# 2001 13th IEEE International Conference on Tools With Artificial Intelligence

# Proceedings

# 13th IEEE International Conference on Tools with Artificial Intelligence

## ICTAI 2001

November 7–9, 2001

Dallas, Texas

*Sponsored by*

IEEE Computer Society
IEEE CS Virtual Intelligence Task Force

*In cooperation with*

The Information Technology Research Institute
Wright State University
AIIS Inc.



IEEE COMPUTER SOCIETY

Los Alamitos, California
Washington · Brussels · Tokyo

IEEE Computer Society Order Number PRO1423
ISBN 0-7695-14~~23-5~~ 170
Library of Congress Control Number 2001096085

*Additional copies may be ordered from:*

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1 714 821 8380
Fax: + 1 714 821 4641
http://computer.org/
csbooks@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1 732 981 0060
Fax: + 1 732 981 9667
http://shop.ieee.org/store/
customer-service@ieee.org

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81 3 3408 3118
Fax: + 81 3 3408 3553
tokyo.ofc@computer.org

Editorial production by Randall Bilof and Lorretta Palagi
Printed in the United States of America by The Printing House

IEEE
COMPUTER
SOCIETY

# Message from the General Chair

## *Welcome to the ICTAI 2001 Meeting*

It is a great pleasure and honor for all of the TAI 2001 organizers and committee members to welcome you to the Thirteenth IEEE International Conference on *Tools with Artificial Intelligence 2001*.

This conference is unique since it involves the development and integration of emerging AI theories with practice and application. This theme will make the conference a success story.

With the support of active international researchers like you, the theme of the TAI meetings has been kept at the forefront of AI technology. We hope that this year's conference activities (including keynote speeches, invited sessions and contributed papers) will generate the same synergy as in the past, and help you generate new ideas and projects.

The success of TAI 2001 was made possible by the hard work and voluntary contributions of many international researchers and volunteers who served in various capacities including reviewers and the members of the Steering Committee, the Program Committee, the Local Arrangement Committee, and other technical committees. Special thanks go to C. V. Ramamoorthy, N. Bourbakis, J. Tsai, B. Wah, D. Moldovan, E. Pontelli, L. Miller, M. K. Rada, A. M. Kelly, R. Bilof, S. Harabagiu, and R. Girju for their dedicated help.

We would also like to thank our keynote speakers and the invited presenters for their participation.

Finally, this meeting would have not been possible without the sponsorship of the IEEE Computer Society and the support from the IEEE Computer Society staff.

**Arvind K. Bansal**
General Chair
*Kent State University*

# Message from the Program Chair

This volume contains the papers accepted for the Thirteenth IEEE International Conference on Tools with Artificial Intelligence. The Program Committee selected 44 papers out of 86 submitted. The only criterion on which the papers were selected was quality of research reported.

The papers cover a broad range of AI topics: multi-agents and distributed AI (7), reasoning (7), planning (3), machine learning (6), data mining (3), natural language processing (6), AI in software engineering (3), genetic algorithms (3), constraint satisfaction (3), and applications (3).

Many computer systems incorporate AI techniques, thus becoming more and more intelligent. We see multi-agent systems able to reason and function in real-world applications, natural language processing techniques that help curb the information explosion, data mining methods that extract useful information from large databases, and other tools that use AI techniques. As companies struggle to become more efficient and competitive, they look toward AI as a mean to build smarter computer systems.

This is truly an international conference as papers came from all over the world. Four distinguished speakers have agreed to give plenary talks. The high quality of papers coupled with the broad international participation and ample time for presentations and discussions are likely to make this event an exciting conference.

The Program Committee members have worked under a tight schedule to review and select the best papers and they deserve our gratitude. I would also like to thank all the TAI committees and the IEEE Computer Society staff that made this proceedings and conference possible. I also wish to acknowledge the hard work and diligence of Roxana Girju who has assisted me with the conference organization.

