

# Knowledge Acquisition for Search and Rescue<sup>1</sup>

Hugh Cottam and Nigel Shadbolt.

University of Nottingham, AI Group, Dept of Psychology,  
University Park, Nottingham NG7 2RD.  
{hdc | nrs}@psyc.nott.ac.uk)

**Abstract** . The field of knowledge level modelling has achieved great success when applied to various domains, yet has thus far largely neglected the generic areas of planning, scheduling and resource allocation. In this paper we outline the development of a knowledge level modelling approach within the domain of planning for Search and Rescue. Existing problem solving models for planning are almost exclusively derived from the analysis of systems. We argue that this makes their suitability for directly assisting knowledge acquisition debatable. Our approach makes a clear distinction between domain derived knowledge level models and those derived from systems. We describe how the combination of these two types of model can achieve definite benefits within the course of KBS development. The paper also highlights important aspects of modelling expert workflow, which reveal key requirements for any effective knowledge intensive system.

- [INTRODUCTION](#)
- [THE SAR DOMAIN](#)
- [KNOWLEDGE ACQUISITION FOR SAR](#)
  - [Observations of the SAR domain](#)
  - [The knowledge intensive nature of the RCCs planning](#)
  - [Generic incident classes](#)
  - [An example of a complex scenario in SAR planning](#)
  - [A Case Study -- The Missing Yacht](#)
  - [Task Decomposition of SAR Planning](#)
  - [Implications for an Embedded Planning Support Aid.](#)
- [KNOWLEDGE LEVEL PROBLEM SOLVING MODELS](#)
  - [Critiquing the Domain PSM Using a System Derived PSM](#)
  - [The Mapping Between the Models](#)
- [COMPARISON WITH EXISTING PSMS FOR PLANNING](#)
- [SUMMARY](#)
- [CONCLUSIONS](#)
- [REFERENCES.](#)

## INTRODUCTION

This paper discusses insights derived from a project commissioned by the Defence Research Agency Flight Systems Division. The project was entitled "Acquiring and Using Planning Knowledge for Search and Rescue" and was motivated by the requirement for the Flight Systems Division of DRA Farnborough to find ways in which their knowledge engineering work within planning system development could be made more efficient and reusable. This was coupled with the mutual interest of the AI Group at Nottingham and AIAI at Edinburgh, in the development of methodologies to assist in knowledge acquisition for planning systems.

A central goal of the project was the development of a generic approach to assist in the reliable capture of knowledge related to planning, scheduling and resource allocation. This would involve the construction of a knowledge level model to describe the structure of problem solving in the SAR domain. The aim was that this model would be suitable for re-use in similar domains. This paper will describe in detail the accomplishment of this goal and its relation to other work on problem solving models for planning. For a more general overview of the SAR project see [Cottam *et al.*, 1995.].

The increased use of intelligent decision support systems has created a demand for efficient acquisition, implementation and maintenance of the knowledge required by such systems. This has led to the construction of methodologies for KBS development that facilitate a generic approach to knowledge acquisition. e.g. KADS [Breuker *et al.*, 1987.] or VITAL [O'Hara *et al.*, 1992]. These methodologies rely heavily upon the concept of a generic problem-solving model (PSM in this context refers to the model used to drive KA not the problem solving method). Generic PSMs resulted from the discovery that when a certain number of PSMs were purged of their domain-specific content, the resulting structures seemed invariant over various domains. Users of knowledge level methodologies have thus built up extensive libraries of generic PSMs, aimed at facilitating the reuse of both knowledge engineering effort and system software itself.

Such generic approaches have achieved great success when applied to various domains, yet have thus far largely neglected the areas of planning, scheduling and resource allocation. This point should be clarified as it might well be argued that there are in fact a number of existing generic PSMs for planning. e.g. *CommonKADS Library for Expertise Modelling* [Breuker & van de Velde, 1994]. The important observation to be made about these existing PSMs for planning is that they are almost exclusively system derived. This raises an important discussion point concerning the origin of the generic PSMs associated with a methodology. The originators of a methodology may deduce the ontological elements for use in such models, yet the structure of generic models must be inferred or validated inductively. There are several established generic PSMs that have originated as a result of system analysis, as opposed to human expert analysis. e.g. the heuristic classification model of Mycin [Clancey, 1985]. However the strength of a model such as heuristic classification is not simply the initial system analysis, but its proven efficacy for knowledge acquisition. This observation has important implications for the construction of generic PSMs within the generic area of planning. Existing planning PSMs have resulted as an attempt to extend the use of knowledge level methodologies, yet are derived from the knowledge level analysis of well known operational planning architectures. Because such planning PSMs model how computers plan rather than how humans plan, their efficacy for human expert knowledge acquisition is debatable, and they may enforce an unsuitable system architecture upon the domain.

Most methodologies for model based knowledge acquisition expect that generic PSMs will require modification once selected for use in a particular domain (or else constructed from several smaller generic PSMs). The model selection/construction/modification process is also the least well supported of the knowledge acquisition stages. This process is aided to a degree by the organisation and classification of the PSM libraries [Valente & Löckenhoff, 1993] [O'Hara, 1993]. Specifically with regard to the CommonKADS planning model library, it is suggested that the availability of knowledge to play certain static roles within a model will give guidance to PSM selection [Valente, 1994] [Barros *et al*, 1996]. This does not however provide any validation of the structure of the PSM. There is a clear danger when applying system derived PSMs to knowledge acquisition. The knowledge engineer may tend to force the characterisation of the domain to fit the "off the shelf" model. These system derived PSMs await validation through their use and refinement in the context of knowledge acquisition.

Our observations on the nature of existing PSMs for planning, led us to the goal of constructing an explicit PSM for the SAR domain from a combination of domain analysis and wider ontological considerations. Thus the structure of the PSM was domain driven, which we considered to be of great importance in order to avoid imposing a pre-conceived PSM upon the domain. Once shorn of its domain specific content this model can be compared to generic PSMs such as those mentioned earlier from the CommonKADS library (see section 5). The applicability of the resulting generic PSM could then be considered for other domains.

It is the intention of this paper to demonstrate that had we attempted to utilise one of the existing CommonKADS models directly, then this would have hindered rather than assisted the knowledge acquisition and would have lead to a system that did not preserve the structure of planning for SAR. We also aim to show that system derived models can be utilised to great effect in knowledge acquisition by using them as a means of critiquing domain derived models.

Initially the ontological issues regarding the entities that would constitute a PSM for planning were addressed, and possible structures of PSMs for planning in the SAR domain were then considered. There is widespread interest in ontologies to support knowledge sharing across a range of domains and, at this time, a rapidly growing interest in the development of ontologies for planning [Tate, 1994]. These ontologies establish a set of consistent terms to describe the entities that constitute a plan and the relationships between such entities. The ontologies thus represent "what" is reasoned about in planning, but do not explicitly represent "how" planning is performed. It was our objective to merge generic planning ontologies with a knowledge level modelling approach in order to formulate a planning PSM for the SAR domain.

