Towards a knowledge-based system prototype for Aeronautical Search and Rescue Operations

Irène Abi-Zeid
Opérations et systèmes de décision
Faculté des Sciences de l'Administration
Université Laval, Québec (Québec), G1V 0A6
Canada
irene.abi-zeid@osd.ulaval.ca

Oscar Nilo
Opérations et systèmes de décision
Faculté des Sciences de l'Administration
Université Laval, Québec (Québec), G1V 0A6
Canada
oscar.nilo@sympatico.ca

Stéphane Schvartz
Département d'informatique et de génie logiciel
Faculté des sciences et de génie
Université Laval, Québec (Québec), G1V 0A6
Canada
stephane.schvartz@bbox.fr

Michael Morin
Département d'informatique et de génie logiciel
Faculté des sciences et de génie
Université Laval, Québec (Québec), G1V 0A6
Canada
michael.morin.3@ulaval.ca

Abstract - The long-term objective of our project is to develop a knowledge-based tool for Search and Rescue (SAR) operations to support a Canadian search mission coordinator in determining the likely location of a missing aircraft overland. In order to attain this objective, we used a knowledge engineering approach to acquire, structure and model SAR experts' knowledge. This knowledge was modeled and implemented in a knowledge-based system prototype. The input to the interactive prototype consists of the known information regarding a given SAR case. Its main output is a set of scenarios describing the various hypotheses on what might have happened to the missing aircraft, why and where, the plausible routes followed, as well as the possibility area, defined as the region most likely to contain the missing aircraft. In this paper, we introduce the knowledge model, present an application example and briefly describe the prototype.

Keywords: Search and Rescue, knowledge-based system.

1 Introduction

Search and Rescue (SAR) is one of the greatest humanitarian activities. Steve Fossett is one famous SAR case. John F. Kennedy Jr. is another famous SAR case. For every famous case, there are dozens of non-famous ones. In Canada, there are thousands of SAR cases every year where tens and sometimes hundreds of lives are lost.[1] Coordinating SAR operations is knowledge intensive. Case prosecution, the main function of a search coordinator, is conducted in response to the occurrence of an incident. Following the receipt of an alert, the coordinator must determine if an emergency situation exists and the degree of emergency. If the missing aircraft is not located within a given period of time, a major search operation is then initiated.

An important step prior to initiating a major search is the definition of the possibility area, the area most likely to contain the search object. For marine searches in Canada, the computer tool CANSARP helps in defining the possibility area by developing different scenarios and probability maps related to the possible location of the search object(s) [2]. For aeronautical searches overland, the current approach in non-mountainous regions, CSAD[2], is based on an empirical distribution of incidents, over six years, where the missing aircraft was located [9]. For mountainous regions, the MVFR[3] method is used. These methods were developed for cases where there is little information to use besides a last known point and a destination. In practice, a search coordinator modifies the initial CSAD or MVFR possibility area as a function of the information available on the missing aircraft, the pilot, the weather at the time of the incident, the terrain topography, etc. [5]. In the absence of a standardised reasoning method, coordinators employ heuristic approaches based on their intuition and experiences.

In recent years, the need for developing a decision support system for overland aeronautical SAR mission planning was identified in Canada. There were two main requirements: a planning module for optimal effort allocation, and a knowledge module for capturing knowledge and expertise. To address these requirements, a decision support system for the optimal resource allocation

[1] Statistical information can be obtained from the National Search and Rescue Secretariat at: http://www.nss.gc.ca.

[2] Canadian Search Area Definition

[3] Mountain Visual Flight Rules

of overland aeronautical SAR, SARPlan, was developed [1]. One of the main inputs to SARPlan is the possibility area, a region manually defined and drawn by the user (Figure 1). As a consequence, SARLoc, a project to develop a tool that can assist the coordinator in defining the possibility area, was initiated to capture the current know-how and knowledge of experienced coordinators whose number is dwindling in the Canadian Forces. Both the SARPlan and SARLoc projects were funded by the National Search and Rescue Secretariat[4] via the National SAR Initiative Fund.

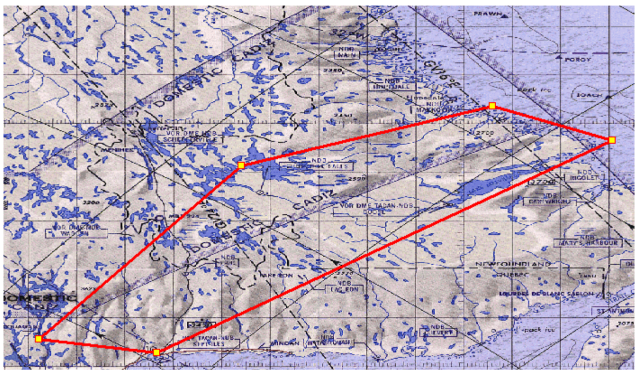

Figure 1: A possibility area in SARPlan

The main objective of the SARLoc project was to capture the expertise of search mission coordinators, to structure this expertise and to develop a knowledge model as a structured method that can guide a coordinator in the definition of the possibility area of a civilian aircraft missing overland in peacetime.

