The Knowledge Engineering Review, Vol. 9:2, 1994, 201–218

Book reviews

The second international conference on artificial intelligence applications on Wall Street—Tactical and strategic computing technologies edited by R Freedman, Software Engineering Press, 1993, pp 311, $65.00 ISBN 0-9-38801-07-4.

The second International Conference on AI Applications on Wall Street took place in New York in April 1993. The first conference, held in 1991, was very successful, with many expert systems and other AI applications being described. The second conference follows two recession-bitten years; the contents of the conference are therefore of great interest, for IT systems which succeed in a cash-strapped economy ought to be good investments for several years to come.

The proceedings contain 42 papers spanning 300 pages, on topics including stock market prediction, trading workstation support, marketing and business strategies and understanding news. The majority of papers describe AI applications, for tasks such as intelligent program stock trading, network management, forecasting returns in bond markets and financial data modelling. The papers are generally informative and well-written; however, the brevity required for a volume where the average paper length is seven pages occasionally leaves the reader feeling that he has tasted something good, but has not been given enough to satisfy.

The balance between applications papers and more theoretical papers is good, especially given that the "theoretical" papers generally have a pragmatic slant to them. For example, one paper provides a useful discussion of recent amendments to the case law affecting the copyright of databases. Some of the applications papers appear to have a tenuous relationship with Wall Street (e.g. "A method to predict television audiences based on neural networks"); however, my opinion is that the inclusion of good quality papers on related topics enhances the proceedings, as long as there are not too many of them. I particularly liked the paper on "An expert system for adoption of innovation decisions", which is applicable not only for Wall Street, but for all industries where new technologies are being introduced.

If there is a theme to be drawn from this year's conference, it is that neural networks have replaced expert systems as the most popular AI technology on Wall Street. It may be the large volumes of financial data which Wall Street has collected, or it may be the comparative ease with which neural networks can be developed; whatever the reason, 15 or more of the application papers mention the use of neural networks, while only six describe expert systems. Other techniques which have been used include fuzzy reasoning, genetic algorithms and case-based reasoning. The conference organizers, having identified this theme, invited a speaker to talk on "Neural, genetic and fuzzy approaches to the design of trading systems", and this helpful introductory paper is included in the proceedings.

In conclusion, these proceedings are well worth purchasing if the reader is looking for an overview, or an introduction, to the use of AI in the world of stocks and shares.

Reviewed by John Kingston, AIAI, University of Edinburgh, UK

Explanation and interaction: The computer generation of explanatory dialogues by Alison Cawsey, MIT Press, USA, 1993, pp 232, £26.95, ISBN 0-262-03202-3

In this book, Alison Cawsey presents the results of her thesis work embodied in a system for the generation of explanatory discourse (EDGE). In doing so, she also provides overviews of the fields of text content planning, organization of discourse and user modelling, all of which contribute to the EDGE computerized tutoring system. Cawsey's work is perhaps unusual in not treating these

fields in isolation, but it is her emphasis on empirical work, both to found various aspects of the EDGE system and to evaluate it as a whole, which is truly distinctive within AI.

In brief, the heart of the EDGE system is a central explanation planning subsystem. This subsystem takes input from three sources: dialogue planning rules as developed within conversation analysis; content planning rules designed to express the required semantic content; and a domain specific knowledge base. It also interacts with a model of the discourse and a model of the user. There is no real language processing within the system, which takes input via a system of menus and uses simple templates for generating text output. The system is developed within the domain of rudimentary electrical circuit analysis, and graphical interactions via circuit diagrams are also supported.

The user model is dynamic, being constantly revised on the basis of user's responses and interruptions. Explanation is goal driven, but the goal may change as the user model is revised. Planning and execution are consequently interleaved, with some reactive strategies also being employed. The result of planning is an agenda of partially ordered explanation goals, with the order of the components of the resultant explanation being determined by discourse context.

Cawsey begins with a brief introduction, outlining the context and goals of her work. Chapter 2 examines the structure of explanatory discourse, with special reference to transcripts of expert/novice tutorials. Chapters 3–5 then address the three major areas (text planning, dialogue planning and user modelling) which Cawsey draws together in the EDGE system. In each case, she begins by surveying the field before focusing on the relevant component of the EDGE system. None of the material is particularly ground-breaking in itself (except, perhaps, Cawsey's incorporation of graphical actions into the dialogue). It is the way which Cawsey pulls together the various fields that characterizes her work. This pulling together is demonstrated in Chapter 6, where detailed examples of the functioning of the complete EDGE system are given.

The concluding chapter frankly evaluates the system, noting its strengths but highlighting its weaknesses. Cawsey is harsh on her own system, listing many inadequacies noted by real users. Some of these she addresses, but others, Cawsey notes, require substantial modifications to the underlying explanation-generation architecture.

Four appendices are also included: a selection of complete human explanations (including expert/novice interactions); a description of the planning rules and domain knowledge; a selection of computer/novice dialogues generated with the system; and a selection of sample screen displays.

The book is clearly structured and well written. Though rather specialized, its appeal is broadened by the overviews given in each of the content chapters. Its target audience thus includes researchers in the fields of text content planning, organization of discourse and user modelling. The book provides sound introductory material for each of these fields, and it should be of special benefit to those willing to consider the relationships between them.

Reviewed by Richard Cooper, Department of Psychology, University College London, UK

Explanation and interaction: The computer generation of explanatory dialogues by Alison Cawsey, MIT Press, USA, 1993, pp 232, £26.95, ISBN 0-262-03202-3.

Explanation systems are commonly designed to present information at some fixed rate, or level of abstraction, with the assumption that the user possesses a sufficient level of prerequisite knowledge and expertise within the domain of the explanation, and so will be able to use well whatever explanation gets provided. But such explanations may be ill-adapted for communicating information to some subset of users because this assumption may either not hold for these users, or may not hold uniformly for these users over the target domain. Thus, bored experts and bewildered beginners (or *vise versa*), and users who follow some but not all of the explanation, with the frequent sense of being lost.

Many explanation systems try to minimize this ineffectiveness by allowing the user to identify his