The Knowledge Engineering Review, Vol. 7:1, 1992, 19–33

Qualitative frameworks for decision support: lessons from medicine

JOHN FOX AND PAUL KRAUSE

Imperial Cancer Research Fund, Lincoln's Inn Fields, London WC2A 3PX, United Kingdom

Abstract

Some weaknesses of current decision support technologies are discussed. Numerical methods have strong theoretical foundations but are representationally weak, and only deal with a small part of the decision process. Knowledge-based systems offer greater flexibility, but have not been accompanied by a clear decision theory. Theoretical development of *symbolic decision procedures* is advocated, an approach to the design of decision support systems based on first-order logic is presented, and work on this approach is reviewed. A central proposal is an extended form of inference called *argumentation*; reasoning qualitatively for and against decision options from generalized domain theories. Argumentation captures a natural and familiar form of reasoning, and contributes to the robustness, flexibility and intelligibility of problem solving, while having a clear theoretical basis. Argumentation was developed initially for medical applications though it may have much wider applicability.

1 Introduction

Computer based decision support systems (DSS) have been the subject of development and discussion in medicine for three decades. In the first 15 years or so the emphasis was primarily on numerical approaches, while more recently medicine was the setting for many early experiments with knowledge based expert systems. It continues to provide fertile ground for developing and evaluating new decision support techniques. The complexity, variety and importance of medical problems is likely to offer an effective forcing function for technical innovations in AI for many years to come.

These same factors, however, have also limited the take-up of decision support systems in medicine. In many decision making situations the nature of the required decision may be ill-defined, the available information unreliable and/or incomplete, and the possible decision options hard to determine ahead of time—yet action may be needed urgently. As in other *safety critical* fields where operating conditions are difficult and the consequences of poor performance substantial, fears are now growing about the potential for catastrophic errors created by poorly designed decision systems and, worse, the potential for catastrophes whose causes cannot be established. Questions of ethical and legal liability are becoming prominent (e.g., Gemignani, 1991) and there have even been calls for restrictions on the deployment of DSSs in safety critical situations (e.g., The Boden Report, 1989).

These concerns have indicated a number of requirements for the design of acceptable DSSs in medicine, and perhaps further afield as well.

Robustness. The system must obviously make or recommend "good" decisions. A criterion of "good" needs to be defined here, but whatever criterion is adopted the system's performance must decline gracefully in non-ideal circumstances.

Flexibility. Situations involving important decisions are rarely simple. In medicine, for example, clinicians are not normally faced with isolated diagnosis decisions (what is wrong with a patient), but many other decisions which potentially interact are usually in play as well (e.g., what investigations are required? What actions are viable? and even "is this decision really necessary?"). Though simple diagnosis systems may be of value in other domains (e.g., fault-finding in equipment), DSSs supporting a single function such as diagnosis may be of limited relevance to a clinician.

Accountability. If decisions lead to errors if must be possible to establish the reasons for those errors. This may take the form of explanations, audit trails, or some other method. However, any account of the behaviour of the system must be sufficiently intelligible that assignment of credit or blame can be effectively achieved, and a skilled human supervisor should be able to intervene effectively in decision making if and when appropriate.

Soundness. It is axiomatic that however a system is engineered we require a well developed and justified theory to underpin the design of critical components.

There have been two primary theoretical traditions in the design of DSSs; statistical decision theory and knowledge based expert systems. We now briefly review experience with these approaches in medicine, with particular attention to the above requirements.

2 "Classical" decision support systems

2.1 Summary

Informally we may define a decision as "a conscious choice between at least two possible courses of action" (Castles et al., 1971), or as "settlement of (question, etc.), conclusion, formal judgement, making up one's mind" (*Concise Oxford Dictionary*). But what formal procedure should be adopted for designing systems that help in making choices or settling questions? The view from classical decision theory is quite unequivocal:

> "First, the uncertainties present in the situation must be quantified in terms of values called probabilities. Second, the various consequences of the courses of action must be similarly described in terms of utilities. Third, that decision must be taken which is expected—on the basis of the calculated probabilities—to give the greatest utility" (Lindley, 1985).

There have been innumerable studies using the "expected utility" approach in designing DSSs for medical diagnosis and treatment (Schwartz & Griffin, 1986). Experiments with numerical techniques for diagnosis and other applications began in the sixties, and produced some early encouraging progress, notably de Dombal's classic work on the diagnosis of abdominal pain (e.g., de Dombal, 1972). Experience has shown that, when Lindley's criteria are satisfied, classical DSSs frequently perform well, and performance degrades smoothly if the quality or availability of data are reduced. Decision theory is well developed and well understood.

2.2 Critique

In practical medicine clear successes are the exception, and the impact of statistical methods is still very limited. By the end of the 1970s it was being realized that mathematical rigour and elegance were not enough for clinical adoption; the great virtues of the quantitative framework had to be set against negative factors, including high costs of development, poor operational flexibility and limited appeal to users. A number of criticisms of expected utility methods are now widely recognized.[1]

---

[1] It should be noted that these criticisms are not universally accepted; there are many who follow Lindley in holding the firm opinion that expected utility is the only correct approach to decision making and that all alternatives are necessarily inferior.

First, the compilation of a comprehensive set of objective probabilities (e.g., frequency-based estimates of the co-occurrence of symptoms and diseases) is often difficult to achieve in general medicine. Precise statistics are not routinely found in the reference literature (authors usually only speak of symptoms being "typical", "common" and so on), and the collection of such statistics takes considerable time and effort. The rate-limiting component of statistical estimation is the appearance of rare conditions, and it is therefore not uncommon for the estimation of a full set of parameters to take of the order of years. (Worse still, parameters of one patient group may not hold for a different group in a different place or different type of clinic.) To simplify development, therefore, it is often proposed to use *subjective* estimates of probabilities, but this has its own difficulties. The psychological processes which yield subjective probabilities, and the proper interpretation of such numbers, are obscure, and errors and biases in estimation are well documented (Kahneman et al., 1982).

