Using Evidential Reasoning To Manage Abstraction and Uncertainty in Planning

Todd Michael Mansell[†]
Defence Science and Technology Organisation
PO Box 4331, Melbourne, 3001, Australia.
email: mansell@drea.dnd.ca [†]

Abstract

This paper describes how evidential reasoning (E-R) has been used to manage uncertainty and abstraction in U-Plan, a planning system capable of utilising uncertain and incomplete information. U-Plan uses a unique possible worlds formalism to describes likely representations of the environment at multiple abstraction levels. The possible worlds are generated from information which may have been collected from disparate sources, and expressed in diverse frames of reference. This information is accompanied by a quantitative measure of belief, that is used to weight the evidence supporting each possible world state. In support of hierarchical planning, each possible world contains a description at a number of levels of abstraction. This allows strategic decisions to be made using a coarse description of the world, and tactical decisions using a detailed description. Compatibility relations from E-R provides the mechanism to represent and manipulate information at varying levels of abstraction. The architecture described in this paper also provides for planning in dynamic environments as U-Plan can rapidly assess when changes to the world, during plan generation, may invalidate existing plans.

1 Introduction to U-Plan

Planning under uncertainty can loosely be described as the composition of a course of action that will achieve the goals presented given uncertain and/or incomplete information, and/or where the effect of the operators on the world are not known with absolute confidence. Throughout this paper the term imprecise information (or evidence) will be used to describe uncertain, incomplete, and/or inexact information.

A general planning strategy capable of generating a course of action given imprecise environmental information is currently not available. One of the reasons for this is planning in general is computationally expensive. Even for very constrained action representations, the problem is known to be NP-hard (Chapman 1987). When incorporating imprecise information about the world, the computational expense is increased, by the need to repeat the planning process for a number of possible worlds.

When planning given imprecise information about the environment it is not possible to construct one initial state that precisely and unambiguously represents the world. U-Plan uses a possible worlds representation, where the available initial information is used to construct every initial possible state (P-state) of the world. Associated with each P-state is a numerical measure of belief specifying the degree to which the evidence supports each P-state as the one that represents the true state of the world. The belief calculus used by U-Plan is evidential reasoning (Lowrance et al 1991), an extension to Dempster-Shafer theory of evidence (Shafer, 1976). The foundations of E-R are expounded in Ruspini (1986) and shown to be sound and complete. E-R has been successfully applied to a number of real world problems where uncertain, incomplete and occasionally inaccurate information is all that is available (Lowrance et al 1991) to characterise the environment.

A hierarchical approach to planning is used as it significantly reduces the search space by first planning at abstract levels, and then expanding these abstract plans into more detailed plans. At the highest abstraction level strategic decisions are made, while at the lowest levels of abstraction, tactical decisions about how best to implement the strategy, are made. In support of hierarchical planning, each P-state is described at a number of predefined abstraction levels, allowing decisions to be made using a state representation at an equivalently detailed level of abstraction.

Hierarchical planning selects an overall strategy before becoming involved with the tactical detail. U-Plan utilises a set of (predefined) goal reduction operators that encode how a planning goal is reduced by the operator's application. What results is a planning hierarchy tree where the goals are broken up into subgoals by the goal reduction operators. This allows us to first make the strategic decisions, which then guides all other decisions down to the tactical implementation of the subgoals. The reduction operators are expressed at various levels of

---

[†] Currently on exchange at:
Defence Research Establishment Atlantic
9 Grove St, Dartmouth, NS, B2Y 3Z7, Canada.

abstraction. The planning algorithm uses the same abstraction level for both operator and P-state.

U-Plan constructs a plan for one P-state at a time, the first plan being constructed for the P-state with the greatest likelihood of representing the true world. Before subsequent plans are constructed, the suitability of reapplying an existing plan to this new possible world is assessed. If an existing plan works for additional P-states then the connection is recorded. If a plan partially works for another P-state (e.g. the strategy works but some of the detail is different), then part of the plan will be adopted, and planning continues from where the plan failed. When a plan exists for every possible world, the operator order of all the plans is combined to obtain a single planning tree that branches when the operator execution order differs. At this point the ability to acquire additional knowledge is used. At each branch, a knowledge acquisition operator can be inserted to determine which action in the planning tree to carry out next.

