

# **Doctoral Mentoring Program**

Sixth International Conference on Autonomous Agents and  
Multiagent Systems (AAMAS 2007)  
<http://www.aamas2007.org>

Honolulu, Hawaii, USA, May 15, 2007

Kate Larson, Iyad Rahwan



## **Preface**

Since its initiation in 2002 as a merger of three successful conferences, AAMAS has become the premier scientific conference for research in autonomous agents and multiagent systems. It provides a single, high-impact, internationally respected archival forum for scientific research in the theory and practice of autonomous agents and multiagent systems. AAMAS-07 is the sixth conference in the AAMAS series, following successful previous conferences in Bologna, Italy (2002), Melbourne, Australia (2003), New York, USA (2004), Utrecht, The Netherlands (2005), and Hakodate, Japan (2006).

One of the key drivers of any scientific field is its young talent; those young researchers who will drive and shape the future of the field. The AAMAS doctoral mentoring program aims to foster and sharpen these talents by providing PhD students with the opportunity to interact closely with established researchers, to receive feedback on their work and to get advice on managing their careers.

In addition to the important role of the mentor, this year, we wanted to emphasize the role of peers in giving advice. To this end, we grouped students into three main groups, and assigned two mentors per group. We encouraged interaction among peers by asking students to review each other's work. We also invited a recent PhD graduate to give a fresh perspective on the process of completing a PhD and starting a career as a young academic.

We look forward to seeing a vibrant and productive interaction on the day. And we eagerly await seeing how the future of the field will be shaped by the doctoral mentoring program's participants.

Kate Larson and Iyad Rahwan  
Waterloo and Dubai, Spring 2007

## Doctoral Mentoring Chairs

### **Kate Larson**

David R. Cheriton School of Computer Science  
University of Waterloo  
Waterloo, Ontario, Canada

### **Iyad Rahwan**

Institute of Informatics	(Fellow) School of Informatics
British University in Dubai	University of Edinburgh
P.O.Box 502216, Dubai	Edinburgh, EH8 9LE
United Arab Emirates	UK

## Invited Speaker

- Sherief Abdallah, British University in Dubai, UAE

## Participants

- Jaesek Ahn, University of Texas, USA
- Reyhan Aydogan, Bogazici University, Turkey
- Brent Eskridge, University of Oklahoma, USA
- Hojjat Ghaderi, University of Toronto, Canada
- Mattijs Ghijsen, University of Amsterdam, The Netherlands
- Andrew Gilpin, Carnegie Mellon University, USA
- Chris Jones, University of Texas, USA
- Sachin Kamboj, University of Delaware, USA
- Kiran Lakkaraju, University of Illinois, Urbana-Champaign, USA
- Raz Lin, Bar Ilan University, Israel
- Antonio, Lopes, ISCTE, Portugal
- Murat Sensoy, Bogazici University, Turkey
- Mandeep Singh, Universitaet Tuebingen, Germany
- Sebastian Stein, University of Southampton, UK
- Matthew Taylor, University of Texas, USA
- Jie Zhang, University of Waterloo, Canada
- Inon Zuckerman, Bar Ilan University, Israel



# Multiagent Team Formation: Role of Attitude and Personality in Decision-Making

Jaesuk Ahn

Laboratory for Intelligent Processes and Systems

University of Texas at Austin

1 University Station C5000, Austin, TX 78712

+1-512-471-5350

jsahn@lips.utexas.edu

## ABSTRACT

In this research, the concepts of attitude and personality trait as well as their roles in agents' decision-making concepts are discussed with regard to job selection and team member selection in multiagent team formation. This research introduces the concept of attitude to guide an agent's behavior and personality trait to model other agents' behavior tendency and identify helpfulness of them given the environmental parameters. It also provides methods to incorporate these attitudes and personality traits in constructing the expected utility function to guide agents in ranking the various alternatives it may have regarding job selection and team member selection. In addition, minority game and reinforcement learning will be applied to alter an agent's attitude to adapt to the changing environments.

## 1. INTRODUCTION

Multiagent systems are increasingly deployed in such problem domains as e-commerce and problem-solving network because agents can form teams to overcome the limitations of individual members [1]. Further, since the prevalent connectivity of the Internet and advances in mobile network technology have created more opportunities for an agent to participate in various problem domains, there is an increasing number of possibilities for the agent to interact with unknown agents to form teams in open environments. In this sense, the ability of the agent to minimize the potential risk due to the other agents' uncertain behavior and ability to respond to changing environmental situation is essential during the decision-making process. When selecting jobs and team members to form a team, therefore, an agent should consider not only matching capabilities of agents, but also uncertain behavior of others, as well as changing environmental conditions.

Accordingly, this paper addresses the problem of the agent's decision-making during the job selection and team member selection phases in the multiagent team formation domain. Normative decision-making theory uses relevant utilities and corresponding probabilities for decision-making process. The utility function aggregates all expected utilities and probabilities into an equation to choose the best option available. In this research, an agent uses this equation as an expected utility function to guide its own behavior in a certain way to maximize its reward while minimizing possible risks. In addition, the notion of attitude has been applied as a kind of norm to dictate individual's behavior pattern across different environments and contexts. Attitude has been defined as the tendency to act in a certain way towards the objects [2]. Among social psychologists, attitude is also known as a good predictor of behavior [2]; thus

attitude can be described as a set of parameters to dictate an individual agent's behavior patterns. The research proposes an attitude model (including attitudes toward proactive behavior, potential risk, and reward) and the method to incorporate an agent's attitudes when constructing expected utility function. Further, the notion of personality trait is defined to model other agents' behavior tendency. The personality models have been used to explain and predict individual behavior across different environment and context [3]. In this sense, personality trait can be served as a way to differentiate other agents and predict one's behavior tendency. Three personality traits (reliability, agreeableness, openness) are proposed based on the Five-Factor model [4]. An agent models the personality traits of other agents based on past interactions and incorporating this model into its own decision process, thereby allowing agents to select and interact with more compatible partners and form a team.

Since agents are often situated in dynamically changing environments, the ability to adapt to the changing environment is important for an agent. In this sense, the next direction of this research is to capture the environmental parameters and an agent's resulting performance and use those factors as part of an evaluation mechanism to alter an agent's attitude. Attitudes are capable of change through the reinforcement or congruity. Given a set of environmental parameters and the previous performance result of an agent, we are planning to develop a learning mechanism to alter an agent's attitudes through the reinforcement, thus allowing an agent to perform well even though the environment has changed.

## 2. ATTITUDE DRIVEN JOB SELECTION (COMPLETE)

The notion of attitude can provide a mechanism to dictate behavioral tendencies of an agent. These attitude parameters are used as a weighted value in agent's utility function to dictate agent's behavior regarding job selection. In this way, attitudes define the priority that an agent places on the various choices it may have regarding job selection. Three attitudes are defined:

- $a_p$ : Attitude toward proactive behavior
  - willingness to be a proactive by being a leader and initiating new jobs to work on. (suggestible vs. assertive)
- $a_{rn}$ : Attitude toward reward
  - willingness to pursue higher rewards (restrained vs. greedy)
- $a_r$ : Attitude toward risk
  - willingness to take risk. (risk-averse vs. risk-blinded)

Attitude toward proactive behavior influence agent's willingness to be a leader, proactively select the jobs, and form a team.

Attitude toward reward and risk are used as a pair to give a priority to either possible reward or potential risk. Given the attitude model, job selection utility (JU) function is used in evaluation phase to rank the available jobs. An agent decides what to work on based on the JU of all the possible choices. The JU of the possible choices are calculated based on their possible consequences including expected payoff and corresponding risks (Eq. 1). Risks are consists of possible penalty for the job failure and variances in the environment (e.g., unassigned tasks in the job, estimated cooperativeness of other agents). Non-leader agents also consider utilization of its capability as a factor of possible reward in addition to the estimated payoff from assigned tasks.

$$JU(\mathbf{j}) = (a_{rw})Reward - (a_r)Risk, \text{ where } \mathbf{j} \text{ is a set of jobs} \quad (1)$$

The preliminary experimental result shows that the changes in the proposed attitudes clearly affect factors for the agents' successful completion of jobs given different sets of environments (e.g. time limit, agent population) [5]. Given the different sets of environments, we are able to see a specific attitude performs better than others. For example, in the environment, where there is limited time to complete jobs, agents with restraining attitude and risk-blinded attitude outperforms the other types. Even if an agent is greedy, the agent is able to perform better by having risk-averse attitude. When the agent population is small or there are fewer proactive agents who want to initiate jobs, agents with assertive attitude can earn more payoffs since there is less competition between agents. However, when competition to get a job among leader agents is high, avoiding the competition and waiting for a job to be offered (being suggestible) is a good strategy.

### 3. PERSONALITY-DRIVEN MEMBER SELECTION (WORK IN PROGRESS)

Although it is difficult to get explicit information about other agents' behavior, it is possible to model and estimate other agents' behavior tendencies by observing their interactions. Such a model will help an agent to identify the most helpful agents in the system to reduce the potential risk of team formation failure. In this sense, personality trait model can represent other agents' behavior tendency so that helps an agent to identify compatible partners for the team. This research proposes three personality traits:

- $p_r$ : Reliability: tendency to fulfill the commitment or be trustworthy
- $p_a$ : Agreeableness: tendency to accept the offer or work together
- $p_o$ : Openness: tendency to interact with anyone or to be free from the previous relationship

$$Helpfulness(agent) = \omega \sum_i \delta_{decay} (observed \text{ personality traits}) \quad (2)$$

Degree of helpfulness is calculated as weighted sum of these three personality trait of other agent (Eq. 2) with time-dependent decay function. Weight values vary according to an agent's own attitudes and time-dependent decay function reflects the timelessness of the information which gives higher priority on the recent information. Given the helpfulness of an agent, team member selection utility function incorporates potential reward (payoff, helpfulness, and utilization) and possible risks (penalty, agent types, and job complexity) and cost associated with it (number of message required to form a team). Similar to JU, an agent's

attitude are used as weight values to give priority either on the reward side or risk side.

## 4. LEARNING ATTITUDE (WORK TO BE DONE)

### 4.1 Role Assignment using Minority Game

Minority game (MG) [6] involves  $N$  agents with bounded rationality forced to make a binary decision without direct interaction among them. At each time step, agents make their decision solely based on the  $m$  most recent outcomes. In MG, agents with minority side only gets payoff. MG gives a simplified model of multiagent system where agents with bounded rationality or limited resources cooperate by taking different roles over time without direct communication. In this sense, MG is suitable for the role assignment problem where agents decide to take either leader or member (non-leader) over time. Attitude toward proactive behavior controls agent's decision to be a leader. Using MG, an agent can adjust the attitude toward proactive behavior and improve their collective behavior in open environment without any form of explicit control over each agent.

### 4.2 Reinforcement Learning

Agents with different attitudes can produce different outcomes given a set of different environmental parameters (e.g., population, types of agents, time limit to complete jobs, payoff, and possible penalty). In this sense, it is desirable to alter agent's attitude depending on the current agent's performance level. A certain attitude can be reinforced to alter agent's current behavior thus improves overall outcome based on the previous performance results. We will incorporate reinforcement learning algorithm to determine the optimal level of each attitude given a set of environmental parameters.

## 5. REFERENCES

- [1] Lerman, K. and Shehory, O. Coalition Formation for Large-Scale Electronic Markets. In Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS-2000), (Boston, MA, USA, July 10-12, 2000), IEEE Computer Society, 167-174.
- [2] Ajzen, I. Nature and Operation of Attitudes. Annual Review of Psychology, 52, 2001, 27-58.
- [3] Talman, S., Hadad, M., Gal, Y. a., and Kraus, S. Adapting to agents' personalities in negotiation. In Proceedings of the 4th international joint conference on Autonomous agents and multiagent systems, (The Netherland, 2005), ACM Press, 383-389.
- [4] Digman, J. M. Personality Structure: Emergence of the Five-Factor Model. Annual Review of Psychology, 41, 1990, 417-440.
- [5] Ahn, J. and Barber, K. S. Role of Attitude in Agent's Decision Process: Job Selection Mechanism in Multiagent Team Formation. TR2007-UT-LIPS-010, The University of Texas at Austin, Technical Report, Austin, 2007.
- [6] Challet, D. and Zhang, Y. C. On the Minority Game: Analytical and Numerical Studies. Physica A, 256, (3-4), 1998, 514-532.

# Content-oriented Composite Service Negotiation in E-commerce

Reyhan Aydoğan  
reyhan.aydogan@boun.edu.tr

Department of Computer Engineering  
Boğaziçi University  
Bebek, Istanbul, Turkey

## ABSTRACT

In online, dynamic environments, the services requested by consumers may not be readily served by the providers. This requires the service consumers and providers to negotiate their service needs and offers. Multiagent negotiation approaches typically assume that the parties agree on service content and focus on finding a consensus on service price. In contrast, this research develops an approach through which the parties can negotiate the content of a service. This calls for a negotiation approach in which the parties can understand the semantics of their requests and offers and learn each other's preferences incrementally over time. Moreover, in some cases, the functionality of the acceptable service that is desired by the consumer cannot be met by just a single service provider. In such cases, there is a need to compose several services to make up the desired service. Therefore, communication and coordination among multiple service providers are required during the negotiation process. As a result, my main research direction involves the topics such as content-oriented negotiation, service composition, AI planning and ontology reasoning for semantics.

## 1. INTRODUCTION

Current approaches to e-commerce treat service price as the primary construct for negotiation by assuming that the service content is fixed [7]. However, negotiation on price presupposes that other properties of the service have already been agreed upon. Nevertheless, many times the service provider may not be offering the exact requested service due to lack of resources, constraints in its business policy, and so on [5]. When this is the case, the producer and the consumer need to negotiate the *content* of the requested service [9].

Most existing negotiation approaches assume that all features of a service are equally important and concentrate on the price [3]. However, in reality not all features may be relevant and the relevance of a feature may vary from consumer to consumer. For instance, completion time of a service may be important for one consumer whereas the quality of the service may be more important for a second consumer. Without doubt, considering the preferences of the consumer has a positive impact on the negotiation process.

For this purpose, evaluation of the service components with different weights can be useful. Some studies take these weights as a priori and uses the fixed weights [6]. On the other hand, mostly the producer does not know the consumer's preferences before the negotiation. Hence, it is more appropriate for the producer to learn these preferences for each consumer.

Furthermore, in some cases, the consumer's request cannot be fulfilled by a single provider since some requests can be complex by their nature. For example, a customer request for getting travel arrangements involves both the hotel reservations and flight schedules. In such cases, negotiation process should be performed among multiple providers. Finding appropriate composite services can be taken as AI planning problem [4]. In addition, using semantic information in both composition and negotiation may help generating various alternative offers; thereby shortening the consensus time.

## 2. COMPLETED RESEARCH

In our completed work [1, 2], we propose an architecture in which both consumers and producers use a shared ontology to negotiate a service. Through repetitive interactions, the provider learns consumers' needs accurately and can make better targeted offers. To enable fast and accurate learning of preferences, we develop an extension to Version Space [8], which is one of the inductive learning approaches that learns concepts from observed examples. Then, we compare our extended Version Space with existing learning techniques. We further develop a metric for measuring semantic similarity between services and compare the performance of our approach using different similarity metrics.

**Preference Learning:** As an alternative to existing negotiation approaches, we propose an architecture in which the service providers learn the relevant features of a service for a particular customer over time. We represent service requests as a vector of service features. We use an ontology in order to capture the relations between services and to construct the features for a given service. By using a common ontology, we enable the consumers and producers to share a common vocabulary for negotiation. The particular service we have used is a wine selling service. The wine seller learns the wine preferences of the customer to sell better targeted wines. The producer models the requests of the consumer and its counter offers to learn which features are more important for the consumer. Since no information is present before the interactions start, the learning algorithm has to be incremental so that it can be trained at run time and can revise itself with each new interaction. For this purpose, we have used inductive learning [8] and we have extended the version space algorithms to take into account the needs of service negotiation context.

**Service Generation:** Even after the producer learns the important

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07, May 14–18 2007, Honolulu, Hawaii.  
Copyright 2007 IFAAMAS .

features for a consumer, it needs a method to generate offers that are the most relevant for the consumer among its set of possible services. In other words, the question is how the producer uses the information that was learned from the dialogues to make the best offer to the consumer. For instance, assume that the producer has learned that the consumer wants to buy a red wine but the producer can only offer rose or white wine. What should the producer's offer contain; white wine or rose wine? If the producer has some domain knowledge about semantic similarity (e.g., knows that the red and rose wines are taste-wise more similar than white wine), then it can generate better offers. However, in addition to domain knowledge, this derivation requires appropriate metrics to measure similarity between available services and learned preferences.

### 3. THE WORK IN PROGRESS

Since a service on its own may not be adequate for the consumer in particular cases, we may need to negotiate *composite* services. Thus, we deal with the composition problem and investigate service composition as a planning problem. For this purpose, we use the input-output parameter relations in order to select the constituent services that make up the composite service. In our approach, we introduce a graph-based composition technique using ontological information. We make use of ontological information between the input-output parameters such that a more specific concept can be used instead of a general concept to make the process more flexible. Our main focus is the relation among input and output parameters. It is assumed that consumer specifies some available input parameters and request some functionalities applied on these parameters to obtain some desired outputs. The aim of our approach is to find the required services to accomplish this task. While finding the required services for the composite service, it is expected to use the additional information that is represented by ontologies such as subclass relations.

Our proposed approach is based on constructing a dependency graph including the service parameters and services themselves. By using this dependency graph, we perform backward chaining as follows. We represent the goal as desired output parameters. We start from these parameter and search backwards to find input parameters. In addition to using semantic information through the search, our approach considers non-functional attributes of the services such as service quality. To consider the quality measures, we find the constituent services by making use of depth first search. After finding the required services, our algorithm generates a plan that shows the execution order of each service.

As far as service composition is concerned, one of the main question is how to obtain the desired output. Therefore, an approach looking for the services that should be called to obtain the desired output is convenient. However, most of the approaches do not consider both semantic information for finding such services and service quality for customer satisfaction. Thus, our approach is more preferable for flexible service composition.

### 4. THE WORK TO BE DONE

In our future work, we plan to construct a negotiation scheme in which there are multiple providers collaborating with each other in order to have a consensus on service content via learning the consumer preferences. Integration of ontology reasoning into the learning algorithm may shorten the consensus time so that hierarchical information can be learned from subsumption hierarchy of relations. Further, by using relationships among features, the producer can discover new knowledge from the existing knowledge.

