PERGAMON

Engineering Applications of Artificial Intelligence 12 (1999) 691–703

Engineering Applications of ARTIFICIAL INTELLIGENCE

www.elsevier.com/locate/engappai

Integrating case-based reasoning and hypermedia documentation: an application for the diagnosis of a welding robot at Odense steel shipyard

Eric Auriol[a,*], Richard M. Crowder[b], Rob McKendrick[c], Roger Rowe[d], Thomas Knudsen[e]

[a]*AcknoSoft, 15 rue Soufflot, 75 0005, Paris, France*
[b]*Department of Electrical Engineering, University of Southampton, Southampton SO17 1BJ, UK*
[c]*Parallel Applications Centre, Chilworth, Southampton SO16 7NS, UK*
[d]*Multicosm Ltd, Chilworth, Southampton SO16 7NS, UK*
[e]*Odense Steel Shipyard Ltd, PO Box 176, DK 5100 Odense C, Denmark*

Received 1 December 1998; accepted 1 July 1999

Abstract

Reliable and effective maintenance support is a vital consideration for the management within today's manufacturing environment. This paper discusses the development of a maintenance system for the world's largest robot welding facility. The development system combines a case-based reasoning approach for diagnosis with context information, as electronic on-line manuals, linked using open hypermedia technology. The work discussed in this paper delivers not only a maintenance system for the robot stations under consideration, but also a design framework for developing maintenance systems for other similar applications. © 1999 Elsevier Science Ltd. All rights reserved.

1. Introduction

Odense Steel Shipyards (OSS) build container ships of a quality and size that cannot be matched by their competitors. OSS achieves this advantage by using a high level of automation in their design and production process. This is typified by the company's widespread introduction of robotic welding systems, in particular the B4 robot station.

The use of high-technology production facilities requires the introduction of effective maintenance systems. The main purpose of these systems is to minimise production costs by ensuring the reliability and production quality of the manufacturing systems. In addition, the maintenance cost should be minimised by effective use of resources.

Any form of maintenance requires the application of knowledge held by people familiar with the system. At OSS these are the operators and the system specialists. The operators are responsible for operating and undertaking maintenance of the robot. At OSS the operators are concerned with the entire system, and will have engineering expertise to solve many problems. Robot specialists will have in depth-knowledge of the robotic systems, and therefore will be called upon when an operator is unable to solve a problem rapidly. It is the case (both at OSS and generally) that there are more operators than robot specialists available for the maintenance of the system, as specialists are shared amongst a number of robot installations. One approach to minimising maintenance costs is to develop a system that enables the operator to undertake more of the maintenance tasks. This will minimise

* Corresponding author. Tel.: +33-1-56-24-22-00; fax: +33-1-56-24-22-66.

E-mail addresses: auriol@acknosoft.fr (E. Auriol), rmc1@soton.ac.uk (R.M. Crowder), rmk@it-innovation.soton.ac.uk (R. McKendrick), r.rowe@multicosm.com (R. Rowe), tdk@oss.dk (T. Knudsen).

0952-1976/99/$ - see front matter
PII: S0952-1976(99)00040-8

the time currently wasted in waiting for robot specialists.

To undertake a maintenance task, reference has to be made to the information stored in documentation. In the case of the robot being considered in the paper, this information consists of a considerable number of loose-leaf documents, drawings ranging from A4 to A0, and proprietary information. This, by its very nature, can be easily damaged, or lost when taken on to the factory floor. The solution is the computer-based integration of a diagnostic system with a document storage and retrieval system.

This paper discusses aspects of the HELPMATE (Hypermedia with Enhanced Linking Providing Maintenance and Analysis Tools for Engineering) project, which is developing a computer-based system capable of supporting the operator to diagnose and repair faults within the B4 welding station. The developed application, *Freya*, required the integration of a case-based reasoning (CBR) system (*KATE* developed by AcknoSoft) with an open hypermedia system (*MicroCosm Pro* developed by Multicosm).

1.1. The target system: the B4 robot station

OSS's B4 robot station is the world's largest robot station for arc welding. The robot station is designed to weld ship hull sections up to $32 \times 22 \times 6$ m, and 400 T in weight. The robot station consists of 12 individual robot gantries, each with eight degrees of freedom, suspended 17 m above the shop floor, and is capable of welding up to 3 km of steel per day. As OSS only has one such installation, any failure will impact on the material flow through the yard.

The B4 robot station was designed by OSS, including a cell-control system called Rob-Ex (Robot Execution system). Rob-Ex is capable of handling the planning and scheduling of the robot's welding jobs, followed by downloading of the post-processed robot program to the local robot controller for execution. Rob-Ex also includes sub-systems that schedule and plan preventative maintenance activities.

While the Rob-Ex system has a fault-detection capability, where the detected fault is communicated to the operator as a code, Rob-Ex does not provide any diagnostic information. The resolution of the fault is down to the operators and specialists. While some faults (e.g., replacement of welding wire) can be resolved with ease, more difficult faults (e.g., a drive or servo failure) will require information from system documentation, or in extreme cases discussion with the B4's system specialist.

1.2. Objectives for OSS

The B4 robot station is a key installation on the yard, and because there are no replacement or backup installations, the work performed on the B4 robot station directly influences the output of the yard. For the B4 robot station, and other installations with the same crucial role, a hypermedia diagnostic maintenance tool, which the HELPMATE technology provides, is essential for the following reasons.

- The reliability of the installation is improved.
- The quality of the work performed by the installation is improved.
- As reliability improves, planning on the entire production line is easier.
- Operators are involved, and are able to perform better, in the daily maintenance, and the sense of responsibility is thereby increased; this is called operator-minded maintenance.
- Time wasted writing for maintenance experts is minimised.
- Experts can be released from existing installations, and be used for development of new technology.
- The maintenance costs are minimised.

2. Inductive learning and case-based reasoning

2.1. Principles

Inductive learning and case-based reasoning (CBR) are different but complementary methods for utilizing past information in solving problems by reasoning from cases. A case is defined as the description of a problem that has been successfully solved in the past, along with its solution. Inductive learning (Quinlan, 1983) creates a general description of past examples, and then applies this description to new data. Case-based reasoning stores past examples, and assigns decisions to new data by relating it to past cases. Inductive learning and CBR are complementary techniques and their integration improves their capabilities (Auriol et al., 1995a; 1995b).

