AgentLink

2000 – 2005

The long term objective of AgentLink II is to put Europe at the leading edge of international competitiveness in the area of agent-based computing. The medium term goals of AgentLink II are:

to gain competitive advantage for European industry by promoting and raising awareness of agent systems technology;

to facilitate improvement in the standard, profile, and industrial relevance of European research in the area of agent-based computer systems;

to promote excellence of teaching and training in the area of agent-based systems;

to provide a widely known, high quality European forum in which current issues, problems, and solutions in the research and development of agent-based computer systems may be debated and resolved.

Jennings, N. R. (2000) On Agent-Based Software Engineering. Artificial Intelligence, 117 (2). pp. 277-296.

2000

Agent-based computing represents an exciting new synthesis both for Artificial Intelligence (AI) and, more generally, Computer Science. It has the potential to significantly improve the theory and the practice of modeling, designing, and implementing computer systems. Yet, to date, there has been little systematic analysis of what makes the agent-based approach such an appealing and powerful computational model. Moreover, even less effort has been devoted to discussing the inherent disadvantages that stem from adopting an agent-oriented view. Here both sets of issues are explored. The standpoint of this analysis is the role of agent-based software in solving complex, real-world problems. In particular, it will be argued that the development of robust and scalable software systems requires autonomous agents that can complete their objectives while situated in a dynamic and uncertain environment, that can engage in rich, high-level social interactions, and that can operate within flexible organisational structures.

Jennings, N. R., Faratin, P., Norman, T. J., O'Brien, P. and Odgers, B. (2000) Autonomous Agents for Business Process Management. Int. Journal of Applied Artificial Intelligence, 14 (2). pp. 145-189.

2000

Traditional approaches to managing business processes are often inadequate for large-scale organisation-wide, dynamic settings. However, since Internet and Intranet technologies have become widespraed, an increasing number of business processes exhibit these properties. Therefore, a new approach is needed. To this end, we describe the motivation, conceptualization, design, and implementation of a novel agent-based business process management system. The key advance of our system is that responsibility for enacting various components of the business process is delegated to a number of autonomous problem solving agents. To enact their role, these agents typically interact and negotiate with other agents in order to coordinate their actions and to buy in the services they require. This approach leads to a system that is significantly more agile and robust than its traditional counterparts. To help demonstrate these benefits, a companion paper describes the application of our system to a real-world problem faced by British Telecom.

Wooldridge, M., Jennings, N. R. and Kinny, D. (2000) The Gaia Methodology for Agent-Oriented Analysis and Design. Journal of Autonomous Agents and Multi-Agent Systems, 3 (3). pp. 285-312.

2000

This article presents Gaia: a methodology for agent-oriented analysis and design. The Gaia methodology is both general, in that it is applicable to a wide range of multi-agent systems, and comprehensive, in that it deals with both the macro-level (societal) and the micro-level (agent) aspects of systems. Gaia is founded on the view of a multi-agent system as a computational organisation consisting of various interacting roles. We illustrate Gaia through a case study (an agent-based business process management system).

Jennings, N. R. (2001) An agent-based approach for building complex software systems. Communications of the ACM, 44 (4). pp. 35-41.

2001

Building high-quality, industrial-strength software is difficult. Indeed, it has been argued that developing such software in domains like telecommunications, industrial control, and business process management represents one of the most complex construction tasks humans undertake. Against this background, a wide range of software engineering paradigms have been devised.Each successive development either claims to make the engineering process easier ority of applications that can feasibly be built.

Zambonelli, F., Jennings, N. R. and Wooldridge, M. (2001) Organisational Rules as an Abstraction for the Analysis and Design of Multi-Agent Systems. International Journal of Software Engineering and Knowledge Engineering, 11 (3). pp. 303-328.

2001

Multi-agent systems can very naturally be viewed as computational organisations. For this reason, we believe organisational abstractions offer a promising set of metaphors and models that can be exploited in the analysis and design of such systems. To this end, the concept of role models is increasingly being used to specify and design multi-agent systems. However, this is not the full picture. In this paper we introduce three additional organisational concepts - organisational rules , organisational structures, and organisational patterns - and discuss why we believe they are necessary for the complete specification of computational organisations. In particular, we focus on the concept of organisational rules and introduce a formalism, based on temporal logic, to specify them. This formalism is then used to drive the definition of the organisational structure and the identification of the organisational patterns. Finally, the paper sketches some guidelines for a methodology for agent-oriented systems based on our expanded set of organisational abstractions.