The preferences of the consumer agent may change during the

negotiation. It would be interesting to extend the learning algorithm to deal with dynamic changes in the requests of the consumer agent. The preferences can be enriched with subtle relations on price. Currently, the producer only tries to learn the consumer's demand without considering its own preferences. Incorporating the business strategies of the producer will allow more realistic scenarios to be tested.

In addition to the content of the service, we should also consider the functionalities of services. For example, there can be some services such that although they give the same inputs and outputs they can vary in functionalities. To address this, we want to focus on the quality of service and cost parameters in detail on services, as well as applying OWL-S so as to make use of preconditions and effects. These are interesting directions that we will pursue in our future work.

### 5. REFERENCES

- [1] R. Aydođan and P. Yolum. Learning consumer preferences for content-oriented negotiation. In *AAMAS Workshop on Business Agents and the Semantic Web (BASeWEB)*, pages 43–52, Hakodate, Japan, May 2006. ACM Press.
- [2] R. Aydođan and P. Yolum. Learning consumer preferences using semantic similarity. In *Proceedings of Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Honolulu, Hawai'i, May 2007.
- [3] L. Busch and I. Horstman. A comment on issue-by-issue negotiations. *Games and Economic Behavior*, 19:144–148, 1997.
- [4] M. Carman, L. Serafini, and P. Traverso. Web service composition as planning. In *ICAPS 2003 Workshop on Planning for Web Services*, 2003.
- [5] J. K. Debenham. Managing e-market negotiation in context with a multiagent system. In *Proceedings Twenty First International Conference on Knowledge Based Systems and Applied Artificial Intelligence, ES'2002: Applications and Innovations in Expert Systems X*, Cambridge, UK, 2002.
- [6] P. Faratin, C. Sierra, and N. R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, 142:205–237, 2002.
- [7] P. Maes, R. H. Guttman, and A. G. Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, 1999.
- [8] T. M. Mitchell. *Machine Learning*. McGraw Hill, New York, 1997.
- [9] M. P. Singh. Value-oriented electronic commerce. *IEEE Internet Computing*, 3(3):6–7, 1999.

# Prioritizing Fuzzy Behavior Hierarchies to Improve Behavior Coordination

Brent E. Eskridge  
Robotics, Evolution, Adaptation and Learning Laboratory (REAL Lab)  
School of Computer Science  
University of Oklahoma  
Norman, Oklahoma 73019-6151, USA  
eskridge@ou.edu

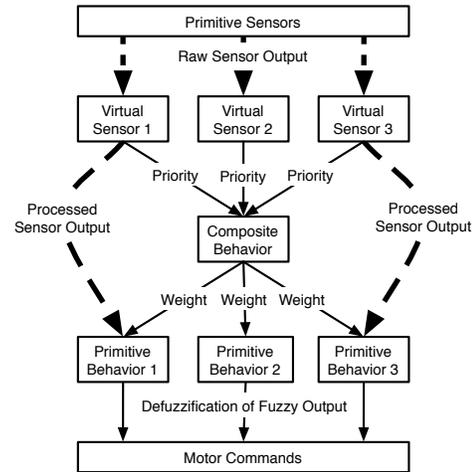
## 1. INTRODUCTION

Behavior-based architectures allow agents to operate real-time in dynamic environments. However, to effectively perform complex tasks, especially those found in multi-agent environments, requires the coordination of a large number of competing behaviors, which can be a difficult proposition. Furthermore, as more behaviors are added, the difficulty of the problem grows exponentially.

This problem has been the focus of much recent research with a wide array of proposed solutions. One approach has been to decompose the control problem into specific tasks [3]. Another approach has been to frame the problem as one of optimization and to apply multiple-objective techniques [4]. While these all address the behavior coordination problem in general, they do not focus on the problem that we are specifically interested in. That is, how can one coordinate a large number of active, competing behaviors in a way that is easy to maintain and develop?

The approach of building a hierarchy of behaviors addresses this problem by abstracting compositions of related behaviors into higher-level, *composite* behaviors [5]. These composite behaviors can then be used to create even higher-level composite behaviors. At each step, only the coordination behaviors at the level immediately below the current level is an issue. The coordination can be further simplified through the use of fuzzy linguistic rules which provide a more intuitive model for implementing both the composite and primitive behaviors. While the combination of these two approaches has addressed the scaling problem, there is still much room for improvement. The difficulty lies in the fact that while the details of a particular behavior are abstracted by the hierarchy, the full information of the agent's state is still required by the hierarchy to coordinate even the highest level composite behaviors.

In an effort to remove this restriction, the idea of a *priority* has been introduced [1]. Priorities provide aggregate abstractions for an agent's low-level state so higher-level composite behaviors only need to consider a few priorities when



**Figure 1: A sample prioritized fuzzy behavior hierarchy. Virtual sensors process raw data and pass priorities, rather than state information, to composite behaviors. Behaviors are implemented as fuzzy rule sets and output either weighting values for other behaviors or control values to be defuzzified. For simplicity, not all sensor output links are shown.**

coordinating lower-level behaviors and not the full state of the agent. Virtual sensors are used to process raw sensor data representing the state of the agent and generate priorities (see Figure 1). These priorities are then used by composite behaviors instead of the agent's state to coordinate lower-level composite behaviors or primitive behaviors. Furthermore, a number of priorities can be used by another virtual sensor to generate a more abstract priority that can be used by an even higher-level composite behavior. Priorities are generated only from the current state of the agent as it pertains to the associated behavior. For example, the priority for the behavior *goal seek* is generated from the relative distance and direction to the goal only. Thus, priorities provide a snapshot of the agent's state from a specific perspective.

The abstractions provided by priorities mean that composite behaviors do not require a significant amount of raw sensor information to coordinate behaviors and can instead use a few priorities. For example, consider the problem of coordinating the primitive behaviors *goal seek* and *avoid obstacle*. In current behavior hierarchies [5], the composite be-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.  
Copyright 2007 IFAAMAS .

havior would need to use the direction and distance to the goal along with the direction and distance to the nearest obstacle to determine how to weight each of the primitive behaviors. However, with priorities, the composite behavior only needs to consider the two priorities when determining how best to coordinate the behaviors.

## 2. COMPLETED WORK

While there are a number of simulation environments currently available for use, such as Player/Stage/Gazebo, none of the ones found provided for both the extensive customization of the agent's architecture and a streamlined execution model necessary for integration with various machine learning techniques that require significant numbers of evaluations. As a result, a custom simulator was developed that satisfies both of these necessary requirements. Once the simulation environment was usable, the first stage of the research was to determine the viability of a priority-based architecture. To this end, a simple priority-based behavior hierarchy was developed for the task of pursuing a prey agent [1]. To further test the abstraction capabilities of priorities, the same pursuit behavior was used in a multi-agent team pursuit environment without modification of the behavior itself<sup>1</sup>. These initial results indicate that a priority-based fuzzy behavior hierarchy is not only capable of handling simple tasks but also allow for the reuse of existing, proven behaviors in contexts different from those in which they were developed. This means that it may be possible to create complex multi-team behaviors within a single-team sandbox to simplify the development process.

After this initial stage of research, our next step was to prove that priority-based architectures provide comparable performance to non-priority-based ones. We also wished to show that developing behaviors is easier when priorities are used. Therefore, an experiment was performed in which machine learning, specifically grammatical evolution, was used to create simple composite behaviors that coordinated two and three primitive behaviors respectively [2]. The results showed that for the given composite behaviors, priority-based architectures performed no worse than non-priority-based ones. This shows that the abstractions provided by priorities do not harm the performance of the agent in the given tasks, an important consideration when deciding if priorities are a viable contribution to the behavior coordination problem. Also, results indicate that priority-based solutions were learned at a faster rate than non-priority-based solutions, but further analysis is required before a definitive conclusion can be made.

## 3. WORK IN PROGRESS

Currently, work is progressing on many fronts. First, a physics engine is being integrated into the simulator to provide a more realistic simulation environment with the expectation that future behaviors can be learned in the simulator and then ported to actual robots with minimal effort. Second, an analysis of the learning rates for priority-based and non-priority-based solutions is being performed to determine if priorities do indeed promote faster learning of composite

behaviors. Lastly, a more complex behavior hierarchy is being developed to more fully test whether or not priorities provide for behavior hierarchies that are deeper and more complex than traditional hierarchies that do not make use of priorities. To increase the complexity of the behavior coordination problem, some of the behaviors used will be multi-agent in nature. This added complexity will further test the priorities ability to abstract many of the low-level details while still retaining the fidelity needed to effectively coordinate behaviors.

## 4. FUTURE WORK

A considerable amount of future work still remains. First and foremost is an investigation into the dynamics of generating the priorities used. Currently, priorities are generated by hand using experience and knowledge of the behaviors being abstracted. This approach would not be scalable to the number or variety of behaviors that will eventually be developed. The goal is for priorities to be learned using mechanisms similar to those which are used to develop the composite behaviors.

Once this goal has been accomplished, work will progress on scaling not only the behavior hierarchy itself, but also the number of agents involved. The end goal is to develop a set of behavior hierarchies for a large number of agent teams in which inter- and intra-team coordination is required for successful completion of the given tasks. It is intended that these hierarchies be so complex that only the use of priorities make their development and effectiveness possible.

## 5. REFERENCES

- [1] B. E. Eskridge and D. F. Hougen. Prioritizing fuzzy behaviors in multi-robot pursuit teams. In *Proceedings of the IEEE Conference on Fuzzy Systems*, pages 6039–6045, Vancouver, BC, 2006.
- [2] B. E. Eskridge and D. F. Hougen. Using priorities to simplify behavior coordination. In *Proceedings of the 2007 International Conference on Autonomous Agents and Multiagent Systems (to appear)*, 2007.
- [3] M. Nicolescu and M. J. Matarić. A hierarchical architecture for behavior-based robots. In *International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 227–233, Bologna, Italy, July 2002.
- [4] P. Pirjanian and M. J. Matarić. Multiple objective vs. fuzzy behavior coordination. In D. Drainkov and A. Saffiotti, editors, *Fuzzy Logic Techniques for Autonomous Vehicle Navigation*, pages 235–253. Springer-Verlag, 2000.
- [5] E. Tunstel. Fuzzy-behavior synthesis, coordination, and evolution in an adaptive behavior hierarchy. In A. Saffiotti and D. Driankov, editors, *Fuzzy Logic Techniques for Autonomous Vehicle Navigation*, volume 61 of *Studies in Fuzziness and Soft Computing Series*, chapter 9, pages 205–234. Physica-Verlag, 2001.

---

<sup>1</sup>The only required changes were the adjustment of the membership functions associated with the existing fuzzy linguistic values.

# Outline of Research Programme

## A Logical Theory of Coordination and Joint Ability

Hojjat Ghaderi  
Department of Computer Science  
University of Toronto  
Toronto, ON M5S3G4, Canada  
hojjat@cs.toronto.edu

### 1. INTRODUCTION

Epistemic ability, i.e. whether an autonomous agent has enough knowledge to achieve a goal, is one of the key notions in studying intelligent behavior. In the last two decades, a considerable amount of research has been devoted to formalizing this notion for *individual* agents. Nevertheless, the concept becomes much more challenging and interesting in multi-agent systems: while each agent may be unable to individually achieve a goal, forming a team may enable them to achieve it as a common goal. A team of agents is *jointly able* to achieve a common goal if despite any incomplete knowledge or even false beliefs that they may have about the world or each other, they still know enough to be able to get to a goal state, should they choose to do so. Unlike in the single-agent case, the mere existence of a working plan is not sufficient since there may be several incompatible working plans and the agents may not be able to choose a share that coordinates with the others'. The coordination of teams of cooperating but autonomous agents is a core problem in multi-agent systems research.

There is a large body of work in game theory [5] dealing with coordination and strategic reasoning for agents. The classical game theory framework has been very successful in dealing with many problems in this area. It can deal with cooperative teams of agents if we assign the same utility function to all team members. However, a major limitation of the framework is that it assumes there is a *complete specification* of the structure of the game including the agents' beliefs. It is also often assumed that this structure is common knowledge among agents. These assumptions often do not hold for team members, let alone for a third party attempting to reason about what the team members can do.

In recent years, there has been a lot of work aimed at developing symbolic logics of games [6, 10] so that more incomplete and qualitative specifications can be dealt with. This can also lead to faster algorithms as sets of states that satisfy a property can be abstracted over in reasoning. However,

this work has often incorporated very strong assumptions. Many logics of games like Coalition logic [6] and ATEL [10] ignore the issue of coordination within a coalition and assume that the coalition can achieve a goal if there exists a joint working plan/strategy profile that achieves the goal. This is only sufficient if we assume that the agents can communicate arbitrarily to agree on a joint plan/strategy profile. As well, most logics of games are propositional, which limits expressiveness.

My thesis proposes a new first-order (with some higher-order features) logic framework to model the coordination of coalitions of agents based on the situation calculus [4, 3, 7]. The proposed formalization of joint ability avoids both of the pitfalls mentioned above: it supports reasoning on the basis of very incomplete specifications about the belief states of the team members and it ensures that team members do not have incompatible strategies. The framework deals with strict uncertainty where no probabilistic information is available and it addresses agents' coordination without requiring arbitrary communication among them. Inspired by concepts from game theory, agents in my framework try to coordinate their actions through discarding dominated plans and keeping only the most rational ones. Some of the technical results are being published in Commonsense 2007 and AAMAS 2007.

### 2. COMPLETED

I have developed a logical formalization of joint ability in the situation calculus. My account of joint ability requires some second-order features of the situation calculus [3], including quantifying over certain functions from situations to actions, that we call *strategies*. To represent belief of agents, I have adopted Shapiro et al. [9]' logic of knowledge for multiple agents which is based on possible world semantics in the situation calculus. Fluent  $B(x, s', s)$  is used to denote that in situation  $s$ , agent  $x$  thinks that situation  $s'$  might be the actual situation. Belief is then defined as an abbreviation:

$$Bel(x, \phi[now], s) \doteq \forall s'. B(x, s', s) \supset \phi[s'].$$

We require  $B$  to be serial, and transitive, so that belief satisfies the modal system *weak* S5. When we need to represent knowledge we will require  $B$  to be reflexive as well (corresponding to the S5 modal system).

As mentioned, the mere existence of a strategy profile that achieves the goal is not sufficient for joint ability since there may be several incompatible such profiles and the agents

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

Copyright 2007 IFAAMAS .

may not be able to choose a share that coordinates with the others'. To ensure the agents are properly coordinated, each one compares her strategies based on her private beliefs. Initially, they consider all available strategies possible. Then they eliminate strategies that are not as good as others given their beliefs about what strategies the other agents have kept. This elimination process is repeated until it converges to a set of preferred strategies for each agent. Joint ability holds if all combinations of preferred strategies succeed in achieving the goal. For  $N$  agents trying to achieve the common goal *Goal*, joint ability is defined as follows:

- $JCan(s)$ : Agents can jointly achieve the goal iff all combinations of their *preferred* strategies *work* together.
- $Works(\vec{\sigma}, s)$ : Strategy profile  $\vec{\sigma}$  works in situation  $s$  if there is a future situation where the common goal *Goal* holds and the strategies prescribe the actions to get there according to whose turn it is.
- $Pref(i, \sigma_i, s) \doteq \forall n. Keep(i, n, \sigma_i, s)$   
Agent  $i$  prefers strategy  $\sigma_i$  if it is *kept* for all levels  $n$ .
- $Keep$  is defined inductively:
  - $Keep(i, 0, \sigma_i, s) \doteq Strategy(i, \sigma_i)$ .  
At level 0, all strategies are kept.
  - $Keep(i, n + 1, \sigma_i, s)$ : For each agent  $i$ , the strategies kept at level  $n + 1$  are those kept at level  $n$  for which there is not a better one ( $\sigma'_i$  is better than  $\sigma_i$  if it is as good as  $\sigma_i$  while  $\sigma_i$  is not as good as it).
- $Strategy(i, \sigma_i)$ : Strategies for an agent are functions from situations to actions such that the action is known to the agent whenever it is the agent's turn to act.
- $AsGoodAs(i, n, \sigma_i, \sigma'_i, s)$ : Strategy  $\sigma_i$  is as good as  $\sigma'_i$  for agent  $i$  if  $i$  believes that whenever  $\sigma'_i$  works with strategies kept by the rest of the agents so does  $\sigma_i$ .

The above definition resembles the iterative elimination of weakly dominated strategies of game theory [5, 1]. Nonetheless, it differs from game theory approaches, however, in a number of ways. Foremost, our framework not only handles incomplete information [2], but also incomplete *specifications* where some aspects of the world or agents including belief/disbelief are left unspecified. Since our proofs are based on entailment, they remain valid should we add more detail to the theory. Second, rather than considering utility functions, our focus is on goal achievability for a team. Moreover, we consider strict uncertainty and assume no probabilistic information is available. Our framework supports a weaker form of belief (as in the weak S5 logic) and allows for false belief. As such, the traditional definition of weak dominance cannot be used and an alternative approach for comparing strategies (as proposed in our work) is needed, one that is based on the private beliefs of each agent about the world and other agents and their beliefs.

To show the plausibility of my proposed definition, I have formalized several simple but interesting examples involving private/distributed knowledge among agents and proved joint ability/inability even from very incomplete specifications. Also, I have proved various general properties of joint ability including its nonmonotonicity w.r.t. the underlying

goal. This is an interesting result as it is in contrast to the single agent case where if an agent can achieve a "hard" goal it can achieve "easier" (i.e. logically entailed) goals as well.

### 3. IN PROGRESS

The definition of joint ability as presented here is w.r.t.  $N$  agents in a single coalition trying to achieve a common goal. It can be straightforwardly generalized to allow some agents to be outside of the coalition (hence not being interested in achieving the goal). It can also be generalized to the cases where agents are partitioned into several coalitions. One simple approach is to assume agents in each coalition adopt the conjunction of the goals of all members. Various scenarios can be considered depending on the assumptions about the beliefs of agents about other coalitions (such as whether they know about their goals, their rationality, etc.).

In addition to these, I am formalizing various complex examples involving more than two agents, sensing actions, and communication. Also, I am formalizing several interesting properties of joint ability including proper conditions for its persistence.

