

Argumentation and risk assessment

Simon Parsons^{†,‡}

[†]Department of Electronic Engineering
Queen Mary and Westfield College
Mile End Road
London E1 4NS
United Kingdom

John Fox[‡] and Andrew Coulson[‡]

[‡]Advanced Computation Laboratory
Imperial Cancer Research Fund
Lincoln's Inn Fields
London WC2 3PX
United Kingdom

Abstract

Over the last ten years we have been involved in the development of a formal framework for decision making and reasoning under uncertainty based on “argumentation”. The latter provides a way of managing uncertainty which differs from probabilistic inference and is particularly valuable in those many practical situations where uncertainty cannot be quantified. Recently we have been applying argumentation as a non-numerical method of risk assessment.

Introduction

Standard decision theory builds on the probabilistic view of uncertainty in reasoning about actions. The costs and benefits of possible outcomes of actions are weighted with their probabilities, yielding a preference ordering on the “expected utility” of alternative actions. However, as many authors have pointed out, the specification of the complete sets of probabilities and utilities required by standard decision theory make the theory impractical in complex tasks which involve common sense knowledge. This realisation has prompted work on qualitative approaches to decision making which attempt to reduce the amount of numerical information required.

We have been involved in work on such qualitative decision making techniques since the early 80s (see (Parsons & Fox 1996) for a review). Our early work was partly concerned with the description of human decision processes (Fox 1980) and partly with the practical development of decision systems for use in medicine (Fox, Barber, & Bardhan 1980). Whilst the qualitative decision procedures we developed proved to have considerable descriptive value and practical promise, our desire to build decision support systems for safety-critical fields such as medicine raised the concern that our early applications were *ad hoc*. In particular we were concerned that they, in common with all other expert systems being built at the time, were not based on a rigorously defined decision theory. As a result we have put considerable effort into developing a theoretical framework for qualitative decision making. The best developed part of this is an approach to uncertainty and belief based on the idea of *argumentation*.

The next section gives a short informal description of argumentation. More extensive accounts may be found in (Fox, Krause, & Ambler 1992; Krause *et al.* 1995; Fox & Parsons 1998).

Argumentation

In a classical logic L , an argument is a sequence of inferences leading to a conclusion. If the argument is correct, then the conclusion is true. An argument:

$$G_1 \dots G_n \vdash St$$

is correct in the logic L if St may be derived using the rules of inference and axioms of L augmented with $G_1 \dots G_n$. Therefore a correct argument simply yields a proposition St . This can be paraphrased as

$$St \text{ is true (in the context } G_1 \dots G_n)$$

In the approach we take, this traditional form of logic based argumentation is extended in two important ways:

1. to allow arguments not only to prove propositions but also to merely indicate support for, or even doubt in, them; and
2. to explicitly record the context in which the deduction holds.

The way we do this is to attach two things to each proposition which is derived—a record of the information used to derive it, and a measure of confidence in the derivation. Thus the result of a derivation is an *argument* of the form:

$$(St : G : Sg)$$

Each argument consists of a triple consisting of a Sentence (St), Grounds (G), which are the formulae used to justify the argument, and a Sign (Sg), which is a number or a symbol which indicates the confidence warranted in the conclusion. The idea of argumentation from a database may thus be summarised by the following schema:

$$\text{Database } \vdash_{ACR} (\text{Sentence} : \text{Grounds} : \text{Sign})$$

In this schema, \vdash_{ACR} is a consequence relation which defines the inference rules by which we may construct

arguments for claims using the information in the database.

The use of confidences rather than logical proofs introduces a slight complication. In classical logic, if we can construct an argument (proof) for *St* then any further arguments for *St* are of no interest since *St* is known to be true. If, however, we only have an indication of support for *St* then it may be the case that additional information casts doubt on *St*. Thus we need to consider every distinct argument concerning *St* and then carry out a process of *aggregation* to combine them. This process is also known as *flattening* since it has the effect of mapping a number of distinct arguments into a single measure.

Argumentation appears to have considerable practical potential. Fox *et al.* (1990) first described the use of argumentation in decision support systems and Das *et al.* (1996) have shown how it can be incorporated in a sound decision making procedure. Argumentation has also provided the basis of an executable agent specification language, PROforma (Fox *et al.* 1997).

