Combining Opinions About the Order of Rule Execution

Jeffrey A. Barnett
Northrop Research and Technology Center
One Research Park
Palos Verdes Peninsula, CA 90274
jbarnett@nrtc.northrop.com

Abstract

How should opinions of control knowledge sources be represented and combined? These issues are addressed for the case where control knowledge is used to form an agenda, i.e., a proposed knowledge source execution order. A formal model is developed in the Dempster/Shafer belief calculus and computational problems are discussed as well. The model is applicable to many other problems where it is desired to order a set of candidates using a knowledge-based approach.

Introduction

Deciding on the order to do things is one of the most important activities performed by an intelligent system. These decisions influence the amount of problem-solving resources utilized and determine the coherence and explainability of system behavior. The decisions are made by control knowledge and it is this knowledge that is responsible for guiding a system's search for problem solutions.

Search is a prevalent problem-solving paradigm used by many intelligent systems (Newell & Simon 1976). Some relevant piece of knowledge is selected by the control knowledge and applied within the current problem-solving state. This step can modify that state and make different pieces of the knowledge base applicable. The selection and application cycle is repeated until the system's termination condition is met.

Without the guidance of proper control knowledge, search wanders aimlessly through the solution space until either resources are exhausted or an answer is uncovered (Pearl 1984 & Nilsson 1980). This is not acceptable in large domains because the chance of stumbling on a solution is very small. Also, it is unlikely that a proposed solution could be defended: it is difficult or impossible to explain why alternative paths were discarded or not explored at all if blind search were substituted for the use of control knowledge.

If the resources consumed by control knowledge did not count as part of the total problem-solving resources used by the system, it would be optimal to determine the piece of knowledge to apply at the beginning of each system cycle. However, these resources do count (Barnett 1984) and the complexity of making control decisions can easily overwhelm many other aspects of system behavior.

For this reason, many systems form an *agenda*, a proposed order to execute or apply the relevant pieces of knowledge. Motivation and techniques to form agendas in rule-based systems are described by (Davis 1980a & 1980b).

Summary

A formal model is developed in the Dempster/Shafer belief calculus (Shafer 1976) so that the opinions of control knowledge sources can be represented and combined in order to form agendas. Contradictory opinions and preference cycles (loops) are dealt with in a straightforward way.

The model is applicable to many problems where it is desired to order a set of candidates using a knowledge-based approach. However, for the sake of concreteness, the discussion is presented as if the problem were that of forming an execution-order agenda for the rule instantiations in a system's conflict set and some flexibility in execution order is assumed.

The model's objective function for agendas is **Pl**, the Dempster/Shafer plausibility measure. However, finding the agenda that maximizes **Pl** is equivalent to solving the weighted feedback edge-set problem which is known to be NP-complete.

Since the same underlying problem is likely to occur in other formulations of weighted voting schemes, an approximation technique is desired. An algorithm, empirically shown to be reasonably accurate and efficient, is described.

The Model

A model to represent and combine the opinions of control knowledge sources is developed in the Dempster/Shafer belief calculus. (The appendix provides a brief introduction to the concepts used below.) The calculus is well-suited to this task because opinions about ordering are naturally expressed as preferences for subsets of the set of the possible agendas.

In the model, primitive opinions are weighted preferences on the execution order of pairs of rule instances. These pairwise preferences are represented by simple support functions. Complex opinions are expressed as sets of primitive opinions and combined with Dempster's rule.

Representing Opinions

Let $R = \{r_1 \ldots r_n\}$ be the collection of rules selected by the retrieval and filtering mechanism of an expert system. Define Θ as the set of possible execution agendas, i.e., the set of permutations of R. Thus, if $R = \{a\ b\ c\}$, then $\Theta = \{abc\ acb\ bac\ bca\ cab\ cba\}$ and the problem of picking an agenda for R is to select the best $\pi \in \Theta$. Therefore, opinions about ordering R are preferences for particular elements or subsets of Θ because these elements and subsets encode order relations among the elements of R.