**Dan Moldovan**
ICTAI 2001 Program Chair
*University of Texas at Dallas*

# Committees

## General Chair

Arvind Bansal, *Kent State University*

## Steering Committee

Nikolaos Bourbakis, *Wright State University*
C. V. Ramamoorthy, *University of California at Berkeley*
Jefferey Tsai, *University of Illinois at Chicago*
Benjamin Wah, *University of Illinois at Urbana–Champaign*

## Program Chair

Dan Moldovan, *Southern Methodist University*

## Program Committee

Farokh Bastani, *University of Texas at Dallas, USA*
Diane Cook, *University of Texas at Arlington, USA*
Hamish Cunningham, *University of Sheffield, UK*
Roxana Girju, *Southern Methodist University, USA*
Forouzan Golshani, *Arizona State University, USA*
Fernando Gomez, *University of Central Florida, USA*
Gopal Gupta, *University of Texas at Dallas, USA*
Babak Hamidzadeh, *University of British Columbia, Canada*
Charles Hannon, *University of Texas at Arlington, USA*
Sanda Harabagiu, *Southern Methodist University, USA*
Rammohana Rao Kotagiri, *University of Melbourne, Australia*
Ioan Alfred Letia, *Technical University of Cluj-Napoca, Romania*
Seven Maiorano, *USA*
Bill Manaris, *College of Charleston, USA*
Teruhisa Miura, *Central Research Institute of Electric Power Industry, Japan*
Tudor Muresan, *Technical University of Cluj-Napoca, Romania*
Ion Alexandru Muslea, *USC/ISI, USA*
Marin Simina, *Tulane University, USA*
Vincent Tam, *The University of Hong Kong, Hong Kong*
Gheorghe Tecuci, *George Mason University, USA*
Klaus Truemper, *University of Texas at Dallas, USA*
Haleh Vafaie, *Logicon Federal Data, USA*
Takashi Yukawa, *NTT, Japan*
Wlodek Zadrozny, *IBM, USA*

## Publicity Chair

Enrico Pontelli, *New Mexico State University*

## Registration Chair

Leona Miller, *Wright State University*

## Wednesday, November 7

| Time slot | Section 1 | Section 2 |
|---|---|---|
| 08:45-9:00 | Welcome | |
| 9:00-10:00 | Invited speaker: *Daniel Cooke*, NASA: "AI for Space Explorations in the Future" | |
| 10:00-10:30 | Coffee Break | |
| | **Multi-Agents** | |
| 10:30-11:00 | Towards Ontological Reconciliation for Agents<br>*L. Sterling* and *K. Lister* | |
| 11:00-11:30 | Transaction Oriented Computational Models for Multi-Agent Systems<br>*K. Ramamohanarao, J. Bailey* and *P. Busetta* | |
| 11:30-12:00 | Resource Coordination in Single Agent And Multiagent Systems<br>*G. Edwin* and *M.T. Cox* | |
| 12:00-12:30 | Smart Cars as Autonomous Agents<br>*N. Bourbakis* and *M. Findler* | |
| 12:30:14:00 | Lunch | |
| | **Distributed Artificial Intelligence** | **Genetic Algorithms** |
| 14:00-14:30 | Interleaved Backtracking in Distributed Constraint Networks<br>*Youssef Hamadi* | Dynamic Load-Balancing by a Genetic Algorithm<br>*William A Greene* |
| 14:30-15:00 | A Multi -Agent System to the Common Management of a Renewable Resource: Application to Water Sharing<br>*Marjorie Le Bars* and *Jean-Marie Attonaty* | Parallel Genetic Programming for Decision Tree Induction<br>*Gianluigi Folino, Clara Pizzuti* and *Giandomenico Spezzano* |
| 15:00-15:30 | Using Software Agents to avoid Collisions Among Multiple Robots<br>*Markus Jaeger* | New Hybrid Genetic Algorithms for the Frequency Assignment Problem<br>*Miguel Alabau, Lhassane Idoumghar* and *Rene Schott* |
| 15:30-16:00 | Coffee Break | |
| | **Reasoning 1** | **Applications** |
| 16:00:16:30 | RaBeCa: A Hybrid Case-Based Reasoning Development Environment<br>*Ricardo Barz Sovat, S. M. Aluísio* and *Andre de Carvalho* | Maintaining Credible Dialogs in a VideoBot System: Special Audio Techniques<br>*Doug DeGroot* |
| 16:30:17:00 | Automatic Knowledge Acquisition from Subject Matter Experts<br>*Mihai Boicu, Gheorghe Tecuci, Bogdan Stanescu, Dorin Marcu,* and *Cristina Cascaval* | Maximizing Paper Spread in Examination Timetabling Using a Vehicle Routing Method<br>*Wee-Kit Ho, Andrew Lim* and *Wee-Chong Oon* |
| 17:00-17:30 | Inconsistent Requirements: An Argumentation View<br>*Laurent Perrussel* and *Pierre-Jean Charrel* | Artificial Neural Networks in Hydrological Watershed Modeling: Surface Flow Contributions from Ungauged Parts of a Catchment<br>*Richard Chibanga, J. Berlamont,* and *J. Vandewalle* |

