The Knowledge Engineering Review, Vol. 13:2, 1998, 175–178
Printed in the United Kingdom

A Report on Expert Assistants at the Autonomous Agents Conference

ANAND S. RAO

Australian Artificial Intelligence Institute, Level 6, 171 La Trobe Street, Melbourne, Victoria 3000, Australia (e-mail: anand@aaii.oz.au)

Abstract

This report summarises the papers presented in the expert assistants area at the *First International Conference on Autonomous Agents*. It describes the relationship between the expert assistants area and the other areas of the conference, and identifies the major research themes within this area.

1 Introduction

The first conference on Autonomous Agents was organised into four major areas – *Synthetic Agents, Autonomous Robots, Software Agents*, and *Expert Assistants*. The area of expert assistants is primarily concerned with building software agents that assist humans in their decision-making. Often these decision-making situations are time-critical (i.e., the humans have to make the decisions within a few seconds) as well as resource-bounded (i.e. the humans do not have complete knowledge nor infinite time in order to make an informed decision). The area can be viewed as an approach to building agents that integrates earlier work in classical artificial intelligence – perception, planning, reasoning, learning, and natural language understanding with the practical concerns of resource-boundedness and time-criticality.

This area is one of the mature areas of autonomous agents. A number of tools and techniques have been developed in this area and complete agents have been built for a number of application domains including, but not restricted to, medical monitoring, industrial control, business process management, computer integrated manufacturing, and air-traffic management. The conference provided a broad cross-section of papers that addressed some of these applications.

2 Relation to other areas

The papers presented at the conference can be compared along two broad dimensions:

1. The *interactivity* dimension, where one moves from systems that require close interaction with humans to systems that are autonomous in the physical or software world. Expert assistants and synthetic agents tend to be high on interactivity, while software agents and autonomous robots tend to be low on interactivity. For example, an expert assistant will provide suggestions to an air-traffic controller as to the best landing sequence and will be guided by the priorities and choices of the human controller. On the other hand, a software agent or autonomous robot can operate, if required, totally autonomously.
2. The *expression* dimension, where one moves from systems that merely manifest interesting and complex behaviours to systems that have to express such behaviours. Expert assistants and software agents are low on expression, while synthetic agents and autonomous robots tend to be high on expression. For example, an expert assistant interacting with a human has to suggest some mechanism for avoiding conflicts in an air-traffic control scenario, while a believable agent interacting with a human has to show a sense of "relief" or be "overjoyed" with averting a collision.

Expert assistants are low on the expression of behaviours and high on the interactivity dimension. The high interactivity has focused research in the different means of encoding the knowledge provided by experts. This has also meant that whenever the system has not had sufficient knowledge to proceed further in its reasoning it has tended to go back to the user for further clarification rather than reasoning from first principles. This contrasts sharply with some of the early work in classical planning, where generic search mechanisms were devised by planners to reason from primitive actions to form hierarchical and abstract plans. It also contrasts with experts systems approaches that tried to exhaustively capture all the domain knowledge in a particular domain. The feature that characterises current generation expert assistants is the way they "guide" their human counterparts towards certain solutions, and are in turn "guided" by the experts during the execution phase.

Expert assistants also happen to be low in expression. In other words, generating complex behaviours from background knowledge is considered important for expert assistants, but the expression of such behaviours in the physical or software world is not usually accorded the same degree of importance. This is accorded greater importance in designing believable agents and autonomous robots. As a result, expert assistants primarily deal with problems related to knowledge acquisition and execution guidance from human experts in an interactive manner and generating complex behaviours that achieve the goals of its human designers and users.

The discussion above has polarised the area of expert assistants as being high in interactivity and low in expression. The axes of reactivity and expression are more realistically viewed as a seamless continuum. Hence, it is possible for expert assistants to be more autonomous and less interactive with human users. Similarly, they may not only generate complex behaviours, but also express such behaviours in the physical and simulated worlds. In other words, the division of agents into the four areas is more a reflection of the research issues that are the primary concern of each of the areas, rather than a fundamental difference in the approaches to building such agents.

3 Research papers

The papers presented at the conference in the area of expert assistants addressed a number of major themes. First, a number of papers discuss expert assistants that have been implemented and trialled in particular application domains. These papers describe the application and the specific techniques used and problems faced in implementing these systems. The second category of papers provide specific techniques, such as negotiation, communication, collaboration, learning, and planning that are embedded in expert assistants and form the core of these agents. The third category of papers address architectural issues in the design and implementation of expert assistants.

3.1 Implemented systems

Rogers et al. discuss a proof-of-concept prototype that has been developed for teleoperation. The teleVIA (tele Visual Interaction Assistant) assists human operators in diagnosing sensing failures at remote sites. Zhou and Franklin describe situated autonomous agents that provide automatic character feature extraction. They show the utility of this approach in recognizing digits. Haynes et al. describe an automated meeting scheduling system that is designed to be adaptive to environmental demands as well as user preferences. Takada et al. describe a schedule management system that incorporates location information and travel times into the schedule management system. Amant and Cohen describe a knowledge-based planning assistant for intelligent data exploration. It helps users navigate through a large space of decisions involved in exploring a dataset. Malheiro and Oliveira discuss a multi-agent approach to environmental management. Miksch et al. describe the Patient Advocate – an intelligent assistant designed to support obstetrics patients at risk of gestational diabetes.

The above papers provide a good cross-section of applications where expert assistants are being trialled or used on a day-to-day basis. The primary distinguishing feature of these systems from other agent systems is the high interactivity between these expert assistants and their human users.