Argumentation and decision making*

Simon Parsons[1][2] and John Fox[2]

[1] Department of Electronic Engineering,
Queen Mary and Westfield College,
University of London,
London E1 4NS,
United Kingdom.
[2] Advanced Computation Laboratory,
Imperial Cancer Research Fund,
P.O. Box 123,
Lincoln's Inn Fields,
London WC2A 3PX,
United Kingdom.

Abstract. This paper summarises our position on the use of symbolic methods for reasoning under uncertainty, and argumentation in particular. Our view is that argumentation offers a complement to numerical methods for reasoning about belief, and a general framework within which many competing approaches can be understood. In applications we have found that argumentation offers a variety of benefits for practical reasoning systems. The presentation is historical, emphasising the reasons which motivated the development of the argumentation framework, drawing primarily on work carried out by researchers at the Imperial Cancer Research Fund over the last fifteen years.

1 The need for symbolic decision making

Work on argumentation at the Imperial Cancer Research Fund arose out of a series of studies on medical decision making, including work on modelling human diagnostic reasoning [12] and comparisons of the relative merits of numerical and symbolic inference techniques in clinical decision making [14, 21, 28]. These studies strongly suggested, contrary to the assumption prevalent at the time, that numerical methods were not the only sound and practical means of making decisions under uncertainty.

The first line of work was in an empirical tradition, concerned with how people make decisions and the strengths and weaknesses of the decision making process. One study investigated diagnostic decision making by medical students and compared two computational models of decision making on a simulated medical diagnosis problem [13]. The models were implemented as sets of production rules, one implementing a statistical model and one a knowledge-based, semi-qualitative model. The results strongly suggested that the latter model gave a

* This is a revised and extended version of a paper which first appeared in the Proceedings of the First Conference on Formal and Applied Practical Reasoning.

better account of human reasoning under uncertainty. This is consistent with a well established finding from psychological research that people do not manage uncertainty in ways which closely resemble normative probabilistic reasoning.

Classical decision theorists have strongly criticised human judgement precisely on the grounds that people do not comply with the requirements of probability theory. In contrast some cognitive scientists have questioned the force of this observation by emphasising the flexibility and other virtues of human reasoning (e.g. [37]). This observation was reinforced by a study which investigated the performance of an expert system based on a model of human reasoning [14]. In a realistic diagnostic problem in gastroenterology it was found that the diagnostic accuracy of the expert system was very similar to that of a Bayesian diagnostic system though the expert system achieved this level of performance using only half the information provided to the probabilistic system. A recent paper has revisited this issue [5].

Other empirical studies suggest that, notwithstanding the quantitative precision of probabilistic evidence analysis, much practical medical decision making can be successfully carried out without precise numerical data. For example [21], a rule-based system for leukaemia diagnosis was developed using the EMYCIN expert system shell. The accuracy of the system, as compared with the decisions of a domain expert, was 64%. When the CFs were limited to just two values, 1.0 or 0.5 (loosely "certain" or "uncertain") a 5% *increase* in correct diagnoses was observed! In another study, decision making using a semiqualitative decision procedure was compared with a probabilistic procedure using data regarding the admission of 140 patients to a coronary care unit. It was found that the symbolic decision model performed at least as well as the probabilistic model as determined by an ROC analysis [28]. Other independent studies have confirmed the finding that practical decision making can be carried out successfully without depending upon precise quantitative data (e.g. [4, 35]).

2 A model for reasoning under uncertainty

Having established these results, the motivation for further work was an interest in building decision support systems which would have a number of advantages over conventionally available technologies:

- They would not depend upon the availability of objective statistical data (which are frequently not available in complex domains like medicine)
- They could make use of other kinds of reasoning than statistical inference (e.g. causal, functional and temporal reasoning) which might be more intuitive than quantitative reasoning.
- They could support all phases of decision making, not just evidence analysis, such as recognising when a decision is needed, what the decision options are, what information is relevant to the choice and so forth.

An initial framework for "symbolic decision theory", using first-order logic for much of the deductive reasoning required and argumentation for the management of uncertainty, was proposed [15, 18], and early evidence for the practicality of the theory came from its use in the Oxford System of Medicine (OSM) [17], a decision support system aimed at general practitioners which provides a generic decision procedure for a range of medical decision tasks. Further evidence for the usefulness of the theory comes from the wide range of possible applications that have now been developed [11] and in particular the program CAPSULE. This latter is a system which gives support to general practitioners in prescribing drugs and which has been shown to be of significant value in this role [39].

