1. Plays as effective multiagent plans enabling opponent-adaptive play selection

@inproceedings{Bowling_2004_4838,
author = "Michael Bowling and Brett Browning and Manuela Veloso",
title = "Plays as effective multiagent plans enabling opponent-adaptive play selection",
booktitle = "Proceedings of International Conference on Automated Planning and Scheduling (ICAPS'04)",
year = "2004",
note = "in press"
}

Abstract

Coordinated action for a team of robots is a challenging problem, especially in dynamic, unpredictable environments. Robot soccer is an instance of a domain where well defined goals need to be achieved by multiple executors in an adversarial setting. Such domains offer challenging multiagent planning problems that need to coordinate multiagent execution in response to other agents that are not part of our team plans. In this work, we introduce the concept of a play as a multiagent plan that combines both reactive principles, which are the focus of traditional approaches for coordinating robot actions, and deliberative principles. We further introduce the concept of a playbook as a method for seamlessly combining multiple team plans. The playbook provides a set of alternative team behaviors which form the basis for our third contribution of play adaptation. We describe how these concepts were concretely implemented in the CMDragons robot soccer team. We also show empirical results indicating the importance of adaptation in adversarial or other unpredictable environments.

2. Research Directions for Service-Oriented Multiagent Systems

@article{Huhns_2005_5239,
author = "Michael N. Huhns and Munindar P. Singh and Mark Burstein and Keith Decker and Edmund Durfee and Tim Finin and Les Gasser and Hrishikesh Goradia and Nick Jennings and Kiran Lakkaraju and Hideyuki Nakashima and H. Van Dyke Parunak and Jeffrey S. Rosenschein and Alicia Ruvinsky and Gita Sukthankar and Samarth Swarup and Katia Sycara and Milind Tambe and Tom Wagner and Laura Zavala",
title = "Research Directions for Service-Oriented Multiagent Systems",
journal = "IEEE Internet Computing",
month = "November",
year = "2005",
volume = "9",
number = "6",
pages = "65-70"
}

Abstract

Today's service-oriented systems realize many ideas from the research conducted a decade or so ago in multiagent systems. Because these two fields are so deeply connected, further advances in multiagent systems could feed into tomorrow's successful service-oriented computing approaches.This article describes a 15-year roadmap for service-oriented multiagent system research.

3. Coordination Advice: A Preliminary Investigation of Human Advice to Multiagent Teams

@inproceedings{Schurr_2004_4892,
author = "N. Schurr and Paul Scerri and M. Tambe",
title = "Coordination Advice: A Preliminary Investigation of Human Advice to Multiagent Teams",
booktitle = "AAAI Spring Symposium on Interaction between Humans and Autonomous Systems over Extended Operation",
year = "2004",
note = "Invited Paper"
}

Abstract

This paper introduces a new area of advice that is specific to advising a multiagent team: Coordination Advice. Coordination Advice differs from traditional advice because it pertains to coordinated tasks and interactions between agents. Given a large multiagent team interacting in a dynamic domain, optimal coordination is a difficult challenge. Human advisors can improve such coordination via advice. This paper is a preliminary look at the evolution of Coordination Advice from a human through three different domains: (i) disaster rescue simulation, (ii) a self-maintaining robotics sensors, and (iii) personal assistants in a office environment. We study how the useful advice a person can give changes as the domains change and the number of agents and roles increase.

4. A formalization of equilibria for multiagent planning

@inproceedings{Bowling_2002_4145,
author = "Michael Bowling and Rune Jensen and Manuela Veloso",
title = "A formalization of equilibria for multiagent planning",
booktitle = "Proceedings of the AAAI-2002 Workshop on Multiagent Planning",
month = "August",
year = "2002"
}

5. Multiagent learning using a variable learning rate

@article{Bowling_2002_4138,
author = "Michael Bowling and Manuela Veloso",
title = "Multiagent learning using a variable learning rate",
journal = "Artificial Intelligence",
year = "2002",
volume = "136",
pages = "215 - 250"
}

6. Simulation level of detail for multiagent control

@inproceedings{Brogan_2002_5195,
  author = "David C. Brogan and Jessica K Hodgins",
  title = "Simulation level of detail for multiagent control",
  booktitle = "In Proceedings of the first international joint conference on Autonomous agents and multiagent systems",
  year = "2002",
  pages = "199 - 206",
  publisher = "ACM Press"
}

Abstract

Many classes of applications require multiagent navigation control algorithms to specify the movements and actions of heterogeneous groups containing thousands of characters. The scale and complexity of these interacting character groups require navigation control algorithms that are both generalizable and specifically tuned to particular character platforms. We propose a technique called simulation level of detail (LOD) that provides a simulation-based interface between navigation control algorithms and the specific mobile characters on which they operate. A simulation LOD efficiently models a character's ability to move given its dynamic state and provides this simplified version of the character to navigation controllers for use in run-time search algorithms that compute locomotion actions. We develop our simulation LOD algorithms on groups of physically simulated human and alien bicyclists and demonstrate reusable controllers that provide improvements in path following and herding tasks.

7. BiddingBot: a multiagent support system for cooperative bidding in multiple auctions

@inproceedings{Ito_2000_3575,
  author = "T. Ito and N. Fukuta and Toramatsu Shintani and Katia Sycara",
  title = "BiddingBot: a multiagent support system for cooperative bidding in multiple auctions",
  booktitle = "Proceedings of the Fourth International Conference on MultiAgent Systems",
  month = "July",
  year = "2000",
  pages = "399 - 400"
}

Abstract

Online auctions are becoming an increasingly important channel for electronic commerce. There exist more than 150 online auction sites on the Internet. It is difficult for users to attend, monitor, and bid at multiple auction sites simultaneously. The authors propose BiddingBot which is a multi agent system that supports users in attending, monitoring, and bidding in multiple auctions. BiddingBot monitors prices of goods in several online auction sites to get reasonable market prices of goods, and uses a new cooperative bidding mechanism to effectively bid in auctions.

8. Interleaving Planning and Execution in a Multiagent Team Planning Environment

@techreport{Paolucci_2000_3274,
  author = "Massimo Paolucci and Onn Shehory and Katia Sycara",
  title = "Interleaving Planning and Execution in a Multiagent Team Planning Environment",
  institution = "Robotics Institute, Carnegie Mellon University",
  month = "January",
  year = "2000",
  number = "CMU-RI-TR-00-01",
  address = "Pittsburgh, PA"
}

Abstract

Agents in a multiagent system may need to share information and services. For this, they need to be able to interleave deliberative planning with execution of actions. The deliberative planning is needed to decide which actions to perform to achieve an objective, whereas execution of some of the actions is needed to make a more informed decision on the other actions and to access services provided by other agents.

HITaP is a planner that interleaves planning and execution: using HITaP an agent can, during planning, gather information by either direct inspection of the domain or by firing queries to other agents and recording their answers. Interleaving planning and execution, as provided by HITaP, plays a crucial role in an agent's ability to construct shared plans with other agents and to manage the negotiation process that leads to agreement with the agent's teammates on these plans.

HITaP is implemented and currently used as planning module for agents in the RETSINA multiagent system. These agents cooperate to solve problems in different domains that range from portfolio management to command and control decision support systems.\footnote{The authors thank Dirk Kalp and Anandeep Pannu for their contribution to the initial implementation of the planner. In addition we are grateful to Josepth Giampapa for the stimulating discussions on team behavior. This research has been sponsored in part by ONR grant.

9. Experience-based Reinforcement Learning to Acquire Effective Behavior in a Multiagent Domain

@inproceedings{Arai_2000_3729,
  author = "Sachiyo Arai and Katia Sycara and Terence Payne",
  title = "Experience-based Reinforcement Learning to Acquire Effective Behavior in a Multiagent Domain",
  booktitle = "The Sixth Pacific Rim International Conference on Artificial Intelligence (PRICAI 2000)",
  year = "2000",
  pages = "125-135",
  publisher = "Springer-Verlag"
}

Abstract

In this paper, we discuss Profit-sharing, an experience-based reinforcement learning approach (which is similar to a Monte-Carlo based reinforcement learning method) that can be used to learn robust and effective actions within uncertain, dynamic, multi-agent systems. We introduce the cut-loop routine that discards looping

behavior, and demonstrate its effectiveness empirically within a simplified NEO (non-combatant evacuation operation) domain. This domain consists of several agents which ferry groups of evacuees to one of several shelters. We demonstrate that the cut-loop routine makes the Profit-sharing approach adaptive and robust within a dynamic and uncertain domain, without the need for predefined knowledge or subgoals. We also compare it empirically with the popular Q-learning approach.