## Thursday, November 8

| Time slot | Section 1 | Section 2 |
|---|---|---|
| 9:00-10:00 | Invited speaker: *Hiroshi Yamaguchi*, NEC: TBA | |
| 10:00-10:30 | Coffee Break | |
| | **Reasoning 2** | |
| 10:30-11:00 | Perfect Sequences for Belief Networks Representation<br>*Radim Jirousek* and *Jirina Vejnarova* | |
| 11:00-11:30 | A Decentralized Model-Based Diagnostic Tool for Complex Systems<br>*Yannick Pencole, Marie-Odile Cordier* and *Laurence Roze* | |
| 11:30-12:00 | Alama/Carne: Implementation of a Time-Situated Meta-Reasoner<br>*K. Purang* | |
| 12:00-12:30 | Towards a General Theory for Qualitative Space<br>*B. A. El-Geresy* and *Alia Abdelmonty* | |
| 12:30:14:00 | Lunch | |
| | **Natural Language Processing 1** | **Planning** |
| 14:00-14:30 | B-Course: A Web Service for Bayesian Data Analysis<br>*Petri Myllymaki, Tomi Silander, Henry Tirri* and *Pekka Uronen* | Fast Suboptimal Planning with Nexus States<br>*Will Briggs* and *Brad Dawson* |
| 14:30-15:00 | An Expert Recommendation System using Concept-based Relevance Discernment<br>*Takashi Yukawa, Kaname Kasahara, Tsuneaki Kato* and *Toshiro Kita* | An Architecture for Planning in Uncertain Domains<br>*Mario E. Asueda* and *Pablo H. Ibarguengoytia* |
| 15:00-15:30 | A Robust Model for Intelligent Text Classification<br>*Roberto Basili* and *Alessandro Moschitti* | A Metaheuristic for the Pickup and Delivery Problem with Time Windows<br>*Haibing Li* and *Andrew Lim* |
| 15:30-16:00 | Coffee Break | |
| | **Software Engineering** | **Data Mining** |
| 16:00-16:30 | LOOP - A Language for LP-based AI Applications<br>*Alin Suciu, Kalman Pusztai, Tudor Muresan* and *Zslot Simon* | Mining First-order Knowledge Bases for Association Rules<br>*Hasan Jamil* |
| 16:30:17:00 | SHAMASH: An AI Tool for Modeling and Optimizing Business Processes<br>*David Camacho, Ricardo Aler, Daniel Borrajo* and *Almudena Sierra-Alonso* | Visualization Support for User-Centered Model Selection in Knowledge Discovery in Databases<br>*Ho Tu Bao* and *Nguyen Trong Dung* |
| 17:00-17:30 | An On-Line Software Repository for Embedded Systems<br>*I-Ling Yen, Latifur Khan, R Prabhakaran, Faroukh Bastani* and *John Linn* | Generation of Propagation Rules for Intentionally Defined Constraints<br>*Slim Abdennadher* and *Christophe Rigotti* |
| 19:00-21:00 | Banquet (Talk: *Doug DeGroot*, Leiden University, Netherlands – "Virtual Immortality and the role of Agent Technology") | |

