

Decision making by intelligent agents: logical argument, probabilistic inference and the maintenance of beliefs and acts

John Fox

Advanced Computation Laboratory,
Cancer Research UK (Previously ICRF),
London WC2A 3PX, UK

Peter McBurney

Department of Computer Science
Liverpool University
Liverpool L69 7ZF

Abstract

PROforma is a language and a technology for designing, implementing, testing and delivering software agents that can operate in dynamic and uncertain environments. The agent specification language is based on R^2L , a logic language that formalises reasoning, decision-making and plan execution using a combination of classical and non-classical logics. *PROforma* can be viewed as an object-oriented layer on top of R^2L that reifies the logic into a small set of “tasks”, notably decisions, plans and I/O tasks. The heart of decision-making in both *PROforma* and R^2L is a logical decision model based on the creation and evaluation of arguments for and against alternative beliefs and actions. In previous papers we have provided proof-theoretic and model-theoretic semantics for a Logic of Argument and suggested that argumentation is a general framework which subsumes many different forms of uncertain reasoning as special cases, including qualitative and non-monotonic logics; quantitative representations like probability and possibility, and “linguistic” representations of belief. Current implementations support a simple monotonic interpretation of argument based decision-making but do not address requirements for defeasibility or probabilistic belief revision. In this paper we discuss how argumentation can provide a framework for integrating these approaches into a unified agent model.

Introduction

PROforma is an executable specification language for modelling decisions, plans and other tasks that an agent needs to have in its repertoire in order to achieve its goals (Fox and Das, 2000). The language is a versatile language for designing agent applications and has been successfully used in a variety of medical applications (e.g. Fox and Thomson, 1998). Software is available for authoring, verifying and testing applications in the language and enacting the specifications as autonomous processes, or as agents supporting human users.

The *PROforma* language is defined around four classes of task:

An *enquiry* is any process that acquires information during the enactment of a plan, whether by a dialogue with another agent, or through a sensor or other device.

An *action* is any act carried out on the agent’s environment, whether a direct call to an effector or a request for action to another agent.

A *decision* is a general procedure for making choices where there is uncertainty about the options, reliability of evidence, soundness of conclusions etc.

A *plan* is a collection of tasks for achieving agent’s goals. Plan tasks can be

enacted in parallel or sequentially over time.

The PROforma language is defined recursively over plans so that an agent can include hierarchical task structures of any complexity.

The language combines features of a *formal specification language* as developed in software engineering with features of *knowledge representation languages* as developed in AI. Alternatively it can be viewed as a hybrid of a *logic programming* language (it supports inference in propositional and predicate logics, together with certain non-classical logics) an *object-oriented* language, in which the objects are tasks, and an agent-oriented language which represents and maintains a complex mental state over time (Fox, 2002).

A PROforma task's behaviour is specified by defining properties associated with each class of task. All tasks include the following properties¹:

Task-name:	<unique ID symbol>
Goal:	(<verb><object>)
Trigger:	<event>
Preconditions:	<condition>
Postconditions:	<assertions>
Autonomy:	'automatic' or 'authorise'

The meanings of these attributes are:

Goal. The situation that the task is intended to bring about e.g. change the state of current knowledge, as in a medical diagnosis decision; change the state of the world, as in a plan or action.

Preconditions. Situational conditions that must be true before a task may be initiated e.g. a patient must be male and of a certain age.

Postconditions. Tasks may have consequences in the world, such as the side effects of drugs.

Triggers. Situations or events that can evoke a task which has not been explicitly scheduled for enactment (see plans below).

Autonomy. A task can be enacted autonomously or require external supervision/authorisation.

Decisions and plans have further attributes that distinguish them from the "generic" task and from each other. Every plan, for instance, has the following attributes as well as the general ones described above:

Plan components	the actions, decisions, sub-plans etc. which make up the plan
Scheduling constraints	ordering constraints on components, if any
Abort conditions	if true plan aborts (fails)
Terminate conditions	if true plan terminates (succeeds)

In contrast a decision has the following distinguishing attributes:

¹ There are also other attributes that are not relevant to the current discussion

Candidates	Set of decision options (choices)
Sources	Data required for decision to be taken
Arguments	Reasons for/against candidates
Commitment	Criteria/constraints for taking decision.

An agent's plans and decisions are interfaced to the outside world via *actions* and *enquiries*. These tasks inherit all the standard properties (above) but their behaviour does not have a specific impact on the present discussion and are not considered further here (see Fox and Das, 2000 for discussion of details).