Second, as with probability, though for different reasons, it is often difficult to assign objective measures of utility to the consequences of decisions (e.g., the utility of life and death; pain or distress). Even in situations where there seems to be an objective scale (e.g., benefit of an action in monetary terms), the relationship between subjective values and objective scales is not at all clear. Furthermore, subjective values may be multidimensional, and are frequently qualitative. For example, a drug treatment may be desirable because it is painless, because it can be taken at home, and because it is low cost, while it may be less desirable than a surgical procedure because the latter has a higher success rate, though it may compromise long-term quality of life. Translation of such concepts into standardized cost–benefit functions is fraught with technical and ethical controversy.

Third, the desire for decision aids to be accountable to, and hence controllable by, human users is not well addressed by an expected utility approach. Quantitative representations of evidence, cost-benefits and preferences are not easy to relate to common sense understanding of decision criteria, and the consequences of selecting one decision option rather than, say, the mathematically preferred one, hard to discern. The need for better explanation provided, of course, part of the motivation for the development of knowledge based techniques.

Finally, expected utility models, and perhaps purely numerical approaches in general, are inflexible. Not only must the decision situation satisfy mathematical assumptions which may be unrealistic, numerical formalisms do not provide a natural basis for automating aspects of decision making which are widely thought to be crucial for flexible patient management, such as deciding what decisions need to be taken, what options should be considered, and so on. As with explanation, AI methods were thought to have promise for improving the sophistication of decision technology.

3 Knowledge based expert systems

3.1 Summary

In the 1970s the decision support community began to experiment with "expert systems", defined as follows by Buchanan & Smith (1988):

> "An expert system is a computer program that (a) reasons with domain specific knowledge that is symbolic as well as mathematical; (b) uses domain-specific methods that are heuristic (plausible) as well as algorithmic (certain); (c) performs as well as specialists in its problem area; (d) makes understandable both what it knows and the reasons for its answers; and (e) retains flexibility."

As many texts on expert systems show, medicine was an early area for investigation of knowledge based techniques. Classic systems such as MYCIN (Shortliffe, 1976), INTERNIST (Pople, 1975) and CASNET (Kulikowski & Weiss, 1973) showed that symbolic techniques for knowledge representation, inference and heuristic reasoning held much promise for DSS design, and the idea of symbolic (i.e., non-quantitative) knowledge representation and other innovations have had great influence.

Among the advantages claimed for the knowledge based approach are that symbolic forms of representation (e.g., *if*... *then*... rules) make it easier to construct knowledge bases from the accounts of the problem solving of experts, and can support intelligible explanations of the processes and conclusions of decision making. Proposals were also made for explicitly representing the expertise of specialists, which was expected to facilitate flexibility by means of *metalevel reasoning*, or reasoning about the processes of inference and decision making themselves (e.g., how to make a diagnosis or treatment decision; what information is relevant, and so forth).

3.2 Critique

The promised benefits of knowledge based techniques have not all been routinely realised, in particular "(d) [understandability] and (e) [flexibility] are less frequently cited and less frequently achieved . . ." (Buchanan & Smith, *op cit*). The idea that systems based on rules can be made more intelligible by literally presenting the rules used in making a recommendation is now largely discounted; explanation will commonly entail an understanding of the deeper justification for the rules as well (e.g., Stenton, 1987; Southwick, 1991). Expert systems are also widely criticized for being "brittle"; if a situation arises for which the appropriate *specific knowledge* has not been encoded, the performance of the system can suddenly collapse. This is particularly worrying if the collapse can occur without any obvious manifestation to the user.

Turning to the question of soundness, early knowledge based systems were criticized on mathematical grounds. While symbolic representational techniques might capture knowledge of informal domains like medicine quite naturally, the formal requirements of sound reasoning are not always adequately addressed. A prominent area of criticism has concerned techniques for incorporating uncertainty handling into symbolic inference (e.g., Cheeseman, 1985). Considerable effort has been put into soundly combining established numerical methods with knowledge based techniques (e.g., Pearl, 1988; Lauritzen & Spiegelhalter, 1988; Kruse & Siegel, 1991), though as we shall see in the next section, the real advantages of these efforts is open to question.

The position on utility seems even less satisfactory for expert systems than for classical methods. The expert system research community does not appear to have systematically addressed the question of how to represent costs, benefits and preferences, either numerically or symbolically.

Critics have also been concerned with certain pragmatic features of "knowledge engineering" which compare unfavourably with established software engineering methodologies (Wilson et al., 1989). Even if the deductive and quantitative mechanisms of an expert system are grounded in well understood logics or calculi, the soundness of applications will still be open to serious question if the development of its knowledge base is constructed with *ad hoc* tools and techniques. Effort is now being put into developing more systematic development methodologies (e.g., Wielinga et al., 1992) and techniques for verifying and validating knowledge bases (Ayel & Laurent, 1991). We are some way from a mature technology, but we watch developments in this area with interest.

We close our critique with perhaps the most serious, yet neglected, criticism of knowledge based decision support systems; they are not based on a general, formal theory of decision making or even a clear statement of what a decision is. Consequently, there are no clear criteria for judging the soundness of a program, or predicting its performance, particularly under adverse conditions. In this respect, knowledge based DSSs compare poorly with classical statistical systems, even if the latter are limited in scope. Those with little interest in theoretical niceties may observe that if we can show empirically that a system leads to improved effectiveness or reduced resource use then we are justified in using it, but in complex fields like medicine such demonstrations are notoriously difficult to obtain. We therefore need *a priori* reasons to be confident that a decision procedure will be effective, robust, comprehensive, etc. In section 5 we outline a qualitative framework which addresses these issues.

