

Preface

This volume contains selected papers from three workshops held in conjunction with the Tenth Australian Joint Conference on Artificial Intelligence (AI'97), held in Perth, Australia, in December 1997: the Second Australian Workshop on Commonsense Reasoning, the Second Workshop on Theoretical and Practical Foundations of Intelligent Agents, and the Third Australian Workshop on Distributed Artificial Intelligence.

From the outset, the organizers of the three workshops planned a series of coordinated sessions which would enable participants to move freely between workshops. The idea was first to cater for the varying interests of the attendants, but, more importantly, to foster interaction amongst the three research communities which, we felt, were becoming more and more disparate in research topics and methodologies.

To this end, we are grateful for the participation of two invited speakers: Hector Levesque from the University of Toronto, who addressed a joint session of all three workshops, and James Delgrande from Simon Fraser University, who addressed a joint session of the Commonsense Reasoning and Intelligent Agents workshops. Levesque's work is particularly appropriate because it represents an effort to apply techniques from commonsense reasoning to the design of practical agent-based systems. Both speakers provided valuable feedback to the presenters of papers in the workshops, and participated in a lively, if somewhat controversial, panel discussion (as panel discussions ought to be!) on the strengths and weaknesses of the research paradigms within the Commonsense Reasoning and Intelligent Agents communities.

The papers selected for publication were revised (in some cases quite substantially) following comments from referees and workshop participants before inclusion in this volume. In addition to an invited contribution from Hector Levesque (with Steven Shapiro and Yves Lespérance), we have solicited a number of papers from researchers whose work covers the combined areas of interest of the three workshops. Randy Goebel from the University of Alberta has worked extensively on the implementation of logical reasoning systems, and is well known for the Theorist system. His paper (with Li-Yan Yuan and Jia-Huai You) is on providing a possible model semantics for disjunctive logic programs. Michael Wooldridge from Queen Mary and Westfield College, University of London, has contributed a paper (with Afsaneh Haddadi) that presents a formal model of multi-agent cooperation which does not require the agents to agree on a complete joint plan before execution can commence. We especially thank these authors for helping to identify research issues at the intersection of the topic areas of the three workshops, and for thus contributing to the coherence of the collection of papers in this volume.

In keeping with the spirit of the joint workshops, the papers in this volume are organized by topic rather than by workshop. There are 17 papers grouped around five topics: formal models of agency, reasoning agents, communication and coordination, social interaction and practical issues for DAI systems. The papers on formal models of agency concern logical issues in the design of agents

and multi-agent systems, while those on reasoning agents concern logical approaches to commonsense reasoning. The papers on communication and coordination all report work on multi-agent systems using an explicit “high level” model of agents and communication, whereas those on social interaction use simpler “low level” agent models and formalize the (sometimes emergent) properties of multi-agent systems using game theory or decision theory. The final section consists of papers focussing on practical issues in the design and implementation of Distributed Artificial Intelligence applications.

Formal Models of Agency

This section contains papers addressing logical issues in the design of intelligent agents.

Shapiro, Lespérance and Levesque’s invited paper presents an extension of the situation calculus to enable representation of both the knowledge of multiple agents and the evolution of situations (and knowledge) through time. The formalism is then used to specify communicative actions (including the speech acts of *request* and *inform*), and these, in turn, are used to specify a meeting scheduling program written in ConGolog, the authors’ concurrent logic programming language for developing agent applications. Thus the paper demonstrates that commonsense reasoning formalisms are sufficiently powerful for specifying a nontrivial multi-agent system.

Wooldridge and Haddadi’s invited paper concerns cooperation in multi-agent systems. Specifically, they present a formal model of cooperation which does not require that the agents agree in advance to a complete joint plan: rather, cooperation arises out of the individual decisions made by each agent at each time point as execution proceeds. By giving each agent the right of veto over the other agents’ proposed actions so that conflicts are avoided, a group of agents can collectively progress towards a common goal.

Wobcke discusses the logic of ability and the related question of what defines an agent. Broadly speaking, the paper is an attempt to reconcile two different intuitions about agents’ abilities to perform simple actions: on the one hand, an ability must be reliable, i.e., able to be exercised at the command of the agent, while, on the other hand, agents are not infallible, and their ability to perform an action is dependent on the context in which the action is executed. The ability to perform an action is treated as the condition that the action (reliably) succeeds in the “normal” course of events, and the paper presents a logical formalization of normality and ability based on situation semantics.