In the proposed model, a *primitive opinion* about the best ordering of R is a pairwise preference written as $a \rightarrow b[w]$, where a and b are elements of R. This is an opinion that a should execute before b and $w \in [0, 1]$ is the strength of that preference.

A primitive opinion is represented by a simple support function. The degree of support is w and the focus is the subset of Θ for which the pairwise preference holds. For example, if $R = \{a\ b\ c\}$, the opinion $a \rightarrow b[w]$ is represented by a simple support function with focus $\{abc\ acb\ cab\}$.

N.B., the opinion "opposite" to $a \rightarrow b[w]$ is $b \rightarrow a[w]$. Unlike certainty factors (Shortliffe 1976), negative degrees of support are not used; rather, support is focused on complementary propositions.

It is easy to imagine other kinds of primitive opinions than those representable in this model, i.e., opinions that support subsets of Θ not allowed herein. However, as shown below, sets of pairwise preferences adequately capture the intent of many types of opinions expressed by control knowledge.

Combining Opinions

The Dempster/Shafer belief calculus provides Dempster's rule to combine sets of belief functions into a single belief function. In particular, the opinions of control knowledge sources are combined by this rule because primitive opinions are represented by simple support functions, a specific kind of belief function.

The Best Agenda

The Dempster/Shafer belief calculus provides decision makers with the **Bel** and **Pl** functions. This section shows that maximizing **Pl** is the better criterion to select the best agenda because **Pl** is more reliable than **Bel** in discriminating among alternatives.

Assume that the total set of primitive opinions expressed by the control knowledge sources is represented by $u_i \rightarrow v_i[w_i]$, where $1 \leq i \leq m$. Let $\pi \in \Theta$ be an agenda and define $\sigma_i(\pi)$ to be satisfied if and only if π is *compatible* with $u_i \rightarrow v_i[w_i]$, i.e., if u_i appears before v_i in π. Then **Pl** is computed by

$$\mathbf{Pl}(\pi) \quad = \quad K \prod_{\neg\sigma_i(\pi)} (1 - w_i), \tag{1}$$

where K is strictly positive and independent of π (see the appendix). Thus, $\mathbf{Pl}(\pi) \neq 0$ unless there is a $u_i \rightarrow v_i[1]$, where $\neg\sigma_i(\pi)$. However, the necessary and sufficient condition that $\mathbf{Bel}(\pi) \neq 0$, where $\pi = r_1 \ldots r_k$, is much stronger:

1. A primitive opinion, $r_i \rightarrow r_{i+1}[w_i]$, where $w_i \neq 0$, exists for each $1 \leq i < k$.
2. No primitive opinion of the form $r_i \rightarrow r_j[1]$ exists for any $1 \leq j < i \leq k$.

Therefore, unless the set of primitive opinions is relatively large, **Bel** will be zero for most or all of the $\pi \in \Theta$. In fact, if there is a $r \in R$ such that no primitive opinion references r, then **Bel** will be zero for every agenda.

Hence, the best agenda is defined, by this model, to be the $\pi \in \Theta$ that maximizes $\mathbf{Pl}(\pi)$ because **Pl** is more stable and reliable at discriminating agenda merit than is **Bel**.

Complex Opinions

The opinions of control knowledge sources are often derived from general knowledge and knowledge of the application rather than specific knowledge about particular rules (Davis 1980a & 1980b). For example, a typical control rule in an investment domain is,

> *If the economy is fluctuating, prefer the rules that analyze risk.*

This control rule is interpreted to prefer to execute investment rules that analyze risk before those that do not given that the system has (can) deduce that the economy is fluctuating. Another example of a control rule in the same domain is,

> *If bonds are being considered, prefer the rules that recommend investments with higher Standard and Poors ratings.*

This example references a ranking (the Standard and Poors index) and groups some of the investment rules (those that recommend bonds) so that the groups can inherit the ranking. Then it prefers to execute the rules so ranked in the specified order.