The research questions that we wished to answer were the following: How do coordinators determine the possibility area? How can we describe and model their knowledge and their problem resolution process? How can we acquire, formalize and represent this knowledge at a conceptual level that can eventually lead to a knowledge-based system (KBS) prototype?

2 Knowledge modeling

In order to find answers to these questions, we positioned our qualitative study in a knowledge engineering context. We used the CommonKADS methodology [10] to help us develop conceptual knowledge models. The role of knowledge models is to describe human tasks that are complex and require a lot of expertise. They allow one to acquire domain knowledge and clarify the structure of a knowledge-intensive task at a level independent of the computer implementation. Although some papers have been published on knowledge engineering in SAR, ([3]; [8]), none have addressed the possibility area problem.

Our knowledge acquisition phase was based on knowledge available in documents, on interviews with domain experts and on participation in training and simulations. These activities helped us identify the main concepts involved in the definition of a possibility area (Figure 2).

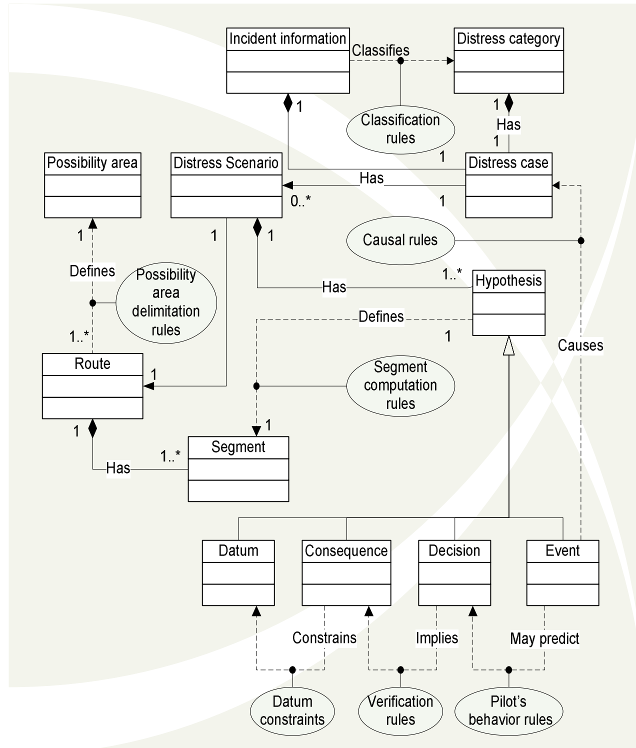

Figure 2: The main concepts for defining a possibility area

A *distress case* contains available *incident information* such as a flight plan, last known position, departure point, destination point, aircraft type and configuration, aircraft autonomy, equipment on board, etc. The available information can help classify the case in a *distress category*. In our prototype, the availability of such information is captured by the interfaces on Figures 3 and 4. Figures 5 and 6 capture the information on the navigation equipment available on board along with the aircraft type.