Lomuscio and Ryan develop the connections between modal logic as developed by logicians (especially Kripke) and the application of models of knowledge to the study of distributed algorithms. In particular, the paper gives a precise translation between Kripke models for logics with multiple modalities (one knowledge modality for each agent) and the “interpreted systems” of Fagin, Halpern, Moses and Vardi, and provides a condition under which the two types of system are isomorphic. The paper thus bridges the gap (in a precise and formal manner) between these two areas of research.

Reasoning Agents

This section contains papers broadly concerned with issues in logical approaches to commonsense reasoning.

Yuan, You and Goebel identify three kinds of logic programming semantics—sceptical, stable and partial-stable—and distinguish two interpretations of default negation: “consistency based” as in Reiter’s default logic, and “minimal-model based” as in McCarthy’s circumscription. The main result of the paper is to furnish a possible model semantics for disjunctive logic programs, introducing the notion of a partial-stable model. Variants of this model coincide with each of the six semantic categories identified.

Nakamatsu and Suzuki present a translation from one class of nonmonotonic ATMS (with out-assumptions used to express nonmonotonic justifications) to the class of annotated logic programs with strong negation (ALPSN). They show a direct correspondence between nonmonotonic ATMS extensions and ALPSN stable models under this translation, and provide an algorithm for implementing the ALPSN model.

Gibbon and Aisbett argue for a reasoning system in which new information is solicited from the user when there is insufficient or inconsistent information, and propose a reasoning architecture which is based on integrating a number of standard inference techniques. Of particular interest is the definition of relative importance (relevance) of information to a query, which enables the system to ask the user questions pertinent to the evaluation of the query.

Communication and Coordination

The papers in this section all draw on explicit models of agency and agent communication to study issues of coordination in multi-agent systems. Two papers (Moulin, Norman and Jennings) discuss issues in communication and negotiation, while another two papers (Ossowski and García-Serrano, C. Zhang and Li) are concerned with coordinated activity arising from multi-agent plans. The papers by Moulin and Ossowski and García-Serrano also present an interesting contrast: both emphasize the role that social relationships play in achieving coordinated activity, but the first paper formalizes this using predefined social roles, while the second uses dependency relationships arising from the structure of multi-agent plans.

Moulin discusses the influence of different social relationships between agents in defining the manner in which they interact. He proposes a modified speech act formalism in which the speaker and hearer fulfil assigned social roles that indicate the power relationship between them. The hearer’s interpretation of a speech act, and hence the appropriate response, depends on this role.

Norman and Jennings present a model of negotiation, considered as a process through which agents reach agreement. The main idea is that each phase of a negotiation involves an agent granting a right to one or more other agents. The other agents may then (legitimately) act on the permission granted to exercise the right, e.g., by accepting a proposal of the first agent, or by offering

a counter-proposal. Granting a right therefore involves an agent in making conditional commitments: commitments to acting appropriately whenever another agent exercises a right. Using a specific language of allowed “moves” in a negotiation, a number of protocols are formalized using the notion of rights.

Ossowski and García-Serrano are concerned with characterizing the “social dependencies” that arise between agents that have possibly conflicting plans. The problem of conflict resolution facing such agents is formalized as a decision problem, and utility theory is used to determine which compromise plans the agents can use to achieve “globally rational” coordinated behaviour. This paper and a related paper by Bui, Venkatesh and Kieronska, also appearing in this volume (see below), address the important question of how utility theory can be generalized to a multi-agent setting.

C. Zhang and Li present an improvement to Katz and Rosenschein’s algorithm for the verification of multi-agent plans. In this work, a plan is represented as a directed acyclic graph whose nodes stand for actions and whose arcs stand for constraints on the order in which actions in the plan can be executed. Each graph, then, generates a set of possible execution orders. The problem of plan verification is to determine whether a particular plan achieves its goal, given only a STRIPS-like description of each action. The algorithm must take into account (in an efficient way) the possible conflicts arising from different execution orders. The algorithm presented here is superior in both time and space complexity to the original algorithm of Katz and Rosenschein.

Social Interaction

The papers in this section are all concerned with the properties of systems of relatively simple agents which may, however, interact in complex ways. This interaction is modelled using the approach of game theory. The authors all adopt an experimental methodology in evaluating their work.

Bui, Venkatesh and Kieronska present a framework in which coordination and learning of teams of decision-theoretic agents can be formalized. The paper is thus related to the work of Ossowski and García-Serrano reported in this volume (see above), although the present paper uses the definition of plan as strategy (from decision theory) rather than the notion of plan as sequence of actions (from Artificial Intelligence). The present paper also focusses on the important problem that agents face in estimating the utility of a team strategy in order to apply the technique, and the authors present experimental results based on a meeting scheduling scenario showing how this function may be learnt as the agents interact with one another over a period of time.