Inductive learning and CBR require that the data be structured, for example by using classes of objects with slots. A standard relational database schema can easily be mapped onto this object model. It allows one to define the vocabulary used to describe the cases. In the current version of Freya, the object model contains 93 different slots, associated with 324 different values. For example, the slot "Display of SR-DSV 5379 on X axis" is associated with five values "OV", "OC", "ST", "OL" and "SLM", which stand for the name of the LEDs displayed by the servo-driver of the X axis.

2.1.1. Inductive learning

In all its various meanings, inductive learning or "induction" has to do with reaching conclusions about a whole class of facts, based on evidence on part of that class. Recorded cases of equipment failure and repair, legal cases and their adjudication, or the credit records of debtors all are a small part of the potential events in each of these areas. The original idea of induction in this sense is to generalise the lessons of past cases into rule, in order to remain consistent with the way human experts think, while providing the input to expert systems that knowledge engineers have failed to do (Lenz et al., 1998; Manago and Auriol, 1998).

The inductive method used in the Freya system is based on a fault tree extracted automatically from the case history. Following the paths of the fault tree generates the rules used for problem solving by the diagnostic system. When the user reaches a leaf of the fault tree (i.e., the conclusion of a rule), the diagnostic system provides both the cases that have been used to create the rule, and the set of solutions associated with these cases.

2.1.2. Case-based reasoning

Instead of building a generalisation of a database as the inductive learning approach does, CBR directly looks for similar cases in the database, with respect to the requirements of the user. When a new problem is defined by the user, CBR recalls similar cases to this new problem, and adapts the solutions that have worked in the past for the current problem. CBR offers flexible indexing and retrieval based on similarity. For applications were safety is important, the conclusions can be further confirmed, or refuted, by entering additional parameters that may modify the similarity values. CBR appeals to those professionals who solve problems by recalling what they did in similar situations. It works well even in domains that are poorly understood, or where rules have many exceptions.

The CBR method used in Freya is fired when the diagnostic system returns several cases with different solutions. Whilst the fault tree orders these solutions according to their relative frequency in the case base, the CBR similarity matching function may suggest another order, based on the relative likelihood of each solution. This type of cooperative integration is common in diagnostic systems (Auriol et al., 1995a; 1995b).

2.2. Application in HELPMATE

The KATE suite of CBR tools was used here to capture the experience in the form of a case base. The INRECA2 methodology (Bergmann et al., 1999) was followed, so that a classical CBR development could be applied:

1. defining an initial data model;
2. creating a questionnaire and acquiring the cases;
3. defining the fault tree and the similarity measures;
4. updating the previous points when additional knowledge/cases are available.

The authors worked in close co-operation with the B4 robot cell specialist at each step of the development. However, it quickly became clear that this approach did not match completely with the specialist's one, since he expressed his knowledge directly in terms of rules. Therefore an additional tool was developed to capture his knowledge in a theoretical failure tree, in order to extract analytical cases automatically from this tree. This resulted in a move from the normal development scheme presented above, to a more integrated one where cases can be described both in an analytical way by the specialist, and in a standard way by using the system and logging its results. Finally, it was decided to apply an inductive learning approach in order to create a common decision tree from both the analytical and the real cases. Additionally, a note was made of statistics of the cases reached when the system is running, so that the measure of similarity between the problem and the different solutions evolves over time in order to favour the most probable solutions. Therefore, the different steps are:

1. defining/updating the data model;
2. creating/updating analytical cases through the theoretical failure tree;
3. creating/updating real cases through analyzing the log reports;
4. merging the analytical and the real cases;
5. creating/updating the fault tree;
6. updating the statistical records of cases to assess the similarity.

When the user starts a consultation, he first browses through the fault tree until the diagnostic system reaches one or several leaves. If several solutions to the query are found, they are ordered by the similarity measure, which combines the likelihood of each solution together with a pure metric between the query and each case.

At the time when the system was put into operation (15 February 1999), the case base contained more than 250 cases. The application covered an estimated 95% of the known faults which generate an error code on Rob-Ex (in excess of 40 error and sub-error codes), giving access to approximately 80 repair procedures. The development of the case base is managed by the service co-ordinator (see Section 6.2), who has been trained in order to be autonomous on the adequate tools (data model builder, theoretical failure tree builder, case extraction, log reports analysis, decision tree building and integration in the existing diagnostic system).

3. Hypermedia documentation

3.1. Principles

Whilst the CBR technology provides the diagnosis, CBR does not provide any of the information required to undertake a repair. In most cases extra information is needed, including the location of the fault, together with the relevant repair and recommissioning procedures. In the case of a large industrial plant this information is currently supplied in paper or electronic format from a large number of suppliers. Thus, even if the fault is known, the operator still has to locate the correct procedure, from what in may robotic systems can be a small library of information.

Hypermedia is particularly suited to engineering applications, as information is conventionally organised into a large number of relatively small documents, with a considerable amount of cross-referencing. In addition, at any one time the user will only require access to a small fraction of the available information resource.

The concept of industrial-strength hypermedia as a solution to information management in manufacturing was initially proposed by Malcolm et al. (1991). Malcolm argued for hypermedia systems that had evolved beyond a stand-alone status to become a technology that integrates resources over the complete engineering enterprise. In a previous industrial hypermedia application, Crowder et al. (1996) demonstrated that by using an open hypermedia system, a common knowledge base can be used for a range of tasks.

In any hypermedia application two sets of data are required, the resource documents and the linking information. It should be noted that a number of industrial information systems have the link information embedded within the resources, resulting in a closed application (Greenough and Fakun, 1998). The use of embedded information will restrict the application of hypermedia systems to industrial applications, as any document change, however minor, could lead to a significant re-authoring exercise. An open hypermedia system permits the development of industrial-strength hypermedia, which incorporates an easily maintained system which can be integrated with existing networks, databases, and as discussed in this paper, knowledge-based system technologies.