U-Plan takes into consideration the dynamic nature of the world, i.e., the world may be altered by events other than those specifically modelled by operators. This is achieved by monitoring the environment and observing those changes caused by exogenous events. A decision is made if the change in the environment is sufficiently different from the established P-state to interrupt planning (discussed in detail in (Mansell 1994b)). This primarily involves ensuring the rank order of evidence supporting the propositions does not alter. If a change in the ordering of any of the propositions that appear in the preconditions or postconditions of the operators constituting the plan change, the plan is re-evaluated.

This planning algorithm has been applied to an air combat domain where the goal is to successfully attack a target aircraft, given only partial information about the target location, type, and status. A number of strategies exist on how the attack should be carried out. Each strategy uses a different method of attack, and therefore has a different probability of success, and a different degree to which it fulfils the goal.

A general description of the U-Plan planning system can be found in Mansell (1993). The result of applying U-Plan to the air combat domain is presented in Mansell (1994b). An outline of the quantitative operator selection algorithm used by U-Plan appears in Mansell and Smith (1994a). A detailed discussion and analysis of U-Plan can be found in (Mansell 1994b, and 1997).

2 Possible Worlds Representation

U-Plan removes a significant restrictions placed on classical planning systems; that is, a single initial state adequately describes the environment. U-Plan constructs a set of possible initial states (P-states) based on the available evidence, (discussed in section 2.1). Each P-state includes a quantitative measure of belief that the P-state accurately describes the true state of the world, (section 2.2). These P-states are grouped according to their descriptions at differing levels of abstraction in a tree-like structure (section 2.3). Planning then takes place for each P-state in the order discussed in section 2.4.

2.1 P-states

When an incomplete model of the world is all that is available, a set of initial states can be used to describe the alternative environments. U-Plan employs a set of initial possible states (P-states) to describe what might be true of the world. A P-state, ps(a), is a complete description of one possible world using propositional statements. Each P-state is described hierarchically with n levels of abstraction, $(ps(a)=\{\ell_1(a) \ldots \ell_n(a)\})$ where n is domain dependent and selected during knowledge engineering (fig 2.3(a)). The level $\ell_i(a)$ is a complete description of a world at the ith level. The highest level of abstraction gives a coarse description of the state of the world. The lowest level gives a detailed view of the world. Intermediate levels provide the description required to make a smooth transition between both extremes.

Information sources provide U-Plan with a set of propositional statements, that represent distinct aspects of the domain. Each propositional statement has associated with it a measure of certainty (U-Plan uses an E-R mass distribution for reasons discussed in section 2.2). The propositional statements are then mapped to the lowest level of abstraction where they are used to generate a set of detailed P-states.

Mapping functions (defined at knowledge engineering time and domain dependent) are then used to construct a representation of the detailed state at the next level of abstraction. This process continues until each state is described at every level of abstraction. A P-state is one possible world with a description at every abstraction level (e.g., P-state, ps(a), is also represented by $\{\ell_3(a),\ell_2(a),\ell_1(a)\}$, giving a description of the same possible world at differing levels of abstraction).

A P-state consists of a set of propositions that syntactically encode specific attributes of the domain, for example, radar status (*Radar(on)*), or the altitude of the aggressor aircraft ($\mathit{Alt}_3(0\ 1)$). These propositions are diverse enabling a description of the environment to be generated at all levels of abstraction. Determining the amount of detail contained in these propositions takes place at knowledge-engineering time according to the strategic and tactical content of the information they represent.

E-R is used to assess the effect of all pieces of available evidence on a hypothesis. A propositional space called the *frame of discernment* is used to define a set of basic statements, exactly one of which may be true at any one time, and a subset of these statements is defined as a *propositional statement*. For example, in the air combat domain, a frame of discernment, θ_A, might be used to represent every type of aircraft, i.e.

$$\Theta_A = \{a_1, a_2, \ldots a_n\} \qquad (2.1)$$

where one of the basic statements a_i might be "the aircraft is a F/A 18". A propositional statement A_i might be "the aircraft is a fighter", that is, the proposition is the subset of θ_A containing all a_j that nominate different types of fighter.

When acting in a complex world, a method for representing and reasoning with information from disparate sources described in different frames of reference must be available. For example, frame θ_A might represent aggressor aircraft type, while frame θ_B might represent aggressor aircraft altitude. E-R uses compatibility relations to characterise interrelationships between different propositional spaces. This allows reasoning to be carried out on information described at different levels of abstraction or on frames of reference with overlapping attributes.

U-Plan uses compatibility relations to manipulate information at different levels of abstraction. U-Plan accepts information about the environment at different levels of abstraction, and compatibility relations are used to generate a description of particular attributes of the world at all levels of abstraction. Compatibility relations are also used to generate a complete description of the possible worlds (called P-states).

