Task-Driven Information Presentation

Nathalie Colineau and Cécile Paris
CSIRO – ICT Centre
Locked Bag 17, North Ryde 1670, NSW, Australia
{nathalie.colineau, cecile.paris}@csiro.au

Abstract

In the context of innovative Airborne Early Warning and Control (AEW&C) platform capabilities, we are building an environment that can support the generation of information tailored to the operator's tasks. To this end, we propose a task-based interaction environment that delivers tailored multimedia information based on the analysis and the recognition of the operator's activity. This aims at supporting the operators and maximising their efficiency by automatically providing them with information relevant to their tasks. The work is being conducted under a program that aims to enhance research and development capabilities in Australia that align with the Commonwealth's areas of interest in AEW&C.

1. Introduction

In the context of innovative Airborne Early Warning and Control (AEW&C) platforms, we investigate a task-driven environment that provides air combat operators with coherent and tailored displays. The environment proposed is based on an analysis and a representation of the operators' tasks and the context in which they are performed to drive the interaction process and provide operators with relevant information.

Indeed, it is crucial to provide AEW&C operators with the information they need to support their decision process. Importantly, as the operators' information needs change at various stages of their task, the information displayed should also change to ensure the most relevant information is always available and prominently displayed, and less relevant information is not allowed to distract the operator. It is thus important to build an environment that can support the generation of displays tailored to the operator's tasks. More specifically, the technology we are building seeks from heterogeneous knowledge sources the appropriate and relevant information to deliver. It then integrates the various elements of information into a coherent message relevant to task and user.

Similarly to *attentive information systems* (e.g., Maglio *et al.*, 2000), our task-driven environment monitors the user's activity and presents relevant information appropriately. There are, however, significant differences between such systems and our approach. Attentive Information systems build a user model by observing users' behaviours and interactions, and are thus able to suggest to users potentially useful information. This includes, for example, related information or information close to their current interest. In contrast, our approach is more strongly task-based, with an explicit aim of helping the operators in performing their tasks. Moreover, our aim is not to tailor information to individuals (i.e., particular operators), but to groups of individuals by identifying their role in the defence organisation (e.g., surveillance operator. fighter controller, etc.) and the tasks they need to perform as part of their role. Our approach thus consists in analysing and modelling the operators' tasks to be able to provide them with more coherent and less cluttered display. Indeed, with highly variable information environments, it is important to not deliver *all* the information that is available, as this can result both in information overload and in inappropriate or irrelevant information. It is therefore important to tailor the operators' display, highlighting information appropriate at a specific point in the task to enable operators to concentrate more easily on what is important, and thus, ultimately, make better decisions.

Our approach is also comparable to the Watson information access system (e.g., Budzik *et al.*, 2001), which provides information in the context of the activities the user is performing. This system observes users interacting with applications and anticipates their information needs using a task model. The Watson system may be described as a smart information retrieval engine that grounds the queries in the context of the user's activity and presents the result as a list. In contrast, our task-driven environment selects, organises and designs the presentation of multimedia information to be integrated on the display. The aim is not only to provide relevant information, but also to present the information coherently through appropriate media and integrate it smoothly with the elements of information already on screen.

This article focuses on how an analysis of the operators' role, task and information needs can be exploited and embodied in a discourse model to help deliver tailored multimedia information.

We first present the task-driven environment we are building to support AEW&C operators. It constitutes the interface between the operator and the application. We then describe more specifically the Virtual Document Planner (VDP) component, which is in charge of automatically generating relevant and coherent multimedia information to the operators. This component is in particular responsible for selecting and organising information coherently deciding how best to integrate the multimedia material into the operators' display. Finally, we explain how the operator's task analysis we performed will be used to organise coherently the multimedia information to be presented and drive information provision and visualisation.

2. Task-Driven Environment

The environment we are developing aims at supporting the air combat operators in their tasks and at enhancing the effectiveness of their interaction with the system. AEW&C operators sometimes need to manually collect and combine information from various sources, including sensors (e.g., Radar, IFF - Identification Friend or Foe, ELINT - Electronic Intelligence, etc) and databases (e.g., AFTN - Aeronautical Fixed Telecommunications Network), to understand the situation and make decisions. We propose to support the operators by automatically providing them with the information relevant to their tasks, organising it and delivering it appropriately.

To this end, we designed a Task-Driven Interaction Module (TDIM), whose role is to provide an easy access to relevant information, deciding how important information should be presented and coordinated within and across media, and ensuring that the display is coherent. The coherence of the display is crucial, especially when combining different modalities (in our case text and graphics). Indeed, one has to ensure that the provision of information is not a simple juxtaposition of information and media, but really integrates each element of information into a coherent whole, explicating the relationships amongst the different elements and media.

Thus, to avoid cluttering the operator's display, it is important to reinforce the cohesion between each element of information.

