Bayesian Methods for Collaborative Decision-Making

Dr. Bruce D'Ambrosio

Robust Decisions Inc.

The Problem

More and more, decision-making is a team activity. Whether a multi-disciplinary design team is choosing a power-source for a next generation, fuel-efficient vehicle, selecting the best vendor or identifying the next business strategy; integration of diverse sets of opinions and priorities is a crucial component of modern decision-making. Unfortunately, until now there has been no good computerized support for team decision-making activity. Informal methods emphasize multiple participants and focus on problem structuring, but lack the ability to provide any analysis or guidance. Formal methods, on the other hand, provide strong analytical capabilities, but only in the context of an individual decision maker.

RDI's patent-pending approach to decision support addresses all the above issues. It is grounded in Bayesian decision theory and offers strong analytical support for both problem structuring AND decision-making in a team-oriented setting. It is instantiated in *Accord*™; the lightweight, windows based decision management system.

Bayesian Decision Theory

Bayesian decision theory has its roots in the work of an obscure 18^{th} century cleric (Rev. Bayes) who worried about how to combine evidence in legal matters. However, its modern form traces to the work of John Von Neumann, mathematician and computer pioneer, in the 1940s; and J. Savage in the 1950s. In Savage's formulation[1], a decision problem has three elements: (1) beliefs about the world; (2) a set of action alternatives; and (3) preferences over the possible outcomes of alternate actions. Given a problem description, the theory prescribes that the optimal action to choose is the alternative that Maximizes the Subjective Expected Utility (MSEU). Bayesian decision theory excels in situations characterized by uncertainty and risk, situations where the available information is imprecise, incomplete, and even inconsistent, and in which outcomes can be uncertain and the decision-maker's attitude towards them can vary widely. Bayesian decision analysis can indicate not only the best alternative to pursue, given the current problem description, but also whether a problem is ripe for deciding and, if not, how to proceed to reach that stage. However, there is a well-known problem in applying Bayesian decision theory to problems involving multiple decision-makers: there was no known sound way to integrate the preferences of multiple decision-makers.

RDI's patent-pending methods solve this problem, extending the application of Bayesian methods to team decision-making. RDI's methods also significantly extend the scope of Bayesian modeling to problem-formulation, previously only available in informal decision-making methods that provide no analytical support.

[1] Savage, Foundations of Statistics, 1955.

Elements of a Bayesian Decision Model
As stated before, a Bayesian decision model, as specified by Savage, has three elements: (1) a set of beliefs about the world; (2) a set of decision alternatives; (3) a *preference* over the possible outcomes of action. So, for example, we might *believe* that house *A* has 2000 square feet, that the seller will probably accept an offer of $300,000, and so on. At the same time, we might have three alternatives available to us: we might be considering house *A*, be considering house *B*, or neither. Finally, we might *prefer* living in a 2500 sq ft. house to living in a 2000 sq ft house. While all of this seems straightforward, successful application to team decision-making involves some subtleties in both belief and preference modeling.

Belief modeling must, first of all, be simple and intuitive. Complex models that require vast amounts on precisely specified information may be theoretically attractive, but are useless to the busy practitioner. We begin with the premise that teams have objectives. These may be ill understood initially, but nonetheless are of prime importance. More specifically, the beliefs we model are the team members' beliefs about *how decision alternatives impact objectives*. These are simple statements of the form "House *A* has at least 2000 Sq Feet" or "Engine *Y* can be produced for under $500 per unit". Often team members will be uncertain about whether such a statement is true or not (e.g., the engine may not have been designed yet, or the house may have only been seen from the outside). RDI's methods provide simple graphical tools for stating not only current belief about a statement, but also the amount of *knowledge* the team member has about the topic. For example, a production engineer looking at a detailed set of blueprints has much more knowledge on which to base his production cost estimate than a design engineer looking at a "back of the envelope" design of the same engine. We model the team-member's estimate as his/her belief about the objective, and the *knowledge* estimate as the relationship between the team-member's belief and the actual state of the matter. This model provides the formal basis for combining beliefs from multiple team members into an overall assessment of the quality of each alternative.

The other major component of RDI's collaborative decision modeling is its **preference model**. A preference model corresponds, roughly, to a set of objectives or criteria that are used to judge the alternative solutions. In its most general form, a decision-theoretic preference model can represent any consistent set of user preferences. However, the general form is very difficult to describe and harder to understand. A typical simplification is to use a simple *additive* model, in which one first decomposes overall preference into a set of *objectives* (e.g., I want a house with at least 2500 sq ft. of living space), and then assigns weights to each objective. This assumes that the team member is willing to trade-off losses in one objective for gains in another. While this is sometimes the case, it does not always hold. For example, you may be unwilling to spend much more than $500/unit for an engine, regardless of how much horsepower it can generate or how good gas mileage it will provide. RDI's' proprietary preference model is non-linear; allowing for incorporation of "must-have" attributes as well as attributes a team member is willing to trade off with other attributes. The model is simple and easy to understand, and requires only a single parameter (weight) per attribute.