### 4. FUTURE WORK

In future work, we will work to generalize the framework in various ways. Supporting sensing actions and simple communication actions should be straightforwardly handled by revising the successor state axiom for belief accessibility as in [8, 9]. We will also examine how different ways of comparing strategies (the *AsGoodAs* order) lead to different notions of joint ability, and try to identify the best. Finally, we will look at how our framework could be used in automated verification and multiagent planning.

### 5. REFERENCES

- [1] A. Brandenburger and H. J. Keisler. Epistemic conditions for iterated admissibility. In *TARK '01*, pages 31–37, San Francisco, 2001. Morgan Kaufmann.
- [2] J. C. Harsanyi. Games with incomplete information played by Bayesian players. *Management Science*, 14:59–182, 1967.
- [3] H. Levesque, F. Pirri, and R. Reiter. Foundations for the situation calculus. *Electronic Transactions on Artificial Intelligence*, 2(3-4):159–178, 1998.
- [4] J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. 1969.
- [5] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1999.
- [6] M. Pauly. A modal logic for coalitional power in games. *J. of Logic and Computation*, 12(1):149–166, 2002.
- [7] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, 2001.
- [8] R. Scherl and H. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144:1–39, 2003.
- [9] S. Shapiro, Y. Lesprance, and H. J. Levesque. Specifying communicative multi-agent systems. In W. Wobcke, M. Pagnucco, and C. Zhang, editors, *Agents and Multiagent systems- Formalisms, Methodologies, and Applications*, pages 1–14. Springer-Verlag, 1998.
- [10] W. van der Hoek and M. Wooldridge. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, 75(1):125–157, 2003.

# Coordination of Multi-Agent Systems in Dynamic Environments

Mattijs Ghijsen

University of Amsterdam, Human Computer Studies Laboratory  
Kruislaan 419, Matrix 1, 1098 VA Amsterdam, The Netherlands  
mattijs@science.uva.nl

## 1. INTRODUCTION

This research project is part of the Interactive Collaborative Information Systems (ICIS) project<sup>1</sup> which aims at developing techniques for making complex information systems more intelligent and supportive in decision making situations. One of the techniques that is being focused on in the ICIS project is Multi-Agent Systems (MAS). The test-case domain in which these MAS techniques are being deployed is disaster management which is (amongst others) characterized by its highly dynamic nature.

The main theme of this PhD project is the relation between the behavior of a MAS and the dynamics in its environment. The characteristics of the environment have a large influence on the behavior and performance of MAS. Vice versa, the behavior and performance of a MAS also influences the environment in which the MAS is embedded. How a MAS behaves is determined by the capabilities of its agents and how these agents are organized. Coordination of the activities of the agents in a MAS is therefore a key element of MAS organization.

The first part of this project will focus on the influence of the environment in which the MAS is embedded and task the MAS is performing on the coordination mechanism of the MAS. The second and third part of this project is focused on the behavior of a MAS in dynamic environments. The second part is a study on how an organization of agents can adapt itself continuously to cope with a dynamic environment. This requires understanding how reorganization behavior can be implemented in an agent and *when* and *how* organizational change should take place. The third part of this research project will focus more closely on information sharing in a dynamic organization. Efficient information sharing helps to achieve a high level of situational and organizational awareness and helps to avoid information overload.

## 2. COORDINATION AND TASK-ENVIRONMENT FACTORS

status: **completed**

Task-environment factors, such as dynamics, uncertainty, task complexity, etc., together with the organization of the MAS determine the efficiency and effectiveness of the MAS

<sup>1</sup>The Interactive Collaborative Information Systems project is supported by the Dutch Ministry of Economic Affairs, grant BSIK03024. <http://www.icis.decis.nl>

[9]. In the RoboCupRescue simulator [7] I have implemented a M.A.S that is capable of enacting three different coordination styles; direct supervision, standardization of skills and standardization of skills extended with a reorganization mechanism. These coordination mechanisms are based on organizational theory by Mintzberg [8]. An experiment has demonstrated the effects of unknown and possibly heterogeneous workload distribution and incomplete environmental information on the performance of the three coordination styles [3]. It was shown that assigning large size tasks to agents reduced the need for complex decision making at the central command and also reduced communication overhead. However, assigning large tasks reduced performance in the case of heterogeneous workload. Restructuring the organization improved performance when assigning large tasks in the case of heterogeneous workload distribution.

status: **in progress**

The experiment in [3] was rather limited in the sense that it included only two task-environment factors and that the coordination styles used are all centralized approaches to coordination. Current work is focused on the development of theory that explains the relation between coordination mechanisms and task-environment factors. The approach is to develop a typology of coordination mechanisms by describing coordination mechanisms in terms of generic domain independent characteristics. The next step is to take a set of task-environment factors and explain the influence of these task-environment factors on the different types of coordination mechanisms. This research results in theory that guides the designer of a MAS in his/her choice for a specific coordination mechanism, based on a set of task-environment characteristics the MAS is expected to encounter.

## 3. REORGANIZATION

status: **completed**

Research in the area of MAS organization has shown that the ability to reorganize will enable a MAS to mitigate or reduce the negative effects of dynamics in the environment [2]. Based on this need for reorganization and the close relation between coordination and reorganization (as found in literature, e.g. [8]), a first version of a framework for agent coordination and reorganization is developed and presented in [4]. This framework enables the agents of a MAS to adapt the organization of the MAS in every aspect of the organization. This includes adapting relations and interaction patterns between agents (e.g. [5]), changes in agent roles (e.g. [6]) and changes in the way the organization is coordinated.

minated (e.g. [1]). The framework has already been used to implement the reorganization mechanism used in [3], which uses changes in authority relations, interaction patterns and agent roles to deal with dynamics in workload distribution.

status: **to be done**

The framework will be used to build agents that are capable of all aspects of reorganization. Using these agents we will be able to conduct experiments in which we investigate *what* triggers organizational change (i.e. uneven load balance, unreliable communication, new tasks or goals, etc.) and *how* the instruments of organizational change should be used to adapt the organization. Furthermore I will need to investigate for which types of problem domains the framework is suited and for which it is not.

This research should result in theory that guides a MAS designer in the choice of which types of organizational change to implement, given a set of task-environment characteristics the MAS is likely to encounter.

#### 4. INFORMATION SHARING

status: **to be done**

Sharing information is an important aspect of coordination. When operating in a dynamic environment, sharing the right information with other agents is vital for the performance of a MAS. To avoid information overload and maintain a high level of situational and organizational awareness, agents have to be careful in deciding what information to share and with which agents to share the information. Additional complexity in this decision is added by the fact that information may become outdated as a result of the dynamics in the environment. In this study I will determine appropriate information sharing strategies, given the organizational design of the M.A.S and the task-environment factors the MAS is coping with.

#### 5. CONCLUSIONS

The research discussed in the previous sections will result in a framework for agent coordination and reorganization. This framework should enable a MAS developer to build agents that operate in dynamic environments by giving the agents the appropriate capabilities to coordinate and reorganize. This PhD project will also result in theory that supports a MAS designer in the choice for specific coordination and reorganization capabilities as this choice will depend on the task-environment factors, the MAS is expected to encounter.

#### 6. REFERENCES

- [1] K. Barber and C. Martin. Dynamic reorganization of decision-making groups. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, pages 513–520, New York, NY, USA, 2001. ACM Press.
- [2] K. Carley. Computational and mathematical organization theory: Perspectives and directions. *Journal of Computational and Mathematical Organizational Theory*, 1995.
- [3] M. Ghijsen, W. Jansweijer, and B. Wielinga. The effect of task and environment factors on M.A.S. coordination and reorganization. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems, short paper, to appear*, 2007.
- [4] M. Ghijsen, W. Jansweijer, and B. Wielinga. Towards a framework for agent coordination and reorganization, agentcore. In *Proceedings of the Workshop on Coordination, Organization, Institutions, and Norms in Agent Systems (COIN @ AAMAS 2007)*. To appear., 2007.
- [5] B. Horling, B. Benyo, and V. Lesser. Using self-diagnosis to adapt organization structures. In *Proceedings of the 5th International Conference on Autonomous Agents*, pages 529–536. ACM Press, June 2001.
- [6] S. Kamboj and K. Decker. Organizational self-design in semi-dynamic environments. In *2007 IJCAI workshop on Agent Organizations: Models and Simulations (AOMS@IJCAI 07)*, 2007.
- [7] H. Kitano, S. Tadokoro, I. Noda, H. Matsubara, T. Takahashi, A. Shinjou, and S. Shimada. Robocup-rescue: Search and rescue for large scale disasters as a domain for multi-agent research. In *Proceedings of IEEE Conference on Man, Systems, and Cybernetics(SMC-99)*, 1999.
- [8] H. Mintzberg. *Structures in fives: designing effective organizations*. Prentice Hall, Englewood Cliffs, N.J., 1993.
- [9] Y. So and E. Durfee. Designing organizations for computational agents. pages 47–64, 1998.

# Game-theoretic approaches to constructing intelligent agents

Andrew Gilpin  
Computer Science Department  
Carnegie Mellon University  
Pittsburgh, PA 15213  
gilpin@cs.cmu.edu

## ABSTRACT

Game theory is the mathematical study of rational behavior in strategic environments. In many settings, most notably two-person zero-sum games, game theory provides particularly strong and appealing solution concepts. Furthermore, these solutions are efficiently computable in the complexity-theory sense. However, in most interesting potential applications in Artificial Intelligence, the solutions are difficult to compute due primarily to the extremely large state-spaces in the environments.

In my thesis, we propose new algorithms for tackling these computational difficulties. In one stream of research, we introduce *automated abstraction algorithms for sequential games of imperfect information*. These algorithms take as input a description of a game and produce a description of a strategically similar, but smaller, game as output. We present algorithms that are lossless (*i.e.* equilibrium-preserving), as well as algorithms that are lossy, but which can yield much smaller games.

We also introduce *specialized optimization algorithms for finding  $\epsilon$ -equilibria in sequential games of imperfect information*. The algorithms are based on recent advances in convex optimization and provide significant improvements over previous algorithms for finding  $\epsilon$ -equilibria.

Using the new automated abstraction algorithms in conjunction with the new algorithms for finding  $\epsilon$ -equilibria enables the application of game theory to much larger games than was previously possible. In particular, we find near-optimal solutions for a four-round model of Texas Hold'em poker, and demonstrate that the resulting player is better than previous computer poker players.

As future work to complete my thesis we propose coming up with worst-case guarantees (both *ex ante* and *ex post*) as well as developing algorithms for finding approximate equilibria in games with many players.

## 1. OVERVIEW

Developing game-playing computer agents that are competitive with humans has long been a major focus of the Ar-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

Copyright 2007 IFAAMAS.

tificial Intelligence community. The most notable successes include Chinook, which defeated the human world champion Dr. Marion Tinsley in 1992; Deep Blue, which defeated Garry Kasparov, the reigning world champion in 1997; and TD-Gammon, the best backgammon-playing program in the world. These are all very impressive applications of AI techniques and have done much to advance the standing of AI in the wider science community. One common property of each of these three games is *perfect information*, *i.e.* at any point in time, both players are fully informed about the state of the world. In contrast, poker and most other card games have the property of *imperfect information*; at most stages of the game, the players are only partially informed of the state of the world. In particular, a poker player does not know what cards the opponent is holding. Due to this difference, the algorithms developed for perfect information games are unhelpful when designing algorithms for games with imperfect information. For this reason, as well as many others, poker is an important area of research for AI [3].

My thesis work concerns the automatic construction of artificially intelligent agents for playing in games with imperfect information. As an application area, my work has focused primarily on poker, though the techniques can be applied much more broadly. In the remainder of this document I describe what I have done so far, and what I am planning on doing next.

## 2. COMPLETED WORK

### 2.1 Lossless automated abstraction

Finding an equilibrium of an extensive form game of imperfect information is a fundamental problem in computational game theory, but previous techniques do not scale to large games. To address this, we introduced the *ordered game isomorphism* and the related *ordered game isomorphic abstraction transformation*. For a multi-player sequential game of imperfect information with observable actions and an ordered signal space, we proved that any Nash equilibrium in an abstracted smaller game, obtained by one or more applications of the transformation, can be easily converted into a Nash equilibrium in the original game. Using this concepts, we developed an algorithm, *GameShrink*, for automatically and exhaustively abstracting the game.

Rhode Island Hold'em is a poker card game that has been proposed as a testbed for AI research. The game tree of Rhode Island Hold'em contains more than 3.1 billion nodes. This game features many characteristics present in full-scale

poker (*e.g.*, Texas Hold'em). Using *GameShrink*, we solved for the optimal (equilibrium) strategies of the game [5, 7]. The optimal player is available for play online at <http://www.cs.cmu.edu/~gilpin/gsi.html>.

## 2.2 Lossy automated abstraction

Unfortunately, even after applying the *GameShrink* algorithm, many games (*e.g.*, Texas Hold'em) are still too large to be solved using existing optimization tools. With this in mind, we developed new automated abstraction algorithms that do not have the same game-theoretic guarantees as the original *GameShrink* algorithm, but still preserve the basic structure of the game.

Our lossy abstraction algorithm works as follows. In each round of the game, there is a limit to the number of strategically different situations that an equilibrium-finding algorithm can handle. Given this constraint, we use clustering to discover similar positions, and we compute the abstraction via an integer program that minimizes the expected error at each stage of the game. Experiments show that this technique leads to a drastic improvement over prior approaches for automatically generating agents, and our agent plays competitively even against the best agents overall [6, 8, 9].

We have also developed a potential-aware automated abstraction technique. It applies to a broad range of sequential imperfect information games. We applied it to Texas Hold'em, and solved the abstracted game using a variant of the excessive gap technique (see Section 2.3). This is, to our knowledge, the first time that all four betting rounds have been abstracted and game-theoretically analyzed in one run (rather than splitting the game into phases). The resulting player beats the prior leading poker programs (*GS2* [9], *Sparbot* [2], and *Vexbot* [1]) with statistical significance.

## 2.3 Equilibrium computation

In addition to developing algorithms for automatically abstracting games, I have also worked on algorithms for computing equilibria in the (possibly abstracted) games. One product of this research is an algorithm for finding  $\epsilon$ -equilibria in two-person zero-sum extensive form games in  $O(1/\epsilon)$  iterations. The algorithm is based on modern smoothing techniques for non-smooth convex optimization [12]. Our implementation of this algorithm has enabled us to compute near-optimal solutions for games that are several orders of magnitude larger than was previously possible [10, 4].

## 3. WORK TO BE DONE

My future work includes developing worst-case *a priori* guarantees on the quality of an abstraction (currently, the quality of an abstraction can only be evaluated by computing an equilibrium and testing the resulting solution), as well as developing algorithms for finding *ex post* guarantees in large games. In addition, we are interested in extending this line of research to games with more than two players, perhaps by extending the *MIP Nash* framework [13] to handle such games, and also we are interested in considering tournament games [11].

## 4. REFERENCES

[1] D. Billings, M. Bowling, N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and

D. Szafron. Game tree search with adaptation in stochastic imperfect information games. In *Proceedings of the 4th International Conference on Computers and Games (CG)*, pages 21–34, Ramat-Gan, Israel, July 2004. Springer-Verlag.

[2] D. Billings, N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and D. Szafron. Approximating game-theoretic optimal strategies for full-scale poker. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 661–668, Acapulco, Mexico, 2003.

[3] D. Billings, A. Davidson, J. Schaeffer, and D. Szafron. The challenge of poker. *Artificial Intelligence*, 134(1-2):201–240, 2002.

[4] A. Gilpin, S. Hoda, J. Peña, and T. Sandholm. Gradient-based algorithms for finding Nash equilibria in extensive form games. Mimeo, 2007.

[5] A. Gilpin and T. Sandholm. Optimal Rhode Island Hold'em poker. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1684–1685, Pittsburgh, PA, 2005. Intelligent Systems Demonstration.

[6] A. Gilpin and T. Sandholm. A competitive Texas Hold'em poker player via automated abstraction and real-time equilibrium computation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Boston, MA, 2006.

[7] A. Gilpin and T. Sandholm. Finding equilibria in large sequential games of imperfect information. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 160–169, Ann Arbor, MI, 2006.

[8] A. Gilpin and T. Sandholm. A Texas Hold'em poker player based on automated abstraction and real-time equilibrium computation. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1453–1454, Hakodate, Japan, 2006. Demonstration Track.

[9] A. Gilpin and T. Sandholm. Better automated abstraction techniques for imperfect information games, with application to Texas Hold'em poker. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Honolulu, HI, USA, 2007. To appear.

[10] S. Hoda, A. Gilpin, and J. Peña. A gradient-based approach for computing Nash equilibria of large sequential games. Manuscript. Presented at *INFORMS*, Pittsburgh, PA, USA, 2006.

[11] P. B. Miltersen and T. B. Sørensen. A near-optimal strategy for a heads-up no-limit Texas Hold'em poker tournament. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Honolulu, HI, USA, 2007. To appear.

[12] Y. Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM Journal of Optimization*, 16(1):235–249, 2005.

[13] T. Sandholm, A. Gilpin, and V. Conitzer. Mixed-integer programming methods for finding Nash equilibria. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 495–501, Pittsburgh, PA, 2005.

# Robust Team Formation Strategies for Dynamic Environments

Chris L. D. Jones

The University of Texas at Austin

Laboratory for Intelligent Processes and Systems

1 University Station C5000, Austin, TX, 78712-0240

+1-512-471-5350

coldjones@lips.utexas.edu

## ABSTRACT

In dynamic market environments, fully autonomous agents may form teams to work on large, multifaceted problems. Factors such as uncertain information, bounded rationality and environmental dynamicism can lead to sudden, unforeseen changes in both solution requirements and team participation. Accordingly, this research examines strategies for robust, bottom-up team formation in dynamic market environment. Robust team formation strategies control how agents select problems to work on and partners to work with, with the aim of maximizing the utility of individual agents even in the event of teammate defections, sudden changes to a problem's solution requirements, and other unforeseen events.

## 1. Problem

Technologies such as the Semantic Web allow humans and agents to easily find and communicate with each other. Extending these ideas to concepts such as Google Answers and the Amazon Mechanical Turk suggests the potential for a huge electronic marketplace, where a constant stream of questions or problems is handled by thousands or millions of humans and self-interested agents, each taking on different roles such as general contractor or service provider, each buying or selling skills, expertise, and other services as needed. At the same time, the comparative openness of such a market may preclude traditional contract-based methodologies with have shaped interactions between agents.

