Modeling Teams of Specialists

Katia P. Sycara
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Michael Lewis
Department of Information Science
University of Pittsburgh
Pittsburgh, PA. 15260

Abstract

Often, complex decisions must be made by a group of specialists rather than a single decision maker. To make an effective decision, the combination of the group's expertise must be brought to bear on the situation. Fusing expertise where individuals have very detailed knowledge in their own areas and much weaker understanding of others is characterized by difficulties: (1) agents cannot communicate their expertise in an intelligible way to non-experts because of differences in vocabulary and conceptual content, (2) the process allows for incorrect inferences, and (3) no one knows what anyone else needs to know. Measures and models developed for single agent decision making do not address the complications posed by the fusion of heterogeneous expertise nor the resource limitations which must be considered in evaluating group decisions. We propose the Team of Specialists (TOS) group decision model which partitions agents' knowledge into expert and naive models to address these inadequacies. TOE models group decision making as a process of model refinement linking communications among agents to modifications of the naive portions of their models. The TOE model gives both a normative reference for evaluating group decisions and a characterization of the process by which they evolve. The model of process and criterion for ordering outcomes provided by TOE are the essential prerequisites for devising decision aids.

1. Introduction

Decision making in a team of specialists is fraught with difficulties [1,5]. A major difficulty is that different specialists lack (1) a shared language for communication and (2) shared perceptions of the task. Evidence supporting this assumption comes from a variety of sources. Case studies of decision making in organizations (e.g., Bond's [2] Lockheed study of aircraft design) have found that specialists do not understand the details of each other's models and language, but through cooperation and interaction are somehow able to produce designs of very complex artifacts, such as aircraft. Bond [3] describes such organizational cooperation as occurring through a series of commitments. Our interest lies in investigating the process through which these commitments are formed.

The focus of the Team of Specialists (TOS) model is to take advantage of the expertise available to the group in evaluating a decision without having to explicitly see or communicate that knowledge. In TOS, the implicit hypothesis is that it is more efficient to share knowledge via evaluation than share the knowledge itself. Issues of consistency, and resolution of factuality are focused on the decision at hand rather than on the logical consistency of the agents' knowledge. Agreement on the ordering of the most favored alternative provides sufficient inter-agent consistency for the model to proceed.

The present model of decision making by a team of specialists extends earlier work [6,7,8,9] characterizing group decision making as a negotiation and refinement process. A team of specialists is considered to be a group of individuals with common goals, each with highly specialized knowledge in a particular area but with less precise knowledge of other areas. The group's decision problem is to arrive at the best decision alternative that their joint expertise allows. In a group decision process of this sort it is neither feasible nor desirable to form a common model incorporating all of the group's expertise. Instead, a normative group decision needs to be one generated by a process that determines a decision as good or better than other possible interactions of comparable length. Because the decisions of one agent impact the decisions of another, a computational model of cooperating specialists cannot simply model each agent. Rather it must augment an agent's problem solving process by incorporating interactions and decision-coordination with the other agents.

2. Aiding Group Decisions

DeSanctis and Gallupe [4] identify three levels of group decision support ranging from shared communications (Level I) to modeling tools such as

Proceedings of the Twenty-Seventh Annual Hawaii International Conference on System Sciences, 1994

spreadsheets or risk analysis (Level II) to computer intervention (such as enforcing Robert's Rules of Order) in group process (Level III). Six years later work in CSCW, groupware, and group decision support retains much this same character with applications such as decision rooms, shared editors, and design history systems predominating. While the modeling tools shared in Level II decision support bear a striking resemblance to normative decision models for single agents they cannot play the same role because heterogeneous groups lack such common models. This inability to specify the decision a heterogeneous group *should* make prevents researchers from following the fruitful normative-descriptive approach which has taught us most of what we know about individual decision making.

The objective of the TOS model is to provide such a normative description of group decision making as a standard for assessing bias (experimental tool) and a reference for use in avoiding bias (decision aid). Decision making by a diverse group of specialists is too complex to be described by normative mathematical theories of operations research, decision theory or game theory. For this reason, our normative model relies on inference and model refinement rather than estimation and parameter change to follow the progress of group decision making. Our characterization of model refinement is not intended as a literal account of cognitive processing but rather as a modeling framework for organizing the knowledge, beliefs, and assumptions agents bring to their problem. TOS does not presume that agents consciously possess or modify mental models of the sort we describe any more than multiattribute utility theory presumes that subjects consciously combine complicated weighted functions to express their preferences. The crucial features of group decision making captured by the TOS model are: distributed expertise, the influence of common sense in expert judgements, and the incremental aggregation of common evaluations.