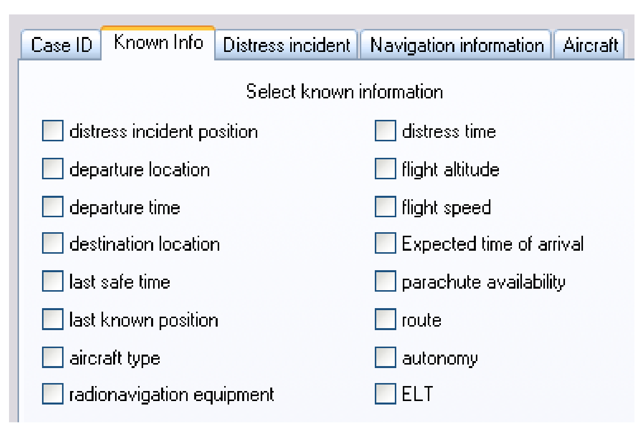

Figure 3: Incident information

[4] http://www.nss.gc.ca/site/index_e.asp

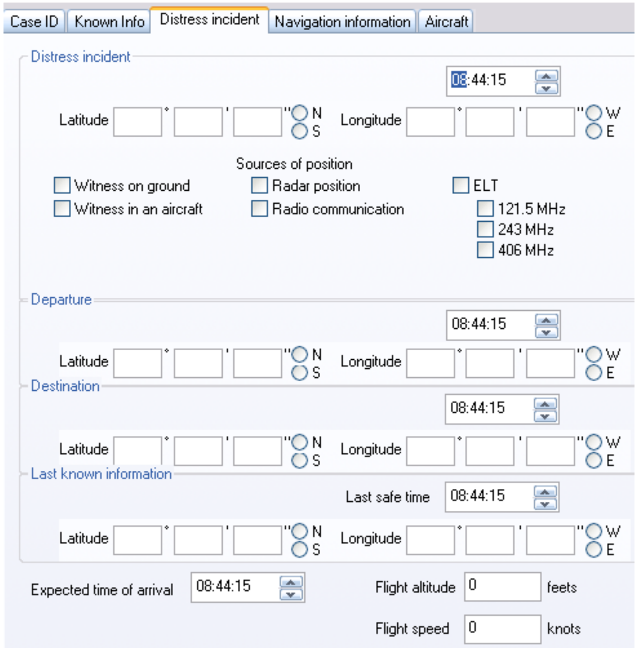

Figure 4: Incident location

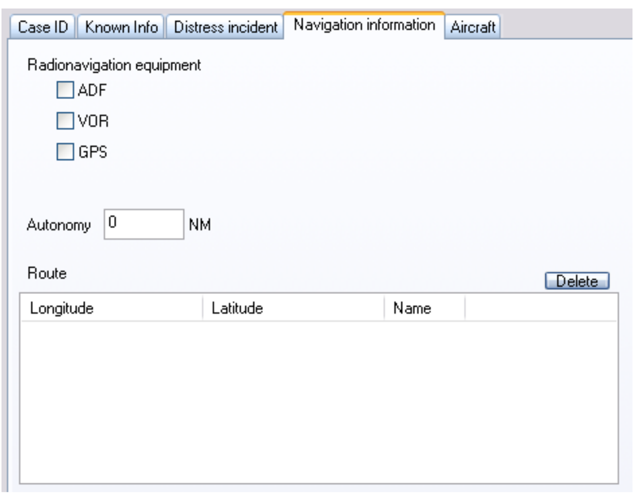

Figure 5: Navigation equipment information

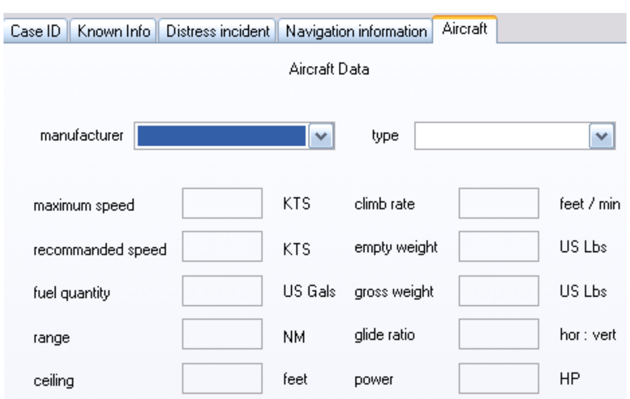

Figure 6: Aircraft information

A *distress scenario* is built for a distress case. It describes the chain of events and is composed of four hypotheses levels:

- ***Event hypotheses*:** These hypotheses relate to a possible event(s) that occurred to the missing aircraft (in peacetime) preventing it from arriving to the planned destination. Knowledge acquisition enabled us to formulate six possible hypotheses: an encounter with a weather barrier, with a topographic barrier, insufficient fuel, a mechanical problem, navigation error, or medical problem.
- ***Decision hypotheses*:** In the face of an unexpected event, the pilot normally reacts and makes a decision on the next course of action. These hypotheses relate to the plausible reactions of the pilot. The occurrence of a sufficiently serious event normally requiring modifications to the flight plan, namely the route initially planned by the pilot. These hypotheses allow the coordinator to presume the pilot's behaviour in reaction to the event in question. The choice of hypothesis relating to the most plausible pilot's decision is based on the information available regarding the pilot's profile, the supposed event, the conditions in which the coordinator believes the flight was carried out, and so forth.
- ***Consequence hypotheses*:** Given an event and a pilot's decision to act, the outcome might be a success or a failure, meaning that the pilot succeeded in executing his plan or not. These hypotheses reflect the possible consequences of the pilot's decision following the occurrence of a problematic event and his actions. The selection of the most plausible consequence, be it the success or failure of the action undertaken by the pilot, is influenced by information such as the pilot's experience and qualifications, the scale of the supposed event, the performance of the aircraft, and so forth.
- ***Datum hypotheses*:** These hypotheses are related to specific geographical positions in the vicinity of the intended route (either initial or modified) towards which the pilot would have been able to head based on the hypotheses previously assumed (*e.g.*, an open area to carry out an emergency landing).

An event-decision-consequence hypothesis triplet defines a path *segment*. A *scenario* is composed of one or many hypothesis triplets. A hypothesis triplet is what the coordinator believes has happened (*event hypothesis*), how the pilot reacted (*decision hypothesis*) and what the result of the pilot’s reaction was (*consequence hypothesis*). *Datum hypotheses* may be associated to a hypothesis triplet. A *route* is associated to a *scenario* and is composed of one or many *segments* including possibly a datum (or many). A resulting *possibility area* is centred on one or many *routes* that the coordinator believes the pilot may have followed.