10. Multiagent Systems

```
@article{Sycara_1998_2589,
  author = "Katia Sycara",
  title = "Multiagent Systems",
  journal = "AI Magazine",
  year = "1998",
  volume = "10",
  number = "2",
  pages = "79-93"
}
```

11. The RETSINA multiagent system (video session): towards integrating planning, execution and information gathering

```
@inproceedings{Sycara_1998_2806,
  author = "Katia Sycara and Anandeep S. Pannu",
  title = "The RETSINA multiagent system (video session): towards integrating planning, execution and information gathering",
  booktitle = "AGENTS '98. Proceedings of the second international conference on Autonomous agents",
  year = "1998",
  pages = "350 - 351"
}
```

12. Modeling Supply Chain Dynamics: A Multiagent Approach

```
@article{Swaminathan_1997_3166,
  author = "Jayashankar Swaminathan and Stephen Smith and Norman Sadeh-Koniecpol",
  title = "Modeling Supply Chain Dynamics: A Multiagent Approach",
  journal = "Decision Sciences",
  month = "April",
  year = "1997"
}
```

13. Multiagent Coordination in Tightly Coupled Task Scheduling

```
@inproceedings{Sycara_1996_2209,
  author = "Katia Sycara and Jyi Shane Liu",
  title = "Multiagent Coordination in Tightly Coupled Task Scheduling",
  booktitle = "1996 International Conference on Multi-Agent Systems",
  year = "1996"
}
```

Abstract

We consider an environment where agents' tasks are tightly coupled and require real-time scheduling and execution. In order to complete their tasks, agents need to

coordinate their actions both constantly and extensively. We present an approach that consists of a standard operating procedure and a look-ahead coordination. The standard operating procedure regulates task coupling and minimizes communication. The look-ahead coordination increases agents' global visibility and provides indicative information for decision adjustment. The goal of our approach is to prune decision myopia while maintaining system responsiveness in real-time, dynamic environments. Experimental results in job shop scheduling problems show that (1) the look-ahead coordination significantly enhances the performance of the standard operating procedure in solution quality, (2) the approach is capable of producing solutions of very high quality in a real-time environment.

14. Cross-Element Validation in Multiagent-Based Simulation: Switching Learning Mechanisms in Agents

@article{repec:jas:jasssj:2003-27-1,
author={keiki takadama and yutaka l. suematsu and norikazu sugimoto and norberto e. nawa and katsunori shimohara},
title={cross-element validation in multiagent-based simulation: switching learning mechanisms in agents},
journal={journal of artificial societies and social simulation},
year=2003,
volume={9},
number={},
pages={},
month={},
note={available at http://ideas.repec.org/a/jas/jasssj/2003-27-1.html}
}

Abstract

The validity of simulation results remains an open problem in multiagent-based simulation (MABS). Since such validity is based on the validation of computational models, we propose a cross-element validation method that validates computational models by investigating whether several models can produce the same results after changing an element in the agent architecture. Specifically, this paper focuses on learning mechanisms applied to agents as one of the important elements and compares three different MABSs employing either an evolutionary strategy (ES), a learning classifier system (LCS), or a reinforcement learning (RL). This type of validation is not based on the between-models addressed in conventional research but on a within-model. A comparison of the simulation results in a bargaining game, one of the fundamental examples in game theory, reveals that (1) computational models are minimally validated in the case of ES- and RL-based agents; and (2) learning mechanisms that enable agents to acquire their rational behaviors differ according to the knowledge representation (i.e., the strategies in the bargaining game) of the agents. Concretely, we found that (2-a) the ES-based agents derive the same tendency in game theory but the LCS-based agents cannot in the case of employing continuous knowledge representation; and (2-b) the same ES-based agents cannot derive the same tendency in game theory but the RL-based agents derive it in the case of employing discrete knowledge representation.

15. Experiments in a Software Aided Multiagent System

@techreport{repec:sce:scecf4:326,
author={chung-ching tai and shu-heng chen},
title={experiments in a software aided multiagent system},
year=2004,
month=aug,
institution={society for computational economics},
type={computing in economics and finance 2004},
note={available at http://ideas.repec.org/p/sce/scecf4/326.html},
number={326}
}

Abstract
Various Computational Intelligence (CI) tools have been devised to depict the essential learning or adapting processes people adopt when they are making decisions. If the learning process people possess is indeed well described by some CI algorithms, then there arise an interesting question: what will people do if they are provided the results computed by the algorithms describing their inner process? This paper attempts to create an experimental environment in which human subjects are provided certain CI tools, which is feasible in modern e-Commerce now and in the future. Observing their decisions may helps us understand more about peoples decision process as well as the rationality human have.

16. Sharing in Large Scale Teams

@inproceedings{Xu_2004_4887,
author = "Y. Xu and Michael Lewis and Katia Sycara and Paul Scerri",
title = "Sharing in Large Scale Teams",
booktitle = "AAMAS'04 Workshop on Challenges in Coordination o f Large Scale MultiAgent Systems",
year = "2004"
}

17. Coordination of Multiple Agents in Distributed Manufacturing Scheduling

@phdthesis{Liu_1996_410,
author = "Jyi Shane Liu",
title = "Coordination of Multiple Agents in Distributed Manufacturing Scheduling",
school = "Robotics Institute, Carnegie Mellon University",
month = "April",
year = "1996",
address = "Pittsburgh, PA"
}

Abstract
Distributed Problem Solving (DPS) is concerned with a group of agents, each with only local views, that cooperate and coordinate to solve a given problem in a coherent and efficient manner. As networked computer systems and practical concurrent processing architectures develop, DPS research has seen increasing realization in real world problems. As opposed to using sophisticated agents in problem solving, DPS approaches based on simple reactive agents have received growing interest. These reactive agents are simple in the sense that they do not have representations of their environments and they act by stimulus and response. They are embedded in the

environment and follow simple patterns of behavior that can be easily programmed. Generating sophisticated group behavior from coordination and integration of activities of simple individuals has potentially great significance.
This thesis presents DPS approaches based on reactive agents, and applied to scheduling in the domain of manufacturing job shops. In general, scheduling deals with organizing possibly related activities over time and limited resources. Scheduling problems occur in many domains, such as manufacturing, transportation, computer processing, communication, services, health care, education, etc. Distributed scheduling typically in concerned with coordinating among resources over processing of related activities in environments where knowledge and control distribution are desired. The objective of this work is the problem solving efficacy of reactive agents in terms of computational cost and solution quality. Our thesis is that multi-agent coordination techniques can provide substantial gains in terms of problem solving efficiency and solution quality.
The main contributions of this research are the coordination schemes and the DPS techniques implemented in the two effective multiagent scheduling systems, CORA and COFCAST. In addition, CORA and COFAST are shown to provide equivilent or superior performance to competing centralized scheduling techniques, and therefore, enrich viable approaches to scheduling problems. In a broader view, this work supplements DPS research with innovative approaches to tightly coupled problems and successful demonstration of competetive performance. This work also provides an essential component for larger-scaled DPS research in manufacturing management and control.

18. Worst-Case-Optimal Anytime Coalition Structure Generation

@inproceedings{Sandholm_1998_2565,
  author = "T. Sandholm and K. Larson and M. Andersson and Onn Shehory and F. Tohme",
  title = "Worst-Case-Optimal Anytime Coalition Structure Generation",
  booktitle = "In Proceedings of AAAI-98",
  month = "July",
  year = "1998",
  pages = "46-53"
}

Abstract
Coalition formation is one of the key problems in multiagent systems. One would prefer a coalition structure that maximizes the sum of the values of the coalitions, but often the number of coalition structures is too large to allow exhaustive search for the optimal one. This paper focuses on nding a worst case bound on how good the optimal coalition structure is compared to the best coalition structure that a nonexhaustive search finds.
We show that none of the previous coalition structure generation algorithms can establish any bound because they search fewer nodes than a threshold that we show necessary for establishing a bound. We present an algorithm that establishes a tight bound within this minimal amount of search, and show that any other algorithm would have to search strictly more. The fraction of nodes needed to be searched approaches zero as the number of agents grows.
If additional time remains, our anytime algorithm searches further, and establishes a progressively lower tight bound. Surprisingly, just searching one more node drops the bound in half. As desired, our algorithm lowers the bound rapidly early on, and

exhibits diminishing returns to computation. It also drastically outperforms its obvious contenders. Finally, we show how to distribute the desired search across self-interested manipulative agents.

19. A survey of Agent-Oriented Software Engineering

@misc{Tveit2001AOSE,
author = "Amund Tveit",
title = "A Survey of Agent-Oriented Software Engineering",
howpublished = "Proc. of the First NTNU CSGS Conference (http://www.amundt.org)",
month = "May",
year = "2001",
URL = "http://www.abiody.com/jfipa/publications/AgentOrientedSoftwareEngineering/"
}

Abstract:
Agent-Oriented Software Engineering is the one of the most recent contributions to the field of Software Engineering. It has several benefits compared to existing development approaches, in particular the ability to let agents represent high-level abstractions of active entities in a software system. This paper gives an overview of recent research and industrial applications of both general high-level methodologies and on more specific design methodologies for industry-strength software engineering.

20. Multi agent collaboration using distributed value functions

@inproceedings{Ferreira_2000_3583,
author = "Enrique Ferreira and Pradeep Khosla",
title = "Multi agent collaboration using distributed value functions",
booktitle = "Proceedings of the IEEE Intelligent Vehicles Symposium (IV 2000)",
month = "October",
year = "2000",
pages = "404 - 409"
}

Abstract
In this paper we present the use of distributed value function techniques to reach collaboration in a multiagent system. We apply this method in two different simulation environments: a mobile robot planning/searching task and an intelligent traffic system in an urban environment. In the case of the intelligent traffic system, results show an improvement with respect to a standard fix-time controller and local adaptive controllers. Trajectories for optimal search in an obstacle environment are obtained in the mobile robot case. Some variations to the actual algorithm are pointed out to suit our cases. We conclude discussing our future work.