In addition, however, it was agreed, in accordance with the view generally held by decision theorists, that any decision procedure which is to be used for practical applications, particularly those like medical applications which have safety implications, must be given a sound theoretical underpinning. While the symbolic approach might be "inspired" by observations of human flexibility, people make mistakes and it is clearly not desirable to emulate those mistakes!

3 Formalising the model

The most contentious element of symbolic decision theory is the use of argumentation as the basic framework for reasoning under uncertainty. The central idea in argumentation about beliefs (which is the area of argumentation which has been formalised so far) is that an argument is a *tentative proof* of a proposition; thus an argument is a proof which can fail if suitably strong arguments against the proposition can be found. The fact that a proposition can have arguments for and against it suggests a divergence between argumentation and classical first order logic in which propositions are true or false. Furthermore, as pointed out in [2], argumentation has many commonalities with intuitionistic logic, suggesting that argumentation might be given a sound basis in category theory since this is possible for intuitionistic logic. The first steps in providing this basis are detailed in [2] which identifies the structure of the space of arguments, along with the kind of operations possible over them. The rest of the formalisation is provided in [1], which also highlights the link between argumentation and Dempster-Shafer theory [36].

With this semantics in mind, it is possible to define a logic LA in which the consequences of a database are a set of arguments, and this is the subject of [3] and [25]. In this work logical formulae are augmented with their proofs and when formulae are combined, the proofs are handled in an appropriate manner. This means that it is possible to determine the validity of formulae derived in the logic based upon the strength of the arguments for and against individual formulae, and that the way in which this is done is in accordance with the category-theoretic semantics. This process has been automated in the Argumentation Theorem Prover [26], and a summary of the formal model is recorded in [24].

4 Alternative semantics

The proof theoretic semantics given in [3] and [25] are not the only possible semantics for argumentation. Indeed, two further interpretations have been provided. The first [7] is an extension of the standard Kripke semantics for modal logic which gives a possible worlds interpretation for what it means for an argument to support a proposition to some degree. The second [31, 32] relates certain types of argumentation to probability theory by taking an argument in favour of a proposition to mean that there is evidence that the probability of the proposition increases (so the proposition becomes more likely to be true). Thus argumentation can capture probabilistic reasoning if required, and so it is possible to claim that, under particular conditions, argumentation is a normative theory for handling uncertainty.

The probabilistic semantics have two further advantages. The first is that, because they tie the notion of an argument securely to well-understood ideas about qualitative probability, it is possible to harness a number of useful results concerning qualitative probability. In particular, it is possible to develop a finer-grained representation of what it means to have an argument for a proposition which allows arguments of different strengths to be accommodated [30, 32, 33]. The second additional advantage of the probabilistic semantics is that it suggests that decision making based upon argumentation can also be understood in terms of the classical decision making paradigm. Thus it offers ways of formalising the concept of arguments about actions and values in addition to the existing formalisation of arguments about beliefs. This last point is explored in [23].

5 A general model of reasoning

In addition to the arguments for argumentation as a symbolic model of decision making, we can argue [20] that it is a model of "practical reasoning" of the kind that humans indulge in every day. It captures many of the modes of commonsense reasoning—finding support for ideas, attacking other ideas, and trying to attack the support of other ideas. It handles contradictions, and should also enable the resolution of conflicting arguments at the meta-level. Furthermore, there is a strong case that argumentation provides a general framework for unifying many methods for reasoning under uncertainty, such as possibility and probability [34] theories. In this role argumentation is less a formalisation of human reasoning than a tool that can enable the use of other formalisms. Argumentation provides a general way of combining logical reasoning with Bayesian probability by using it to construct a network of influences between relevant variables. Indeed, argumentation is sufficiently general as to underlie symbolic as well as quantitative formalisms [27].

However, it is possible to do more with argumentation than just provide a framework for using established formalisms, instead, as is discussed in [9], it is possible to handle inconsistent information—a problem beyond the scope of many established approaches to handling uncertainty. That is, it is possible to have certain arguments for both a proposition p and its negation $\neg p$. This inconsistency enables LA to provide a ranking over the propositions for which arguments may be proposed. In particular, arguments for propositions are allocated different classes of acceptability [8, 10], the allocation depending on factors such as whether an argument is based on a consistent database or whether there are any counter-arguments. This approach can be used to give a purely logical approach to uncertainty that ranks propositions only on the structure of the arguments for them, and it can be augmented by the use of preference relations over subsets of the database.