## Friday, November 9

| Time slot | Section 1 | Section 2 |
|---|---|---|
| 9:00-10:00 | Invited speaker: *Raymond Mooney*, UT. Austin: "Machine Learning for Recommender Systems" | |
| 10:00-10:30 | Coffee Break | |
| | **Natural Language Processing 2** | |
| 10:30-11:00 | CorefDraw - A Tool for Annotation and Visualization of Coreference Data<br>*Sanda Harabagiu* , *Razvan Bunescu* and *Stefan Trausan-Matu* | |
| 11:00-11:30 | Word Semantics for Information Retrieval: Moving One Step Closer to the Semantic Web<br>*Rada Mihalcea and S.I. Mihalcea* | |
| 11:30-12:00 | High Precision Logic Form Transformation<br>*Vasile Rus* | |
| | | |
| 12:00-13:30 | Lunch | |
| | **Machine Learning 1** | **Constraint Satisfaction** |
| 13:30-14:00 | An Effective Method for Generating Multiple Linear Regression Rules from Artificial Neural Networks<br>*Rudy Setiono* and *Arnulfo Palacio Azcarraga* | A New Hybrid Method for Solving Constraint Optimization Problems in Anytime Contexts<br>*Samir Loudni* and *Patrice Boizumault* |
| 14:00-14:30 | Data Flow Coherence Criteria in ILP Tools<br>*Smaranda Muresan, Tudor Muresan* and *Rodica Potolea* | Pickup and Delivery with Time Windows: Algorithms and Test Case Generation<br>*H. C. Lau* and *Z. Liang* |
| 14:30-15:00 | Comparing Keyword Extraction Techniques for WEBSOM Text Archives<br>*Arnulfo Palacio Azcarraga* and *Teddy Yap, Jr.* | Successive Search Method for Valued Constraint Satisfaction and Optimization Problems<br>*Mohamed Tounsi* and *Philippe David* |
| 15:00-15:30 | Coffee Break | |
| | **Machine Learning 2** | |
| 15:30-16:00 | Developing Colaborative Golog Agents by Reinforcement Learning<br>*Ioan Alfred Letia* and *Doina Precup* | |
| 16:00-16:30 | Combinatorial Optimization through Statistical Instance-Based Learning<br>*Orestis Telelis* and *Panagiotis Stamatopoulos* | |
| 16:30:17:00 | Empirical Study of a Stacking State-Space<br>*Agapito Ledezma, Ricardo Aler* and *Daniel Borrajo* | |

# Table of Contents

## 13th IEEE International Conference on Tools with Artificial Intelligence—ICTAI 2001

## Session 6: Natural Language Processing and Information Retrieval



## Session 7: Software Engineering



## Session 8: Constraint Satisfaction

## Session 9: Applications

# Towards Ontological Reconciliation for Agents

Leon Sterling & Kendall Lister
Intelligent Agent Laboratory:
http://www.cs.mu.oz.au/agentlab
Department of Computer Science and Software Engineering
The University of Melbourne, Victoria, 3010, Australia
leon,krl@cs.mu.oz.au

## Abstract

*This paper addresses issues faced by agents operating in large-scale multi-cultural environments. We argue for systems that are tolerant of heterogeneity. The discussion is illustrated with a running example of researching and comparing university web sites, which is a realistic scenario representative of many current knowledge management tasks that would benefit from agent assistance. We discuss efforts of the Intelligent Agent Laboratory toward designing such tolerant systems, giving a detailed presentation of the results of several implementations.*

## 1 Introduction

Useful knowledge systems inevitably incorporate vast amounts of information. The evolution of the computer as a data processing device, and computer networks as communication media, has provided the technical means to aggregate enormous quantities of information. Our capacity for accumulation, storage and reproduction of data and information has out-paced our ability to perceive and manipulate knowledge. These three observations are not new – Vannevar Bush identified just such a glut of knowledge and information over fifty years ago. He proposed a technological solution in the form of the memex, an enlarged intimate supplement to memory that anticipated the hypertext systems of today [3].

The development of a (pseudo-)global communication infrastructure that provides means for the publication, comparison and aggregation of apparently limitless amounts of data, i.e. the WWW, has changed the way we can manipulate information. We have created a potential to ask questions as individuals conducting our daily lives that previously would have been dismissed as infeasible unless one had the resources of a dedicated organisation. For example, with the entry cost of publishing a web site effectively negligible, the university that does not do so is the exception rather than the rule. Consequently, thousands of descriptions of courses, programs and facilities are available for us to peruse. There immediately arises the need for comparison. It is natural to ask reasonable and seemingly simple questions such as "Which faculties offer courses in applied machine vision?" or "Which campuses provide accommodation facilities for post-graduate students?".

