# AAMAS 2011

# The 10th International Conference on Autonomous Agents and Multiagent Systems

May 2–6, 2011 • Taipei, Taiwan

Conference Abstracts

IFAAMAS

**IFAAMAS**
International Foundation for Autonomous Agents and Multiagent Systems
www.ifaamas.org



ISBN-10: 0-9826571-8-8
ISBN-13: 978-0-9826571-8-8

# Contents

# Main Program – Best Papers

## Wednesday

### Best Papers Session I

*Room: 101C*
*Chair: Peter McBurney*
13:30 – 15:10

**Agent-Based Control for Decentralised Demand Side Management in the Smart Grid**
*Sarvapali D. Ramchurn, Perukrishnen Vytelingum, Alex Rogers, Nicholas R. Jennings*
(Poster B 39)

*Best Innovative Application Paper Nominee*
Central to the vision of the smart grid is the deployment of smart meters that will allow autonomous software agents, representing the consumers, to optimise their use of devices and heating in the smart home while interacting with the grid. However, without some form of coordination, the population of agents may end up with overly-homogeneous optimised consumption patterns that may generate significant peaks in demand in the grid. These peaks, in turn, reduce the efficiency of the overall system, increase carbon emissions, and may even, in the worst case, cause blackouts. Hence, in this paper, we introduce a novel model of a Decentralised Demand Side Management (DDSM) mechanism that allows agents, by adapting the deferment of their loads based on grid prices, to coordinate in a decentralised manner. Specifically, using average UK consumption profiles for 26M homes, we demonstrate that, through an emergent coordination of the agents, the peak demand of domestic consumers in the grid can be reduced by up to 17% and carbon emissions by up to 6%. We also show that our DDSM mechanism is robust to the increasing electrification of heating in UK homes (i.e., it exhibits a similar efficiency).

**Deploying Power Grid-Integrated Electric Vehicles as a Multi-Agent System**
*Sachin Kamboj, Willett Kempton, Keith S. Decker*
(Poster G 38)

*Best Innovative Application Paper Nominee*
Grid-Integrated Vehicles (GIVs) are plug-in Electric Drive Vehicles (EDVs) with power-management and other controls that allow them to respond to external commands sent by power-grid operators, or their affiliates, when parked and plugged-in to the grid. At a bare minimum, such GIVs should respond to demand-management commands or pricing signals to delay, reduce or switch-off the rate of charging when the demand for electricity is high. In more advanced cases, these GIVs might sell both power and storage capacity back to the grid in any of the several electric power markets – a concept known as Vehicle-to-Grid power or V2G power. Although individual EDVs control too little power to sell in the market at an individual level, a large group of EDVs may form an aggregate or coalition that controls enough power to meaningfully sell, at a profit, in these markets. The profits made by such a coalition can then be used by the coalition members to offset the costs of the electric vehicles and batteries themselves. In this paper we describe an implemented and deployed multi-agent system that is used to integrate EDVs into the electricity grid managed by PJM, the largest transmission service operator in the world. We provide a brief introduction to GIVs and the various power markets and discuss why multi-agent systems are a good match for this application.

**Multi-Agent Monte Carlo Go**
*Leandro Soriano Marcolino, Hitoshi Matsubara*
(Poster B 40)

*Best Paper Nominee*
In this paper we propose a Multi-Agent version of UCT Monte Carlo Go. We use the emergent behavior of a great number of simple agents to increase the quality of the Monte Carlo simulations, increasing the strength of the artificial player as a whole. Instead of one agent playing against itself, different agents play in the simulation phase of the algorithm, leading to a better exploration of the search space. We could significantly overcome Fuego, a top Computer Go software. Emergent behavior seems to be the next step of Computer Go development.

**Towards a Unifying Characterization for Quantifying Weak Coupling in Dec-POMDPs**
*Stefan J. Witwicki, Edmund H. Durfee*
(Poster G 39)

*Best Student Paper Nominee*

Researchers in the field of multiagent sequential decision making have commonly used the terms "weakly-coupled" and "loosely-coupled" to qualitatively classify problems involving agents whose interactions are limited, and to identify various structural restrictions that yield computational advantages to decomposing agents' centralized planning and reasoning into largely-decentralized planning and reasoning. Together, these restrictions make up a heterogeneous collection of facets of "weakly-coupled" structure that are conceptually related, but whose purported computational benefits are hard to compare evenhandedly. The contribution of this paper is a unified characterization of weak coupling that brings together three complementary aspects of agent interaction structure. By considering these aspects in combination, we derive new bounds on the computational complexity of optimal DecPOMDP planning, that together quantify the relative benefits of exploiting different forms of interaction structure. Further, we demonstrate how our characterizations can be used to explain why existing classes of decoupled solution algorithms perform well on some problems but poorly on others, as well as to predict the performance of a particular algorithm from identifiable problem attributes.

**GUARDS - Game Theoretic Security Allocation on a National Scale**
*James Pita, Milind Tambe, Christopher Kiekintveld, Shane Cullen, Erin Steigerwald*
(Poster G 40)

*Best Innovative Application Paper Nominee*

Building on research previously reported at AAMAS conferences, this paper describes an innovative application of a novel gametheoretic approach for a national scale security deployment. Working with the United States Transportation Security Administration (TSA), we have developed a new application called GUARDS to assist in resource allocation tasks for airport protection at over 400 United States airports. In contrast with previous efforts such as ARMOR and IRIS, which focused on one-off tailored applications and one security activity (e.g. canine patrol or checkpoints) per application, GUARDS faces three key issues: (i) reasoning about hundreds of heterogeneous security activities; (ii) reasoning over diverse potential threats; (iii) developing a system designed for hundreds of end-users. Since a national deployment precludes tailoring to specific airports, our key ideas are: (i) creating a new game-theoretic framework that allows for heterogeneous defender activities and compact modeling of a large number of threats; (ii) developing an efficient solution technique based on general purpose Stackelberg game solvers; (iii) taking a partially centralized approach for knowledge acquisition and development of the system. In doing so we develop a software scheduling assistant, GUARDS, designed to reason over two agents – the TSA and a potential adversary – and allocate the TSA's limited resources across hundreds of security activities in order to provide protection within airports. The scheduling assistant has been delivered to the TSA and is currently under evaluation and testing for scheduling practices at an undisclosed airport. If successful, the TSA intends to incorporate the system into their unpredictable scheduling practices nationwide. In this paper we discuss the design choices and challenges encountered during the implementation of GUARDS. GUARDS represents promising potential for transitioning years of academic research into a nationally deployed system.

## Best Papers Session II

*Room: 101D*
*Chair: Liz Sonenberg*
13:30 – 15:10

**On the Outcomes of Multiparty Persuasion**
*Elise Bonzon, Nicolas Maudet*
(Poster B 41)

*Best Paper Nominee*

In recent years, several bilateral protocols regulating the exchange of arguments between agents have been proposed. When dealing with persuasion, the objective is to arbitrate among conflicting viewpoints. Often, these debates are not entirely predetermined from the initial situation, which means that agents have a chance to influence the outcome in a way that fits their individual preferences. This paper introduces a simple and intuitive protocol for multiparty argumentation, in which several (more than two) agents are equipped with argumentation systems. We further assume that they focus on a (unique) argument (or issue) –thus making the debate two-sided–

but do not coordinate. We study what outcomes can (or will) be reached if agents follow this protocol. We investigate in particular under which conditions the debate is pre-determined or not, and whether the outcome coincides with the result obtained by merging the argumentation systems.

**Arbitrators in Overlapping Coalition Formation Games**
*Yair Zick, Edith Elkind*
(Poster R 38)

*Best Student Paper Nominee*

Overlapping Coalition Formation (OCF) games are cooperative games where the players can simultaneously participate in several coalitions. Capturing the notion of stability in OCF games is a difficult task: a player may deviate by abandoning some, but not all of the coalitions he is involved in, and the crucial question is whether he then gets to keep his payoff from the unaffected coalitions. In related work the authors introduce three stability concepts for OCF games – the conservative, refined, and optimistic core – that are based on different answers to this question. In this paper, we propose a unified framework for the study of stability in the OCF setting, which encompasses the concepts considered previously as well as a wide variety of alternative stability concepts. Our approach is based on the notion of an arbitrator, which can be thought of as an external party that determines payoff to deviators. We give a complete characterization of outcomes that are stable under arbitration. In particular, our results provide a criterion for the outcome to be in the refined or optimistic core, thus complementing previously results for the conservative core, and answering questions left open previously. We also introduce a notion of the nucleolus for arbitrated OCF games, and argue that it is non-empty. Finally, we extend the definition of the Shapley value to the OCF setting, and provide an axiomatic characterization for it.