2.1. TDIM Overview

The TDIM constitutes the interface between the operator and the system. The TDIM developed jointly with the CSIRO Business Intelligence (BI) group is composed of three components as shown in Figure 1:

- The Graphical User Interface (Operator-GUI), which allows the operator to interact with the system (i.e., consulting and accessing data);
- The Task-Recognition Parser (TRP, developed by Simon Williams of the BI group), which analyses the operator's activity to be able to both determine the current task in progress and predict the next one; and,
- The Virtual Document Planner (VDP), which delivers multimedia information tailored to the operators' tasks.

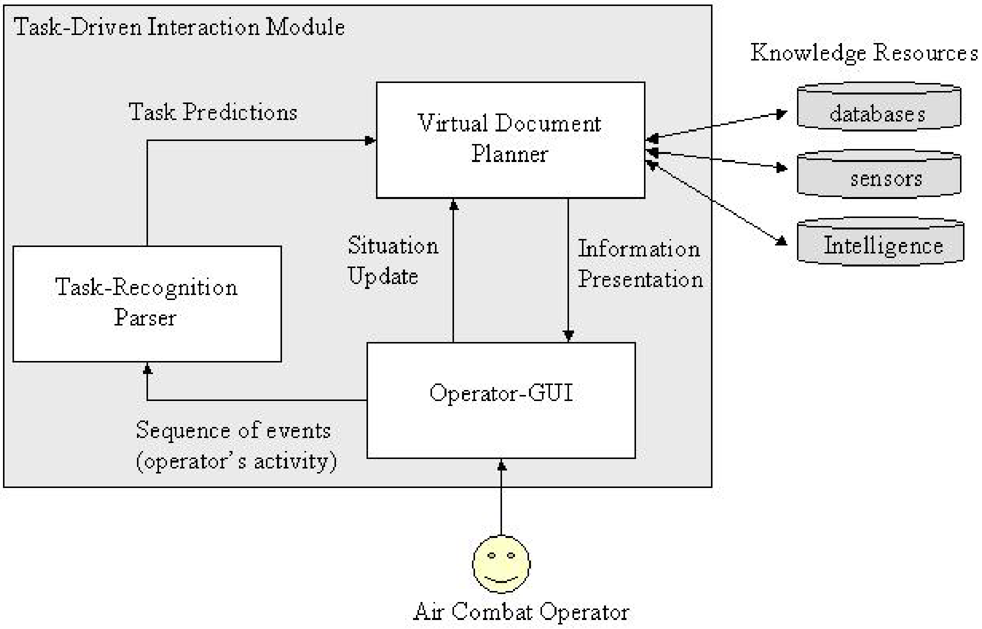

Figure 1. Architecture overview for the Task-Driven Interaction Module (TDIM)

2.2. Virtual Document Planner

The Virtual Document Planner (VDP) was originally designed to produce documents dynamically, integrating several data sources and customising the content for a user (Colineau and Wan, 2001). We are now extending this platform to handle both textual and graphical information. The architecture we use is a typical Natural Language Generation (NLG) architecture, where the linguistics resources are separate from the engines (cf. Reiter and Dale, 2000).

The engines use discourse rules, which select and organise the content, design rules, which determine an appropriate way to realise the content and the structure of the presentation, and the linguistic and graphical resources (such as the vocabulary and syntax of the application domain, icons, ready-made graphics, etc.).

The specific engines that the VDP employs are *planning engines*, and the resources are represented as plans, as in (Moore and Paris, 1993). The VDP, shown in Figure 2, employs a discourse generation approach. In this approach, a *communicative goal* (or *discourse goal*) is given as input to the document planner. The communicative goal is to answer an information need of the user, the operator in this case. The planning engines then try to achieve this goal, using their set of discourse, and design rules, using a traditional goal/sub-goal decomposition. In addition, this approach exploits *rhetorical relations* (or *coherence relations*), based on Rhetorical Structure Theory (RST) (Mann and Thompson, 1988), to guarantee the coherence of the resulting document. Coherence relations must hold between sibling subgoals in the goal decomposition. They indicate how the various discourse segments and piece of information fit together to achieve the communicative goal. Generation proceeds in four main steps:

- The content planner selects and organises the content, with explicit rhetorical relationships between discourse elements using a library of discourse plans. The content planner retrieves information to achieve the top-level discourse goal, that of answering the user's information need. Then, the chosen discourse elements are assigned to specific media;
- The modules dedicated to media/modality (e.g., the sentence planner or the graphics planner) enrich the discourse structure with media specific information, using design rules;
- The specific media realisers (e.g., the sentence realiser and the graphics realiser) produce the corresponding media objects. These may have been retrieved from databases or generated from first principles;
- Finally, the layout realisation module merges media objects and defines the document or the display layout.

