Foresight Report on IJCAI-03

Brandon Bennett
School of Computing, University of Leeds, Leeds LS2 9JT, UK
brandon@comp.leeds.ac.uk

Abstract

This document reports on the 18th International Joint Conference on Artificial Intelligence (IJCAI-03), held in Acapulco during 9th–15th August 2003.

1 Introduction

The Biennial IJCAI conference is the largest international conference on Artificial Intelligence. As such, its participants come from a wide spectrum of different sub-fields including many specialisms of traditional symbolic AI as well as fuzzy, statistical and bio-inspired approaches. IJCAI 2003 took place in Acapulco, a thriving tourist centre on the Pacific coast of Mexico. It was attended by over 1000 registered delegates from around the world.

The scale of the event makes it impossible for an individual to take in the full scope of scientific research presented at IJCAI. During most of the conference there was a choice of eight different technical papers as well as an invited talk. Hence the current report represents a personal and selective view of the conference.

2 Invited Talks

The conference incorporated an extensive programme of invited talks which gave overviews of some of the most rapidly developing fields in AI. This section describes a selection of these.

2.1 The Keynote Talk

Computer Vision: AI or non-AI problem, Takeo Kanade (Carnegie Mellon University)
The conference began with the keynote presentation given by professor Takeo Kanade. The talk sought to expose underlying problems that account for the difficulty of computer interpretation of visual scenes. In the inception of this field of study, the problem was considered as one of knowledge representation and manipulation. But it has subsequently become a specialised and largely separate area of research. It was noted that early successes relied upon heuristic algorithms that work very well for many cases but do not solve the problem in general. They require extensive *add hoc* tinkering to cope with unusual cases.

Prof Kanade went on to explain that researchers from different backgrounds tend to have very different ideas of the nature of this difficulty. For electrical engineers, the problem is normally seen as one of manipulating the massive amount of data that must be handled in order to process a visual image. Geometers point to the degeneracy of the transform from 3-dimensional reality to a 2-dimensional image, which means that the 3-dimensional geometry

is not recoverable from the 2-dimensional image. Physicists are concerned with further kinds of degeneracy due to the interaction of physical factors such as light intensity and reflexivity, whose parameters cannot be recovered from a 2-dimensional image.

Attention was drawn to the particular problem of dealing with reflective highlights, which cause serious problems for image interpretation. Kanade described an effective method of identifying highlights by physical modelling of reflectance. This suggests that a detailed physical model is required to implement certain aspects of visual processing. He also described his *hallucination* method, a statistical technique which enables one to derive precise images from low resolution input, based on the distribution of sample images over a parameterised space of possibilities. The effectiveness of this approach supports the view that visual processing has a significant statistical component.

Kanade noted that successes in computer vision primarily solve problems constructing geometrical world-models from image data, whereas the important problem of relating visual images to high-level scene descriptions is largely neglected. Moreover, it is clear that the capacity for scene description, coupled with an ability to utilise world knowledge, can greatly enhance the interpretation of visual scenes. A good example of this is that when we see cars on a road from a distance, we may not be able to tell their orientation from the visual image alone. But because we know that cars drive (in any given country) on one side of the road, we can often work out a car's orientation from its position on the road.

The talk was concluded with an assertion that visual recognition is in large part a knowledge problem after all; and, in order to tackle aspects of the problem that depend on high-level knowledge, computer vision and AI should remarry.

2.2 AI and the Web — Special Track

The invited talk schedule included a special track of three presentations on on uses and requirements for AI in the World Wide Web:

Mehran Sahami (Google and Stamford University) described current work at Google on analysing the content of web information in order to develop a more sophisticated search-engines that can better distinguish genuinely relevant web pages from those that are irrelevant or bogus.

Craig Knoblock (University of Southern California) talked about the use of intelligent software agents to automatically gather information from the web that is relevant to a user's specific needs. The proposed approach first constructs *wrappers* which provide semantically structured and agent-friendly interfaces to the information contained in web sites. Constructing this meta-information is done using a machine learning approach. Further learning is then used to establish linking information which enables the resources of several web sites to be combined. Once these interfaces have been constructed, software agents gather relevant information by means of automated planning techniques. A query is transformed into an executable plan which obtains relevant information through a sequence of web queries and data manipulation operations.

Jiming Liu (Kong Kong Baptist University and Web Intelligence Consortium) presented an ambitious programme for developing a futuristic *Wisdom Web* populated by intelligent and helpful software agents with highly sophisticated information processing capabilities (as well as having a sense of humour). Few details were given of exactly how this vision might be realised.