Payne, T. R., Singh, R. and Sycara, K. (2002) Calendar Agents on the Semantic Web. IEEE Intelligent Systems, 17 (3). pp. 84-86.

2002

The emergence of the Semantic Web2 has simplified and improved knowledge reuse on the Web. Agents can parse Semantic Web pages to elicit relevant information. They can now understand and reason about information and use
it to meet users' needs and provide assistance through reference to ontologies, axioms, and languages such as DARPA Agent Markup Language. The emerging symbiosis of knowledge about users' desires, preferences, and habits with information garnered from the Semantic Web results in superior agent-based assistance than that provided by existing agents.

Reed, C., Norman, T. J. and Jennings, N. R. (2002) Negotiating the semantics of agent communication languages. Computational Intelligence, 18 (2). pp. 229-252.

2002

This paper presents a formal framework and outlines a method that autonomous agents can use to negotiate the semantics of their communication language at run-time. Such an ability is needed in open multi-agent systems so that agents can ensure they understand the implications of the utterances that are being made and so that they can tailor the meaning of the primitives to best fit their prevailing circumstances. To this end, the semantic space framework provides a systematic means of classifying the primitives along multiple relevant dimensions. This classification can then be used by the agents to structure their negotiation (or semantic fixing) process so that they converge to the mutually agreeable semantics that are necessary for coherent social interactions.

He, M., Jennings, N. R. and Leung, H. (2003) On agent-mediated electronic commerce. IEEE Trans on Knowledge and Data Engineering, 15 (4). pp. 985-1003.

2003

This paper surveys and analyzes the state of the art of agent-mediated electronic commerce (e-commerce), concentrating
particularly on the business-to-consumer (B2C) and business-to-business (B2B) aspects. From the consumer buying behavior perspective, agents are being used in the following activities: need identification, product brokering, buyer coalition formation, merchant brokering, and negotiation. The roles of agents in B2B e-commerce are discussed through the business-to-business transaction model that identifies agents as being employed in partnership formation, brokering, and negotiation. Having identified the roles for agents in B2C and B2B e-commerce, some of the key underpinning technologies of this vision are highlighted. Finally, we conclude by discussing the future directions and potential impediments to the wide-scale adoption of agent-mediated e-commerce.

Jennings, N. R. and Bussmann, S. (2003) Agent-based control systems. IEEE Control Systems, 23 (3). pp. 61-74.

2003

Modern control systems must meet increasingly demanding requirements stemming from the need to cope with significant degrees of uncertainty, as
well as with environments, and to provide greater flexibility. This,in turn, means that control systems software is highly complex in that it invariably has a large number of interacting parts. This complexity requires that state-of-the-art software engineering methods and techniques be employed. In this article, we will argue that analyzing, designing, and implementing such complex software systems as a collection of interacting, autonomous, flexible components (i.e., as agents) affords software engineers several significant advantages over contemporary methods.

Luo, X., Jennings, N. R., Shadbolt, N., Leung, H. and Lee, J. H. (2003) "A fuzzy constraint based model for bilateral multi-issue negotiations in semi-competitive environments". Artificial Intelligence Journal, 148 (1-2). pp. 53-102.

2003

This paper develops a fuzzy constraint based model for bilateral multi-issue negotiation in trading environments. In particular, we are concerned with the principled negotiation approach in which agents seek to strike a fair deal for both parties, but which, nevertheless, maximises their own payoff. Thus, there are elements of both competition and cooperation in the negotiation (hence semi-competitive environments). One of the key intuitions of the approach is that there is often more than one option that can satisfy the interests of both parties. So, if the opponent cannot accept an offer then the proponent should endeavour to find an alternative that is equally acceptable to it, but more acceptable to the opponent. That is, the agent should make a trade-off. Only if such a trade-off is not possible should the agent make a concession. Against this background, our model ensures the agents reach a deal that is fair (Pareto-optimal) for both parties if such a solution exists.

Agent technologies will boost business efficiency

2004

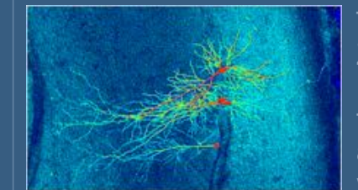

This year's Agent Technology Conference, taking place in Zurich on Friday 1 October, will bring together world-leading agent technology experts and representatives from traditional industries to report on how they are using autonomous computer agents to perform functions in the workplace. The conference is co-ordinated by the School of Electronics and Computer Science, working with colleagues at the University of Liverpool.

Can Computers Argue?

