be generic and may interface with different protocol software to assess the risk of different genetic conditions.

The next stage of development of this intelligent system is intended to assist the patient to make an informed decision as to her monitoring, prophylactic, treatment and lifestyle options. For instance, a woman's risk of breast cancer at any age may be impacted if she chooses to have subsequent children. When faced with personal decisions of a serious nature, most of us weigh the arguments for and against each possible action in some way, including the likely consequences and our evaluations of their expected impacts on us. Motivated by this example, we propose the use of a formalism, akin to argumentation formalisms, for the representation and manipulation of the likely consequences of the patient's decision choices. Argumentation formalisms have been used for some time in expert systems, especially in medical and legal domains, and the ICRF has been a pioneer in this regard [7] especially in the use of argumentation in risk assessment [14, 22]. This paper outlines our position on the use of such formalisms for breast cancer risk counselling systems by specifying at a high level the requirements the formalism will need to satisfy in this particular domain.

2 Making Medical Decisions

In designing an intelligent system to assist patients to make important medical decisions, we need to understand how such decisions are made. Traditionally, at least in Western cultures, most medical decisions were made by doctors and medical personnel, on behalf of their patients. This practice has changed over the last quarter-century, at least in the English-speaking world. One catalyst for change was the publication in 1984 of Katz's book [13], which argued for doctors to respect the decision-making autonomy of patients. Katz urged doctors to share decision-making with patients, and to engage in a genuine dialogue with patients and their families to achieve this. Another catalyst was the rise of malpractice litigation, in which patients increasingly have taken legal remedy against what they perceive as professional incompetence, malfeasance or disrespect of their individual rights. It has thus become common, at least in Western countries, for patients to be asked by medical or hospital personnel for their "informed consent" before an operation or major course of treatment is undertaken.[1]

As Schneider [25] has noted, this approach implicitly assumes patients are deliberative and rational decision-makers, in the sense of normative decision theory [17]. However, little is known about how patients actually make medical decisions, as this topic has not been the subject of much research. What limited research has been undertaken appears to indicate that patients typically make decisions more quickly than is required, often without gathering a great deal of information, and often using what may be termed "gut instinct" [5, 12, 26]. In addition to making decisions quickly, patients appear to use processes which do not conform to the normative models of classical decision theory: for instance, they typically do not consider the consequences of all the available options, or they resist any revision of the initial decision in the light of new information [19, 23, 26]. Given this lack of conformance to normative decision theory, it is not surprising that patients also use a diversity of approaches to such decisions. In his recent review of the limited research on the subject, Schneider concluded [25, p.97]:

[1] That this is a recent phenomonen can be seen in the practice of medical personnel and researchers as late as the 1960s and 1970s to apply treatments without the consent of patients or experimental subjects [9].

"In sum, patients often do not make decisions the way students of decisions, or even most people, would advise, and certainly not the way the law of informed consent contemplates or autonomists expect when they speak of 'substantial understanding' *and 'substantial* absence of control *by others,' of 'the informed exercise of a choice' with the 'opportunity to evaluate knowledgeably the options available.' "*

Saying that patients make decisions in ways which are inconsistent with normative decision theory is not to criticize these decision-methods. We have, as Schneider [25, p.92] also notes, few means to measure the quality of medical decisions. Selecting one procedure or course of treatment usually precludes the selection of alternatives, and so comparison of results of alternative decision options at an individual level are impossible. The diversity and complexity of individual circumstances and medical aetiologies make comparisons at an aggregated level also problematic. Any intelligent system, then, which seeks to support individual patients in making major medical decisions needs to allow for a diversity of approaches to decision-making by those patients. Moreover, we believe it is important ethically that such systems not impose any particular approaches to decision-making on patients: the systems must be neutral regarding the decision-processes used, the criteria used to evaluate options, and the outcomes selected by patients.