**Learning the Demand Curve in Posted-Price Digital Goods Auctions**
*Meenal Chhabra, Sanmay Das*
(Poster R 39)

*Best Student Paper Nominee*

Online digital goods auctions are settings where a seller with an unlimited supply of goods (e.g. music or movie downloads) interacts with a stream of potential buyers. In the posted price setting, the seller makes a take-it-or-leave-it offer to each arriving buyer. We study the seller's revenue maximization problem in posted-price auctions of digital goods. We find that algorithms from the multi-armed bandit literature like UCB, which come with good regret bounds, can be slow to converge. We propose and study two alternatives: (1) a scheme based on using Gittins indices with priors that make appropriate use of domain knowledge; (2) a new learning algorithm, LLVD, that assumes a linear demand curve, and maintains a Beta prior over the free parameter using a moment-matching approximation. LLVD is not only (approximately) optimal for linear demand, but also learns fast and performs well when the linearity assumption is violated, for example in the cases of two natural valuation distributions, exponential and log-normal.

**Ties Matter: Complexity of Voting Manipulation Revisited**
*Svetlana Obraztsova, Edith Elkind, Noam Hazon*
(Poster B 42)

*Best Paper Nominee*

In their groundbreaking paper, Bartholdi, Tovey and Trick argued that many well-known voting rules, such as Plurality, Borda, Copeland and Maximin are easy to manipulate. An important assumption made in that paper is that the manipulator's goal is to ensure that his preferred candidate is among the candidates with the maximum score, or, equivalently, that ties are broken in favor of the manipulator's preferred candidate. In this paper, we examine the role of this assumption in the easiness results of Bartholdi et al. We observe that the algorithm presented in Bartholdi et al extends to all rules that break ties according to a fixed ordering over the candidates. We then show that all scoring rules are easy to manipulate if the winner is selected from all tied candidates uniformly at random. This result extends to Maximin under an additional assumption on the manipulator's utility function that is inspired by the original model of Bartholdi et al. In contrast, we show that manipulation becomes hard when arbitrary polynomial-time tie-breaking rules are allowed, both for the rules considered in Bartholdi et al, and for a large class of scoring rules.

**Designing Incentives for Boolean Games**
*Ulle Endriss, Sarit Kraus, Jérôme Lang, Michael Wooldridge*
(Poster R 40)

*Best Paper Nominee*
Boolean games are a natural, compact, and expressive class of logic-based games, in which each player exercises unique control over some set of Boolean variables, and has some logical goal formula that it desires to be achieved. A player's strategy set is the set of all possible valuations that may be made to its variables. A player's goal formula may contain variables controlled by other agents, and in this case, it must reason strategically about how best to assign values to its variables. In the present paper, we consider the possibility of overlaying Boolean games with taxation schemes. A taxation scheme imposes a cost on every possible assignment an agent can make. By designing a taxation scheme appropriately, it is possible to perturb the preferences of the agents within a society, so that agents are rationally incentivised to choose some socially desirable equilibrium that would not otherwise be chosen, or incentivised to rule out some socially undesirable equilibria. After formally presenting the model, we explore some issues surrounding it (e.g., the complexity of finding a taxation scheme that implements some socially desirable outcome), and then discuss possible desirable properties of taxation schemes.

# Main Program – Full Papers

## Wednesday

### Session A1 – Robotics

*Room: 101A*
*Chair: Maria Gini*
*10:30 – 12:10*

**Who Goes There? Selecting a Robot to Reach a Goal Using Social Regret**
*Meytal Traub, Gal A. Kaminka, Noa Agmon*
(Poster R 41)

A common decision problem in multi-robot applications involves deciding on which robot, out of a group of $N$ robots, should travel to a goal location, to carry out a task there. Trivially, this decision problem can be solved greedily, by selecting the robot with the shortest expected travel time. However, this ignores the inherent uncertainty in path traversal times; we may prefer a robot that is slower (but always takes the same time), over a robot that is expected to reach the goal faster, but on occasion takes a very long time to arrive. We make several contributions that address this challenge. First, we bring to bear economic decision-making theory, to distinguish between different selection policies, based on risk (risk averse, risk seeking, etc.). Second, we introduce social regret (the difference between the actual travel time by the selected robot, and the hypothetical time of other robots) to augment decision-making in practice. Then, we carry out experiments in simulation and with real robots, to demonstrate the usefulness of the selection procedures under real-world settings, and find that travel-time distributions have repeating characteristics.

**Exploration Strategies Based on Multi-Criteria Decision Making for Search and Rescue Autonomous Robots**
*Nicola Basilico, Francesco Amigoni*
(Poster R 42)

Autonomous mobile robots are considered a valuable technology for search and rescue applications, where an initially unknown environment has to be explored to locate human victims. In this scenario, robots exploit exploration strategies to autonomously move around the environment. Most of the strategies proposed in literature are based on the idea of evaluating a number of candidate locations according to ad hoc utility functions that combine different criteria. In this paper, we show some of the advantages of using a more theoretically-grounded approach, based on Multi-Criteria Decision Making (MCDM), to define exploration strategies for robots employed in search and rescue applications. We implemented some MCDM-based exploration strategies within an existing robot controller and we experimentally evaluated their performance in a simulated environment.

**Simulation-based Temporal Projection of Everyday Robot Object Manipulation**
*Lars Kunze, Mihai Emanuel Dolha, Emitza Guzman, Michael Beetz*
(Poster G 41)

Performing everyday manipulation tasks successfully depends on the ability of autonomous robots to appropriately account for the physical behavior of task-related objects. Meaning that robots have to predict and consider the physical effects of their possible actions to take. In this work we investigate a simulation-based approach to naive physics temporal projection in the context of autonomous robot everyday manipulation. We identify the abstractions underlying typical first-order axiomatizations as the key obstacles for making valid naive physics predictions. We propose that temporal projection for naive physics problems should not be performed based on abstractions but rather based on detailed physical simulations. This idea is realized as a temporal projection system for autonomous manipulation robots that translates naive physics problems into parametrized physical simulation tasks, that logs the data structures and states traversed in simulation, and translates the logged data back into symbolic time-interval-based first-order representations. Within this paper, we describe the concept and implementation of the temporal projection system and present the example of an egg-cracking robot for demonstrating its feasibility.

**Online Anomaly Detection in Unmanned Vehicles**
*Eliahu Khalastchi, Gal A. Kaminka, Meir Kalech, Raz Lin*
(Poster B 43)

Autonomy requires robustness. The use of unmanned (autonomous) vehicles is appealing for tasks which are dangerous or dull. However, increased reliance on autonomous robots increases reliance on their robustness. Even with validated software, physical faults can cause the controlling software to perceive the environment incorrectly, and thus to make decisions that lead to task failure. We present an online anomaly detection method for robots, that is lightweight, and is able to take into account a large number of monitored sensors and internal measurements, with high precision. We demonstrate a specialization of the familiar Mahalanobis Distance for robot use, and also show how it can be used even with very large dimensions, by online selection of correlated measurements for its use. We empirically evaluate these contributions in different domains: commercial Unmanned Aerial Vehicles (UAVs), a vacuum-cleaning robot, and a high-fidelity flight simulator. We find that the online Mahalanobis distance technique, presented here, is superior to previous methods.

**Tree Adaptive A***
*Carlos Hernández, Xiaoxun Sun, Sven Koenig, Pedro Meseguer*
(Poster B 44)

Incremental heuristic search algorithms can solve sequences of similar search problems potentially faster than heuristic search algorithms that solve each search problem from scratch. So far, there existed incremental heuristic search algorithms (such as Adaptive A*) that make the h-values of the current A* search more informed, which can speed up future A* searches, and incremental heuristic search algorithms (such as D* Lite) that change the search tree of the current A* search to the search tree of the next A* search, which can be faster than constructing it from scratch. In this paper, we present Tree Adaptive A*, which applies to goal-directed navigation in unknown terrain and builds on Adaptive A* but combines both classes of incremental heuristic search algorithms in a novel way. We demonstrate experimentally that it can run faster than Adaptive A*, Path Adaptive A* and D* Lite, the top incremental heuristic search algorithms in the context of goal-directed navigation in unknown grids.