6 Argumentation and autonomous agents

The flexibility and generality of argumentation has led to it being used at the heart of a model of autonomous agency [6, 16], a model which extends the classical decision making paradigm to include reasoning about which plans to adopt as well as what to believe. This, in turn, has resulted in an interest in handling medical treatment protocols and the development of the tool PRO*forma* [19] for building decision support systems to help in following such protocols.

In all of the variants of argumentation that have been discussed so far, the argumentation mechanism has been used within a single agent. This agent, then, argues internally about which course of action is the best way of resolving its predicament. However, argumentation is also possible in scenarios involving a number of agents—indeed it is quite natural to think of an argument as involving two or more participants. Applying argumentation in a multi-agent scenario gives the agents involved a mechanism for resolving their differences and reaching suitable compromises, and this is especially important where the agents are completely autonomous and so have no particular reason to work together. The use of argumentation for negotiation is explored in [29], and a minimal framework to support such argumentation is defined in [38].

7 Summary

To summarise, empirical evidence for the usefulness of a symbolic theory of decision making has led to the development of a formal model based on first-order logic combined with the use of argumentation for handling uncertainty. The versatility of the model is suggested by its wide practical applicability, while the justification for the use of argumentation is based upon its proven practical flexibility and its well-developed formal semantics. The model now forms the basis of a general technology for decision support systems.

References

1. Ambler, S. (1992) A categorical approach to the semantics of argumentation, Technical Report 606, Department of Computer Science, Queen Mary and Westfield College.
2. Ambler, S. and Krause, P. (1992) Enriched categories in the semantics of evidential reasoning, Technical Report 153 Advanced Computation Laboratory, Imperial Cancer Research Fund.

3. Ambler, S. and Krause, P. (1992) The development of a "Logic of Argumentation", *Proceedings of the International Conference on Information Processing and the Management of Uncertainty*, Palma.
4. Chard, T. (1991) Qualitative probability versus quantitative probability in clinical diagnosis: a study using a computer simulation, *Medical Decision Making*, **11**, 38–41.
5. Cooper, R. and Fox, J. (1997) Learning to make decisions under uncertainty: the contribution of qualitative reasoning, *Proceedings of the 19th Annual Conference of the Cognitive Science Society.*
6. Das, S., Fox, J., Elsdon, D. and Hammond, P. (1997) Decision making and plan management by autonomous agents: theory, implementation and applications, *Proceedings of Autonomous Agents '97*, Marina del Rey, CA.
7. Das, S., Fox, J. and Krause, P. (1996) A unified framework for hypothetical and practical reasoning (1): theoretical foundations, *Proceedings of the International Conference on Formal and Applied Practical Reasoning*, Bonn.
8. Elvang-Gøransson, M. and Hunter, A. (1993) Argumentative logics—reasoning with classically inconsistent information, *Data and Knowledge Engineering*, **16**, 125–145.
9. Elvang-Gøransson, M., Krause, P. and Fox, J. (1993) A logical approach to handling uncertainty, *Proceedings of the Workshop on Modelling Problems in Control and Supervision of Complex Dynamic Systems*, Lyngby, Denmark.
10. Elvang-Gøransson, M., Krause, P. and Fox, J. (1993) Dialectic reasoning with inconsistent information, *Proceedings of 9th Conference on Uncertainty in Artificial Intelligence*, Washington, D. C.
11. Fox, J. and Das S.(1996) A unified framework for hypothetical and practical reasoning (2): lessons from clinical medicine, *Proceedings of the International Conference on Formal and Applied Practical Reasoning*, Bonn.
12. Fox, J. (1987) Making decisions under the influence of knowledge, in *Modelling Cognition*, P. Morris ed., John Wiley & Sons Ltd.
13. Fox, J. (1980) Making decisions under the influence of memory, *Psychological Review*, **87**, 2, 190–211.
14. 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.
15. Fox, J., Clark, D. A., Glowinski, A. Gordon, C. and O'Neil, M. J. (1990) Using predicate logic to integrate qualitative reasoning and classical decision theory, *IEEE Transactions on Systems, Man and Cybernetics*, **20**, 347–357.
16. Fox, J., Das, S. Elsdon, D. and Hammond, P. (1995) Decision making and planning by autonomous agents: an architecture for safety critical applications, *Proceedings of Expert Systems '95*, Cambridge.
17. Fox, J., Glowinski, A. Gordon, C. Hajnal, S. and O'Neil, M. (1990) Logic engineering for knowledge engineering: design and implementation of the Oxford System of Medicine, *Artificial Intelligence in Medicine*, **2**, 323–339.
18. Fox, J., Glowinski, A. J., O'Neil, M. J. and Clark, D. A. (1988) Decision making as a logical process, *Proceedings of Expert Systems '88*, Cambridge.
19. Fox, J., Johns, N., Lyons, C., Rahmanzadeh, A., Thomson, R., and Wilson, P. (1997) PRO*forma*: a general technology for clinical decision support systems, *Computer Methods and Programs in Biomedicine*, to appear.
20. Fox, J., Krause, P. and Ambler, S. (1992) Arguments, contradictions and practical reasoning, *Proceedings of the European Conference on Artificial Intelligence*,