2.3 Other Invited Talks

Quantum Information: Fundamentals and Applications Anton Zeilinger (Vienna University)
Professor Zeilinger's talk described recent developments in the emerging field of quantum information theory. The fundamental concepts in the study of of quantum information are quantum superposition, entanglement and the objective randomness of the individual event. The presentation explained how these concepts enable the design of devices and computational systems with a variety of novel and powerful functionality. Applications include quantum cryptography, teleportation, quantum communication and quantum gates for future quantum computers.

Zeilinger reported successes in achieving instantaneous transmission of information based on the principle of quantum entanglement. This phenomenon means that operations performed on a quantum system in one location can have an immediate effect on the state of another system at a remote location. Details of experimental realisations of this possibility were given, where data was transmitted across the river Danube. However, a problem remains in that the transmitted data contains a certain ambiguity that can only be resolved by means of additional information that cannot (as yet) be transfered instantaneously.

Self-Reconfiguring Robots: Challenges and Successes (Daniela Rus, Dartmouth University)
This talk presented ongoing work on the construction of versatile robots by using self-reconfiguration. Such robots consist of hundreds of small modules autonomously organize and reorganise as geometric structures to best fit the terrain on which the robot has to move, the shape of the object the robot has to manipulate, or the sensing needs for the given task. Self-reconfiguration allows large collections of small robots to actively organise as the most optimal geometric structure to perform useful coordinated work. This leads to versatile robots that can support multiple modalities of locomotion and manipulation. Rus proposed that these kinds of robot constitute large scale distributed systems, which can be regarded as forming a network. The talk discusses many practical challenges that have been faced in designing and building working robots. Several short videos were presented which showed the systems in action using a variety of (sometimes bizarre) forms of locomotion.

Intelligent Information Integration (Mauizio Lenzerini, University of Rome)
This invited tutorial session examined the important problem of combining information residing at different sources and providing a user with a unified view of the whole body of information. The tutorial gave a powerful theoretical analysis of the problem and the different methods that can be used to solve it. A very general technique is to construct some kind of global knowledge description language (*schema*) which can be used to map between different sources. However, given such a global schema there are two opposite viewpoints that one can take on modelling the mapping to local data language. One is the global-as-view (GAV) approach, where one maps the concepts in the global schema to views over the sources. The other is the local-as-view (LAV) approach, where one maps the data sources into views over the global schema. Professor Lenzerini discussed the pros and cons of the two approaches and advocated a new approach (GLAV) which combines the advantages of both perspectives.

A number of other invited talks gave insightful overviews of a number of areas:

- *Intelligent Systems in Travel and Tourism, Hannes Werthner (University of Trento)*
- *Data Integration: Successes and Challenges, Along Halevy (University of Washington)*

- *Constraint Satisfaction, Databases and Logic, Phokion Kolaitis (University of California)*
- *Automated Verification = Graphs, Automata and Logic, Moshe Vardi (Rice University)*
- *Automatic Theorem Proving: past, present and future (Andrei Voronkov, Manchester University)*
- *User Interfaces: an AI Challenge, Daniel Weld (University of Washington)*

3 Technical Programme

3.1 Overview

To give an idea of the scope of the conference and its areas of particular focus, I list here the numbers of papers in each of the topics that were used to classify the proceedings:

AI and Data Integration 4	Non-monotonic Reasoning 8
AI and the Internet 3	Ontologies and Foundations 3
Art and Creativity 3	Perception 3
Automated Reasoning 3	Planning 6
Belief Revision and Update 3	Probabilistic Inference 4
Case-based Reasoning 3	Probabilistic Planning 4
Causality 3	Qualitative Reasoning 4
Cognitive Modelling 2	Reasoning about Actions and Change . 6
Cognitive Robotics 2	Resource-bounded Reasoning 3
Constraints 15	Robotics 6
Decision Theory 5	Satisfiability 3
Description Logics 7	Scheduling 3
Diagnosis 6	Search 5
Information Extraction 6	Spatial Reasoning 2
Knowledge Representation 5	Temporal Reasoning 3
Learning 19	User Modelling 2
Multi-agent Systems 26	Vision 2
Natural Language 4	

One must be careful in interpreting these figures, since many papers span several topics and some categories are much broader than others. In particular, one should note that the topics of symbolic knowledge representation and automated reasoning are sub-divided into several specialised domains and techniques, whereas the categories of *Learning* and *Multi-Agent Systems* each cover a spectrum of different approaches. In fact the *Multi-Agent Systems* papers were further classified into 10 sub-fields. But it is clear that, the large numbers of papers on *Learning* and *Multi-Agent Systems* correspond to a real growth in interest since the previous IJCAI, especially in the latter area. It is also evident that a large part of current work in automated reasoning for AI systems concentrates on the approaches of *Constraints* and *Description Logics*.