The processes of negotiation and iterative refinement of solutions modeled in the PERSUADER [6,7] provide our framework for describing the extended discussions and exchange of ideas necessary to coordinate decisions among specialists. The PERSUADER, acts as a mediator between disagreeing parties by modeling both agents and searching for favorable inter agent tradeoffs. TOS extends this approach to more than two agents in non adversarial problem solving. In TOS group decisions are improved through the elimination of unfavored alternatives and the progressive discovery of favorable inter-agent tradeoffs.

Although TOS is a model of the influence of interactions on individual agents, its presumption that agents possess a common naive understanding and incorporate changes through refinements to this common understanding provides a kernel for PERSUADER like aiding. By refining a naive model in parallel with its agents, an intelligent aid could maintain an aggregated decision model for identifying favorable inter-agent interactions. Aiding on this basis preserves the defining features of TOS decision making: 1) agents' knowledge is too extensive to be efficiently shared and 2) decision aiding and evaluation reference the group's current state of knowledge rather than some unachievable aggregated state of total knowledge.

3. Agent Models

Each agent:

Has a model of his individual unique expertise, called the "*expert model*", characterized by detailed knowledge about some particular aspect of the task.

Has a naive understanding of aspects of the problem outside of his area of expertise, called the "*naive model*". The naive model characterizes weak commonly held beliefs such as, "the more expensive a material is, the more durable it will be."

Develops through interaction with other agents a more comprehensive model of the problem at hand which incorporates elements of others' expertise and defines a common vocabulary that the two agents can use to communicate in an intelligible way.

For a group to make a common decision, the potentially different models formed through interactions must converge to a coherent[1] set of evaluations of the global problem. Shared models evolve from the naive models by incremental modifications which make them conform to *justifications* and evaluations supplied by other agents. The naive model supplies both communication and inference capabilities by providing a common language, an inference mechanism for underdetermined evaluations, and an initial model for modification through communications.

4. Representation of Agent Models

[1]Coherence requires consistency in private knowledge among agents. If our experts included both Keynesian and supply side economists, for example, a decision on tax policy in accord with the full expertise of the group would be incoherent and agent evaluations could not be expected to converge.

In this section we detail the representations that allow the expression of the expert, naive and partially refined models. The overall model of an agent consists of a public *description space* used to characterize decision alternatives, a public *decision space* containing evaluative variables, and *naive* or *private* mappings linking descriptions to decisions. Examples of a decision alternative is a vector of description variables that have been instantiated to particular values in their domain. Decision variables are "aggregate" variables in that they refer to a decision alternative rather than a description variable (attribute). Agents are assumed to have a single naive common sense model of relations among decision variables, such as "selling price, manufacturing cost, and quality covary", which are represented in the decision space.

In more detail, the basic parts of the model are as follows:

1. *public and well defined description space* that consists of attributes of a decision alternative described in the public language l_0. In design, these attributes are attributes of the artifact, such as its dimensions, material, components, connections among components etc. Each description variable has a domain of values that respect appropriate constraints. For a turbine blade, a description consists of the vector of attributes/variables [root-radius, blade-length], and a design alternative could be [root-radius=35 in, blade-length=76 in].

2. *public and well defined decision space* modeled by influences among the decision variables represented as a directed acyclic graph. The language that describes the decision space is L_0. Edges of the graph linking two decision variables represent the relationship between them in terms of how one affects (positively or negatively) the achievement of the other. For example, in aircraft design, aerodynamic efficiency positively affects lower operation costs.

An influence $V_i^+ \rightarrow V_j$ means that V_j increases with increasing V_i. $V_l^- \rightarrow V_k$ means that V_k decreases with increasing V_l. Examples of influences are structural-soundness, of$^+$ $\rightarrow$reliability price$^-$ $\rightarrow$saleability. In the current model we assume that relations among decision variables are directly proportional in their ranges so that a change in one will effect an increase or decrease of a corresponding size in another. As agent models are refined the (naive) influence of direct relations among decision variables decreases and is replaced by indirect relations through the joint influence of attributes in determining their values. For example, if quality and price were initially related only through the decision space they might later come to be partially related through the attribute "material" which has direct relations to both quality and cost. This phasing out of initial naive approximations in favor of more precise determination by attributes is what we mean by *refinement*. The model requires that values of decision variables be completely determined and that determination by attributes take precedence over determination by other decision variables. This allows the influences represented in the decision space to serve as an error term to the refinement process adjusting their weights as needed to preserve full determination. In a refined agent model in which no direct influences are left in the decision space, these effects could only be recovered as a relation among descriptions. Refinement involving closer approximation of $V_{i,j}$ and $B_{i,j}$ may continue even after every influence has been eliminated from the decision space.