Both examples exhibit execution-order preferences that induce partial orders on R, the domain rules. A simple representation captures the intent of such control knowledge.

Let the P_i, where $1 \leq i \leq m$, be predicates with domain R and assume that the weights $s_{ij} \in [0, 1]$ are given. The partial order preference is realized as the collection of all primitive opinions of the form $a \rightarrow b[s_{ij}]$, where $a, b \in R$, $P_i(a)$, $P_j(b)$, and $i < j$. Thus, the control knowledge source is represented by

π	$\mathbf{Pl}(\pi)$ Model	Simple	Chunk
abc	1	1	1
acb	$1-s$	$1-s$	$1-s$
bac	$1-s$	$1-s$	$1-s$
bca	$(1-s)^2$	$1-s$	$1-s$
cab	$(1-s)^2$	$1-s$	$1-s$
cba	$(1-s)^3$	$(1-s)^2$	$1-s$

Figure 1: Example with 3 interpretations.

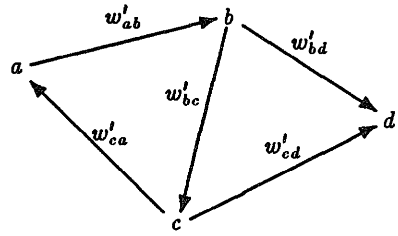

Figure 2: Example with a loop.

its P_i and s_{ij}. It should be noted that the P_i may need to access variables in the problem-solving state of the expert system, e.g., "economy fluctuation" in the first example in this section.

The next section presents an example and considers alternative realizations of the execution-order preferences of control knowledge sources.

Alternative Interpretations

Let $R = \{a\ b\ c\}$ and assume that a control knowledge source prefers that execution be in alphabetically order. In the notation of the previous section, a, b, and c are the only rules that, respectively, satisfy P_1, P_2, and P_3. Additionally, let $s_{ab} = s_{ac} = s_{bc} = s$. Thus, the opinion of this control knowledge source is expressed by $a \rightarrow b[s]$, $a \rightarrow c[s]$, and $b \rightarrow c[s]$.

With these assumptions, Figure 1 shows the value of **Pl** for each $\pi \in \Theta$. The column labeled "Model" lists the **Pl** values computed by the model. The six $\pi \in \Theta$ split into four groups because **Pl** awards values that depend on whether π agrees with 0, 1, 2, or 3 of the primitive opinions.

An alternative realization of the complex opinion that a, b, and c should execute in the stated order is to combine only the two primitive opinions $a \rightarrow b[s]$ and $b \rightarrow c[s]$, i.e., do not take the closure of the transitive preference relation. This alternative results in the **Pl** values shown in the figure under the heading "Simple".

A third alternative is to form a single simple support function that focuses only on the singleton set $\{abc\}$. The **Pl** values that result are shown in the column titled "Chunk".

A problem with the second and third interpretations is that they are less discriminating then the model's. The third alternative is the most insensitive: minor disagreements such as acb, with only b and c out of order, are awarded the same **Pl** value as total disagreements such as cba where everything is backwards.

Since there can, in general, be many knowledge sources expressing ordering opinions, it is not a good idea to employ an all-or-nothing interpretation in domains where "half a loaf is better than nothing". The approximation computation described below is applicable with both the "Model" and the "Simple" interpretations but not "Chunk" because the latter is not based on pairwise preferences.

The Optimization Problem

Finding the best agenda means finding the $\pi \in \Theta$ that maximizes $\mathbf{Pl}(\pi)$. This is shown to be the weighted feedback edge-set problem which is known to be NP-complete (Garey & Johnson 1979).