2004

The answer is yes, and not only that: they can also evaluate what will be the most successful strategy for conflict resolution, including re-formulating their action, or evading confrontation. Argument is used by computer agents only as the last resort.

Middleton, S. E., Shadbolt, N. R. and De Roure, D. C. (2004) Ontological User Profiling in Recommender Systems. ACM Transactions on Information Systems (TOIS), 22 (1). pp. 54-88.

2004

We explore a novel ontological approach to user profiling within recommender systems, working on the problem of recommending on-line academic research papers. Our two experimental systems, Quickstep and Foxtrot, create user profiles from unobtrusively monitored behaviour and relevance feedback, representing the profiles in terms of a research paper topic ontology. A novel profile visualization approach is taken to acquire profile feedback. Research papers are classified using ontological classes and collaborative recommendation algorithms used to recommend papers seen by similar people on their current topics of interest. Two small-scale experiments, with 24 subjects over 3 months, and a large-scale experiment, with 260 subjects over an academic year, are conducted to evaluate different aspects of our approach. Ontological inference is shown to improve user profiling, external ontological knowledge used to successfully bootstrap a recommender system and profile visualization employed to improve profiling accuracy. The overall performance of our ontological recommender systems are also presented and favourably compared to other systems in the literature.

Prisoner's Dilemma Competition won by ECS team

2004

A solution which allows computer agents to collude, rather than compete with each other, has won this year's recreation of Axelrod's classic Iterated Prisoner's Dilemma competition.

The winning solution was devised by a team from the University of Southampton's School of Electronics and Computer Science (ECS) which met the competition's 20th anniversary challenge to resolve the original dilemma in a noisy environment where moves could be misinterpreted.

ALADDIN: Autonomous Learning Agents for Decentralised Data and Information Networks

2005 – 2010

This project aims to develop techniques, methods and architectures for modelling, designing and building decentralised systems that can bring together information from a variety of heterogeneous sources in order to take informed actions. To do this, the project needs to take a total systems view on information and knowledge fusion and to consider the feedback that exists between sensing, decision making and acting in such systems. Moreover, it must be able to achieve these objectives in environments in which: control is distributed; uncertainty, ambiguity, imprecision and bias are endemic; multiple stakeholders with different aims and objectives are present; and resources are limited and continually vary during the system's operation.

Computer agents to aid disaster recovery scenarios

2005

An ambitious five-year research programme is aiming to find solutions for some of the most complex and challenging problems that we currently face. Disaster recovery, particularly the restoration and maintenance of decentralised data and information systems, has been chosen as the real-world application for the ALADDIN project, which comprises a team of experts led by Professor Nick Jennings of the School of Electronics and Computer Science

De Roure, D., Jennings, N. R. and Shadbolt, N. R. (2005) The Semantic Grid: Past, Present and Future. Proceedings of the IEEE, 93 (3). pp. 669-681.

2005

Grid computing offers significant enhancements to our capabilities for computation, information processing and collaboration, and has exciting ambitions in many fields of endeavour. In this paper we argue that the full richness of the Grid vision, with its application in e-Science, e-Research or e-Business, requires the 'Semantic Grid'. The Semantic Grid is an extension of the current Grid in which information and services are given well-defined meaning, better enabling computers and people to work in cooperation. To this end, we outline the requirements of the Semantic Grid, discuss the state of the art in achieving them, and identify the key research challenges in realising this vision.

Professor Nick Jennings elected FREng

2005

Professor Nick Jennings of the School of Electronics and Computer Science has been elected a Fellow of the Royal Academy of Engineering, the UK's national academy for the promotion of engineering and technology. Professor Jennings is Deputy Head of School (Research) and Head of the Intelligence, Agents, Multimedia Group. He is Director of the newly established £5.5 million BAE Systems/EPSRC Strategic Partnership in Distributed Data and Information Systems.

ECS double success at AAMAS 06

2006

Teams from ECS carried off two major awards at AAMAS 06, the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, held this month in Hakodate, Japan.

New models to improve the reliability of virtual organisations

2006

Models of virtual organisations that operate within reliable and trustworthy parameters are being developed in ECS. As the market for virtual organisations grows there is a need to ensure that computerised agents can be trusted.

Teacy, W. T. L., Patel, J., Jennings, N. R. and Luck, M. (2006) TRAVOS: Trust and Reputation in the Context of Inaccurate Information Sources. Autonomous Agents and Multi-Agent Systems, 12 (2). pp. 183-198.