For the naive model, we assume that the relations among decision variables are linear and that multiple influences contribute equally to determining their values. These assumptions simplify propagation of new decision variable values and estimation of whether or not a new proposal increases profits. Figure 1 shows three such naive relations, quality$^+$ $\rightarrow$production-cost, quality$^+$ $\rightarrow$unit-sales and unit-cost$^-$ $\rightarrow$unit-sales. For example, if a new proposal increases the product's quality by x%, then (assuming the relations in figure 1 and production-cost$^+$ $\rightarrow$unit-cost) unit-cost must also increase by x% and unit-sales must also decrease by the same percentage. Since unit-sales$^+$ $\rightarrow$profit, in this naive estimation, the new proposal leaves profits the same as the previous proposal. The initial naive relations are refined and updated as a result of the group's interactions. For example, the quality agent may say that while the change from stamping to machining increases quality by x%, unit-cost will decrease only by {x/2}%. This statement establishes a new relation between the attribute, manufacturing process, and the decision variables quality and production-cost replacing the previous naive relation between the two decision variables. The value of a decision variable is a function of one or more attributes in the description space and the influences of other decision variables. For example structural-soundness[root-radius=35, blade-length=76] = 8.5 (on an arbitrary scale 0 to 10).

In the current version of the model, we make the assumption that the value of a decision variable represents the utility of the decision alternative with respect to the particular decision variable. The value that an agent assigns to a decision variable for a particular alternative may depend on its private knowledge. This representation provides both the multi-attribute utility values used to

Figure 1. Naive relations

evaluate alternatives and a public representation relating decision variables used to determine the naive inferences needed to refine agent models.

3. *specialized/expert "black-box" knowledge* modeled as private functions of arbitrary complexity relating attributes to decision variables. The private knowledge of each agent α_i is expressed in its private language *lamda*$_i$. In engineering design, specialized knowledge can be represented in terms of qualitative and quantitative relations and equations. This choice allows multiple attributes to jointly influence multiple decision variables creating "hidden paths" not represented in the public influence diagram.

4. *naive mappings* between description and decision variables modeled as publically defined functions $v_{i,k}$ (relating attribute a_k to decision variable V_i) and expressed in the language λ_0 Λ^{λ_0} for non-expert agents. An attribute can be of relevance to more than one decision variable and the domain of a decision variable is a vector of more than one attributes. The relevance of an attribute, a_j, to a decision variable V_i is its "contribution" and is expressed by a weight coefficient $B_{i,j}$. Shared models are formed through the refinement of this naive knowledge. Refinements are plausible inferences defined as changes to naive portions of an agent's model. Refinements could change the coefficients $B_{i,k}$, the functions $v_{i,k}$ and through them the decision variables to which an attribute "contributes".

5. *expert mappings* between description and decision variables. The form of these mappings is determined by the agent's expert knowledge and is expressed in the agent's private language λ_i.

Figure 2 shows the architecture of the mental model of an agent. To calculate the value of a decision variable that is not within its area of expertise, an agent uses the publically known weighted sum of the functions $v_{i,k}$. Since refinements are limited to the naive portions of agents' models and result from public communications, these relations remain public under refinement. To calculate a decision variable within its domain of expertise, an agent uses its private expert knowledge. In the figure, the function ϕ expresses the expert mapping of the agent's private knowledge to decision variable V_3. Each agent's model is similar to that shown in figure 2, except that the "black-box" private knowledge would involve different attributes and decision variables.

Therefore, an agent's expert model consists of (a) the collection of qualitative and quantitative relations within its "black-box" along with (b) functions, such as ϕ that allow expert mappings between the "black-box" and decision variables. An agent's naive model consists of (a) the decision and description variables, (b) naive mappings within the decision space, and (c) naive mappings between description and decision variables.