To answer questions like these, we could fairly easily compile a list of university web sites. Each site could be visited in turn. Through browsing or searching we could uncover the information relevant to answer our question. We could then compare the results of our research from each site to formulate an answer. Many people perform this very task every day.

The question that interests this paper is why our computers can't do this for us yet. The followup question is how to approach the issue of enabling our computers to search and filter information to ask specific questions. While search engines are becoming increasingly powerful, and answer some specific questions well, they do not have anywhere near the generality needed.

The example of university service descriptions is useful to study. Finding information from university sites is a real problem. Universities as institutions tend naturally to develop and often then actively promote their individuality. Their local culture flavours their presentation of information that must then be reconciled with information from other institutions that apply their own cultural characteristics to their publications. If we are to manage knowledge from a variety of sources effectively, we will need the assistance of software that is culturally aware and is capable of negotiating the conflicts that arise when such heterogeneous knowledge is juxtaposed.

## 2 Organisational Culture & Communication

Today's reality is that knowledge from large numbers of heterogeneous sources must be integrated in such a way that any differences in representation and context can be effectively reconciled. The ability to work with knowledge from incongruous sources is becoming increasingly necessary [15] as the focus of information processing moves beyond intra-organisational interaction and begins to transgress borders, whether departmental, corporate, academic or ethnic. Organisations, whether companies, universities, industries, or nations, develop unique cultures as they grow. Organisational culture is considered to be both constructive and inhibitive as far as the day to day operations of the organisation are concerned. In the context of knowledge management organisational culture creates significant barriers to inter-organisational communications and transactions.

Organisational cultures arise as individual organisations develop mechanisms, procedures and representations for dealing with the issues that they face. Inevitably, because these cultures are generally developed in isolation, each organisation arrives at different solutions to what are often very similar problems. To work effectively in an organisation, individuals often must disregard their personal approach to a situation in lieu of an agreed common understanding shared by the other members of the group. We do this naturally when we work together on a problem. Teamwork and the ability to understand another's point of view are recognised as desirable qualities. Such qualities are also becoming desirable in software as agents play an increasing role in our communication and collaboration.

There are also disadvantages in requiring every member of an organisation to follow a centralised doctrine. Standardisation can result in inflexibility and a reduced ability to adapt and cope with a wide variety of situations effectively. Lack of flexibility is exacerbated when organisations attempt to interact with external groups. People inevitably find that even when they think that they are working in similar domains and facing common problems, they are unable to communicate effectively. Cultural differences between organisations, differences that arise as each organisation attempts to codify its individual approach to various situations, create impediments. The streamlining that appeared necessary for efficient operation within each organisation now stands as a barrier to interoperation and sharing of resources. Exactly such an incongruity also manifests in the information and knowledge generated by organisations. The problems faced by software agents negotiating such data are analogous.

When we suppress our own intuitive understanding of a situation and attempt to adopt a standardised, agreed upon approach, we increase our ability to interact with others who have similarly adapted their individual understanding to that of the group or community. But we also lose something in the process: context and generality. An efficient understanding of a situation is like a model. The more closely it describes a particular situation, the less effectively it describes a general class of situations. As we move from a general conceptualisation of a situation rich with semantic flexibility to a specific understanding, we tend to eschew context. The very generality that gives us the ability to deal with many varied and new situations is a barrier to communication. At the same time that ambiguity allows adaptation, it prohibits individuals from establishing the certainty of agreement that is necessary for confidence that each understands the other. Standardisation of practices and understandings does not create a panacea for the difficulties of communication and collaboration, as organisations discover. On a small scale, adoption of standardised approaches helps individuals to cooperate and achieve goals too large for a single person. On a larger scale, the effort required to establish and prescribe global standards and common approaches grows rapidly beyond feasibility as the number of participants and the amount of data being manipulated increases. As our ability to communicate and interact across cultural borders increases, so does our desire to do so. If our software tools are to scale, they must be provided with reconciliation capabilities.

## 3 Our Software Colleagues

Computers can be viewed as an extreme example of co-workers with poor teamwork and communication skills. When specifying a task for a software application or agent, we must specify every step in precise detail, detail that will generally remain constant throughout the life of the software. Humans are able to adjust the level of abstraction at which they conceptualise a particular situation. Computers by contrast have the capacity only for comparatively very low levels of abstraction. As machines that follow explicit instructions to the letter, their operation is analogous to the most procedural organisational standards, and unsurprisingly they adapt to new situations with great difficulty.