## **THE SAR DOMAIN**

Planning for search and rescue is based at the Rescue Co-ordination Centre (RCC) at Pitreavie near Edinburgh. The RCC have responsibility for the support of military flying, yet their most common role is in support and co-ordination of civilian emergencies. The RCC's geographical area of responsibility extends from a line South of Birmingham northwards as far as Iceland, and out into the North Atlantic and North Sea. The RCC have direct responsibility for the allocation, application and co-ordination of military assets for SAR (this includes SAR helicopters, RAF Nimrods and RAF mountain rescue teams). They may however have to co-ordinate with a number of civilian emergency authorities such as fire, police, ambulance, coastguard and civilian mountain rescue teams. They might also take responsibility for overall co-ordination of a rescue incident that includes the allocation and application of these civilian rescue assets. A rescue incident can vary in scale from retrieving a walker with a sprained ankle to handling a large aircraft crash.

Figure 1 shows the layout of the RCC main information sources. These consist of magnetic boards and are described from left to right. The bases map indicates the positions of SAR bases, main hospitals, and county police borders and refuelling sites. The UK Airborne Assets map indicates position of incidents and airborne assets using magnetic icons that are moved by hand. It may also be used to indicate the position of such things as radio locator beacons. The UK SAR Assets Status board indicates the status of SAR helicopters, Nimrods and mountain rescue teams. When assets become airborne the icon representing them is moved to the appropriate position on the UK Airborne Assets Map. Air ambulances are also kept track of on these displays as the RCC is often required to co-ordinate with them. An Active Danger Area map displays the position of areas that are dangerous for flying. A helicopter with a rescue callsign takes priority over all other air traffic. Active Danger Area board displays information corresponding to the active danger area map, such as the times at which the area will be active or the height to which the area is considered of danger.

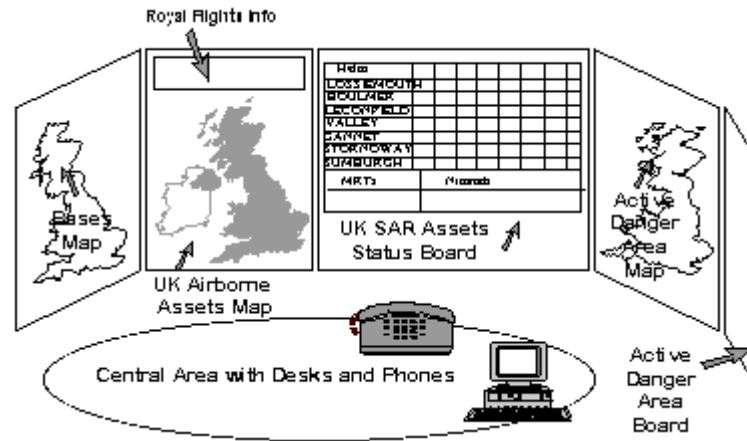


Figure 1: Front View of RCC Information Sources

## KNOWLEDGE ACQUISITION FOR SAR

The course of knowledge acquisition within the project largely followed the lifecycles of KBS described in methodologies such as KADS [Tansley & Hayball, 1993] and VITAL [Motta *et al.*, 1994]. The lifecycle departs from the norm at certain stages owing to the lack of a generic PSM on which to base the modelling process. Knowledge acquisition commenced with extensive tutorial KA involving basic documentation of the domain and was followed by domain experts giving in depth explanations of individual rescue incidents and the decisions associated with the handling of those incidents. These interviews were taped and transcribed for analysis. A number of video recordings were made of actual incidents in progress and the RCC's handling of them. These recordings proved invaluable in eliciting an explicit structure describing problem solving for SAR. The video medium proved particularly effective at capturing the temporal relationships prevalent within planning at the RCC. The videos enabled us to rigorously document both the RCC's information sources and the manner in which they interacted with these during the course of their workflow. An information stores document listed all the data stores used by the RCC during the course of their SAR work.

### Observations of the SAR domain

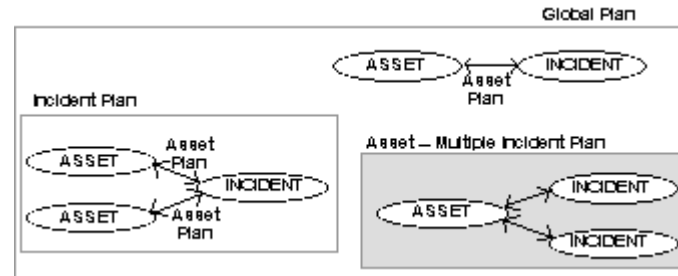
Tutorial KA was followed by detailed analysis that concentrated on the distinguishing features peculiar or characteristic of the problem solving performed within the RCC. These observations arose partly through the visits to the RCC, but also through extensive analysis of the taped tutorial interviews and the workflow video footage. This initial analysis resulted in various system proposals. The acceptance/rejection of these further focused the next stages of KA.

An initial observation of the RCC is that most of their information sources are either paper based or magnetic boards. These types of representation tend to be used for those information items within the RCC that are frequently changing. A computer database is used containing information of a more static nature; e.g. geographical positions of hospitals, landing sites, and decompression chambers.

A fundamental observation of the domain was that SAR assets can be divided into two distinct categories. Those that are directly controlled and those that can only be controlled indirectly. The indirectly controlled assets are almost exclusively directed by other authorities, and the RCC's actions associated with these types of asset are limited to attempting control via the appropriate authority. Any communication with or feedback from these assets takes place via the authority. The observation, led to the conclusion that such constraints upon the control of an asset correlate closely with the degree of detailed planning the RCC perform upon it. Similarly, the more indirect the RCC's control of an asset the more they must hypothesise about it's actions.

The three directly controlled assets for which detailed planning is possible, are helicopters, Nimrods and mountain rescue teams. Helicopters are the most heavily utilised of these assets, and are crucial in the majority of the RCC's work. It was agreed that the later stages of KA and the actual demonstrator should reflect this bias.

Figure 2 indicates four distinct levels at which planning in the SAR domain could take place. One of these is shaded as it does not actually arise within the domain. The three levels at which plans can be considered are at the asset/resource level, the incident level and the global level. The level that does not arise at the RCC is the asset to multiple incident level. This multitasking of assets does not occur due to the geographical separation of incidents. However the position of an asset due to it's application to a particular incident may be taken advantage of. This efficient dovetailing of asset plans will be seen in the example explained in section 4.4.



**Figure 2: Plans for SAR**

## The knowledge intensive nature of the RCCs planning

An early impression gained of planning for SAR is that the situations with which the RCC are dealing can soon become very complex with many factors that must be considered when making a decision. Even for an incident requiring a single helicopter many factors influence plan details and resource allocation: proximity of base to incident, weather, terrain between base and incident, necessity to refuel or collect rescue workers, deciding if an extensive search will be required, state of base readiness, suitability and existence of landing sites.