Many intelligent decision support systems aim to embody human expertise about some domain, and aim to outperform human decision-makers; competitive comparison with human decision makers is often then a test of the effectiveness of the system. Our system for genetic risk counselling is different. We do not seek to out-perform a human decision-maker, but to assist her to make decisions according to the manner in which she would have done so anyway. We desire that the system will help inform her about her options and their consequences, and help her structure the decision, better than might otherwise be the case. The success of the system, then, will not be measured by out-performing the patient as a decision-maker; instead, success will be measured by the extent to which the patient feels comfortable about the decision process, and the extent to which she and her GP or risk counsellor continue to use the system over the course of her life. In the next section, we commence the process of specifying such a system by operationalizing these ideas of value-neutrality.

3 Towards a Specification Framework

Suppose at a given time t we are faced with a finite number of alternative actions, $A, B, C, \ldots, Z$. Suppose that each of these actions has a finite number of possible consequences, $A_1, A_2, \ldots, A_n$ for action A. The consequences and their number may be different for the different actions. Suppose further that each consequence has some value to us, as the decision-maker, $val(A_i)$, perhaps a utility or a loss. Suppose further that we have some measure of the likelihood of each consequence, given the relevant decision. We denote this by $pr(A_i|A)$. Nothing in the formalism we are presenting constrains the expression of consequences, values or likelihoods to be in quantitative terms; indeed, these variables will most often take values from qualitative dictionaries in the genetic counselling application.

Thus, we associate with each decision A, a list of consequences, their likelihoods and their consequence values, and write them all as a tuple (called a "consequence-tuple"), thus:

$$(A, A_1, pr(A_1|A), val(A_1), A_2, pr(A_2|A), val(A_2), \ldots, A_n, pr(A_n|A), val(A_n))$$

Note that the same action may have different consequences or the consequences may have different likelihoods or values when undertaken at different times or subsequent to different prior events. For example, the success rate of an operation may decline with the age of the patient beyond a certain age. Hence, we should properly associate each element of this tuple with a particular time t, thus (called a "timed consequence-tuple") :

$$(t, A_t, A_{t,1}, pr(A_{t,1}|A_t), val(A_{t,1}), \ldots, A_{t,n}, pr(A_{t,n}|A_t), val(A_{t,n}))$$

This formalism is similar in algebraic terms to formalisms we have previously used for argumentation [7, 21], with the consequences of an action represented rather than the grounds for a claim. Indeed, two of us have previously developed a formalism to represent and manipulate the qualitative values of actions, and a decision-calculus to enable choice between different actions on the basis of their expected values [8]. These formalisms were there called Logics of Value (LV) and of Expected Value (LEV). Because we intend our current approach (as explained below) to be neutral regarding the decision calculus used, we term the formalism presented here simply a Logic of Consequence (LC). In order to deploy this formalism in an intelligent system, we need to define its syntactical rules of manipulation and explore its formal properties. This paper commences this process by first articulating what is required of such a formalism in the domain of genetic counselling for risk.

Bearing in mind the comments of the previous Section, we propose that a Logic of Consequence for this domain must have the following properties.

Preference-Neutral: People's psychological preferences for, and reactions to, different types of treatment can vary markedly. Indeed, the same is true even of tests and monitoring programs: not everyone may wish to undertake certain medical tests, or acquire the knowledge which such tests provide, and they may have good, rational reasons for their preferences [2, 25]. It is crucially important, therefore, that the risk counselling system not impose a set of preferences on the patient or her advisor, nor that it be perceived by either of them as doing so. This means the LC must be preference-neutral. Indeed, successful elicitation and incorporation of the patient's own preferences and valuations will be a key element of system effectiveness. Researchers in the area of consumer behaviour have devoted considerable effort to developing effective means of eliciting consumer preferences and we will utilize this work [1, 10, 16].

Decision Calculus-Neutral: As noted earlier, people make medical decisions in many different ways, according to many different criteria, not all in accord with normative decision theory. We believe this is something to be celebrated, not decried. For decisions as important as the life-critical ones involving cancer, it is essential that individual patients are able to make their own decisions regarding monitoring, treatment and associated lifestyle choices, in their own time and in their own manner. For a decision-support tool, then, to actually support different means of decision-making and different criteria of optimality, the underlying LC must be neutral with regard to the decision-calculus. It must not prescribe or impose, nor be perceived to impose, a normative decision approach on the patient or her advisors.