Furthermore, many problems submitted to a sophisticated problem-solving market are likely to be complex, requiring multiple skills from multiple providers. Many of these problems are also likely to be novel, undertaken with incomplete information about the problem and a limited understanding of the solution requirements, both of which will almost certainly change as the problem is worked on. Real world problems also occur in dynamic environments, where unexpected changes occur to both the problem and the personnel involved in solving it.

In this type of market environment, many issues arise. Even given a way to accurately determine which agents possess what skill set, how can an agent seeking to maximize its own utility select the best jobs to work on, and the best agents to partner with? How do the partners available to work on different possible jobs influence which job an agent should pursue? How should a team of agents be structured to handle difficulties such as sudden changes to a problem, or defections from the team? And how can these varying requirements be balanced against each other?

Work on task selection and coalition formation in the multi-agent systems community has addressed different pieces of the problem of robust team formation in dynamic environments: task

selection [1, 2], coalition formation in dynamic environments [3, 4], bottom-up team formation between autonomous agents [5], team formation in environments with imperfect information [6, 7, 8]. However, relatively little work has been done on methodologies that address these issues simultaneously. Furthermore, various incompatibilities in basic assumptions (e.g. selfish vs. non-selfish agents, static vs. dynamic environments) make merely combining the above work infeasible in most cases.

Therefore, to address these questions, this research proposes and explores a set of strategies for robust, bottom-up team formation between autonomous agents in a dynamic environment. These strategies control how an agent selects problems to work on and partners to work with, with the overall goal of maximizing the utility of individual agents even in the event of teammate defections and/or sudden changes to a problem's solution requirements.

Thus far, this research has focused on searching for primary strategies for forming robust teams in open, dynamic market environments. These strategies break down into three main types:

- "Risk-averse" strategies, which minimize agents' risk of failure in a dynamic environment by minimizing the length of time needed for a team to complete a given job.
- "Fault tolerant" strategies, which build teams capable of solving problems even in the event of team-member failure or defection.
- "Task adaptive" strategies, which build teams capable of solving problems even as previously unforeseen skills or resources are added to the solution requirements.

An initial exploration of these strategies and their relative performance in an open, dynamic market environment was conducted in [9]. Future plans for this research can be divided into two main tasks: in-depth exploration of each of the three strategy types given above, and research on how the different strategy types may be combined to build more balanced, successful strategies overall.

## 2. Risk-averse strategies

Initial work on these strategies has focused on defining a job selection heuristic and a task selection heuristic which work in combination to form a complete strategy. More particularly, the initial job selection heuristic proposed in [9] merely selects jobs which have minimal time to completion, while the initial team

selection heuristic selects teams which can fulfill their assigned jobs the fastest.

Future work in this area will compare minimal time-to-completion strategies with strategies that seek to minimize the perceived risk of a given job or team. Future work in this area will also explore the reward weightings needed for agents executing risk-averse strategies to select complex, high-priority jobs rather than simple, low-priority jobs.

### 3. Fault tolerant strategies

Existing work in fault-tolerant team formation in multi-agent systems primarily focuses on the technical challenges of building a robust agent team, such as replicating agents as needed at runtime [10] or creating tightly-bound teams where agents are all in direct contact with one another [11]. However, such research does not touch on the problems that arise when unique, autonomous agents choose to drop out of a given team – e.g., providing sufficient skills to continue working on a problem when a task force member in a business environment leaves to pursue another opportunity. Accordingly, my initial work on these strategies examined a team-selection heuristic which preferred teams composed of multi-skilled agents with overlapping abilities to teams where the multiple skills of member agents did not overlap.

Future work in exploring these fault tolerant strategies seeks to examine algorithms for determining the maximum possible amount of skill overlap given an initial pool of multi-skilled agents, as well as quantifying the degree of skill redundancy needed to achieve a given level of reliability and a given probability that an agent will fail or defect. In addition, future work will relax the distributed set covering problem (DSCP) assumption of [1], and examine market-based mechanisms for allowing agents to distribute themselves between multiple jobs and multiple teams to maximize skill redundancy and team reliability.

### 4. Task adaptive strategies

Task adaptive strategies operate in a manner similar to fault tolerant strategies, but whereas multi-skilled agents work towards building redundant skills in a fault tolerant strategy, multi-skilled agents work towards building auxiliary skills under a task adaptive strategy. More specifically, task adaptive strategies are built around a team selection heuristic which attempts to select teams that have excess skills which are not currently required by the problem requirements, but may be required in the future if the problem suddenly changes.

Future work in exploring task adaptive strategies is expected to cover a similar range as work on fault tolerant strategies, examining algorithms for determining the maximum possible amount of auxiliary skill coverage possible given an initial pool of multi-skilled agents, and examining market-based mechanisms for allowing agents to distribute themselves between multiple jobs and multiple teams to maximize skill redundancy and team reliability.

### 5. Hybrid strategies

Once risk-averse, fault tolerant and task adaptive strategies have been better explored as outlined above, this research will seek to determine the most effective balances between the three strategies to create high-performing hybrid strategies. More particularly, using a real-world domain system currently under development, various hybrid agent strategies will be tested against each other and evaluated to determine which strategies allow an agent to maximize its utility.

- [1] Klusch, M. and Gerber, A. Dynamic Coalition Formation among Rational Agents. *IEEE Intelligent Systems* 17, 3 (May/June 2002) 42-47.
- [2] Kraus, S, Shehory, O. and Taase, G. Coalition Formation with Uncertain Heterogeneous Information. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '03)*. (Melbourne,3Australia, July 14-18). Pages 1-8.
- [3] Lerman, K., and Shehory, O. Coalition Formation for Large-Scale Electronic Markets. In *Proceedings of the Fourth International Conference on Multi Agent Systems (ICMAS '00)* (Boston, MA, July 7-12, 2000). Pages 167 – 174.
- [4] Martin, C. *Adaptive Decision Making Frameworks for Multi-Agent Systems*. Ph.D. Thesis, University of Texas, Austin, TX, 2001.
- [5] Modi, P. J., Jung, H., Tambe, M., Shen, W., and Kulkarni, S. A Dynamic Distributed Constraint Satisfaction Approach to Resource Allocation. In *Proceedings of the Seventh International Conference on Principles and Practice of Constraint Programming (CP-2001)* (Paphos, Cyprus, November 26 – December 1, 2001). Pages 685-700.
- [6] Nair, R., Tambe, M., and Marsella, S. Team Formation for Reformation. In *Proceedings of the AAAI Spring Symposium on Intelligent Distributed and Embedded Systems*. 2002.
- [7] Rathod, P. and desJardins, M. Stable Team Formation Among Self-Interested Agents. In *Working Notes of the AAAI-2004 Workshop on Forming and Maintaining Coalitions in Adaptive Multiagent Systems* (San Jose, California, July 26, 2004).
- [8] Sandholm, T.W., and Lesser, V. Coalitions among Computationally Bounded Agents. *Artificial Intelligence* 94, 1-2 (July 1997) 99-137.
- [9] Jones, C.L.J., Barber K. S. Bottom-up Team Formation Strategies in a Dynamic Environment. TR2007-UT-LIPS-004, The University of Texas at Austin, Austin, TX, 2007.
- [10] Guessoum, Z., Briot, J-P., Faci, N. Towards Fault-Tolerant Massively Multiagent Systems. Massively Multi-Agent Systems I, First International Workshop, MMAS 2004, Kyoto, Japan, December 10-11, 2004
- [11] Tomic, P. T., Agha, G. A. Maximal Clique Based Distributed Coalition Formation for Task Allocation in Large-Scale Multi-agent Systems. Massively Multi-Agent Systems I, First International Workshop, MMAS 2004, Kyoto, Japan, December 10-11, 2004

# Use of Organizational Self-Design for the design and coordination of multiagent systems

Sachin Kamboj  
Department of Computer and Information Sciences  
University of Delaware  
Newark, DE 19716  
kamboj@cis.udel.edu

In my research, I intend to focus on the organizational design of a subset of multiagent systems — specifically ones in which the environment is dynamically, albeit slowly, changing. I intend to show that most current approaches to organizational design either model the organization at design time, assuming a static environment, or generate a new organization on the fly, at run time, for each new instance of the problem and that such approaches are inefficient and fail to correctly model the dynamics of a slowly changing environment.

Organizational Self-Design (OSD) [1, 5, 9], in which the agents are responsible for designing their own organizations at run-time, is particularly suited to such environments. I am working on extending the work done by Gasser, Ishida et.al., on OSD, to more complex task environments for worth oriented domains. In particular, I am using TÆMS [2, 6] as the underlying model for a domain, which allows for (a) the representation of a much larger class of problems and problem characteristics than can be represented using existing approaches and (b) the evaluation of alternative coordination protocols and scheduling algorithms to be used in conjunction with OSD.

As a part of my thesis proposal, I have formally defined the OSD approach as a mapping from the task structure of the problem being solved to the organizational structure of the agents responsible for solving the problem. I have also designed and written a simulator/test-bed for running OSD experiments and for evaluating my approach. In particular, I have compared my OSD approach to the Contract Net Protocol for both static and dynamic environments. I used statistical testing to show that our approach works significantly better than the one-off task allocation strategy used by the Contract Net Protocol, given the same task/environmental conditions. (*COMPLETED*)

Furthermore, I have evaluated various task/resource allocation heuristics that can be applied to OSD, using a variety of performance criteria. In particular, I have shown how the variation in task allocation strategy affects the performance

of the organization and the tradeoffs being made when using different strategies. I have also evaluated the robustness of our OSD approach using the citizens approach [4]. (*COMPLETED*)

Finally, I have written a graphical user interface (GUI) that allows (a) the viewing of changes to the organizational structure over time; and (b) the progress of tasks through the organization on a timeline. This allows not only the easy debugging of the simulator (by viewing the changes taking place at any instance) but also the development of alternative approaches to OSD (by viewing situations in which the approach does not perform optimally). (*COMPLETED*)

Future goals of my research are outlined below:

- I would like to evaluate the survivalist approach[7] to robustness. The survivalist approach might actually be better than the citizen approach (used so far) for higher probabilities of agent failure, as the replicated agents may be processing the task structures in parallel and can take over the moment the original agents fail — thus saving time around tight deadlines. Also, I strongly believe that the optimal organizational structure may vary, depending on the probability of failure and the desired level of robustness. For example, one way of achieving a higher level of robustness in the survivalist approach, given a large numbers of agent failures, would be to relax the task deadlines. However, such a relaxation would result in the system using fewer agents in order to conserve resources, which in turn would have a detrimental effect on the robustness. Therefore, towards this end, I have begun exploring the robustness properties of task structures and the ways in which the organizational design can be modified to take such properties into account (*IN PROGRESS*).
- I would like to evaluate the effect of communication delays on the type of organizational structure that can be generated. Small communication delays would allow the generation larger/deeper organizations (consisting of more agents) than can be generated when the communication delay is significant. I would like to be able to quantify the tradeoff between the amount of communication delay and the size/type of organization that can be generated. (*IN PROGRESS*)
- I plan to evaluate the effect of non-local interdependencies between tasks on the organizational structures of the agents. For example, some tasks may have to be

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawaii, USA.  
Copyright 2007 ACM 1-59593-303-4/06/0005 ...\$5.00.

done by the same agent, limiting the ability to spawn a new agent for one of the tasks. (For example, the same agent that read a paper must write the review for it.) Other interdependencies (for example enable and facilitate relationships) might have similar effects on the organizational structure. I would like to be able to quantify what those effects are and how they affect the organizational structure. (*IN PROGRESS*)

- I would like to evaluate the interplay between coordination and organization. Do certain coordination mechanisms preclude the use of certain organizational structures, or vice versa. What coordination protocols are suitable for an arbitrary type of organizational structure. (*TO BE DONE*)
- I would like to evaluate the tradeoff between cloning[3, 8] and OSD. i.e., when is it better to clone and when is it better to divide the task structures? (*TO BE DONE*)
- I plan to evaluate the effect of changes in multiple optimizing criteria (for example changes in the desired quality or maximum cost) on the organizational structures that can be generated. (*TO BE DONE*)
- I would like to show how existing workflow languages might benefit from the use of organizational modeling, especially the kind of modeling being done by our approach. (*TO BE DONE*)
- Finally, I would like to do a case study to test my approach on a practical application. (*TO BE DONE*)

## 1. REFERENCES

- [1] D. Corkill and V. Lesser. The use of meta-level control for coordination in a distributed problem solving network. *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 748–756, August 1983.
- [2] K. S. Decker. Environment centered analysis and design of coordination mechanisms. *Ph.D. Thesis, Department of Computer Science, University of Massachusetts, Amherst*, May 1995.
- [3] K. S. Decker, K. P. Sycara, and M. Williamson. Cloning for intelligent adaptive information agents. In *Revised Papers from the Second Australian Workshop on Distributed Artificial Intelligence*, number 3-540-63412-6, pages 63–75, London, UK, 1997. Springer-Verlag.
- [4] C. Dellarocas and M. Klein. An experimental evaluation of domain-independent fault handling services in open multi-agent systems. *Proceedings of the International Conference on Multi-Agent Systems (ICMAS-2000)*, July 2000.
- [5] L. Gasser and T. Ishida. A dynamic organizational architecture for adaptive problem solving. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 185–190, Menlo Park, California, 1991. AAAI Press.
- [6] V. Lesser, K. Decker, T. Wagner, N. Carver, A. Garvey, B. Horling, D. Neiman, R. Podorozhny, M. NagendraPrasad, A. Raja, R. Vincent, P. Xuan, and X. Zhang. Evolution of the gpgp/taems domain-independent coordination framework. *Proceedings 1st International Conference on Autonomous Agents and Multi-Agent Systems (Plenary Lecture/Extended Abstract)*, pages 1–2, 2002.
- [7] O. Marin, P. Sens, J. Briot, and Z. Guessoum. Towards adaptive fault tolerance for distributed multi-agent systems. *Proceedings of European Research Seminar on Advances in Distributed Systems (ERSADS 2001)*, May 2001.
- [8] O. Shehory, K. Sycara, P. Chalasani, and S. Jha. Agent cloning: an approach to agent mobility and resource allocation. *IEEE Communications Magazine*, 36(7):58–67, 1998.
- [9] Y. So and E. Durfee. An organizational self-design model for organizational change. In *AAAI-93 Workshop on AI and Theories of Groups and Organizations: Conceptual and Empirical Research*, pages 8–15, Washington, D.C., July 1993.

# Language Coverage and Multi-agent Agreement Problems

Kiran Lakkaraju  
Department of Computer Science  
University of Illinois, Urbana-Champaign  
klakkara@uiuc.edu

## 1. INTRODUCTION

Agents in many different situations need to be able to communicate amongst themselves in order to exchange knowledge and coordinate actions to pursue goals that require multi-agent coordination or teamwork. Communication of knowledge between agents is hindered, however, by the different ways knowledge can be represented by different agents.

Each agent in a MAS might have their own way of representing its knowledge. For instance, some agents might use a connectionist model (such as a recurrent neural network) to represent facts about the world, and another agent could use a knowledge base of first order logic sentences. The heterogeneous encodings of knowledge in different agents makes it impossible for agents to directly communicate their knowledge to others.

To overcome this problem agents must develop a *language*, a mapping from internally encoded knowledge to a sign from a shared system of signs. Each agent must be able to translate their internally encoded knowledge into a sign, and each agent must be able to translate a sign into their internal knowledge representation. We will use the term *language* to denote this mapping from signs to knowledge (or meanings), as well as the particular form of the signs used by the agents.

Developers of MAS can decide upon a language for the agents in the system to use but this only works when the developers know exactly what knowledge needs to be encoded and shared by agents. In a dynamic system with new tasks, agents being added/removed, and different types of knowledge a static language might not be able to handle changes in the environment. A better option would be to have the agents develop a language amongst themselves that would reflect the priorities and needs of the agents. Necessarily, all agents must agree upon the same language in order to communicate with each other.

Not only must all agents agree upon a single language, they must agree upon a *good* language. A good language is one that can represent all the knowledge that needed to be communicated as well as do this efficiently. We call this the *language convergence* problem—how to get a set of agents

to autonomously create a good language that is shared by all the agents in the population.

My work focuses on the problem of language convergence. In the next section I will elaborate on my objective.

## 2. LANGUAGE CONVERGENCE AS A MULTI-AGENT AGREEMENT PROBLEM

What does it mean for an agent to converge upon a language? A language has several components. For our purposes, a language consists of a finite set of *words* which can be combined to form *sentences*. We consider words to be the smallest linguistic unit that has meaning.

The ways in which words can be combined to form legal sentences in the language is called the *syntax* of the language. Every sentence in the language is assigned a *meaning*. The method by which meaning is associated with a sentence is called the *semantics* of the language. The syntax and semantics of a language are closely related, oftentimes the meaning assigned to a sentence depends upon how the sentence is organized, and vice versa.

For agents to agree upon a language they must agree upon several things. First of all they must agree upon the words that can be used in the language. Secondly, they must agree upon the syntax of the language, how to organize the words into sentences. Thirdly, they must agree on what a sentence *means*.

We can see that language convergence involves agreement at several different levels, at the lexical level (word agreement), then at the syntactic level (sentence structure agreement), and then at the semantic level (agreement on what a sentence means).

From this perspective, language convergence seems like a Multi-Agent Agreement problem (MAP). In an MAP multiple agents must navigate a space of possible states ( a *potential agreement space*) and eventually converge on the same state. Agreement problems have been studied before in the MAS community. Work in distributed transaction processing is one instance of a MAP, other instances involve meeting time scheduling, flocking/swarming, negotiation, team decision processes and synchronization of coupled oscillators for instance.

The hypothesis of my thesis is that the language convergence problem can be viewed as an instance of a Multi Agent Agreement problem, where agents must agree upon a language by searching through the language space. The complexity of the language convergence problem (agreement on multiple levels) lends itself to be considered a *model problem*—a problem through which insight on other MAP prob-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 Doctoral Mentoring Program May 13 2007, Honolulu, Hawai'i, USA.  
Copyright 2007 IFAAMAS .

lems can be gained.