Many of these constraints are known in advance or can be ascertained prior to making a decision. Even in mundane situations planning and resource allocation can significantly alter the response capability of the RCC. It is easy to see how logistical problems can become very complex when many assets are involved. The degree of complex interaction between factors makes the construction of effective solutions a knowledge intensive task requiring much expertise. However, the problem solving is structured enough to allow a planning system to advise and inform a human expert in this process. Many of the benefits that could be derived from a KBS approach to supporting the RCC's problem solving would be found in plan visualisation, complexity management and the automation of certain mundane tasks such as log keeping.

## Generic incident classes

The primary decision made at RCC with relation to an incident is the initial classification of that incident. There are 6 explicit classifications used by RCC, yet it is our assertion that we can usefully refine these incident classifications. The finer grained classifications are made through a combination of the characteristics of the incident and the RCC's intended handling of that incident. They are best described as generic classes that correspond to the RCC's view of an incident. Figure 3 illustrates a hierarchical refinement of generic incident classes for three high level classifications (as already used by RCC), mountain, maritime, and air crash. We believe that for each of these generic classes we can identify a plan template. This describes a series of high level goals associated with planning the application of assets to an incident. A number of these templates have been identified.

Mountain	No search	Single Helo
"	"	Single Helo + MRI
"	"	Single MRI
"	Search required	MRI's + Helos
Maritime	No search	Helos only
"	"	Helos + Lifeboats
"	"	Lifeboats only
"	"	Nimrods + Helos
"	"	Nimrods + Helos + Lifeboats
"	Search required	Nimrods only
"	"	Nimrods + Helos
"	"	Nimrods + Helos + Lifeboats
Air crash	No search	Military
"	"	Civilian
"	Search required	Military
"	"	Civilian

## An example of a complex scenario in SAR planning

The example corresponds to the generic class for a Maritime incident involving search and requiring the application of Nimrods, helicopters and lifeboats (see Figure 3). It is an incident that actually occurred during one of the early visits to the RCC and was captured on video. It demonstrates how a situation can become complex due to a number of relatively mundane incidents occurring in parallel. This has the effect of seriously straining the capabilities of the available assets and necessitating careful decisions with regard to resource allocation and plan management. The developing situation is pictured in Figure 4. At the beginning of the scenario there is one air ambulance that is presently unserviceable in a railway yard at Wick. This can be seen as the top grey circular icon in Figure 4 marked with a cross. Rescue 137 (this is the callsign of the helicopter) is in the process of taking a Marine with head injuries to hospital in Inverness. These injuries were sustained falling from a climb in Glencoe. Rescue 131 has been requested by the coastguard to assist in going to the aid of a sinking trawler near Montrose. Incidents in Figure 4 are marked as star shaped icons. Rescue 132 is airborne but is only out on a training sortie. Rescue 131 and 132 are both based at RAF Boulmer in Northumberland. Boulmer is only required to have one helicopter on call, and to be able to scramble another within an hours notice. At this point the RCC receive a call about a new incident involving a fallen climber in the Cheviots. At the same time they learn that Rescue 131 which is on its way to the stricken trawler near Montrose is no longer required and can return to base.




Figure 4: Example of complex scenario

It is at this point that the situation becomes complex and requires careful decision making. The RCC need to find a way of handling the Cheviots incident. However, this presents difficulties. The obvious choice is to send Rescue 131, yet they find that they are unable to establish contact with 131 and are unable to ascertain its position. Rescue 132 could be sent but they do not have a full crew on board (meaning that they have no winchman or winch operator). Rescue 137 cannot be sent because it hasn't finished taking the injured Marine to hospital and it will also require to refuel after this. There is an air ambulance that is just finishing a job taking someone from the island of Mull to hospital in Glasgow. The RCC therefore decide to use this asset. However, this decision represents a gamble as it is possible that the fallen climber in the Cheviots may require lifting off with a winch. The air ambulance is not equipped with a winch so this would represent difficulties. The RCC deal with this by reasoning that the Cheviots tend to be quite rolling hills so the helicopter will probably be able to land. They also cover the alternative situation by seeking to route Rescue 131 back over the Cheviots. This however represents difficulties in itself as they continue to have problems establishing communications with 131.

The example shows how "run of the mill" situations quickly develop into complex ones that stretch the limited SAR resources at the RCC's disposal. It can be seen how complexities escalate in the context of a large incident such as an aircraft crash. The example also illustrates that the RCC are normally working in the context of unreliable communications with airborne helicopters.

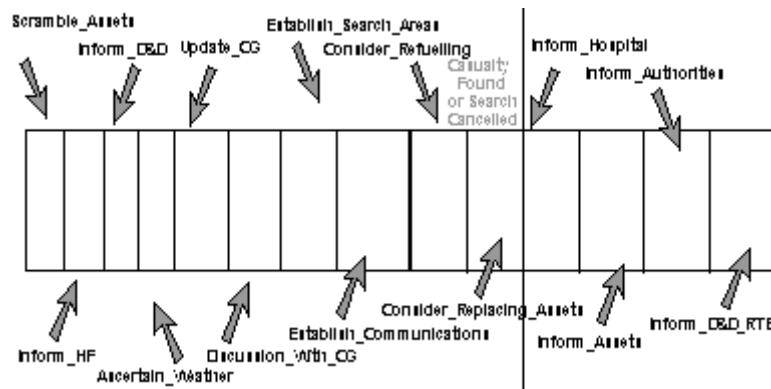
## A Case Study -- The Missing Yacht

The knowledge that has been acquired to support the handling of maritime incidents requiring search and involving Nimrods, helicopters and lifeboats, represents the population of the domain derived problem solving model described in section 4. This specific knowledge has been acquired from a combination of sources. We have used the actual RCC documented records of this type of incident, video taped interviews of the RCC explaining case histories of this type of incident and a transcription of an explanation of another maritime search incident involving just Nimrods. The knowledge represented within this latter type of incident has certain common features with the generic incident class with which we are concerned. This demonstrates the fact that there may be considerable overlap of specific knowledge

between generic incident classes. As would be expected there are similarities across the maritime class of incidents. However this overlap also occurs in less obviously related incident classes. We shall use this example to point out some distinguishing features of the incident that alter the RCC's handling of it from the manner in which they would handle incidents that initially appear similar.

The incident involves a yacht, the "CRUSADER" which left Portree harbour at approximately 1100 in the company of another yacht, the "BERLIN". The "BERLIN" made for shore due to bad weather. The "CRUSADER" however has not been seen since, and cannot be raised on any frequency. This was reported to the RCC by Stornaway coastguard at 0010. The last sighting position is uncertain. It is not known whether the yacht is in difficulty or not. If the yacht is in difficulty it is not known where and when this began. This is an important factor, as if the yacht were drifting from a datum position the search area can be calculated according to wind and tides. In this instance they must also search the coast line in order to see if the yacht has taken shelter. A characteristic of this incident is the unconstrained nature of the search that is required.