Since K in Equation 1 is strictly positive, simple algebra demonstrates that the $\pi \in \Theta$ that *minimizes*

$$\mathbf{Pl}'(\pi) \quad = \sum_{\neg\sigma_i(\pi)} w'_i, \qquad (2)$$

is the best agenda, i.e., the $\pi \in \Theta$ that *maximizes* **Pl** in Equation 1. The $w'_i = -\log(1-w_i)$ are the weights of evidence, in the terminology of (Shafer 1976), and are positive because $w_i \in [0,1]$. Hence, $\mathbf{Pl}'(\pi)$ is just a sum of a positive weight for each $u_i \rightarrow v_i[w_i]$ that is not compatible with π.

One would expect a similar formulation, with perhaps different weight semantics, for any weighted preference scheme used to determine optimal agendas—those that are incompatible with the least vote weight are valued most.

An example is shown graphically in Figure 2. The elements of R are the nodes and each directed labeled arc represents a pairwise preference, e.g., the arc labeled w'_{ab} directed from a to b represents $a \rightarrow b[w_{ab}]$. Thus, the example shows five primitive opinions that contain a preference loop between a, b, and c formed by the arcs labeled w'_{ab}, w'_{bc}, and w'_{ca}.

Let $\pi = abcd$, then $\mathbf{Pl}'(\pi) = w'_{ca}$ because the arc directed from c to a is the only one that is inconsistent with π. Since there is a cycle, every $\pi \in \Theta$ is penalized by a weight on at least one of the arcs in that cycle.

In the general case, every agenda is penalized by the weight of at least one arc in each directed cycle in the preference graph. Therefore, the best agendas are those that are compatible with the graph that remains after the least total weight has been removed on a set of arcs that cut each directed cycle. Finding a minimum-weight deletion is the weighted feedback edge-set problem.

Again, consider the example with a loop shown in Figure 2. The best you can do is to accept one of the penalties, w'_{ab}, w'_{bc}, or w'_{ca}. Assume that the minimum of the three is w'_{ab}. Then the orders $bcad$ and $bcda$ both have this minimum penalty and, hence, both maximize plausibility, i.e., both are optimal.

PROCEDURE FIND-AGENDA(w')

1. Set π to a random permutation of R.
2. Visit the elements of π in left-to-right order. Move each visited element to the position in π that minimizes $\mathbf{Pl}'(\pi)$. If any element is moved by this step, continue with the next step. Otherwise, halt and return π.
3. Visit the elements of π in right-to-left order. Move each visited element to the position in π that minimizes $\mathbf{Pl}'(\pi)$. If any element is moved by this step, continue with the previous step. Otherwise, halt and return π.

END FIND-GOOD-AGENDA;

Figure 3: Minimization algorithm.

A graphical representation of a set of primitive opinions provides a simple method to check their consistency. Let R be the graph's nodes as above. However, only include those edges that represent primitive opinions of the form $a \rightarrow b[1]$. The total set of control knowledge opinions is consistent if and only if this restricted graph is free of directed cycles.

An Approximation

Since determining the $\pi \in \Theta$ that minimizes Equation 2 is an NP-complete problem, an approximation technique is necessary if the above model is to be used for applications with more than a few agenda items. Unfortunately, a search for previous work on such approximations has not been fruitful.

Therefore, several simple approximation techniques were programmed and empirically tested, by comparison to each other, and to actual optimal results for small problems. The test cases were generated randomly and exact values computed, when possible, by exhaustive search.

Based on the empirical evidence, one algorithm appears to be efficient and accurate enough to be useful. The core of that approximation is shown in Figure 3. Given the $n \times n$ matrix, w', it is possible to find and move an element to its optimal place, relative to the current order of π, in $O(n)$ time, where $n = |R|$. Thus, each application of step 2 and step 3 is $O(n^2)$.

Steps 2 and 3 alternate because it is usually possible to prune a substantial part of a step 3 after a step 2 and vice versa. On the other hand, if either step is directly repeated, pruning is not available.