A PROforma process is specified by composing tasks into collections of prepared plans that are to be carried out under particular contingencies. Tasks can be enacted sequentially (explicitly scheduled), in parallel, or reactively in response to events, or a mixture of these modes. A simple example is ERA (Early Referrals Application) which is being used in the UK National Health Service to assist doctors in deciding whether patients with suspected cancer should be referred for urgent investigation. ERA can be found at www.infermed.com/wap/era.² A graphical representation of the ERA task structure is shown in figure 1. A PROforma specification for these tasks is included in the Appendix.

This ERA referral guideline consists of a single plan containing a sequence of component tasks: an enquiry that acquires patient data, a decision that follows the data collection, and two alternative actions that follow the decision.

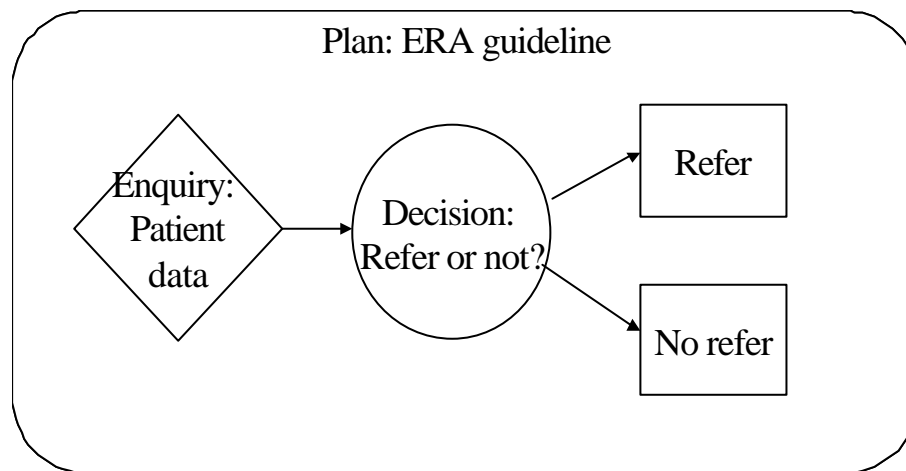


Figure 1: Task structure of an ERA referral guideline

ERA is based on a policy document "Referral Guidelines for Suspected Cancer published by the UK Department of Health"³. In this document the decision rules are

² To run one of the examples select a specific cancer by clicking on "referral form" for one of the cancers indicated. A patient data form will be displayed; once this has been completed click on the OK button to submit the data for processing by a PROforma server on the ICRF web site. ERA will return its recommendations on the basis of the data provided. Please note this is a demonstration and is not appropriate for practical clinical use without a professional clinical training.

³ <http://www.doh.gov.uk/cancer/referral.htm>

set out in informal English as in the following conditions for *urgent referral for suspected breast cancer*:

- Patients with a discrete lump in the appropriate age group (e.g. age > 30).
- Signs which are highly suggestive of cancer such as:
 - Ulceration
 - Skin nodule
 - Skin distortion
 - Nipple eczema
 - Recent nipple retraction or distortion (< 3 months)

These criteria involve uncertainty (e.g. the risk associated with individual signs and symptoms) but do not explicitly quantify it (only using phrases like “highly suggestive”). This is common in medical guidelines and, despite the lack of precision and vagueness of meaning, qualitative terms are natural for doctors, easy to understand and easy to formalise.

Decision making in *PROforma*

It is straightforward to capture such simple criteria as the premises of rules in logical languages. In the *PROforma* decision procedure rules are viewed as “arguments” that may support or oppose decision options (e.g. the patient does or does not require urgent referral) but are not necessarily conclusive. Arguments have a slightly different syntax as well as a weaker interpretation than *if ... then...* rules in traditional logic programs. An example of an argument “*for*” urgent referral in ERA is:

```
argument :: for,                                     [1]
  ( nipple_changes includes Discharge
    and nipple_disc_features include 'Large volume'
    and nipple_disc_features includes 'Bilateral' );
```

If arguments can be shown to be true then they support can support one or more decision option (in this case the option is “the patient should be urgently seen by a cancer specialist”) but a single argument may not be decisive. Expert system rules are often associated with measures of confidence (e.g. certainty factors or other quantitative coefficients) but this is not the case here. Each argument merely represents a qualitative assertion that there is a *reason to believe* or a *reason to act* (Fox and Das, *op cit* chapter 4).