3.2. Application in HELPMATE

The B4 robot station information consists of a considerable number of loose-leaf documents, drawings ranging from A4 to A0, and proprietary information, including engineering drawings, location drawings and test schedules. To speed the processing of text-based material, use is made of optical character-recognition software capable of processing documents in Danish. During the development of this application, at times it was considered quicker to redraw, or re-enter the information in the correct format, than convert the documents to electronic format, particular if the quality was poor. However some concern has been expressed over the maintenance of the audit trail in this situation.

In all, over one hundred documents were converted to populate the resource base of the pilot system. Following the construction of the resource base, the link base can be constructed. MicroCosm is supplied with a set of tools that ensure that authoring is an efficient process. Three different types of approach to linking are required.

- Structural linking of the documents, for example indexes to sections of manuals. This is a routine process that can easily be automated. The hierarchical structure of many technical documents comes to the aid of the industrial author, as most chapters, sections, and subsections, etc. are formatted using heading styles, or at least have a different format from the rest of the text.
- The robot specialist who is experienced in the subject area makes manual links. This turns the information into a cognitive and pedagogical structure that is easy to navigate.
- The robot specialist makes the links between the solutions provided by the KATE system and the information resources provided.

The development of the application is of necessity an iterative process. The robot specialist was asked which documents were linked, and then the developers implemented this and confirmed with by the robot specialist. In this way the operators can capture the organisational memory for use in the cell.

The number of documents incorporated in the delivered system was more than 500, and the number of links was more than 2000.

4. Integration between CBR and hypermedia

The Freya system required the technology partners to extend their basic software, KATE and MicroCosm Pro, to function in an integrated manner as a hypermedia diagnostic system interfacing to the Rob-Ex cell control system. The generic integrated part has been called KateCosm, and is designed to be reused for other similar applications. The KateCosm architecture is based onto three reusable objects:

- the KATE filter;
- the MicroCosm filter interface for creating indirect links (set of functions);
- the KATE system book.

These different parts are detailed in the following.

4.1. The KateCosm architecture

4.1.1. The MicroCosm Pro architecture

MicroCosm Pro integrates the various required functions into a single application system. MicroCosm systems contain "filters" that carry out actions (e.g., Compute Links) requested by the user or requested by other filters: MicroCosm "viewers" present resources such as documents or video clips to the user and capture the actions the user wishes to apply to the resource. The communication is via an internal message protocol that can be extended to other system components via the Windows DDE system. A new *KATE filter* is added to MicroCosm to integrate Rob-Ex and KATE into the MicroCosm system. Being a filter, it can receive messages from other parts of MicroCosm, and it can send messages to be handled by other components. The filter accepts messages from the Rob-Ex filter, to pass control codes from the Rob-Ex system into KATE. The filter embodies a new set of common functions (a Windows DLL) that serves as the MicroCosm side of the code interface to KATE. The filter receives MicroCosm messages from KATE, indicating the links to be made available, and generates the MicroCosm messages that cause MicroCosm to present the links (typically by a Microcosm Results box).

Microcosm Pro is an environment that supports reading and authoring from the same configuration and via the same user interfaces (though it is always possible to install MicroCosm so that authoring of hypermedia links is unavailable for particular users). MicroCosm supplies "linkbase filters" to create and edit the links. The request to add a link is made via Viewer's "Action" menu item while the documents are in the viewer: the request is passed to the linkbase filter which in turn communicates with the user to capture further details about the intended new link (a link title, for instance).

MicroCosm's link creation process supports Button Links — the standard hypertext concept in the WWW or Windows Help, and "Generic Links". The latter are associations between a phrase, the source selection, and a destination that may be in any document. But rather than applying just to a single occurrence of the phrase in a given document, a generic link applies to any occurrence of the phrase anywhere in the document set. Consequently the link has to be authored once only — but is effective from any number of documents.

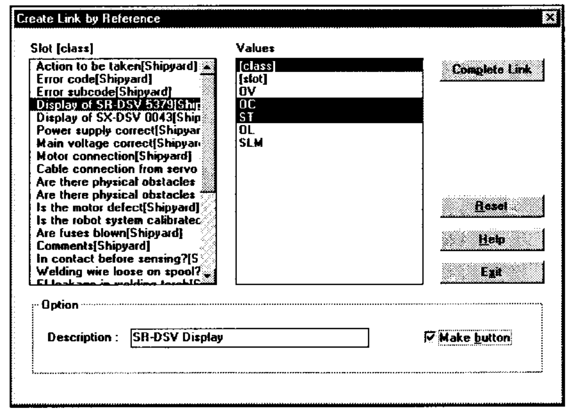

Fig. 1. MicroCosm interface for authoring indirect links. When building a link, the user has to start from the document to be linked. Then he chooses in the Kate object (class, slot or value) that has to be linked to the document. The link is not done directly to the name of the chosen object, but to its internal identifier. As a consequence, changing the name of an object, or adding a new name in another language does not affect the links. MicroCosm Pro allows multiple linking: the same object can be linked to several documents.

4.1.2. Extension of the indirect links

The generic link idea is taken further in KateCosm to tailor links to the CBR environment — the "Indirect Link", or *Linking by Reference*. These indirect links have Source Selections that refer to a KATE data model, expressed by the Classes, Slots and Values used in the Case. This offers a number of benefits — thus Natural Language can be switched just by substituting a new table, and diagnostic cases can be rearranged without having to re-author the associated hypermedia links. KateCosm has a *MicroCosm Filter User Interface* for authoring indirect links (see Fig. 1).