Empirical testing, with n varying from 3 to 100, showed that the entire algorithm is $O(n^z)$, where z is between 2.6 and 3. Several cases were tested with $n = 200$ and the results were compatible with this analysis. Exponent variation does not appear to depend very much on n or on the average degree of the preference graph. Rather, it is most strongly affected by the fraction of the total arcs that agree with the best agendas.

This algorithm is a local optimizations that employs a random starting point. Since different starting points can find different local optima, it may pay to run the algorithm several times and keep the best solution.

Empirical evidence suggests that the need for multiple applications increases when a relatively large fraction of the w_{ij} are zero, the distribution of non-zero w_{ij} values is flat, and therefore, the variance between different solutions tends to be largest. The variance in $\mathbf{Pl}'$ for different starting points was never observed to be more than a few percent.

Replanning

Sometimes new information becomes available while an execution agenda is being pursued. That raises questions about how to test the impact on the unexecuted part of the current agenda and how to economically reorder that portion if and when it seems appropriate to do so.

The algorithm described in the previous section has a property that makes it well-suited to address such replanning problems: every subagenda (a contiguous sequence) in π is a local optimum in the sense that it cannot be improved by moving any single element within that subagenda. In particular, the unexecuted portion of π, a tail sequence, is a local optimum unless some w'_{ij} changes, where both r_i and r_j are in the tail of π.

If a change in opinions affecting rules that have not been executed occurs, a modification of the algorithm shown in Figure 3 is used. The changes are to restrict w' to the tail and not execute step 1. In other words, start with an agenda that is probably close to reasonable rather than start from a random point.

The full algorithm with multiple starting points only needs to be considered if there is substantial change in the value of $\mathbf{Pl}'$ calculated for the tail. Empirical testing shows that this type of replanning, where only a few opinions change, is very economical. Typically, steps 2 and 3 do not iterative—one or two applications are sufficient.

It is possible to add new candidates to the tail of π or remove some that no longer belong in the conflict set. In these cases, simply place the new candidates at the end of π, appropriately grow and/or restrict w' to reflect the actual slate of candidates in the tail, and rerun the algorithm using the modified π as the starting point.

In all of these cases, the approximation algorithm commends itself as a diagnostic to determine the probable impact of the changed opinions. In addition, it can do the full replanning when it is appropriate to do so. The fact that subagendas of locally optimal agendas are themselves locally optimal, provides a measure of stability.

Discussion

The model developed here appears to have sufficient generality to solve ordering problems in many domains. Its use of the Dempster/Shafer belief calculus makes the formulation straightforward because (1) it is desired to invest belief in subsets of the possible agendas and (2) the simple support functions provided by the calculus make it easy to do so.

Finding the best agenda, as defined above, is an NP-complete problem. However, the availability of a reasonably efficient and accurate approximation partially mitigates this fact. Further, the underlying computation just determines the agenda that is incompatible with the least vote weight, an idea that seems to be very natural.

Acknowledgements

Thanks to Dan Geiger and the referees for many valuable comments that improved this paper.

Appendix: The Belief Calculus

This is a brief introduction to the Dempster/Shafer Belief Calculus. The interested reader is referred to (Shafer 1976) for more detailed explanations of the concepts that are involved.

Let Θ be a set of mutually exclusive and exhaustive possibilities and interpret the subsets of Θ as disjunctions of their elements. The function, $\mathbf{m} : 2^{\Theta} \rightarrow [0,1]$, is called a *basic probability assignment* or *mass* function if $\mathbf{m}(\emptyset) = 0$ and

$$\sum_{S \subseteq \Theta} \mathbf{m}(S) = 1.$$

In the Dempster/Shafer belief calculus, $\mathbf{m}$ plays a role similar to a density function in ordinary probability theory. The difference is that the domain is 2^{Θ}, the subsets of Θ, rather than its elements.