4 The quantitative/qualitative tradeoff in decision support

Criticisms such as these have stimulated a great deal of work on rigorous and precise underpinnings for knowledge based DSSs. However, although there are important ways in which knowledge engineering techniques must be improved, it can be argued that some problems have received rather too much attention while other more critical issues have been neglected.

In recent years the work on uncertainty in AI has been increasingly dominated by a preoccupation with the correct use of probabilities and other quantitative measures of belief in knowledge based systems (the proceedings of the annual AAAI conferences on this topic may be consulted). Expert systems are typically viewed as a knowledge base, a logical engine and some method of calculating uncertainty, and it has been assumed that their performance is strongly influenced by the last. Careful work within this framework by Pearl (1988), Lauritzen & Spiegelhalter (1988) and others has provided a much better understanding of the theoretical relationship between probabilistic and symbolic inference, but little attention has been given to the actual importance of quantitative precision in practical decision making.

The need for precision has in fact been overstated. In medical diagnosis, for example, decision accuracy can be quite insensitive to the precise probabilities of symptoms. For example Fox et al. (1980) compared the performance of probabilistic and rule-based diagnosis systems using a database of patients who had been reliably diagnosed as having one of five different gastrointestinal conditions. The first system computed the probability of each of these conditions using a table of objective symptom/disease probabilities compiled from the database. The second used a set of categorical rules, which excluded quantitative information about the associations between symptoms and diseases, to interpret patterns of symptoms such as *if age(elderly) and weight_loss(present) then maybe(cancer)*. Notwithstanding the *ad hoc* nature of such rules, the diagnostic accuracy of the system was comparable to that of the probabilistic system (c 70%). Chard (1991) has obtained a similar result, finding that a qualitative approach to diagnosis in gynaecology only gave different results from a probabilistic system in those cases with a rare diagnosis, and concluded that so long as it could be ensured that a purely symbolic approach could pick up the less common conditions, then it would suffer no performance disadvantage. O'Neil & Glowinski (1990) showed that the performance of a simple unweighted decision rule was indistinguishable from that of a precise probabilistic calculation in classifying patients as having had or not had a heart attack. These are not isolated results; insensitivity of performance to precision of numerical parameters has been frequently reported; Dawes (1979) provides a classic discussion.

The preoccupation with quantitative aspects of decision making has resulted in the neglect of aspects of decision making that are of greater practical significance. An observation that encourages this view is that decision making does not consist merely of the ability to choose the most likely from among a number of hypotheses or the most preferred from among a number of options for action:

> "It is important to note that decision support is far more than option selection. Decisions that may have nothing to do with option selections may well provide valuable decision support, while those responsible for option selection may rely completely upon other decisions for information and guidance" (Andriole, 1989).

Figure 1 illustrates some of the activities involved in decision making that classical theory does not address, and expert systems have only done so in *ad hoc* ways. With the exception of option selection, there is little formal theory to underpin most of these activities. To put it another way, the elegant clarity of the passage from Lindley quoted above is perhaps only possible because so much of this complexity is ignored.

In taking a broader view of decision making than numerical option selection it is also instructive to look at what is known about human decision behaviour. Human decision making skills have had a bad press in the decision theory community, in large part because of their rough and ready use of probabilities. Yet psychological studies have shown that while human decision makers can be quantitatively imprecise they are highly flexible in their procedures. Shanteau (1987) describes how experts know what to attend to in a busy environment, how to adapt to changing task conditions, when to make exceptions to general rules, etc. Experts know a lot about what they know, and they can make decisions about their own decision making; which decisions to make, and

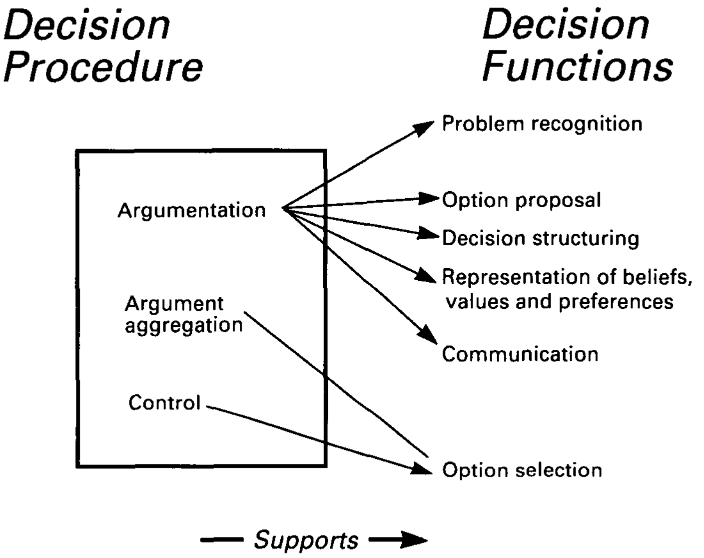

Figure 1 A range of functions required for decision support systems can be provided flexibly by a symbolic decision procedure with a small number of components

which not to, what is relevant to a specific decision, and so forth. Finally, they frequently have good communication skills and abilities to articulate their decision processes. The adaptiveness, reflectiveness and articulacy of able human decision makers are capabilities that we would surely wish to be able to mimic in DSSs, if we can do so in a theoretically sound way.