## Session B1 – Distributed Problem Solving I

*Room: 101B*
*Chair: Ed Durfee*
*10:30 – 12:10*

**Quality Guarantees for Region Optimal DCOP Algorithms**
*Meritxell Vinyals, Eric Shieh, Jesus Cerquides, Juan Antonio Rodriguez-Aguilar, Zhengyu Yin, Milind Tambe, Emma Bowring*
(Poster G 42)

$k$- and $t$-optimality algorithms provide solutions to DCOPs that are optimal in regions characterized by its size and distance respectively. Moreover, they provide quality guarantees on their solutions. Here we generalise the $k$- and $t$-optimal framework to introduce $\mathcal{C}$-optimality, a flexible framework that provides reward-independent quality guarantees for optima in regions characterised by any arbitrary criterion. Therefore, $\mathcal{C}$-optimality allows us to explore the space of criteria (beyond size and distance) looking for those that lead to better solution qualities. We benefit from this larger space of criteria to propose a new criterion, the socalled size-bounded-distance criterion, which outperforms $k$- and $t$-optimality.

**Distributed Algorithms for Solving the Multiagent Temporal Decoupling Problem**
*James C. Boerkoel, Edmund H. Durfee*
(Poster G 43)

Scheduling agents can use the Multiagent Simple Temporal Problem (MaSTP) formulation to efficiently find and represent the complete set of alternative consistent joint schedules in a distributed and privacy-maintaining manner. However, continually revising this set of consistent joint schedules as new constraints arise may not be a viable option in environments where communication is uncertain, costly, or otherwise problematic. As an alternative, agents can find and represent a temporal decoupling in terms of locally independent sets of consistent schedules that, when combined, form a set of consistent joint schedules. Unlike current algorithms for calculating a temporal decoupling that require centralization of the problem representation, in this paper we present a new, provably correct, distributed algorithm for calculating a temporal decoupling. We

prove that this algorithm has the same theoretical computational complexity as current state-of-the-art MaSTP solution algorithms, and empirically demonstrate that it is more efficient in practice. We also introduce and perform an empirical cost/benefit analysis of new techniques and heuristics for selecting a maximally flexible temporal decoupling.

**Decomposing Constraint Systems: Equivalences and Computational Properties**
*Wiebe van der Hoek, Cees Witteveen, Michael Wooldridge*
(Poster R 43)

Distributed systems can often be modeled as a collection of distributed (system) variables whose values are constrained by a set of constraints. In distributed multi-agent systems, the set of variables occurring at a site (subsystem) is usually viewed as controllable by a local agent. This agent assigns values to the variables, and the aim is to provide distributed methods enabling a set of agents to come up with a global assignment (solution) that satisfies all the constraints. Alternatively, the system might be understood as a distributed database. Here, the focus is on ensuring consistency of the global system if local constraints (the distributed parts of the database) change. In this setting, the aim is to determine whether the existence of a global solution can be guaranteed. In other settings (e.g., P2P systems, sensor networks), the values of the variables might be completely out of control of the individual systems, and the constraints only characterize globally normal states or behavior of the system. In order to detect anomalies, one specifies distributed methods that can efficiently indicate violations of such constraints. The aim of this paper is to show that the following three main problems identified in these research areas are in fact identical: (i) the problem of ensuring that independent agents come up with a global solution; (ii) the problem of ensuring that global consistency is maintained if local constraint stores change; and (iii) the problem of ensuring that global violations can be detected by local nodes. This claim is made precise by developing a decomposition framework for distributed constraint systems and then extracting preservation properties that must satisfied in order to solve the above mentioned problems. Although satisfying the preservation properties seems to require different decomposition modes, our results demonstrate that in fact these decomposition properties are equivalent, thereby showing that the three main problems identified above are identical. We then show that the complexity of finding such decompositions is polynomially related to finding solutions for the original constraint system, which explains the popularity of decomposition applied to tractable constraint systems. Finally, we address the problem of finding optimal decompositions and show that even for tractable constraint systems, this problem is hard.

**Decentralized Monitoring of Distributed Anytime Algorithms**
*Alan Carlin, Shlomo Zilberstein*
(Poster B 45)

Anytime algorithms allow a system to trade solution quality for computation time. In previous work, monitoring techniques have been developed to allow agents to stop the computation at the "right" time so as to optimize a given time-dependent utility function. However, these results apply only to the single-agent case. In this paper we analyze the problems that arise when several agents solve components of a larger problem, each using an anytime algorithm. Monitoring in this case is more challenging as each agent is uncertain about the progress made so far by the others. We develop a formal framework for decentralized monitoring, establish the complexity of several interesting variants of the problem, and propose solution techniques for each one. Finally, we show that the framework can be applied to decentralized flow and planning problems.

**Consensus Acceleration in Multiagent Systems with the Chebyshev Semi-Iterative Method**
*Renato L.G. Cavalcante, Alex Rogers, Nicholas R. Jennings*
(Poster B 46)

We consider the fundamental problem of reaching consensus in multiagent systems; an operation required in many applications such as, among others, vehicle formation and coordination, shape formation in modular robotics, distributed target tracking, and environmental modeling. To date, the consensus problem (the problem where agents have to agree on their reported values) has been typically solved with iterative decentralized algorithms based on graph Laplacians. However, the convergence of these existing consensus algorithms is often too slow for many important multiagent applications, and thus they are increasingly being combined with acceler-

ation methods. Unfortunately, state-of-the- art acceleration techniques require parameters that can be optimally selected only if complete information about the network topology is available, which is rarely the case in practice. We address this limitation by deriving two novel acceleration methods that can deliver good performance even if little information about the network is available. The first proposed algorithm is based on the Chebyshev semi-iterative method and is optimal in a well defined sense; it maximizes the worst-case convergence speed (in the mean sense) given that only rough bounds on the extremal eigenvalues of the network matrix are available. It can be applied to systems where agents use unreliable communication links, and its computational complexity is similar to those of simple Laplacian-based methods. This algorithm requires synchronization among agents, so we also propose an asynchronous version that approximates the output of the synchronous algorithm. Mathematical analysis and numerical simulations show that the convergence speed of the proposed acceleration methods decrease gracefully in scenarios where the sole use of Laplacian-based methods is known to be impractical.

## Session C1 – Game Theory I

*Room: 101C*
*Chair: Jeff Rosenschein*
*10:30 – 12:10*

### Information Elicitation for Decision Making

*Yiling Chen, Ian A. Kash*
(Poster R 44)

Proper scoring rules, particularly when used as the basis for a prediction market, are powerful tools for eliciting and aggregating beliefs about events such as the likely outcome of an election or sporting event. Such scoring rules incentivize a single agent to reveal her true beliefs about the event. Othman and Sandholm introduced the idea of a decision rule to examine these problems in contexts where the information being elicited is conditional on some decision alternatives. For example,"What is the probability having ten million viewers if we choose to air new television show X? What if we choose Y?" Since only one show can actually air in a slot, only the results under the chosen alternative can ever be observed. Othman and Sandholm developed proper scoring rules (and thus decision markets) for a single, deterministic decision rule: always select the action with the greatest probability of success. In this work we significantly generalize their results, developing scoring rules for other deterministic decision rules, randomized decision rules, and situations where there may be more than two outcomes (e.g. less than a million viewers, more than one but less than ten, or more than ten million).

### Stable Partitions in Additively Separable Hedonic Games

*Haris Aziz, Felix Brandt, Hans Georg Seedig*
(Poster R 45)

An important aspect in systems of multiple autonomous agents is the exploitation of synergies via coalition formation. In this paper, we solve various open problems concerning the computational complexity of stable partitions in additively separable hedonic games. First, we propose a polynomial-time algorithm to compute a contractually individually stable partition. This contrasts with previous results such as the NP-hardness of computing individually stable or Nash stable partitions. Secondly, we prove that checking whether the core or the strict core exists is NP-hard in the strong sense even if the preferences of the players are symmetric. Finally, it is shown that

verifying whether a partition consisting of the grand coalition is contractual strict core stable or Pareto optimal is coNP-complete.