A P-state description is first generated at the lowest level of abstraction where the world is described at its most detailed. Compatibility relations are used to generate a description of this P-state at each of the higher levels of abstraction. As a P-state is a complete description of the world, it must bring together evidence from every frame of reference. In most cases, one can assume the frame that brings all the frames together is the cross product of the individual frames. In the air combat domain, the common frame generated using compatibility relations between frames at the same level of abstraction, is equivalent to the cross product of those individual frames of discernment (as the frames are mutually exclusive).

An example of two compatibility relations are given in figures 2.1 and 2.2. These figures detail the compatibility mappings for aircraft altitude and type/intent respectively between levels of abstraction. For example, initial evidence may identify the type of aircraft as being a fighter. The compatibility relation in figures 2.2. can be used to identify the aircraft to be either fighter-1 or fighter-2, and the intent as being either fighter-cover or air-superiority (i.e., Type3(f1 $\wedge$ f2) and Intent1(fighter-cover $\wedge$ air-super)).

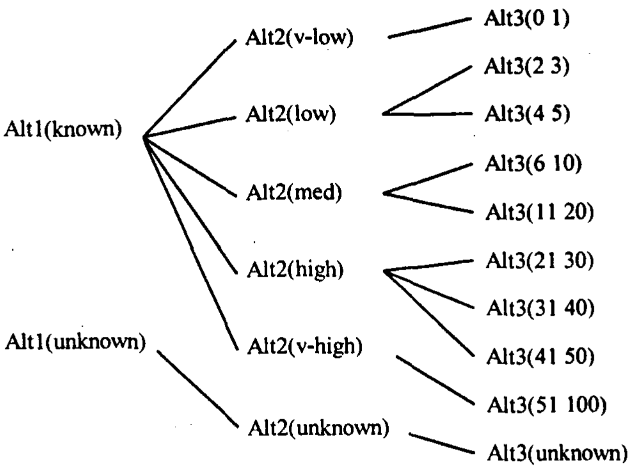

Figure 2.1: The compatibility relation used to map altitude to the neighbouring levels of abstraction.

Compatibility relation only link proposition at one level, to the levels directly above or below it. (That is, a compatibility relation can not be constructed to link level 1 propositions directly to level 3 propositions). A compatibility relation linking level n to n+1 can be thought of as a coarsening of the frame (e.g., in figure 2.1 the compatibility relation for $Alt_3((0\ 1))$ is coarsened to $Alt_2(V\text{-}Low)$). Similarly, a compatibility relation linking level n to n-1 can be thought of as a refining of the frame, (e.g., in figure 2.2 the compatibility relation for $Intent_1(Air\text{-}Super)$ is refined to $Type_2(fighter)$ and $Type_2(fight\text{-}bomb)$). To ensure each level is a complete description of the possible world, a compatibility relation must exist for every proposition at every level. This can result in some compatibility relation adding no new information for those propositions whose representation does not change from one level to the next (for example, Alt$_3$(unknown) $\leftrightarrow$ Alt$_2$(unknown)).

The lowest level (i.e., level 3) representation of the P-state contains all the information needed to do the low level tactical planning, (e.g. velocity, position, target-location, etc.). One would expect to find the quantitative data required to make tactical decisions within the domain. The propositions (and their masses) that represent level 3 information will originate directly from the knowledge sources, or as a result of the refining of propositions at the second level of abstraction. This information will be used at the lowest level of plan generation (i.e. the leaf nodes of the strategy hierarchy) for such things as deduction and geometric calculation.

The propositions at intermediate levels (i.e., level 2) of the P-state represents information with both strategic and tactical content. They can result from being directly

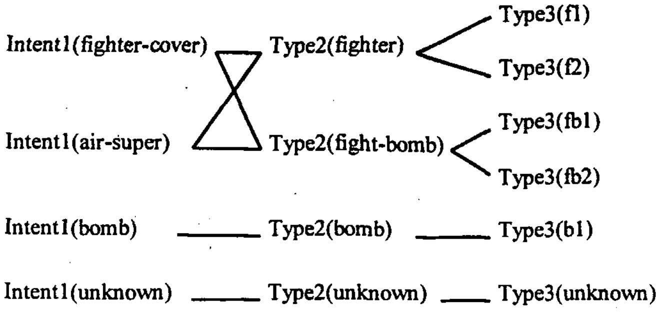

Figure 2.2: This figure shows the compatibility relation used to link aircraft type and intent across the levels of abstraction.

inferred by some knowledge source; or they can be the result of the coarsening of propositions at the next lowest level; or by the refining of propositions at the next highest level. The information conveyed here is likely to be both qualitative (e.g. $ALT_2(high)$) and quantitative, (e.g. $HEAD_2(0)$).