In naive models, relations and contributions may be either indifferent (same contribution across attributes or same relation across attribute values) or ordered with respect to their contribution or relation to the decision variables. For example, in a naive model of the turbine blade root-radius and blade-length contribute equally to all decision variables, such as structural-soundness, cost etc. As another example, a naive model might hold unit sales which contributes to profit to be proportional to quality and inversely proportional to price with price proportional to quality in its decision space. If two materials, plastic and steel, were ordered with respect to quality, then this model would be indifferent to the choice because the contribution to profit of choosing steel (via quality and unit sales) is balanced by the adverse impact of quality (via price and unit sales) on profit. This sketchy knowledge of alternatives and their evaluation is shared (outside of individual areas of specialization) by all members of the group. The indifference

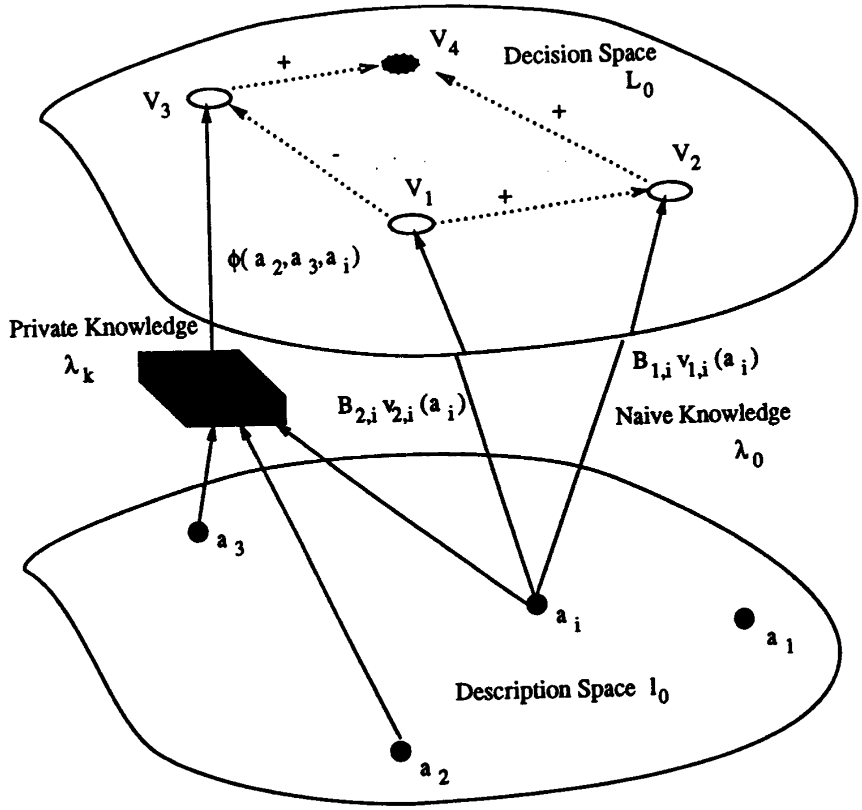

Figure 2. Agent Models

of the naive model expresses the uncertainty of agents outside of their areas of expertise while still characterizing the common sense world knowledge they bring to the situation regardless of their specialization.

6. Communications

Each agent uses its private language λ_i to generate or evaluate proposals. However, communication among agents is restricted to the public languages L_0, l_0, and λ_0. This restriction allows agents to communicate their expertise only within the context of the group's problem and in terms that are intelligible to other agents in the group. So, although each agent's expertise is private to it, the common vocabulary is the medium for making public relevant portions or results of the expertise in the form of suggestions, justifications, and objections. In the turbine blade design example, terms such as Swirl-Coefficient and Axial-Velocity (see Figure 3) are private to the structural engineering agent. Terms such as Structural-Soundness and Blade-Efficiency belong to the common public vocabulary and are used for intelligible communication among the design agents. For example, the marketing agent understands the concepts of Blade-Efficiency and Structural-Soundness and how they relate to marketability, a decision variable within his area of expertise. In Figure 3, the shaded portions indicate the private expertise of the aerodynamics and structural agents, whereas the unshaded portion indicates terms to express goals and issues in the public vocabulary. Communications are the means by which various parts of an agent's model get updated. A structural engineer justifying a proposal to thicken the turbine blade, for example, might report that increasing the root radius of a turbine blade by 5 inches would double its structural soundness. The engineering models behind this observation would remain private but the entire group could now benefit from his more valid estimate of the relation between a description variable, root radius, and a decision variable, structural soundness. Because other agents lack an understanding of the engineering judgement which underlies this pronouncement, they must continue to rely on naive inference (proportionality) to evaluate other alternatives. For example, another agent might infer that "if five inches doubles