In this argumentation approach arguments can be “aggregated” to yield an overall measure of confidence in a decision option, and then a preference ordering on the set of options can be determined. In the simplest scheme each argument is treated as having equal weight (for or against a candidate) and the ordering is a direct function of the ratio of supporting to opposing arguments. Despite this simplicity argumentation schemes have been found to be effective and robust for many medical decisions (see Fox et al, 2001 for a recent view). Argument systems are straightforward to set up and maintain, and easy for non-technical users to understand.

The argumentation model has emerged as a solution to a set of practical problems in the design of agent applications and decision support systems. Applications are important because they provide settings in which to evaluate theoretical proposals empirically, but they are also important because they challenge theory. Medicine is arguably one of the most difficult areas of application for computer science and AI, not least because of the many different forms and sources of uncertainty that must be addressed. Argumentation can be given a clear and sound logical semantics providing some comfort that it can be used without incurring the pathological behaviour for which “*ad hoc*” decision procedures have been criticised (e.g. in some expert systems).

Argumentation is also a framework for uncertainty management that has theoretical and philosophical interest. In our view most, if not all, approaches to handling uncertainty can be viewed as special cases of a general argumentation model. In the next section we present a formalisation of the general model and then return to PROforma as a platform for discussing some technical and practical requirements.

A formal account of argument and uncertainty

Argumentation is viewed as a general framework for logical inference and as an abstraction of a number of different kinds of uncertainty calculus (e.g. Krause et al, 1995; Fox and Das, 2000). In this framework an argument has the *form* of a logical rule, but not its *force*. For example, a doctor may have a number of reasons to believe that a patient has had chicken pox as a child but cannot be certain. The difference between a logical procedure based on argumentation and a conventional proof procedure is that the conclusion of an argument cannot be guaranteed to be correct.

A logic of argument (LA) is summarised with the following consequence relation:

$$\text{Knowledge base} \cup \text{Data} \quad \vdash_{LA} \quad (\text{Claim} : \text{Grounds} : \text{Qualifier}) \quad [2]$$

Here the *knowledge base* is a collection of facts and rules about a domain (e.g. knowledge of diseases and their symptoms and treatments) together with a set of *data* such as a set of facts about a particular patient.

The turnstile \vdash is the *meta-symbol* representing a logic (here LA) under which we may construct valid arguments about alternative claims. There are multiple options for these axioms but, whatever the choice, the argument term has the same general form: it is a triple of elements consisting of a *Claim* which the argument addresses; the *Grounds* that justify the argument (a subset of facts and rules instantiated with data) and a *Qualifier* that represents the confidence warranted by the argument in the claim.

The qualifier may take a number of forms. It can be qualitative, as when an argument “supports” a claim or “opposes” it (ditto *for, against*), meaning that the argument increases or decreases our confidence in a claim but without indicating by how much. Alternatively the qualifier can be numeric, drawn from the [0,1] interval say, so probabilistic and other conventional measures of confidence may be used to quantify the force of arguments and the confidence we have in a claim.

The axioms of the base logic may vary. They could be those of classical propositional or predicate calculus for example or, as is often appropriate in medicine, Intuitionistic Logic (IL). The main difference here is that classical logic includes the axiom of the “excluded middle” (EM: $P \vee \neg P$) i.e. exactly one of P or not-P is true but not both. Since doctors may not know all the possible diseases or drugs that are relevant to a particular case dropping this axiom is often useful in modelling clinical reasoning. Furthermore there may be arguments for both P and not-P which are simultaneously valid, which cannot be accommodated in classical logic.

In their simplest form arguments are constructed by a mechanical proof procedure that generates a sequence of proof steps ending in a conclusion (e.g. “because the patient is elderly and has lost weight then the diagnosis is cancer”, “because the patient may have cancer and cancer is a life-threatening and progressive disease we should refer him/her for urgent investigation”). However, although the form of the argument is similar to a proof, it does not have the conclusive force of a proof; in short a claim may be deductively valid but is *not necessarily true in the real world*. The premises for the argument may be unreliable for example (e.g. a laboratory test is unreliable) or the claim may later prove to be false as further data and arguments are considered.

We resolve this issue in decision making in two main ways.