**Complexity of Coalition Structure Generation**
*Haris Aziz, Bart de Keijzer*
(Poster R 46)

We revisit the coalition structure generation problem in which the goal is to partition the players into exhaustive and disjoint coalitions so as to maximize the social welfare. One of our key results is a general polynomial-time algorithm to solve the problem for all monotonic coalitional games provided that player types are known and the number of player types is bounded by a constant. As a corollary, we obtain a polynomial-time algorithm to compute an optimal partition for weighted voting games with a constant number of weight values and for coalitional skill games with a constant number of skills. We also consider well-studied and well-motivated coalitional games defined compactly on combinatorial domains. For these games, we characterize the complexity of computing an optimal coalition structure by presenting polynomial-time algorithms, approximation algorithms, or NP-hardness and inapproximability lower bounds.

**Equilibrium Approximation in Simulation-Based Extensive-Form Games**
*Nicola Gatti, Marcello Restelli*
(Poster B 47)

The class of simulation-based games, in which the payoffs are generated as an output of a simulation process, recently received a lot of attention in literature. In this paper, we extend such class to games in extensive form with continuous actions and perfect information. We design two convergent algorithms to find an approximate subgame perfect equilibrium (SPE) and an approximate Nash equilibrium (NE) respectively. Our algorithms can exploit different optimization techniques. In particular, we use: simulated annealing, cross entropy method, and Lipschitz optimization. We produce an extensive experimental evaluation of the performance of our algorithms in terms of approximation degree of the optimal solution and number of evaluated samples. Finding approximate NE and SPE requires exponential time in the game tree depth: an SPE can be computed in game trees with a small depth, while the computation of an NE is easier.

**Maximum Causal Entropy Correlated Equilibria for Markov Games**
*Brian D. Ziebart, J. Andrew Bagnell, Anind K. Dey*
(Poster B 48)

Motivated by a machine learning perspective–that game-theoretic equilibria constraints should serve as guidelines for predicting agents' strategies, we introduce maximum causal entropy correlated equilibria (MCECE), a novel solution concept for general-sum Markov games. In line with this perspective, a MCECE strategy profile is a uniquely-defined joint probability distribution over actions for each game state that minimizes the worst-case prediction of agents' actions under log-loss. Equivalently, it maximizes the worst-case growth rate for gambling on the sequences of agents' joint actions under uniform odds. We present a convex optimization technique for obtaining MCECE strategy profiles that resembles value iteration in finite-horizon games. We assess the predictive benefits of our approach by predicting the strategies generated by previously proposed correlated equilibria solution concepts, and compare against those previous approaches on that same prediction task.

## Session D1 – Multiagent Learning

*Room: 101D*
*Chair: Karl Tuyls*
*10:30 – 12:10*

**Learning Action Models for Multi-Agent Planning**
*Hankz Hankui Zhuo, Hector Muñoz-Avila, Qiang Yang*
(Poster G 44)

In multi-agent planning environments, action models for each agent must be given as input. However, creating such action models by hand is difficult and time-consuming, because it requires formally representing the complex relationships among different objects in the environment. The problem is compounded in multi-agent environments where agents can take more types of actions. In this paper, we present an algorithm to learn action models for multi-agent planning systems from a set of input plan traces. Our learning algorithm `Lammas` automatically generates three kinds of constraints: (1) constraints on the interactions between agents, (2) constraints on the correctness of the action models for each individual agent, and (3) constraints on actions themselves. `Lammas` at-

tempts to satisfy these constraints simultaneously using a weighted maximum satisfiability model known as MAX-SAT, and converts the solution into action models. We believe this to be one of the first learning algorithms to learn action models in the context of multi-agent planning environments. We empirically demonstrate that `Lammas` performs effectively and efficiently in several planning domains.

**Theoretical Considerations of Potential-Based Reward Shaping for Multi-Agent Systems**
*Sam Devlin, Daniel Kudenko*
(Poster G 45)

Potential-based reward shaping has previously been proven to both be equivalent to Q-table initialisation and guarantee policy invariance in single-agent reinforcement learning. The method has since been used in multi-agent reinforcement learning without consideration of whether the theoretical equivalence and guarantees hold. This paper extends the existing proofs to similar results in multi-agent systems, providing the theoretical background to explain the success of previous empirical studies. Specifically, it is proven that the equivalence to Q-table initialisation remains and the Nash Equilibria of the underlying stochastic game are not modified. Furthermore, we demonstrate empirically that potential-based reward shaping affects exploration and, consequentially, can alter the joint policy converged upon.

**Evolving Subjective Utilities: Prisoner's Dilemma Game Examples**
*Koichi Moriyama, Satoshi Kurihara, Masayuki Numao*
(Poster G 46)

We have proposed the utility-based Q-learning concept that supposes an agent internally has an emotional mechanism that derives subjective utilities from objective rewards and the agent uses the utilities as rewards of Q-learning. We have also proposed such an emotional mechanism that facilitates cooperative actions in Prisoner's Dilemma (PD) games. However, this mechanism has been designed and implemented manually in order to force the agents to take cooperative actions in PD games. Since it seems slightly unnatural, this work considers whether such an emotional mechanism exists and where it comes from. We try to evolve such mechanisms that facilitate cooperative actions in PD games by conducting simulation experiments with a genetic algorithm, and we investigate the evolved mechanisms from various points of view.

**Cooperation through Reciprocity in Multiagent Systems: An Evolutionary Analysis**
*Christian Hütter, Klemens Böhm*
(Poster B 49)

The Service Game is a model for reciprocity in multiagent systems. Here, agents interact repeatedly by requesting and providing services. In contrast to existing models where players are matched randomly, players of the Service Game may choose with whom they play. The rationale behind provider selection is to choose a provider that is likely to perform a task as desired. We develop a formal model for provider selection in the Service Game. An evolutionary process based on a genetic algorithm allows us to incorporate notions of bounded rationality, learning, and adaptation into the analysis of the game. We conduct a series of experiments to study the evolution of strategies and the emergence of cooperation. We show that cooperation is more expensive with provider selection than with random matching. Further, populations consisting of discriminators and defectors form a bistable community.

**Distributed Cooperation in Wireless Sensor Networks**
*Mihail Mihaylov, Yann-Aël Le Borgne, Karl Tuyls, Ann Nowé*
(Poster G 47)

We present a game-theoretic self-organizing approach for scheduling the radio activity of wireless sensor nodes. Our approach makes each node play a *win-stay lose-shift* (WSLS) strategy to choose when to schedule radio transmission, reception and sleeping periods. The proposed strategy relies only on local interactions with neighboring nodes, and is thus fully decentralized. This behavior results in shorter communication schedules, allowing to not only reduce energy consumption by reducing the wake-up cycles of sensor nodes, but also to decrease the data retrieval latency. We implement this WSLS approach in the OMNeT++ sensor network simulator where nodes are organized in three topologies —line, grid and random. We compare the performance of our approach to two state-of-the-art scheduling protocols, namely S-MAC and D-MAC, and show that the WSLS strategy brings significant gains in terms of energy savings, while at the same time reduces communication delays.

In addition, we show that our approach performs particularly well in large, random topologies.

## Session A2 – Logic-Based Approaches I

*Room: 101A*
*Chair: Wiebe van der Hoek*
*15:30 – 16:30*

### A Framework for Coalitional Normative Systems

*Jun Wu, Chongjun Wang, Junyuan Xie*
(Poster B 50)

We propose coalitional normative system (CNS), which can selectively restrict the joint behavior of a coalition, in this paper. We extend the semantics of atl and propose Coordinated ATL (CO-ATL) to support the formalizing of CNS. We soundly and completely characterize the limitation of the normative power of a coalition by identifying two fragments of COL-ATL language corresponding to two types of system properties that are unchangeable by restricting the joint behavior of such a coalition. Then, we prove that the effectiveness checking, feasibility and synthesis problems of CNS are PTIME-complete, CP-complete and FNP-complete, respectively. Moreover, we define two concepts of optimality for CNS, that is, minimality and compactness, and prove that both minimality checking and compactness checking are CONP-complete while the problem of checking whether a coalition is a minimal controllable coalition is DP-complete. The relation between NS and CNS is discussed, and it turns out that CSS intrinsically consists of a proper subset of CNSs and some basic problems related to cns are no more complex than that of NS.