Traditional computational paradigms require that computer-mediated representations of information and knowledge be exact and literal; for a computer to process information requires simplistic structuring of data and homogeneous representations of concepts. In order to maintain consistency during processing, traditional approaches require that each participant in a system, whether human or software, subscribes to a common understanding of the concepts within the system. In other words, traditional knowledge systems require the adoption of an absolute ontological world-view; deviation from a priori agreed terms and understandings results in a breakdown in communication and loss of consistency through the system.

Ontological homogeneity has worked well for systems with little direct human interaction, when the computers can be left to sort out technical details and humans can work at a level removed. Isolating the technical details of a system from those areas with which humans interact permits engineering of the technical aspects to create an optimised environment. The World Wide Web is an example of a large-scale system where the level at which humans interact with the system is quite separate from the level at which machines interact with each other. We write web pages and read them by navigating along hypertextual paths, while machines manage domain name resolution, protocol selection, transmission of data and rendering of text and images.

The gap between the activities of humans and machines is highlighted by the problems that occur when we try to make machines work closer to our level as we attempt to automate various functions that we currently perform manually. The example of this most recognisable to the ordinary web user is searching for information, an obviously difficult problem that has yet to be solved to our satisfaction. But a more far-reaching problem is that of integrating the vast quantities of information available in such a way that we can seamlessly assimilate whatever sources of data are most appropriate to the task at hand, whatever that task may be.

## 4 Reconciling Conceptualisations

The ability to manipulate concepts at varying levels of detail and to match the level of detail to the needs of the situation at hand is an effective tool for processing knowledge and communicating. Being able to subsume detail within conceptual units of knowledge allows us to overcome the natural limits of our processing capacity. Although there appear to be cognitive limits on the number of concepts we can articulate at any given time, we have the critical ability to 'chunk' collections of knowledge into single units [11, 5], effectively providing a capacity to search through information webs both widely and deeply as necessary.

When the scope of an information or data handling task becomes too great for us to process in a reasonable amount of time, we conscript computers to assist us with storage, recall and simple processing. By handing low-level information processing tasks to machines, humans are freed to consider issues at higher levels of abstraction. To continue to increase the assistance provided by computers as we work, our tools must be elevated to higher levels of abstraction.

As knowledge travels through progressively lower levels of abstraction, its context degrades as generality is replaced by specificity and logical operability. Humans require some specificity in order to communicate successfully. The desired degree of consistency of conceptualisations determines the extent of specificity that is necessary. It has been suggested that consensus between participants is not always necessary for successful collaboration [1, 12].

Humans are capable of identifying mismatches of understanding in our communications and negotiating shared perspectives as we interact with others [2]. Human natural language is neither precise nor predictable, and this seems to reflect the way that we understand the world through our internal representations and conceptualisations. When we express ourselves in natural language, we often encounter confusion and difficulty as others attempt to understand us. This requires us to explore alternative expressions, searching for representations that others understand. We do this naturally, and our attention is drawn to the process only when it fails. However we are generally capable of finding enough common ground for communication of knowledge to proceed. We are often even able to convey basic information without a common language, as any tourist who has managed to gain directions to a restaurant or train station with much waving of hands can attest.

Computer mediated communication removes many of the mechanisms that we use to assist our process of reconciling conceptual differences during interpersonal communication, and generally leaves us at best with spoken or written language. We use the term *ontological reconciliation* for the process of matching conceptual differences. Anecdotal evidence documents the detrimental affects on effective communication of using a 'low bandwidth' medium such as a telephone or a 'high latency' medium such as the post or e-mail. The effects of limited representation of concepts are exacerbated when computers are no longer just the communication medium but also themselves participants in the communication and knowledge manipulation. In order for the processing power of computers to be utilised, knowledge must be reduced to a representation suitable for logical operations. Fitting knowledge to logical representations is largely a subjective process. Decisions must be made about how to express complex concepts in relatively constrained languages; these decisions are made by people whose choices of representation and expression are influenced by their own cultural background. Consequently, as context is lost problems then arise as other organisations with different cultures, or even just individuals with different conceptualisations, attempt to understand the logical representation and rebuild the original knowledge.