In my review of the technical papers I use a much coarser division of the papers into subject and pick out a selection of representative papers from each of these. The choice of papers considered is somewhat biased towards covering the talks I actually attended but I have included some brief descriptions of other papers in the proceedings which report work in particularly original research directions.

3.2 Knowledge Representation and Reasoning

Knowledge representation has always been a core area of AI. However, as the field has progressed, it has been increasingly divided into the study of specialised domains. Some separation has also arisen between the development of knowledge representation formalisms and the study of computational inference techniques. However, in approaches such as constraint satisfaction and description logics, capabilities of representation and inference have in some cases been brought together to yield powerful and practical reasoning systems.

3.2.1 Representation and Ontology

Several papers dealt with fundamental issues in the development of knowledge representation formalisms.

Tucking RCC in CYC's Ontological Bed, Pierre Grenon (University of Leipzig, IFOMIS)
This paper examined fundamental issues that arise when one attempts to incorporate a formal axiomatic theory of some particular domain into a large general-purpose ontology. His investigation focused on the particular case of integrating the spatial formalism known as the *Region Connection Calculus* into the huge CYC ontology. This example illustrates how achieving integration may lead either to nontrivial changes in the embedding ontology or to the choice of a peculiar interpretation for the theory to be integrated.

Layered Mereotopology, M. Donnelly (University of Leipzig, IFOMIS)
This paper gives a deep ontological analysis of the types of spatial relations that can hold between entities that are of radically different kinds. This general *mereology* (theory of 'parthood' relations) is applicable to many practical domains, such as geography and medicine. In biology, for instance we may be concerned with the relationship of material entities (bodily organs, tumours, food etc.) to various cavities within the human body. A particular case might be a gall-stone located within the bladder. Here we have a material object whose spatial extension lies within a cavity; and yet we would not want to consider the gall-stone as a 'part' of the bladder cavity. Donnelly's paper develops a theory of spatial relationship, which distinguishes mere spatial coextensionality from the actual sharing of a 'part'.

3.2.2 Automated Reasoning

The problem of reasoning with arbitrary 1st-order theories has been studied in symbolic AI since its inception but still remains one of the most unyielding problems.

Practical Partition-Based Theorem Proving for Large Knowledge Basis, B. MacCartney, S. McIlraith (Stanford University), E. Amir (Berkeley), T. Uribe (Melno Park)
This was the only paper that tackled reasoning with unrestricted 1st-order languages. It employs a *partition-based* approach, where a theory is divided into several more tractable parts. The bulk of reasoning is done within these partitions but this is supplemented by message passing between the partitions. Efficiency is further improved by applying sophisticated order-restrictions to the resolution inference mechanism.

3.2.3 Constraint Satisfaction

Constraint satisfaction is a very active area of research (with 15 technical papers on constraints accepted for IJCAI-03). The problems considered are typically decidable but intractable in the general case. Hence, a major research topic is the development of heuristic algorithms which can dramatically improve average-case performance of constraint consistency checking. Several papers take forward preference-based techniques for guiding search. Others look at the use of symmetry analysis to reduce the size of the search space.

A number of papers looked at the decomposition of a CSP into a number of simpler interacting problems, a technique that is well established in 1st-order reasoning, but more novel within the CSP framework. Another line of research that seems to be increasingly studied is the processing of *soft* constraints — i.e. conditions which are desirable but need not necessarily hold.

3.2.4 Description Logics

After constraint-based formulations, the next most popular (7 papers in IJCAI-03) formalism for implementing reasoning systems is Description Logics (DL). The basic idea of this approach is to use a kind of restricted 1st-order logic, which is both decidable and also well suited to representing conceptual knowledge.

Most of the papers concentrate on the computational properties of various extensions of the basic DL system and on optimising inference in these systems. One paper (by S. Schlobach and R. Cornet) looked at supporting debugging of DL systems, which can be a serious problem in large DL knowledge-bases. The focus in DL research is very much on the formalisms themselves rather than the actual information manipulated with these systems. However, one paper (actually in the Spatial Reasoning section of the proceedings) applied a modified DL framework to reasoning about a concrete domain of spatial relationships.