Firstly, arguments can have different precedences; conclusive argument may take precedence over merely indicative arguments, for example. If we know that a patient has breast cancer then we can argue, conclusively, that she has cancer, but if a patient is elderly and has lost weight we can argue that this supports a diagnosis of cancer but is not by itself sufficient to conclude this as a diagnosis.

Secondly, although arguments are logical in nature claims are not only true or false, but generally are ordered in terms of their relative merit. The ordering is determined by the aggregation of arguments for and against each claim. As data are acquired new arguments for or against each claim may become valid, resulting in changes to the aggregation and potentially changing the ordering.

We call the PROforma decision procedure “semi-quantitative”, in that the mechanisms of argumentation are entirely logical, while the mechanism of aggregation is based on relative and/or absolute force of the arguments. Despite its simplicity the semi-quantitative approach has proved to be highly effective in many practical applications.

Non-monotonicity and argumentation

From the point of view of non-monotonic logic argument aggregation is a kind of hybrid of classical monotonic reasoning and non-classical non-monotonic logic. As more data are acquired (e.g. a patient’s symptoms) arguments are constructed from the data that may represent *reasons to believe* in different interpretations, or *reasons to act* in particular ways. These arguments simply “pile up” monotonically. Observing a patient is elderly and has lost weight may support a diagnosis of cancer while a negative biopsy opposes this diagnosis, but both arguments are simultaneously valid, and will remain valid even as other arguments are formulated. The arguments remain

valid even if conclusive arguments against cancer are subsequently made; this apparent contradiction is resolved by the aggregation process not in the production of arguments themselves.

Issues of non-monotonicity only enter in when we take a decision, when we *commit* to a belief (e.g. the patient actually does have cancer rather than any of the alternatives) or to an act (e.g. to begin chemotherapy). Only if the agent has reason to change to change its mind about the diagnosis or about the treatment does it have to consider the possible downstream consequences (e.g. another decision about follow-up tests may no longer be valid).

We have found that for many medical applications it is quite practical to ignore issues of non-monotonicity. Typically we can arrange that all relevant data about a patient have been acquired before any commitment is made, so that at the point of decision all arguments have been derived and aggregated and the most preferred decision option selected. Under these circumstances it will not normally be necessary to “decommit” from the option.

This is not, however, guaranteed in all possible situations. We cannot for example guarantee that any data we have acquired will prove to be correct – a hospital laboratory may inform us that its equipment is faulty some days after telling us that a test was normal for example. Or alternatively it may not be possible to arrange that all data are available before starting an urgent treatment. We may be forced to act on a tentative diagnosis, while we wait several days for the definitive test results to become available which could show the tentative diagnosis to be wrong.

The argumentation model can be extended to accommodate some of these problems by means of a suitably adapted belief-maintenance scheme. Let Δ be a database of beliefs associated with a decision; K the set of claims between which we are choosing, and A the set of arguments that support these claims $\{(K_i:A_j: +)\}$. Then the conditions under which an argument A_j for a claim K_i is valid are as follows

$$\begin{aligned} \text{Valid}(K_i:A_j) &\text{ iff } \exists \delta \in \Delta \text{ and Provable}(K_i: \delta: +) & [3] \\ \text{Valid}(K_i:A_j) &\text{ iff } \exists K_j \in K, \delta \in \Delta \text{ and Provable}(K_i: K_j \& \delta: +) \end{aligned}$$

Which is to say $(K_i:A_j)$ is a valid argument (under some logical consequence relation such as LA) if and only if there is a provable line of reasoning from a ground belief which supports a claim K_i or an indirect argument which depends upon other Valid arguments . It is a relatively simple programming task to create an argument maintenance system that will automatically add and delete argument terms as data are added or deleted from database Δ , as when we discover that the laboratory equipment is faulty or that a default assumption about the patient is false.

A claim remains *acceptable* while the following condition holds.

$$\text{Acceptable}(K_i) \text{ iff } \exists A_j \in A \text{ and Valid}(K_i:A_j) \quad [4]$$

Again it is a simple programming task to maintain a preference ordering over the acceptable claims.

Argument and claim maintenance are not currently supported in the *PROforma* technology though ideally it would be part of the functionality of the decision task manager. This is a subject of continuing research. For example, an interesting refinement of the maintenance scheme outlined is based on the observation that the set of “acceptable” claims based on the existence of at least one valid argument is only one criterion of acceptability. We may, for example, require a more stringent criterion, such as “there must not only be a supporting argument but also there should be no excluding counter-arguments” for a claim to be acceptable.