Argumentation and risk

Because it can be applied in the absence of detailed numerical estimates of uncertainty, argumentation seems to be a useful way of assessing risks. The StAR project developed software for identifying the risk of carcinogenicity associated with chemical compounds. (Fox 1997; Krause, Judson, & Patel 1998). In this domain environmental and epidemiological impact statistics are often unavailable, so argumentation provides an alternative method for reasoning about risks from general scientific knowledge. The approach is to build arguments, based on whatever information is available, for or against the carcinogenicity of the chemical in question, and to use the interaction between these arguments to estimate the gravity of the risk. Thus if there is one argument that a chemical might be carcinogenic (because it contains some functional group which is known to cause cancer in rats) then there is a risk that the chemical might cause cancer in humans. However, if there is a second argument which defeats the first (by, for instance, pointing out that the cancer-causing mechanism in rats involves an enzyme which is not present in humans) then the risk is considered to be lower. An HMSO report on micro-biological risk assessment identifies StAR as a major new approach to this important problem (Health and Safety Commission 1996).

The demonstrator system produced by the StAR project is a prototype for a computer based assistant for the prediction of the potential carcinogenic risk due to novel chemical compounds. A notion of hazard identification is taken as a preliminary stage in the assessment of risk. The hazard identification used here draws heavily on the approach taken in the expert system DEREK, which is used for the qualitative prediction of possible toxic action of chemical compounds (Sanderson & Earnshaw 1991). DEREK is able to detect chemical sub-structures within molecules, known as structural




Figure 1: The StAR Demonstrator: Example 1

alerts, and relate these to a rule-base linking them with likely types of toxicity. In the demonstration, the structural alerts have been taken from a U.S. FDA report identifying sub-structures associated with various forms of carcinogenic activity (Federal Drug Administration 1986).

The user of the carcinogenicity risk adviser presents the system with the chemical structure of the compound to be assessed, together with any additional information which may be thought relevant (such as possible exposure routes, or species of animal that will be exposed to the chemical). The chemical structure may be presented using a graphical interface. The database of structural alerts is then searched for matches against the entered structure. If a match is found, a theorem prover tries to construct arguments for or against the hazard being manifest in the context under consideration. Having constructed all the relevant arguments, a report is generated on the basis of the available evidence¹.

For the first screen (Figure 1), the user has entered a relatively simple structure based on an aromatic ring. The system has identified that it contains an alert for epoxides (the triangular structure to the top right). Whilst constructing arguments, the system has recognised that the LogP value is relevant in this case, and so queries the user for this information (loosely, the value of LogP gives a measure of how easily the substance will be absorbed into tissue). The functional group for epoxides is indicative of a direct acting carcinogen, and the value of LogP supplied by the user is supportive of the substance being readily absorbed into tissue. Hazard recognition plus supportive evidence, with no arguments countering potential carcinogenic activity, yields the classification of a "probable human carcinogen" (the result might be different for different animals). Figure 1 shows the summary report. The query box is illustrated

¹For ease of presentation, the examples use a simplified database, and some of the following assessments may be chemically or biologically naive.




Figure 2: The StAR Demonstrator: Example 2




Figure 3: The StAR Demonstrator: Example 3

in this screen image, although it would normally have been closed by this stage.

The second example (Figure 2) involves a structure which contains an alert for peroxisome proliferators. The top-most screen contains a simple non-judgemental statement to this effect. The lower screen contains the summary of the argumentation stage of analysis. Here, evidence is equivocal—as explained in Figure 3, there are arguments for and against carcinogenicity so no overall conclusion can be reached.

Genetic risk

We are currently applying similar technology to the problem of estimating genetic risk. In particular, we




Figure 4: The RAGs Demonstrator: Family Tree

have developed a PROforma application called RAGs (Risk Assessment in Genetics), which is designed to assist a family doctor (GP) in evaluating the genetic risk of breast cancer for a given patient. The interface to the system is pictured in Figures 4 and 5. The GP enters the relevant personal details of a patient and her family, gradually building a graphical “family tree” on the screen. Once such a pedigree has been created, RAGs assesses the person’s genetic risk of breast cancer, based on the known incidence of cancer in her family. An appropriate on-screen report is then generated.

The risk calculation in RAGs is performed using argumentation. Domain knowledge was provided by Jon Emery at the Imperial Cancer Research Fund’s General Practice Research Group who generated a set of rules for building arguments about a person’s increased or decreased genetic risk. For example, one rule is that a first degree relative with breast cancer increases the presenting patient’s genetic risk. If appropriate these rules generate arguments about a patient, and a simple total risk score is computed. Based on this score, the patient is put into a high, medium or low genetic risk category, and the appropriate referral advice is given.

Figure 4 shows the family tree for an imaginary case, where Karen is the patient whose risk is being assessed. The heavy line shows the highest risk path of inheritance, and Figure 5 shows the advice generated by the system in response to this case, along with the arguments which underlie it. The fact that these arguments can be used as explanations for the GP and patient is perhaps the most interesting thing about the demonstrator. These explanations can be related directly to the graphical image of the family tree, thus providing an exceptionally clear account of what the system is doing.