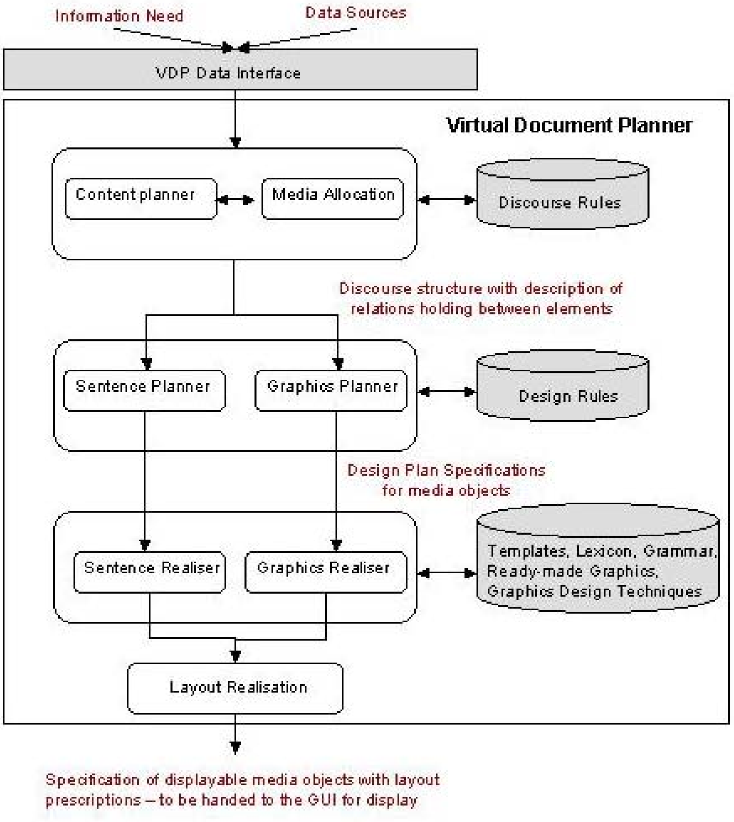

Figure 2. Multimedia Presentation System Architecture

2.3. Scenario of Interaction

When the operator starts to interact with the system, the Task Recognition Parser (TRP) analyses the operator's activity (i.e., the sequence of interactions performed by the operator with the system).

It does so by using a representation of the task and its decomposition, and temporal and logical relationships amongst the sub-tasks. The TRP thus determines the current (sub)task and is able to predict the next one. This is important so that the system can respond appropriately, which, in this case, means providing the information required at this stage.

Moreover, having the TRP analysing continuously the operator's interactions allows also the system to detect when the operator switches from one task to another. Indeed, the situation may require the operator to interrupt or abort the current task, and thus move to a higher priority task.

The task predicted by the TRP is thus sent to the VDP, which converts it into an information need, based on the operator's task analysis we performed. The VDP then collects the data considered appropriate and relevant to answer the information need, selecting and combining information from various sources. As explained above, it organises the content retrieved and allocates it across the media (i.e., textual box, graphical display, etc.), taking into account the data and task characteristics.

The information thus retrieved, organised and allocated to media is then ready to be presented. So, finally, it is sent to the graphical user interface, which, in turn, displays the presentation. To guarantee the continuity and the consistency of the display, the VDP has to provide smooth transitions, by gradually transforming the current display into a new one.

3. Task-Based Multimedia Information Delivery

The research question underlying multimedia information delivery is to determine what information should be displayed, in which context and for which purpose. The core idea behind our work is to drive information provision based on a representation and recognition of the operator's tasks. These tasks are, in turn, used to determine the operator's information needs.

To ensure that the system would not be based on ad-hoc resources, we collected information about the operator's tasks and information needs. The aim of this analysis was to inform the TDIM and more specifically the Task Recognition Parser and the Virtual Document Planner. More specifically, the task analysis was used to determine the task sequences to be formalised in the TRP grammar, and to build the VDP knowledge sources (mainly the discourse rules that structure and organise the content to be presented).

3.1. Operator's Task Analysis

Understanding the tasks and the environment of air combat operators is a crucial step to be able to provide them with relevant and appropriate information. We thus analysed the role of the air combat operators, and, more specifically, the roles of surveillance operator and fighter controller. Indeed, these two roles are likely to represent the new role of the AEW&C operator. The objective of the task analysis was to collect information that will inform us about:

- The distinction between the surveillance operator and the fighter controller roles;
- The operators' respective tasks; and
- The operators' information needs.

3.1.1. Source of Information

The primary source of information for our task analysis was observations and interviews with the Royal Australian Air Force (RAAF) surveillance operators and fighter controllers. Although Wedgetail operators are not currently available for study, as the aircraft has not been yet delivered, we had the opportunity to visit the Eastern Region Operation Centre (EASTROC) at RAAF Williamtown. A study of RAAF air defence tasks was conducted during visits to the 3rd Reporting and Controlling Unit (3CRU) and the Surveillance and Control Training Unit (SACTU) at Williamtown. We conducted interviews and observed surveillance operators, fighter controllers and supervisors, in group and individual situations. We took notes and sketches about their operational role, and its challenges were discussed. The operators introduced the two interfaces they currently use and described features of AEW&C interfaces. One surveillance operator-pair and six fighter controller operators were observed using the ground-based interfaces during actual missions.