5 A generalized decision model for dilemmas

We identify much of the adaptiveness, reflectiveness and articulacy of able human decision makers with their ability to carry out sophisticated symbolic reasoning, and to represent their decision processes explicitly. In this section we describe a qualitative framework that can take us further towards emulating these capabilities than purely quantitative methods, yet can be clearly formulated. To introduce the approach consider the decision scenario in which a problem solving agent has a "mission" to ensure that its environment remains in some desired state (call this "normal"). To do this the agent must (1) make observations of its environment, and (2) act appropriately when abnormal changes occur. A general problem solving strategy runs as follows:

> Whenever an observation is made raise a goal to explain it. This may be done by consulting a library of specific knowledge, by deducing possible solutions from general knowledge, by attempting to simulate a causal process that gave rise to the observation, or in some other way. If a single satisfactory explanation is found then the goal is resolved, and a new goal to determine the appropriate action can be raised. Here, as before, candidate actions can be retrieved, deduced, identified by simulation, etc., and if a single acceptable action is identified it can be taken.

The difficulty arises when the agent is able to generate more than one solution in either step; it has a "dilemma" of what to believe or do, so a decision has to be made between the alternatives. A dilemma can only be resolved by obtaining new information. Consequently, the agent must raise goals to acquire this information (from sensors, external databases, internal computations, etc.). As information is acquired it is used to formulate *arguments* for and against the current decision options and, as a side-effect, to argue for options not so far considered.

Information acquisition may also lead to dilemmas (about how to get the information or how to interpret it) which require new decision goals to be raised; if so the decision procedure is recursively applied. A decision procedure will either terminate without resolving the dilemma (e.g., no more

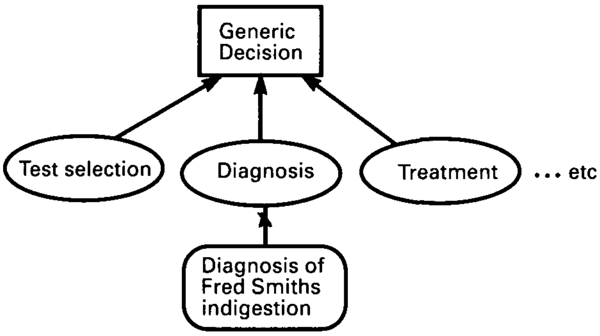

Figure 2 Decisions can be viewed as instances of general classes of decision like diagnosis and treatment decisions, which are themselves subclasses of a general class of "dilemma resolving decisions"

relevant information sources can be identified), or the arguments in favour of one option will become persuasive.

The requirements for a theory capable of dealing with such complex decision making fall into three main groups:

1 The theory must provide well defined conditions for establishing when a decision is required and when it can be taken. The former is simply the occurrence of a dilemma, the latter is discussed later.
2 The theory should prescribe a sound procedure for marshalling and evaluating relevant information. This must include:
   (a) identification of possible decision options
   (b) identification of information sources that bear on the options
   (c) construction of arguments for and against the options
   (d) assessment of the relative merits of the options.
3 If actions compete for resources, the existence of multiple decision or information acquisition goals will entail a requirement for scheduling the necessary actions.

Rather than embedding a decision procedure in a conventional algorithm, symbolic methods allow us to explicitly represent decisions, knowledge sources, reasoning strategies, representations, etc., and to reason about the control and inference processes involved. Decisions are treated as objects in a conventional class hierarchy (Figure 2). Here a generic class of decisions is defined by a set of attributes which are inherited by subclasses, such as diagnostic decisions and treatment decisions in medicine, but the values of the attributes are specific to each class. In order to diagnose Fred Smith's stomach pain, for example, we consult general knowledge of *how to make a decision* and more specific knowledge of *how to make a diagnosis decision*.

The framework is presented as three schemas. The first two define the data-structures to be maintained when any specific decision is being taken, and the third summarizes the inference mechanisms required to maintain them.

The first schema defines the data space associated with a specific decision

(*Dilemma*, *Topology*, *Labelling*, *Merit*, *Result*) [1]

where:

Dilemma is a description of the decision, i.e., a goal for which there is more than one possible solution that could satisfy it.

Topology is a directed acyclic graph consisting of: (a) a set of input or information nodes (e.g., test results, symptoms); (b) a set of output or option nodes (e.g., diagnoses, treatment options); (c) a

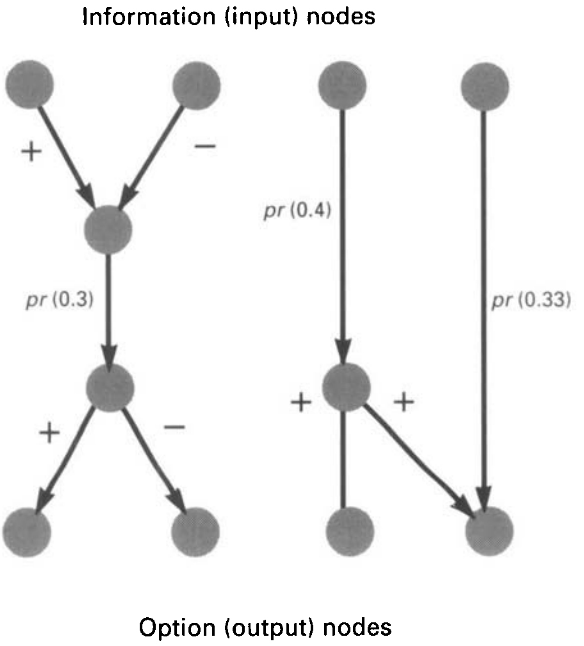

Figure 3 The topology of a dilemma. Logical argumentation can be used to construct a DAG relating information sources and decision options. A labelling operation may provide qualitative or quantitative or a mixture of both types of belief coefficient

set of (typed) edges between nodes (Figure 3). For simplicity we shall focus on edges between input and output nodes, which summarize "arguments" from input data to candidate decision options (the precise nature of an argument is described in a moment). However, input/input and output/output edges, representing relationships between data sources or between decision options, and arguments involving multiple steps (Figure 3) can also be accommodated (Fox et al., 1990a).