### Practical Argumentation Semantics for Socially Efficient Defeasible Consequence

*Hiroyuki Kido, Katsumi Nitta*
(Poster G 48)

An abstract argumentation framework and the semantics, often called Dungean semantics, give a general framework for nonmonotonic logics. In the last fteen years, a great number of papers in computational argumentation adopt Dungean semantics as a fundamental principle for evaluating various kinds of defeasible consequences. Recently, many papers address problems not only with theoretical reasoning, i.e., reasoning about what to believe, but also practical reasoning, i.e., reasoning about what to do. This paper proposes a practical argumentation semantics specic to practical argumentation. This is motivated by our hypothesis that consequences of such argumentation should satisfy Pareto optimality because the consequences strongly depend on desires, aims, or values an individual agent or a group of agents has. We dene a practical argumentation framework and two kinds of extensions, preferred and grounded extensions, with respect to each group of agents. We show that evaluating Pareto optimality can be translated to evaluating preferred extensions of a particular practical argumentation framework. Furthermore, we show that our semantics is a natural extension of Dungean semantics in terms of considering more than one defeat relation. We give a generality order of four practical argumentation frameworks specied by taking into account Dungean semantics and Pareto optimality. We show that a member of preferred extensions of the most specic one is not just Pareto optimal, but also it is theoretically justied.

### Taming the Complexity of Linear Time BDI Logics

*Nils Bulling, Koen V. Hindriks*
(Poster G 49)

Reasoning about the mental states of agents is important in various settings, and has been recognized as vital for teamwork. But the complexity of some of the more well-known agent logics that facilitate reasoning about mental states prohibits the use of these logics in practice. An alternative is to investigate fragments of these logics that have a lower complexity but are still expressive enough for reasoning about the mental states of (other) agents. We explore this alternative and take as our starting point the linear time variant of BDI logic ($\text{BDI}_{\text{LTL}}$). We summarize some of the relevant known complexity results for e.g. LTL, $\text{KD45}_n$, and $\text{BDI}_{\text{LTL}}$ itself. We present a tableau-based method for establishing complexity bounds, and provide a map of the complexity of (various fragments of) $\text{BDI}_{\text{LTL}}$. Finally, we identify a few fragments that may be usefully applied for reasoning about mental states.

## Session B2 – Agent-Based System Development I

*Room: 101B*
*Chair: John Thangarajah*
*15:30 – 16:30*

**Scenarios for System Requirements Traceability and Testing**
*John Thangarajah, Gaya Jayatilleke, Lin Padgham*
(Poster B 51)

Scenarios in current design methodologies, provide a natural way for the users to identify the inputs and outputs of the system revolving around a particular interaction process. A scenario typically consists of a sequence of steps which captures a particular run of the system and satisfies some aspect of the requirements. In this work we add additional structure to the scenarios used in the Prometheus agent development methodology. This additional structure then facilitates both traceability and automated testing. We describe our process for mapping the scenarios and their steps to the initial detailed design, where we then maintain the traceability as the design develops. The structured action lists that we define for both scenarios and their variations provides the basis for facilitating automated testing of system behavior. We describe how we use the newly defined structure within the scenarios to facilitate testing, describing how we automate test case generation, execution and analysis.

**Kokomo: An Empirically Evaluated Methodology for Affective Applications**
*Derek J. Sollenberger, Munindar P. Singh*
(Poster G 50)

The introduction of affect or emotion modeling into software opens up new possibilities for improving user experience. Yet, current techniques for building affective applications are limited, with the treatment of affect in essence handcrafted in each application. The multiagent middleware Koko attempts to reduce the burden of incorporating affect modeling into applications. However, Koko can be effective only if the models it needs to function are suitably constructed. We propose Kokomo, a methodology that employs expressive communicative acts as an organizing principle for affective applications. Kokomo specifies the steps needed to create an affective application in Koko. A key motivation is that Kokomo would facilitate the construction of an affective application by engineers who may lack a prior background in affective modeling. We empirically evaluate Kokomo's utility through a developer study. The results are positive and demonstrate that the developers who used Kokomo were able to develop an affective application in less time, with fewer lines of code, and with a reduced perception of difficulty than developers who worked without Kokomo.

**Programming Mental State Abduction**
*Michal Sindlar, Mehdi Dastani, John-Jules Ch. Meyer*
(Poster G 51)

Many multi-agent system applications involve software agents that reason about the behavior of other agents with which they interact in cooperation or competition. In order to design and develop those systems, the employed programming languages should provide tools to facilitate the implementation of agents that can perform such reasoning. This paper focuses on BDI-based programming languages and proposes a nonmonotonic reasoning mechanism that can be incorporated into agents, allowing them to reason about observed behavior to infer others' beliefs or goals. In particular, it is suggested that the behavior-generating rules of agents are translated into a nonmonotonic logic programming framework. A formal analysis of the presented approach is provided and it is shown that it has desirable properties.

## Session C2 – Social Choice Theory

*Room: 101C*
*Chair: Vincent Connitzer*
*15:30 – 16:30*

**Possible And Necessary Winners In Voting Trees: Majority Graphs Vs. Profiles**
*Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, Toby Walsh*
(Poster G 52)

Given the preferences of several agents over a common set of candidates, voting trees can be used to select a candidate (the winner) by a sequence of pairwise competitions modelled by a binary tree (the agenda). The majority graph compactly represents the preferences of the agents and provides enough information to compute the winner. When some preferences are missing, there are various notions of winners, such as the possible winners (that is, winners in at least one completion) or the necessary winners (that is, winners

in all completions). In this generalized scenario, we show that using the majority graph to compute winners is not correct, since it may declare as winners candidates that are not so. Nonetheless, the majority graph can be used to compute efficiently an upper or lower approximation of the correct set of winners.

**Tight Bounds for Strategyproof Classification**
*Reshef Meir, Shaull Almagor, Assaf Michaely, Jeffrey S. Rosenschein*
(Poster R 47)

Strategyproof (SP) classification considers situations in which a decision-maker must classify a set of input points with binary labels, minimizing expected error. Labels of input points are reported by self-interested agents, who may lie so as to obtain a classifier more closely matching their own labels. These lies would create a bias in the data, and thus motivate the design of truthful mechanisms that discourage false reporting.

We here answer questions left open by previous research on strategyproof classification, in particular regarding the best approximation ratio (in terms of social welfare) that an SP mechanism can guarantee for $n$ agents. Our primary result is a lower bound of $3 - \frac{2}{n}$ on the approximation ratio of SP mechanisms under the shared inputs assumption; this shows that the previously known upper bound (for uniform weights) is tight. The proof relies on a result from Social Choice theory, showing that any SP mechanism must select a dictator at random, according to some fixed distribution. We then show how different randomizations can improve the best known mechanism when agents are weighted, matching the lower bound with a tight upper bound. These results contribute both to a better understanding of the limits of SP classification, as well as to the development of similar tools in other, related domains such as SP facility location.

**A Double Oracle Algorithm for Zero-Sum Security Games on Graphs**
*Manish Jain, Dmytro Korzhyk, Ondřej Vaněk, Vincent Conitzer, Michal Pěchouček, Milind Tambe*
(Poster G 53)

In response to the Mumbai attacks of 2008, the Mumbai police have started to schedule a limited number of inspection checkpoints on the road network throughout the city. Algorithms for similar security-related scheduling problems have been proposed in recent literature, but security scheduling in networked domains when targets have varying importance remains an open problem at large. In this paper, we cast the network security problem as an attacker-defender zero-sum game. The strategy spaces for both players are exponentially large, so this requires the development of novel, scalable techniques. We first show that existing algorithms for approximate solutions can be arbitrarily bad in general settings. We present RUGGED (Randomization in Urban Graphs by Generating strategies for Enemy and Defender), the first scalable optimal solution technique for such network security games. Our technique is based on a double oracle approach and thus does not require the enumeration of the entire strategy space for either of the players. It scales up to realistic problem sizes, as is shown by our evaluation of maps of southern Mumbai obtained from GIS data.