Since MAPs have been studied in many domains we initially thought that solutions to MAPs should follow a general pattern. Unfortunately this was not the case. We discovered that each MAP domain has its own idiosyncrasies of problem description that are sometimes hard to apply to other domains. A MAP *lingua franca* was needed in order to compare, contrast and elicit the underlying concepts of MAPs. In order to compare and contrast MAPs we developed the *Distributed Optimal Agreement Framework* (DOA).

By mapping MAPs to the DOA framework we discovered how MAPs differed on dimensions including complexity of the potential-agreement space, state-to-state accessibility, agent interaction topologies, state evaluation measures, and type of solution needed (to name a few). This has allowed us to see what types of situations research in MAPs has currently tackled and how it can be relevant to the language convergence problem.

To summarize, my thesis proposes that language convergence is an instance of a MAP; and that through the study of language convergence and other MAPs general approaches to the study of MAPs can be developed.

To study this hypothesis I must show how language convergence problem can be viewed as an MAP. The research programme listed below goes over the steps to achieve this.

### 3. OUTLINE OF RESEARCH PROGRAMME

- **(complete) Develop a framework within which we can compare and contrast different multi-agent agreement problems and identify the key properties of multi-agent agreement problems.** We developed a framework (the Distributed Optimal Agreement framework) for describing a wide variety of multi-agent agreement problems in many domains. The framework captures aspects of interaction between agents, complex accessibility relations between states in the potential agreement space, and agreement on an optimal agreement state. This work was published in [1].
- **(complete) Develop a domain and task for which agents must autonomously create a shared language.** To study language convergence we focus on situations where agents need to convey information about a scene in an artificial grid world to another agent. Each scene consists of two objects that can vary in terms of shape and color. The two objects are placed on a 5 by 5 grid. The task of the agents is to describe the scene to other agents in the population. The goal is for all the agents to develop a scene description language by which they can communicate to each other, unambiguously, what objects are in a scene and how they are positioned.

This domain is interesting because it provides a foundation for further study of more complex domains. Coordinated planning relies on an agents ability to describe the state of the world to other agents in the system. It is a crucial cognitive ability for an agent to sense and describe a scene to another agents. To describe a scene agents will have to convey information about the size, shape and location of a pair of objects

to another agents. This will require encoding complex relationships, such as “blue square is one block east of the red triangle” which contain numerous individual entities and concepts.

- **(in progress) Develop a model of language that allows us to describe complex languages that have the expressibility to represent the domain described above.**

A *Language Model* is a model for representing a set of language. As such, it should be able to describe the syntax and semantic rules of a language. For instance, a language model could specify the syntax of a language by specifying a a context free grammar. The semantics could be represented as a table, with rows as sentences and columns as possible meanings.

In addition, our language model must be able to capture how we can modify a language (for instance by adding new words, specifying new meanings to words / sentences, admitting new organizations of words in sentences, etc.). The language model must allow incremental changes in the syntax and semantics of languages.

Currently we are pursuing a simplified variant of *Fluid Construction Grammars* (FCG) [2] for our language model. FCGs are suited to our needs because they have a uniform representation for syntactic and semantic rules, and easily allow for incremental language by creating new *constructions*.

- **(to be done) Model agreement upon a language as a multi-agent agreement problem.** To talk about language convergence as a multi-agent agreement problem we must determine how to map the process of an agent searching through a state space to the idea of an agent searching through a language space.

One of the crucial elements involved in defining a search space is the concept of distance between states. Knowledge of how distant two agents are in terms of language allows them to take proper action. Defining a distance metric over languages is a difficult endeavor.

For instance, how do you define the distance between two languages that differ in word ordering, from Subject-Verb-Object (SVO) (like English) to Subject-Object-Verb (SOV) (like Hindi)?

- **(to be done) Develop algorithms for reaching convergence over language space by applying general approaches from MAP** Developing algorithms for language convergence will provide insight into general patterns of solutions for MAPs.

### 4. REFERENCES

- [1] Kiran Lakkaraju and Les Gasser. A unified framework for multi-agent agreement. In *Proceedings of AAMAS '07*, 2007.
- [2] L. Steels and J. De Beule. Unify and merge in fluid construction grammar. In Nehaniv L. Lyon, C. and A. Cangelosi, editors, *Emergence and Evolution of Linguistic Communication*, Lecture Notes in Computer Science. Springer-Verlag, Berlin, 2006.

# Cooperation and Competitiveness of Intelligent Agents in Multi-Agent Systems

Raz Lin  
Computer Science Department  
Bar-Ilan University  
Ramat-Gan, Israel  
linraz@cs.biu.ac.il

## 1. INTRODUCTION

Software agents offer great advantages for people. They can be used for training, as well as for assisting in many environments and tasks. In this work we study the design of agents in three distinct contexts.

First, we consider the design of automated negotiators. Various tasks in day-to-day life require negotiation capabilities. Frequently negotiations are conducted by ordinary people who have not acquired formal training in this field. The aim of our work is not to replace humans in negotiations, but rather to develop an automated negotiation agent that will enable training in real daily life negotiations, such as e-commerce. We show that it is indeed possible to design an efficient agent for bilateral negotiations with humans. Our results indicate that our agent plays better than humans as well as better than other automated negotiation agents.

Second, motivated by the recent emergence of fourth generation network technology and the expected trends which include both an increase in the number of end users and convergence of a variety of services, such as IP Multimedia Subsystems, email and instant messaging, we investigate the effects of agents in the cellular network domain. Still, there is a continual need for efficient resource utilization and network maintenance. We have developed general distributed algorithms and heuristics that allows short-term self-adjustment capabilities. Using simulations on a real cellular network deployed in a large central European country we have shown that if agents use distributed negotiations global performance of the network can be effectively improved. In addition, our development is one of the first integrated simulation environments for cellular networks with an agent oriented paradigm. While most simulations use theoretical models, we have built a realistic reconstruction of a deployed cellular network, while demonstrating the benefits of reactive intelligent agents in settings as close to reality as possible. Thus, our simulation environment can serve as a test-bed for numerous aspects of artificial intelligence and agent-based mechanisms in cellular networks, far beyond our proposed negotiation tool.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.  
Copyright 2007 IFAAMAS .

In the third context, we focus on self-interested agents in complex networks such as transportation networks. Vehicle-To-Vehicle (V2V) communication is already on-board among some car manufactures, enabling the collaboration between different cars on the road. In this work we investigate the attraction of being a selfish agent in vehicular networks. That is, we investigate the benefits achieved by car owners, who tamper with on-board devices and incorporate their own self-interested agents in them for their own benefits. Using simulations which models a real vehicular network in a large city we investigated behaviors that might be taken by self-interested agents. Our simulations indicate that the self-interested agents have only limited success achieving their goal, even if no counter-measures are taken. This is as opposed to the greater impact inflicted by self-interested agents in other domains (e.g., E-Commerce). Some reasons for this are the special characteristics of vehicular networks and their dynamic nature. While the self-interested agents spread lies, they cannot choose which agents with whom they interact. Also, by the time their lies reach other agents, they might become irrelevant, as more recent data has reached those agents.

## 2. AN AUTOMATED AGENT FOR BILATERAL NEGOTIATIONS WITH BOUNDED RATIONAL AGENTS WITH INCOMPLETE INFORMATION

Many tasks in day-to-day life involve interactions among several people. Many of these interactions involve negotiating over a desired outcome. Often the negotiation can be as simple and ordinary as haggling over a price in the market or deciding what show to watch on TV, though it can also involve issues over which million of lives are at stake, such as resource allocation and interstate disputes.

In this research we propose a model of an automated negotiation agent capable of negotiating with bounded rational agents under conditions of incomplete information. We have tested the agent against people in two distinct domains, in order to verify that its model is generic, and thus can be adapted to any domain as long as the negotiators' preferences can be expressed in additive utilities. One of the domains was based on an international crisis, and the subjects had to play a role that was outside of their normal experience. On the other hand, the second domain was more related to the subjects experience, thus they were better able to identify with the negotiations.

Our results [4] indicate that the agent reached more agree-

ments and played more effectively than its human counterparts. Moreover, in most of the cases, the automated agent played significantly better than the human counterparts.

### 3. TOWARDS THE FOURTH GENERATION OF CELLULAR NETWORKS: IMPROVING PERFORMANCE USING DISTRIBUTED NEGOTIATIONS

We have developed a novel programmatic approach to efficiently and intelligently distribute resources in a dynamic cellular network, using local negotiations. Our proposed mechanism is reactive and facilitates parallel self-adaptation efforts, leading to dynamics that improve overall network performance. The local nature of the negotiations is performed as part of the adaptation process enables frequent changes in the network's parameters with a negligible coordination overhead. This mechanism is important when a frequent global optimization is infeasible or substantially costly. The proposed local negotiation mechanism is incorporated into a simulated network based on cutting-edge industry technologies.

Our simulations were run on a simulation of a real cellular network deployed in a large central European country. Our current results [1, 2] show that our mechanism enables improvement in the performance of large cellular networks with dynamic changes. Thus far we have shown success in multi-attribute optimization (a framework to negotiate over several arguments whose success has been demonstrated via simulations) and have demonstrated the approach on a large network of 300 base stations.

Our future work will focus on the development of algorithms for improving the performance of cellular networks using distributed negotiations when fault components are detected. We will also develop dynamic algorithms for online bandwidth allocation between different service providers. The success of our research could motivate the creation of standards (e.g. IEEE) to enable the dynamical sharing of resources and negotiation over bandwidth between different service providers.

### 4. ON THE BENEFITS OF CHEATING BY SELF-INTERESTED AGENTS IN VEHICULAR NETWORKS

Along this line of research we investigated the effects caused by self-interested agents in a transportation network. As more and more cars are equipped with GPS and Wi-Fi transmitters, it becomes easier to design systems that will allow cars to interact autonomously with each other, e.g., regarding traffic on the roads. Indeed, car manufacturers are already equipping their cars with such devices. Though, currently these systems are a proprietary, we envision a natural evolution where agent applications will be developed for vehicular systems, e.g., to improve car routing in dense urban areas. Nonetheless, this new technology and agent applications may lead to the emergence of self-interested car owners, who will care more about their own welfare than the social welfare of their peers. These car owners will try to manipulate their agents such that they transmit false data to their peers. Using a simulation environment, which models a real transportation network in a large city, we demonstrate the benefits achieved by self-interested agents if no

counter-measures are implemented [3].

Motivated by our results, we argue that a way to distinguish between false messages and correct ones is in order. This leads to two main problems, for which we propose several mechanisms. The first is the distributed reputation mechanism that should enable each agent to propagate information, including its belief regarding others in the network. The second issue is how each agent can form and update its belief about others in the network. Distributed reputation has been investigated in the literature in the context of mobile ad-hoc networks (MANET), Peer-to-Peer networks and Byzantine networks. However, gossip in VANET is quite different from those environments, due to its unique characteristics. In particular, an agent cannot choose with whom to interact and the communication and direct data exchange of an agent is only done with a small portion of the total agents in the network. Our future research will involve the development of a distributed reputation mechanism for gossip in vehicular networks and the incorporation of a belief mechanism formation within the agents. We assume that the identity of the senders cannot be forged (e.g., using public keys). We aim to use the mechanism in order to test interesting key questions in respect to our domain:

- Does the protocol assist in detecting agents that propagate false information and how well does it characterize them or false characterize true telling agents?
- What happens if there are special agents guaranteed to tell the truth (e.g., ambulance/police agents; public transportation agents)?
- What is the minimum amount of information needed to be propagated for the protocol to perform well?

We will test our proposed method based on an existing simulation tool, which simulates the vehicular network of Jerusalem, with 50 junctions and 150 roads. This simulation tool will allow us to record various data about the simulation itself and each of the agents (e.g., journey length, propagation of data, etc.). The results will also be compared to simulations in which no distributed reputation or belief formation is used.

### Acknowledgments

This work is done under the supervision of Prof. Sarit Kraus, Computer Science Department, Bar-Ilan University, Israel.

### 5. REFERENCES

- [1] R. Lin, D. Dor-Shifer, S. Kraus, and D. Sarne. Local negotiation in cellular networks: From theory to practice. In *IAAI*, 2006.
- [2] R. Lin, D. Dor-Shifer, S. Rosenberg, S. Kraus, and D. Sarne. Towards the fourth generation of cellular networks: Improving performance using distributed negotiation. In *MSWiM*, pages 347–356, 2006.
- [3] R. Lin, S. Kraus, and Y. Shavitt. On the benefits of cheating by self-interested agents in vehicular networks. In *AAMAS*, 2007.
- [4] R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry. An automated agent for bilateral negotiation with bounded rational agents with incomplete information. In *ECAL*, pages 270–274, 2006.

# Task Decomposition and Delegation Algorithms for Coordinating Unstructured Multi Agent Systems

António Luís Lopes

We, the Body, and the Mind Research Lab of  
ADETTI-ISCTE

Avenida Forças Armadas, Ed. ISCTE  
1600-082 Lisboa, Portugal

antonio.lopes@we-b-mind.org

Luís Miguel Botelho

We, the Body, and the Mind Research Lab of  
ADETTI-ISCTE

Avenida Forças Armadas, Ed. ISCTE  
1600-082 Lisboa, Portugal

luis.botelho@we-b-mind.org

## ABSTRACT

This paper describes the initial PhD research work for developing an innovative process for agent coordination in unstructured distributed environments. In this proposal, an agent that receives a request will use an AI-based planning algorithm to decompose the received request into several subtasks; it will handle those for which it has know-how and available resources; and it will use decentralized networks computing methods to delegate the rest of the request to other agents that can possibly handle it. These, in turn, will recursively apply the same process until the initial request is completely solved or until some specified termination condition holds.

## 1. INTRODUCTION

The main goal of the research described in this paper is to develop a robust agent architecture that enables agents to freely participate in distributed problem solving in open scalable agent societies. In such societies, agents receiving requests from other agents would be capable of using their own capabilities to handle the parts of the problems for which they have competence and resources, and deliver the remaining parts to other agents that can possibly handle them more appropriately. Agent architectures with such mechanisms would enable the development of more autonomous robust artificial agents capable of acting in open flexible, dynamic and scalable agent societies. This work combines artificial intelligence, distributed problem solving and computational paradigms from unstructured decentralized networks in order to address its two main challenges: partial task decomposition and coordinated task delegation over unstructured decentralized environments. One of the objectives is to overcome the limitation of current planning algorithms by developing a planning algorithm that is capable of creating plans containing action slots for undetermined actions, enabling an agent that is trying to solve a problem, to delegate the part for which it cannot contribute, to other agents. All known planning algorithms must know all available planning operators (all possible actions) to generate their plans. However we need an algorithm that generates a plan in which part of the planned actions are known to the algorithm whereas others are only action slots for yet to be determined actions.

In order to provide an infrastructure in which agents can cooperate in this kind of environments, this work also addresses the problem of coordinated task delegation over unstructured decentralized networks. The objective is to create a task distribution algorithm that guarantees the consistency of the distributed problem solving. Current work in this area [8] involves

either centralized components (which compromise scalability) or the excessive use of communication for synchronizing agents. Another objective of this work is to avoid unnecessary communication used in synchronization (either to avoid having tasks being performed twice or to avoid execution errors in a specific sequence of actions) through optimization of task delegation based on previous interactions.

## 2. RELATED WORK

Distributed problem solving has been defined as the cooperative process of solving problems by a decentralized collection of knowledge sources located in a number of different processor nodes [2]. The domain can be characterized as a collection of agents, each with different skills and limited knowledge, which acts to solve problems in a distributed fashion. The main challenges are [5]: describing the problem; decomposing and allocating sub-problems to agents, designing mechanisms for agents to interact; and ensuring coherence, that is, making sure that the agents make decisions and perform tasks that get them closer to the solution. The focus of this research is on the decomposition of tasks and corresponding coordination mechanisms that ensure the means for agents to cooperate in problem solving, using only their own capabilities.

Planning algorithms use methods to decompose tasks recursively into smaller subtasks, until they reach primitive tasks that can be performed directly. STRIPS-like [4] representations have been the most common logic-based representations for discrete planning problems. Another well-known representation and planning algorithm is based on Hierarchical Task Networks (HTN) [3], which introduced powerful ideas such as task decomposition. However, HTN planning systems require full knowledge of the available planning operators (all agents' capabilities) and involve centralized solutions such as service composition agents. These two requirements have been deemed unacceptable in partially inaccessible large-scale environments. Partial task decomposition, i.e. the possibility for an agent to contribute only to a part of the proposed problem instead of having to contribute with a complete plan, provided within a distributed problem solving coordinating infrastructure, may be the answer to this challenge. The main advantage of this approach lies on the fact that the entire decomposition process emerges as the result of a collection of local decomposition processes and that an agent does not need to know the other agents' capabilities.

Efficient coordination among large numbers of heterogeneous agents promises to revolutionize the way in which some complex tasks can be performed. However, state-of-the-art coordination

approaches [6] are not capable of achieving efficient and effective coordination when a team is very large, since centralized discovery and composition processes need to be fuelled with agents' new capabilities in continuously growing environments, dramatically decreasing the performance of the entire system. Also, these approaches spend large amounts of time finding complete solutions that will fail, because they can only solve part of the problem. Recent work focusing on scalable coordination [10] illustrates that exponential search spaces, excessive communication demands, localized views, and incomplete information of agents pose major problems for large scale systems. The main limitations are based on the fact that these approaches rely on partial, dynamic centralization [1][6], such as brokers [7]. Distributed constraint-based algorithms [9] have high communication requirements that get dramatically worse as the team size increases. Therefore, there is a clear need for alternative methods for large-scale coordination of agents' cooperation in distributed problem solving.

### 3. TECHNICAL APPROACH

One of the focuses of this research is the development of the coordination infrastructure that will enable agents to cooperate in an unstructured decentralized environment. Two possible approaches are considered for the task delegation coordination infrastructure:

**One-phase-approach:** Delegated tasks will be decomposed into subtasks that will be immediately executed or recursively delegated until the whole task is completed. This approach requires a great amount of synchronization to avoid having two agents performing the same task or tasks that depend on others that were not performed yet. Besides, it may and will happen that agents will perform subtasks that are part of a plan that will later fail.

**Two-phase-approach:** Previously, agents are asked if they will ensure the task proper execution. The initial request is decomposed into subtasks, some of which can be handled by the agent that receives the request. Then, the agent inquires other agents if they would be capable of executing the remaining subtasks. These, using an analogous process, will proceed until an execution path is found that ensures the completion of the whole initial task. This approach does not require the described synchronization but it will possibly have a longer execution time for the entire task. Besides, in highly dynamic environments, an agent that has previously ensured the execution of a given subtask may no longer be capable of executing it in a latter moment.