Labelling is a set of coefficients, one for each edge in the topology. These coefficients represent, loosely, the "force" of arguments. Such coefficients may be quantitative (e.g., conditional probabilities in the interval [0, 1] or qualitative (e.g., terms such as "supports", "detracts").

Merit is a set of pairs {(*Node*, *Value*)}, one for each output node of *Topology*.

Value is the result of aggregating the coefficients of arguments that lead into the node.

Result is either value "open", meaning no decision can yet be made, or an identifier indicating that the arguments/merit of an output node satisfy decision criteria associated with the class.

The second schema identifies a number of data-structures which are required for dynamic control of a decision

(*Dilemma*, *Class*, *Supergoal*, *Status*) [2]

where:

Dilemma is as above and *Class* is a specific class of dilemmas (such as diagnosis dilemmas, treatment dilemmas) of which it is an instance. *Supergoal* is the problem solving context in which the dilemma was encountered, and *Status* is one of {A, S, T}, indicating the decision process is active, suspended or terminated.

The third schema identifies a set of functions that operate over the data structures of [1] and [2]

(*T-Function*, *L-Function*, *A-Function*, *D-Function*, *C-Function*) [3]

where:

T-Function is a function whose domain is the application knowledge base (KB). It adds arguments to *Topology* on the basis of facts retrieved from the knowledge base, deductions, or other forms of computation.

L-Function is a function with domain *Topology* ∪ *KB* and range *Labelling*. Its role is to provide labels representing the strength of the arguments within the decision topology (Figure 3).

A-Function is a function with domain *Topology* ∪ *Labelling* and range *Merit* which aggregates the arguments for each option to yield a quantitative measure of the relative merit of the options.

D-Function is a function with domain *Topology* ∪ *Merit* and range *Result*. Its role is the central one of determining when a specific option satisfactorily resolves the dilemma, or taking the decision is otherwise justified.

C-Function defines the conditions for the control status of a decision to change (between A, S and T).

6 Application of the model

Taking a simple example of medical diagnosis we now discuss how the generic model can be applied to a specific problem. To begin with we assume that a medical decision maker (MDM) has a responsibility to establish the cause of a patient's complaints, and to decide on appropriate action.

Suppose a young patient, Fred Smith, is complaining of stomach pain after meals and loss of weight. The MDM raises a goal to explain these symptoms and, after consulting its knowledge base, argues that the symptoms could be caused by gastric cancer or a gastric ulcer. Hence there is a dilemma, and a goal is raised to resolve it. Since the dilemma concerns the aetiology of abnormal observations it is classified as a diagnostic dilemma. Knowledge of diagnostic dilemmas indicates that relevant classes of information include personal and life-style information, symptoms that can be caused by any diseases that are under consideration, and tests that can confirm the presence of a disease. Using this information as a guide the MDM establishes from its knowledge base that the patient's age, appetite and an endoscopy test are relevant to distinguishing the two diseases. Knowing the patient is young argues against both gastric cancer and gastric ulcer, but also yields an argument in favour of duodenal ulcer which is more common among young patients. The MDM now raises goals to investigate Fred's appetite, and to order an endoscopy. Fred's appetite is normal, which argues against cancer. He is sent for the test and no further action is taken pending results. These arrive a few days later and are completely negative. The MDM concludes there is no organic disease and prescribes antacids to relieve the symptoms.

The MDM has the simple goals of explaining observations and taking actions when necessary. If an observation (presence of pain) has a single explanation, or a desired state (absence of pain) can be achieved in only one way then no decisions are necessary. However, a dilemma is encountered here, in trying to explain Fred's symptoms. The framework can be applied to resolving the dilemma as follows:

Generating the topology. The diagnosis topology is developed progressively, finally including age, weight loss, indigestion, appetite and endoscopy (input nodes), and gastric cancer, gastric ulcer, duodenal ulcer and non-organic disease (output nodes); and certain causal and empirical arguments by which the symptoms are held to support or weaken the various explanatory hypotheses.

The function which generates the topology (*T-function*) can be modelled as an extended form of deduction called argumentation. In everyday decisions arguments are often routine, and can therefore be represented in precompiled (if . . . then . . .) typical of first-generation expert systems. More generally, however, arguments may need to be *constructed* to explain new observations about which we have little specific knowledge. To do this we exploit *theories*, expressed in first-order logic (FOL), rather than special case rules. Theories are generalizations about the world, and thereby justify the construction of specific arguments in particular situations.

Argumentation is the process of arriving at a proposition P by means of theories which can be represented as sets of facts and inference rules in FOL (Fox & Clarke, 1991; Fox & Krause, 1991). The basic idea is this. If we can construct a proof that links an input node to an output node, then the argument is in the topology. If either node is absent from the topology then it is added. Formally we can define the construction of proofs as follows:

$$\text{Theory} \cup \text{Data} \vdash (\text{Proposition}, \text{Grounds}) \qquad [4]$$

Paraphrasing loosely we say that given some data (e.g., symptoms) and a general theory about dependencies among states of the world (here a knowledge base about gastrointestinal conditions), we may deduce a proposition (e.g., that a disease is a candidate diagnosis) with a proof consisting of a sequence of deductions. (Working in the other direction we can argue from a theory and a hypothesis that some as yet unobserved symptom is relevant to the decision.) This is conventional deduction, except for the explicit representation of the *grounds* of the argument; the theory that warrant it, the assumptions made, and perhaps even the complete proof. The grounds can serve functions ranging from explanation to truth-maintenance and many other forms of metalevel reasoning.