4.1.3. Use of the indirect links in KateCosm

The MicroCosm Pro message model ties the KateCosm software components together. Events from Rob-Ex are presented to other MicroCosm filters as MicroCosm messages. MicroCosm messages are communicated from KATE to MicroCosm to indicate the hypermedia links to be offered to the user at each stage during the consultation. In the latter case the messages indicate the Table indexes of the hypertext links needed at a given stage in the consultation. The communication between KATE and MicroCosm is ensured by a common set of functions embedded in a Windows 'DLL' (or 'Dynamic Linked Library'). This DLL contains functions like:

SendMessageToMCM /* For example, does a link exist for a specific slot? */

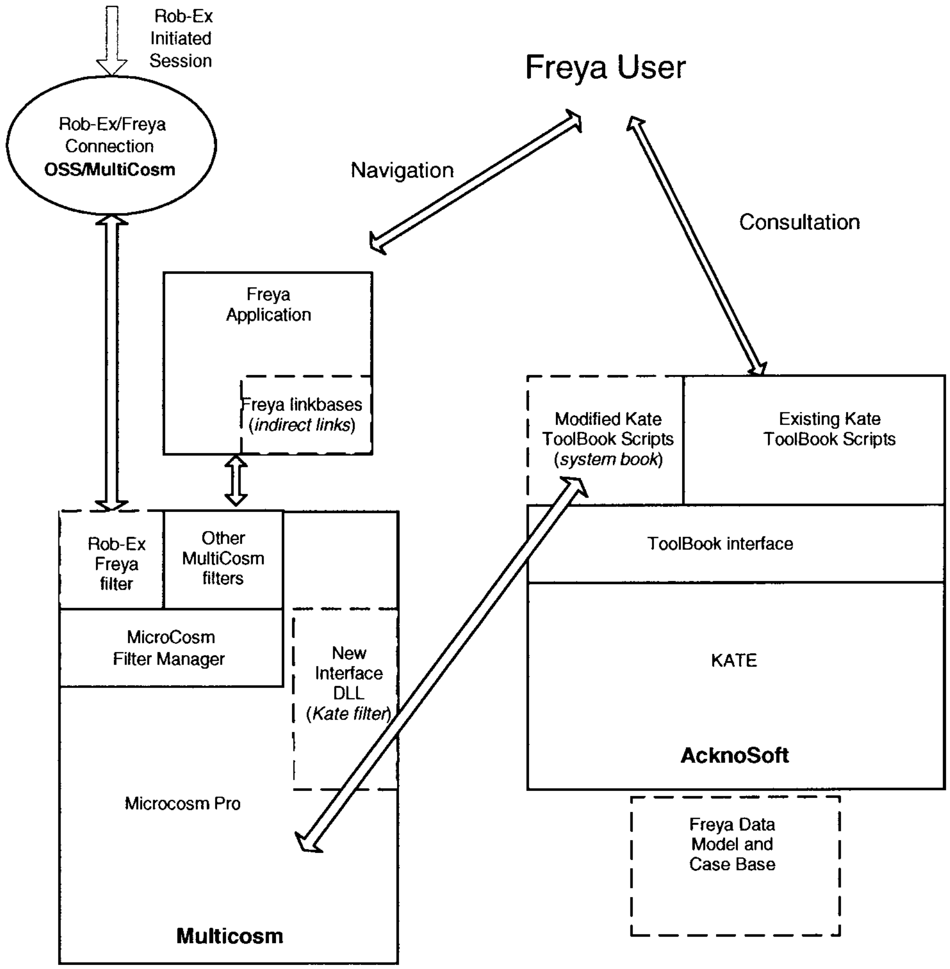

Fig. 2. Freya Architecture. The Freya user starts the navigation in the MicroCosm environment and can follow the links found in the Freya linkbases. From there he can start a consultation and launch the Freya system book. When browsing the fault tree, the indirect links between the Freya data model and the Freya linkbases are displayed automatically and the user can follow them thanks to the Kate filter.

SendMessageToKATE /* For instance, exit KATE */
SendShowLinks /* Displays the linked document */

A KATE CBR application can be written in one of two ways — using a scripting language (e.g., ToolBook) to develop all the CBR case screens, or by programming these screens directly using the C interfaces to KATE. KateCosm takes advantage of the ToolBook approach to minimise application development effort and complexity. KATE offers a ToolBook System script, the *System Book*, to integrate KATE with other software. The KateCosm integration exploits this script, which has been extended to cope with the new functionality required, by calling the adequate functions of the common DLLs. The preceding pseudo-code enables one to highlight (in blue) the slots in the KATE model that have a link in the MicroCosm linkbase, and to display the link if any.

Example.

Code = SendMessageToMCM(*slot*)
IF Code = 0
  set storkeColor of *slot* to black
ELSE
  set strokeColor of *slot* to blue
  get SendShowLinks(*slot*)
END

4.2. Freya architecture

Freya is the application based on the KateCosm architecture, specifically designed to diagnose the B4 robot station. Therefore, the main components of KateCosm are retrieved, applied to the welding robot

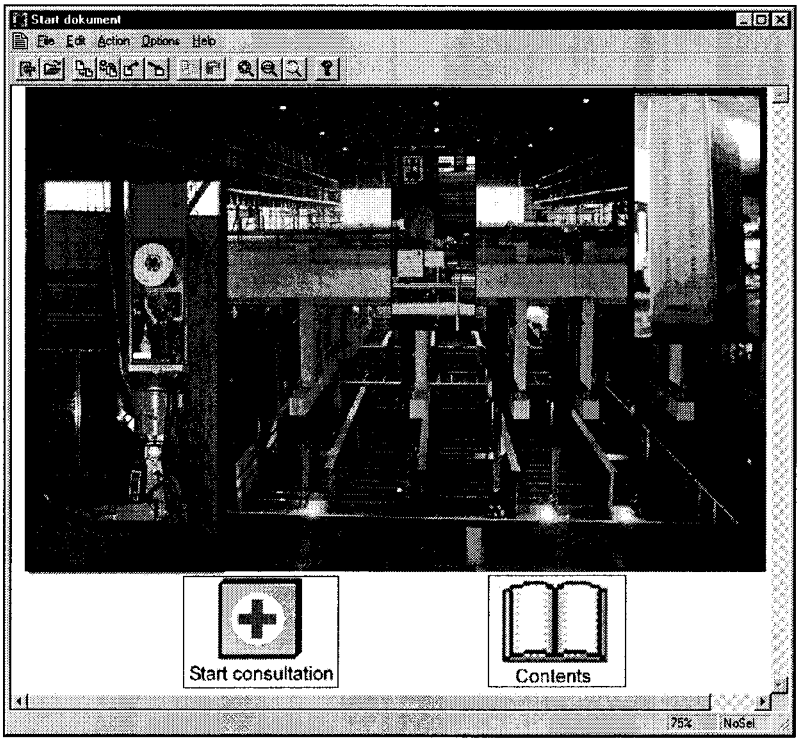

Fig. 3. Freya start-up document. The picture shows the B4 hall and short-cuts to the robot documentations. From this document, the user can either browse through the documentation, or start a consultation for diagnostic support.

specific data model, case base and hypermedia documentation. Fig. 2 presents the architecture.