Both approaches can be used in non-cooperative environments, where agents that receive a task may first want to negotiate possible rewards for its decomposition, delegation and execution. If this happens, the agent that negotiates task delegation will have to choose the best deal. We will analyse and implement both approaches so that comparative data can be acquired on their relative strengths and weaknesses.

After deploying and testing the task delegation coordination infrastructure, the focus of the research will be on the definition and development of the task decomposition algorithm. The main challenge is to provide a way for agents to produce partially-specified plans, that is, plans that contain actions that are known to the agent, which it can perform and slots for actions yet

to be determined, which will be delegated to other agents. The generated plans will contain totally refined parts and parts yet to be refined. The major contribution of this algorithm is the ability to consider both the agent concrete capabilities, which will be known to the algorithm, and the other agents' abstract capabilities, which will not be known to the algorithm. That is, an algorithm capable of generating plans containing slots for actions that are not known to the algorithm. The plan will only become totally refined by the recursive approach of delegating the non-refined parts to other agents, based on the coordinated task delegation infrastructure described in the previous sub-section.

### 4. REFERENCES

- [1] Davin, J., and Modi, P. J., Impact of problem centralization in distributed constraint optimization algorithms. In *Proceedings of the Fourth international Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 1057-1063, 2005
- [2] Decker, K. S., Distributed problem solving: A survey, *IEEE Trans. on Systems, Man, and Cybernetics*, 17(5):729-740, 1987.
- [3] Erol, K., Nau, D. and Hendler, J., HTN Planning: Complexity and Expressivity, In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1123-1128. Menlo Park, Calif.: American Association for Artificial Intelligence, 1994.
- [4] Fikes, R. E. and Nilsson, N. J., STRIPS: A new approach to the application of theorem proving, *Artificial Intelligence Journal*, 2:189-208, 1971.
- [5] Harandi, M. T. and Rendon, G., Classification and organizational issues in distributed problem solving, In *Proceedings of the 1998 ACM Symposium on Applied Computing*, pp. 65-69, 1998.
- [6] Helin, H., Klusch, M., Lopes, A., Fernandez, A., Schumacher, M., Schuldt, H., Bergenti, F., and Kinnunen, A., CASCOS: Context-Aware Service Co-ordination in Mobile P2P Environments, *Multiagent System Technologies, Lecture Notes in Computer Science, Vol. 3550 / 2005, ISSN: 0302-9743*, pp. 242-243, 2005.
- [7] Hunsberger, L. and Grosz, B., A combinatorial auction for collaborative planning, In *Proceedings of the Fourth International Conference on MultiAgent Systems*, pp. 151-158, 2000.
- [8] desJardins, M. E., Durfee, E. H., Ortiz, C. L. Jr., Wolverton, M. J., Survey of research in distributed, continual planning, *Artificial Intelligence Magazine. Vol. 20, no. 4*, pp. 13-22, 1999.
- [9] Mailler, R. and Lesser, V., Solving distributed constraint optimization problems using cooperative mediation, In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, 2004.
- [10] Scerri, P., Vincent, R. and Mailler, R., Comparing three approaches to large scale coordination, In *Proceedings of the Workshop on Challenges in the Coordination of Large Scale Multi-Agent Systems of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, 2004

# A Framework For Ontology-Based Service Selection

Murat Şensoy  
murat.sensoy@boun.edu.tr

Department of Computer Engineering  
Boğaziçi University  
Bebek, 34342, Istanbul, Turkey

## ABSTRACT

Previous approaches to service selection are mainly based on capturing and exchanging the ratings of consumers to providers. However, ratings reflect tastes of the raters. Therefore, service selection using ratings may mislead the consumers having a taste different than that of the raters. We propose to use experiences instead of the ratings. Experiences are the representation of what is requested by a consumer and what is received at the end. Unlike ratings, experiences do not reflect the opinion of the others. In fact experiences reflect the actual story between consumers and providers concerning a service demand. Using experiences, the consumer models the services of a provider for a specific service demand and selects the provider that is expected to satisfy the consumer the most. Our simulations show that proposed approach significantly increases the overall satisfaction of the service consumers. In order to represent their experiences, consumers use ontologies. However, as the consumers' service needs evolve over time, the concepts in their ontologies may not be sufficient to represent their service demands and their experiences. If each consumer creates new concepts on its own to represent its evolving service needs and adds them to its local ontology, this results in highly different ontologies. Hence, the consumers could not understand each other any more using the concepts in their ontologies. Accordingly, we propose an approach through which consumers can cooperatively update their ontologies and teach one another concepts from their ontologies. Our simulations show that the proposed approach leads to a society of consumers with different but overlapping ontologies. Moreover, mutually accepted concepts emerge based on the interactions of the consumers.

## 1. INTRODUCTION

Previous approaches to service selection are mainly based on ratings. *Reputation systems* enable consumers to rate the service providers in a centralized location. The ratings of the consumers are then aggregated to decide whether a service provider will act as expected [6]. E-bay [4] uses such a reputation system. Distributed approaches to service selection consider *trust* among entities [8, 7, 5]. Trust captures a trustor's expectation from a trustee for a

particular service. Most formalizations of trust depend on ratings. However, ratings should be evaluated within their scope. Scope of a rating is the context in which the rater experienced the service. In different contexts, a service has different values for a consumer (i.e., it is rated differently by the consumer). Therefore, scope of ratings should be considered while using the ratings.

Although scope of the ratings are expressed clearly and considered carefully during service selection, ratings may still mislead service consumer during their service decisions. That is, ratings reflect the subjective opinion of the raters. While using ratings, consumers make decisions depending on the satisfaction criteria of the raters. If satisfaction criteria of the raters are similar to that of the consumer, then the consumer may make satisfactory service decisions; otherwise the consumer comes up with unsatisfactory service decisions. Therefore, rating-based approaches suffer from the inherent subjectiveness of ratings. In order to get rid of subjectiveness originated from the others' satisfaction criteria, a service selection approach should not primarily depend on ratings. Hence, main goal of this research is to develop a context-aware service selection framework so that consumers could make the context of their service demands explicit and make service decisions primarily on their own satisfaction criteria.

## 2. COMPLETED RESEARCH

Even if their service interests are the same, different consumers rate the same service differently depending on their satisfaction criteria. That is why ratings reflect the satisfaction criteria and subjective opinion of the raters. The main question at this point is how to get rid of subjectiveness of ratings in service selection. In our completed research [2, 1], we propose that, instead of ratings, consumers can record their interactions with service providers in a great detail within an experience structure. An experience contains a consumer's service demand and the supplied service in response to this service demand. In our approach, instead of sharing their ratings, consumers share their experiences. An experience expresses the story between the consumer and the provider regarding a specific service demand. Hence, any consumer receiving an experience can evaluate the service provider according to its own satisfaction criteria using the detailed data in the experience. This approach may get rid of the subjectiveness of the rating-based approaches.

Although ratings can simply be represented using plain numbers, representation of consumers' experiences requires representational power of ontologies. In our completed research, we assume that consumer agents share a static shared ontology through which the consumers can represent their service needs and past experiences in the best way possible. Consumers use experiences to decide on a service provider. For this purpose, consumers first collect ex-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07, May 14–18 2007, Honolulu, Hawaii.  
Copyright 2007 IFAAMAS .

periences related to their current service demands from other consumers using a P2P protocol. Then, information in the experiences is used for the modeling of providers' behaviors for different service demands (i.e., using machine learning methods such as parametric classification [2] or case-based reasoning [1]). After modeling service providers, the provider having the highest probability of supplying satisfactory service is chosen.

Using various simulations, we compared performance of the proposed approach and that of various rating-based service selection approaches. The performed simulations show that the use of experiences improves the decisions of the service consumers and increases the overall satisfaction significantly compared with the rating-based service selection approaches.

### 3. THE WORK IN PROGRESS

Although the proposed experience-based service selection approach is successful in identifying service providers, it has a major weakness: assumption of a static, shared ontology among agents. This assumption cannot account for the fact that a consumer's service needs may evolve over time and new concepts may be necessary for the consumer to describe its evolving service needs. Hence, a service selection approach should be able to accommodate this, since in many e-commerce settings, individuals learn new concepts and services from different sources, add them to their ontology, and further form service requests that are based on these new concepts.

Accordingly, we propose a distributed approach for evolution and maintenance of local ontologies for service selection [3]. In this approach, a consumer queries its neighborhood to learn a suitable concept if the concepts in its local ontology are not sufficient to describe its current service need. If such a concept is not known by its neighbors, the consumer creates the concept, and teaches it to its neighborhood. This way, the consumer prohibits its future communication problems by informing its neighbors about the created concept before using it. Furthermore, if the neighbors find the new concept useful, they can use it in their forthcoming interactions with others. This way, the semantics of new concepts circulate and get established in the consumer society. This interaction-based learning scheme leads to cooperative evolution of service ontologies. When a consumer learns a useful concept from its neighbors, it can directly use it or create another concept that builds on the learned concept. Hence, more accurate concepts that describe the service needs and experiences of the consumers are cooperatively and iteratively created. Integration of this approach into our previously proposed experience-based service selection approach is in progress.

We evaluated our approach using simulations. Our simulations show that consumers' ontologies evolve considerably so that at the end of the simulations each consumer is able to represent its service needs and past experiences as concisely as possible. The evolutions are not only due to individual effort but mostly a result of cooperation. That is, most of the concepts used by each consumer are devised by others in the society. Further, the consumers learn a small portion of the concepts that emerge in the society, making sure that most of the learned concepts are useful for them in representing their service needs. Through cooperation, consumers can have different but evolving ontologies, yet they can communicate with those that are similar to them to represent their service needs.

### 4. THE WORK TO BE DONE

The proposed approaches highly depend on domain dependent ontologies. These ontologies are used to represent service demands and the provided services. In our research, ontologies are repre-

sented using OWL, but standard service description languages such as WSDL or OWL-S are not used. In order to support interoperability and for compliance with current Web services standards, the service ontologies should be represented using OWL-S.

In this work, we have assumed that consumers exchange their experiences honestly. However, experiences are produced by consumers and some consumers may want to defame some providers because of personal or commercial reasons. This situation imposes the requirement of using a trust mechanism so that some consumers could be ranked as trustful and others could be ranked as distrustful (cheater or slanderer). As a result, we plan to integrate a collaborative trust mechanism into our service selection framework.

As a future work, we also plan to deal with the incentive problem: i.e., willingness of participants to provide information such as experiences. Sharing experiences is important for better decision making and it creates reciprocity between agents. However, additional mechanisms are needed to encourage agents for sharing their experiences and other valuable information.

### 5. REFERENCES

- [1] M. Şensoy and P. Yolum. A comparative study of reasoning techniques for service selection. In *Proceedings of the Fifth International Workshop on Agents and Peer-to-Peer Computing (AP2PC)*, pages 91–102, 2006.
- [2] M. Şensoy and P. Yolum. A context-aware approach for service selection using ontologies. In *Proceedings of Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 931–938, 2006.
- [3] M. Şensoy and P. Yolum. Cooperative evolution of service ontologies. In *Proceedings of Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2007.
- [4] eBay. <http://www.ebay.com>, 1995.
- [5] T. D. Huynh, N. R. Jennings, and N. Shadbolt. Fire: An integrated trust and reputation model for open multi-agent systems. In *Proceedings of 16th European Conference on Artificial Intelligence*, pages 18–22, 2004.
- [6] J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 475–482, 2002.
- [7] P. Yolum and M. P. Singh. Engineering self-organizing referral networks for trustworthy service selection. *IEEE Transactions on Systems, Man, and Cybernetics*, A35(3):396–407, 2005.
- [8] B. Yu and M. P. Singh. Emergence of agent-based referral networks. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1268–1269, 2002.

# A Hybrid Declarative Framework for Multi-Agent Systems (HDFMAS)

Mandeep Singh  
Universität Tübingen  
Wilhelm-Schickard-Institut Für Informatik  
Sand 13, D-72076  
Tübingen, Germany  
singh@informatik.uni-tuebingen.de

## ABSTRACT

The first part of this paper describes the concerned research area and goal of my Ph. D. thesis. The second part gives the outline of the research programme.

## 1. GOAL STATEMENT

This thesis deals with issues faced by developers in the areas of *agent architectures* and *agent languages* as introduced by Wooldridge and Jennings [9]. In heterogeneous distributed environments information may be stored and managed in various approaches, for example in tables of databases or in data objects of some programming language. The integration of heterogeneous information systems, databases or application software to facilitate the information exchange, is one of the major interoperability problems [6]. In an agent dialogue, even if all the agents understand the same declarative language, evaluation of queries is difficult to handle due to heterogeneous information systems. In heterogeneous environments agents should be capable of exchanging complex information, such as their questions, intentions, plans, and strategies. As already pointed out in [1] an *Agent Communication Language(ACL)* must be declarative and have at least a small number of primitives that are necessary to support agent interoperability.

As ACLs lack in their declarative nature, developed software agents for distributed systems are forced to be dedicated and have frozen code for pre-planned execution, all this hampers the agents to function autonomically in open heterogeneous systems. The *Foundations of Intelligent Physical Agents(FIPA)* ACL gives the programmers the freedom to set any expression as content. Taking this advantage KQML [1] can be used to send a declarative expression as message's content. In this way interoperability can be augmented to some extent, however transformation of KQML to other specific languages fails at the moment.

If we follow the *Model Driven Development(MDD)* me-

thodology, transformations/translators can be used to convert declarative expression in to or from a specific language expression, after that a language-specific interpreter can be used for information processing and evaluation. The goal of this work is to develop a *Hybrid Declarative Framework for Multi Agent Systems(HDFMAS)*, while considering the approaches used in METATEM [5], MALLETT [4], Tapir [8], KQML, and *Agent Intensional Programming Language(AIPL)*. HDFMAS will offer the agent developers an abstraction of information query languages, which will reduce the hard coding and enhance the automation process during agent interaction in runtime. Clearly applicable areas of this framework are heterogeneous information systems. It will support the integration of *Web Service Applications* with *Multi-Agent Systems(MAS)*. Even *intentional systems* [9] will directly benefit from this framework in delegation scenarios.

Figure 1 shows some details of HDFMAS and illustrates an application scenario. In this scenario agents still communicate in the ACL format, however they use a declarative language to query information and tell their intentions. The declarative query engine handles the request and uses the translations and interpreters to transform and interpret declarative expression. In the same way the response can be sent as a declarative expression and interpreted with different interpretation rules.

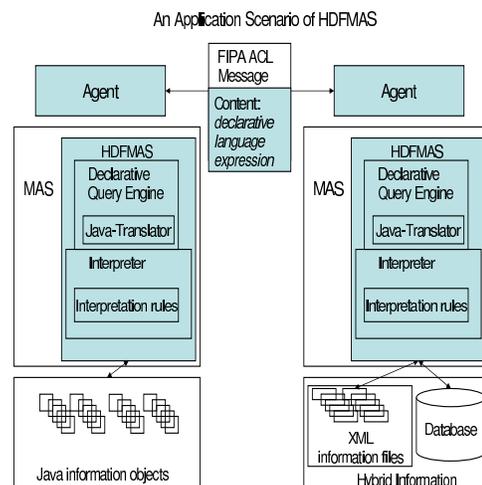


Figure 1: Using HDFMAS with MAS

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.  
Copyright 2007 IFAAMAS .

## 2. RESEARCH PROGRAMME

The research programme is roughly divided into *Analysis*, *Implementation* and *Testing*. The *Analysis Phase* and the *Implementation Phase* are currently in progress.

### 2.1 Analysis Phase

This phase will continue to analyse the state-of-the-art of different technologies that are relevant to this thesis, until the goal of Ph.D. is not reached, so that the existing concepts can be integrated with enhancements all the time with the framework. The planned surveys are *Multi-Agent Systems*, declarative languages, policy definitions languages, management definition languages, query languages, and the role of logic in declarative languages. Survey reports will help to design the concept of HDFMAS for the best possible integration to MASs and declarative languages. MASs like JADE, JACK, and JASON will be analysed from concept, design, and implementation point-of-view. The analysis of definition languages like KQML, AIPL, KAoS, Rei, Ponder and MDL will give the abstraction of these languages, so that the abstraction can be used to declare the content of messages during agent interaction. Ontologies and ACLs are unseparable companions when talking about agent communication, so concepts of *Knowledge Sifter* [7], *Knowledge Interchange Format*(KIF), and *Web Ontology Language*(OWL) are also partly of interest. This analysis will identify the role of logic in declarative languages, which will help to interpret the intention of agent's message.

### 2.2 Implementation Phase

The information gathered about the technologies in the analysis phase will govern the conceptual design and architecture. In this phase a generic architecture of the framework will be designed. The next step will be the realisation of a specific architecture with technologies and standards such as FIPA, JADE, SQL, and J2EE as proof-of-concept. The query-translator-prototypes will be written in Java, which will convert the Java query expression into Prolog query expressions and SQL query expression. To allow the agents to evaluate the logical expressions interpreter-prototypes will be developed in Prolog. The declarative language of the HDFMAS is the core part of this thesis, so the reports of the survey from the analysis will be compared very carefully with the survey [2] to choose or develop a new declarative language for the framework. The first platform-specific in the sense of *Model Driven Architecture*(MDA) implementation will produce translators for Java, SQL, Prolog, XML, and use JADE as an agent platform. The interpreters and interpretation rules are also part of the first version. Interpreters are needed to understand the intension of agents, for example, in purchasing scenarios different intentions of buyer and seller can be identified. The ontology interface of JADE will be used to access the structural information in runtimes.

### 2.3 Testing/Verification Phase

In this phase test scenarios will be developed to exploit the functionality of this framework and test cases will be build to test the code. The verification of the framework is planned to be done with CASP [3].