In the following phase of our project, we used these concepts to elaborate a reasoning model for determining the possibility area. This structured procedure is described by Figure 7.

First, a distress case must be **classified** in a distress category based on the available information and on "distress category classification rules". The *distress category 1- certainty* is a simple situation where the distress position is known and no further investigation is necessary. The *distress category 3- total uncertainty* corresponds to a situation where crucial information is missing such as the pilot's destination. For these two extreme categories, there exist, in SAR manuals such as the NSM [5], precise procedures that can be followed to delimit a possibility area. The most interesting case is the one where partial information is available, such as distress position is unknown, destination is known, intended trajectory either known or unknown, which leads to a *category 2- partial uncertainty*. This is the most frequent situation where diagnosis, the most knowledge-intensive sub-task, is required. The application of the reasoning model associated with diagnosis leads to distress scenarios defined by hypotheses and associated routes.

Given a hypothesis triplet, the diagnosis process can be terminated or can continue if the coordinator believes some other event has occurred. In this case, another hypothesis triplet can be constructed by the coordinator and a new segment added to the presumed route. In addition, it might be pertinent to determine specific positions where the pilot might have crashed or landed in the vicinity of the routes constructed. These positions are also included in the possibility area. For example, if the coordinator believes that during the flight the pilot has figured out that he will not be able to reach his destination for whatever reason (lack of fuel, mechanical problem, physiological problem, etc.), the coordinator may look for spots where the pilot might have landed. A possible landing spot is a *datum*. Obviously, the type of aircraft will determine the possible landing spots available for landing in a given region. For an aircraft on wheels, secondary airfields, highways, long straight portions of roads are possible landing spots. For an aircraft on floats, lakes with sufficient length are possible landing spots.

Taking into account possible landing spots is relevant when the coordinator tries to determine alternative courses for the aircraft. In such cases possible landing spots that could have been reached by the aircraft should be identified because they correspond to alternative destinations that the pilot might have tried to reach. In cases of mechanical, electrical failure of the aircraft or of pilot's physical problems, one can assume that the pilot will try to land as soon as possible in order to avoid major damages if possible. This datum hypothesis level allows one to complete the diagnosis by finding possible emergency landing or crash points along the segments associated to each hypothesis triplet.

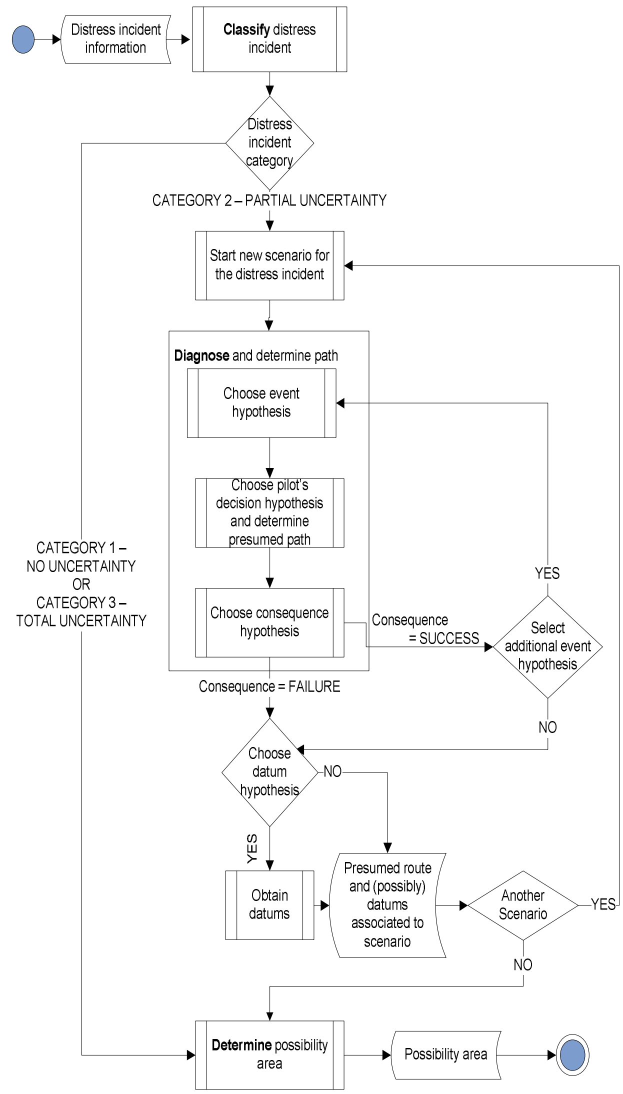

Figure 7: The reasoning model

Once the coordinator is satisfied with the scenario, the possibility area can then be created. The possibility area (for *distress category 2*) is delimited using a rectangular area of 10 Nautical Miles (NM) in width centered on the routes presumably followed by the pilot including the segments to the datum points. Note that the user might wish to study another distress scenario consisting of different hypotheses and associated routes. In this case, the two possibility areas can be merged. For the other categories, the possibility area is delimited as recommended in the NSM [5].