The data model, and therefore the case base, is available in two languages: Danish (for the end-users) and English (for the demonstrations and papers like this one). Thanks to the indirect link mechanism, the same link, is used whichever the current data language used (see Section 4.1.2).

The interface is available in two languages as well (Danish and English). This feature is supported by standard resource files.

Some additional components have been requested by OSS for improving the access to the system.

- The wrapper. This additional piece of software is the start-up screen of Freya. The wrapper contains the list of Freya users, including their names, surnames, photos and access rights (with the possibility to modify the size or colours of the pictures, to create new links etc.). When starting a Freya session, the user clicks on the icon that corresponds to his/her name. Everything he/she does is then logged for further analysis. Depending on the access rights, a password may be required. Two default user types are proposed: "super" (password protected) and "test".
- The tool bar indicates who is logged in, and for how long. The tool bar is always visible. When a user wants start a new session, he has to select "Genstart" ("New start") from the tool bar in order to log out the previous user. The tool bar is also used to display messages from Rob-Ex.
- The Rob-Ex Freya filter captures information generated dynamically by the Rob-Ex tool and sends it to Freya. Right now, the only captured information comprises the error codes provided by the robots. They are displayed in the tool bar.

Note that the wrapper and the toolbar are important since there is only one PC available in the hall.

5. Current state of the application

A small-scale pilot system has been used at OSS since the summer of 1998, to obtain initial user feedback, and to let the operators gain confidence in the technology. The feedback has been highly positive, with the users taking a degree of ownership of the project. This can be achieved by actively encouraging the users to voice their concerns and criticism, and feeding this back into the design process. As will be appreciated, the factory floor is a hostile environment for any piece of computer equipment, in particular when the operator's input device is subject to dirt and grease. While Freya could be kept in a supervisor's office, to gain maximum benefits from the system, it was clear from the outset that a machine on the factory floor

Fig. 4. The tool bar indicates the name of the current user (MT, Morten Tellefsen) and the duration of use. It can display messages from Rob-Ex.

would be required. For the interface device a modified tracker ball is used, these are easy to use and do not suffer from the problems experienced by the mouse or touch screens.

A first full-scale pilot application was delivered in January 1999, to obtain feed-back from the super-users. The first version was installed in the B4 hall in February 1999 and is regularly evaluated internally and updated with respect to the users' wish list (fixing bugs, updating the diagnostic process and documentation).

5.1. Typical session: starting point

A typical session starts when the operator selects his name from the user's list presented in the wrapper. If an existing session has already been opened by another operator, the current state is stored in a history file so that the previous operator can restart his session at his point of use (it should be recognised that only one PC is available in the B4 hall for Freya).

Then the start-up application screen appears (Fig. 3). The operator is presented with two options, to use the system as an information resource by following the various links proposed in the start-up screen, or to follow a diagnostic procedure by selecting the "Consultation" link.

As soon as an operator is selected, the tool bar starts recording the time spent (Fig. 4). The only way to interrupt a current session is to click on the 'Genstart' button of the tool bar, which stands for 'New start'.

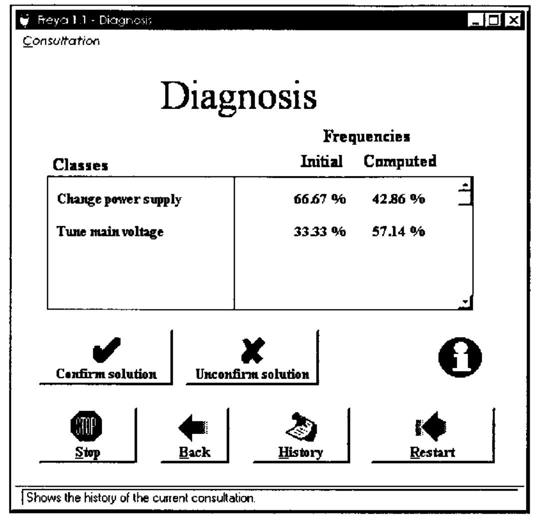

Fig. 6. Once the operator reaches a set of possible solutions, he can compare the initial probability of each solution as provided in the case base and the computed frequency based on real cases. He can confirm the solution (and hence, update the case statistics) or unconfirm it (then he has to provide the correct solution).

5.2. Typical session: diagnostic support system

When starting a diagnostic procedure the operator follows a question-and-answer dialogue session, using the user interface as shown in Fig. 5. Once the answer is selected, selection of '*OK*' takes the user to the next stage in the decision tree. Other possibilities are offered.

- If the question cannot be answered easily, he can

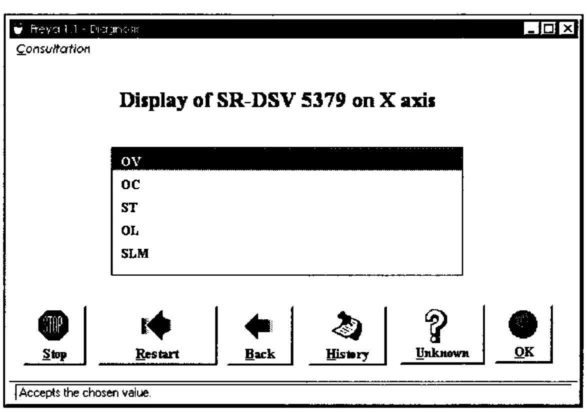

Fig. 5. Typical session during the diagnostic process. At each level of the decision process, the operator can display the list of solutions reached so far, go back and modify a previous level, select one of the proposed values or select "Unknown" to expand the search.