## APPENDIX

### A. REFERENCES

- [1] V. S. Alagar, J. Paquet, and K. Wan. *Intentional Programming for Agent Communication*. 2005.
- [2] R. Bordini, L. Braubach, M. Dastani, A. Seghrouchni, J. Gomez-Sanz, J. Leite, G. O'Hare, A. Pokahr, and A. Ricci. A survey of programming languages and platforms for multi-agents systems. *Informatica*, 30, 2006.
- [3] R. H. Bordini, W. Visser, M. Fischer, C. Pardavila, and M. Wooldridge. *Model checking multi-agent programs with CASP*, volume 2725/2003. 2003.
- [4] X. Fan, J. Yen, M. Miller, and R. Volz. *The Semantics of MALLET- An Agent Teamwork Encoding Language*. 2005.
- [5] M. Fischer. A survey of concurrent metatemp- the language and its applications. pages 480–505, 1994.
- [6] J. Galliers. A strategic framework for multi-agent cooperative dialogue. 1988.
- [7] L. Kerschberg, M. Chowdhury, A. Damiano, H. Jeong, S. Mitchell, J. Si, and S. Smith. Knowledge sifter:ontology-driven search over heterogeneous databases. page 431.
- [8] G. W. King, M. S. Atkin, and D. L. Westbrook. Tapir:the evolution of an agent control language. *First Conference on Autonomous Agents and Multiagent Systems*, 2002.
- [9] M. J. Wooldridge and N. R. Jennings. Agents theories, architectures, and languages: A survey. *Intelligent Agents*, 94.

# Flexible Provisioning of Service Workflows (Thesis Summary)

Sebastian Stein  
School of Electronics and Computer Science  
University of Southampton  
Southampton, SO17 1BJ, UK  
ss04r@ecs.soton.ac.uk

## 1. PROBLEM OUTLINE

As computer systems are becoming increasingly distributed and complex in nature, services are emerging as a promising paradigm for offering processes and resources on demand to consumers. By relying on machine-readable service descriptions, software agents in such systems are able to discover and invoke remote services without human intervention at run-time. This allows them to achieve complex goals, often by combining many atomic services into larger workflows.

Current work on service-oriented computing has so far concentrated on knowledge representation techniques, reasoning over service descriptions and appropriate communication protocols. Such work typically assumes that service providers are reliable and that they publish accurate and truthful service descriptions. However, this assumption is unrealistic in large distributed systems, such as the Web, large-scale computational Grids or peer-to-peer systems. These systems are inherently dynamic and uncertain — services may fail frequently due to hardware problems, bugs or network failures. Furthermore, the availability of services can change rapidly as providers leave or enter the system, and competition for resources can result in uncertain execution times. Additionally, service providers are increasingly likely to be autonomous agents with their own decision-making mechanisms. This means that they may prioritise service requests or even ignore them if that is in their interest.

The resulting uncertainty and unreliability must be addressed when designing software agents that execute complex workflows in service-oriented systems. This issue is particularly critical in scenarios where workflows represent a real value to the consumer, where they must be completed within a limited amount of time, and where service providers demand remuneration. In these cases, any service failures or even delays can jeopardise the overall outcome and result in losses to the consumer.

In my thesis, I am addressing these challenges by considering the *provisioning* of service workflows. During provisioning, the service consumer decides which providers to invoke for the constituent tasks of an abstract workflow. This allows it to choose more reliable providers where necessary and balance the overall cost of executing a workflow with its value. Furthermore, the consumer can address service failures and uncertainty proactively by provisioning multiple providers in parallel for particularly failure-prone tasks, and, when necessary, it can re-provision providers reactively in the case of failure.

To make these decisions autonomously at run-time, my work uses decision theory to maximise the consumer agent's expected utility. However, as the required calculations are generally intractable, I focus on efficient and scalable heuristics that make fast decisions in realistic settings. To evaluate the work, I conduct thorough empirical experiments and compare it to various benchmarks.

## 2. RESEARCH PLAN

The overall time-plan for my PhD is shown in Figure 1 and contains six high-level tasks that represent substantial pieces of research work. Each of these corresponds to one chapter in the final thesis and deals with different aspects of the overall research problem. The next sections describe these in more detail, starting with the background work (tasks 1 and 2), the research completed so far (tasks 3 and 4) and finally my future work (tasks 5 and 6).

### 2.1 Background Work

During the initial period of the thesis, I conducted a thorough literature review of related work. This included general background reading on service-oriented computing, Web services, Grid computing, Semantic Web technologies and multi-agent systems, as well as more specific work on service selection, composition and quality-of-service issues. This allowed me to identify the current state of the art and formulate my overall research problem.

Based on the literature review and the research problem, I then devised an abstract system model to serve as the basis for my research. In particular, this enabled me to formalise the behaviour of service providers and the types of uncertainty my work addresses (initially, this included service failures and uncertain execution times). Furthermore, I developed a simulation for evaluating my work using a number of benchmarks strategies.

### 2.2 Completed Research

Having concluded the background work, I began to consider a realistic workflow scenario based on systems where services are invoked “on demand” (i.e., without advance reservation, as is common for current Web services). To address service uncertainty in this case, I developed a heuristic strategy that provisions service providers flexibly using some domain knowledge about the workflow tasks. In particular, the strategy decides when to provision multiple providers for unreliable tasks, and when to re-provision providers that seem to have failed (even when providers do not signal their failures). These decisions are made using a local search tech-

Task	Progress	2005				2006				2007				2008		
		Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2
Literature Review	complete	██████████														
System Model and Implementation	complete	██████████														
Provisioning On-Demand Services	complete	██████████														
Heterogeneous Services	complete	██████████														
Advance Negotiation	in progress	██████████				██████████										
System Dynamism	to be done	██████████				██████████				██████████						
Thesis Write-Up	to be done	██████████				██████████				██████████				██████████		

Figure 1: Time-plan for thesis (dotted line indicates current position).

nique and a heuristic utility prediction function that estimates the failure probability, duration and cost of a workflow execution. In empirical experiments, I have shown that the strategy achieves an approximately 350% improvement over current approaches that do not address uncertainty.

The initial strategy was published in [1]. This was followed by further work on the sensitivity of the strategy in the presence of inaccurate information [7] and on an improved prediction mechanism [5]. To put the work into the context of related work on Semantic Web services, [4] explores the use of an appropriate ontology to express performance information. An extended journal article about the strategy with further results is currently under review [6].

During the next stage of my research, I began looking at a more complex scenario, where services are offered by heterogeneous providers. This is typical in large multi-agent systems, where many providers may offer the same type of service at different levels of reliability and price. To this end, I changed the system model and developed a new strategy to provision multiple heterogeneous providers for each task of a workflow. This required substantial modifications of the heuristic utility function and I had to consider different local search techniques to deal with the larger search space. These results are published in [2], with an extended version under review [3].

### 2.3 Future Work

I have recently started expanding my work to include more advanced interaction models that are commonly found in multi-agent systems. This part of my thesis is currently at an early stage, but I plan to consider systems where services can be reserved in advance and where various service parameters (such as cost, reliability and time of execution) are determined dynamically through a negotiation process. In these systems, service consumers will need to decide whether and when to negotiate the use of services, considering that advance reservation might afford a higher reliability, but also result in increased costs and less flexibility. Addressing this problem will form a substantial and important extension of my work, because these negotiation techniques are beginning to emerge in areas such as high-performance and Grid computing, where expensive computational resources are sold on dynamic markets.

Furthermore, I plan to consider more dynamic systems. In my work to date, I have assumed that the population of

service providers stays constant throughout the provisioning and invocation of a single workflow. However, this is not generally the case in open systems, where providers may leave or enter at any time, thus creating new opportunities or service shortages at critical moments. In this part of the thesis, I will extend the system model to include such dynamism and develop more adaptive strategies.

### 3. CONCLUSION

As described in the previous sections, I have already made significant progress with my work plan, and established a solid basis from which to move forward. Moreover, I have fulfilled all mandatory parts of the PhD programme successfully and in time. This included a number of taught courses, a progress report and a transfer thesis (both of which were assessed formally during oral vivas). Hence, my thesis is currently on track for completion in autumn 2008.

### 4. REFERENCES

- [1] S. Stein, N. R. Jennings, and T. R. Payne. Flexible provisioning of service workflows. In *Proc. 17th European Conf. on Artificial Intelligence (ECAI), Riva del Garda, Italy*, pages 295–299. IOS Press, 2006.
- [2] S. Stein, N. R. Jennings, and T. R. Payne. Provisioning heterogeneous and unreliable providers for service workflows. In *Proc. 6th Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS), Honolulu, Hawai'i, USA, 2007*. (poster).
- [3] S. Stein, N. R. Jennings, and T. R. Payne. Provisioning heterogeneous and unreliable providers for service workflows. In *Proc. 22nd AAAI Conf. on Artificial Intelligence (AAAI), Vancouver, Canada, 2007*. (submitted).
- [4] S. Stein, T. R. Payne, and N. R. Jennings. Flexible provisioning of semantic web service workflows using a QoS ontology. In *Proc. 5th Int. Semantic Web Conf. (ISWC), Athens, Georgia, USA, online supplement, 2006*. (poster).
- [5] S. Stein, T. R. Payne, and N. R. Jennings. An effective strategy for the flexible provisioning of service workflows. In *Proc. Workshop on Service-Oriented Computing: Agents, Semantics, and Engineering (SOCASE), Honolulu, Hawai'i, USA, volume 4504 of Lecture Notes in Computer Science*. Springer, 2007.
- [6] S. Stein, T. R. Payne, and N. R. Jennings. Flexible provisioning of semantic web service workflows. In *ACM Transactions on Internet Technology*, 2007. (submitted).
- [7] S. Stein, T. R. Payne, and N. R. Jennings. Sensitivity analysis of flexible service provisioning. Technical Report ECSTR-IAM07-001, University of Southampton, 2007.

# Autonomous Inter-Task Transfer in Reinforcement Learning Domains

Matthew E. Taylor  
Department of Computer Sciences  
The University of Texas at Austin  
Austin, Texas 78712-1188  
mtaylor@cs.utexas.edu

In *reinforcement learning* (RL) [12] problems, agents take sequential actions with the goal of maximizing a reward signal, which may be time-delayed. In recent years RL tasks have been gaining in popularity as learning methods able to handle complex problems. RL algorithms, unlike many machine learning approaches, do not require correctly labeled training examples and thus may address a wide range of difficult and interesting problems. If RL agents begin their learning *tabula rasa*, mastering tasks may be slow or infeasible. A significant amount of current research in RL thus focuses on improving the speed of learning by exploiting domain expertise with varying degrees of autonomy.

My thesis will examine one such general method for speeding up learning: *transfer learning*. In transfer learning problems, a *source task* can be used to improve performance on, or speed up learning in, a *target task*. An agent may thus leverage experience from an earlier task to learn the current task. A common formulation of this problem presents an agent with a pair of tasks and the agent is told explicitly to train on one before the other. Alternately, in the spirit of *multitask learning* [3] or *lifelong learning* [17], an agent could consult a library of past tasks that it has mastered and transfer knowledge from one or more of them to speed up the current task.

Transfer learning in RL is an important topic to address at this time primarily for three reasons. Firstly, RL techniques have, in recent years, achieved notable successes in difficult tasks which other machine learning techniques are either unable or ill-equipped to address (e.g. TDGammon [16], elevator control [4], Keepaway [11], and Server Job Scheduling [18]). Secondly, classical machine learning techniques are sufficiently mature that they may now easily be leveraged to assist with transfer learning. Thirdly, promising initial results show that not only are such transfer methods possible, but they can be very effective at speeding up learning.

Past research on transfer between reinforcement learning tasks have formulated multiple ways in which the source and target tasks may differ:

1. Transition function: Effects of agents' actions differ [8]
2. Reward structure: Agents have different goals [9]
3. Initial state: Agents start in different locations over time [2]
4. State space: Agents act in different environments [5]
5. Actions: Agents have different available actions [7, 14, 10]
6. State variables: Agents' state descriptions differ [7, 14, 10]

When physical or virtual agents are deployed, any mechanism that allows for faster learned responses to a new task has the potential to greatly improve their efficacy. Thus, any transfer method that is able to handle the above differences could potentially be utilized by such agents to increase their adaptability and performance when an agent must perform a new task.

With motivations similar to those of *case based reasoning* [1], where a symbolic learner constructs partial solutions to the current task from past solutions, a primary goal of transfer learning is to autonomously determine how a current task is related to a previously mastered task and then to automatically use past experience to learn faster. My thesis focuses on the following question:

Given a pair of related RL tasks that have different state spaces, different applicable actions, and/or different representative state variables, how and to what extent can agents transfer knowledge from the source task to learn faster in the target task, and what, if any, domain knowledge must be provided to the agent?

The primary contribution of this thesis will be to address the above question, demonstrating a series of techniques that are able to successfully transfer knowledge between tasks with varying degrees of similarity and given domain knowledge. There are many ways of formulating and addressing the transfer learning problem, but we distinguish this work in three ways:

1. Our methods focus on allowing differences in the action space, the state, and state variables between the two tasks, increasing their applicability relative to many existing transfer methods. However, we will show that they are also applicable when the transition function, reward function, and/or initial state differ.
2. Our methods are competitive with, or are able to outperform, other transfer methods with similar goals.
3. Our methods are able to *learn* relationships between pairs of tasks without relying on human domain knowledge, a necessity for achieving autonomous transfer.

I now enumerate the components we consider necessary to address the main question posed in this thesis.

1. **Problem Definition:** Our transfer problems will focus on using a *source task* to speed up learning in a *target task* and I will define the scope of such problems in a RL setting.
2. **Performance Metrics:** In order to measure the efficacy of our methods I have defined two transfer-specific metrics. I argue that the two metrics are appropriate for the RL domains considered and focus on the performance speedup due to transfer, rather than the performance of a particular underlying TD or policy search *base learning algorithm*.
3. **Oracle-Enabled Transfer:** One class of transfer methods considered utilize inter-task mappings. Inter-task mappings describe relations between state variables and actions in the source and target tasks; they are used so that learned knowledge in the source task can apply to a target task even when

the state and action spaces have changed. I first assume that an oracle provides mappings that are complete and correct.

- (a) Transfer via Inter-Task Mapping for Value Function Methods: Complete [14]
- (b) Transfer via Inter-Task Mapping for Policy Search Methods: Complete for neural network action selectors [15]
- (c) Transfer via Rules: In progress

4. **Learning Task Relationships:** I also consider pairs of tasks where no oracle exists and the inter-task mapping must be learned. Constructing such relationships is the primary difficulty when transferring between disparate tasks, but I plan to leverage a variety of existing machine learning techniques to assist with this process. I will demonstrate the effectiveness of these relationship-learning methods on pairs of related tasks and then combine them with the above transfer methods to achieve autonomous transfer.

- (a) Learning mappings via classification with strong assumptions: Complete [15]
- (b) Learning mappings via classification with weak assumptions: To be done
- (c) Learning mappings via semantic knowledge (either learned or provided): To be done

5. **Empirical Validation:** To validate our transfer methods, I will fully implement them in at least three domains. Success in different domains and with different implementations, which have different qualitative characteristics, will show that our methods have broad applicability as well as significant impact.

- (a) Robosoccer Keepaway Domain: Complete [14, 13]
- (b) Server Job Scheduling Domain: Complete [15]
- (c) General Game Playing [6]: To be done

In addition to these goals, I am considering at least two supplemental goals, but am actively searching for more goals so that my thesis can more fully develop our understanding of transfer:

1. **Inter-Domain Transfer:** I informally define a *domain* to be a setting for a group of semantically similar *tasks*. While many methods exist to transfer between domains, none have been shown to work between domains. In addition to showing that inter-domain transfer is feasible, I would like to show that such transfer can be done autonomously.
2. **Effects of Task Similarity on Transfer Efficacy:** All the RL tasks I consider can be parameterized and thus it is possible to make the source and target tasks more or less similar. For instance, preliminary results in the Keepaway domain show that transfer is able to improve learning, compared to learning without the benefit of transfer, when the players in the two tasks have pass actuators with different accuracies, but transfer is more beneficial when the players in both tasks have actuators with the same capabilities. Observing, and ideally predicting, how transfer degrades as the source and target tasks become more dissimilar should lead to a better understanding of the proposed transfer methods. Such heuristics could be used to determine if two tasks are “similar enough” that transfer could provide any benefit. Defining a similarity metric for tasks based on these heuristics would also potentially allow us to *construct* a source task for a given target task.

## 1. REFERENCES

- [1] A. Agnar and P. Enric. Case-based reasoning: Foundational issues, methodological variations, and system approaches, 1994.
- [2] M. Asada, S. Noda, S. Tawaratsumida, and K. Hosoda. Vision-based behavior acquisition for a shooting robot by using a reinforcement learning. In *Proc. of IAPR/IEEE Workshop on Visual Behaviors-1994*, pages 112–118, 1994.
- [3] R. Caruana. Learning many related tasks at the same time with backpropagation. In *Advances in Neural Information Processing Systems 7*, pages 657–664, 1995.
- [4] R. H. Crites and A. G. Barto. Improving elevator performance using reinforcement learning. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*, pages 1017–1023, Cambridge, MA, 1996. MIT Press.
- [5] F. Fernandez and M. Veloso. Probabilistic policy reuse in a reinforcement learning agent. In *Proc. of the 5th International Conference on Autonomous Agents and Multiagent Systems*, pages 720–727, 2006.
- [6] M. Genesereth and N. Love. General game playing: Overview of the aaii competition. *AI Magazine*, 26(2), 2005.
- [7] R. Maclin, J. Shavlik, L. Torrey, T. Walker, and E. Wild. Giving advice about preferred actions to reinforcement learners via knowledge-based kernel regression. In *Proceedings of the 20th National Conference on Artificial Intelligence*, 2005.
- [8] O. G. Selfridge, R. S. Sutton, and A. G. Barto. Training and tracking in robotics. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 670–672, 1985.
- [9] S. P. Singh. Transfer of learning by composing solutions of elemental sequential tasks. *Machine Learning*, 8:323–339, 1992.
- [10] V. Soni and S. Singh. Using homomorphisms to transfer options across continuous reinforcement learning domains. In *Proceedings of the Twenty First National Conference on Artificial Intelligence*, July 2006.
- [11] P. Stone, R. S. Sutton, and G. Kuhlmann. Reinforcement learning for RoboCup-soccer keepaway. *Adaptive Behavior*, 13(3):165–188, 2005.
- [12] R. S. Sutton and A. G. Barto. *Introduction to Reinforcement Learning*. MIT Press, 1998.
- [13] M. E. Taylor and P. Stone. Behavior transfer for value-function-based reinforcement learning. In F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. P. Singh, and M. Wooldridge, editors, *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 53–59, New York, NY, July 2005. ACM Press.
- [14] M. E. Taylor, P. Stone, and Y. Liu. Value functions for RL-based behavior transfer: A comparative study. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*, July 2005.
- [15] M. E. Taylor, S. Whiteson, and P. Stone. Transfer via inter-task mappings in policy search reinforcement learning. In *The Sixth International Joint Conference on Autonomous Agents and Multiagent Systems*, May 2007.
- [16] G. Tesauro. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 6(2):215–219, 1994.
- [17] S. Thrun. Is learning the  $n$ -th thing any easier than learning the first? In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems*, volume 8, pages 640–646. The MIT Press, 1996.
- [18] S. Whiteson and P. Stone. Evolutionary function approximation for reinforcement learning. *Journal of Machine Learning Research*, 7(May):877–917, 2006.