Nimrods, helicopters and lifeboats can all be used effectively in the search. The coastguard have already requested a helicopter and 3 lifeboats for the search. The lifeboats are an effective resource for the search of the coastline, but would be little use in searching a large area of open sea. The maritime search, is potentially a large area, so a Nimrod is requested by RCC. Because the coastguard are the callout authority they will have already defined the search area. The coastguard have overall control of the search organisation in these situations. They calculate the search area, and this is then often re-calculated by RCC, which acts as a check. Normal procedure is to define a search area of what are considered the most probable areas for the location of the casualty. This area is then searched extensively, and if the casualty is not located then the search area is expanded.



**Figure 5: Plan Template for Maritime Incident (Lifeboats, Helos & Nimrods)**

We can clearly discriminate between maritime incidents requiring search and those that merely require a rescue. Within the maritime search category, the incident categorisation can be further refined according to a combination of the resources to be employed and the type of search that will be carried out. Searches that involve large areas of sea cannot be effectively aided by the use of lifeboats, so these types of search require the use of either helicopters or Nimrods, or both. An incident of this generic class necessitates a search of large sea areas and areas of coastline, thus requiring the use of helicopters, Nimrods and lifeboats. The lifeboats are effective for this coastline search. The search area is very unconstrained, as there is a large time window in which the casualty may have got into difficulty. It also means that the RCC must investigate avenues of what the casualty may have done if it is not in difficulty.

There is also a refuelling sites consideration, both in the planning of the search and its execution. The RCC (or the coastguard) may attempt to contact shipping for information on the yacht and also for refuelling facilities for helicopters. Shipping can also be used to provide on scene weather reports. Similarly oil rigs are potential refuelling sites and sources of weather information.

The plan template is divided into a number of high level goals which are ordered according to the likely temporal ordering of events within the course of applying the assets to the incident. This is an approximate ordering of the high level goals and it is likely that subgoals within different divisions will occur in parallel or in different order to that defined by the template. The plan template for the incident class applicable to the yacht incident would be that for an offshore maritime search, involving Nimrods, Helos, and Lifeboats. A diagram of this template can be found in Figure 5. The template could be decomposed to a finer level or the decomposition can take place during the application of the template and associated rules to the incident. We considered it a better reflection of RCC problem solving to keep the goals in the plan template at a high level. As can be seen in the diagram there may be certain important events that occur within the course of an incident, such as the casualty being found or the search being cancelled. This represents a point that changes the set of possible goals associated with the handling of the incident. In this case it can be seen that the high

level goals do not change if the search is cancelled or if the casualty is found. It will however change the actions required for achieving the goals.

The modularisation of the rules in the knowledge base corresponds to the inference steps described in the domain PSM. The structure unearthed during the knowledge acquisition was preserved through the design stage of the planning aid.

**Goal decomposition rules** decompose high level goals into lower level ones. They may include conditions that enforce a context sensitivity on goal decomposition.

```
e.g. RCC_Responsible_For_Informing_Hospital = TRUE
Casualty_Found = TRUE
Hospital_Treatment_Required = TRUE
-> Inform_Hospital_Of_Injuries
    Inform_Hospital_Of_ETA
    Inform_Other_Relevant_Parties
```

**Action assembly rules** assemble low level actions into higher level ones. Often these rules will mirror a similar goal decomposition rule, though this is not always the case. There may also be a context sensitivity to action assembly.


```
e.g. Informed_Helo_Bases_To_Scramble
Sitrep_Faxed_To_Helo_Bases
Helo_Crew_Verbally_Briefed
-> Helos_Scrambled
```

**Goal action translation rules** describe how goals translate into actions. In the demonstrator they require user input to inform the system that the actions have been completed. This might involve a choice point if there are several different actions that achieve the goals. When the user wishes to attend to a particular goal, the goal will either decompose into sub-goals or translate into one or more actions. A third possibility that corresponds with the O-Plan model is that other goals might be placed on the agenda (backward chaining). The conditions of the goal action translation rules may also require other actions to have been completed. In this manner temporal dependencies can be represented in the system.

```
e.g. Helo_Crew_At_Base = FALSE
Helo_Airborne = FALSE
-> Verbally_Brief_Helo_Crew_By_Portable_Phone
```

## Task Decomposition of SAR Planning

Figure 6 shows the high level tasks identified in the RCC's workflow. This task decomposition was necessary to gain a wider picture of the RCC's work flow, including the interaction and nature of the high level tasks outlined. Tasks 2, 3 & 4 in Figure 6 were considered to be knowledge intensive. Task 4 is the RCC making decisions about the actions that resources should take when applied to an incident. This is where detailed planning takes place and it is here that we concentrated our KA for the actual PSM construction.



**Figure 6: High Level Task Decomposition**

In addition to our KA goals, the task decomposition and workflow analysis also act as a means of defining user requirements and of assigning a focus for the deployment of KBS support. It became apparent during this analysis that there were key knowledge structures in the RCC's reasoning that had no explicit representation. A good example of this is the RCC's plans, which have no representation outside the minds of the operators. The reason for this is that the plans are simply changing too rapidly to be practically recorded in a paper based environment. A system incorporating multiple views upon a consistent maintainable knowledge-base was a clear requirement for any KBS support at the RCC. This would then form a framework around which plans could be represented and visualised. It also became clear to us at this stage that an appropriate solution would be a mixed initiative rather than a "black box" planner. The system should follow the RCC's natural style of problem solving as far as possible, and should act as an assistant rather than taking control away from the operator.



## **Implications for an Embedded Planning Support Aid.**

Our knowledge acquisition revealed several key characteristics of the SAR domain. Resource allocation, plan calculation and plan execution in the SAR domain can become complex with many affecting factors, resulting in knowledge intensive problem solving. In many situations the results of this allocation and planning become an input for the other areas of the RCC's reasoning. The target system would be aimed specifically at supporting the process of resource action planning. The functionality of the proposed system would include:

1. Reminding the user of the factors they should take into account when formulating a plan. It would then offer suggestions as to how these factors could be combined in order to satisfy the goals (these suggestions would be in the form of plans).
2. Providing an easily assimilable visualisation (visual record) of the current state of all assets, and the current plan. This would facilitate handover between shifts. Easy access would also be available to all decisions made and explanations of these.
3. Facilitating record keeping. This includes both updating of appropriate knowledge and log keeping. The log keeping would be aided as the system could easily printout all decisions and events concerning incidents. This could then act as a framework in which to interrogate the history of decisions and events.