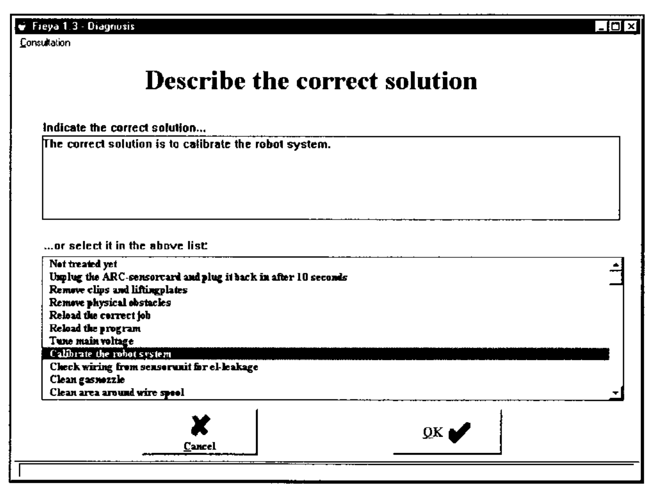

Fig. 7. If the solution proposed by the system is not correct, the operator can provide his own expertise.

```
*************************************************
Session opened on 14/05/99 at 11:33:28 by Borkur Antonson.
*************************************************
Information provided during consultation:
Error code = 2
Error subcode = 5
Display of SR-DSV 5379 on X axis = OL
Are there physical obstacles for the robot = ???
Are there physical obstacles inside the gear = Yes
DIAGNOSIS NOT CONFIRMED:
Replace robot gears                     50 %     75 %
Remove physical obstacles               50 %     25 %
CORRECT DIAGNOSIS IS: Calibrate the robot system
The correct solution is to recalibrate the robot.
*******************************************************
Session closed on 05/14/99 11:34 AM at 11:34:37 by Borkur Antonson.
*******************************************************
```

Fig. 8. The log report can be further analysed for improving the system.

select '?'. This will expand the number of solutions proposed by the system.

- It is possible to modify information provided previously during the consultation, by selecting the '*Back*' or '*History*' buttons.

At the end of the consultation, the system retrieves, with their respective probabilities, the most probable solution based on the operator's knowledge of the problem and on case statistics (Fig. 6). The operator has the opportunity to confirm the solution proposed, which changes the statistics of the case base and therefore the computed frequency of solutions, or to unconfirm it. In the latter case, he is required to indicate what would be the correct solution (Fig. 7).

All operators' actions are traced into a log report, which is regularly analysed by the B4 specialist in order to update the case base (Fig. 8).

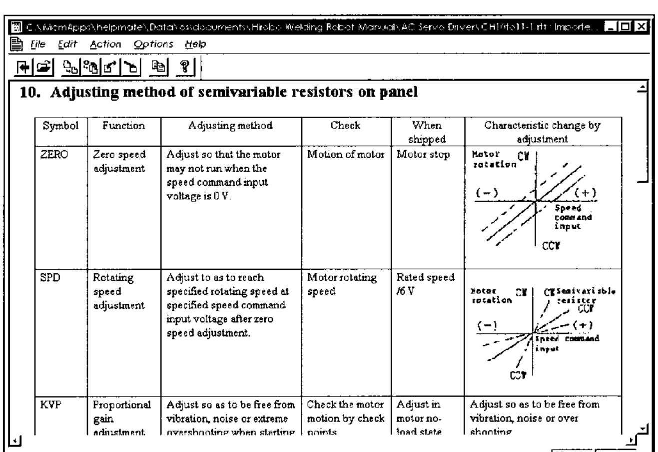

Fig. 9. Typical information resource. The documents are linked to each other thanks to the MicroCosm linkbase.

5.3. Typical session: browsing through the information resources

When using the system as an information resource the operator is taken to an index page, and by following links can navigate through the resources to locate the information. The operator can browse through the information resources from the start-up screen. What is much more interesting for him is that, at any place during the diagnostic process, a point may be reached where information is required. A link then takes the user to the open hypermedia resource base, allowing the repair to be undertaken. The KATE filter automatically paints the active links in blue (this is not shown in the black-and-white screenshots depicted here). Fig. 9 shows a typical page of information.

Remark. *Recall that the figures presented above have an English interface for the sake of clarity. In the system installed in the B4 hall, the user interface and the data model are in Danish. This does not affect the linking procedure thanks to the indirect links. Additionally, this multi-language approach will make it easy to evolve to other languages in the future.*

6. Evaluation

This section describes the current results of this ongoing project, both in terms of the development efforts involved, and in terms of the improvement in performance and easiness of use of the diagnostic process.

Table 1
The type and number of Freya users

HELPMATE users	No. of users	Reserve users	No. of super-users within each user type
Operators[a]	12	5	1
Maintainers[b]	3	1	1
Foreman[c]	1	1	
Service co-ordinator[d]	1	0	1

[a] As shown in the Table, there are 12 full-time operators and five reserve operators. The welding robot installation in hall B4 is operated 24 h daily, by three separate shifts. There are four operators per shift. The operators' main tasks are to perform daily robot welding operations, to repair minor faults and to co-ordinate major repairs.

[b] There are three (not full-time) maintainers and one reserve maintainer. The maintainers' main tasks are to perform routine preventative maintenance and major repairs.

[c] There is only one foreman responsible for the robot system. The foreman's main tasks are to co-ordinate and plan the production of assemblies at the robot station in shop B4.

[d] There is one service co-ordinator for co-ordinating and planning the routine and emergency servicing of all robot installations in the yard. He is also responsible for helping other shipyards, using the OSS robot systems, to detect robot faults. The service co-ordinator's role has not become fully established, because the position is relatively new.

6.1. Main dates and efforts involved

The HELPMATE project started in September 1997 and will end up in November 1999. Most of the time during the first year has been spent on technical work involved in building the bridge, called KateCosm, between the case-based reasoning and the multimedia tools. In parallel, a pilot system was developed and delivered in June 1998 in order to collect the end-users' requirements. The first version of the system was delivered in January 999 and installed after testing in the B4 hall on the 15 February 1999. A second version was delivered in April 1999. During the development process, an end-user's "wish list" has been continuously maintained in order to improve the system in the desired way.