For example, some patients may have religious or other reasons not to undergo surgery, despite knowing the consequences of not doing so, and this preference may dominate all their choices. Alternatively, many people are likely to avoid options with large downsides, such as painful or invasive treatment. An effective risk counselling system will respect such diversity of decision-styles and approaches.

Uncertainty Formalism-Neutral: Likewise, the risk counselling system and the underlying LC must be neutral with respect to the type of uncertainty formalism used. In the representation defined above, the elements of the consequence tuples may be quantitative or qualitative. For instance, the consequences may be in the form of prose lists, their likelihoods elements of a qualitative dictionary such as: {*Likely, Unlikely*}, and their values elements of a dictionary such as: {*Desirable, Not Desirable, Unacceptable*}. When uncertainties and valuations can be represented in quantitative terms, classical decision theory [17, 24] can be used to determine the best option according to some optimality or satisficing criteria, for example selecting that option which maximizes expected value. For certain consequences, such as the possibility of surviving for a given time following chemotherapy, accurate quantitative estimates are available. However, even if available, quantitative information may not be the best format for communications to ordinary people [6, 27]. Accordingly, the LC must be capable of coherent manipulation and presentation of qualitative information sources.

Tradeoff-enabled: The consequences of medical tests and treatments may not be perceived as uniformly positive or uniformly negative by patients, and indeed, they may not be so by any objective measure. Accordingly, choosing between options on the basis of their consequences is likely to require trade-offs between different options, none of which is uniformly preferable to all other options. The LC will therefore require an ability to represent and manage such trade-offs. Consumers appear to make similar trade-offs when deciding between products with multiple features, and marketing theory has developed multi-attribute models to represent such consumer purchase decisions [15, 16]. We believe these approaches provide a useful template for the cancer risk domain, although none to our knowledge have yet been constructed using qualitative dictionaries.

Time-enabled: In the formalism defined above, we noted the fact that decision-options and their consequences will vary by time. Indeed, to be supportive of the full complexity of decision-making in this domain, the advisory system will need to represent a rich variety of decision-time interactions, including: the same decision-option (e.g. to have a mastectomy) when selected at different times; the same option after different prior events (e.g. a mastectomy before and after having a child); different options at the same and at different times; and different sequences of options through time.

Normative decision theories are not good models, in general, of people's decisions when these involve time. For example, some people do not appear to use a constant discount factor when comparing events in the immediate future with those in the longer-term future [20]. These are failures of modelling, not of human decision-making and, as indicated above, we believe an effective risk counselling system and the underlying LC will need to support people's actual behaviour, not impose a normative model upon them.

Decision-sequence-enabled: An effective risk counselling system will act throughout the period when the patient is at risk of cancer, and through any monitoring and treatment periods, and subsequent to these. This means the system needs to be able to support chains of decisions, i.e. sequences of separate decisions made at different time points. The consequence tuples of such decisions will depend upon choices made at earlier times, and it is therefore natural to represent such chains as branches in a tree structure. The LC underlying the risk counsel system needs to be able to support such a structure, and will need to include natural mechanisms for comparing one branch (i.e. one decision sequence) with another.

Debate-enabled: Because of the criticality of the decisions involved, many patients will discuss their decision options with family members and others, as part of their decision-process. The opinions of these others may well differ from that of the patient, and the patient and her risk advisor may wish to compare these. Elsewhere [18], two of us have proposed the use of a form of dialectical argumentation in intelligent systems to model and manipulate competing views of claims of chemical carcinogenicity. We believe such an approach would be useful here, and so the risk counselling system and its underlying LC need to be capable of representing and manipulating such debates.

A related issue is that a patient's family, although they may be party to her decision, are not all likely to be present at her risk counselling sessions with her GP. One advantage of argumentation-based systems is their ability to be transparent about the reasons for a suggested conclusion, even through a long chain of reasoning. The justification for, and the reasoning-path of, a suggested conclusion can be presented along with the conclusion. If such transparency can be represented in printed form, then this has obvious advantages for decision support where the decision-maker needs to discuss the decision with others outside the session or where the decisions are jointly made.