## **KNOWLEDGE LEVEL PROBLEM SOLVING MODELS**

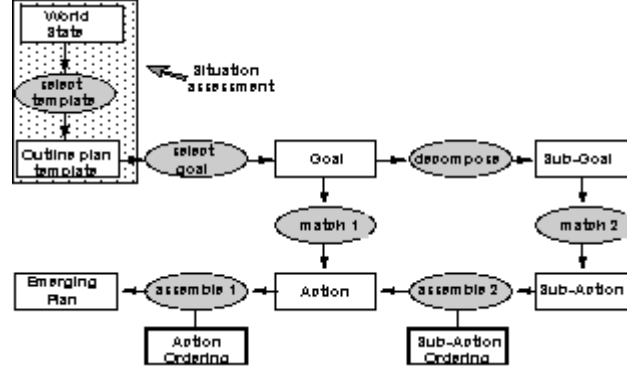
Within our discussion of planning for SAR (task 4), we shall refer to plan templates, executable plans, goals and actions. A plan template consists of a set of partially ordered high level goals that define the requirements of an incident. The template is generic to a generic class of incident e.g. maritime rescue requiring search and involving lifeboats, helicopters and Nimrod aircraft. There would be a certain set of goals associated with the handling of this type of incident. However another generic incident class would be associated with a different set of goals e.g. a mountain rescue not requiring search and involving a mountain rescue team and a helicopter. Planning within the RCC can be viewed as the process of moving from a plan template to an executable plan. An executable plan is a set of ordered physical actions to be taken at the RCC. The reasoning process of the RCC enables this transition between plan template and executable plan. We identified a library of plan templates which were then indexed according to a hierarchical organisation of generic incident classes. The first stage of RCC problem solving is situation assessment in order to define the incident class and to enable the selection of a plan template.

Planning for SAR is a progressive task that spans the temporal duration of a particular resource's application to an incident. This process involves the use of heuristic expert knowledge in order to make planning decisions in a domain where future data and constraints on planning are unpredictable. Due to this unpredictability, the decomposition of high level goals, and instantiation into planned physical actions, is usually not performed until the situation demands it. The RCC often resort to this least commitment strategy when planning. In this manner, the maximum amount of factual knowledge about the current situation is gathered before decision making. Oftentimes the RCC must hold back from taking physical actions, because if they wait for a small amount of time their factual knowledge of the situation will have increased so as to make a more effective decision possible.

### **Critiquing the Domain PSM Using a System Derived PSM**

Figure 7 shows a knowledge level model representing the inference layer for problem solving within task 4, from the perspective of a single search and rescue incident. The model represents the inference types and domain roles that exist within task 4 reasoning. The shaded boxes represent support knowledge for a particular inference. This is a KADS like model that is expressed in domain terminology. It was constructed through a lengthy process of second stage KA involving taped structured interviews, video tape analysis, protocol analysis, incident documentation and structured analysis of specific incident cases. The model represents the process of reasoning from an high level plan to an executable plan; converting high level goals into an ordered set of physical actions to be carried out by the RCC. A detailed description of Figure 7 is given in section 4.2.






**Figure 7: Domain Problem Solving Model for asset utilisation in SAR**

The domain PSM (Figure 7) was validated via a process of lengthy discussion with multiple domain experts. This discussion was based upon relating the model back to specific incidents in SAR, in order to confirm that all cases could be characterised accurately within the model. A knowledge level methodology advises that the next stage in the development lifecycle should be the population and refinement of the PSM with domain knowledge. At this point we departed from the suggested course of development. Rather than launching directly into the domain knowledge acquisition, we wanted first to consider operationalisation issues.

The reason for this is a common problem occurring in KBS development that relies upon a domain inferred PSM. Either during the design process or the actual implementation, it often becomes apparent that there are vital knowledge elements missing from the original PSM. This potential incompleteness of the PSM is a recognised problem in domain driven knowledge acquisition [Ford & Bradshaw, 1993]. Our proposed method of overcoming this was to select a generic system based planning PSM, and attempt to establish a mapping into this from the domain based PSM. The rationale behind this was that the system PSM would be complete and sufficient, because it represented an operationalised architecture. If a clear mapping existed between the elements of the PSMs, then the system PSM may aid us in anticipating any omissions in the domain PSM, that would compromise its ability to produce decisions.

The system PSM that we selected for this purpose was derived from the Open Planning Architecture (O-Plan) [Currie & Tate, 1991]. The O-Plan system is a generic computational planning architecture. The generic nature of the system and the fact that initial CommonKADS models of O-Plan had been proposed [Kingston *et al*, 1996], made this an attractive option for our purpose. We established the existence of a mapping from the ontological elements in the domain PSM to corresponding elements in the O-Plan PSM (Figure 8). This approach successfully aided us in the identification of knowledge omissions in our domain PSM, prior to the completion of KA. This two model approach presents a number of advantages:

- comparison between domain and system based PSMs facilitates the early identification of omissions in the domain PSM, without enforcing a system structure upon it.
- the explicit mapping between the two types of PSM provides a mapping between domain-specific terminology and generic planning terminology.



**Figure 8: System Problem Solving Model**

Both the system PSM (Figure 8) and the domain inferred PSM (Figure 7) possess inference steps which allow the transition from one form of knowledge to another. In this section we describe the inference steps and knowledge roles in the domain based PSM, and how these items map to the O-Plan based PSM. Although the two models initially bear little resemblance to each other, there is in actual fact a clear mapping between both the ontological elements and inference steps that are depicted in the domain PSM to corresponding items within the O-Plan PSM. Both structures represent the matching of goals to possible actions. Both PSMs also facilitate the decomposition of goals and identify the selection of the next goal as an important inference step in the planning process. The comparison showed that the O-Plan PSM had a richer representation for the selection of goals, highlighting the necessity for knowledge that supports this inference step

in the architecture of any intended system. The comparison had therefore successfully identified weaknesses in the domain based PSM.

## The Mapping Between the Models

The following is an explanation of the knowledge roles and inference steps in the domain PSM (Figure 7), accompanied by their mappings into the O-Plan PSM (Figure 8):

**4.2.1 Template and Goal Selection: Select template** -- The input to this is the world state, and the output is a plan template consisting of high level goals that correspond to a generic type of incident. They have a partial temporal ordering, yet it is only when the RCC plan the application of resources to an incident that these goals are more fully defined and ordered. In complex situations it may be the case that no plan template exists and one will have to be constructed. This then requires a knowledge of temporal constraints on the goals that make up a plan template. This template maps to the initial agenda in the O-Plan PSM. The plan template is a set of high level goals to be resolved, and the agenda consists of a set of issues to be resolved. In this case, the goals in the domain map to outstanding issues in the O-Plan PSM. The plan template in the domain PSM is selected at the commencement of planning for an incident.