# An Incentive Mechanism for Eliciting Fair Ratings of Sellers in E-Marketplaces

Jie Zhang

David R. Cheriton School of Computer Science  
University of Waterloo  
Waterloo, ON, Canada N2L 3G1  
j44zhang@uwaterloo.ca

## 1. RESEARCH PROGRAM

Our research is within the subfield of modeling trust and reputation in multi-agent systems for electronic commerce. More specifically, we are interested in addressing two problems that may arise in trust and reputation models where buying agents elicit opinions about selling agents from other buyers (known as advisors) in the marketplace:

- Unfair ratings of sellers provided to buyers
- Developing incentives for buyers to report their ratings of sellers

To explain, the ratings provided by advisors are possibly unfair. Buyers may provide unfairly high ratings to promote the seller. This is referred to as “ballot stuffing” [1]. Buyers may also provide unfairly low ratings, in order to cooperate with other sellers to drive a seller out of the marketplace. This is referred to as “bad-mouthing”. Besides the problem of unfair ratings, rating submission is voluntary in most trust management systems. Buyers do not have direct incentives to provide ratings because, for example, providing reputation ratings of sellers requires some effort [3]. Providing fair ratings for a trustworthy seller may also decrease the chance of doing business with the seller because of competition from other buyers.

## 2. PROGRESS TO DATE

We first seek to develop a model that addresses unfair ratings. Our proposal is to adopt a personalized approach that allows a buyer to estimate the reputation (referred to as private reputation) of an advisor based on their ratings for commonly rated sellers. When the buyer has limited private knowledge of the advisor, the public reputation of the advisor will also be considered, based on all the ratings for the sellers ever rated by the advisor. Finally, the trustworthiness of the advisor will be modeled by combining the weighted private and public reputations, where the weights are determined based on the estimated reliability of the private reputation, using probabilistic reasoning.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 Doctoral Mentoring Program, May 13, 2007, Honolulu, Hawaii, USA.

Equipped with the richer method for modeling trustworthiness of advisors in terms of private and public reputation, we are then interested in embedding this reasoning into a framework where there is as well incentive for being honest. Other researchers have also been working on developing incentive reputation mechanisms to encourage honesty in the reporting from buyers, in order to diminish concerns about unfair ratings. Two types of mechanisms have been developed, side payment mechanisms [4] and credibility mechanisms [5]. Side payment mechanisms offer side payment to buyers that fairly rate results of business with sellers. In these mechanisms, providing fair ratings for business results is a Nash equilibrium. Credibility mechanisms measure agents' credibility. The credibility of a buyer and a seller in a business will be decreased if their ratings about the business result are different.

We, however, begin with a novel insight that advisors may be motivated to provide honest ratings when asked by other buyers if advisors that are honest are rewarded by sellers through more profitable transactions. This idea is supported by the work in the field of evolutionary game theory, such as the work of Gintis et al. [2]. They argue that an agent's altruism in one context signals “quality” of the agent that will benefit from increased opportunities in other wider contexts. We use our personalized approach to create a social network of buyers. Each buyer in the society retains a neighborhood of the most trustworthy buyers, as advisors. In our mechanism, we also allow sellers to explicitly model the reputability of buyers, based on the number of neighborhoods they belong to in the society. Buyers that provide fair ratings of sellers are likely to be neighbors of many other buyers and can be considered reputable. This is also supported by Gintis et al. [2] through the model of a multi-player game. They argue that agents reporting honestly provide benefit to others and will further be preferred by others as allies. These agents will be able to attract a larger audience to witness their feedback (also known as increasing “broadcast efficiency”). Sellers in our system will increase quality and decrease prices of products for more reputable buyers, in order to build their own reputation. This therefore creates an incentive for buyers to provide fair ratings of sellers.

To date, we have developed a specific personalized model for representing the trustworthiness of advisors and sellers. One main idea that we use is to model the ratings that arrive according to their time windows. This helps to avoid the situation where advisors may untruthfully rate selling agents a large number of times and deal with changes of agents' behavior. Similarly, the personalized approach allows a buyer

to model the private reputation of a seller based on its own ratings for the seller. If the buyer does not want to rely fully on its personal experience with the seller, it will ask for advisors' ratings of the seller. It then can derive a public reputation of the seller from ratings provided by them. The trustworthiness of the seller is modeled by combining the weighted private and public reputation values, using forgetting and discounting factors. We have carried out some experiments based on simulations to illustrate the effectiveness of our approach. For example, experimental results indicate that our approach can effectively model the trustworthiness of advisors even when buyers do not have much experience with sellers. Also, our approach is still effective when the majority of advisors provide large numbers of unfair ratings.

We have also begun the specification of the incentive mechanism. Consider the scenario in an electronic marketplace where a buyer  $B$  wants to buy a product  $p$ . We assume that the buying and selling process is operated as a procurement auction. The buyer  $B$  sends the request to a central server. The request contains information about a set of non-price features  $\{f_1, f_2, \dots, f_m\}$  of the product, as well as a set of weights  $\{w_1, w_2, \dots, w_m\}$  that correspond to how important each non-price feature is. The central server forwards the request to sellers in the marketplace. A seller  $S \in \bar{S}$  sets the price and values for the non-price features of  $p$ . To gain profit from each possible transaction, the seller may not include in its bid the true cost of producing  $p$  with certain non-price features. The best potential gain the seller can offer the buyer from the transaction is as follows:

$$V'(p) = \sum_{i=1}^m w_i D(f_i) - C(p) \quad (1)$$

where  $D()$  is a function to convert descriptive non-price feature values to numeric values and  $C(p)$  is the cost for  $S$  to produce  $p$ . We define the distribution function for  $V'(p)$  as  $F(V')$ . A symmetric Bayes-Nash equilibrium can be derived. The equilibrium bidding function can be derived as follows:

$$P^*(p) = C(p) + \frac{\int_{V_L - C_H}^{V'(p)} F(x) dx}{F(V')} - V_D(R) \quad (2)$$

where  $V_D(R)$  is the valuation of discount for the buyer  $B$  with reputation  $R(B)$ ,  $V_L$  is the lower bound of the value for the non-price features of  $p$ , and  $C_H$  ( $V_L \geq C_H$ ) is the higher bound of the cost for the seller to produce  $p$ .

Our mechanism allows the central server to maintain a fixed number of neighbors for each buyer from which the buyer can trust and ask advice about sellers. The central server models the trust value a buyer has of another buyer (an advisor) through the personalized approach. The seller  $S$  periodically acquires neighbor list information of buyers from the central server. It then counts for each buyer the number of neighborhoods. Suppose that there are  $N_B$  other buyers considering the buyer  $B$  as one of their neighbors. The reputation of  $B$  can be calculated as follows:

$$R(B) = \begin{cases} \frac{N_B}{\theta} & \text{if } N_B < \theta; \\ 1 & \text{otherwise.} \end{cases} \quad (3)$$

The value of  $\theta$  depends on the total number of buyers in the marketplace. As can be seen from Equations 2 and 3, buyers that are neighbors of many other buyers will be offered more discount by sellers. Our mechanism also allows sellers to

see how they have been rated by buyers, allowing sellers to reward those buyers deemed to be honest.

We have carried out preliminary experiments based on simulations to prove that both honest buyers and sellers are able to gain better profit in marketplaces operating with our mechanism.

### 3. FUTURE RESEARCH

Our research has two contributions, a personalized approach for buying agents to effectively model trustworthiness of other buyers and a novel incentive mechanism to elicit fair ratings of selling agents in electronic marketplaces. We are aware that many current social reputation models do not effectively allow for both public and private reputation modeling. For the future, we want to develop strategies for effectively comparing our model to other competing approaches. We may also learn more about how best to perform this modeling as we continue to make use of it for the problem of developing incentives for honesty in e-marketplaces.

For the incentive mechanism, one main direction for the future is to develop our mechanism in more detail. We will seek a more comprehensive approach for modeling buyers' reputation based on the social network topology. We are particularly interested in exploring how to demonstrate that our approach copes with collusion, whereas other incentive mechanisms do not (as noted by other researchers). Our mechanism allows sellers to view the ratings provided by buyers and can in this way detect dishonesty. It also allows buyers to maintain a list of trustworthy other buyers as their neighbors. If a buyer colludes, it can be excluded from neighborhoods and will not be rewarded by sellers. Sellers that collude will also not profit because buyers can make informed decisions about which sellers to do business with, based on advice from their neighbors. To prove the above expectations, we will develop experiments using agents that strategically collude. We seek to develop as well definitive comparisons with competing models. It may also be useful to determine how robust our model is to buyers and sellers leaving the marketplace.

### 4. REFERENCES

- [1] C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, Minneapolis, MN, October 2000.
- [2] H. Gintis, E. A. Smith, and S. Bowles. Costly signaling and cooperation. *Journal of Theoretical Biology*, 213:103–19, 2001.
- [3] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. (*to appear*). *Decision Support Systems*, 2005.
- [4] R. Jurca and B. Faltings. An incentive compatible reputation mechanism. In *Proceedings of the IEEE Conference on E-Commerce, Newport Beach, CA, USA*, 2003.
- [5] R. Jurca and B. Faltings. Confess: An incentive compatible reputation mechanism for the online hotel booking industry. In *Proceedings of the IEEE Conference on E-Commerce, San Diego, CA, USA*, 2004.

# Agent Models for Enhanced Interactions in Cooperative and Competitive Environments - Research outline

Inon Zuckerman  
Department of Computer Science  
Bar-Ilan University  
Ramat-Gan, 52900, Israel  
zukermi@cs.biu.ac.il

## 1. INTRODUCTION

Multi-agent Systems, as the name suggests, deals with environments in which there are several agents with potential to interact. The field of Multi-Agents Systems began its rapid development since the development of distributed, interconnected computer systems, such as the Internet. Such interconnected settings, when one computer agent interacts with another, required researching interactions such as coordination, cooperation and competition.

As stated in the title, my research is guided toward the development of agent models for enhanced interactions both in cooperative and competitive environment. Specifically, I will look into two problems, each in the other side of the cooperativeness/competitiveness spectrum.

In my first problem I will try to build a formal model for an agent who is situated in an adversarial setting. Over the years there has been extensive research to formalize different sorts of cooperative behavior which tried providing formal specification for the development of cooperative agents (Kraus and Grosz's *SharePlans* [2], and Cohen and Levesque's *Joint Intentions* [3] models being the most influential ones). However, not all social interactions are collaborative or cooperative, it might be the case when agents find themselves in a competitive scenario with other agents who competes with them on the same goal, and might even initiate explicit actions to prevent them from completing their goals.

In this part of my research plan I will build a formal model for different types of adversarial environments and mental attitudes and behavioral axioms which will be used as a design guideline for building adversarial agents.

The second problem will deal with agents coordination when communication is not possible or is undesirable (e.g. malfunctioning communication device, or a need to act undetected). Schelling [4] showed that in some situations, human players were able to achieve coordination with high percentage of success, even in cases when game theory predicted a very low successful coordination percentages. In such cases, some of the possible coordination solutions were

conceived as prominent solutions by the players, which led to their selection and coordination success. Such prominent solutions were coined *Focal Points* by Schelling.

I plan to explore the notion of *Focal Points* and try to model and use them to enhance interactions in various human-agent and agent-agent coordination tasks.

## 2. MODELING ADVERSARIAL ENVIRONMENTS

Since the emergence of the multiagent systems field of research an extensive research was made to formalize, using various logical languages, different types of teamwork and cooperative problem solving interactions. Those formalizations meant to deal with the mental states of agents, to reason about their actions in different scenarios, or generally speaking to act as a guidelines for designing such agents.

However, as human nature dictates, when interactions are guided by interests, participants may clash and form competitions and rivalries, which in turn might lead to disagreements, uncooperative interactions, or even proactive actions to damage your adversary. We will denote environments which contain agents that have conflicting interests as *Adversarial Environments* and the conflicting agents will be named *adversaries*.

Our aim in this chapter is to create a formal adversarial model which will describe the mental states of agents who are situated in such environments. The model will define the mental states of agents as well as key components and axioms that will serve as a basic reference for implementation of such agents. Using the model, we will be able to create agents that will behave better in adversarial situations, consequently achieving superior results than the typical agents in competitive scenarios. Our formalization will build upon the *SharedPlans* model, which provides an extensive set of definitions, predicates and modal logic operators for collaborative planning.

### 2.1 What has been done

We began by looking into various competitive scenarios and classified them according to key properties, which resulted in different class of adversarial environments: from the simplest case of interacting with one adversary, to the most complex environments containing not only multiple adversaries but also *friends* and *foes*.

In the next step (published in [6]) we took the simplest of our environments, a *Zero-sum* and *Simple* environment (the formal definitions of the terms are found in the paper), and formalized the mental attitudes and behavior of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

Copyright 2007 IFAAMAS .

a single agent who is situated in such environment. Our empirical evaluation was done by an analysis of completed *Connect-Four* games which were played by humans. The game was solved by [1], where he showed that the opening player (White) has an optimal strategy and can always win in principle. Therefore, we isolated and analyzed the games where the Black player won and found interesting findings which conform to our proposed behavioral axioms.

## 2.2 Future directions

The first step we are now taking is conducting a more robust experimental evaluation of our axioms as proposed in [6]. We have decided on two experimental domains, one for a bilateral adversarial interaction and the other for a multilateral interaction. We will conduct comparison evaluation of agents which were implemented with and without our proposed axiomatic guidelines.

The next step will be to go back to the drawing board and work on expanding the formalism to more complex adversarial environment, where the next step would be to work on non zero-sum interactions, which opens a new world of behaviors like alliances, negotiations and more.

## 3. FOCAL POINTS COORDINATION

Agents often need to coordinate their actions in a coherent manner. Sometimes, achieving coherent behavior is the result of explicit communication and negotiation. However, communication is not always possible, for reasons as varied as high communication costs, the need to avoid detection, damaged communication devices, or language incompatibility.

Schelling [4] called coordination-without-communication scenarios *tactic coordination games*, and named these games’ “prominent solutions” *focal points*. A classic example is the solution most people choose when asked to divide \$100 into two piles, of any sizes; they should attempt only to match the expected choice of some other, unseen player. Usually, people create two piles of \$50 each, and that is what Schelling dubbed a focal point<sup>1</sup>. Coordination can be achieved using focal point where both coordinating sides recognizes the prominence of a solution, and reasons that the coordinating partner will also be able to recognize the focality of that situation.

Our aim in this research is to explore the notion of *focal points* and see how it can be augmented into computerized agents to enhance both Human-Agent and Agent-Agent coordination.

### 3.1 What has been done

In our first focal points work [5] we looked into the scenarios where an agent needs to coordinate with a human partner, but he cannot know in advance who his partner is, which means that he cannot train for a specific predefined partner. Instead, we wanted to train an agent who will be successful in coordination tasks with a *general* human partner.

We employ learning algorithms (a decision learning tree and an artificial neural network) to help our agent discover coordination strategies. we designed three experimental domain (named *Pick the pile*, *shape matching* and *candidate*

<sup>1</sup>In Schelling’s experiment 36 out of 41 had split the money to piles of fifty-fifty.

*selection games*, see [5] for detailed explanation) and conducted large scale experiments<sup>2</sup> using human subjects to get their unique answers to randomized instances of those domains. Those instances were used to create an automated agent that performs well when faced with a new human partner in a newly generated environment. However, applying machine learning on raw domain data results in classifiers having poor performance (due to several problems: e.g. similar instances with different classifications). Instead, we propose the usage of a *Focal Point Learning* approach: we preprocess raw domain data, and place it into a new representation space, based on focal point properties. Given our domain’s raw data  $O_i$ , we apply a transformation  $T$ , such that  $N_j = T(O_i)$ , where  $i, j$  are the number of attributes before and after the transformation, respectively. Our *Focal Point Learning* approach resulted in agents which coordinate 40% to 80% better than the other learning methods with a general human partner

### 3.2 Future directions

There are three different directions which I plan taking from here: the first being a direct continuation of the first paper, which I will try to use focal points with a reinforcement learning algorithm, such that it will provide a better default coordination results before the agent’s algorithm will adjust to its specific human partner. We are looking for a real world application which such technique would be beneficial for real contribution.

The other possible future directions are of a more theoretic orientation, where I will try to find some relationship (which I hypothesis to exist) between the emergence of focal points and minimal model in knowledge base. Another direction will be to delve into the focal points properties and try to provide them with a more accurate description.

## 4. ACKNOWLEDGMENT

This work is jointly supervised by: Prof’ Sarit Kraus and Prof’ Jeffrey S. Rosenschein.

## 5. REFERENCES

- [1] V. Allis. A knowledge-based approach of Connect-Four — the game is solved: White wins.
- [2] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.
- [3] H. J. Levesque, P. R. Cohen, and J. H. T. Nunes. On acting together. In *Proc. of AAAI-90*, pages 94–99, Boston, MA, 1990.
- [4] T. Schelling. *The Strategy of Conflict*. Oxford University Press, New York, 1963.
- [5] I. Zuckerman, S. Kraus, and J. S. Rosenschein. Using focal point learning to improve tactic coordination in human-machine interactions. In *IJCAI 2007*, pages 1563–1568, Hyderabad, India, January 2007.
- [6] I. Zuckerman, S. Kraus, J. S. Rosenschein, and G. A. Kaminka. An adversarial environment model for bounded rational agents in zero-sum interactions. In *AAMAS 2007*, Honolulu, Hawaii, May 2007. To appear.

<sup>2</sup>in the *Pick the Pile* game we gathered data of over 3000 games instances, played by more than 250 different subjects from all over the world