21. Defining and using ideal teammate and opponent agent models: a case study in robotic soccer

@inproceedings{Stone_2000_3818,
author = "P. Stone and Patrick Riley and Manuela Veloso",
title = "Defining and using ideal teammate and opponent agent models: a case study in robotic soccer",
booktitle = "Proceedings of the Fourth International Conference on MultiAgent Systems 2000",
month = "July",
year = "2000",
pages = "441 - 442"

}

Abstract
A common challenge for agents in multiagent systems is trying to predict what other agents are going to do in the future. Such knowledge can help an agent determine which of its current action options are most likely to help it achieve its goals. Ideally, an agent could learn a model of other agents' behavior patterns via direct observation of their past actions. However, that is only possible when agents have many repeated interactions with one another. We explore the use of agent models in an application where extensive interactions with a particular agent are not possible, namely robotic soccer. In robotic soccer tournaments, such as RoboCup (Kitano et al., 1997), a team of agents plays against another team for a single, short (typically 10-minute) period. The opponents' behaviors are usually not observable prior to this game and there are not enough interactions during the game to build a useful model. We introduce "ideal-model-based behavior outcome prediction" (IMBBOP). This technique predicts an agent's future actions in relation to the optimal behavior in its given situation, This optimal behavior is agent-independent and can therefore be computed based solely on a model of the world dynamics. IMBBOP does not assume that the other agent will act according to the theoretical optimum, but rather characterizes its expected behavior in terms of deviation from this optimum.

22. Multi-agent reinforcement learning for planning and scheduling multiple goals

@inproceedings{Arai_2000_3576,
  author = "Sachiyo Arai and Katia Sycara and Terence Payne",
  title = "Multi-agent reinforcement learning for planning and scheduling multiple goals",
  booktitle = "Proceedings of the Fourth International Conference on MultiAgent Systems",
  month = "July",
  year = "2000",
  pages = "359 - 360"
}

Abstract
Recently, reinforcement learning has been proposed as an effective method for knowledge acquisition of multiagent systems. However, most research on multiagent systems applying a reinforcement learning algorithm, focus on a method to reduce complexity due to the existence of multiple agents and goals. Although these pre-defined structures succeeded in lessening the undesirable effect due to the existence of multiple agents, they would also suppress the desirable emergence of cooperative behaviors in the multiagent domain. We show that the potential cooperative properties among the agent are emerged by means of profit-sharing (J. Grefenstette, 1988; K. Miyazaki et al., 1994) which is robust in the non-MDPs.

23. Agent Interoperation Across Multiagent System Boundaries

@inproceedings{Giampapa_2000_3281,
  author = "Joseph Andrew Giampapa and Massimo Paolucci and Katia Sycara",
  title = "Agent Interoperation Across Multagent System Boundaries",
  booktitle = "Proceedings of the Fourth International Conference on Autonomous Agents (Agents 2000)",

month = "June",
year = "2000",
publisher = "Association for Computing Machinery",
address = "ACM, 1515 Broadway, New York, NY, 10036, USA",
note = "ISBN: 1-58113-230-1"
}

Abstract
Recently the number of autonomous agents and multiagent systems (MAS) that have been developed by different developers has increased. Despite efforts for the creation of standards (eg. in communication languages, registration protocols etc.), it is clear that at least in the near term heterogeneous agents and MASs will be prevalent. Therefore, mechanisms that allow agents and/or MASs to interoperate and transact are needed. In this paper we report on a case study and lessons learned of an interoperator agent we developed. We discuss requirements for interoperation mechanisms, resulting challenges and our design decisions and implementation of the RETSINA-OAA InterOperator\footnote{The authors thank David Martin, Adam Cheyer and Didier Guzzoni at SRI for their support and help. Marcus Huber of OGI was instrumental in explaining the operation of the OAA-KQML Bridge. In addition we are especially grateful to Martin Van Velsen for his help with the graphical display of the RETSINA and OAA systems and with the data collection for the evaluation section.

24. Communication in Domains with Unreliable, Single-Channel, Low-Bandwidth Communication

@inproceedings{Stone_1998_3175,
author = "Peter Stone and Manuela Veloso",
title = "Communication in Domains with Unreliable, Single-Channel, Low-Bandwidth Communication",
booktitle = "First International Workshop on Collective Robotics (CRW'98)",
year = "1998"
}

Abstract
In most multiagent systems with communicating agents, the agents have the luxury of using reliable, multi-step negotiation protocols. They can do so primarily when communication is reliable and the cost of communication relative to other actions is small. Conversely, this paper considers multiagent environments with unreliable, high-cost communication. This paper presents techniques for dealing with the obstacles to inter-agent communication in such environments. A successful prototype system is fully implemented and tested in the simulated robotic soccer domain.

25. A Market-Oriented Programming Environment and its Application to Distributed Multicommodity Flow Problems

@article{ wellman93marketoriented,
author = "Michael P. Wellman",
title = "A Market-Oriented Programming Environment and its Application to Distributed Multicommodity Flow Problems",
journal = "Journal of Artificial Intelligence Research",
volume = "1",
pages = "1-23",
year = "1993"

}
Abstract
Market price systems constitute a well-understood class of mechanisms that under certain conditions provide effective decentralization of decision making with minimal communication overhead. In a market-oriented programming approach to distributed problem solving, we derive the activities and resource allocations for a set of computational agents by computing the competitive equilibrium of an artificial economy. Walras provides basic constructs for defining computational market structures, and protocols for deriving their corresponding price equilibria. In a particular realization of this approach for a form of multicommodity flow problem, we see that careful construction of the decision process according to economic principles can lead to efficient distributed resource allocation, and that the behavior of the system can be meaningfully analyzed in economic terms.

26. Multi-Agent Reinforcement Learning: Independent vs. Cooperative Agents

@incollection{ tan97multiagent,
author = "Ming Tan",
title = "Multi-Agent Reinforcement Learning: Independent vs.\ Cooperative Learning",
booktitle = "Readings in Agents",
publisher = "Morgan Kaufmann",
address = "San Francisco, CA, USA",
editor = "Michael N. Huhns and Munindar P. Singh",
pages = "487--494",
year = "1997"
}
Abstract
Intelligent human agents exist in a cooperative social environment that facilitates learning. They learn not only by trialand
-error, but also through cooperation by sharing instantaneous information, episodic experience, and learned knowledge. The key investigations of this paper are, "Given the same number of reinforcement learning agents, will cooperative agents outperform independent agents who do not communicate during learning?" and "What is the price for such cooperation?".

27. Using RoboCup to Teach Multiagent Systems and the Distributed Mindset

@misc{ vidal02using,
author = "J. Vidal and P. Buhler",
title = "Using robocup to teach multiagent systems and the distributed mindset",
text = "J. M. Vidal and P. Buhler. Using robocup to teach multiagent systems and the distributed mindset. In Proceedings of the 33rd ACM Technical Symposium on Computer Science Education, pages 3--7, 2002.",
year = "2002" }
Abstract
We present our experiences using the RoboCup soccerserver simulator and Biter, our own agent platform, for the teaching of a graduate multiagent systems' class. The RoboCup simulator and Biter are both described. We argue that the combination of RoboCup and Biter forms an e_ective platform for the teaching of multiagent systems

and the distributed mindset. Results from two semesters using these tools are presented. These results con_rm our claims. Finally, we characterize this work within the framework provided by the STEELMAN Draft of the Computing Curricula 2001 initiative.

28. Multiagent Planning as Control Síntesis

@misc{ seow-multiagent,
author = "Kiam-Tian Seow and Chuan Ma and Makoto Yokoo",
title = "Multiagent Planning as Control Synthesis"
year = 2004 }

Abstract
This paper proposes a new multiagent planning approach to coordination synthesis that views distributed agents as discrete-event processes. The connection between discreteevent control synthesis and coordination planning is first established, thereby enabling the exploitation of the vast body of knowledge and associated software synthesis tools from
`The Supervisory Control of Discrete-Event Systems' for automatic coordination synthesis of distributed agents. Importantly, these coordinating agents designed collectively generate a behaviour guaranteed not to contradict any specified inter-agent constraint, is nonblocking and optimal. A simple planning methodology is proposed in terms of procedures supported by CTCT, an existing, freely available design tool developed based on the control synthesis framework. A simple example illustrates the use of the CTCT-based methodology to synthesize coordination modules for distributed agents. Discussions in relation to previous work examine the relative significance of the new multiagent planning framework.

29. The Application of Agent Technology to Health Care

@misc{ nealon02application,
author = "J. Nealon and A. Moreno",
title = "The application of Agent Technology to Healthcare",
text = "J. Nealon, and A Moreno The application of Agent Technology to Healthcare, 1st International Joint Conference on Autonomous Agents and Multiagent Systems, Bologna, July 2002.",
year = "2002" }

Abstract
In this paper we briefly describe the initial work of the proposed AgentCities Working Group on Health Care. Issues that are particularly relevant to the adoption of an agentoriented approach to developing software systems in health care are discussed: the characteristics of health care domains, agent-based solutions, and the research and development challenges inherent in employing agents to solve problems in health care. We conclude that multi-agent systems do have an increasingly important role to play in health care domains in that they significantly enhance our ability to model, design and build complex, distributed health care software systems.