During task analysis, we captured the operator's activity using a semi-formal representation. We analysed their respective tasks and the way they perform their job. In particular, we tried to make explicit the underlying procedures, the information used (or accessed, requested), and the way they interact in general to do the tasks (interaction with the system, with pilots, with other operators, etc.). We also consulted additional information, including presentations describing the future AEW&C capabilities, and documentation describing the technology and procedures used by US Navy air defence operators. These served to complete information gathered during the interviews. Then, from this analysis, we formalised the operators' tasks, building a task model.

3.1.2. Task Characteristics

RAAF air defence activities involve a lot of people who have a specific role that contributes to air defence. Key people in the process include the surveillance operators, the fighter controllers, commanders, and various analysts, technicians and pilots. To understand how surveillance operators and fighter controllers contribute to air defence, we analysed the specificity of their roles and their task characteristics.

The air defence organisation relies on information about the battle space to complete their job. The primary source of information is real-time information from sensors such as radars that indicate the location of airborne objects in the area. This radar information can be difficult to interpret in its raw form. A "tracker" system can, over time, associate radar returns as coming from an individual aircraft. This aircraft can then be established as a *track*. The raw radar picture and track symbols are the basic display seen by air defence operators. In addition to the sensor's return, other information may come from a variety of data sources. It is the role of the surveillance operator to collect and integrate this additional information to facilitate the interpretation of the battle space. The combined output of the surveillance operator's work is the *Recognised Air Picture*. It is considered as the best understanding of the current situation in the battle space and is used by fighter controllers and command and control (C^2) personnel to complete their own tasks. Each of these "end-users" has different requirements of the recognised air picture. The requirements are a result of the different tasks they perform. The fighter controllers, for example, use the recognised air picture to exercise command and control over an airborne battle and, in particular, to provide a focused picture of the situation tailored to the pilot's needs.

The environment we propose is likely to best support tasks with large volume of information to be managed, potentially with less time criticality. Table 1 summarises the most important task characteristics for each role (i.e., surveillance operator and fighter controller).

Table 1. Role demands on operators

Features \ Roles	*Surveillance Op.*	*Fighter controller*
Multi-tasking	Some	Very High
Time-criticality	Some	Very High
Info Management	Very High	Very High
Info Volume	Very High	Some

From this analysis, we decided to focus on the surveillance operator role and one of the tasks of the fighter controller role, the one which requires them to handle various information to get an understanding of the situation.

The surveillance operators and fighter controllers have to accomplish a wide range of sub-tasks, including:

- Interactive tasks, i.e., tasks performed with the system;
- Dialogues tasks, i.e., tasks involving exchanges with other operators (e.g., exchanges with the pilots); and,
- Cognitive tasks (e.g., tasks involving interpretation).

We decided to focus on the interactive tasks as these tasks are performed through the system and therefore are likely to be best supported by our environment. In particular, we analysed and formalised two sub-tasks:

- The elaboration and maintenance of the air picture. This task constitutes 80% of the surveillance operator's role and consists in providing track classification (e.g., classify the track as hostile, friendly, neutral or unknown) and identification (e.g., type of aircraft, weapons carried, etc.); and
- The understanding of the situation by the fighter controller. This task consists in maintaining the fighter controllers' situation awareness (i.e., help the operators understand the situation to allow them to interpret and predict events).

We chose these two tasks for several reasons: first, these tasks involve an interaction with the system (in contrast with cognitive tasks or dialogue tasks); second, these tasks are knowledge intensive, namely they require a large volume of information to be accessed and managed, and this is where an approach such as ours can make a difference and contribute the most.

The operator's information needs have been identified during the interviews and observations. They comprise all the information used or accessed to perform the tasks. They may include sensor information (e.g., track information displayed on screen), data sources (e.g., flight plan information, AFTN, NOTAMS – Notices to Airmens), mission briefing or background knowledge.

3.1.3. Task Model

The task model we developed is a user-oriented model that provides a high-level description of the tasks' decomposition, rather than a fine-grained functionality specification. Our aim was to identify the main tasks of the operators and how these tasks are performed.

The task model shown in Figure 3 represents the air defence task corresponding to the surveillance operator and fighter controller roles. It is composed of two main sub-tasks that may be done in parallel: *maintain air picture* (i.e., provide a clear picture of the current situation in the battle space) and *control fighters*. In the former sub-task, the operator has to categorise tracks, maintain tracks and sometimes watch for a specific track. In the latter sub-task, the operator has to maintain his/her situation awareness (decomposed further into a sequence of tasks), maintain the pilots' situation awareness and direct friendly fighters. The tasks we focused on and that have been formalised in the task model are noted in black. The tasks in grey are simply mentioned here but have not been formalised further for the reasons mentioned above. To formalise the operator's tasks, we used TAMOT, a graphical Task Model Editor that we developed during previous work on task modelling (Lu *et al.,* 2002). The formalism used is the Diane+ notation (Tarby and Barthet, 1996).