One of our objectives in the knowledge modeling phase was to propose a structured method to assist the coordinator in determining whether the pilot flew the direct route from the last known point to the planned destination or whether he followed an alternative route due to various unexpected events. We wished to guide the coordinator in generating various possible scenarios of what happened to the aircraft, where and why. By following this modeling approach, we were reproducing, albeit in a structured manner, the reasoning followed by the coordinator who tries to get a mental picture of what the pilot saw during his flight, and to guess what might have been his reactions. A full description of the knowledge model can be found in [12] and in [11].

3 Application example

We present an application example for defining a possibility area based on the reasoning model that we developed. The information available in this case indicates *good weather conditions at the departure airfield with noticeable degradation 90 NM to the north along the planned flight route. Skies are cloudy with a ceiling varying between 1500 ft (1000 ft at some places) and 9000 ft. Visibility is reduced to 1 NM in some places. The last radio contact occurred at 10:34, about 48 NM after takeoff, which establishes the last known point (LKP). Radar data, available until 10:29, estimate the aircraft's speed at 100 knots. It is known that the pilot had completed 350 hours of flight including about 40 hours during night and about 30 hours using instruments. According to his next of kin, he felt more and more confident in difficult weather conditions. He was known to have a GPS (Global Positioning System) on-board that he had programmed to get to his destination as well as to an alternate airfield where he could refuel.*

Based on the model developed, the reasoning could proceed as follows:

1. The coordinator confirms that he believes that the aircraft has met with a weather barrier 42 NM after his LKP (*event hypothesis*). A first segment is then drawn from the departure point to the LKP.

2. The available information concerning the pilot's profile suggests that he has attempted to cross the weather barrier (*decision hypothesis*). Since this hypothesis is retained by the coordinator, only segments crossing the weather barrier are drawn. If the pilot has decided to cross the barrier, it is likely that he has used his GPS and attempted to follow his direct route to destination.

3. Given that the pilot has little experience flying with instruments, it is possible that he tried to fly below the clouds' ceiling in order to be able to see the ground and, that he has crashed in a mountain. Therefore, the coordinator retains the hypothesis that the pilot did not succeed in crossing the weather barrier, (*consequence hypothesis*). This hypothesis triplet indicates that the presumed route stops at the exit point of the weather barrier.

4. The closest mountain to the route, presumably followed by the pilot, is considered a *datum*.

5. Figure 8 shows the planned route, its whereabouts, the weather barrier (in grey), three mountains along the route (triangles) as well as the resulting possibility area (rectangle).

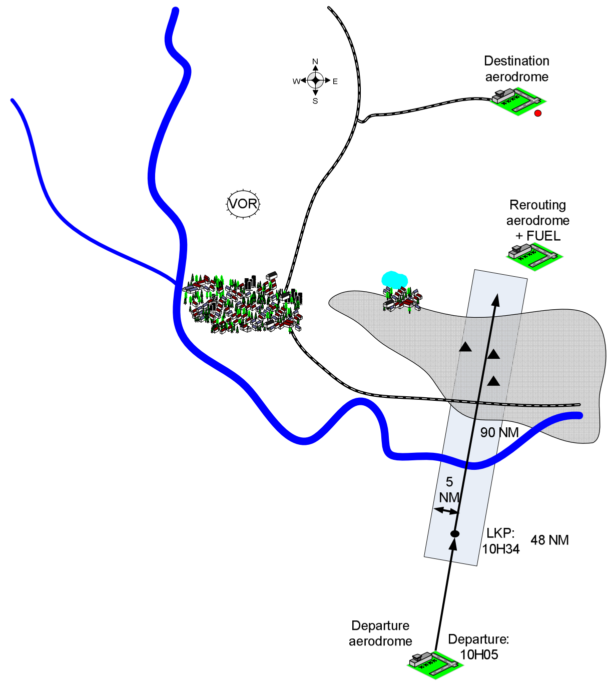

Figure 8: An example of a possibility area (not to scale)

4 KBS prototype

The reasoning model described in the previous section is meant to guide the coordinator in delimiting the possibility area by providing him with a structured procedure to follow and by identifying the decision points where he must intervene. Nonetheless, we wished to verify the feasibility of implementing this method in an advisor tool. We therefore designed an interactive prototype meant to act as a wizard. The wizard guides the coordinator in the problem resolution and structuring process. The approach we followed is rule-based where knowledge and expertise captured from the coordinator are translated into rules in the expert system component of the prototype. It was developed in the object oriented philosophy and implemented in Visual C++; it uses Oracle and Mapobjects. The knowledge bases are implemented in CLIPS [4]. The prototype contains a geographic information system module (GIS) for maps and geo-referenced data.