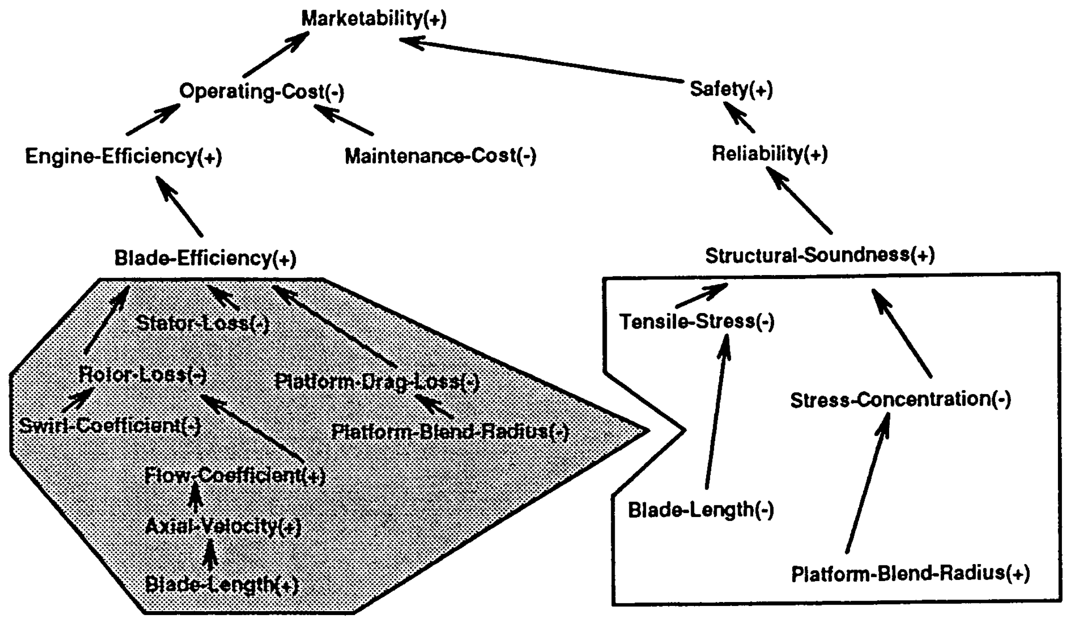

Figure 1. Public and private models

structural soundness then ten inches should quadruple it." Because a blade is a relatively uniform solid object, however, structural soundness is more closely proportional to cross sectional area (the square of its radius) making this an underestimate. This new misconception could only be rectified by further communications from the structural engineer. Because all agents' judgements are flawed in this way, they can only arrive at decisions reflecting their joint expertise through cycles of communication and updates to their individual naive models.

Information communicated by agents is classified as *evaluation* or *justification*. An expression of preference among alternatives is an evaluation. An expression relating attribute values to decision variables is a justification. Communications will often contain information of both sorts. The statement: "We should not change from alloy to composite materials because the manufacturing cost would be too high," for example, consists of the evaluation[2]: profit(design-composite) < profit(design-alloy) and the justification: manufacturing-cost(design-composite) > manufacturing-cost(design-alloy). Note that the evaluation is re-expressed in terms of the implicit and publically known relation between manufacturing cost and profit. This is done to clarify the status of evaluation as an ordering of alternatives with respect to the decision criterion rather than a relation between attributes and a particular decision variable. As suggested by the example, the distinction between these forms of information lies in the relation referenced rather than the apparent phrasing. The statement: "We should choose the highest quality design because profits depend on quality" contains no justification because it does not relate attributes to decision variables. This communication conveys, instead, the agent's "evaluation" of the contribution of quality to profits. In TOS, the assumption of a common naive model makes this statement vacuous (devoid of information) since the "evaluation" is already known to other agents.

Factual statements such as: "Selection of steel instead of plastic would increase material costs by 30%" are justifications because they relate attribute and decision variable values even though they lack an evaluation to be "justified". To make an evaluation between the alternatives design-steel and design-plastic, as a result of hearing this statement, an agent must perform the inference material-cost(steel) > material-cost(plastic), material-cost and manufacturing-cost are related in a positive manner (+), manufacturing-cost and profits are related in an inverse (-) manner, hence profit(design-steel) < profit(design-plastic).

Justifications may be expressed in any of four forms, each of which conveys a different type of information about the relation between attributes, private knowledge, and decision variables.

[2] Note that since profit is the decision criterion, the statement profit(design-composite) > profit(design-alloy) is equivalent to utility(design-composite)> utility(design-alloy).