There remains a second problem with preference models: as we mentioned earlier, there are strong theoretical reasons why it is impossible to combine preferences from multiple decision makers.[2] How, then, can RDI analyze and make recommendations in team decision-making situations? Simple – we don't combine team-members preferences! We will discuss this further below.

Decision Model Analysis

Team members gain significant advantage simply from structuring a decision problem. Bayesian modeling permits RDI to go further by providing analytical support to the decision-making process. We combine Bayesian methods, which provide basic analyses, with expert-systems-based methods for advanced guidance on how to proceed.

RDI's solutions include the following Bayesian analysis methods:

Subjective expected utility is the standard decision-theoretic analysis used to identify the optimal decision. It is, essentially, the same analysis you would use to reject an offer to pay $1.00 for a 1-in-a-million chance to play the lottery and win $10,000.

Marginal Value of Information is an analysis of the *value* of additional information. For example, given everything that you now know, how much would it be worth to know more about whether a specific engine can get 50 mpg? Well, suppose the engine is currently your best choice, and even if it can't get 50 mpg, it would still be your best choice, due to its low pollution and low cost per unit. Then it isn't worth anything to know the engine's gas mileage! Why not? Because the information won't change your decision. Marginal value of information is a key element of our support for the decision *process* – guidance on whether or not the problem is ready for a decision, and if not, where to focus further effort. RDI calls this "what to do next" analysis.

Probability of being best is RDI's proprietary solution to the problem of combining preferences from multiple team members' evaluation. Simply stated, RDI avoids combining team members' evaluations by sampling from the entire space of preference functions bounded by the individual member preferences. If the same choice is best everywhere in the space, then the choice is clear. If there are some parts of the space in which a different choice is preferred, then the analysis measures the spatial volume in which each choice is preferred. As a result of the semantics of the model, this volume can be interpreted as a *probability* that the choice is the best under a (unknown) consensus preference model.

The Basic Mathematics

This section develops the basic mathematics behind RDI's decision support.

It may seem that the *alternative/criterion* representation for a decision problem is rather simplistic and ad-hoc. However, support for this representation comes from extensive research into modeling decision-making processes in design. In addition, there is a

[2] Briefly, Arrow's theorem shows that there is no way to combine multiple individual preference functions in a way that yields results compatible with intuition in all cases.

fairly straightforward mapping to an influence diagram, as shown in Figure 1. It is this graphical representation from which our model of argumentation is derived.

Figure 1 contains representations of the alternatives available, the criteria by which alternatives will be judged, the relative importance of the criteria, and design team member opinions on the likelihood that various alternatives meet various criteria.

Diagram Semantics

In Figure 1 the box labeled "Decision" takes as values the alternatives for resolving the issue represented by the diagram. The circle labeled $S(C_c|A_a)$ represents the satisfaction of criterion C_c given alternative A_a and will be called a *satisfaction* node. While we show only one, there will be one for each alternative/criterion combination. In our most basic evaluations we allow only Boolean ({*satisfied, unsatisfied*}) satisfaction levels[3]. Therefore, knowledge and confidence are about the certainty that the alternative will satisfy the criterion, not the degree to which satisfaction is achieved[4]. The pair of two node chains hanging from $S(C_c|A_a)$ represents opinions posted by participants. There can be any number of such chains hanging from each of the

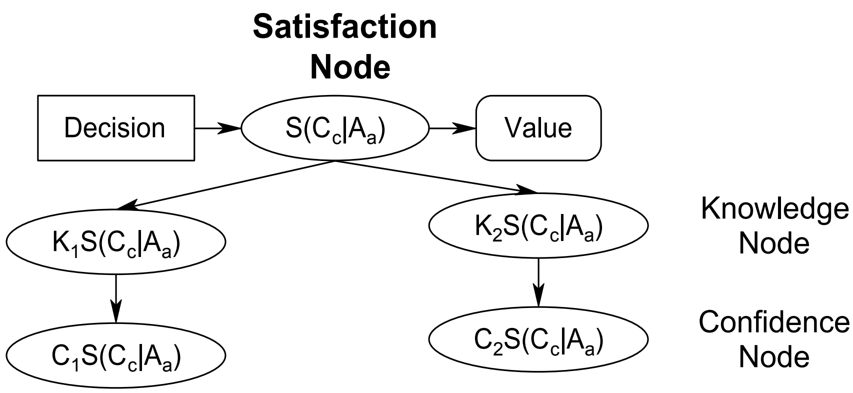

Fig. 1. Influence diagram

$S(C_c|A_a)$ satisfaction nodes, one for each opinion. The higher of the two circles represents the state of participant knowledge about the ability of the alternative to meet the criterion, and the lower is a diagram artifact used to encode probabilistic evidence. The upper node (we will call this a knowledge node) takes the same values as the original satisfaction node, namely {satisfied, unsatisfied}. We will denote these nodes as $K_pS(C_c|A_a)$, where *a* is the specific alternative being addressed, *c* is the criterion, and *p* is the participant. The lower node takes a single value, *true*[5].