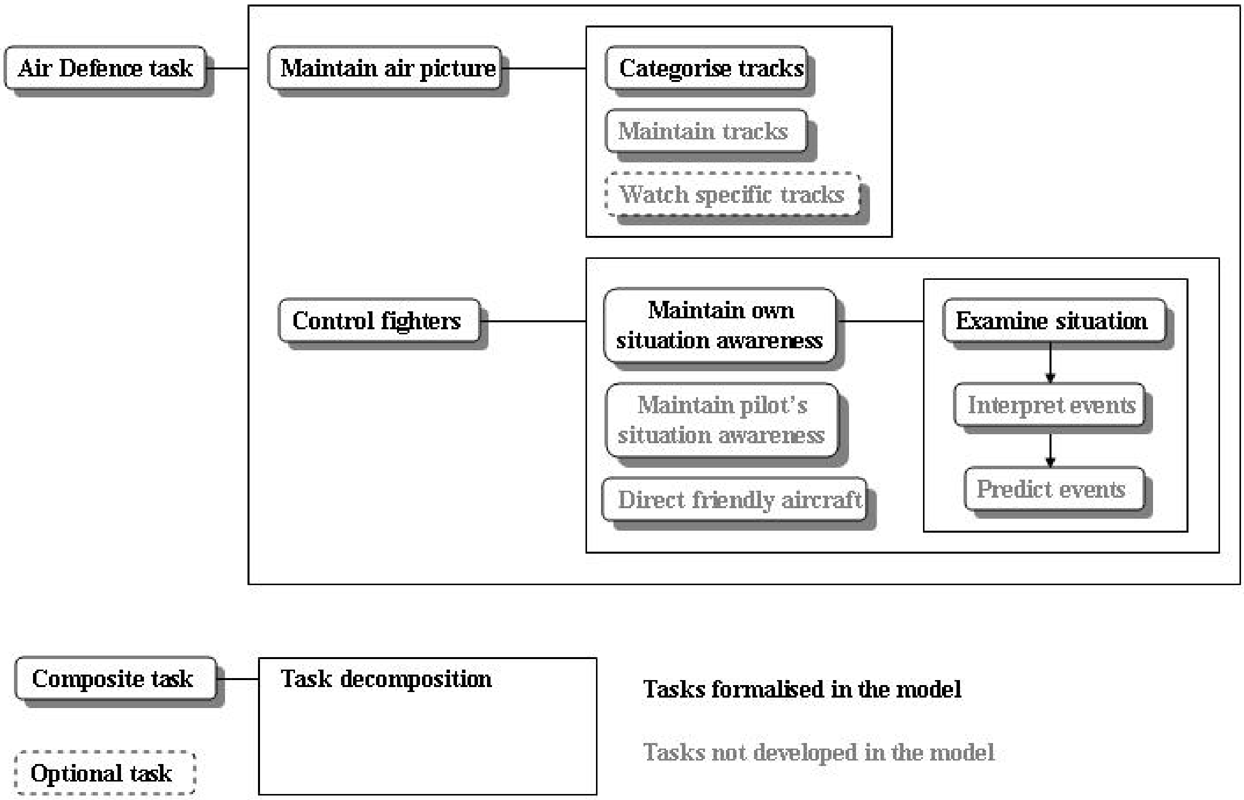

Figure 3. High-level decomposition of the "Air defence" task model

3.2. Construction of the Linguistic Resources

Once the operators' activity has been analysed, formalised into a task model and the information needs identified, we need to represent this information under the different resources that the VDP will access during the content and the design planning process. Two main resources need to be built: the discourse plan operators to organise the content and the design rules to display this content.

3.2.1. The discourse structure

The discourse rules are used to plan the presentation content to satisfy a specific communicative goal, which represents the purpose of the presentation. To design these rules, we first need to represent the discourse structure to be produced, and then, from this representation, to derive the discourse rules that will be used by the content planner. To represent the discourse structure, we started from the task model representation and the information needs identified. The task model represents the tasks that need to be done, their decompositions, and the temporal relationships amongst them. In contrast, the discourse structure represents the information needed to support the task. In particular, it represents the information that needs to be conveyed (i.e., the information content), the purpose for which it is conveyed (i.e., the communicative goal of each information displayed), and how each piece of information is related to and coherent with another piece of information. In our case, having identified the task to be supported and the information the operators accessed at each point, we needed to analyse how to organise each piece of information, emphasising its role with respect to the task and to the whole discourse structure (i.e., the other information presented).