The propositions described at the highest level represent the strategic information required for early goal/subgoal selection. Most of the propositions at this level will represent the coarsening of information stored in some lower level of the P-state's hierarchy. However, some directly inferred information may make up this level's set of propositions (depending on the domain). The information conveyed at this level is intended to be qualitative.

So, for example, if an information source reported an enemy aircraft approaching at an altitude of 900 ft, this could be represented by a level 3 operator - $ALT_3(<1K)$. A compatibility relation is used to produce a level 2 description as $ALT_2(v\text{-}low)$, and similarly, at level 1 description is produced, $ALT_1(known)$.

One of the properties of compatibility relations is that they can map many propositions to one proposition. By being able to represent the world with fewer numbers of predicates for decision making at higher levels of abstraction, a more efficient representation and search can be made.

2.2 P-state Ranking

Information acquired in a real-world situation provides evidence about the possible states of the world. This information is typically uncertain and incomplete. E-R (Shafer, 1976 and Lowrance, *et al* 1991) is one way of handling such evidence, using an interval to explicitly capture what is known as well as what is not known (i.e., uncertainty).

Dempster-Shafer (D-S) theory is a mathematical theory of evidence conceived by Dempster (1968) and elaborated by Shafer (1976). E-R (an extension to D-S theory) reasons about information that is typically uncertain, incomplete and error-prone. E-R differs from classical probability theory in that it allows its measure of belief to be associated with disjunctions of events rather than requiring probabilities to be distributed across a set of individual events. The result is that one need not assume that all data are available and it provides a means for combining data captured at various levels of abstraction.

In much the same way that one may, given sufficient information, assign probabilities to situations and compute probabilities for all possible combinations of situations of interest, one may assign values (known as masses) to one's beliefs in all possible propositional statements in a frame of discernment and use these to compute the evidence supporting a hypothesis and the evidence refuting the hypothesis. The process of assigning masses, $m_A(A_i)$, is called a *mass distribution*. Masses have the property:

$$\sum_{A_i \subseteq \Theta_A} m_A(A_i) = 1 \qquad (2.2)$$

where the domain of A_i is the set of all possible subsets of θ_A, i.e. the power set 2^{Θ_A}.

In U-Plan, one obtains a mass distribution across all frames of discernment. Compatibility relations (section 2.1) are then used to describe which elements from two frames can be true simultaneously; allowing propositional statements to be addressed jointly. For example, Θ_B might represent altitude of a target aircraft, and the compatibility relation maps Θ_A and Θ_B to a new frame $\Theta_{A,B}$ which (in this case) is the cross-product of the two sets.

Information about belief in a hypothesis A_j is contained in what is called the *evidential interval*, which constitutes a measure of *support*, Spt(A_j), and *plausibility*, given by:

$$\mathrm{Spt}(A_j) = \sum_{A_i \subseteq A_j} m_A(A_i) \qquad (2.3)$$

$$\mathrm{Pls}(A_j) = 1 - \mathrm{Spt}(\Theta_A - A_j), \qquad (2.4)$$

$$[\mathrm{Spt}(A_j), \mathrm{Pls}(A_j)] \subseteq [0, 1]. \qquad (2.5)$$

Stated simply, the support for a hypothesis A_j is the sum of the masses of all propositions that are subsets of A_j (including A_j itself). And, the *plausibility*, Pls(A_j) is the degree to which the evidence fails to support its negation. The difference between support and plausibility represents the residual ignorance, or uncertainty, $U_\theta(A_j)$=Pls(A_j)-Spt(A_j). The evidential interval is illustrated by:

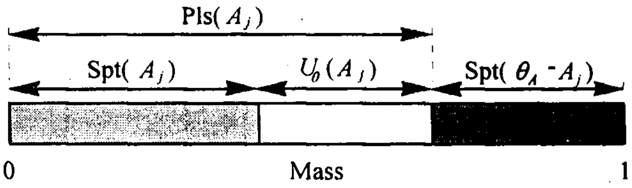

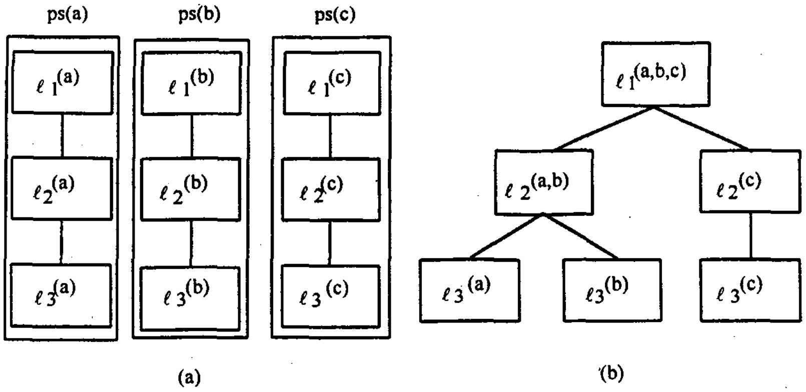

Figure 2.3: (a) Depicts how 3 sample P-states have representations at 3 abstraction levels. $\ell_n(x)$ represents P-state, x, at abstraction level, n. (b) An example of how 3 initial P-states may be grouped in tree form.

The evidential interval is usually represented by [Spt(A_j), Pls(A_j)], (where actual numerical values are calculated for each A_j of interest).

A measure of support and plausibility is calculated for each initial P-state at every level of abstraction using equations (2.3 and 2.4) based on the mass distributions of the initial evidence.

2.3 P-state Grouping

U-Plan groups together equivalent initial P-states according to their hierarchical levels, ie., the P-states with the same state description at a particular abstraction level are grouped together.

Figure 2.3(b) demonstrates how initial P-states may be grouped in tree form. In this example the set of P-states from figure 2.3(a) are used. At the lowest level of abstraction the set of possible worlds are distinct, represented as the leaf nodes of the tree, $\{\ell_3(a),\ell_3(b),\ell_3(c)\}$. Let us assume, when viewing the world in a more coarse light, ie. at a higher level of abstraction, $\ell_2(a)$ and $\ell_2(b)$ are identical. In this case they would be grouped together to give $\ell_2(a,b)$. At the highest level of abstraction $\ell_1(a,b)$ and $\ell_1(c)$ might also be identical resulting in the state $\ell_1(a,b,c)$.

2.4 P-state Selection

The selection of the initial P-state to begin planning involves choosing the P-state with greatest support[1] at the highest level of abstraction, (for example $\ell_1(a,b,c)$). The node in the P-state tree that is a child of this initial P-state with the greatest support is then selected (e.g. $\ell_2(a,b)$ or $\ell_2(c)$). This selection process continues from highest to lowest level of abstraction. The result is an initial P-state with a description at all levels of abstraction.

The P-states are chosen in this manner in an attempt to allow the possible world with the greatest support to be planned first. This does not guarantee the plan will have the greatest support when planning is complete, or that the best plan will be constructed first. The usefulness of this strategy becomes apparent when attempting to use all or part of previously constructed plans during planning for other P-states (Mansell 1994b). The effectiveness of this approach relies on a suitable representation of the domain and the reduction operators.

If one particular P-state has an outstandingly high degree of evidential endorsement, that P-state generally dominates the high level representation of the world with the greatest evidential interval. This usually results in a plan that has the greatest evidential endorsement. If no single P-state stands out, the high level state first planned for should produce the plan strategy likely to be picked up by a collection of P-states that represent a significant portion of the set of P-states. (For example, if the first plan generated uses the *Cutoff-intercept* manoeuvre as its strategy, then that strategy is likely to be reused enough times to result in it having the greatest evidential endorsement when a plan exists for every possible world). However, this method does not guarantee the best plan will be generated first, although it appears to achieve the desired results for information that can be represented in a Markov Gallery (Lowrance *et al* 1991) (i.e., conditionally independent information).

3 The Dynamic World

The assumption that the world remains static during planning has been a point of criticism among classical planners, and an argument for the use (at least partially) of reactive planning techniques (Georgeff and Lansky 1987). U-Plan is designed to operate in a dynamic environment, that is, where the body of information available to describe the state of the world may be constantly changing. Of particular interest to the U-Plan system is when changes to the balance of information affect the P-states, and how *significant* those changes may be on existing plan, (i.e., do they necessitate replanning).

The Planning and execution of actions are carried out by U-Plan in a dynamic environment. U-Plan uses a separate module to monitor changes in the dynamic environment and will interrupt the planning process if the world changes significantly. To define what a significant change in a proposition held in a P-state and the true state of the

---

[1] The selection of the initial P-state is based on the selection of the best E-R interval. A variety of techniques dealing with interval based decision making exists.

environment, we must first look at how the initial information may change.