These forms are: 1. *value ordering*- The communication expresses an ordering on values of a decision variable as a function of attribute values. Value orderings have the form: relation(attribute_value-1,attribute_value-2) -> Order(v(attribute_value-1),v(attribute_value-2)), where $v_{i,j}$ is the function relating values of attribute a_i to decision variable V_j. The earlier statement: "We should not change from alloy to composite materials because the manufacturing cost would be too high," is an instance of value ordering because it orders the nominal values "alloy" and "composite" of the attribute, material, by their relation to the decision variable, manufacturing cost. Value ordering communications refine agent models by modifying function $v_{i,j}$.

2. *contribution determining*- The communication expresses the contribution of an attribute to a decision variable. Contribution statements may be either absolute or relative. The most common absolute contribution statements assert that for values under consideration an attribute makes no contribution to a particular decision variable. A statement of this sort would be: "The choice between composite and alloy materials will not affect reliability." An example of a relative contribution justification would be: "Choice of wheel attachments is more important to the reliability of a tricycle than choice of frame material." Contribution justifications do not provide information about the relation between values of attributes and decision variables but instead characterize the extent of the attribute's contribution. Contribution determining justifications refine agent models by modifying the weight $B_{i,j}$ which determines the contribution of attribute variable, a_i, to decision variable, V_j.

3. *value determining*- The communication expresses both the contribution and the relation between attribute values and values of a decision variable. Value determining expressions have the form: relation(attribute_value-1,attribute_value-2) -> V(attribute_value-1)=KV(attribute_value-2) where K is the ratio between the two aggregate values. The earlier statement: "Selection of steel instead of plastic would increase material costs by 30%" is an example of a value determining communication: material-cost(steel)=1.3material-cost(plastic). Value determining communications may affect either the function, $v_{i,j}$, its weight, $B_{i,j}$, or both. The refinement which occurs depends on the prior values of $B_{i,j}$, the range of $v_{i,j}$, and the range and weights of other attributes with non-zero contributions to V_i. These adjustments are made according to the calculus of propagating values of decision variables based on the linearity relations and their refinements, subject to the constraints:

$$\max(V_j) = \sum_k B_{k,j} max(v_{j,k} a_k))$$

$C = \sum_k B_{k,j}$ where C is a constant reflecting the relative scaling of weights

4. *conjunct labeling* - Justifications may involve terms from an agent's private language. In these situations the public version of the term from the private language serves to "label" a relation involving multiple attributes and decision variables expressed in the public language. A communication of this sort about a proposed tricycle design might be: "We should change to a ball-bearing and race for the headset (attribute-1) if the frame is plastic (attribute-2) and the bearing is metal (attribute-3) because otherwise torsion will cause the bearing to slip and weaken the frame (private knowledge) making the tricycle unreliable (decision variable)". The justification in this case is the public expression: reliability(headset=ball-bearing-and-race $\wedge$ frame=plastic $\wedge$ bearing=metal) < reliability(headset=bearing $\wedge$ frame=plastic $\wedge$ bearing=metal) which is labeled by the private term "torsion". Labeled justifications are consistent with our contention that it is neither feasible nor desirable for specialists to develop detailed models of one another's expertise. In this example the technical meaning of the word "torsion" and its use to describe forces affecting mechanical devices remains private to the communicating agent. In the public communication the term "torsion" serves only to label a public expression relating these three attribute values to the decision variable, reliability. Conjunct labeling reifies private knowledge by expressing relations between multiple attributes/decision variables in public form for the values of some particular alternative. Agent models are refined by adding this new composite attribute to their description space.

Justifications are further classified as expert or naive. An expert justification is a justification which involves decision variables within an agent's domain of expertise. A rejection of a proposed turbine blade design by the structural engineer which referenced a change in the decision variable, structural soundness, for example, would be classified as an expert justification. If the structural engineer's justification for rejection had involved the decision variable, manufacturing cost, instead, it would have been classified as a naive justification.

Refinements are naive in much the same way as the influences in the decision space they replace. The model again assumes proportionality in the absence of more precise specification and uses the determination of weights and values in previously considered alternatives to anchor the evolving model. Functions $v_{i,j}$ and labeled conjuncts are learned in a piece-wise linear fashion and are revised to maintain fit to previous alternatives when contribution weights are adjusted. Orderings of alternatives enter the model as constraints which are converted to values in accordance with the proportionality and indifference assumptions of the model. These values, like the naive ones they replace, are treated as "second class citizens" by the refinement which anchors the evolving model to directly determined values. Because values of decision variables are computable as weighted sums of functions of attributes, evaluation of alternatives is always possible.