Thus, the discourse structure is built by representing the information need for each task, and defining a set of corresponding communicative goals. Instead of temporal and logical relationships amongst sub-tasks, we now have rhetorical relationships amongst discourse segments. Figure 3 and Figure 4 illustrate respectively the task model we started from and the discourse structure we derived. For the clarity of the presentation, only the high levels of each representation are shown. The discourse structure (Figure 4) corresponds to the information required to support the air defence task. It is a hierarchical structure, with high-level goals being decomposed into lower-level goals, until primitive actions are reached. These primitive actions may be speech acts (for text) or graphical acts (for graphical representations). In addition to the representation of the discourse structure through its communicative goal/sub-goal decomposition, we use rhetorical relations to indicate how siblings in the hierarchy are related to each other and thus, together, form a coherent whole.

Borrowed from the Rhetorical Structure Theory (RST) (Mann and Thompson, 1988), the rhetorical relations link discourse sub-goals together. Originally used to describe text coherence, the Rhetorical Structure Theory has been widely exploited since its inception to represent discourse structure in text and multimedia information generation (e.g., André and Rist, 1990, 1995; Hovy, 1993; Reichenberger *et al*., 1995; Bateman *et al*., 2001). This theory has also been used to drive the generation process in both text and multimedia presentation systems, in particular to indicate how the content to be presented is organised in a coherent way. Finally, the theory has been employed as a useful means to represent the discourse constraints between each discourse elements (or communicative goal/sub-goal). The discourse structure in Figure 4 can be read as follows:

- At the top level, we have identified two main goals: provide information to support the operator to monitor the situation and provide information to support the operator to build and maintain the air picture. The first goal is the main goal, and, in RST terms, it is considered to be the *nucleus.* This is noted in the representation as a strait link from its parent goal. The second goal is considered as a secondary goal, or, in RST terms, as a *satellite*. The satellite is linked to the nucleus by a rhetorical relation. In this case, the rhetorical relation is one of *enablement.* This indicates that the goal represented by the satellite is intended to aid in performing the goal represented by the nucleus. Thus, helping the operator to maintain an accurate air picture will enable him/her to best monitor the situation;
- If we go down the hierarchy, the satellite (i.e., *support the operator maintaining the air picture*) is decomposed further into a list of nuclei, namely a list of goals equally important that are linked by an unconstrained rhetorical relationship *joint.*
  This indicates that the three communicative sub-goals will equally contribute to the satisfaction of the parent goal. Only one of the three nuclei has been decomposed further currently: the communicative goal that will support the track categorisation;
- If we return to the main nucleus (i.e., *support the operator monitoring the situation*), it is, in turn, decomposed into two main goals: *support the operator to control the fighter* (the nucleus); and, *maintain the operator's situation awareness* (the satellite). The satellite is again linked to the nucleus by the rhetorical relation *enablement*;
- Finally, the satellite "*maintain the operator's situation awareness"* is decomposed into two goals: *support the operator's understanding of the situation* (the nucleus); and, *provide the operator with situation assessment* (the satellite).