**Select Goal** -- represents the selection of a goal from the plan template. This is often simple, corresponding to the default ordering of goals defined in the plan template. However, in a complex or rapidly changing situation, the selection of goals becomes more knowledge intensive. It is here that our model comparison was informative, as it suggested the existence of certain types of control knowledge affecting goal selection. It is clear that in the O-Plan PSM there is a much richer representation of the knowledge affecting which goal/issue to resolve next. These are described in the O-Plan PSM as three possible expansions of the *match-3* inference step, which is marked in bold in Figure 8 to indicate its importance. The three expansions represent three different ways in which O-Plan can attempt the resolving of issues. The expansions are depicted in the O-Plan PSM as three separate inference structures. Two of these structures have clear mappings to the domain PSM (the decompose expansion is described in the next section, Figure 9). The third does not and represents knowledge about goal selection that is not accounted for in the domain PSM. In O-Plan this knowledge drives a backward chaining search process that decides which issues to resolve when the present issue's conditions are not satisfied. Issues are selected in order to achieve actions that will satisfy the original unsatisfied conditions. This backward chaining process caused by interaction between issues had not been considered in the domain PSM. We had merely considered basic dependencies between goals, of the form Goal A must be satisfied before Goal B can be considered. The comparison between PSMs suggested a deeper form of knowledge about interdependencies between conditions necessary to resolve goals and the actions that satisfy them. This knowledge will provide important support to the user when incidents become complex and the default ordering of goals described by the plan template may not be applicable. The SAR domain has examples of the need for this "backward chaining" e.g. if the RCC are in charge of an incident involving mountain rescue in poor visibility the default ordering of the high level goals may not apply. The plan template for this incident places "scramble resources" before "ascertain weather". The high level goal "scramble resources" decomposes to "scramble helicopter", "scramble mountain rescue team" or both, depending upon the world state. A condition of this decomposition will be that visibility at the incident scene is sufficient for a helicopter to operate. If the visibility condition is unknown then the high level goal "ascertain weather" will be initiated in order to effect an action that will satisfy this condition.




Figure 9: Expansion of match-3 for issue decomposition

**4.2.2 Goal Decomposition: Decompose** -- This decomposes a goal into a sub-goal. Similarly in O-plan issues may be decomposed into sub-issues. There is a clear correspondence with one expansion of the *match-3* inference step in the O-Plan PSM (Figure 9). The degree to which goals are decomposed by RCC varies. Some high level goals match to high level actions that consist of an invariant ordered set of physical actions. The existence of such invariant high level actions and their suitability to satisfy high level goals are factors affecting the degree to which goals must be decomposed. In some cases goals will have to be decomposed to the granularity of physical actions. It would seem that decomposition increases when an incident and its associated goals are out of the ordinary. Intuitively the level of reasoning increases in the exceptional cases.

**4.2.3 Matching Goals to Actions:** These inference steps involve finding actions that can fulfil (or help to fulfil) goals. There will generally be multiple actions that can fulfil a particular goal. The match therefore depends upon the present world state. There is a mapping between these two inference types and the third method of expanding match-3 in the O-Plan PSM.

**Match 1** -- The input to this is a high level goal. The output is a high level action. This step will require a knowledge of high level actions that satisfy high level goals, and the conditions that make the match valid.

**Match 2** -- The input to this is a low level goal, corresponding to a physical action. The output is a lower level action. This is similar to the Match 1 inference type, though it is actually matching to a physical action.

**4.2.4 Assembling the Executable Plan: Assemble 1** -- The input to this is a high level action, corresponding to an high level goal. This high level action consists of a set of ordered physical actions or of sub-actions. The output is the current executable plan. As discussed earlier the executable plan is gradually formulated throughout the course of the resources application to the incident. This is supported by a knowledge of action ordering, though a lot of this ordering will have been decided at the goal ordering stage.

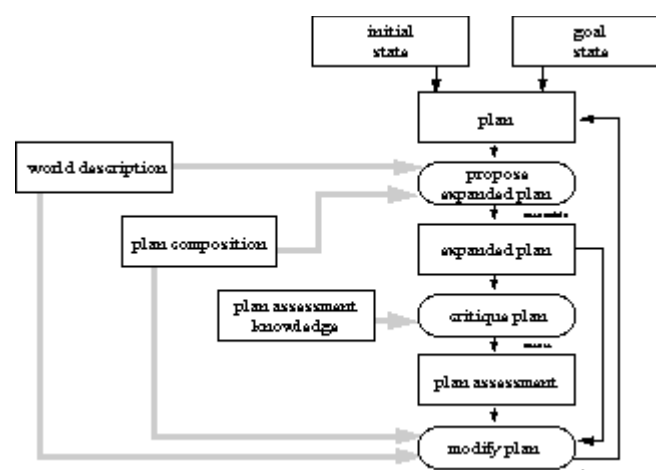
**Assemble 2** -- The input to this is a set of lower level actions; the lowest level being physical actions. The output is a higher level action; the highest being those that actually make up the executable plan (i.e. those that correspond to the high level goals).

There is no replication of these assembly inference steps in the O-Plan PSM, although typically any set of actions in O-Plan have ordering constraints instantiated during planning and are therefore assembled implicitly within the plan. The explicit assembly of actions is however important to the RCC, as it serves to summarise what has been done (or what is intended to be done), in order to achieve goals. This reflects the nature of the RCC's planning, which proceeds in small chunks corresponding to the goals in the template. There may be activity in several chunks at once, though this tends to be the exception rather than the norm.

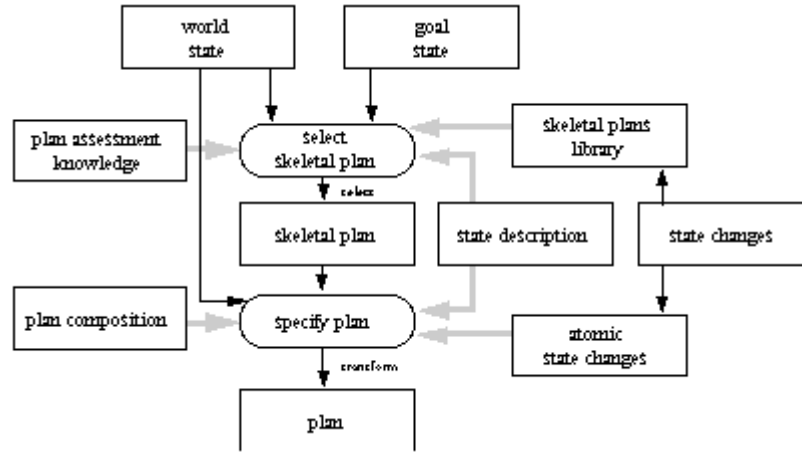
There are clearly three inference steps in Figure 8 (besides the "backward chaining" expansion of *match-3*) that are not represented in the domain PSM. The *modify-6* step describes how the actions carried out in the plan modify the world state (world state is part of current plan state as it includes constraints). This could be included in the domain PSM as a step from emerging plan back to world state. The *modify-5* and *specify-4* steps describe how intended actions in the plan cause new issues to arise, and this therefore modifies the current plan state. These are not represented in the domain PSM, and this form of knowledge has not been observed in the domain. This is probably due to the previously mentioned least commitment strategy of the RCC decreasing the amount of intended actions that exist within their plans. This type of knowledge may be important in forms of SAR incident that we have not yet witnessed, and we regard this as an important area for future KA.