Vienna.

21. Fox, J., Myers, C. D., Greaves, M. F., and Pegram, S. (1985) Knowledge acquisition for expert systems: experience in leukemia diagnosis, *Methods of Information in Medicine*, **24**, 65–72.
22. Fox, J., O'Neil, M., Glowinski, A. J. and Clark, D. (1988) A logic of decision making, *Illinois Interdisciplinary Workshop on Decision Making*, Urbana, Illinois.
23. Fox, J. and Parsons, S. (1997) On using arguments for reasoning about actions and values, *AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, Stanford.
24. Krause, P., Ambler, S., Elvang-Gørannson, M. and Fox, J. (1994) A logic of argumentation for reasoning under uncertainty, *Computational Intelligence*, **11**, 113–131.
25. Krause, P., Ambler, S. and Fox, J. (1993) The development of a "Logic of Argumentation", in *Advanced Methods in Artificial Intelligence*, B. Bouchon-Meunier, L. Valverde and R. R. Yager eds., Springer-Verlag, Berlin.
26. Krause, P., Ambler, S. and Fox, J. (1993) ATP user manual, Technical Report 187, Advanced Computation Laboratory, Imperial Cancer Research Fund.
27. Krause, P. J., Fox, J. and Ambler S. (1992) Argumentation as a unifying concept for reasoning under uncertainty, Technical Report 166, Advanced Computation Laboratory, Imperial Cancer Research Fund.
28. O'Neil, M. and Glowinski, A. (1990) Evaluating and validating very large knowledge-based systems, *Medical Informatics*, **3**, 237–251.
29. Parsons, S., Sierra, C. and Jennings, N. R. (1998) Agents which reason and negotiate by arguing, *Journal of Logic and Computation*, to appear.
30. Parsons, S. (1997) On order of magnitude reasoning and qualitative probability, Proceedings of the Florida AI Research Symposium, Daytona Beach, FL.
31. Parsons, S. (1997) Normative argumentation and qualitative probability, *Proceedings of the International Joint Conference on Practical Reasoning*, Bonn.
32. Parsons, S. (1996) Defining normative systems for qualitative argumentation, *Proceedings of the International Conference on Formal and Applied Practical Reasoning*, Bonn.
33. Parsons, S. (1995) Refining reasoning in qualitative probabilistic networks, Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence, Montreal.
34. Parsons, S. and Fox, J. (1994) A general approach to managing imperfect information in deductive databases, *Proceedings of the Workshop on Uncertainty in Databases and Deductive Systems*, Ithaca, NY.
35. Pradhan, M., Henrion, M., Provan, G., del Favero, B., and Huang, K. (1996) The sensitivity of belief networks to imprecise probabilities: an experimental investigation, *Artificial Intelligence*, **85**, 363–397.
36. Shafer, G. (1976) *A mathematical theory of evidence*, Princeton University Press.
37. Shanteau, J. (1987) Psychological characteristics of expert decision makers, in *Expert Judgement and Expert Systems*, J. Mumpower ed., NATO ASI Series, vol F35.
38. Sierra, C., Jennings, N. R., Noriega, P. and Parsons, S. (1997) A framework for argument-based negotiation, *Proceedings of the 4th International Workshop on Agent Theories, Architectures and Languages*, Providence, RI.
39. Walton, R. (1996) An evaluation of CAPSULE, a computer system giving advice to general practioners about drug prescribing, *Journal of Informatics in Primary Care*, March, 2–7.