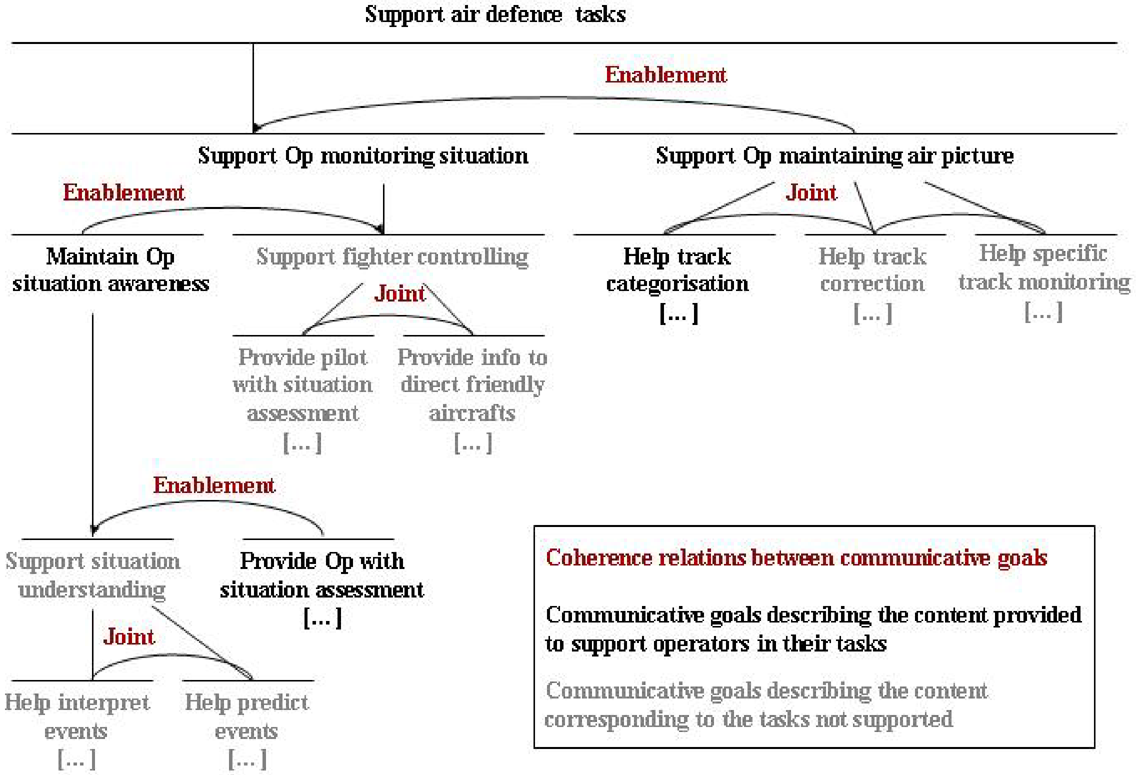

Figure 4. The "Air defence" discourse structure

3.2.2. The presentation structure

The other linguistic resources used by the VDP are the design rules. They specify how the content selected during the content planning stage will be presented on the display (e.g., as textual or graphical information). These rules are built by taking into several criteria to decide on the most appropriate way to communicate information:

- The data characteristics (e.g., Levie et Lentz, 1982; Larkin et Simon, 1987; Hegarty *et al.*, 1990);
- The purpose of the presentation (e.g., Mackinlay, 1986; Casner, 1991); and
- The combination of modalities (e.g., Chandler and Sweller, 1991; Mayer, 2001).

The way the design of the presentation structure takes place in the multimedia generation process varies from one approach to another. For some researchers (e.g., Feiner and McKeown, 1993; Arens *et al.*, 1993), there is a strict separation between the data selection and its organisation on the one hand, and the media/modality allocation on the other hand. Once the discourse tree has been entirely specified, the combination of modalities and media is chosen. For others (e.g., Maybury, 1993; André and Rist, 1995), these two steps are strongly linked and should be performed in parallel to integrate both the constraints that the content may impose on the modality or medium choice, and, similarly the constraints that the modality or the medium may impose on the content selection.

However, whatever the approach taken, the VDP has to deal with several modalities and decide on the best way to realise the content when presenting the information to the operators. Thus, the next step for us consists in enriching the current discourse structure with information for the presentation of each discourse sub-goal. Each presentation specification is then sent to the media-specific generators to realise the content i.e., the media objects to be displayed on screen.

4. Discussion

Building such an environment may raise some issues we discuss here. For example, one may consider that re-enforcing the operator's focal attention on the current task by displaying prominently the information considered by the system as relevant may constitute an hindrance for the operator in noticing new information or new situations. Indeed, this raises the issue of finding a good balance between having everything available all time, which may result in information overload and cluttered displays, or directing the operator's attention to specific information by highlighting it. To maintain operator's situation understanding when focusing on a particular task, several ideas are envisaged:

- Ensure that crucial information is always visible or easily available. Based on the task analysis and the identification of the information need, the system should focus on information not only relevant to the current task, but also related to the operator's role in general to enable them to perform their work;
- Take into account contextual information, such as new situation, new events, crucial parameter changes. In this case, the TDIM environment would not only be driven by the task, but also by the context. For example, the system may be sensitive to new tracks entering the battle space, or to tracks suddenly heading to friendly forces or friendly ground bases. In that case, the system would generate a new task, which would consist in providing information related to this event. However, on such issues, we have to remain careful as the system aim is mainly to support operators in making decisions, but not necessarily to interpret the situation for the operators.

Another issue we would like to raise is that of evaluation. Evaluations may be conduct at several levels: a) at the component level, and b) at the global level (i.e., the whole environment). At the component level, we may envisage to assess the accuracy and reactivity of the Task-Recognition Parser in analysing and predicting the operator's tasks and the appropriateness of the information delivered by the Virtual Document Planner with respect to the relevance to the task and the adequacy of the presentation. At the global level, we may envisage to assess our task-based environment by comparing it with another platform without the TDIM. In that case, the sensor data will be still available through the radar display, but the other data sources will be accessed in their raw form i.e., as currently done, namely as separate data but not automatically extracted and integrated into the display. Such evaluations would provide useful inputs for the design of the new Airborne Early Warning and Control (AEW&C) platforms. However, we will need to wait for a pilot system to be implemented and for the AEW&C operator's role to be completely defined to evaluate the real impact of such a technology to maximise the operator's efficiency.

5. Conclusion

Future AEW&C operators will deal with a large amount of information. At any point in time, however, they do not necessarily need access to all available information. We propose to aid the operators to focus on important information and thus make better decisions by displaying only crucial data and highlighting its importance. To do so, we have set an environment driven by the operator's task, which aims at collecting, organising and appropriately presenting information that the operators need to perform their task.