## COMPARISON WITH EXISTING PSMS FOR PLANNING

In this section we compare our domain PSM developed for knowledge acquisition in SAR (Figure 7) with two of the planning models available in the CommonKADS library [Valente, 1994]. These are high level models of the inference structures of a non-linear planner (Figure 10) and a skeletal planner (Figure 11). We also compare the CommonKADS planning models to the O-Plan PSM, in order to consider how effective these models would have been if used in the critiquing role described above.



**Figure 10: CommonKADS function structure for non-linear planner**



**Figure 11: CommonKADS function structure for skeletal planner**

Initially the SAR domain seemed to match a form of non-linear planning such as described in Figure 10. However with hindsight the concept of plan templates became very important within the knowledge acquisition. In actual fact planning for SAR lies in-between non-linear planning and the skeletal planning described in Figure 11. At the RCC non-linear planning takes place around a framework of default skeletal plans (plan templates). This is caused by particular types of rescue incident possessing a common structure. The goals and their ordering are thus implicitly defined by the type of incident, yet more detailed explicit knowledge may override this when necessary. Had we used either of the models below as the basis of a PSM for knowledge acquisition, we would have neglected one of the major defining characteristics of the expert problem solving that we were attempting to model.

Because planning for SAR is a continuous process the initial state and goal state are constantly changing. Planning takes place by selecting a high level goal which then becomes decomposed into sub-goals. They continue being decomposed until appropriate actions can be assigned to the sub-goals. This takes place in the context of the actions that have already been executed and the high level goals that are still to be addressed. Goals and the actions that will achieve the goals thus become very prevalent within the RCC's problem solving. Each high level goal is a planning cycle in itself. Both the non-linear planner model and the skeletal planner model are not broken down as far as the level of actions and goals, though it would be reasonable to refine either of these models to include these entities. However we believe that if this approach had been used, too much emphasis would have been placed on knowledge roles such as plan assessment knowledge, initial states and goal states. In the context of the RCC these knowledge roles are much less prevalent than would be suggested by PSMs such as Figures 10 and 11. By less prevalent we mean that although these roles are present within the domain they are seldom actually utilised in problem solving. Therefore it would not be fruitful for us to attempt to base our PSM around such secondary knowledge roles.

The key characteristic of the domain PSM is that it represents planning in situations of high unpredictability. Because this effectively limits the amount of planning that can be usefully carried out the knowledge roles that are most prevalent in problem solving are different from those in domains of greater predictability. It is possible that the system PSMs shown in Figures 10 and 11 would be more applicable in such domains. Planning is sometimes viewed as a sub-task of design and it would be useful to consider PSMs for design in some planning domains; particularly considering the wealth of analysis that has been done in this area (e.g. VT elevator design [Marcus *et al*, 1988]). However we believe that the high unpredictability within SAR planning makes it quite unlike such design tasks. It would correspond to a design task that required the continuous production of partial sub-task solutions within an environment of constantly changing constraints.

Off the shelf system models are useful in informing us about the knowledge roles and inference steps that are necessary for the operationalisation of the knowledge in a system. The O-Plan PSM is the most suitable in this instance as it is already decomposed to the level of goals and actions (called issues and activities in the O-Plan terminology), enabling us to establish a mapping between the domain PSM and the system PSM.

## SUMMARY

If we consider the generic PSM as a framework that the knowledge engineer must map their domain level knowledge into, then it is vital that the structure of this framework is as close to that of the domain as possible. The greater the difference between the structure of the generic PSM and the structure of problem solving in the domain then the greater the difficulty that will be encountered by the knowledge engineer in mapping the domain entities into the PSM. The crux of our approach is to arrive at a PSM that is not just sufficient to characterise the problem solving, but whose problem solving structure reflects accurately that in the domain. Thus our goal in knowledge engineering is emulation of the

domain. We accept that this may not always be a priority in the knowledge engineering process. If the purpose of using a knowledge modelling technique is just to support modular KBS design then utilising a system derived PSM may prove more useful. System derived PSMs are also very important in situations where a pre-defined operational system is to be deployed in a particular domain [Major *et al*, 1994], though a two model approach might present advantages in this situation.

Because our PSM had to be inferred directly from the domain this changes the emphasis of KA, and also the emphasis of the methodology developed as a result of the project. The methodology developed concentrates on supporting the earlier stages of KA. Support is provided in the form of techniques and advice on the construction and critiquing of domain inferred PSMs in planning type domains. The methodology also provides a new generic PSM for planning, that is hoped will have applicability in a number of other areas. Many domains that involve some form of resource allocation appear to have a similar structure in their problem solving to the PSM we have uncovered.

KA was also focused and directed by decisions on the intended functionality of the demonstrator system. Two of these guiding principles were (i) that the system should follow the Rescue Co-ordination Centre (RCC)'s natural idiom (*ie*: provide a similar interface/environment) as far as possible, and (ii) that the system should act as an assistant, supporting the operator, and not as a "*black box*" which takes control away from the operator.

The completed demonstrator system supports visualisations that correspond with those in the RCC's current environment. It also provides additional explicit representations of the activity plan for an incident. These representations include a TODO list and a PERT chart. The TODO list provides the operator with a view of the "*active edge*" between tasks which it is possible to perform now and those which are not yet possible. Thus the operator is free to choose the order in which he performs tasks, except where the knowledge elicitation identified a necessary ordering between tasks. The PERT chart provides a good overview of the plan, and is therefore especially useful in (i) tutorial situations, and (ii) handing over at the end of a shift. Both the TODO list and the PERT chart visualisations support the hierarchical exploration of plans. This allows the operator to view plans in more or less detail. Some of the operators' mundane work is removed by the system automatically logging events that are activated through the system, and by automatically updating all visualisations as information changes. A single underlying model for the various representations and visualisations ensures that consistency is maintained.

The major achievement in the development of the demonstrator system was the single underlying model that provides a basic planning infrastructure. This intentionally remains hidden from the user yet effectively facilitates all the visualisations that comprise the outward appearance of the system, and enable the maintenance of consistency across visualisations.

## CONCLUSIONS

The work described represents the construction and demonstrated use of a domain driven knowledge level modelling approach to KBS development in a planning domain. We make a definite distinction between problem solving models that are inferred from the domain and those that have been derived from systems. Indeed one insight derived from the project has been that the limited number of existing PSMs for planning are almost exclusively system derived models. We recognise that current model libraries (such as CommonKADS) do not claim to be comprehensive, and this system bias in planning PSMs reflects the manner in which knowledge level approaches have been utilised so far in the generic area of planning. PSMs have been harnessed within planning in order to facilitate the knowledge level analysis of systems, as opposed to the knowledge level acquisition of human expertise.

