

## Preface

These proceedings are of the 1st. UK Workshop on Case-Based Reasoning, organised by AI-CBR and the British Computer Society Specialist Group on Expert Systems. AI-CBR, an Internet forum for discussion and dissemination of Case-Based Reasoning (CBR) research, was founded in February 1994. In less than a year it has grown to have several hundred members from most parts of the world and has become a focus for researchers in CBR. This workshop was organised as an attempt to let the UK CBR community meet each other. Some of us know each other from other conferences and professional relationships, but for many this workshop was an opportunity to meet people who had previously only been e-mail names.

A decision was taken early in the planning of the workshop that it should be of a general nature, i.e., of both theoretical and practical value. To this end the event was divided into three broad areas:

1. papers on theoretical aspects of CBR,
2. presentations by CBR software tool vendors, and
3. papers on CBR applications.

It was hoped that this mix would be of interest to everyone. In particular, the event was the first opportunity for people to see comparative demonstrations of the following leading CBR software tools:

- Inference demonstrated CBR Express, their successful help-desk application. Some researchers criticise tools like CBR Express for being *too* simple. However, many of us can learn from Inference's philosophy. Their tools provide *business solutions*, not just *technology*. It is thus no surprise that this US company is currently a market leader in the provision of case-based support to customer help desks with over 13,000 licences sold world-wide.
- In contrast Rob Milne of Intelligent Applications showed an undoubtedly more powerful tool from the US called ReMind. ReMind, in contrast to Inference's products is packed with technology and Rob did well to show most of its functionality in 40 minutes. ReMind is powerfull and flexible tool that has been used successfully in a variety of applications, which were briefly illustrated.
- ISoft are a French company with a relatively new tool called ReCall. As usual the Europeans at the event shamed us with excellent command of English. ReCall in contrast to the US products is object-oriented. ReCall represents cases in an object hierarchy rather than the flat-file representation of the older US products. ReCall's objects have feature slots and slot monitors or deamons. However, although OOPS has proven itself to be a powerful representational and programming paradigm in both AI and conventional systems it is yet to be proven if case-bases benefit from the application of OOPS.

- Finally, Acknosoft are demonstrating KATE, a powerful toolkit that has evolved out of collaborative European research. This is another product that represents data as structured objects. KATE-INDUCTION can build decision trees from training cases whilst KATE-CBR dynamically builds paths to similar cases respecting the wishes of the user.

In 1991, Ian Graham, chair of the British Computer Society Expert Systems conference, wrote that “*CBR promises to become one of the major new methods for developing expert systems*”. Many of us would agree with this comment. Three years later, Derek Sleemen in his key note address to the same conference, commented that CBR appeared to be the *hot-topic*. Whilst it is gratifying for one’s own research area to be in the spotlight it carries with it a danger. Few of the older amongst us need reminding of the AI-Winter, which many believe was partially caused by the hype surrounding expert systems in the late seventies and early eighties. We are in a dilemma. We have to hype CBR in order to attract research funds and to interest companies and their managers in the technology. Yet by hyping CBR we may be raising expectations that could return to haunt us. For example, Fari Marir in his paper points out that CBR does not remove the knowledge elicitation bottleneck, as claimed by some, it merely eases it.

I would like to suggest that in the UK we co-ordinate a careful strategy. We need more good business solutions, like the ones Inference offer and the tool described by Chris Price, to show to industry that here is a low risk technology that can really deliver. At the same time we need research that pushes the limits of CBR within the laboratory and that explores potential solutions. We must try not to make our mistakes in public. For example, we should learn from the lessons of the expert system RI/XCON. It was a ground breaking rule-based system that really worked and saved DEC money. However, as its rule-base grew to 7000 rules and more it became nearly impossible to maintain. Therefore, issues of size and complexity of case-bases should be explored in the laboratory and not in the field.

To sum up, I believe that CBR is, for better or worse, a *hot topic*. It is up to the UK research community in collaboration with our European and international colleagues to ensure that we live up to the hype. This workshop has shown that we have mature software tools and inventive and articulate researchers to use them. There are also no shortage of researchers investigating novel CBR techniques for the next generation of tools. I therefore hope that this is the first in a regular series of CBR Workshops that help all in the community produce better work.

Finally, I would like to thank all those who helped in the preparation for and organisation of this event. In particular Farhi Marir who dealt with the administration, Srinath Perera who collated these proceedings, the Research & Graduate College for providing such a good venue, the British Computer Society Specialist Group on Expert Systems for supporting the event, and EPSRC whose funding (research project no. GR/J42496) helped give me the time to organise it.

*Dr. Ian Watson*  
*University of Salford*