In everyday language a particular disease may only be a “possible” diagnosis if there are arguments in support of it being present *and* no arguments that rule it out. Another possible acceptability condition is that a claim is only acceptable if it is “plausible” meaning the aggregated supporting arguments have greater force than the aggregated opposing arguments. If arguments are also weighted with probabilities (see next section) we could have the acceptability condition that the probability of a claim must be greater than some value, such as 0.5, for it to be considered as an option in the decision process.

Argumentation also appears to offer a promising approach to providing interpretations for common “linguistic” uncertainty terms. Linguistic representations of uncertainty have considerable value in building practical medical systems as clinical users find terms like “possible”, “probable” useful and natural. Elvang-Gøransson et al (1995) have formally explored logical interpretations of such terms in analysing the idea of acceptability of arguments. Glasspool et al (1998) have shown empirically that it can provide at least as good an account of people’s use of such terms as more traditional probabilistic explanations. The relationship between argument patterns and linguistic uncertainty terms and the relationship to modal and default logics is also discussed by Fox and Das (2000, chapter 4).

For practical agent technologies, such as *PROforma*, argument maintenance is only part of the problem however, because decision-making does not only entail decisions among alternative *beliefs* but also commitments to *actions* in the world. Suppose a *PROforma* agent decides that a patient is suffering from cancer, and that the required treatment is chemotherapy, a substantial medical process involving extensive planning, ordering of services and commitment of resources. If the agent later finds that the patient data on which these decisions were made were incorrect (a common enough event) then chemotherapy may no longer be the most preferred or an acceptable treatment.

The difficulty is not that an agent’s beliefs about the diagnosis become invalid, the argument maintenance and aggregation systems will take care of that, but rather that commitments will have been made by the agent that are no longer under its control (e.g. making an appointment for the patient to be seen urgently by a cancer specialist). The agent must not only revise its beliefs but also take any action that is necessary to *reverse its commitments*.

In general an arbitrary number of commitments can arise as a result of decision-making and we require a general solution to this problem. This is not simply a matter of issuing “cancellation notices” for all commitments erroneously made, for by the

time the error is identified any number of consequences may already be irreversible. However, simply cancelling every incomplete order is not always sufficient either. For example a cancellation could have greater costs than leaving the commitment in place so the challenge is to seek a minimum-cost solution.⁴

Arguments and probabilities

On this decision model the decision process is a hybrid of classical monotonic reasoning (for argumentation) and non-classical non-monotonic logic (for commitment and de-commitment). It can also be viewed as a hybrid of logical and quantitative methods, in that in “weighing the pros and cons” of a claim we are effectively applying a linear decision rule with uniform weights on the arguments. It is an established result in the decision analysis literature that such rules are frequently as effective as more sophisticated decision procedures that use complex weighting schemes (e.g. Bayesian and expected utility decision models).

In medicine it appears that precise weighting of arguments is frequently ignored in practice (the NHS cancer referral criteria above leave uncertainty implicit) and attempting to be precise about the uncertainty involved frequently produces little or no improvement in medical outcome (Fox et al, 2001).

Despite this observation there are surely domains and applications where decision-making requires precise weighting of evidence, as when some symptoms are more strongly diagnostic of some diseases than others. *PROforma* supports a simple integer weighting scheme for this purpose, which allows an application designer to place differential weights on arguments, which can be taken into account by the built-in aggregation functions in establishing the preference orderings over the decision options. While this scheme is practical it is subject to many of the standard criticisms of *ad hoc* decision procedures (e.g. the argument that belief revision methods which do not satisfy the probability axioms will be “incoherent” (Lindley, 1984).

The use of an argumentation framework *per se* is not the primary weakness since the scheme in [2] is sufficiently general to accommodate different logics (including classical truth-functional logic) and other representations of uncertainty (including probability) while remaining within the argumentation meta-theory. This implies that if an application requires a probabilistic decision procedure it ought to be possible to capture it within the argumentation framework.

The way to achieve this in *PROforma* is by defining a new class of argument-based decision called a *Bayesian decision* which respects the probability axioms and constraints (figure 2).

⁴ As health services and airlines know patients and passengers often adopt the least-cost solution of doing nothing, having to make provision for a substantial proportion of “no shows”. We do not think that do nothing is always a reasonable solution to this problem, but await a general solution to this difficult challenge, where “non-monotonicity meets the real world”.