6. Acknowledgement

We wish to thank Robert Tot for his work on the task analysis and modelling and the surveillance operators and fighter controllers from the RAAF who participate to the task analysis interviews. We wish also to thank Simon Williams for his work on the Task-Recognition Parser. We acknowledge the support of Boeing (Research Agreement WT-SIDA 010.2) and CSIRO.

7. References

André, E. and Rist, T. (1990). Towards a plan-based synthesis of illustrated documents. In *Proceedings of 9th ECAI*, Stockholm, Sweden, 25-30. Also DFKI Research Report RR-90-11.

André, E. and Rist, T. (1995). Generating coherent presentations employing textual and Visual material. In *Artificial Intelligent Review*, *Special volume on the Integration of Natural Language and Vision Processing*, 9 (2/3). 147-165.

Arens, Y., Hovy, E. and van Mulken, S. (1993). Structure and rules in automated multimedia presentation planning. In *Proceedings of the 13th International joint Conference on Artificial Intelligence (IJCAI'93)*, Chambery, France. 1253-1259.

Bateman, J., Kamps, T., Kleinz, J. and Reichenberger, K. (2001). Towards constructive text, diagram, and layout generation for information presentation. In *Computational Linguistics, Vol. 27, No. 3, September 2001*. 409-449.

Budzik, J., Hammond, K.J. and Birnbaum, L. (2001). Information access in context. In *Knowledge-Based Systems 14*, 37-53, Elsevier.

Casner, S. (1991). A task-analytic approach to the automated design of graphic presentations. In *ACM Transactions on Graphics, Vol. 10, No. 2, April 1991*. 111-151.

Chandler, P., Sweller, J. (1991). Cognitive load theory and the format of instruction. *Cognition and Instruction,* 8, 293-332.

Colineau, N. and Wan, S. (2001). Mobile delivery of customized information using Natural Language Generation. In *Monitor (Special Issue on Wireless Communication Special)*, 26(3), Sept.-Nov. 2001, 27-31.

Feiner, S. K. and McKeown, K. R. (1993). Automating the Generation of Coordinated Multimedia Explanations. In M. Marbury (Ed.), *Intelligent Multimedia Interfaces,* chapter 5, 117-138. AAAI Press / The MIT Press, Menlo Park, CA.

Hegarty, M., Carpenter, P.A. and Just, M.A. (1990). Diagrams in the comprehension of scientific texts. Barr, R., Kamil, M.L., Mosenthal, P. and Pearson, P.D. (éds). *Handbook of reading research*, New York: Longman, 641-668.

Hovy, E. (1993). Automated discourse generation using discourse structure relations. *Artificial Intelligence* 63, *Special Issue on Natural Language Processing*.

Larkin, J.H. and Simon, H.A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65-99.

Levie, W.H. and Lentz, R. (1982). Effects of text illustration: a review of research. *Educational Communication and Technology Journal*, 30(4), 195-232.

Mackinlay, J.D. (1986). Automating the design of graphical presentations of relational information. In *ACM Transactions on Graphics Volume 5, No.2*, April 1986, 110-141.

Maglio, P., Barett, R., Campbell, C. and Selker, T. (2000). SUITOR: An attentive information system. In *the Proceedings of the International Conference on Intelligent User Interfaces (IUI'00)*, New Orleans, Louisiana, USA.

Mann, W. and Thompson S. (1988), Rhetorical Structure Theory: Towards a Functional Theory of Text Organization. *Text*, 8(3), 243-281, 1988.

Maybury, M. (1993). Planning multimedia explanations using communicative acts. In M. Marbury (Ed.), *Intelligent Multimedia Interfaces,* chapter 2, 59-74. AAAI Press / The MIT Press, Menlo Park, CA.

Mayer, R.E. (2001). *Multimedia learning*. Cambridge: Cambridge University Press.

Moore, J. and Paris, C (1993), Planning text for advisory dialogues: Capturing intentional and rhetorical information. In *Computational Linguistics,* 19(4), 651-694.

Lu, S., Paris, C. and Vander Linden, K. (2002), Tamot: Towards a Flexible Task Modelling Tool, In *Proceedings of Human Factors 2002*, Melbourne, Australia, 25-27.

Reichenberger, K., Rondhuis, K., Kleinz, J. and Bateman, J. (1995). Effective presentation of information through page layout: a linguistically-based approach. In *Electronic Proceedings of the ACM Workshop on Effective Abstractions in Multimedia, San Francisco, California*.

Reiter, E. and Dale, R. (2000), *Building Natural Language Generations Systems*. Cambridge University Press 2000.

Tarby, J-C. and Barthet, M-F. (1996), The Diane+ method. In *Proceedings of CADUI'96*. Namur, Belgium, 5-7 June, 95-119.