Let's return to our running example of university web sites. We accept that universities must deal with teaching and research. Most universities offer undergraduate degrees in the areas of engineering, arts, science and commerce. But when it comes to describing their activities, where one university may use the word *course* to refer to a particular degree program, another will use *course* to mean an individual subject within a degree; a third institution may use *course* to describe a particular stream or program within a degree. Some institutions will say *unit* where others say *subject* or *class*. Due to their own individual organisational cultures, different institutions use different vocabularies to describe their activities. The researcher wishing to compare the services provided by different universities will generally quickly identify the differences. Through an understanding of the knowledge domain of university activities and services, a researcher will be able to translate between terms, usually assimilating them into their own personal ontological understanding, which itself will be shaped by personal experiences. If they are from a university that uses *course* to mean a unit of teaching and *program* to describe an undergraduate degree, they will probably translate the descriptions from other institutions into this ontology. If they are not from a particular university, they will probably draw on whatever experience they have of academic institutions, and if they have none, they may build their own ontology from the collection of university representations.

To create software agents that can handle this level of ontological complexity would seem to be very difficult. Why then is it preferable to simply agreeing upon a global ontology to which all agents subscribe, a centralised language of understanding and representation, or even a global directory of multiple re-usable ontologies from which agents select as necessary? Ontology creation itself is very difficult. It requires the ability to define many concepts precisely and consistently. It requires the ability to predict appropriate assumptions and generalisations that will be acceptable to most, if not all, people. It also requires universal access and distribution infrastructure, and a well-established and accepted knowledge representation format. It requires some way to address the desire for agents and humans to interact at variable levels of abstraction as particular situations demand. It requires constant maintenance to ensure freshness and currency, yet also must provide backward compatibility for old agents. It requires that agent developers familiarise themselves with the prescribed knowledge representation formats, ontologies and protocols and adapt their own development efforts to suit them. These issues make a global ontology infrastructure unsuitable as the sole approach, and it is our belief that effort spent adding tolerance of heterogeneity to systems will provide greater benefit as we begin to introduce agents to our multi-cultural world.

In addition to the practical benefits, one of our strongest desires for tolerance of heterogeneity for software systems is rooted unashamedly in idealism: humans manage to resolve ontological differences successfully, in real time and 'on the fly'. This ability gives us much flexibility and adaptability and allows us to specialise and optimise where possible and yet generalise and compromise when necessary. Therefore, it seems both feasible and desirable to have as a goal a similar capability for software agents.

If we are to make effective use of multi-cultural data from heterogeneous sources, we need ways and means to reconcile the differences in representation. If we are to work efficiently to solve large information problems, we need the assistance of automated mechanisms. To achieve both, we need systems that are tolerant of heterogeneity.

Reconciling ontological differences requires understanding the difference between concepts and their representations; in semiotic terms, appreciating the difference between the signified and the signifier. Reconciling ontological differences means reading multiple texts that represent identical, similar or related concepts and being able to work with them at the concept level rather than at the level of representation.

For databases or XML documents, ontologicial reconciliation might be as simple as realising that two fields in different data sources actually contain the same class of data. On the other hand, it might be as complex as deciding that articles from an economics magazine and an automotive magazine are discussing different topics even though they both have 'Ford' and 'analysis' in their titles, something that current search technologies would be unlikely to realise.

As the number of data sources available to us and our ability to access them on demand and in real time is increasing, the overhead of pre-constructing a complete ontology for a given interaction becomes less and less viable. Large scale interconnectedness and increased frequency of data transactions across organisational and cultural borders leads to a reduction in the useful life of any context constructed for a particular transaction. Just as we are able to establish contexts and construct suitable local ontologies as needed for particular interactions, if we want to be able to include software agents in our higher level communication and knowledge management they will need to be capable of similar conceptualisation.

## 5 IAL Developments

The Intelligent Agent Laboratory at the University of Melbourne has been working for a number of years on knowledge representation and manipulation for information agents [13, 14]. When considering how best to structure knowledge for information agents, two questions arise: what types of knowledge should be pre-defined and what should be left to be learned dynamically? Research at the Intelligent Agent Laboratory addresses these questions in both theory and practice; the remainder of this paper describes three recent projects.