In the dynamic and imprecisely described environment, the view of the world held in the P-state can become incompatible with the true description of the world in two ways:

1. The propositional statements used to capture an initial piece of evidence may change at a later date, invalidating the statement in the P-state. This is an easily detected change in the environment that could either invalidate a plan (triggering replanning), or have no effect on the plan.
2. The degree of evidence supporting a particular proposition in the initial P-state may change in time. When this occurs the mass distribution (and consequently the support and plausibility) for the propositions may result. The effect in this situation is a change in the mass distribution for the generated plan(s).

It is the second condition that is of interest. That is, when is a change in the environment significant enough to alter the order in which two plans are ranked (using their evidential interval). Mansell (1994b) has shown that the rank order of two plans does not alter until the rank order of the evidence supporting the propositional statements changes. For example, if two propositional statements are given the following mass distribution:

$m(Alt_3((0\ 1))) = 0.4 \qquad m(Alt_3((2\ 3))) = 0.2$

$m(Alt_3((0\ 1),(2\ 3))) = 0.4$

Lets assume a change in the environment sees an alteration in the mass distributed among the Alt_3 predicates (i.e., a change in the degree of uncertainty in the environment). Then the rank order of the plans produced for the super-plan will not change unless the mass attributed to $Alt_3(0\ 1)$ falls below the mass attributed to $Alt_3(2\ 3)$.

4 Summary and Discussion

Evidential reasoning has proved invaluable to the development of U-Plan due to its ability to managing uncertain information that is expressed in multiple levels of abstraction. E-R's compatibility relations are ideally suited to representing the interrelationships between propositions at different levels of abstraction essential to U-Plan. This formalism also allows U-Plan to gather information from disparate sources using pertinent frames of reference, and link these to a common frame. Compatibility relations are then used to construct consistent representations of the world at different levels of abstraction.

U-Plan has been successfully applied to an air-combat post mission analysis domain and hazard action response domain, and is currently being applied to a ship manoeuvre recommendation domain (Mansell, 1994b, and Mansell, 1997). The hierarchical P-state architecture described in this paper does not prove to be a computational burden on U-Plan as P-state generation is done once, prior to planning, and the dynamic world monitoring is distributed onto other systems. In fact the contrary has been observed; i.e, U-Plan is more efficient when the domain is abstracted (using 2 and 3 levels) compared to when the same domain is not hierarchically organised (Mansell, 1994b). This results because the number of propositions that describe the world at higher levels of abstraction are reduced, allowing U-Plan to function more efficiently at the higher levels of abstraction. In addition, the representation of operators are more concise as fewer propositions are required to describe preconditions, postconditions, effects, probability of success, etc.

Bibliography

D. Chapman (1987), Planning for Conjunctive Goals. *Artificial Intelligence,* 32: 333-377.

A. Dempster (1968), A Generalization of Bayesian Inference, *Journal of the Royal Statistical Society, Series B*, 30, 205-247.

M. Georgeff and A. Lansky (1987), Reactive Reasoning and Planning. In *Proceedings of the 6th AAAI*, 677-682.

J. Lowrance, T. Strat, L. Wesley, T. Garvey, E. Ruspini and D. Wilkins (1991), *The Theory, Implementation, and Practice of Evidential Reasoning*. Technical report, Artificial Intelligence Centre, SRI International, California.

T. M. Mansell (1993), A Method for Planning Given Uncertain and Incomplete Information. *The Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*, Washington DC, 350-358.

T. M. Mansell and G. Smith (1994a), Operator Selection While Planning Under Uncertainty. *The Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*, Seatlle, Washington.

T. M. Mansell (1994b), *An Approach To Planning Given Uncertain and Incomplete Information.* PhD Thesis, The University of Melbourne, Australia.

T. M. Mansell (1997), An Approach To Planning Given Uncertain and Incomplete Information, *Computational Intelligence*, (in Submission).

E. H. Ruspini (1986), *The Logical Foundations of Evidential Reasoning*. Technical Note No. 408, Artificial Intelligence Centre, SRI International, Menlo Park, California.

G. Shafer (1976), *Mathematical Theory of Evidence.* Princeton University Press, Princeton, New Jersey, USA.

D. Wilkins (1988). *Practical Planning: Extending the Classical AI Planning Paradigm*. Los Altos, California: Morgan Kaufmann.