## Session D2 – Preferences and Strategies

*Room: 101D*
*Chair: Stephane Airiau*
*15:30 – 16:30*

**Modeling Social Preferences in Multi-player Games**
*Brandon Wilson, Inon Zuckerman, Dana Nau*
(Poster G 54)

Game-tree search algorithms have contributed greatly to the success of computerized players in two-player extensive-form games. In multi-player games there has been less success, partly because of the difficulty of recognizing and reasoning about the inter-player relationships that often develop and change during human game-play. Simplifying assumptions (e.g., assuming each player selfishly aims to maximize its own payoff) have not worked very well in practice. We describe a new algorithm for multi-player games, Socially-oriented Search (SOS), that incorporates ideas from Social Value Orientation theory from social psychology. We provide a theoretical study of the algorithm, and a method for recognizing and reasoning about relationships as they develop and change during a game. Our empirical evaluations of SOS in the strategic board game Quoridor show it to be significantly more effective against players with dynamic interrelationships than the current state-of-the-art algorithms.

**A Study of Computational and Human Strategies in Revelation Games**
*Noam Peled, Ya'akov (Kobi) Gal, Sarit Kraus*
(Poster R 48)

Revelation games are bilateral bargaining games in which agents may choose to truthfully reveal their private information before engaging in multiple rounds of negotiation. They are analogous to real-world situations in which people need to decide whether to disclose information such as medical records or university transcripts when negotiating over health plans and business transactions. This paper presents an agent-design that is able to negotiate proficiently with people in a revelation game with different dependencies that hold between players. The agent modeled the social factors that affect the players' revelation decisions on people's negotiation behavior. It was empirically shown to outperform people in empirical evaluations as well as agents playing equilibrium strategies. It was also more likely to reach agreement than people or equilibrium agents.

**Efficient Heuristic Approach to Dominance Testing in CP-nets**
*Minyi Li, Quoc Bao Vo, Ryszard Kowalczyk*
(Poster R 49)

CP-net (Conditional Preference Network) is one of the extensively studied languages for representing and reasoning with preferences. The fundamental operation of dominance testing in CP-nets, i.e. determining whether an outcome is preferred to another, is very important in many real-world applications. Current techniques for solving general dominance queries is to search for improving flipping sequence from one outcome to another as a proof of the dominance relation in all rankings satisfying the given CP-net. However, it is generally a hard problem even for binary-valued, acyclic CP-nets and tractable search algorithms exist only for specific problem classes. Hence, there is a need for efficient algorithms and techniques for dominance testing in more general problem settings. In this paper, we propose a heuristic approach, called DT*, to dominance testing in arbitrary acyclic multi-valued CP-nets. Our proposed approach guides the search process efficiently and allows significant reduction of search effort without impacting soundness or completeness of the search process. We present results of experiments that demonstrate the computational efficiency and feasibility of our approach to dominance testing.

## Session A3 – Distributed Problem Solving II

*Room: 101A*
*Chair: Matt Taylor*
*16:30 – 17:30*

**Resource-Aware Junction Trees for Efficient Multi-Agent Coordination**
*N. Stefanovitch, A. Farinelli, Alex Rogers, Nicholas R. Jennings*
(Poster R 50)

In this paper we address efficient decentralised coordination of cooperative multi-agent systems by taking into account the actual computation and communication capabilities of the agents. We consider coordination problems that can be framed as Distributed Constraint Optimisation Problems, and as such, are suitable to be deployed on large scale multi-agent systems such as sensor networks or multiple unmanned aerial vehicles. Specifically, we focus on techniques that exploit structural independence among agents' actions to provide optimal solutions to the coordination problem, and, in particular, we use the Generalized Distributive Law (GDL) algorithm. In this settings, we propose a novel resource aware heuristic to build junction trees and to schedule GDL computations across the agents. Our goal is to minimise the total running time of the coordination process, rather than the theoretical complexity of the computation, by explicitly considering the computation and communication capabilities of agents. We evaluate our proposed approach against DPOP, RDPI and a centralized solver on a number of benchmark coordination problems, and show that our approach is able to provide optimal solutions for DCOPs faster than previous approaches. Specifically, in the settings considered, when resources are scarce our approach is up to three times faster than DPOP (which proved to be the best among the competitors in our settings).

**Bounded Decentralised Coordination over Multiple Objectives**
*Francesco M. Delle Fave, Ruben Stranders, Alex Rogers, Nicholas R. Jennings*
(Poster B 52)

We propose the bounded multi-objective max-sum algorithm (B-MOMS), the first decentralised coordination algorithm for multi-objective optimisation problems. B-MOMS extends the max-sum message-passing algorithm for decentralised coordination to

compute bounded approximate solutions to multi-objective decentralised constraint optimisation problems (MO-DCOPs). Specifically, we prove the optimality of B-MOMS in acyclic constraint graphs, and derive problem dependent bounds on its approximation ratio when these graphs contain cycles. Furthermore, we empirically evaluate its performance on a multi-objective extension of the canonical graph colouring problem. In so doing, we demonstrate that, for the settings we consider, the approximation ratio never exceeds 2, and is typically less than 1.5 for less-constrained graphs. Moreover, the runtime required by B-MOMS on the problem instances we considered never exceeds 30 minutes, even for maximally constrained graphs with 100 agents. Thus, B-MOMS brings the problem of multi-objective optimisation well within the boundaries of the limited capabilities of embedded agents.

**Communication-Constrained DCOPs: Message Approximation in GDL with Function Filtering**
*Marc Pujol-Gonzalez, Jesus Cerquides, Pedro Meseguer, Juan Antonio Rodriguez-Aguilar*
(Poster B 53)

In this paper we focus on solving DCOPs in communication constrained scenarios. The GDL algorithm optimally solves DCOP problems, but requires the exchange of exponentially large messages which makes it impractical in such settings. Function filtering is a technique that alleviates this high communication requirement while maintaining optimality. Function filtering involves calculating approximations of the exact cost functions exchanged by GDL. In this work, we explore different ways to compute such approximations, providing a novel method that empirically achieves significant communication savings.

## Session B3 – Agent-Based System Development II

*Room: 101B*
*Chair: Mehdi Dastani*
*16:30 – 17:30*

**AgentScope: Multi-Agent Systems Development in Focus**
*Elth Ogston, Frances Brazier*
(Poster R 51)

Multi-agent systems form the basis of many innovative large-scale distributed applications. The development of such applications requires a careful balance of a wide range of concerns: a detailed understanding of the behaviour of the abstract algorithms being employed, a knowledge of the effects and costs of operating in a distributed environment, and an expertise in the performance requirements of the application itself. Experimental work plays a key role in the process of designing such systems. This paper examines the multi-agent systems development cycle from a distributed systems perspective. A survey of recent experimental studies finds that a large proportion of work on the design of multi-agent systems is focused on the analytical and simulation phases of development. This paper advocates an alternative more comprehensive development cycle, which extends from theoretical studies to simulations, emulations, demonstrators and finally staged deployment. AgentScope, a tool that supports the experimental stages of multi-agents systems development and facilitates long-term dispersed research efforts, is introduced. AgentScope consists of a small set of interfaces on which experimental work can be built independently of a particular type of platform. The aim is to make not only agent code but also experimental scenarios, and metrics reusable, both between projects and over simulation, emulation and demonstration platforms. An example gossip-based sampling experiment demonstrates reusability, showing the ease with which an experiment can be defined, modified into a comparison study, and ported between a simulator and an actual agent-operating system.

**Agent Programming with Priorities and Deadlines**
*Konstantin Vikhorev, Natasha Alechina, Brian Logan*
(Poster R 52)

We present AgentSpeak(RT), a real-time BDI agent programming language based on AgentSpeak(L).

AgentSpeak(RT) extends AgentSpeak intentions with deadlines which specify the time by which the agent should respond to an event, and priorities which specify the relative importance of responding to a particular event. The AgentSpeak(RT) interpreter commits to a priority-maximal set of intentions: a set of intentions which is maximally feasible while preferring higher priority intentions. We prove some properties of the language, such as guaranteed reactivity delay of the AgentSpeak(RT) interpreter and probabilistic guarantees of successful execution of intentions by their deadlines.