The overall effort required for the HELPMATE project was estimated as follows:

- technical integration between the case-base and the hypermedia tools — 7 months;
- model building and case authoring — 6 months;
- hypermedia preparation — 7 months;
- system development — 9 months;
- organisation (training, operation, tuning, evaluation ...) — 10 months.

In practice, it appeared that the technical part (integration and system development) required more effort than expected. For the integration part, this is due to the fact that it was desirable to create a generic integration procedure between KATE and MicroCosm, beyond the target application. For the system development part, this is due to the two-step approach chosen (pilot + several versions delivered). The user requirements that arose from the pilot were much more complex and difficult to complete than had been expected when we started the project. Additionally, new tools were required, both for the development part (e.g., the tool for building analytical cases from fault trees) and for the application part (e.g., the wrapper, and the tool for logging the end-users' actions). On the other hand, the case authoring was easier to complete, thanks to the pilot system. It was recognised that analytical cases were needed, which were available directly from the super-users in the form of fault trees. The analysis of cases provided by the end-users was eased thanks to the logging facility.

6.2. Freya users

There are a number of different types of users of Freya. Table 1 lists the various categories, and indicates the number of users in each category. "Normal" users can consult the system. They can confirm or unconfirm a solution proposed by the diagnostic tool. Their actions are logged into a log file so that it is possible to create new cases or to modify existing cases by analyzing the content of the log files. More generally, all actions of the users (cases used, solutions obtained, documents browsed ...) are logged for further analysis. The super-users have different rights in Freya. They can create new cases or validate existing ones, modify the data model, create new links between the documents, etc. Also, the super-users do the initial testing of a new version of Freya. Among the super-users, the service co-ordinator has a specific role: he is responsible for keeping the case base, the document base, the data model and (more globally) the diagnostic system up to date.

6.3. Results of the evaluation

The main goal of the second phase of the project, which started with the release of the first version of Freya in the B4 hall in February 1999, aims at evaluating the results of the system. The system is used in daily maintenance by the users, and has been evaluated in several ways.

- By questionnaires to ensure that the system fulfills the specifications stated in the user requirements. This is a very objective method of evaluation.
- By interviews to enable the users to state their more subjective opinion of the complete system, and the separate functionalities.
- By automatically logging the use of the system. This will show the extent to which the system is used, and which parts of the system are used by which users.

The evaluation is being performed in several steps during the evaluation period, and the result of the evaluation are used to refine the system.

It is too early in the evaluation phase to have a clear measurement of the quantitative gains associated to the use of Freya. However, according to the Odense Steel Shipyard, the initial results of the evaluation show that Freya will contribute to obtaining a high level of reliability, and will thereby enable efficient planning of complicated technological robot gantry systems. The results of the HELPMATE application are today mainly qualitative, and in the short time of use the system has proved to:

- Involve the robot operators in the maintenance process by creating greater competence in diagnosing and repairing minor failures in the system. The robot operators do not need to refer systematically to the maintainer each time a failure occurs.
- Reduce robot down time caused by repetitive faults and no cause found, which can affect the cost effectiveness of the gantry system.
- Exchange and expand knowledge and experience concerning maintenance. This is particularly true of the maintainers, the foreman and the service co-ordinator.
- Accelerate operator training. The introduction of Freya required regular teaching sessions by the service co-ordinator to the operators. As a consequence, operators working during the day shifts could share their experience with those working during the night shift.
- Create an on-line help system for the operators and give the operators easy access to pertinent information.

The evaluation of the system will continue until the end of the project in October 1999, but it has already been decided to use the system in production and maintenance after the end of the project.

7. Comparison with existing similar work

There is a growing number of applications within the CBR market. In the recent book about the methodology for building CBR applications (Bergmann et al., 1999), the authors present seventeen examples of running applications. The most convincing CBR systems appear to be in the help-desk areas, particularly for troubleshooting complex equipment, and not in the manufacturing process. In a recent paper where Gunasekaran (Gunasekaran and Love, 1999) summarizes the current state of the art across manufacturing, it is clear that the majority of current research is being used to support manufacturing (e.g., training, design information) rather than (as in this paper) be directly integrated into the manufacturing process. In (Althoff et al., 1998), the authors present the CBR as a means of managing the knowledge in the organizations. The proposed scheme is viewed as a solution for a "learning software organization", which is wide enough to include most CBR systems. However the presented work is more orientated towards software engineering best practice and business processes, rather than specific implementation in a manufacturing process.

Today many CBR systems are available that are integrated with hypermedia resources. One example is the Analog Devices applications (see Bergmann et al., 1999 chapter III) that gives access to data sheets of operational amplifiers (either through PDF files on a CDROM, or through HTML files on the WEB) of retrieved cases. Most CBR tools allow one to define a "link slot" that enables a user to connect a case together with its multimedia reference. However, the KateCosm approach extends this standard relation in several ways.

- The linking service provided by Microcosm is open to any source document (Word, Excel, PDV, AVI, AutoCAD, etc) in a native format.
- It allows multiple linking from a single source.
- The links are not tied to specific cases, but to the case data model itself. New cases may use these links without any additional work.
- The links are language-independent. A unique link may be defined, say, in English, and reused when the data model is turned into Danish or German.

Finally, the integration between KATE and Microcosm is defined at a generic level, and not only for the Freya application. The authors are not aware of any CBR system that seamlessly integrates an open hypermedia system at a generic level, as is proposed in the KateCosm solution.

8. Conclusion

The HELPMATE project aims to take the diagnostic capabilities of case-based reasoning and augment this with open hypermedia linking technology to provide the operator with important context information regarding the problems and solutions presented to him. This has been achieved through a seamless integration of two industrial products, KATE and MicroCosm Pro.[1] It is a major challenge, not only to provide a single application for the specific need of the B4 robot station at OSS, but also to develop a pragmatic and repeatable process in order to ensure that the system receives acceptance by the people responsible for robot maintenance. In the future, Freya is expected to be used for other robot stations and other types of equipment, such as cutting machines and cranes. The industrial partners of HELPMATE will sell the KateCosm solution as an add-on of their own specific product.