Reasoning about distances, Frank Wolter (University of Liverpool) and Michael Zakharyaschev (King's College London)
The paper presents a novel and expressive logic-based formalism intended for reasoning about numerical distance relationships. Computational properties were investigated and gives a tableau-based satisfiability-checking algorithm. The reasoning problem is shown to be EXPTIME-complete.

3.2.5 Probabilistic Inference

Probabilistic accounts of inference (such as *Fuzzy Logic*) have been a major strand of AI for several decades. However, this work tends to be reported at more specialised conferences, and only 4 papers on probabilistic inference were presented at IJCAI-03. An innovation that seems to be feeding into much recent work in the area is the introduction of *relational* probabilities, as can be achieved within a Probabilistic Relational Model (PRM).

Dynamic Probabilistic Relational Models, S. Sanghai, P. Domingos, D. Weld (University of Washington)
This paper proposes the use of PRM's as state descriptions representing time slices of dynamical processes modelled by means of dynamic Bayesian networks. It develops an inference procedure using a refinement of the well-known *particle filtering* method.

3.2.6 Meta-Logical Properties

Evaluating Significance of Inconsistencies, Anthony Hunter (University College London)
By means *quasi-classical* logic (a form of *paraconsistent* logic) the paper analyses the significance different types of inconsistency that may occur in logical knowledge bases. The formalism is applied to the problem of managing heterogeneous sources of knowledge.

3.2.7 Planning

Because of its clear utility and relative efficiency compared to more general modes of inference, planning is a core technique of symbolic AI. However, when applied to complex domains, planning by means of naive planning algorithms soon becomes intractable. Hence there is considerable work on enhancing planning algorithms.

As with other modes of reasoning, the division of a planning problem into a collection of smaller, loosely interacting problems has been found to be very effective. In this vein, the paper *Factored Planning, E. Amir and B. Engelhardt (Berkley)* looks at automatically factoring planning problems into components. Another approach (taken in the paper by *G. Armano, G. Cherchi and E. Vargiu*) is to parameterise plans in such a way as to achieve generality while keeping effect axioms simple.

Incorporating probabilistic models into planning algorithms has led to a growing field of Probabilistic Planning. The following paper describes a powerful combination of probabilistic planning and generalisation, which can also be applied to multi-agent situations.

Generalising Plans to New Environments in Relational MDPs (Carlos Guestrin, Daphne Koller, Chris Gearhart and Neal Kanodia, Stanford University)
The authors suggest that generalisation can greatly increase planning efficiency. Thus their goal is to provide a mechanism by which plans developed for a particular environment can be transfered to new but similar environments, with minimal or no re-planning. The approach is based on a new framework of Relational Markov Decision Processes (RMDPs). These model a set of similar environments by representing objects as instances of different classes. Plans are optimised in terms of a class-based function. This is evaluated by sampling the space of possible instantiations of the classes in order to produce a function that approximates to an optimal evaluation over the whole space of possibilities. The viability of the approach is demonstrated by application to a real strategic computer war game.

3.2.8 Other Reasoning Techniques

Work on many other approaches to reasoning was also presented. These papers presented extensions and refinements of many well established techniques. In particular there were several papers on non-monotonic reasoning using logic programming or default logic approaches.

3.3 Multi-Agent Systems

A large number of AJCAI papers come under the Multi-Agent Systems heading. But this area covers many more specific research topics and a wide variety of different approaches,

Approximating Game-Theoretic Optimal Strategies for Full-scale Poker (D. Billings, N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and D. Szafron, University of Alberta)

This work in the area of multi-agent game playing received the award for the best technical paper of IJCAI-03. The research gives a method of computing strategies for playing the full game of *Texas Hold'em* poker, which approximate game-theoretically optimal strategies. The approach uses a combination of abstraction techniques which enable the vast space of possibilities which arise in the game to be reduced to a far more manageable model which still captures the essential structure of the game. Linear programming solutions to the abstracted game are then used to create substantially improved poker-playing programs, able to defeat strong human players and be competitive against world-class opponents.