The decision making process can be seen as hill-climbing where the group starts with an initially proposed (perhaps randomly generated) decision alternative and iteratively adapts it to arrive at an improving best decision. This search is satisficing rather than optimizing. Because the distributed evaluation function is not fixed, both the evaluation of alternatives and the choice of alternatives to evaluate are determined by the history of interactions. Whether or not justifications are exhaustively exchanged (full expertise), evaluation depends jointly on alternatives already considered, degree of refinement of naive models, and the alternative itself. Within the model it is possible for a shallowly explored alternative to resurface later for evaluation at greater depth. This sort of behavior is not backtracking in a strict sense because the alternative has changed with respect to the agents' evaluations. As a conventional search problem, the model would require the power set of possible histories and alternatives as its search space. There is no guarantee that a group would not stop at a local maxima or even stop short of it if evaluation were sufficiently favorable. The problem addressed by this model is not how to escape local maxima but how to detect the hills which may be invisible to individual agents.

In the TOS model, we assume that maximizing profits is the *decision criterion.* Each current decision alternative is taken as a baseline that must be improved by the next acceptable proposal (i.e. inferior proposals will be rejected). Each new proposed decision alternative contains at least one change in an attribute value, and possibly results in new values for a subset of the decision variables, or possible update of various mappings between variables in the description and decision space. Under the restriction that only expert justifications lead to refinement, communications of sufficient length will cause all agent models to converge to a single evaluation for any particular alternative. Convergence follows from the anchoring of refinements to attribute and decision variable values associated with an alternative. Model refinement with respect to this alternative is simply an inefficient mechanism for exchanging these values. Even absolute convergence for a single alternative contributes little to improving decision making which requires agreement in the ordering of alternatives rather than precise agreement on the utility of any particular alternative. Our focus is not on the role of communication per se in providing accurate evaluations of particular alternatives but rather on the role model revision plays in directing search. As agent models are refined, their evaluative ordering on alternatives will change causing alternatives previously dismissed as infeasible to become practical while causing undesirable ones to become newly attractive, thus allowing proposing new alternatives which increase the value of the decision criterion. These revisions and reevaluations occur within individual agent models taking advantage of individual expertise and an improving approximation of problem relevant aspects of other agents' models. An exact control procedure for using agent models to direct search is not dictated by TOS. Plausible choices include bidding among agents for proposing the next alternative (if private reevaluation results in many agents' estimation of differential increases in the value of the decision criterion), round robin selection of the next proposer, or joint partial specification of the next alternative by multiple agents, each of which specifies values for a subset of description variables.

5. An Example

Consider the decision situation for a team of specialists in a manufacturing enterprise tasked with concurrently engineering the design of a tricycle. The team objective is to arrive at a tricycle design that will maximize profits, under certain assumptions relating design attributes to cost, price, and ease of selling the tricycle. The group's goal may be expressed as:

profit = unit_sales x (unit_price - unit_cost)

Tricycle sales, price and cost can be expressed as functions of high level attributes, such as tricycle, performance, style, durability, ease of use, reliability, structural soundness etc. Since there is no precise mathematical model of design evaluation (there are too many variables that interact in nonlinear and unpredictable ways) or design saleability, the group's goal of "optimizing the design" gets operationalized to "using profits as a decision criterion, find a design acceptable to all concerned agents". This is as optimal a decision as the group can give since it is a decision that takes into consideration the fused expertise of the group. The decision problem then is to evaluate alternative designs, negotiate on suggestions for design modifications and arrive at a design agreeable to all.

Let us suppose that there are three agents involved: a designer, a manufacturing agent and a sales agent. In his expert model, the designer knows the precise (true) relations among design attributes and decision variables such as performance, durability etc. For example, a designer's expert model predicts that high grade plastic makes the tricycle lighter thus leading to higher performance, whereas heavy steel tubing leads to lower performance. In terms of strength of the tricycle frame, braced and welded frame leads to higher frame strength, bolted the next highest, and integral the lowest. In terms of reliability, using cotter pins and caps to hold the rear wheels and pedals of a tricycle together results in much higher reliability (the designer's model may include precise equations or empirical results from which precise numbers could be derived), than using press-on caps, since press on caps are likely to start falling off after a short time of tricycle use.