Carlsson and Johansson also discuss conflict resolution using a game theoretic framework. They present an experimental comparison of different strategies applied in an iterated setting to three games that have (formally) similar payoff matrices: the Prisoner’s dilemma, a chicken game and a hawk-and-dove game. The authors investigate the evolutionary stability of the strategies; a strategy is evolutionary stable if (roughly) once all members of a population adopt the strat-

egy, any mutations of the strategy are inferior and hence the strategy remains dominant in the population.

Yoshida, Inuzuka, Naing, Seki and Itoh also consider game theoretic models in a multi-agent setting, extending the model to one in which teams of agents compete with each other. Each member of a team adopts the same strategy, and contributes to the overall success of the team as measured by repeated iterations of the game between two randomly selected team members. This allows teams of agents with differing strategies to be compared.

Practical Issues for DAI Systems

The final section in this volume contains papers concerned broadly with issues in developing Distributed Artificial Intelligence systems.

C. Zhang and Luo present a translation from the EMYCIN model of uncertainty based on certainty factors, to Bayesian networks based on probability theory. This would enable existing systems using the EMYCIN model to be combined with systems using Bayesian networks. Thus both this paper and the paper by M. Zhang (see below) address issues of combining solutions in distributed expert systems.

D.M. Zhang, Alem and Yacef present a multi-agent framework for the design of instruction systems which can be used in conjunction with simulators to enable users to acquire skills in dynamic domains. The framework allows for separate agents responsible for simulating the domain, monitoring the user's progress to learn a model of the user, selecting new goals for instruction, and generating new learning scenarios, amongst other functions. They present an application of the model to developing a system for training air traffic controllers.

M. Zhang presents a case-based strategy for combining different solutions to problems in a distributed expert system. A heuristic for comparison of solutions is described and then evaluated experimentally on a sample data set.

Acknowledgements

The editors would like to thank the other organizers of the three workshops: Abhaya Nayak (Commonsense Reasoning), Lawrence Cavedon and Anand Rao (Intelligent Agents), and Dickson Lukose (Distributed Artificial Intelligence), for their efforts in arranging and coordinating the workshops. We would also like to thank the programme committee members and additional reviewers for each of the workshops, the organizers of the conference, and Griffith University (especially Abdul Sattar), the University of New South Wales (especially Norman Foo) and the University of Newcastle (especially Mary-Anne Williams) for funding the visits to Australia of the invited speakers. Finally, we would once again like to thank the invited speakers and all the workshop presenters and participants for contributing to the success of the three workshops.

Programme Committee (Commonsense Reasoning)

Grigoris Antoniou	Griffith University, Australia
Joseph Davis	University of Wollongong, Australia
Peter Eklund	University of Adelaide, Australia
Krishna Rao Madala	Griffith University, Australia
Pavlos Peppas	Macquarie University, Australia
Anand Rao	Australian Artificial Intelligence Institute, Australia
Mary-Anne Williams	University of Newcastle, Australia

Workshop Organizers

Abhaya Nayak	University of New South Wales, Australia
Maurice Pagnucco	University of New South Wales, Australia

Programme Committee (Intelligent Agents)

Hector Levesque	University of Toronto, Canada
Jörg Müller	Zuno Limited, UK
Hideyuki Nakashima	Electrotechnical Laboratory, Japan
Abdul Sattar	Griffith University, Australia
Munindar Singh	North Carolina State University, USA
Liz Sonenberg	University of Melbourne, Australia
Michael Wooldridge	Queen Mary and Westfield College, UK

Workshop Organizers

Lawrence Cavendon	Royal Melbourne Institute of Technology, Australia
Anand Rao	Australian Artificial Intelligence Institute, Australia
Wayne Wobcke	University of Sydney, Australia

Programme Committee (Distributed Artificial Intelligence)

Keith Decker	University of Delaware, USA
Yves Demazeau	IMAG-Leibniz Grenoble, France
Rose Dieng	INRIA Sophia Antipolis, France
Mark d'Inverno	University of Westminster, UK
Norbert Glaser	INRIA Sophia Antipolis, France
Victor Lesser	University of Massachusetts, USA
Michael Luck	University of Warwick, UK
Bernard Moulin	Laval University, Canada
Timothy Norman	Queen Mary and Westfield College, UK
Douglas Norrie	University of Calgary, Canada
Jeffrey Rosenschein	Hebrew University of Jerusalem, Israel
Toshiharu Sugawara	NTT Basic Research Laboratory, Japan
Minjie Zhang	Edith Cowan University, Australia

Workshop Organizers

Dickson Lukose	University of New England, Australia
Chengqi Zhang	University of New England, Australia