Labelling the arguments. Arguments have "strength" or "force". This is represented in the model by a function (the *L-Function*) which can be defined formally as

$$\text{L-Function}_R\text{: Proposition} \times \text{Grounds} \rightarrow \text{Proposition} \times \text{Grounds} \times \text{Label}_R \qquad [5]$$

The L-Function is a class of functions which maps arguments to arguments with labels representing their strength. In general labels can be quantitative or qualitative, and in various representations *R*. For example, in traditional expert systems rules are supplied by a knowledge base designer along with an explicit probability, possibility, certainty factor, etc.; here the L-Function is some external process of estimating the values, whether objective or subjective.

In the example the labelling of arguments is entirely qualitative, i.e., the arguments are simply for or against the alternative hypotheses. Here the L-Function simply says that if an argument can be constructed for P then it *increases* belief in P, while if it supports not(P) then it *reduces* belief in P. The underlying semantics can be viewed simply as the sign of the change in belief in P, in some representation *R*, that is warranted by the argument. Such labels correspond to derivatives, and can therefore be defined for any uncertainty representation *R*; probability, possibility, belief function, etc. Wellman (1990) has described a qualitative basis of this type for probability, and Parsons (1991) discusses representations which can accommodate probabilities, possibilities and belief functions.

Since the grounds for an argument are represented explicitly, we can reason about arguments as well as their conclusions. One use of this capability is to construct "meta-arguments" about the credibility of other arguments. Let *Args* be some set of arguments about a proposition, then

$$\text{Metatheory} \cup \text{Args} \vdash ((\text{Proposition}, \text{Grounds}, \text{Label}), \text{MGrounds}, \text{MLabel}) \qquad [6]$$

This provides a basis for comparing the relative strength of arguments and sets of arguments, specifically we can construct meta-arguments which rebut or discount other arguments. Suppose we have an argument that Fred may have gastric cancer. This will have a certain strength, but if there are arguments for doubting one of the assumptions of this argument then it will be weakened, and if the generality of the theory that warrants the argument is questioned if will be weaker still. This is a style of reasoning with which we are intuitively familiar; it is entirely qualitative but permits a labelling in which arguments are not assumed to have equal force.

Aggregating the arguments. Given an exclusive set of possible decision options the decision maker needs to be able to weigh up their relative merits. The recommendation for how to do this from classical decision theory is clear: "Numbers have . . . been associated with the events and with the consequences. The final state in the argument [*sic*] is to associate numbers with the decisions, in such a way that the best decision is that with the highest number" (Lindley, 1985, p. 57). Our example was restricted to the use of qualitative signs, arguments simply increased or decreased belief in diagnoses, but we did not give a precise figure for the degree of the change in belief. However, although argumentation is based on qualitative reasoning it does not exclude a quantitative result. In classical FOL a single proof makes further proofs redundant, but in argumentation each proof provides further information by yielding new grounds for increases or decreases in belief. Formally

$$\text{A-Function}_R\text{: P(Proposition} \times \text{Grounds} \times \text{Label}_R) \rightarrow (\text{Conclusion} \times \text{Merit}_R) \qquad [7]$$

(where P stands for Power set). A simple function for computing the merit of a possible decision option simply determines the proportion of supporting arguments in the total set of arguments (Fox, 1991). In other words, all arguments are treated as having equal weight. On the surface this seems unsatisfactory, but recall our earlier review of some of the considerable literature indicating that decision procedures which ignore weights of evidence can nevertheless be effective. Intuitively, of course, we often think of arguments as having different degrees of force, but this can still be captured within a qualitative framework by meta-level reasoning about the relative persuasiveness of arguments, as discussed in the last section.

Taking the decision. There is rather more to decision taking than simply ordering the alternative options and picking the one with the highest value. We cannot be comfortable with the idea that decisions, particularly important decisions, should be taken simply on the basis that some figure of merit has been exceeded. For example, when we are accumulating information over time we may, by chance, obtain a run of findings that favour a certain option and our threshold of merit is exceeded, yet further data acquisition may reverse the position. We appear to need stronger criteria than simple ordering on the basis of some numerical measure of merit.

A more satisfactory class of decision functions (D-Functions), which yields a decision on the graph, is easy to appreciate, but the precise properties such functions should have are open to discussion. Intuitively we want the function to define exactly those conditions where a set of arguments is "persuasively" in favour of an option. In the example this was trivial; we allowed endoscopy to give a definitive answer. This is frequently adequate, but in medicine it is also common that there is no definitive information. In these conditions we may need to take symbolic constraints into account before taking a decision, by examining the topology as well as the merit of the options. For example, we might act on the most meritritious diagnosis only if there is no further information that, if available, would change the ordering of options on their merit. Similarly, we might only believe a diagnosis if it can also explain ("cover") all the observed symptoms. Finally, we might allow multiple decisions to influence each other. For example, we may not be able to distinguish two diseases on the available information (duodenal ulcers and gastric ulcers have similar symptoms), but we may be able to resolve the dilemma by arguing that it does not matter; duodenal and gastric ulcers can both be effectively treated with drugs of a type called H2-antagonists.

Finding general and convincing criteria for judging when we should be "persuaded" of a decision option is a poorly understood business. It may be that general criteria for all types of decision do not exist, and decision rules are specific to particular classes of decision and/or particular application domains.

Controlling the decision process. For brevity we shall say little about control here, since it adds little to what we have already said about qualitative methods in decision making. In the example the control status of the diagnosis decision changed twice, from active to suspended and back, while the results of endoscopy were awaited. It is not difficult to define some set of control rules for decision making, such as "resolve any dilemma arising during decision making before seeking further information", or to program a scheduler embodying such rules. It is interesting to note, however, that control itself involves decisions, and an attractive possibility would be to implement the scheduler itself with a symbolic decision procedure like that described here, so that it could argue the pros and cons of trying to solve goals guided by general or domain specific knowledge.