Acknowledgements

The authors acknowledge the European Commission for funding HELPMATE under grant number ESPRIT IV, project 25282. They acknowledge as well the reviewers of the initial version of this paper, whose comments have greately helped to clarify and to improve its content.

References

Althoff, K.-D., Bomarius, F., Tautz, C., 1998. Using Case-Based Reasoning for building Learning Software Organizations. In: Proceedings of European Conference on Artificial Intelligence, Workshop on Building, Maintaining, and Using Organizational Memories (OM-98) Brighton, 24–28 August.

Auriol, E., Manago, M., Althoff, K.-D., Wess, S., Dittrich, S., 1995a. Integrating induction and case-based reasoning: methodological approach and first evaluations. In: Haton, J.P., Keane, M., Manago, M. (Eds.), Advances in Case-Based Reasoning. Springer–Verlag, Heidelberg.

Auriol, E., Wess, S., Althoff, K.-D., Manago, M., Traphöner, R., 1995b. INRECA: a seamlessly integrated system based on inductive inference and base-based reasoning. In: Veloso, M., Aamodt, A. (Eds.), Proceedings of ICCBR 95, First International Conference on Case-Based Reasoning. Springer–Verlag, Neidelberg.

Bergmann, R., Breen, S., Göker, M., Manago, M., Wess, S. (Eds.), 1999. Developing Industrial Case-Based Reasoning Applications — The INRECA Methodology. Springer.

Crowder, R.M., Wall, W., Heath, I., Bernard, R., Gaskall, D., 1996. A hypermedia maintenance information system. IEE Computing and Control Engineering Journal 7 (3), 109–113.

Greenough, R., Fakun, D., 1998. An innovative information system to support team based maintenance. In: Proceedings of Third International Conference: Managing Innovative Manufacturing, Nottingham, UK, July.

Gunasekaran, A., Love, P., 1999. A review of multimedia technology. Manufacturing. Computers in Industry 38, 65–76 January.

Lenz, M., Auriol, E., Manago, M., 1998. Diagnosis and decision support. In: Lenz, M., Bartsch-Spörl, B., Burkhard, H.-D., Wess, S. (Eds.), Case-Based Reasoning Technology — From Foundations to Applications. Springer.

Malcolm, K., Poltrock, S., Schuler, D., 1991. Industrial strength hypermedia: requirements for a large engineering enterprise. In: Proceedings of Hypertext '91.

Manago, M., Auriol, E., 1998. Applications of inductive learning and case-based reasoning for troubleshooting industrial machines. In: Michalski, R.S., Bratko, I., Kubat, M. (Eds.), Machine Learning and Data Mining — Methods and Applications. John Wiley and Sons.

Quinland, J.R., 1983. Learing efficient classification procedures and their application to chess end games. In: Michalski, R.S., Carbonell, J.G., Mitchell, T.M. (Eds.), Machine Learning 1: An Artificial Intelligence Approach. Morgan Kaufmann, Redwood City CA.

[1] Further information on KATE can be found at http://www.acknosoft.com/ and further information on MicroCosm can be found at http://www.multicosm.com/.

Eric Auriol completed his PhD (Integration of Induction and Case-Based Reasoning Technologies for Industrial Applications) in 1995 at the French Institut de Recherche en Informatique et Automatique (INRIA) under the supervision of Professor Edwin Diday and at AcknoSoft. He spent 16 months at the University of Kaiserslautern in Professor Michael Richter's research group. Author of 30 publications on the integration of CBR and Data Mining for service support systems, he is now project manager at AcknoSoft, and is actively involved in developing and installing CBR applications in engineering.

Richard Crowder is the senior lecturer in Robotics and Control in the Department of Electrical Engineering, University of Southampton, UK. He received a BSc in Engineering and a PhD in Engineering from the University of Leicester. Following graduation he was employed for 11 years in the machine-tool industry, developing specialist machine-tool systems for aircraft engine manufacturing. He joined the University of Southampton in 1984. His current research interests are in to the application of information technology to manufacturing industry. Current research projects include the development of industrial-strength hypermedia information systems to support process line operation; this work is being funded by the EPSRC in collaboration with a number of major companies. A related activity is the development of specialist robotic systems, including research into the control of flexible, multi-fingered end-effectors for unstructured environments. He has published extensively, both at conferences and in learned journals. He is the author of a text book on the application of advanced drive systems. Dr Crowder is an associate members of the Multimedia Research Group which is based in the Department of Electronics and Computer Science at the University of Southampton. He is a Member of the Institution of Electrical Engineers and a Chartered Engineer.

Roger Row is the Development Director of Multicosm, and one of the original founders of the Company in 1994. He has led the commercialisation of the Microcosm software technology, originated at the University of Southampton, on which the Microcosm Pro product is based. He has a BSc (Eng) from the University of Southampton and is a Member of the British Computer Society. He has spent much of his 35 years in the IT Industry in System

Architecture, with extensive experience of Languages, Computer Architecture, International Standardisation, and Software Product development both in the UK and in the US. He directed the development and launch of world-wide products for IBM, and was responsible for a number of advances in IBM computer and software architecture. As Presentation Management Systems Manager during the 1980s, he managed the group of programmers responsible for the software architecture spanning all IBM system for graphic, textural and image data and for the various software implementations at the presentation level. The first products in this series were for mainframes and midrange systems (GDDM, the Graphical Data Display Manager). These were followed by the development and launch, jointly with Microsoft, of the first three releases of the OS/2 Presentation Manager and by IBM's own 32-bit implementation of the OS/2 Presentation Manager.

Thomas Knudsen is since 1997 group manager for mechanical research and development automation at Odense Steel Shipyard. He has been involved since 1994 in the yard in the engine room design of 6600 TEU container vessels, in the mechanical and electrical development of robot joints for manipulators used in shipbuilding, and in the project management for robot welding system installed in various parts of Europe. He has a BSc Eng (Mechanical systems and processes). Since 1999, he is committee member of Danish Industrial Robot Association (DIRA).