Figure 2: Arguments as networks

Bayesian decisions inherit the usual properties of the standard decision task outlined earlier, though we rename them with more familiar terminology from the Bayesian Network literature).

Input nodes = Data sources for a decision (e.g. patient data)
 Output nodes = Decision options or candidates (e.g. diagnoses)
 Paths =Arguments

The paths in Bayesian networks are equivalent to arguments in standard *PROforma* decisions. Paths from input to output nodes are equivalent to implications from premises to conclusions. Paths may include internal nodes representing Boolean combinations of data, possibly drawing on external facts in a knowledge base. The paths in this graph are (IN1&EN1-> ON1); (IN2->ON2) and (IN2 & IN3 & EN2) -> ON2).

PROforma arguments are associated with simple qualifiers like “for” and “against”, or simple integer weights to indicate the force of an argument with respect to particular decision options (claims). The equivalent to these in a Bayesian decision is a quantitative weight drawn from [0,1].

Path weight

The weight represents the conditional probability of a candidate being the most preferred given that all input nodes and nodes external to the path are true.

Prior belief

Each candidate will have a coefficient in [0,1] which represents the prior belief in the candidate.

Updating function

The revision function that is to be applied to the current belief in each candidate to yield a posterior probability for the candidate for each true argument. The updating function is a special class of aggregation function.

The practicality of implementing a Bayesian decision has been demonstrated by Vicens (2001) who has implemented an application for differential diagnosis of

abdominal pain based on a classical Bayesian decision support system originally developed by de Dombal et al (1972).

However, the model remains unsatisfactory in certain respects. If we opt for a standard Bayesian belief updating function then this places strict constraints on the structuring of the network. For example, there must be a complete set of dependency paths from all input nodes to all output nodes. Vicens encoded all possible paths in the abdominal pain diagnosis system mentioned above. This was possible because she had available a complete conditional probability table for symptoms and diseases which had been compiled from objective clinical data by de Dombal (*op cit*).

In practice it will be more typical that objective information of this kind is not available, so we will need to draw upon subjective clinical opinion in modelling a medical decision. Using this method, however, the decision graph is likely to be incomplete. As we saw with the ERA rules an expert diagnostician will typically assert statements like “the presence of swollen glands suggests glandular fever” while making no comment on (or possibly even denying “the absence of swollen glands suggests an absence of glandular fever (and therefore the presence of something else)”. Indeed it is a standard result from the psychology of reasoning and decision making that people generally focus on things that are true and pay insufficient information to information and relationships that are false (e.g. Legrenzi et al, 1993). In other words if we rely entirely on knowledge acquisition from human experts there may well be paths missing from the graphs.

A standard Bayesian updating function cannot accept an incomplete topology. One possible solution is to associate a “graph normalisation function” with Bayesian decisions that autonomously infers all the links and estimates associated probabilities missing from the graph. However, this approach may lead to problems when the results are presented to the medical experts as the automatically-added paths may not have any evidence to support them. An alternative may be to use a non-probabilistic representation of uncertainty. For example, we may view those edges for which there is no evidence as being constrained by the assignment of uncertainty measures given to those edges for which there is evidence. This view means that the uncertainty values appropriate for the former edges (i.e., those with no evidence) may be seen as lying inside some range of values, between a minimum value and a maximum value. Such an approach leads naturally to Belief Function (Dempster-Shafer) Theory, where the minimum and maximum values are given by Belief and Plausibility values respectively (Shafer 1976). Belief Function Theory is a generalization of probability theory, with a weaker set of axioms.

Conclusions

In this paper we have presented an agent representation language and logical argumentation system that have been used successfully to represent uncertain beliefs and guide decision-making in medical domains. Although this application domain is very specific, it has proven to be a rich source of general theoretical problems. The most interesting conclusion is that argumentation appears to be a general framework within which many other uncertain reasoning techniques can be modeled as special cases, since it can accommodate different base logics (e.g. classical and intuitionistic logics) and various representations of uncertainty (quantitative, qualitative and

linguistic). As examples we have discussed non-monotonic reasoning and probabilistic belief revision. The medical domain continues to raise new challenges. For example, it reveals the apparent absence of a well-developed theory of defeasible practical reasoning, a theory of decision-making that formalises commitments and decommitments to actions that have real-world consequences, which would be the action analogue of non-monotonic logics of beliefs.