7 Discussion

In the introduction we identified a number of requirements that a decision support system should satisfy, and considered the adequacy of classical decision theory and expert systems with respect to these requirements. The main criticisms of classical theory were that it does not address important aspects of decision making, such as whether a decision is required, or how to construct the decision graph dynamically, while knowledge based decision support systems compare poorly with classical ones because they have not been supplied with a sound decision theory on which to construct them. In this paper we have presented a theoretical framework for knowledge based DSSs which has more general applicability than the classical numerical approach. We now review the four criteria of robustness, flexibility, accountability and soundness introduced above to assess the adequacy of the proposals.

7.1 Robustness

Practical reasoning must cope with lack of knowledge and lack of quantitatively precise knowledge. Argumentation has been developed specifically to address these problems. It takes advantage of concepts from first-order logic to represent and argue from general theories about the world (e.g., general theories of medical pathophysiology) without depending upon comprehensive collections of special-case rules (e.g., explicit statements about exactly which symptoms can be caused by exactly which diseases under what circumstances).[2]

In addition, arguments do not depend upon the availability of precise coefficients such as prior and conditional probabilities, but only on the ability to construct arguments for and against decision options. If precise information is available, or the knowledge base includes a mixture of precise and imprecise information, we can make use of semiqualitative aggregation functions which accommodate both (Parsons & Fox, 1991; Parsons, 1991).

Argumentation has much in common with the forms of qualitative reasoning which were developed for reasoning about physical systems when precise system parameters are unavailable (Coiera, 1991; Travé-Massuyès, 1992), though here the emphasis is on decision making rather than simulation.

Another aspect of robustness is that we may have to cope with inconsistency in knowledge bases. Most non-trivial bodies of knowledge, particularly those that are developing rapidly such as medical knowledge, can contain inconsistencies so some such capability is highly desirable. Contradictions cannot normally be tolerated within classical deduction, but argumentation helps to address the problem. By labelling propositions with grounds for belief, we are able to capture a statement such as "we have reason to believe P is true on such and such an argument, but on this other argument we have reason to believe it is false". At the metalevel there is simply no inconsistency. On the contrary, metalevel reasoning provides us with a basis for comparing arguments and counter-arguments and, as we discuss elsewhere, reasoning about contradictions explicitly (Fox & Krause, 1991).

7.2 Flexibility

The inflexibility of decision algorithms based exclusively on expected utility methods derives in large part from the assumption that a human "decision analyst" is necessary to decide on what

---

[2] The development of "deep" expert systems which exploit generic forms of reasoning have in part been motivated by similar concerns (Keravnou & Washbrook, 1990).

decisions need to be taken, what information is relevant, what decision options should be considered, and so forth. These commitments must be made before such a system can be deployed, but, as we have seen, a large part of the skill of decision making is the ability to formulate a decision dynamically, as information becomes available (Shanteau, 1987). In the symbolic approach we can contemplate mechanising these things. First we identify a situation requiring a decision simply as any problem solving situation where a dilemma arises, where more than one solution is possible. Second, we have described how argumentation can be used to contruct the decision graph dynamically, based on general theories of medicine (or whatever domain we are interested in) as well as specific knowledge.

7.3 Accountability

We have also emphasized the importance of good communication between human supervisors and/or auditors wishing to examine, and potentially to intervene in, any aspect of the decision process. The argumentation approach opens up a wide range of possibilities here, including the ability to account for the decision options that are being considered and why, the information that is potentially relevant to a choice, conclusions and their arguments and counter-arguments, and so on. The framework offers opportunities for providing a wide range of explanations and summaries in the user interface of a DSS in practice as well as in theory (Fox et al., 1990b).

We observed that probabilities, utilities, etc., are particularly obscure to practising clinicians. This led us to attempt to develop a symbolic scheme for representing beliefs and values. Attempts to develop symbolic uncertainty schemes are not new; countless proposals have been made to view ordinary terms like *possible*, *probable* as numerical; they are typically represented as probability intervals. On the whole these attempts have been unsatisfactory (Clark, 1991). Argumentation permits a different approach. We define beliefs in a proposition with respect to the *pattern of arguments* for and against the proposition, which appears to offer an intelligible and natural approach (Fox, 1990).

7.4 Soundness

The outstanding requirement concerns the soundness of symbolic decision procedures. Probabilistic reasoning must obey certain mathematical constraints, and if it does so it can be relied upon to provide efficient and in certain senses "rational" orderings on options. Knowledge based systems, most obviously rule-based ones, typically exploit logical deduction, though frequently import *ad hoc* features. Current trends are to combine deduction and probabilistic inference more carefully, but as we have seen, many aspects of decision making have been neglected in these developments. Argumentation is a framework for using conventional deduction in constructing arguments and well understood calculi for representing and aggregating beliefs (whether they appear in qualitative or quantitative form), and is able to support many of the requirements of practical decision making.

8 Conclusions

The symbolic decision framework presented here addresses the demands for robustness, flexibility and intelligibility required for medical decision support systems, while preserving a claim on theoretical soundness. Its development has been driven by the concerns of a particular application domain, but the techniques appear to be applicable to other fields where "dilemmas" are among the challenges to practical knowledge engineering.

Acknowledgements

We would like to thank Simon Ambler, Dominic Clark, Mike Clarke, Mirko Dohnal, Mike O'Neil and Simon Parsons, who have all contributed to the development of the work described here.

References

Andriole, SJ, 1989. *Handbook of Decision Support Systems* Tab Books.

Ayel, M and Laurent, J-P, eds., 1991. *Validation, Verification and Test of Knowledge* John Wiley.