2006

In many dynamic open systems, agents have to interact with one another to achieve their goals. Here, agents may be self-interested, and when trusted to perform an action for another, may betray that trust by not performing the action as required. In addition, due to the size of such systems, agents will often interact with other agents with which they have little or no past experience. There is therefore a need to develop a model of trust and reputation that will ensure good interactions among software agents in large scale open sys-tems. Against this background, we have developed TRAVOS (Trust and Reputation model for Agent-based Virtual OrganisationS) which models an agent s trust in an interaction partner. Specifically, trust is calculated using probability theory taking account of past interactions between agents, and when there is a lack of personal experience between agents, the model draws upon reputation information gathered from third parties. In this latter case, we pay particular attention to handling the possibility that reputation information may be inaccurate.

Dash, R. K., Vytelingum, P., Rogers, A., David, E. and Jennings, N. R. (2007) Market-Based Task Allocation Mechanisms for Limited Capacity Suppliers. IEEE Transactions on Systems, Man, and Cybernetics - Part A, 37 (3). pp. 391-405.

2007

This paper reports on the design and comparison of two economically-inspired mechanisms for task allocation in environments where sellers have finite production capacities and a cost structure composed of a fixed overhead cost and a constant marginal cost. Such mechanisms are required when a system consists of multiple self-interested stakeholders that each possess private information that is relevant to solving a system-wide problem. Against this background, we first develop a computationally tractable centralised mechanism that finds the set of producers that have the lowest total cost in providing a certain demand (i.e. it is efficient). We achieve this by extending the standard Vickrey-Clarke-Groves mechanism to allow for multi-attribute bids and by introducing a novel penalty scheme such that producers are incentivised to truthfully report their capacities and their costs.

ECS disaster simulation system wins championship prize

2007

An ECS-designed system which relies on computerised agents to act on its behalf during emergency scenarios has been awarded a RoboCupRescue championship prize.

ECS to develop automated agents for the stock market

2007

An ECS team which won an automated trading agent competition believe their entry may form the basis for an automated stock market.

IAM agent repeats trust success in ART competition

2007

For the second year running an ECS team has won the Agent Reputation and Trust (ART) competition at the International Joint Conference on Autonomous Agents and Multiagent Systems, held this month in Hawaii.

Rahwan, T. and Jennings, N. R. (2007) An Algorithm for Distributing Coalitional Value Calculations among Cooperating Agents. Artificial Intelligence Journal, 171 (8-9). pp. 535-567.

2007

The process of forming coalitions of software agents generally requires calculating a value for every possible coalition which indicates how beneficial that coalition would be if it was formed. Now, instead of having a single agent calculate all these values (as is typically the case), it is more efficient to distribute this calculation among the agents, thus using all the computational resources available to the system and avoiding the existence of a single point of failure. Given this, we present a novel algorithm for distributing this calculation among agents in cooperative environments. Specifically, by using our algorithm, each agent is assigned some part of the calculation such that the agents' shares are exhaustive and disjoint.

Control and Management of Autonomous Mobile Sensors

2008 – 2010

This project addresses the challenge of developing effective and computational efficient inference and coordination algorithms in order to allow multiple mobile (and stationary) sensors to form agile teams such that they can efficiently represent, explore and search challenging, uncertain and dynamic environments. Working in collaboration with the University of Oxford, the project will combine fundamental theory, algorithms and methodologies from the fields of multi-agent systems, decentralised control and Bayesian inference to allow physically distributed autonomous sensors to make effective, timely and coordinated decisions.

Instant Knowledge

2008 – 2011

The 'Instant Knowledge' research project is jointly funded by MobileVCE, the Technology Strategy Board (TSB) and the EPSRC, carrying out research into the field of machine learning and intelligent agents. The project itself centres around the notion of a personal distributed environment involving mobile devices (such as phones, PDA's, laptops etc.) and using context-sensitive information to extract and share knowledge in a useful way; whilst simultaneously addressing any security and privacy issues.

Computerised agents to cope with disasters

2008

A system which will use networks of computerised agents to cope with disaster scenarios such as outbreaks of fires, will be outlined by an ECS researcher on Guy Fawkes night.

Farinelli, A., Rogers, A., Petcu, A. and Jennings, N. R. (2008) Decentralised Coordination of Low-Power Embedded Devices Using the Max-Sum Algorithm. In: Seventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-08), 12-16 May 2008, Estoril, Portugal. pp. 639-646.

2008

This paper considers the problem of performing decentralised coordination of low-power embedded devices (as is required within many environmental sensing and surveillance applications). Specifically, we address the generic problem of maximising social welfare within a group of interacting agents. We propose a novel representation of the problem, as a cyclic bipartite factor graph, composed of variable and function nodes (representing the agents' states and utilities respectively). We show that such representation allows us to use an extension of the max-sum algorithm to generate approximate solutions to this global optimisation problem through local decentralised message passing.