This observation is important if we are to maintain a level of independence between analysis and implementation within KBS development. Domain derived models do not presently exist for the support of system development within the generic area of planning. Our work merges a knowledge level modelling approach with the work that has been done on ontologies for planning, in order to formulate a generic approach to the acquisition and utilisation of knowledge for planning systems. The approach was tested and refined through the development of a knowledge-based system for the support of planning for search and rescue.

The distinction that we have made between domain and system based models, led us to investigate possible benefits that could be gained by exploring the mappings between these models. In the context of the SAR demonstrator development, we discovered that the comparison of these models enabled us to identify omissions in the domain model. It also enabled us to identify specific areas for future KA. We believe this twin model approach may have more general applicability for KBS development.

During a KBS development lifecycle there must be iteration between the developmental stages. Later stages of development iteratively inform and revise the earlier stages. It is expected that this iterative cycle will improve the final

product. The price paid is that a large amount of iteration in the development lifecycle increases the effort expended. Approaches that enable the detection of shortcomings in the earlier stages represent a potentially large saving in development effort. The two model approach described demonstrates such a saving in the case of domain driven knowledge level modelling for planning.

This approach offers a high degree of support and advice for the early stages of KA in planning domains. It principally supports the construction of a high level PSM that reflects the structure of problem solving in a planning domain (this construction may be replaced by generic PSM selection as a comprehensive library of generic domain derived PSMs is evolved). It also offers a high degree of support for converting the completed PSM into the high level system design (including the appropriate system knowledge structures).

The project has produced a new generic model for planning and has discovered a set of generic plans for SAR in the form of plan templates. It has made important observations on the derivation of PSMs, and clearly outlined the distinction between those that are system derived and those that are derived from the analysis of expert problem solving. This distinction was explored in depth within the context of the SAR demonstrator development. The project has discovered both the implications of the distinction and methods by which we can effectively exploit these to our advantage.

The early stages of KA support contained in the methodology as it currently stands require further validation and refinement through the practical application of the methodology in other planning domains. We need to derive more domain PSMs in other planning domains. A process of comparison can then take place both between these domain PSMs and with system PSMs. Such a comparison might be informative in suggesting modifications to present planners in order to more effectively deploy them in selected domains.

## REFERENCES.

Breuker, J. and van de Velde, W.(Eds.). (1994) *CommonKADS Library for Expertise Modelling*. IOS Press, Amsterdam, September 1994.

Breuker, J., Wielinga, B., Van Someren, M., De Hoog, R., Schreiber, G., De Greef, P., Bredeweg, B., Wielemaker, L., Billault, J-P., Davoodi, M. and Hayward, S. (1987) *Model driven knowledge acquisition: interpretation models*. KADS-I Project Deliverable, University of Amsterdam, Holland 1987.

Barros, L., Valente, A. and Benjamins, R. (1996) *Modelling Planning Tasks*. In Proceedings of the Third International Conference on Artificial Intelligence Planning Systems (AIPS 96) Edinburgh, Scotland. pages 11-18. 1996.

Clancey, W.J. (1985) *Heuristic classification*. Artificial Intelligence, 27: 289-350, 1985.

Cottam, H., Shadbolt, N., Kingston, J., Beck, H. and Tate, A. (1995) *Knowledge Level Planning in the Search and Rescue Domain*. In Bramer, M.A., Nealon, J.L. and Milne, R., (eds.) Research and Development in Expert Systems XII: Proceedings of BCS SGES Expert Systems' 95, pages 309-326. SGES Publications. 1995.

Cottam, H. and Shadbolt, N. (1996) *Domain and System Influences in Problem Solving Models for Planning*. In Proceedings of the Ninth European Knowledge Acquisition Workshop (EKAW 96) Nottingham, UK. pages 354-369, 1996.

Currie, K.W. and Tate, A. (1991) *O-Plan: The Open Planning Architecture*. Artificial Intelligence, 52 (1): 49-86, 1991. Also available as AIAI-TR-67.

Ford, K.M. and Bradshaw, J.M.(Eds.). (1993) *Knowledge acquisition as modelling*. New York: Wiley. 1993. Kingston, J.K. (1995) *CommonKADS Models for Planning tasks*. forthcoming AIAI technical report, 1995.

Kingston, J., Shadbolt, N., and Tate, A. (1996) *CommonKADS Models for knowledge-based Planning*. to appear in Proceedings of AAAI 1996.

Marcus, S., Stout, J. & McDermott, J. (1988) VT: An Expert Elevator Designer That Uses Knowledge-Based Backtracking. AI Magazine, Vol. 9, No. 1. pages 95-114, 1988.

Major, N., Cupit, J. & Shadbolt, N. (1994) *Applying the REKAP methodology to situation assessment*. in the Proceedings of the European Knowledge Acquisition Workshop, EKAW 94. Brussels. Also published in Voss, H. and Studer, R. (Eds.) Proceedings of the 4th International KADS Meeting, Bonn, Germany. Arbeitspapiere der GMD, 832, March 1994.

Motta, E., O'Hara, K., Shadbolt, N., Stutt, A. & Zdrahal, Z. (1994). *A VITAL Solution to the Sisyphus II Elevator Design Problem*. Proceedings of the Eighth Banff Knowledge Acquisition for Knowledge Based Systems Workshop, Banff, Alberta, Canada. 1994.

O'Hara, K., Shadbolt, N. R., Laublet, P., Zacklad, M. & Le Roux, B. (1992). *Knowledge acquisition methodology*. VITAL deliverable DD212. Nottingham University, UK. 1992.

O'Hara, K. (1993). *A Representation of KADS-I Interpretation Models Using a Decompositional Approach*. In C. Lockenhoff, D. Fensel & R. Studer (eds.) 3rd KADS Meeting (Siemens AG, Munich) 147-69. 1993.

Tate, A. (1994) *"Plan Ontology", a paper to the Workshop on Ontology Development and Use*. San Diego, California, USA, November 1994. (Also available as AIAI Technical Report AIAI-TR-161, and through URL: <ftp://ftp.aiai.ed.ac.uk/pub/document/1994/94-krsl-plans.html>)

Valente, A. (1994) *Planning models for the CommonKADS library*. ESPRIT Project P5248 KADS-II KADS-II/M2.3/UVA/56/1.0, University of Amsterdam, 1994.

Valente, A. and Löckenhoff, C. (1993) *Organization as guidance: A library of assessment models*. In Proceedings of the Seventh European Knowledge Acquisition Workshop (EKAW 93) pages 243-262, 1993.

[+] The work described has been done under contract to the Defence Research Agency Flight Systems Division, Farnborough (Contract No: AS/897/U). The project was a joint effort between the AI Group at Nottingham and AIAI at Edinburgh. Parts of the work reported have been published previously [Cottam *et al.*, 1995.][Cottam and Shadbolt, 1996].