30. Sistema Multiagente para Monitorización Inteligente Domiciliaria de Pacientes con Patologías Cardiovasculares

@misc{
author = "A. Gabriel Agostini ",

title = "Sistema Multiagente para Monitorización Inteligente Domiciliaria de Pacientes con Patologías Cardiovasculares ",
text = "Open Discussion Track. VIII Iberoamerican Conference on Artificial Intelligence.",
pages = "205--210",
year = "2002" }

Abstract

Actualmente existe un considerable número de personas con enfermedades cardiovasculares que sufren complicaciones en ámbitos donde el auxilio médico puede llegar demasiado tarde. Por otro lado la medicina cardiológica ha incorporado recientemente herramientas diagnósticas y pronósticas potentes y no invasivas basadas en el procesamiento de señales obtenidas del sistema cardiovascular utilizando herramientas matemáticas avanzadas. El presente trabajo propone un sistema multiagente que vincula estas herramientas y técnicas de inteligencia artificial para realizar la monitorización inteligente de pacientes con patologías cardiovasculares, como así también, para coordinar las acciones y gestionar los recursos de las tres partes fundamentales que deben intervenir en caso de una emergencia, a saber, el propio paciente y su familia, su médico de cabecera, y la institución de salud respectiva. Se utilizan sistemas de razonamiento basado en casos y basado en reglas para decidir el accionar de acuerdo a la situación de cada paciente.

31. Aplicación Clínica de la Biblioteca de Funciones VFCLab para el Análisis de la Variabilidad de la Frecuencia Cardíaca

@misc{
author = "A. Gabriel Agostini and L. Gamero and P. Rumi",
title = "Aplicación Clínica de la Biblioteca de Funciones VFCLab para el Análisis de la Variabilidad de la Frecuencia Cardíaca",
text = "XVII Congreso Brasilero de Ingeniería Biomédica.",
year = "2000" }

Abstract

The cardiology medicine has recently incorporate powerful non invasive diagnostic tools based in the processing of signals obtained from the cardiovascular system. The heart rate variability (HRV) is one of them and it has very valuable information referred to cardiovascular pathologies, the state of the autonomic nervous system and to the prognostic of the sudden cardiac death. In this work is presented a clinic application of a library of functions developed in Matlab (Matlab Toolbox) for the mathematical analysis of the HRV. The functions of this library contain tools for time and frequency domain analysis, as well as, tools for the study of the non linear dynamics of the cardiovascular system. A group of patient with characteristic cardiovascular pathologies is used for the clinic application. The results are compared with those of other works and they are presented numeric and graphically.

32. VFCLab: Biblioteca de Funciones en Matlab para el Análisis de la Variabilidad de la Frecuencia Cardíaca

@misc{
author = "A. Gabriel Agostini and L. Gamero and P. Rumi",
title = " Biblioteca de Funciones en Matlab para el Análisis de la Variabilidad de la Frecuencia Cardíaca ",
text = "XII Congreso Argentino de Bioingeniería.",
year = "1999" }

Abstract
El análisis de la variabilidad de la frecuencia cardíaca (VFC) constituye una de las herramientas más prometedoras para el estudio y diagnóstico de patologías cardiovasculares y del sistema nervioso autónomo, y para el pronóstico de la muerte súbita cardíaca. En este trabajo se presenta una biblioteca de funciones para analizar la VFC en el dominio temporal y frecuencial, como así también en el estudio de dinámicas no lineales. Se describen las funciones más importantes y se efectúa una aplicación práctica sobre un registro de Holter de 24 horas de un paciente sano. Se presentan los resultados en forma numérica y gráfica.

33. An efficient algorithm for multiagent plan coordination

@inproceedings{1082599,
author = {Jeffrey S. Cox and Edmund H. Durfee},
title = {An efficient algorithm for multiagent plan coordination},
booktitle = {AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems},
year = {2005},
isbn = {1-59593-093-0},
pages = {828--835},
location = {The Netherlands},
doi = {http://doi.acm.org/10.1145/1082473.1082599},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract
The multiagent plan coordination problem arises whenever multiple agents plan to achieve their individual goals independently, but might mutually benefit by coordinating their plans to avoid working at cross purposes or duplicating effort. Although variations of this problem have been studied in the literature, there is as yet no agreement over a general characterization of the problem. In this paper, we describe a general framework that extends the partial-order, causal-link plan representation to the multiagent case, and that treats coordination as a form of iterative repair of plan flaws that cross agents. We show, analytically and empirically, that this algorithmic formulation can scale to the multiagent case better than can a straightforward application of the most advanced single-agent plan coordination technique, highlighting fundamental differences between single-agent and multiagent planning

34. Developing multiagent systems: The Gaia methodology

@article{958963,
author = {Franco Zambonelli and Nicholas R. Jennings and Michael Wooldridge},
title = {Developing multiagent systems: The Gaia methodology},
journal = {ACM Trans. Softw. Eng. Methodol.},
volume = {12},
number = {3},
year = {2003},
issn = {1049-331X},
pages = {317--370},
doi = {http://doi.acm.org/10.1145/958961.958963},
publisher = {ACM Press},

address = {New York, NY, USA},
}
Abstract
Systems composed of interacting autonomous agents offer a promising software engineering approach for developing applications in complex domains. However, this multiagent system paradigm introduces a number of new abstractions and design/development issues when compared with more traditional approaches to software development. Accordingly, new analysis and design methodologies, as well as new tools, are needed to effectively engineer such systems. Against this background, the contribution of this article is twofold. First, we synthesize and clarify the key abstractions of agent-based computing as they pertain to agent-oriented software engineering. In particular, we argue that a multiagent system can naturally be viewed and architected as a computational organization, and we identify the appropriate organizational abstractions that are central to the analysis and design of such systems. Second, we detail and extend the Gaia methodology for the analysis and design of multiagent systems. Gaia exploits the aforementioned organizational abstractions to provide clear guidelines for the analysis and design of complex and open software systems. Two representative case studies are introduced to exemplify Gaia's concepts and to show its use and effectiveness in different types of multiagent system.

35. Learning Communication Strategies in Multiagent Systems
@article{590837,
author = {Michael Kinney and Costas Tsatsoulis},
title = {Learning Communication Strategies in Multiagent Systems},
journal = {Applied Intelligence},
volume = {9},
number = {1},
year = {1998},
issn = {0924-669X},
pages = {71--91},
doi = {http://dx.doi.org/10.1023/A:1008251315338},
publisher = {Kluwer Academic Publishers},
address = {Hingham, MA, USA},
}
Abstract
In this paper we describe a dynamic, adaptive communication strategy for multiagent systems. We discuss the behavioral parameters of each agent that need to be computed, and provide a quantitative solution to the problem of controlling these parameters. We also describe the testbed we built and the experiments we performed to evaluate the effectiveness of our methodology. Several experiments using varying populations and varying organizations of agents were performed and are reported. A number of performance measurements were collected as each experiment was performed so the effectiveness of the adaptive communications strategy could be measured quantitatively.
The adaptive communications strategy proved effective for fully connected networks of agents. The performance of these experiments improved for larger populations of agents and even approached optimal performance levels. Experiments with non-fully connected networks showed that the adaptive communications strategy is extremely effective, but does not approach optimality. Other experiments investigated the

ability of the adaptive communications strategy to compensate for “distracting” agents, for systems where agents are required to assume the role of information routers, and for systems that must decide between routing paths based on cost information.

36. A distributed framework for solving the Multiagent Plan Coordination Problem

@inproceedings{1082598,
author = {Jeffrey S. Cox and Edmund H. Durfee and Thomas Bartold},
title = {A distributed framework for solving the Multiagent Plan Coordination Problem},
booktitle = {AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems},
year = {2005},
isbn = {1-59593-093-0},
pages = {821--827},
location = {The Netherlands},
doi = {http://doi.acm.org/10.1145/1082473.1082598},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

We examine whether and how the Multiagent Plan Coordination Problem, the problem of resolving interactions between the plans of multiple agents, can be cast as a Distributed Constraint Optimization Problem (DCOP). We use ADOPT, a state-of-the-art DCOP solver that can solve DCOPs in an asynchronous, parallel manner using local communication between individual computational agents. We then demonstrate how we can take advantage of novel flaw-assignment strategies and plan coordination algorithms to significantly improve the performance of ADOPT on representative coordination problems. We close with a consideration of possible advances in framing our DCOP representation of the Multiagent Plan Coordination Problem.

37. Towards Adaptive Workflow Enactment Using Multiagent Systems

@article{1047630,
author = {Paul A. Buhler and Jos\&\#233; M. Vidal},
title = {Towards Adaptive Workflow Enactment Using Multiagent Systems},
journal = {Inf. Tech. and Management},
volume = {6},
number = {1},
year = {2005},
issn = {1385-951X},
pages = {61--87},
doi = {http://dx.doi.org/10.1007/s10799-004-7775-2},
publisher = {Kluwer Academic Publishers},
address = {Hingham, MA, USA},
}

Abstract

Advances in Information Technology have created opportunities for business enterprises to redesign their information and process management systems. The redesigned systems will likely employ some form of workflow management system.