More evidence that autonomous agents are emerging from the laboratory

2008

For the second year running an ECS dissertation on computerised agents has won the BCS Distinguished Dissertation Competition, indicating that agents are becoming a reality.

New release of computerised agent simulator

2008

A new simulator which will lead to the creation of more effective computerised agents in disaster scenarios has been made available by an ECS research team.

Payne, T. R. (2008) Web Services from an Agent Perspective. IEEE Intelligent Systems, 23 (2).

2008

The growth of automated services on the Web has facilitated the emergence of a new software development paradigm. This paradigm is based on the composition of disparate services to achieve large-scale, reusable solutions for diverse domains. Borrowing heavily from research in object-oriented software engineering and multiagent systems, the Web service paradigm supports the construction, publication, provision, integration, and utilization of course-grained, software services in open, heterogeneous environments. This article examines this Web service paradigm from an open multiagent systems perspective and contrasts the formally grounded, knowledge-centric view of agents with the pragmatic, declarative, bottom-up approach adopted by Web services.

ALADDIN scoops Engineering Award for "stunning collaboration"

2009

The ALADDIN project, which uses agent-based technology to tackle natural and man-made disasters has just received the Aerospace and Defence award in this year's Engineer Technology & Innovation Awards.

Agents which haggle and resolve conflict

2009

A new series of algorithms which enables computerised agents to haggle and to resolve conflict have been devised by a team led by Professor Nick Jennings.

Chalkiadakis, G., Elkind, E. and Jennings, N. (2009) Simple Coalitional Games with Beliefs. In: Proc 21st Int. Joint Conf on AI (IJCAI), Pasadena, USA. pp. 85-90.

2009

We introduce coalitional games with beliefs (CGBs), a natural generalization of coalitional games to environments where agents possess private beliefs regarding the capabilities (or types) of others. We put forward a model to capture such agent-type uncertainty, and study coalitional stability in this setting. Specifically, we introduce a notion of the core for CGBs, both with and without coalition structures. For simple games without coalition structures, we then provide a characterization of the core that matches the one for the full information case, and use it to derive a polynomial-time algorithm to check core nonemptiness. In contrast, we demonstrate that in games with coalition structures allowing beliefs increases the computational complexity of stability-related problems. In doing so, we introduce and analyze weighted voting games with beliefs, which may be of independent interest. Finally, we discuss connections between our model and other classes of coalitional games.

ECS researcher's computerised negotiating agents go commercial

2009

Computerised agents developed by an ECS academic which will negotiate the best online deals for buyers and sellers will be fully operational by the end of the year.

Intelligent Decentralised Energy-Aware Systems (iDEaS)

2009 – 2011

This project aims to explore the issues associated with the decentralised control, operation and management of future generation electricity networks. It is targetted at scenarios in which micro-generation and storage capabilities are ubiquitous, where intelligent sensing devices allow users to make informed choices about the control of devices in their home, and where producers and consumers are connected via a series of dynamically negotiated supply contracts. This is an industrially funded project from a Hampshire-based company.

Rahwan, T., Ramchurn, S., Jennings, N. and Giovannucci, A. (2009) An Anytime Algorithm for Optimal Coalition Structure Generation. Journal of Artificial Intelligence Research (JAIR), 34 . pp. 521-567.

2009

Coalition formation is a fundamental type of interaction that involves the creation of coherent groupings of distinct, autonomous, agents in order to efficiently achieve their individual or collective goals. Forming effective coalitions is a major research challenge in the field of multi-agent systems. Central to this endeavour is the problem of determining which of the many possible coalitions to form in order to achieve some goal. This usually requires calculating a value for every possible coalition, known as the coalition value, which indicates how beneficial that coalition would be if it was formed. Once these values are calculated, the agents usually need to find a combination of coalitions, in which every agent belongs to exactly one coalition, and by which the overall outcome of the system is maximized. However, this coalition structure generation problem is extremely challenging due to the number of possible solutions that need to be examined, which grows exponentially with the number of agents involved. To date, therefore, many algorithms have been proposed to solve this problem using different techniques—ranging from dynamic programming, to integer programming, to stochastic search — all of which suffer from major limitations relating to execution time, solution quality, and memory requirements. With this in mind, we develop an anytime algorithm to solve the coalition structure generation problem.

2001 2002 2003 2004 2005 2006 2007 2008 2009 2010

Agents Research in ECS