The prototype allows the user to build various scenarios of what happened to the missing aircraft, where and why (hypotheses on events, pilot's decisions, and likely outcome). To each scenario is associated the likely route that the pilot might have followed given the assumed hypotheses.

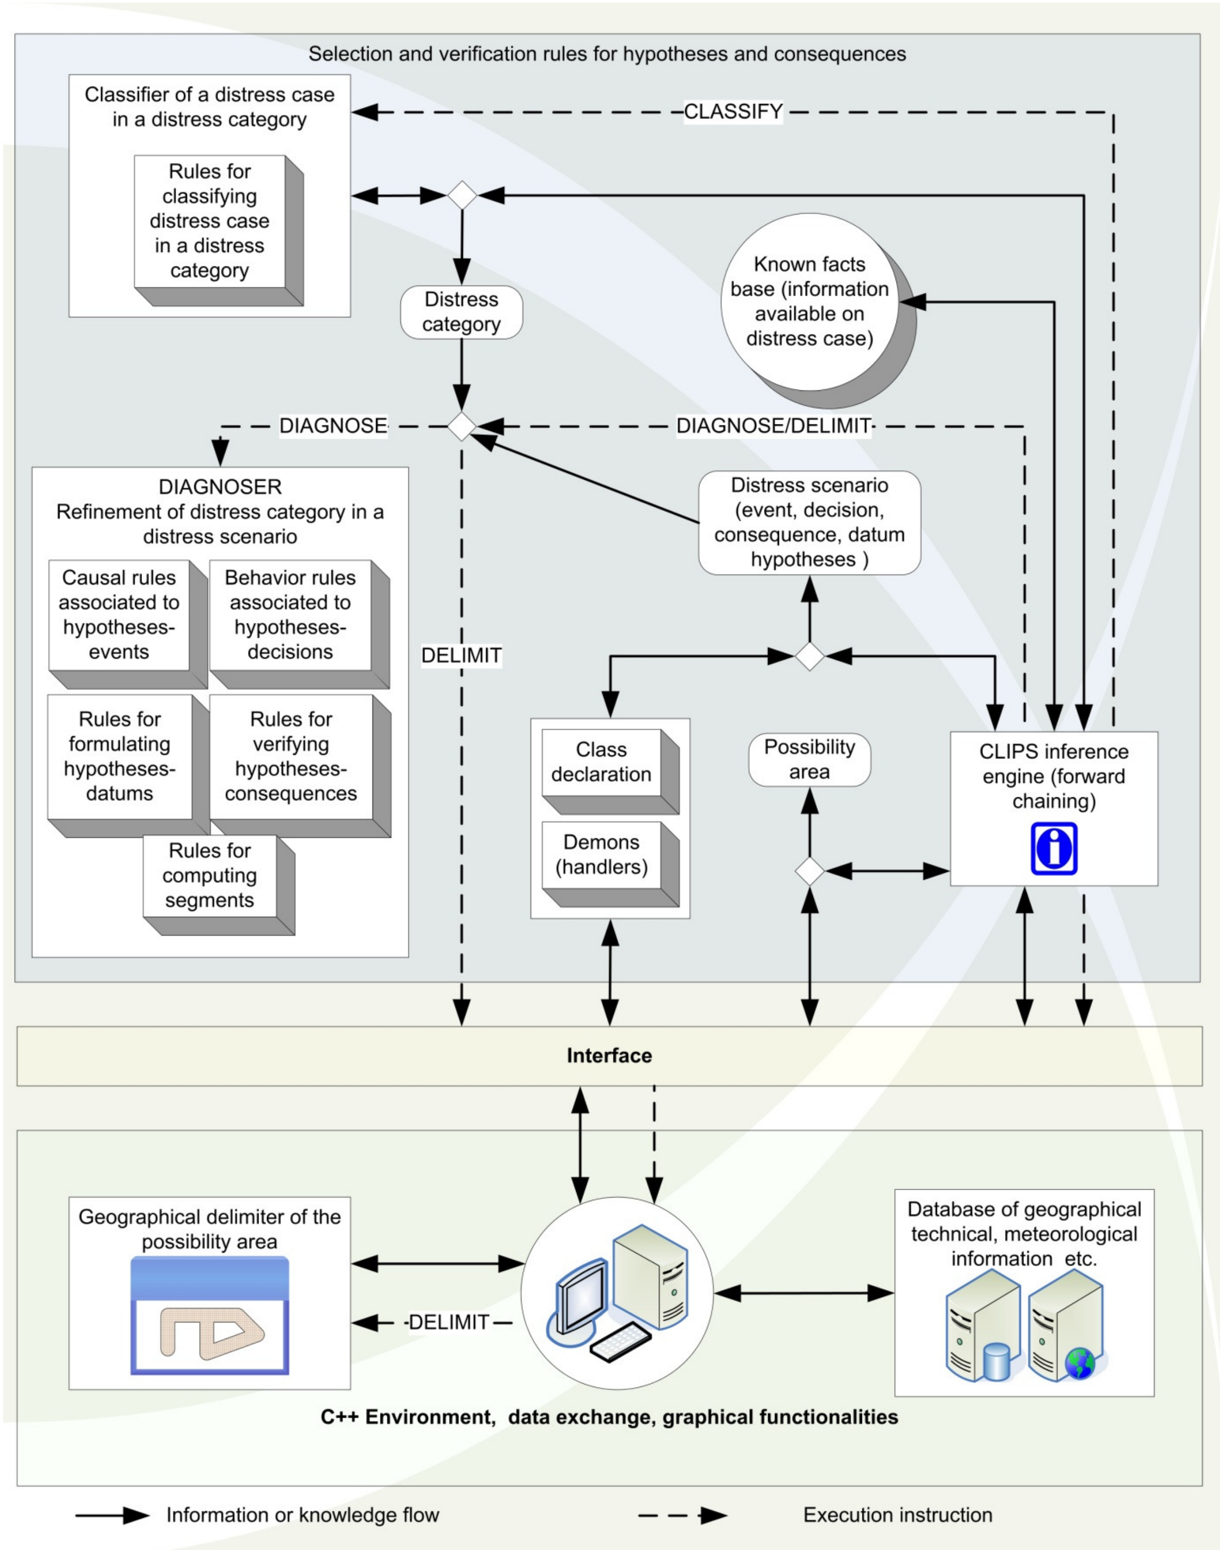

Figure 9: Architecture and prototype output

In addition to the various scenarios and possible routes, the prototype will delimit a possibility area likely to contain the search object. The main spatial factors considered in the GIS are the visual landmarks (runway, watercourse, road, railway, coast line, river, major cities, etc.) and the terrain topography. Other factors taken into account are the weather, the airplane's autonomy, and the popular visual flight rules routes (airways, reporting points).

Figure 9 shows the main components of the prototype. There are two main sub-tasks, classification and diagnosis, implemented in a **Classifier** and **Diagnoser** respectively.

The outcome of the first sub-task is a distress category, and the outcomes of the second sub-task are a distress scenario and a route. The possibility area is the outcome of the procedure **Delimit**.

The knowledge bases contain all the classes declaration and the demons that index the different hypotheses retained and that allow control transfer between sub-tasks. For example, the demons are started when objects such as distress category, distress scenario, segment route or a hypothesis are created. The implementation in C++ handles graphical functions as wells as the GIS.

The inference structure for the classification sub-task is implemented using the classification rules that assign a distress case to a distress category. The inference structure for the diagnosis task is implemented using causal rules identifying events that may lead to an abnormal situation (event hypothesis), behavior rules to link an event hypothesis to a decision hypothesis, and rules for selecting possible consequences and verifying whether the three hypotheses retained are consistent. There are also rules for defining a datum and a segment.

The application begins by capturing the available information. The CLIPS inference engine fires up the **Classifier** rules sequence to determine the distress category. At this stage, either the **Diagnoser** sub-task or the **Delimiter** module is activated depending on the distress category. *Distress categories 1* or *3* activate the geographical **Delimiter** immediately while *distress category 2* activates the **Diagnoser**. The **Diagnoser** is the most knowledge-intensive