Let us concentrate on the design choice of cotter pins versus press on caps. Before the group meeting, each agent has different evaluations of the two designs. Suppose that the designer's expert knowledge rates the cotter pin design twice as reliable as the press-on caps design. On the other hand, the precise relations between cost and other decision variables are not in the expert model of the designer. The designer's naive model considers that sales and cost are linear with reliability. This leads him to infer: (1) the cost of a cotter pin design will be double the cost of a press-on cap design, (2) the sales for a cotter pin design will be double the sales for a press-on design, and (3) he should be indifferent to the design choice. Similarly, the manufacturing engineer knows (from his expert knowledge) that drilling round stock to make the hole for the cotter pins is 3% more expensive than fitting the press-on caps. The expert model of the manufacturing agent does not contain precise knowledge of the relation between press on caps (or cotter pins) and reliability (or other high level design attributes). His naive model considers cost and sales linear with reliability. This leads him to infer: (1) a cotter design is 3% more expensive than a press-on design, (2) a cotter design will sell 3% more than a press on design and (3) he should be indifferent to the choice. The expert model of the sales agent predicts that *appearance of ruggedness* (for a device used by children, and which will suffer a lot of wear and tear) multiplies sales by a factor of three. The sales agent has no expert knowledge about either the relative cost for manufacturing cotter pins versus press-on caps or about relative reliability of the two designs. Thus, he is initially indifferent to the choice.

The group meets. Drawings for the two designs are displayed and discussed. The group interaction over the design choice of press-on caps versus cotters may proceed as follows:

Manufacturing Agent: Do we want to use press-on caps or cotters for the wheels and pedals? The drill press operations will add another 3% to manufacturing costs. (This is a value determining communication). [The design agent learns that cotter pin cost is only 3% more rather than double. He, therefore, updates the (naive) relations between fasteners (attribute relevant to decision variables), cost and reliability. On the other hand, the naive relation of linearity between reliability and sales predicts that the cotter design should almost double profits, thus furthering the group's goal of profit maximization and making the design agent strongly prefer the cotter design. The sales agent learns the relative cost of cotters, cost(cotters)=1.03cost(press-ons) and updates the cost decision variable].

Design Agent: In that case, I think we should use cotters. The press-on caps are likely to start falling off after 6 months to a year while the cotter pins will hold the wheels on for 10 years. [The manufacturing agent now learns that the cotter design is much more reliable, reliability(cotters)>10reliability(press-ons), and updates his model so that cost is now indirectly related to reliability through the fastener attribute. The sales agent learns that cotter design is much more reliable than press-on design and updates the relevant part of his naive model. However, his expert model tells him that it is the appearance of ruggedness that sells the product. By looking at the designs, the sales agent finds out that a buyer cannot see whether a cotter or a press-on has been used].

Sales Agent: I don't think we should use cotters. A buyer can't see that there is a cotter under the cap and therefore it has no effect on appearance. On the other hand, cotters are 3% more expensive.

The design and manufacturing agents substitute this direct relation between the fastener attribute and sales for the previous relation between decision variables and now agree because with this modification, their shared model predicts that the increased production cost associated with cotters will be detrimental to the goal of increasing profits because there is no offsetting influence of this form of reliability on unit sales.

References

1. Alder, M.R. and Simoudis, E. "Integrating distributed expertise". *Proceedings of the 10th International Workshop on DAI.* Banderra,TX (1990).

2. Bond, A. and Ricci, R. "Cooperation in aircraft design". *Proceedings of the MIT-JSME Workshop on Cooperative Product Development.* Cambridge, MA (1989).

3. Bond, A. "PROJECTS: A normative model of collaboration in organizations". *Proceedings of the 10th International Workshop on DAI.* Banderra, TX (1990).

4. DeSanctis, G. and Gallupe, B. "A foundation for the study of group decision support systems". *Management Science*, 33,5 (1987), 589-609.

5. Lander, S. and Lesser, V. "A framework for cooperative problem-solving among knowledge-based systems". *Proceedings of the MIT-JSME Workshop on Cooperative Product Development.* Cambridge, MA (1989).

6. Sycara, K. "Multi-agent compromise via negotiation". In M. Huhns and L. Gasser (Ed), *Distributed Artificial Intelligence, Vol II.* Pittman Publishing Ltd. and Morgan Kaufmann (1989).

7. Sycara, K. "Problem restructuring in negotiation". *Management Science* 37,10 (1991), 1248-1268.

8. Sycara, K. and Lewis, M. "Modeling group decision making and negotiation in concurrent product design". *Systems Automation: Research and Applications* 1,3 (1991), 217-238.

9. Sycara, K. and Lewis, M. "Forming shared mental models". *Proceedings of the 13th Annual Conference of the Cognitive Science Society.* Chicago, IL (1991), 400-405.