References

- Ambler S "A categorical approach to the semantics of argumentation" *Mathematical structures in computer science*, 6 (2), 167-188.
- Das S K, Fox J, Hammond P, Elsdon D "A flexible architecture for autonomous agents" *J. Exper. Theoretical Artificial Intelligence*, 9 (4), 407-440, 1997
- De Dombal, T "Computer Aided Diagnosis of Abdominal Pain", *British Medical Journal*, 2(1), 9-13, 1972
- Elvang-Goransson M, Krause P J, Fox J "Acceptability of arguments as logical uncertainty" in M Clarke, R Kruse, S Moral (eds) *Symbolic and quantitative approaches to reasoning und uncertainty, ECSQARU93*, Lecture notes in computer science 747, Berlin: Springer, 1993.
- Fox J, Krause P and Elvang-Goransson M, "Argumentation as a general framework for uncertain reasoning" *Proc. 9th Conference on Uncertainty in AI*, Washington 1993.
- Fox, J. and Parsons, S. "Arguments about beliefs and Actions". in A. Hunter and S. Parsons (eds.) *Applications of uncertainty formalisms*, Springer Verlag, 1998.
- Fox J and Thomson R "Decision support and disease management: a logic engineering approach" *IEEE Transactions in Biomedicine*, 2 (4), 217-228, 1998.
- Fox J and Das S *Safe and Sound: Artificial Intelligence in Hazardous Applications*, AAAI and MIT Press, 2000
- Fox J, Glasspool D, Bury J "Quantitative and qualitative approaches to reasoning under uncertainty in medical decision making" in S Quaglini, P Barahona, S Andreassen (eds) *Artificial Intelligence in Medicine, Lecture notes in artificial intelligence 2101*, Berlin: Springer, 2001.
- Krause P, Ambler S, Elvang-Goransson M, Fox J "A logic of argumentation for reasoning under uncertainty" *Computational Intelligence*, 11 (1), 113-131, 1995.
- Legrenzi, P., Girotto, V. and Johnson-Laird, P.N. (1993). Focussing in Reasoning and Decision Making. *Cognition*, 49, 37-66.
- Pearl J *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* Morgan-Kaufmann, 1988.
- Shafer, G A mathematical theory of evidence, Princeton University Press: Princeton NJ, 1976.
- Vicens, C Dissertation for masters degree in decision analysis, London School of Economics, 2001.

Appendix

```
/** PROforma Guideline: Suspected breast cancer */
/** 25/10/2000, (simplified by JF for presentation, 1/5/2001 **/

plan :: Breast ;
  caption :: 'Breast' ;
  component :: Clinical_information ;
  component :: Referral_decision ;
    schedule_constraint :: completed(Clinical_information) ;
  component :: No_two_week_referral ;
    schedule_constraint :: completed(Referral_decision) ;
  component :: Two_week_referral ;
    schedule_constraint :: completed(Referral_decision) ;
end plan .

decision :: Referral_decision ;
  caption :: 'Referral decision' ;
  choice_mode :: single ;
  support_mode :: symbolic ;
  candidate :: Two_week_referral ;
    argument :: for, ( tissue_changes includes 'Discrete lump' and age >= 30 ) ;
    argument :: for, ( skin_changes includes Ulceration ) ;
    argument :: for, ( skin_changes includes Nodule ) ;
    argument :: for, ( skin_changes includes Distortion ) ;
    argument :: for, ( nipple_changes includes Eczema ) ;
    argument :: for, ( nipple_changes includes 'Retraction_or_distortion' ) ;
    recommendation :: Netsupport( Referral_decision, Two_week_referral ) >= 1 ;
  candidate :: No_referral ;
    recommendation :: Netsupport( Referral_decision, Two_week_referral ) < 1 ;
end decision .

action :: No_two_week_referral ;
  caption :: 'No two week referral' ;
  precondition :: result_of( Referral_decision ) = No_referral ;
  procedure :: 'No referral' ;
end action .

action :: Two_week_referral ;
  caption :: 'Two week referral' ;
  precondition :: result_of( Referral_decision ) = Two_week_referral ;
  procedure :: 'Referral, though not necessarily urgent' ;
end action .
```

PROforma specification of the ERA breast cancer referral guideline (excluding data definitions and with a reduced number of decision options to improve clarity). The specification covers the tasks in figure 1, consisting of a plan that contains sequence of atomic tasks: an enquiry to collect patient data, a decision to determine urgency, and two alternative actions, one of which will depend on the result of the referral decision.