**Rich Goal Types in Agent Programming**
*Mehdi Dastani, M. Birna van Riemsdijk, Michael Winikoff*
(Poster B 54)

Goals are central to the design and implementation of intelligent software agents. Much of the literature on goals and reasoning about goals in agent programming frameworks only deals with a limited set of goal types, typically achievement goals, and sometimes maintenance goals. In this paper we extend a previously proposed unifying framework for goals with additional richer goal types that are explicitly represented as Linear Temporal Logic (LTL) formulae. We show that these goal types can be modelled as a combination of achieve and maintain goals. This is done by providing an operationalization of these new goal types, and showing that the operationalization generates computation traces that satisfy the temporal formula.

## Session C3 – Bounded Rationality

*Room: 101C*
*Chair: Edith Elkind*
*16:30 – 17:30*

**Expert-Mediated Search**
*Meenal Chhabra, Sanmay Das, David Sarne*
(Poster R 53)

Increasingly in both traditional, and especially Internet-based marketplaces, knowledge is becoming a traded commodity. This paper considers the impact of the presence of knowledge-brokers, or experts, on search-based markets with noisy signals. For example, consider a consumer looking for a used car on a large Internet marketplace. She sees noisy signals of the true value of any car she looks at the advertisement for, and can disambiguate this signal by paying for the services of an expert (for example, getting a Carfax report, or taking the car to a mechanic for an inspection). Both the consumer and the expert are rational, self-interested agents. We present a model for such search environments, and analyze several aspects of the model, making three main contributions: (1) We derive the consumer's optimal search strategy in environments with noisy signals, with and without the option of consulting an expert; (2) We find the optimal strategy for maximizing the expert's profit; (3) We study the option of market designers to subsidize search in a way that improves overall social welfare. We illustrate our results in the context of a plausible distribution of signals and values.

**Using Aspiration Adaptation Theory to Improve Learning**
*Avi Rosenfeld, Sarit Kraus*
(Poster B 55)

Creating agents that properly simulate and interact with people is critical for many applications. Towards creating these agents, models are needed that quickly and accurately predict how people behave in a variety of domains and problems. This paper explores how one bounded rationality theory, Aspiration Adaptation Theory (AAT), can be used to aid in this task. We extensively studied two types of problems - a relatively simple optimization problem and two complex negotiation problems. We compared the predictive capabilities of traditional learning methods with those where we added key elements of AAT and other optimal and bounded rationality models. Within the extensive empirical studies we conducted, we found that machine learning models combined with AAT were most effective in quickly and accurately predicting people's behavior.

**Less Is More: Restructuring Decisions to Improve Agent Search**
*David Sarne, Avshalom Elmalech, Barbara J. Grosz, Moti Geva*
(Poster G 55)

In many settings and for various reasons, people fail to make optimal decisions. These factors also influence the agents people design to act on their behalf in such virtual environments as eCommerce and distributed operating systems, so that the agents also act sub-optimally despite their greater computational capabilities. In some decision-making situations it is theoretically possible to supply the optimal strategy

to people or their agents, but this optimal strategy may be non-intuitive, and providing a convincing explanation of optimality may be complex. This paper explores an alternative approach to improving the performance of a decision-maker in such settings: the data on choices is manipulated to guide searchers to a strategy that is closer to optimal. This approach was tested for sequential search, which is a classical sequential decision-making problem with broad areas of applicability (e.g., product search, partnership search). The paper introduces three heuristics for manipulating choices, including one for settings in which repeated interaction or access to a decision-maker's past history is available. The heuristics were evaluated on a large population of computer agents, each of which embodies a search strategy programmed by a different person. Extensive tests on thousands of search settings demonstrate the promise of the problem-restructuring approach: despite a minor degradation in performance for a small portion of the population, the overall and average individual performance improve substantially. The heuristic that adapts based on a decision-maker's history achieved the best results.

## Session D3 – Virtual Agents I

*Room: 101D*
*Chair: Jonathan Gratch*
*16:30 – 17:30*

### Culture-related Differences in Aspects of Behavior for Virtual Characters Across Germany and Japan

*Birgit Endrass, Elisabeth André, Afia Akhter Lipi, Matthias Rehm, Yukiko Nakano*
(Poster B 56)

Integrating culture as a parameter into the behavioral models of virtual characters to simulate cultural differences is becoming more and more popular. But do these differences affect the user's perception? In the work described in this paper, we integrated aspects of non-verbal behavior as well as communication management behavior into the behavioral models of virtual characters for the two cultures of Germany and Japan in order to find out which of these aspects affect human observers of the target cultures. We give a literature review pointing out the expected differences in these two cultures and describe the analysis of a multi-modal corpus including video recordings of German and Japanese interlocutors. After integrating our findings into a demonstrator featuring a German and a Japanese scenario, we presented the virtual scenarios to human observers of the two target cultures in an evaluation study.

### Controlling Narrative Time in Interactive Storytelling

*Julie Porteous, Jonathan Teutenberg, Fred Charles, Marc Cavazza*
(Poster G 56)

Narrative time has an important role to play in Interactive Storytelling (IS). The prevailing approach to controlling narrative time has been to use implicit models that allow only limited temporal reasoning about virtual agent behaviour. In contrast, this paper proposes the use of an explicit model of narrative time which provides a control mechanism that enhances narrative generation, orchestration of virtual agents and number of possibilities for the staging of agent actions. This approach can help address a number of problems experienced in IS systems both at the level of execution staging and at the level of narrative generation. Consequently it has a number of advantages: it is more flexible with respect to the staging of virtual agent actions; it reduces the possibility of

timing problems in the coordination of virtual agents; and it enables more expressive representation of narrative worlds and narrative generative power. Overall it provides a uniform, consistent, principled and rigorous approach to the problem of time in agent-based storytelling. In the paper we demonstrate how this approach to controlling narrative time can be implemented within an IS system and illustrate this using our fully implemented IS system that features virtual agents inspired by Shakespeare's The Merchant of Venice. The paper presents results of an experimental evaluation with the system that demonstrates the use of this approach to co-ordinate the actions of virtual agents and to increase narrative generative power.

**ESCAPES - Evacuation Simulation with Children, Authorities, Parents, Emotions, and Social comparison**
*Jason Tsai, Natalie Fridman, Emma Bowring, Matthew Brown, Shira Epstein, Gal A. Kaminka, Stacy Marsella, Andrew Ogden, Inbal Rika, Ankur Sheel, Matthew E. Taylor, Xuezhi Wang, Avishay Zilka, Milind Tambe*
(Poster G 57)

In creating an evacuation simulation for training and planning, realistic agents that reproduce known phenomenon are required. Evacuation simulation in the airport domain requires additional features beyond most simulations, including the unique behaviors of first-time visitors who have incomplete knowledge of the area and families that do not necessarily adhere to often-assumed pedestrian behaviors. Evacuation simulations not customized for the airport domain do not incorporate the factors important to it, leading to inaccuracies when applied to it. In this paper, we describe ESCAPES, a multiagent evacuation simulation tool that incorporates four key features: (i) different agent types; (ii) emotional interactions; (iii) informational interactions; (iv) behavioral interactions. Our simulator reproduces phenomena observed in existing studies on evacuation scenarios and the features we incorporate substantially impact escape time. We use ESCAPES to model the International Terminal at Los Angeles International Airport (LAX) and receive high praise from security officials.