sub-task. It is the main part of the CLIPS application. It will begin by proposing plausible events that may have caused the SAR incident. The user has the possibility to override the system in the choice of the event hypothesis. The event chosen by the user as the most probable one is then processed using knowledge bases containing the behavior rules. The decision and consequence hypotheses retained by the user condition the construction of the associated segment. The outcome of the **Diagnoser** is a distress scenario, consisting of the various hypotheses retained, the segments associated to the hypotheses, the datum point(s) and all the distress case information. The presence of the distress scenario in the facts-base activates the **Delimiter** procedure that will produce the final possibility area.

The rules were implemented in different knowledge bases in order to reduce the possibility for conflicts between rules. As a matter of fact, this modular approach makes the prototype's maintenance easier. The modification or introduction of a rule is limited to the corresponding pertinent rule base.

4.1 Validation

In general, validation is conducted at two levels: external and internal. An external validation consists in determining "whether this is the right model" according to the requirements, while an internal validation consists of checking "whether the model is right", whether it accomplishes what is expected of it. Internal validation is sometimes called model verification [6].

Throughout the modelling process, we were in regular contact with an expert coordinator, who was able to keep us on track and validate our approach at various stages. Initially the approach resulted in three major iterations at the knowledge model level. Subsequently, during a half-day meeting, the expert validated the definitive reasoning model that we had developed, namely the multi-hypotheses approach, the various hypotheses levels, the types of hypotheses, as well as the inferences diversity and chaining order. To help with the validation process, we simulated representative incidents based on actual cases. We therefore were able to validate that the application of the proposed method enables one to construct scenarios and associated routes. Our objective was not to reproduce exactly the way coordinators solve the possibility area problem since, as mentioned previously, they reason using their own implicit heuristics. Nonetheless, the expert confirmed, with the help of a simulation, that a structured reasoning method, such as the one we propose can guide the reasoning process in a clear and rigorous manner.