Friendly User Interfaces: Neither GPs nor their patients are typically expert computer users. As the user interface will be their sole form of interaction with the system, the design of the interface will be crucial to the success of the system for risk counselling. The user interfaces, therefore, need to be intuitive, robust and friendly. As with any new technology, they will need to be pilot-tested with risk advisors and their patients before final release. We have already undertaken such testing for the cancer risk prediction system component of RAGs [4]. In another domain, Verheij [28], designing intelligent systems for legal applications, has argued that new kinds of user interfaces are needed for argumentation systems.

Updatable and connected: Finally, our vision for the risk counselling system is that it would be used by a GP over the course of a patient's life, being updated with new information as it comes to hand (such as future test results and lifestyle changes), along with updates to the patient's preferences and valuations as these change. For this reason, the system will need to be readily linkable to other patient information management systems used by the GP and relevant health authorities. It is even possible to imagine such a system being web-accessible[2], so that the patient may update personal and preference information, and seek automated risk counsel, outside the GP's surgery.

This list of requirements effectively marks a series of high-level philosophical positions within which the risk counselling system and the underlying Logic of Consequences can be defined. Having defined this high-level framework, we are currently engaged in the task of developing the detailed LC.

4 Conclusions and Further Work

The ICRF is engaged in the development of theoretically-grounded intelligent decision support systems to advise General Practitioners and their patients about their individual risks of breast cancer and their options arising therefrom. These options include actions of monitoring, prophylaxis and treatment, as well as lifestyle choices available to the patient. Because much of the information available to populate such a system is qualitative in nature, and because, even when quantitative information is available, it is often not well understood by ordinary people, we believe strongly that a qualitative formalism is appropriate for representing and manipulating the information in such advice systems. We have proposed the use of a Logic of Consequences (LC) for this task, akin to Logics of Argumentation, but with the consequences of actions taking the place of the grounds for claims. We are currently attempting to define a theoretically-sound LC for the breast cancer advice system, and this paper has identified at

[2]In much the same way that the demonstration version of the Capsule system developed at the ICRF is (see `http://www.infermed.com`).

a high-level the key requirements of such a prognostic formalism.

Several of these requirements may be contentious. Some may argue that expert systems need to adopt normative approaches to decision-making and that, even in conditions of great ignorance, quantitative uncertainty formalisms (such as Bayesian techniques) are appropriate [11]. We disagree strongly with such views. The application we have in mind must assist ordinary people in making informed choices about their health-care, treatment and lifestyle options in a context of potentially great anxiety. The people making the decisions must comprehend, and be comfortable with, the entire process of decision making and also the outcomes of that process. This means the risk counselling system must be consonant with, and support, the way patients actually make decisions of this sort. After all, it is their lives which are at stake.

Acknowledgments

Project RAGs is funded by the UK Social and Economic Research Council (SERC) and the ICRF, and we are grateful for this support. One author (PM) is supported by the UK Engineering and Physical Sciences Research Council (EPSRC), and we also gratefully acknowledge this assistance.

References

[1] D. A. Aaker, V. Kumar, and G. S. Day. *Marketing Research.* John Wiley and Sons, New York City, NY, USA, 1995.

[2] J. E. Brody. Cancer gene tests turn out to be far from simple. *New York Times (New York City, NY, USA)*, 17 August 1999.

[3] A. S. Coulson, D. W. Glasspool, J. Fox, and J. Emery. Computerized genetic risk assessment and decision support in primary care. *Informatics: The Journal of Informatics in Primary Care*, 1999 *(In press)*.

[4] J. Emery, R. Walton, A. Coulson, D. Glasspool, S. Ziebland, and J. Fox. Computer support for recording and interpreting family histories of breast and ovarian cancer in primary care (RAGs): qualitative evaluation with simulated patients. *British Medical Journal*, 319:32–36, 1999.

[5] C. H. Fellner and J. R. Marshall. Kidney donors — the myth of informed consent. *American Journal of Psychiatry*, 126:1245, 1970.