Workflow management systems exactly enact business processes described in a process description language. Unfortunately, such strict adherence to the prescribed workflow makes it impossible for the system to adapt to unforeseen circumstances. We firmly believe that the historic trajectory of software development paradigms and IT advancements will establish multiagent systems as the workflow enactment mechanism of the future. In this paper we provide a critical survey of workflow, workflow description languages, web services and agent technologies. We propose that workflow description languages and their associated design tools can be used to specify a multiagent system. Specifically, we advance the idea that the Business Process Execution Language for Web Services (BPEL4WS) can be used as a specification language for expressing the initial social order of the multiagent system, which can then intelligently adapt to changing environmental conditions.

38. Conjectural Equilibrium in Multiagent Learning

@article{306216,
author = {Michael P. Wellman and Junling Hu},
title = {Conjectural Equilibrium in Multiagent Learning},
journal = {Mach. Learn.},
volume = {33},
number = {2-3},
year = {1998},
issn = {0885-6125},
pages = {179--200},
publisher = {Kluwer Academic Publishers},
address = {Hingham, MA, USA},
}

Abstract

Learning in a multiagent environment is complicated by the fact that as other agents learn, the environment effectively changes. Moreover, other agents‘ actions are often not directly observable, and the actions taken by the learning agent can strongly bias which range of behaviors are encountered. We define the concept of a conjectural equilibrium, where all agents‘ expectations are realized, and each agent responds optimally to its expectations. We present a generic multiagent exchange situation, in which competitive behavior constitutes a conjectural equilibrium. We then introduce an agent that executes a more sophisticated strategic learning strategy, building a model of the response of other agents. We find that the system reliably converges to a conjectural equilibrium, but that the final result achieved is highly sensitive to initial belief. In essence, the strategic learner‘s actions tend to fulfill its expectations. Depending on the starting point, the agent may be better or worse off than had it not attempted to learn a model of the other agents at all.

39. Multiagent systems engineering of organization-based multiagent systems

@inproceedings{1082967,
author = {Scott A. DeLoach},
title = {Multiagent systems engineering of organization-based multiagent systems},
booktitle = {SELMAS '05: Proceedings of the fourth international workshop on Software engineering for large-scale multi-agent systems},
year = {2005},
isbn = {1-59593-116-3},

pages = {1--7},
location = {St. Louis, Missouri},
doi = {http://doi.acm.org/10.1145/1082960.1082967},
publisher = {ACM Press},
address = {New York, NY, USA},
}
Abstract
In this paper, we examine the Multiagent Systems Engineering (MaSE) methodology and its applicability to developing organization-based multiagent systems, which are especially relevant to context aware systems. We discuss the inherent shortcomings of MaSE and then present our approach to modeling the concepts required for organizations including goals, roles, agents, capabilities, and the assignment of agents to roles. Finally, we extend MaSE to allow it to overcome its inherent shortcomings and capture the organizational concepts defined in our organizational model.

40. Learning Coordination Strategies for Cooperative Multiagent Systems
@article{306213,
author = {F. Ho and M. Kamel},
title = {Learning Coordination Strategies for Cooperative Multiagent Systems},
journal = {Mach. Learn.},
volume = {33},
number = {2-3},
year = {1998},
issn = {0885-6125},
pages = {155--177},
publisher = {Kluwer Academic Publishers},
address = {Hingham, MA, USA},
}
Abstract
A central issue in the design of cooperative multiagent systems is how to coordinate the behavior of the agents to meet the goals of the designer. Traditionally, this had been accomplished by hand-coding the coordination strategies. However, this task is complex due to the interactions that can take place among agents. Recent work in the area has focused on how strategies can be learned. Yet, many of these systems suffer from convergence, complexity and performance problems. This paper presents a new approach for learning multiagent coordination strategies that addresses these issues. The effectiveness of the technique is demonstrated using a synthetic domain and the predator and prey pursuit problem.

41. Organizing Multiagent Systems
@article{1090752,
author = {Javier V\&\#225;zquez-Salceda and Virginia Dignum and Frank Dignum},
title = {Organizing Multiagent Systems},
journal = {Autonomous Agents and Multi-Agent Systems},
volume = {11},
number = {3},
year = {2005},
issn = {1387-2532},
pages = {307--360},

doi = {http://dx.doi.org/10.1007/s10458-005-1673-9},
publisher = {Kluwer Academic Publishers},
address = {Hingham, MA, USA},
}

Abstract
Despite all the research done in the last years on the development of methodologies for designing MAS, there is no methodology suitable for the specification and design of MAS in complex domains where both the agent view and the organizational view can be modeled. Current multiagent approaches either take a centralist, static approach to organizational design or take an emergent view in which agent interactions are not pre-determined, thus making it impossible to make any predictions on the behavior of the whole systems. Most of them also lack a model of the norms in the environment that should rule the (emergent) behavior of the agent society as a whole and/or the actions of individuals. In this paper, we propose a framework for modeling agent organizations, Organizational Model for Normative Institutions (OMNI), that allows the balance of global organizational requirements with the autonomy of individual agents. It specifies global goals of the system independently from those of the specific agents that populate the system. Both the norms that regulate interaction between agents, as well as the contextual meaning of those interactions are important aspects when specifying the organizational structure.

42. Multiagent Planning as Control Síntesis

@inproceedings{1018851,
author = {Kiam-Tian Seow and Chuan Ma and Makoto Yokoo},
title = {Multiagent Planning as Control Synthesis},
booktitle = {AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems},
year = {2004},
isbn = {1-58113-864-4},
pages = {972--979},
location = {New York, New York},
doi = {http://dx.doi.org/10.1109/AAMAS.2004.188},
publisher = {IEEE Computer Society},
address = {Washington, DC, USA},
}

Abstract
This paper proposes a new multiagent planning approach to coordination synthesis that views distributed agents as discrete-event processes. The connection between discrete-event control synthesis and coordination planning is first established, thereby enabling the exploitation of the vast body of knowledge and associated software synthesis tools from ýThe Supervisory Control of Discrete-Event Systemsý for automatic coordination synthesis of distributed agents. Importantly, these coordinating agents designed collectively generate a behaviour guaranteed not to contradict any specified inter-agent constraint, is nonblocking and optimal. A simple planning methodology is proposed in terms of procedures supported by CTCT, an existing, freely available design tool developed based on the control synthesis framework.Asimple example illustrates the use of the CTCT-based methodology to synthesize coordination modules for distributed agents. Discussions in relation to

previous work examine the relative significance of the new multiagent planning framework.

43. Performance models for large scale multiagent systems: using distributed POMDP building blocks

@inproceedings{860623,
author = {Hyuckchul Jung and Milind Tambe},
title = {Performance models for large scale multiagent systems: using distributed POMDP building blocks},
booktitle = {AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems},
year = {2003},
isbn = {1-58113-683-8},
pages = {297--304},
location = {Melbourne, Australia},
doi = {http://doi.acm.org/10.1145/860575.860623},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

Given a large group of cooperative agents, selecting the right coordination or conflict resolution strategy can have a significant impact on their performance (e.g., speed of convergence). While performance models of such coordination or conflict resolution strategies could aid in selecting the right strategy for a given domain, such models remain largely uninvestigated in the multiagent literature. This paper takes a step towards applying the recently emerging distributed POMDP (partially observable Markov decision process) frameworks, such as MTDP (Markov team decision process), in service of creating such performance models. To address issues of scale-up, we use small-scale models, called building blocks that represent the local interaction among a small group of agents. We discuss several ways to combine building blocks for performance prediction of a larger-scale multiagent system.We present our approach in the context of DCSPs (distributed constraint satisfaction problems), where we first show that there is a large bank of conflict resolution strategies and no strategy dominates all others across different domains. By modeling and combining building blocks, we are able to predict the performance of five different DCSP strategies for four different domain settings, for a large-scale multiagent system. Our approach thus points the way to new tools for strategy analysis and performance modeling in multiagent systems in general.

44. Multiagent Modeling and Simulation of Hydraulic Management of the Camargue

@article{1068006,
author = {Bernard Espinasse and Nathalie Franchesquin},
title = {Multiagent Modeling and Simulation of Hydraulic Management of the Camargue},
journal = {Simulation},
volume = {81},
number = {3},
year = {2005},
issn = {0037-5497},

pages = {201--221},
doi = {http://dx.doi.org/10.1177/003754970505317l},
publisher = {Society for Computer Simulation International},
address = {San Diego, CA, USA},
}
Abstract
Modeling and simulation of human-influenced ecosystems require the integration of natural and decision-making processes. In such ecosystems, natural resource management aims at protecting natural areas while enabling human activities that contribute to these area characteristics. The object of this research is the modeling and simulation of the hydraulic management of the Camargue ecosystem with a multiagent system. The authors first present the problematics of this hydraulic management, its objectives, and a formalization of the different decision processes centered on a social contract and associated with two main phases: the contract elaboration phase and the contract realization phase. They present conceptual modeling and multiagent modeling of these two phases with details on agent models, agent behaviors, and negotiation models.