The *belief* function, $\mathbf{Bel} : 2^{\Theta} \rightarrow [0,1]$, and the *plausibility* function, $\mathbf{Pl} : 2^{\Theta} \rightarrow [0,1]$, are defined to play roles similar to distribution functions.

$$\begin{aligned}
\mathbf{Bel}(S) &= \sum_{T \subseteq S} \mathbf{m}(T) \\
\mathbf{Pl}(S) &= 1 - \mathbf{Bel}(\neg S) \\
&= \sum_{T \cap S \neq \emptyset} \mathbf{m}(T)
\end{aligned}$$

Thus, $\mathbf{Bel}(S) \leq \mathbf{Pl}(S)$ for all $S \subseteq \Theta$ and, therefore, $\mathbf{Bel}$ and $\mathbf{Pl}$ are sometimes referred to, respectively, as the *lower* and *upper probability measures*. Both are available to decision makers.

The calculus provides *Dempster's rule* to combine several belief functions into one. Let $m_1 \ldots m_n$ be the mass functions associated with n belief functions, then Dempster's rule defines, $\mathbf{m}$, the mass function for their combination to be $\mathbf{m}(\emptyset) = 0$ and

$$\begin{aligned}
\mathbf{m}(S) &= K \times \sum_{S_1 \cap \ldots \cap S_n = S} \prod_{1 \leq i \leq n} \mathbf{m}(S_i) \\
K^{-1} &= \sum_{S_1 \cap \ldots \cap S_n \neq \emptyset} \prod_{1 \leq i \leq n} \mathbf{m}(S_i),
\end{aligned}$$

for nonempty $S \subseteq \Theta$. This combination is defined whenever K is.

A *simple support function* is a belief function for which there is at most one $S \neq \Theta$ such that $\mathbf{m}(S) \neq 0$. The trivial simple support function is the one where $\mathbf{m}(\Theta) = 1$ and $\mathbf{m}(S) = 0$ for all $S \neq \Theta$.

Other simple support functions are parameterized by a $F \subset \Theta$ and a $w \in (0,1]$, where $\mathbf{m}(F) = w$ and $\mathbf{m}(\Theta) = 1 - w$. The subset F is called the *focus* of the simple support function and w is called its degree of support.

A simple support function is a mechanism to place committed belief on the single hypothesis represented by its focus. The remaining weight, placed directly on Θ, is uncommitted since Θ represents the universally true proposition.

Let Dempster's rule be used to combine the n simple support functions with the foci F_i and degrees of support s_i. Then (Barnett 1991) shows that

$$\mathbf{Pl}(\pi) = K \times \prod_{\substack{1 \leq i \leq n \\ \pi \notin F_i}} (1 - s_i),$$

for the combined belief function. This formula is the justification for Equation 1 above.

References

Barnett, J.A. 1991. Calculating Dempster-Shafer Plausibility. *IEEE Trans. PAMI.* In Press.

Barnett, J.A. 1984. How much is control knowledge worth: A primitive example. *Artificial Intelligence* 22:77–89.

Davis, R. 1980a. Meta-rules: reasoning about control. *Artificial Intelligence* 15:179–222.

Davis, R. 1980b. Content reference: reasoning about rules. *Artificial Intelligence* 15:223–239.

Garey, M.R., and Johnson, D.S. 1979. *Computers and Intractability.* Freeman Press.

Newell, A., and Simon, H.A. 1976. Computer science as empirical inquiry: symbols and search. *CACM* 19(3):113–126.

Nilsson, N.J. 1980. *Principals of Artificial Intelligence.* Tioga Press.

Pearl, J. 1984. *Heuristics.* Addison Wesely.

Shafer, G. 1976. *A Mathematical Theory of Evidence.* Princeton University Press.

Shortliffe, E.H. 1976. *Computer-Based Medical Consultations: MYCIN.* Elsevier.