Boden, M, 1989. *Benefits and risks of knowledge-based systems*, Report of Council for Science and Society, Oxford University Press.

Buchanan, BG and Smith RG, 1988. "Fundamentals of Expert Systems" *Ann. Rev. Comput. Sci.* **3** 23–58.

Castles, FG, Murray, DJ and Potter, DC, eds., 1971. "Introduction" to *Decisions, Organisations and Society* Penguin.

Chard, T, 1991. "Qualitative probability versus quantitative probability in clinical diagnosis: a study using a computer simulation" *Medical Decision Making* **11** 38–41.

Cheeseman, P, 1985. "In defence of probability" In: *Proceedings of 9th International Joint Conference on AI* Morgan Kaufman, 1002–1009.

Clark, DA, 1990. "Verbal uncertainty expressions: a critical review of two decades of research" *Current Psychology: Research and Reviews* **9**(3) 203–235.

Coiera, E, 1992. "The qualitative representation of physical systems" *The Knowledge Engineering Review* **7**(1).

Dawes, RM, 1979. "The robust beauty of improper linear models in decision making" *American Psychologist* **34** 571–582.

de Dombal, FT, 1975. "Computer assisted diagnosis of Abdominal pain" In: J Rose and JH Mitchell, eds., *Advances in Medical Computing* Churchill Livingstone.

Fox, J, Barber, DC and Bardhan, KD, 1980. "Alternatives to Bayes: A quantitative comparison with rule-based diagnosis" *Methods of Information in Medicine* **19**(4) 210–215.

Fox, J, 1990. "Automating assistance for safety critical decisions" *Phil. Trans. Roy. Soc.* **B 327** 555–567.

Fox, J, 1991. "Decision theory and autonomous systems" In: MG Singh and L Travé-Massuyès, eds., *Proceedings of IMACS workshop on Decision Support Systems and Qualitative Reasoning* North Holland.

Fox, J and Clarke, M, 1991. "Towards a formalisation of arguments in decision making" In: *Proceedings of AAAI Spring Symposium on Argumentation and Belief* AAAI.

Fox, J, Clark, DA, Glowinski, AJ and O'Neil, M, 1990a. "Using predicate logic to integrate qualitative reasoning and classical decision theory" *IEEE Transactions on Systems, Man, and Cybernetics* **20** 347–357.

Fox, J and Krause, P, 1991. A formal basis for argumentation and practical reasoning, ICRF technical report.

Fox, J, Gordon, C, Glowsinski, AJ and O'Neil, M, 1990b. "Logic engineering for knowledge engineering: the Oxford system of medicine" *Artificial Intelligence in Medicine* **2** 323–339.

Gemigniani, MC, 1991. "Some legal aspects of expert systems" *Expert Systems with Applications* **2**(4) 269–284.

Kahneman, D, Slovic, P and Tversky, A, eds., 1982. *Judgement Under Uncertainty: Heuristics and Biases* Cambridge University Press.

Keravnou, E and Washbrook, J, 1989. "What is a deep expert system?" *The Knowledge Engineering Review* **4**(3) 205–233.

Kruse, R and Siegel, P, eds., 1991. *Symbolic and Quantitative Approaches to Uncertainty* Springer-Verlag.

Kulikowski, CA and Weiss, S, 1973. "An interactive facility for the inferential modelling of disease" *Proc. Princeton Conf. on Information Sciences and Systems*, 524.

Lauritzen, SL and Spiegelhalter, D, 1988. "Local computations with probablities on graphical structures and their application to expert systems" *J Roy. Statist. Soc* **B 50** (2) 157–224.

Lindley, DV, 1985. *Making Decisions (2nd ed)* John Wiley.

O'Neil, M and Glowinski, A, 1990. "Evaluating and validating very large knowledge-based systems" *Medical Information* **15** 237–251.

Parsons, S and Fox, J, 1991. "Qualitative and interval algebras for robust decision making under uncertainty" In: MG Singh and L Travé-Massuyès, eds., *Proceedings of IMACS workshop on Decision Support Systems and Qualitative Reasoning* North Holland.

Parsons, S, 1991. "Qualitative methods for integrating uncertainty handling formalisms", Technical Report, Department of Electrical Engineering, Queen Mary and Westfield College, London.

Pearl, J, 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann.

Pople, HE, 1975. "DIALOG: A model of diagnostic logic for internal medicine" *Proc. IJCAI-77*, 814–818.

Shanteau, J, 1987. "Psychological characteristics of expert decision makers" In: J Mumpower, ed., *Expert Judgement and Expert Systems* vol F35, NATO ASI Series.

Shortliffe, EH, 1976. *Computer-based medical consultations: MYCIN.* Elsevier, New York.

Schwartz S and Griffin, T, 1986. *Medical Thinking: The psychology of medical judgement* Springer-Verlag.

Southwick, RW, 1991. "Explaining reasoning: an overview of explantion in knowledge based systems" *The Knowledge Engineering Review* **6**(1) 1–20.

Stenton, SP, 1987. "Dialogue management for co-operative knowledge based systems" *The Knowledge Engineering Review* **2**(2) 99–122.

Travé-Massuyès, L, 1992. "Qualitative reasoning over time: history and current prospects" *The Knowledge Engineering Review* **7**(1).

Wellman, MP, 1990. "Fundamental concepts of qualitative probabilistic networks" *Artificial Intelligence* **44** 257–303.

Wielinga, R, Schreiber, AT and Breuker, J, 1992. "KADS: A modelling approach to knowledge engineering" *Knowledge Acquisition: Special issue on KADS* (to appear).

Wilson, M, Duce, D and Simpson, D, 1989. "Life cycles in software and knowledge engineering: a comparative review" *The Knowledge Engineering Review* **4**(3) 189–204.