45. Planning, learning and coordination in multiagent decision processes
@inproceedings{1029710,
author = {Craig Boutilier},
title = {Planning, learning and coordination in multiagent decision processes},
booktitle = {TARK '96: Proceedings of the 6th conference on Theoretical aspects of rationality and knowledge},
year = {1996},
isbn = {1-55860-417-9},
pages = {195--210},
location = {The Netherlands},
publisher = {Morgan Kaufmann Publishers Inc.},
address = {San Francisco, CA, USA},
}
Abstract
There has been a growing interest in AI in the design of multiagent systems, especially in multiagent cooperative planning. In this paper, we investigate the extent to which methods from single-agent planning and learning can be applied in multiagent settings. We survey a number of different techniques from decision-theoretic planning and reinforcement learning and describe a number of interesting issues that arise with regard to coordinating the policies of individual agents. To this end, we describe multiagent Markov decision processes as a general model in which to frame this discussion. These are special n-person cooperative games in which agents share the same utility function. We discuss coordination mechanisms based on imposed conventions (or social laws) as well as learning methods for coordination. Our focus is on the decomposition of sequential decision processes so that coordination can be learned (or imposed) locally, at the level of individual states. We also discuss the use of structured problem representations and their role in the generalization of learned conventions and in approximation.

46. Multiagent coordination by Extended Markov Tracking
@inproceedings{1082539,

author = {Zinovi Rabinovich and Jeffrey S. Rosenschein},
title = {Multiagent coordination by Extended Markov Tracking},
booktitle = {AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems},
year = {2005},
isbn = {1-59593-093-0},
pages = {431--438},
location = {The Netherlands},
doi = {http://doi.acm.org/10.1145/1082473.1082539},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract
We present here Extended Markov Tracking (EMT), a computationally tractable method for the online estimation of Markovian system dynamics, along with experimental support for its successful contribution to a specific control architecture. The control architecture leverages EMT to simultaneously track and correct system dynamics.Using a widespread extension of the Markovian environment model to multiagent systems, we provide an application of EMT-based control to multiagent coordination. The resulting coordinated action algorithm, in contrast to alternative approaches, does not eliminate interference among agents, but rather exploits it for purposes of synchronization and implicit information transfer. This information transfer enables the algorithm to be computationally tractable. Experiments are presented that demonstrate the effectiveness of EMT-based control for multiagent coordination in stochastic environments.

47. Knowledge in multiagent systems: initial configurations and broadcast

@article{359527,
author = {Alessio R. Lomuscio and Ron van der Meyden and Mark Ryan},
title = {Knowledge in multiagent systems: initial configurations and broadcast},
journal = {ACM Trans. Comput. Logic},
volume = {1},
number = {2},
year = {2000},
issn = {1529-3785},
pages = {247--284},
doi = {http://doi.acm.org/10.1145/359496.359527},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract
The semantic framework for the modal logic of knowledge due to Halpern and Moses provides a way to ascribe knowlegde to agents in distributed and multiagent systems. In this paper we study two special cases of this framework: full systems and hypercubes. Both model static situtations in which no agents has any information about another agent's state. Full systems and hypercubes are an appropriate model for the initial configurations of many systems of interest. We establish a correspondence between full systems and hypercube systems and certain classes of Kripke frames. We show that these classes of systems correspond to the same logic.

Moreover, this logic is also the same as that generated by the larger class of weakly directed frames. We provide a sound and complete axiomatization, S5WDn of this logic, and study its computational complexity. Finally, we show that under certain natural assumptions, in a model where knowledge evolves over time, S5WDn characteristics the properties of knowledge not just at the initial configuration, but also at all later configurations. In this particular, this holds for homogeneous broadcast systems, which capture settings in which agents are intially ignorant of each others local states, operate synchronously, have perfect recall, and can communicate only by broadcasting.

48. Best-Response Multiagent Learning in Non-Stationary Environments

@inproceedings{1018798,
author = {Michael Weinberg and Jeffrey S. Rosenschein},
title = {Best-Response Multiagent Learning in Non-Stationary Environments},
booktitle = {AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems},
year = {2004},
isbn = {1-58113-864-4},
pages = {506--513},
location = {New York, New York},
doi = {http://dx.doi.org/10.1109/AAMAS.2004.75},
publisher = {IEEE Computer Society},
address = {Washington, DC, USA},
}

Abstract

This paper investigates a relatively new direction in Multiagent Reinforcement Learning. Most multiagent learning techniques focus on Nash equilibria as elements of both the learning algorithm and its evaluation criteria. In contrast, we propose a multiagent learning algorithm that is optimal in the sense of finding a best-response policy, rather than in reaching an equilibrium. We present the first learning algorithm that is provably optimal against restricted classes of non-stationary opponents. The algorithm infers an accurate model of the opponentýs non-stationary strategy, and simultaneously creates a best-response policy against that strategy. Our learning algorithm works within the very general framework of n-player, general-sum stochastic games, and learns both the game structure and its associated optimal policy.

49. Self-stabilization as multiagent systems property

@inproceedings{545144,
author = {Petra Funk and Ingo Zinnikus},
title = {Self-stabilization as multiagent systems property},
booktitle = {AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems},
year = {2002},
isbn = {1-58113-480-0},
pages = {1413--1414},
location = {Bologna, Italy},
doi = {http://doi.acm.org/10.1145/545056.545144},
publisher = {ACM Press},
address = {New York, NY, USA},

}
Abstract
Self-Stabilization is an important concept for distributed computing and communication networks. It describes a system's ability to recover automatically from unexpected failure. It is also an important issue for multiagent systems, as they are distributed and communicative systems. Therefore, we investigate self-stabilization for multiagent systems and how it can be achieved. We define the functionality of an agent society and the task to be achieved in terms of roles which agents may fulfill. Self-Stabilization can thus be achieved by ensuring role adopting and role fulfillment within an agent society.

50. Online learning about other agents in a dynamic multiagent system

@inproceedings{280839,
author = {Junling Hu and Michael P. Wellman},
title = {Online learning about other agents in a dynamic multiagent system},
booktitle = {AGENTS '98: Proceedings of the second international conference on Autonomous agents},
year = {1998},
isbn = {0-89791-983-1},
pages = {239--246},
location = {Minneapolis, Minnesota, United States},
doi = {http://doi.acm.org/10.1145/280765.280839},
publisher = {ACM Press},
address = {New York, NY, USA},
}

51. Issues in Multiagent System Development

@inproceedings{1018845,
author = {Mehdi Dastani and Joris Hulstijn and Frank Dignum and John-Jules Ch. Meyer},
title = {Issues in Multiagent System Development},
booktitle = {AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems},
year = {2004},
isbn = {1-58113-864-4},
pages = {922--929},
location = {New York, New York},
doi = {http://dx.doi.org/10.1109/AAMAS.2004.156},
publisher = {IEEE Computer Society},
address = {Washington, DC, USA},
}
Abstract
Methodologies for multiagent system development should assist the developer in making decisions about those aspects of the analysis, design and implementation, that are crucial for multiagent systems, namely, social and cognitive concepts (e.g. norms and goals). In this paper, we review existing agent-oriented methodologies. We conclude that there is a big gap between the analysis and design models and the implementation. We identify some open issues for multiagent system development. We introduce our vision of a development methodology for multiagent systems,

based on the OperA analysis models and the agent-oriented programming language 3APL.

52. A case for adaptation of the distributed environment layout in multiagent applications

@inproceedings{1082974,
author = {Koenraad Mertens and Tom Holvoet and Yolande Berbers},
title = {A case for adaptation of the distributed environment layout in multiagent applications},
booktitle = {SELMAS '05: Proceedings of the fourth international workshop on Software engineering for large-scale multi-agent systems},
year = {2005},
isbn = {1-59593-116-3},
pages = {1--8},
location = {St. Louis, Missouri},
doi = {http://doi.acm.org/10.1145/1082960.1082974},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

A multiagent application consists of an environment and a number of agents. The environment contains information that the agents use and manipulate to do their work. When a multiagent system is distributed over a number of different hosts (i.e. more than one execution platform is used), the environment has to be distributed as well. The distributed layout of the environment can influence the performance of agents and of the system.In this paper we discuss when a distributed environment can dynamically change its distribution layout. Our focus is on a distributed environment in which mobile agents move around and are aware of the distributed nature of the system.Changes to the layout of the distribution are not only triggered by the agents (like other, application-specific actions), but they can also be triggered by external events, and the environment itself has the ability to pro-actively change its distribution layout over the different hosts when this improves the behavior and efficiency of the application.Using a specific application (solving distributed constraint satisfaction problems) as an example, we indicate the usefulness of changes to the distribution layout and how they can be incorporated easily into a multiagent application design. It turns out that for some problems, the improvement in efficiency can be more than 30%.

53. Using game days to teach a multiagent system class

@inproceedings{971378,
author = {Leen-Kiat Soh},
title = {Using game days to teach a multiagent system class},
booktitle = {SIGCSE '04: Proceedings of the 35th SIGCSE technical symposium on Computer science education},
year = {2004},
isbn = {1-58113-798-2},
pages = {219--223},
location = {Norfolk, Virginia, USA},
doi = {http://doi.acm.org/10.1145/971300.971378},
publisher = {ACM Press},

address = {New York, NY, USA},
}
Abstract
Multiagent systems is an attractive problem solving approach that is becoming ever more feasible and popular in today's world. It combines artificial intelligence (AI) and distributed problem solving to allow designers (programmers and engineers alike) to solve problems otherwise deemed awkward in traditional approaches that are less flexible and centralized. In the Fall semester of 2002, I introduced a new game-based technique to my Multiagent Systems class. The class was aimed for seniors (with special permission) and graduate students in Computer Science, covering some breadth and depth of issues in multiagent systems. One of the requirements was participation in four Game Days. On each Game Day, student teams competed against each other in games related to issues such as auction, task allocation, coalition formation, and negotiation. This article documents my designs of and lessons learned from these Game Days. The Game Days were very successful. Through role-playing, the students were motivated and learned about multiagent systems.