[3] Methods not described here allow for measured satisfaction levels.

[4] For some criteria the degree to which satisfaction can be achieved can be measured. This will be the topic of a future paper.

[5] It is more usual, perhaps, to simply represent *confidence* as a likelihood statement on the knowledge node. However, we find the explicit graphical representation useful. These can be interpreted as standard belief net nodes that have been observed.

The conditional probability distribution for the knowledge node given the actual satisfaction has two degrees of freedom. We reduce this to a single degree by assuming symmetry to simplify knowledge acquisition. That is, we assume

$P(K_pS(C_c|A_a)=yes|S(C_c|A_a)=yes) = P(K_pS(C_c|A_a)=no|S(C_c|A_a)=no)$.

This single degree of freedom is the *knowledge* the participant has about the alternative/criterion pair, because this single parameter encodes how accurately the participant's belief reflects the actual world state. The complete distribution for a knowledge node, then, is:

$S(C_c\|A_a)$	$P(K_pS(C_c\|A_a)=yes\|S(C_c\|A_a))$	$P(K_pS(C_c\|A_a)=no\|S(C_c\|A_a))$
Yes	$K_{c,a,p}$	$1 - K_{c,a,p}$
No	$1- K_{c,a,p}$	$K_{c,a,p}$

We allow $K_{c,a,p}$ to range between 0.5 and 1.0, where 1.0 represents perfect knowledge and 0.5 represents complete ignorance, and use the textual scale described earlier to acquire the $K_{c,a,p}$ value.

We will refer to the lower node as the *Confidence* node, $C_pS(C_c|A_a)$. The confidence node has only one value and all that matters is the ratio of the probabilities for that value given $K_pS(C_c|A_a)$ (i.e., this node holds the user stated *likelihood ratio,* normalized to a 0-1 range). We acquire this as the "probability that the alternative will satisfy the criterion, given the participants state of knowledge." That is, we treat the participant as a making a noisy or soft observation (report) on his or her belief. We encode this as a pair of numbers constrained to sum to one, as follows:

$K_pS(C_c\|A_a)$	$C_p(S(C_c\|A_a)$
Yes	$C_{c,a,p}$
No	$1- C_{c,a,p}$

Note that this model assumes uncorrelated evidence from team members, and thus is optimized for multi-disciplinary teams. While modeling correlation among opinions is straightforward, it is an extra burden on the team that outweighs the advantages in most situations. Since this is a dynamic model, information presented by one team member can affect the knowledge and confidence of others. This is taken into account by changing the K_p and C_p values. In the computerized instantiation such a change adds a new record to the database, effectively recording the argumentation history.

Alternative Evaluation

Given the above semantics, the expected value of an alternative is:

$$EV(A_a) = \Sigma_c W(C_c)P(S(C_c|A_a)=yes)$$

where

$W(C_c)$ is the weight assigned to criterion C by the participant;

and

$$P(S(C_c|A_a)=yes) = \alpha \Pi_p \left(C_{c,a,p}K_{c,a,p} + (1- C_{c,a,p})(1- K_{c,a,p})\right)$$

and α is a normalization factor:

$$\alpha = 1/(\Pi_p (C_{c,a,p}K_{c,a,p} + (1- C_{c,a,p})(1- K_{c,a,p})) + \Pi_p (C_{c,a,p}(1- K_{c,a,p}) + (1- C_{c,a,p})K_{c,a,p}))$$

The alternative with the highest satisfaction value is the "best" as judged by the team.

$$EV(Decision) = \max_a EV(A_a)$$

Our approach to aid the team in planning what to do next is to compute value of further exploration of each alternative/criterion pair. This is accomplished by calculating $EV(Decision \mid S(C_c|A_a)=yes)$ and $EV(Decision \mid S(C_c|A_a)=no)$. The first is found as it was for EV(Decision), but with a pair of nodes added indicating perfect knowledge and confidence that alternative *a* will satisfy criterion *c*(K=1.0, C=1.0). This perfect knowledge calculation clearly shows the highest satisfaction achievable if the knowledge in each of the alternative/criterion pairs is as high as it can be. Another way to look at this calculation is that it is as if a new team member was added to the team. For each attribute of each alternative this person is "the" expert and has confidence that the alternative in question perfectly meets the criterion. This calculation shows how this person would change the satisfaction and possibly the team's decision.

Similarly, we also compute $EV(Decision \mid S(C_c|A_a))=no$, the situation with *c*(K=1.0, C=0.0). Here the expert has told the team that there is no way the alternative can meet the criteria and so the lowest possible satisfaction is calculated.

This analysis serves as a basis for other calculations in the Accord™ software solution.

Summary
Our unique, patented formulation of Bayesian decision methods provides the first complete, integrated answer to needs in technical and business team decision-making.