Constitutive Rules for Agent Communication Languages, Jeremy Pitt (Imperial College London)
This paper examines the communicative capacity of agent languages in terms of the *Speech Act* theory of Searle. He argues that communication between software agents in terms of 'method calls' (*via* a standard API) corresponds to essentially *perlocutory* speech acts, where the meaning of the act is completely reducible to its effect on the receiving agent. But by using a more sophisticated "Agent Communication Language," agents may also communicate propositional content and *illocutionary force* (e.g. requesting, asserting, promising).

3.4 Learning

The many papers on Machine Learning cover a wide spectrum of different approaches an application domains. These range from applying inductive logic programming (ILP) to the game of minesweeper to performing text classification using Gaussian weighting. The primary tasks studied are on the one hand *classification* and on the other *rule learning*. In the former case, statistical analysis is the primary tool, whereas the latter case symbolic approaches such as ILP are also employed. Learnt classifiers take various forms such as Bayesian nets and decision tree as well as statistical clusterings.

3.5 AI and the Internet

While 4 out of the 13 invited talks dealt with intelligent applications on the internet, this emphasis was not reflected in the technical programme, with only 3 out of the 189 accepted papers being specifically focused on internet applications. But this is largely due to the level of theoretical abstraction of IJCAI papers, which typically analyse information processing problems in a rather general setting. Hence, it would be a mistake to conclude that AI is not relevant to or concerned with the internet. Indeed many papers directly address internet applications, especially those in the areas of information extraction, learning, and multi-agent systems

3.6 Foundations of Artificial Intelligence

IJCAI papers are very much concerned with the technical details of achieving intelligent behaviour rather than philosophical questions on the nature of intelligenc. However, there was one paper that examined the fundamental question of what is AI.
What is Artificial Intelligence: Psychometric AI as an Answer, Selmer Bringsjord and Bettina Schimanski (Rensselaer Polytechnic Institute)

This paper proposed that psychometric testing be used a basis to evaluate AI systems. In other words the intelligence of computers should be tested in much the same way as we do

with humans, by requiring them to solve a series of test problems stated in natural language text. Clearly the ability to solve such problems involves a wide range of intelligent capabilities, including natural language understanding, world knowledge and reasoning.

4 Other Conference Events

The IJCAI exhibition included stalls from publishers, AI software developers and robot manufacturers. Several small robots could be seen meandering around the hall in an erratic but thankfully non-threatening manner.

Dominating one corner of the hall was a large installation of panels and scaffolding that enclosed an area strewn with waste paper, junk and contorted manikins. This was the venue for the *Robot Rescue* competition. The objective of this contest was to use AI controlled robots to search the environment, identify human casualties and direct rescuers to their location. Some robots were very effective at this task, while others were amusingly incompetent.

Another lively event was the *Trading Agent* competition, in which software agents trade (virtual) commodities *via* an electronic brokering system.

5 Future Research Directions

Given the wide-open nature of the field, it is perhaps unwise to attempt to predict the future of AI research. However, to round off this report, some speculative final remarks are in order.

The current state of evolution of AI research might be characterised as one of *consolidation without unification*. That is, there are many sub-fields each of which has build up a substantial body of foundational work and is in the process of refining and systematising this material. Within these areas, developments are often highly technical, which means there is difficulty in transferring results and innovations between different AI sub-fields. Such specialised work undoubtedly contributes to the enormous long-term project of realising the aspiration of building a system having human-like intelligence. However, it seems unlikely that any large breakthrough will come directly from any of the existing well-established sub-fields.

On the other hand one might imagine that a vehicle for dramatic progress might come from outside the traditional framework of AI — for example through a breakthrough on quantum computing (as suggested by Anton Zeilinger in his invited talk). However, though such a development might well provide powerful algorithmic tools for AI, it seems less likely that it would lead directly to a corresponding breakthrough in the understanding of intelligence itself.

In my opinion the best opportunities for new AI research lie in the integration and combination of the disparate sub-fields of the discipline. Already we are seeing pioneering work that incorporates methods spanning several hitherto separate approaches within the field. In other areas a need for integration is widely acknowledged, but the way to achieve it not yet clear (for instance the use of high-level knowledge representation in computer vision — as advocated in the keynote talk of Takeo Kanade). Such interlinking will certainly strengthen AI as a whole and may eventually lead to a new generation of more powerful and flexible AI software systems.

References

All the papers referred to in this report can be found in the IJCAI-03 proceedings:

Gottlob, G. and Walsh, T. (eds): 2003, *Proceedings of the 8th International Joint Conference on Artificial Intelligence (Acapulco, Mexico)*, Morgan Kaufmann, San Francisco.