54. Bumping strategies for the multiagent agreement problem
@inproceedings{1082533,
author = {Pragnesh Jay Modi and Manuela Veloso},
title = {Bumping strategies for the multiagent agreement problem},
booktitle = {AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems},
year = {2005},
isbn = {1-59593-093-0},
pages = {390--396},
location = {The Netherlands},
doi = {http://doi.acm.org/10.1145/1082473.1082533},
publisher = {ACM Press},
address = {New York, NY, USA},
}
Abstract
We introduce the Multiagent Agreement Problem (MAP) to represent a class of multiagent scheduling problems. MAP is based on the Distributed Constraint Reasoning (DCR) paradigm and requires agents to choose values for variables to satisfy not only their own constraints, but also equality constraints with other agents. The goal is to represent problems in which agents must agree on scheduling decisions, for example, to agree on the start time of a meeting. We investigate a challenging class of MAP - private, incremental MAP (piMAP) in which agents do incremental scheduling of activities and there exist privacy restrictions on information exchange. We investigate a range of strategies for piMAP, called "bumping" strategies. We empirically evaluate these strategies in the domain of calendar management where a personal assistant agent must schedule meetings on behalf of its human user. Our results show that bumping decisions based on scheduling difficulty models of other agents can significantly improve performance over simpler bumping strategies.

55. Multiagent teamwork: analyzing the optimality and complexity of key theories and models

@inproceedings{544946,
author = {David V. Pynadath and Milind Tambe},
title = {Multiagent teamwork: analyzing the optimality and complexity of key theories and models},
booktitle = {AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems},
year = {2002},
isbn = {1-58113-480-0},
pages = {873--880},
location = {Bologna, Italy},
doi = {http://doi.acm.org/10.1145/544862.544946},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract
Despite the significant progress in multiagent teamwork, existing research does not address the optimality of its prescriptions nor the complexity of the teamwork problem. Thus, we cannot determine whether the assumptions and approximations made by a particular theory gain enough efficiency to justify the losses in overall performance. To provide a tool for evaluating this tradeoff, we present a unified framework, the COMmunicative Multiagent Team Decision Problem (COM-MTDP) model, which is general enough to subsume many existing models of multiagent systems. We analyze use the COM-MTDP model to provide a breakdown of the computational complexity of constructing optimal teams under problem domains divided along the dimensions of observability and communication cost. We then exploit the COM-MTDP's ability to encode existing teamwork theories and models to encode two instantiations of joint intentions theory, including STEAM. We then derive a domain-independent criterion for optimal communication and provide a comparative analysis of the two joint intentions instantiations. We have implemented a reusable, domain-independent software package based COM-MTDPs to analyze teamwork coordination strategies, and we demonstrate its use by encoding and evaluating the two joint intentions strategies within an example domain.

56. Collaborative Multiagent Adaptation for Business Environmental Scanning Through the Internet

@article{966607,
author = {Rey-Long Liu},
title = {Collaborative Multiagent Adaptation for Business Environmental Scanning Through the Internet},
journal = {Applied Intelligence},
volume = {20},
number = {2},
year = {2004},
issn = {0924-669X},
pages = {119--133},
doi = {http://dx.doi.org/10.1023/B:APIN.0000013335.15447.b4},
publisher = {Kluwer Academic Publishers},
address = {Hingham, MA, USA},
}

Abstract

Businesses should promptly respond to their dynamic environments. Environmental scanners are thus essential for the businesses to discover and monitor environmental information of interest (IOI). In this paper, we explore user-centered, continuous and resource-bounded environmental scanning (UCRES). Upon receiving information preferences of managers, new IOI should be continuously detected in a timely and complete manner without consuming too much resource (e.g. bandwidths of computer networks and services of information servers). We develop a multiagent framework AESA to tackle the challenges of UCRES. Each agent is a simple entity. All agents collaboratively adapt their population and resource consumption to several dynamic aspects of UCRES: information preferences of individual users, resource limitation of environmental scanning, distribution of IOI in the environments, and update behaviors of the IOI. The delivery of AESA to businesses may constantly provide a larger amount of important and timely IOI without exhausting the Intranet and the Internet communities.

57. A multiagent approach for musical interactive systems

@inproceedings{860669,
author = {Rodolfo Daniel Wulfhorst and Lauro Nakayama and Rosa Maria Vicari},
title = {A multiagent approach for musical interactive systems},
booktitle = {AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems},
year = {2003},
isbn = {1-58113-683-8},
pages = {584--591},
location = {Melbourne, Australia},
doi = {http://doi.acm.org/10.1145/860575.860669},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

The present work proposes a model for a Multiagent System capable to deal with music. Through such a Musical Multiagent System many Computer Music problems may appear. This proposal is based on a community of agents that interact through musical events (MIDI), simulating the behavior of a musical group. The result is the implementation of an intelligent accompaniment system, where agents "listen" and interact with each other. Agents with musical cognition were able to play their instruments with synchronism and satisfy their internal goals.

58. JAFMAS: a multiagent application development system

@inproceedings{280782,
author = {Deepika Chauhan and Albert D. Baker},
title = {JAFMAS: a multiagent application development system},
booktitle = {AGENTS '98: Proceedings of the second international conference on Autonomous agents},
year = {1998},
isbn = {0-89791-983-1},
pages = {100--107},
location = {Minneapolis, Minnesota, United States},
doi = {http://doi.acm.org/10.1145/280765.280782},
publisher = {ACM Press},

address = {New York, NY, USA},
}

59. Partitioned multiagent systems in information oriented domains

@inproceedings{301155,
author = {Claudia V. Goldman and Jeffrey S. Rosenschein},
title = {Partitioned multiagent systems in information oriented domains},
booktitle = {AGENTS '99: Proceedings of the third annual conference on Autonomous Agents},
year = {1999},
isbn = {1-58113-066-X},
pages = {32--39},
location = {Seattle, Washington, United States},
doi = {http://doi.acm.org/10.1145/301136.301155},
publisher = {ACM Press},
address = {New York, NY, USA},
}

60. Tools and techniques for performance measurement of large distributed multiagent systems

@inproceedings{860711,
author = {Aaron Helsinger and Richard Lazarus and William Wright and John Zinky},
title = {Tools and techniques for performance measurement of large distributed multiagent systems},
booktitle = {AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems},
year = {2003},
isbn = {1-58113-683-8},
pages = {843--850},
location = {Melbourne, Australia},
doi = {http://doi.acm.org/10.1145/860575.860711},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

Performance measurement of large distributed multiagent systems (MAS) offers challenges that must be addressed explicitly in the agent infrastructure. Performance data is widely distributed and voluminous, and poor data collection can impact the operation of the system itself. However, performance metrics are essential to internal system function, e.g., autonomous adaptation to dynamic environments, as well as to external assessment. In this paper we describe the tools, techniques, and results of performance characterization of the Cougaar distributed agent architecture. These techniques include infrastructure instrumentation, plugin-based instrumentation of agents, and dynamic control of metric collection. We introduce multiple redundant "channels" for metric delivery, each serving separate quality of service requirements. We present our techniques for instrumenting the agent society, justify the metrics chosen, and describe the tools developed for collecting these metrics. We also present results from distributed agent societies comprising hundreds of agents.

61. Addressing the complexity of patient monitoring with multiagents and modular logic

@inproceedings{860794,
author = {Susan L. Mabry and Timothy Etters and Troy Schneringer and Naomi Edwards and Caleb Hug},
title = {Addressing the complexity of patient monitoring with multiagents and modular logic},
booktitle = {AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems},
year = {2003},
isbn = {1-58113-683-8},
pages = {1058--1059},
location = {Melbourne, Australia},
doi = {http://doi.acm.org/10.1145/860575.860794},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

While there is a recognized need among medical clinicians for assistance in monitoring and diagnostic tasks, the domain is extremely complex. Teams of distributed agents and interchangeable logic modules can address practical issues in a fielding of MAS technology. IM-Agents employ dynamic inference through a hybrid approach of fuzzy logic, connection networks, trend analysis and qualitative logic. Decision functionality in test cases involves emergency trauma scenarios, with particular focus on stabilizing hemorrhagic shock.

62. Intelligent agents for patient monitoring and diagnostics

@inproceedings{952585,
author = {Susan L. Mabry and Troy Schneringer and Timothy Etters and Naomi Edwards},
title = {Intelligent agents for patient monitoring and diagnostics},
booktitle = {SAC '03: Proceedings of the 2003 ACM symposium on Applied computing},
year = {2003},
isbn = {1-58113-624-2},
pages = {257--262},
location = {Melbourne, Florida},
doi = {http://doi.acm.org/10.1145/952532.952585},
publisher = {ACM Press},
address = {New York, NY, USA},
}

This paper describes Intelligent Monitor Agents (IM-Agents), a project that utilizes intelligent multiagents to assist patient health care. Clinicians must simultaneously audit and interpret an overwhelming amount of information regarding a patient during the process of monitoring, performing diagnostics and determining therapeutic intervention. Multiple IM-Agents coordinate as a team with each agent performing specialized monitor and diagnostic tasks. Fuzzy logic, connection networks, trend analysis and qualitative logic methods are used within the dynamic decision modules of all IM-Agents. The agents are autonomous, interactive, mobile and capable of performing dynamic intelligent inference during execution. Decision

functionality is demonstrated in preliminary prototype test cases involving emergency trauma scenarios, particularly focusing on stabilizing hemorrhagic shock.