We verified the CLIPS rules triggering behaviour (internal prototype validation), which enabled us to validate the sequencing in its reasoning process. The prototype output allowed us to check that the triggered CLIPS rules complied with the knowledge model expectations. This verification allowed us to ensure that the rules base did not contain any anomalies (redundancy, conflicts, circular rules, unused rules). As mentioned in [7], it is easier to verify a small rules base. Since the number of rules implemented in our prototype is fairly small (approximately 150 rules), we conducted the verification by a careful examination of the structure of the rules base and by applying tests to observe the rules behavior [12]. However, we emphasize that our objective in developing the prototype was to demonstrate the feasibility of implementing the knowledge model and to illustrate its applicability to potential users, and not to develop an operational system.

5 Conclusion and future research

The main contributions of our work are the development of a structured approach for defining the possibility area of aeronautical SAR incidents overland. This was the result of concerted efforts to capture, organize, structure, reuse, and make available the expertise and knowledge of experienced search mission coordinators. This is very valuable for the intellectual capital and knowledge assets of the SAR organization. By providing a systematic approach to scenario definition and analysis of what might have happened to the missing aircraft and where, we wished to enhance the ability to estimate the search object's location, resulting in a smaller and more accurate search area. This project also allowed us to develop a unique interactive advisor system prototype that is a combination of an expert system and a geographical decision support system for aeronautical SAR incidents. In addition to enhancing the operational environment, such a tool could be useful in a training context to support new coordinators.

Our knowledge model has some limitations in because the individual hypotheses must have been previously identified during the knowledge acquisition phase. Furthermore, the model we proposed does not allow the coordinator to allocate a degree of certainty to the various scenarios derived. This however could be corrected by adding nuances to the hypotheses and asking the coordinator to provide subjective confidence levels. For example, in a given case, instead of retaining the hypothesis of running out of fuel, the coordinator may state that the "out of fuel" hypothesis is hardly plausible, quite plausible, or very plausible.

Future work includes the continued development of a user-friendly geographical knowledge-based system prototype. In addition to software development, we intend to enhance the knowledge model by an in-depth exploration of the hypotheses already defined at each level, and by enriching the knowledge base with new hypotheses as new incident types occur.

References

[1] Abi-Zeid, I. and Frost, J., SARPlan: A Decision Support System for Canadian Search and Rescue Operations. *European Journal of Operations Research*, Vol. 162/3, pp. 630-653, 2005.

[2] Canadian Coast Guard College, *A brief guide to CANSARP*, TP 11718E, 1993.

[3] Cottam, H. and Shadbolt N., Knowledge Acquisition for Search and Rescue planning. *International Journal of Human-Computer Studies*, Vol. 48, pp. 449-473, 1998.

[4] Giarratano, J., *CLIPS: User Guide version 6.20*, http://clipsrules.sourceforge.net/documentation/v624/ug.pdf, accessed on November 30, 2009, 2002.

[5] NSM, *National Search and Rescue Manual*, http://www.casaraontario.ca/~webmaster1/Documents/NSM_2000.pdf, accessed on November 30, 2009, 2000.

[6] Mosqueira-Rey, E. and Moret-Bonillo V., Validation of intelligent systems: a critical study and a tool, *Expert Systems with Applications*, Vol. 18, pp. 1–16, 2000.

[7] O'Keefe, R. M. and O'Leary, D. E., Expert system verification and validation: a survey and tutorial, *Artificial Intelligence Review*, Vol. 7, pp. 3-42, 1993.

[8] Rodriguez-Rodriguez, A. and Alemàn-Flores M., A Framework for the Search and Rescue Domain, *Proceedings 14th International Conference on Applications of Prolog*, pp. 305-316, 2001.

[9] Saunders, P. H., *A new search method based on an analysis of air distress cases: 1981-1986*, ATGOR Staff Note 2/87, Department of National Defence Canada, 1987.

[10] Schreiber, A. Th., Akkermans, J. M., Anjewierden, A. A., de Hoog, R., Shadbolt, N. R., Van de Velde, W. and Wielinga, B. J., *Knowledge Engineering and Management The CommonKADS Methodology*, MIT Press, 2000.

[11] Schvartz, S., Abi-Zeid, I. and Tourigny N., Knowledge Engineering for Modelling Reasoning in a Diagnosis Task: Application to Search and Rescue, *Canadian Journal of Administrative Sciences,* Vol. 24, pp. 196-211, 2007.

[12] Schvartz, S., *Modèles et prototype d'un système à base de connaissances pour la localisation d'aéronefs portés disparus*, M. Sc. thesis, Département d'informatique et de génie logiciel, Université Laval, Québec, 2005.