The results given by the RAGs program have been

The following information applies ONLY to the highlighted path.

This patient is at moderate risk of being a gene carrier because, on the highest-risk path of inheritance found by the program:

- * The mother of the presenting patient is affected, which indicates an increased risk level.
- * One first-degree relative (FDR) is affected (Each affected FDR indicates an additional risk factor).
- * More than one second-degree relative is affected, indicating a moderate increase in risk.
- * The combination of more than one breast and one ovarian cancer indicates a considerable increase in risk level.

However, this is balanced to some extent by the following factors which indicate lower risk level:

- * The oldest affected second-degree relative has an age of onset between 50 and 60. Genetic predisposition is more likely to be associated with lower ages of onset, and this age indicates a moderate reduction in risk level.
- * Genetic predisposition is less likely in a person over 40 who has not developed cancer.

General Explanation Referral Advice Reasons for Advice

Figure 5: The RAGs Demonstrator: Advice

compared with those given by a commercial pedigree drawing program, Cyrillic, which calculates genetic risk as a numerical probability. Cyrillic represents the state of the art for probabilistic genetic risk calculations. The comparison showed that RAGs categorizes patients into the three genetic risk categories in agreement with Cyrillic.

Summary

This paper has briefly introduced our approach to argumentation, a formal mechanism for handling uncertainty in a largely symbolic way, and has discussed its use in risk assessment. In particular, it has discussed the use of argumentation in two projects at the Imperial Cancer Research Fund. In the first, argumentation is used to predict the risk of carcinogenicity of novel chemical compounds. In the second, it is used to assess the risk of genetic disposition to breast cancer. We are also in the process of setting up a project which will further develop argumentation as an approach to risk assessment, by using it as a mechanism for exploring possible (as opposed to probable) risks, and as a basis for risk mitigation and management.

Acknowledgments

The authors would like to thank all their colleagues who have contributed to the work described here, including Jon Emery, David Glasspool, Phillip Judson, Paul Krause and Nicola Pappas.

References

- Das, S. K.; Fox, J.; and Krause, P. J. 1996. A unified framework for hypothetical and practical reasoning (1): theoretical foundations. In *Formal and Applied Practical Reasoning*, 58–72. Berlin, Germany: Springer Verlag.
- Federal Drug Administration. 1986. *General principles for evaluating the safety of compounds used in food-producing animals: Appendix 1. Carcinogen structure Guide*. FDA.
- Fox, J., and Parsons, S. 1998. Arguments about beliefs and actions. In Hunter, A., and Parsons, S., eds., *Applications of Uncertainty Formalisms*. Berlin: Springer Verlag.
- Fox, J.; Barber, D.; and Bardhan, K. D. 1980. Alternatives to Bayes? A quantitative comparison with rule-based diagnostic inference. *Methods of Information in Medicine* 19:210–215.
- Fox, J.; Clark, D. A.; Glowinski, A.; and O'Neil, M. 1990. Using predicate logic to integrate qualitative reasoning and classical decision theory. *IEEE Transactions on Systems, Man and Cybernetics* 20:347–357.
- Fox, J.; Johns, N.; Lyons, C.; Rahmzadeh, A.; Thomson, R.; and Wilson, P. 1997. Proforma: a general technology for clinical decision support systems. *Computer Methods and Programs in Biomedicine* 54:59–67.
- Fox, J.; Krause, P.; and Ambler, S. 1992. Arguments, contradictions and practical reasoning. In *Proceedings of the 10th European Conference on Artificial Intelligence*, 623–627. Chichester, UK: John Wiley & Sons.
- Fox, J. 1980. Making decisions under the influence of memory. *Psychological Review* 87:190–211.
- Fox, J. 1997. Will it happen? can it happen? *Science and Public Affairs* Winter '97:45–48.
- Health and Safety Commission. 1996. *Advisory Committee on Dangerous Pathogens (UK), Microbiological risk assessment, Interim report*. HMSO.
- Krause, P.; Ambler, S.; Elvang-Gøransson, M.; and Fox, J. 1995. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence* 11:113–131.
- Krause, P.; Judson, P.; and Patel, M. 1998. Qualitative risk assessment fulfills a need. In Hunter, A., and Parsons, S., eds., *Applications of Uncertainty Formalisms*. Berlin: Springer Verlag.
- Parsons, S., and Fox, J. 1996. Argumentation and decision making: a position paper. In *Formal and Applied Practical Reasoning*, 705–709. Berlin, Germany: Springer Verlag.
- Sanderson, D. M., and Earnshaw, C. G. 1991. Computer prediction of possible toxic action from chemical structure; the DEREK system. *Human & Experimental Toxicology* 10:261–273.