63. A first step towards providing health-care agent-based services to mobile users

@inproceedings{544883,
author = {Antonio Moreno and David Isern},
title = {A first step towards providing health-care agent-based services to mobile users},
booktitle = {AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems},
year = {2002},
isbn = {1-58113-480-0},
pages = {589--590},
location = {Bologna, Italy},
doi = {http://doi.acm.org/10.1145/544862.544883},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

The aim of the Agent Cities project is to build agent-based information services for mobile users. In the framework of this project it is interesting to consider the possibility of deploying agents that provide Health Care-related services. In this paper we report the design and construction of a multi-agent system (MAS) that is composed of agents that provide medical services. The MAS contains agents that allow the user to search for medical centres satisfying a given set of requirements, to access his/her medical record or to make a booking to be visited by a particular kind of doctor.

64. A Multi-agent system architecture for monitoring medical protocols

@inproceedings{337580,
author = {T. Alsinet and R. B\&\#233;jar and C. Fernanadez and F. Many\&\#224;},
title = {A Multi-agent system architecture for monitoring medical protocols},
booktitle = {AGENTS '00: Proceedings of the fourth international conference on Autonomous agents},
year = {2000},
isbn = {1-58113-230-1},
pages = {499--505},
location = {Barcelona, Spain},
doi = {http://doi.acm.org/10.1145/336595.337580},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

The control of the right application of medical protocols is a key issue in hospital environments. For the automated monitoring of medical protocols, we need a domain-independent language for their representation and a fully, or semi, autonomous system that understands the protocols and supervises their application. In this paper we describe a specification language and a multi-agent system architecture for monitoring medical protocols. We model medical services in hospital

environments as specialized domain agents and interpret a medical protocol as a negotiation process between agents. A medical service can be involved in multiple medical protocols, and so specialized domain agents are independent of negotiation processes and autonomous system agents perform monitoring tasks. We present the detailed architecture of the system agents and of an important domain agent, the database broker agent, that is responsible of obtaining relevant information about the clinical history of patients. We also describe how we tackle the problems of privacy, integrity and authentication during the process of exchanging information between agents.

65. Agent-oriented context-aware platforms supporting communities of practice in health care

@inproceedings{1082736,
author = {Luiz Olavo Bonino da Silva Santos and Renata S. S. Guizzardi and Marten van Sinderen},
title = {Agent-oriented context-aware platforms supporting communities of practice in health care},
booktitle = {AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems},
year = {2005},
isbn = {1-59593-093-0},
pages = {1287--1288},
location = {The Netherlands},
doi = {http://doi.acm.org/10.1145/1082473.1082736},
publisher = {ACM Press},
address = {New York, NY, USA},
}

Abstract

This paper presents and discusses the use of an agent-oriented context-aware platform to support communities of practice (CoPs) in the health care domain. Our work is based on a scenario where CoPs are applied in a hospital to enhance the knowledge sharing among the staff members who share interests and goals. Here, we test the support of an agent-oriented modeling language (AORML) for the analysis of the proposed application for the test case scenario.

66. An agent based framework for virtual medical devices

@inproceedings{544896,
author = {Zeljko Obrenovic and Dusan Starcevic and Emil Jovanov and Vlada Radivojevic},
title = {An agent based framework for virtual medical devices},
booktitle = {AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems},
year = {2002},
isbn = {1-58113-480-0},
pages = {659--660},
location = {Bologna, Italy},
doi = {http://doi.acm.org/10.1145/544862.544896},
publisher = {ACM Press},
address = {New York, NY, USA},
}

In this paper we present the telemedical environment based on VMDs implemented with Java mobile agent technology, called aglets. The agent based VMD implementation provides ad-hoc agent interaction, support for mobile agents and different user interface components in the telemedical system. We have developed a VMD agent framework with four types of agents: data agents, processing agents, presentation agents, and monitoring agents. Data agents abstract data source, creating uniform view on different types of data, independent of data acquisition device. Processing agents produce derived data, such us FFT power spectrum, from raw data provided by the data agents. Presentation agents supply user interface components using a variety of user data views. User interface components are based on HTTP, SMS and WAP protocols. Monitoring agents collaborate with data and processing agents providing support for data mining operations, and search for relevant patterns. Typical example is monitoring for possible epileptic attacks. We have applied VMDs to facilitate distributed EEG analysis. We have found that the flexibility of distributed agent architecture is well suited for the telemedical application domain. This flexibility is particularly important in the case of an emergency, enabling swift system reconfiguration on the fly.

67. Agent-Based Modelling: A Case Study in HIV Epidemic

@inproceedings{1053958,
author = {Eyob Teweldemedhin and Tshilidzi Marwala and Conrad Mueller},
title = {Agent-Based Modelling: A Case Study in HIV Epidemic},
booktitle = {HIS '04: Proceedings of the Fourth International Conference on Hybrid Intelligent Systems (HIS'04)},
year = {2004},
isbn = {0-7695-2291-2},
pages = {154--159},
doi = {http://dx.doi.org/10.1109/ICHIS.2004.16},
publisher = {IEEE Computer Society},
address = {Washington, DC, USA},
}

Abstract

This research presents an agent-based, bottom-up modelling approach to develop a simulation tool for estimating and predicting the spread of the Human Immunodeficiency Virus (HIV) in a given population. HIV is mainly a sexually transmitted disease (STD) causing a serious problem to human health. The virus is transmitted from an infected person to another who was previously healthy through different biological, social and environmental factors. The research develops the simulation tool by modelling these factors by agents. Although research has and is being conducted to estimate and predict the spread of the HIV epidemic, the proposed research seeks to investigate the spread using a different approach. The previous models used a top-down modelling approach. They are built from the general characteristics and behaviours of the population. They have not explored the potential use of agent technology. This research attempts to investigate the flexibility that the multiagent system offers. Agent-based models are close to the situations that exist in a given real system that consists of autonomous components interacting with each other. The modelling approach has the advantage of observing the interaction made between agents, which is a difficult task in the top-down modelling approach. The research investigates the performance of the tool and presents the first results obtained.

68. A Multiagent Architecture for Developing Medical Information Retrieval Agents

@article{940863,
author = {Steven Walczak},
title = {A Multiagent Architecture for Developing Medical Information Retrieval Agents},
journal = {J. Med. Syst.},
volume = {27},
number = {5},
year = {2003},
issn = {0148-5598},
pages = {479--498},
doi = {http://dx.doi.org/10.1023/A:1025668124244},
publisher = {Plenum Press},
address = {New York, NY, USA},
}

Abstract

Information that is available on the world wide web (WWW) is already more vast than can be comprehensibly studied by individuals and this quantity is increasing at a staggering pace. The quality of service delivered by physicians is dependent on the availability of current information. The agent paradigm offers a means for enabling physicians to filter information and retrieve only information that is relevant to current patient treatments. As with many specialized domains, agent-based information retrieval in medical domains must satisfy several domain-dependent constraints. A multiple agent architecture is developed and described in detail to efficiently provide agent-based information retrieval from the WWW and other explicit information resources. A simulation of the proposed multiple agent architecture shows a 97% decrease in information overload and an 85% increase in information relevancy over existing meta-search tools (with even larger gains over standard search engines).

69. Multi-Agent Autonomic Architecture and Its Application in E-Medicine

@inproceedings{946376,
author = {Huaglory Tianfield},
title = {Multi-Agent Autonomic Architecture and Its Application in E- Medicine},
booktitle = {IAT '03: Proceedings of the IEEE/WIC International Conference on Intelligent Agent Technology},
year = {2003},
isbn = {0-7695-1931-8},
pages = {601},
publisher = {IEEE Computer Society},
address = {Washington, DC, USA},
}

Abstract

Autonomic computing is an entirely new philosophy for the development of computing systems. This paper first presents the basic concept and principle of autonomic computing. Then, an architecture is proposed for autonomic computing

systems. Finally, design issues of amulti-agent autonomic architecture and the collaborative environment for e-medicine are presented.

70. Bioinformatics integration and agent technology

@article{1022384,
author = {K. A. Karasavvas and R. Baldock and A. Burger},
title = {Bioinformatics integration and agent technology},
journal = {J. of Biomedical Informatics},
volume = {37},
number = {3},
year = {2004},
issn = {1532-0464},
pages = {205--219},
doi = {http://dx.doi.org/10.1016/j.jbi.2004.04.003},
publisher = {Elsevier Science},
address = {San Diego, USA},
}

Vast amounts of life sciences data are scattered around the world in the form of a variety of heterogeneous data sources. The need to be able to co-relate relevant information is fundamental to increase the overall knowledge and understanding of a specific subject. Bioinformaticians aspire to find ways to integrate biological data sources for this purpose and system integration is a very important research topic. The purpose of this paper is to provide an overview of important integration issues that should be considered when designing a bioinformatics integration system. The currently prevailing approach for integration is presented with examples of bioinformatics information systems together with their main characteristics. Here, we introduce agent technology and we argue why it provides an appropriate solution for designing bioinformatics integration systems.