[6] J. Fox and D. W. Glasspool. Assessing and communicating personal risk: a cognitive engineering approach. Technical Report 353, Advanced Computation Laboratory, Imperical Cancer Research Fund, London, UK, 1999.

[7] J. Fox, P. Krause, and S. Ambler. Arguments, contradictions and practical reasoning. *Proceedings of the European Conference on Artificial Intelligence 1992, Vienna, Austria*, 1992.

[8] J. Fox and S. Parsons. Arguing about beliefs and actions. In A. Hunter and S. Parsons, editors, *Applications of Uncertainty Formalisms*, pages 266–302. Springer Verlag (Lecture Notes in Artificial Intelligence 1455), Berlin, Germany, 1998.

[9] J. Green. *The Social Construction of the XYY Syndrome*. PhD thesis, University of Manchester, Manchester, UK, 1983.

[10] P. E. Green and A. M. Krieger. Conjoint analysis with product-positioning applications. In J. Eliashberg and G. L. Lilien, editors, *Handbooks in Operations Research and Management Science, Volume 5: Marketing*, pages 467–516. North-Holland, Amsterdam, The Netherlands, 1993.

[11] K. S. Van Horn. If it ain't Bayesian, it's *broken*: an examination of Cox's Theorem. Unpublished manuscript, 1999.

[12] I. L. Janis. The patient as decision-maker. In W. D. Gentry, editor, *Handbook of Behavioral Medicine*, page 333. Guilford, 1984.

[13] J. Katz. *The Silent World of Doctor and Patient*. The Free Press, New York City, NY, USA, 1984.

[14] P. Krause, J. Fox, and P. Judson. An argumentation based approach to risk assessment. *IMA Journal of Mathematics Applied in Business and Industry*, 5:249–263, 1994.

[15] K. J. Lancaster. A new approach to consumer theory. *Journal of Political Economy*, 74:132–157, 1966.

[16] G. L. Lilien, P. Kotler, and K. S. Moorthy. *Marketing Models*. Prentice-Hall, Englewood Cliffs, NJ, USA, 1992.

[17] D. V. Lindley. *Making Decisions*. John Wiley and Sons, London, UK, second edition, 1985.

[18] P. McBurney and S. Parsons. Truth or consequences: using argumentation to reason about risk. In *BPS Workshop on Practical Reasoning, 21 December*, London, UK, 1999.

[19] R. Nisbett and L. Ross. *Human Inference: Strategies and Shortcomings of Social Judgment*. Prentice-Hall, Englewood Cliffs, NJ, USA, 1980.

[20] T. O'Donoghue and M. Rabin. Doing it now or later. *American Economic Review*, 89 (1):103–124, 1999.

[21] S. Parsons and J. Fox. Argumentation and decision-making: a position paper. In *Formal and Applied Practical Reasoning*, pages 705–709. Springer Verlag, Berlin, Germany, 1996.

[22] S. Parsons, J. Fox, and A. Coulson. Argumentation and risk assessment. In *AAAI Spring Symposium on Predictive Toxicology of Chemicals: Experiences and Impact of AI Tools*, pages 60–63, Stanford, CA, 1999.

[23] P. F. Pierce. Deciding on breast cancer treatment: a description of decision behavior. *Nursing Research*, 42:20–23, 1993.

[24] H. Raiffa. *Decision Analysis: Introductory Lectures on Choices Under Uncertainty.* McGraw-Hill, New York, 1968.

[25] C. E. Schneider. *The Practice of Autonomy: Patients, Doctors and Medical Decisions.* Oxford University Press, Oxford, UK, 1998.

[26] R. G. Simmons, S. K. Marine, and R. L. Simmons. *Gift of Life: The Effect of Organ Transplantation on Individual, Family, and Societal Dynamics.* Transaction Books, 1987.

[27] L. C. van der Gaag, S. Renooij, C. L. M. Witteman, B. M. P. Aleman, and B. G. Taal. How to elicit many probabilities. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence 1999*, 1999.

[28] B. Verheij. Automated argument assistance for lawyers. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Law, Oslo, Norway*, pages 43–52, New York City, NY, USA, 1999. ACM.