# Thursday

## Session A4 – Agent Communication

*Room: 101A*
*Chair: Michael Rovatsos*
*10:30 – 12:10*

**Commitments with Regulations: Reasoning about Safety and Control in** REGULA
*Elisa Marengo, Matteo Baldoni, Cristina Baroglio, Amit K. Chopra, Viviana Patti, Munindar P. Singh*
(Poster R 54)

Commitments provide a flexible means for specifying the business relationships among autonomous and heterogeneous agents, and lead to a natural way of enacting such relationships. However, current formalizations of commitments incorporate conditions expressed as propositions, but disregard (1) temporal regulations and (2) an agent's control over such regulations. Thus, they cannot handle realistic application scenarios where time and control are often central because of domain conventions or other requirements.
We propose a new formalization of commitments that builds on an existing representation of events in which we can naturally express temporal regulations as well as what an agent can control, including indirectly as based on the commitments and capabilities of other agents. Our formalization supports a notion of commitment safety. A benefit of our consolidated approach is that by incorporating these considerations into commitments we enable agents to reason about and flexibly enact the regulations.
The main contributions of this paper include (1) a formal semantics of commitments that accommodates temporal regulations; (2) a formal semantics of the notions of innate and social control; and (3) a formalization of when a temporal commitment is safe for its debtor. We evaluate our contributions using an extensive case study.

**Specifying and Applying Commitment-Based Business Patterns**
*Amit K. Chopra, Munindar P. Singh*
(Poster R 55)

Recent work in communications and business modeling emphasizes a commitment-based view of interaction. By abstracting away from implementation-level details, commitments can potentially enhance

perspicuity during modeling and flexibility during enactment. We address the problem of creating commitment-based specifications that directly capture business requirements, yet apply in distributed settings. We encode important business patterns in terms of commitments and group them into methods to better capture business requirements. Our approach yields significant advantages over existing approaches: our patterns (1) respect agent autonomy; (2) capture business intuitions faithfully; and (3) can be enacted in real-life, distributed settings. We evaluate our contributions using the Extended Contract Net Protocol.

**On the Verification of Social Commitments and Time**
*Mohamed El-Menshawy, Jamal Bentahar, Hongyang Qu, Rachida Dssouli*
(Poster G 58)

Social commitments have been widely studied to represent business contracts among agents with different competing objectives in communicating multi-agent systems. However, their formal verification is still an open issue. This paper proposes a novel model-checking algorithm to address this problem. We define a new temporal logic, CTLC, which extends CTL with modalities for social commitments and their fulfillment and violation. The verification technique is based on symbolic model checking that uses ordered binary decision diagrams to give a compact representation of the system. We also prove that the problem of model checking CTLC is polynomial-time reducible to the problem of model checking CTLK, the combination of CTL with modalities for knowledge. We finally present the full implementation of the proposed algorithm by extending the MCMAS symbolic model checker and report on the experimental results obtained when verifying the NetBill protocol.

**Information-Driven Interaction-Oriented Programming: BSPL, the Blindingly Simple Protocol Language**
*Munindar P. Singh*
(Poster B 57)

We present a novel approach to interaction-oriented programming based on declaratively representing communication protocols. Our approach exhibits the following distinguishing features. First, it treats a protocol as an engineering abstraction in its own right. Second, it models a protocol in terms of the information that the protocol needs to proceed (so agents enact it properly) and the information the protocol would produce (when it is enacted). Third, it naturally maps traditional operational constraints to the information needs of protocols, thereby obtaining the desired interactions without additional effort or reasoning. Fourth, our approach naturally supports *shared nothing* enactments: everything of relevance is included in the communications and no separate global state need be maintained. Fifth, our approach accommodates, but does not require, formal representations of the meanings of the protocols. We evaluate this approach via examples from the literature.

**On Topic Selection Strategies in Multi-Agent Naming Game**
*Wojciech Lorkiewicz, Ryszard Kowalczyk, Radoslaw Katarzyniak, Quoc Bao Vo*
(Poster B 58)

Communication is a key capability of autonomous agents in a multiagent system to exchange information about their environment. It requires a naming convention that typically involves a set of predefined names for all objects in the environment, which the agents share and understand. However, when the agents are heterogeneous, highly distributed, and situated in an unknown environment, it is very unrealistic to assume that all the objects can be foreseen in advance, and therefore their names cannot be defined beforehand. In such a case, each individual agent needs to be able to introduce new names for the objects it encounters and align them with the naming convention used by the other agents. A language game is a prospective mechanism for the agents to learn and align the naming conventions between them. In this paper we extend the language game model by proposing novel strategies for selecting topics, i.e. attracting agent's attention to different objects during the learning process. Using a simulated multi-agent system we evaluate the process of name alignment in the case of the least restrictive type of language game, the naming game without feedback. Utilising proposed strategies we study the dynamic character of formation of coherent naming conventions and compare it with the behaviour of commonly used random selection strategy. The experimental results demonstrate that the new strategies improve the overall convergence of the alignment process, limit agent's overall demand on memory, and scale with the increasing number of the interacting agents.

## Session B4 – Game Theory and Learning

*Room: 101B*
*Chair: Alex Rogers*
*10:30 – 12:10*

**Reaching Correlated Equilibria Through Multi-agent Learning**
*Ludek Cigler, Boi Faltings*
(Poster B 59)

Many games have undesirable Nash equilibria. For example consider a resource allocation game in which two players compete for an exclusive access to a single resource. It has three Nash equilibria. The two pure-strategy NE are efficient, but not fair. The one mixed-strategy NE is fair, but not efficient. Aumann's notion of correlated equilibrium fixes this problem: It assumes a correlation device which suggests each agent an action to take. However, such a "smart" coordination device might not be available. We propose using a randomly chosen, "stupid" integer coordination signal. "Smart" agents learn which action they should use for each value of the coordination signal. We present a multi-agent learning algorithm which converges in polynomial number of steps to a correlated equilibrium of a wireless channel allocation game, a variant of the resource allocation game. We show that the agents learn to play for each coordination signal value a randomly chosen pure-strategy Nash equilibrium of the game. Therefore, the outcome is an efficient correlated equilibrium. This CE becomes more fair as the number of the available coordination signal values increases. We believe that a similar approach can be used to reach efficient and fair correlated equilibria in a wider set of games, such as potential games.

**Sequential Targeted Optimality as a New Criterion for Teaching and Following in Repeated Games**
*Max Knobbout, Gerard A.W. Vreeswijk*
(Poster B 60)

In infinitely repeated games, the act of teaching an outcome to our adversaries can be beneficial to reach coordination, as well as allowing us to 'steer' adversaries to outcomes that are more beneficial to us. Teaching works well against followers, agents that are willing to go along with the proposal, but can lead to miscoordination otherwise. In the context of infinitely repeated games there is, as of yet, no clear formalism that tries to capture and combine these behaviours into a unified view in order to reach a solution of a game. In this paper, we propose such a formalism in the form of an algorithmic criterion, which uses the concept of targeted learning. As we will argue, this criterion can be a beneficial criterion to adopt in order to reach coordination. Afterwards we propose an algorithm that adheres to our criterion that is able to teach pure strategy Nash Equilibria to a broad class of opponents in a broad class of games and is able to follow otherwise, as well as able to perform well in self-play.

**On the Quality and Complexity of Pareto Equilibria in the Job Scheduling Game**
*Leah Epstein, Elena Kleiman*
(Poster B 61)

In the well-known scheduling game, a set of jobs controlled by selfish players wishes each to minimize the load of the machine on which it is executed, while the social goal is to minimize the makespan, that is, the maximum load of any machine. We consider this problem on the three most common machines models, identical machines, uniformly related machines and unrelated machines, with respect to both weak and strict Pareto optimal Nash equilibria. These are kinds of equilibria which are stable not only in the sense that no player can improve its cost by changing its strategy unilaterally, but in addition, there is no alternative choice of strategies for the entire set of players where no player increases its cost, and at least one player reduces its cost (in the case of strict Pareto optimality), or where all players reduce their costs (in the case of weak Pareto optimality). We give a complete classification of the social quality of such solutions with respect to an optimal solution, that is, we find the Price of Anarchy of such schedules as a function of the number of machines, $m$. In addition, we give a full classification of the recognition complexity of such schedules.

**Game Theory-Based Opponent Modeling in Large Imperfect-Information Games**
*Sam Ganzfried, Tuomas Sandholm*
(Poster R 56)

We develop an algorithm for opponent modeling in large extensive-form games of imperfect information. It works by observing the opponent's action frequencies and building an opponent model by combining information from a precomputed equilibrium strat-

egy with the observations. It then computes and plays a best response to this opponent model; the opponent model and best response are both updated continually in real time. The approach combines game-theoretic reasoning and pure opponent modeling, yielding a hybrid that can effectively exploit opponents after only a small number of interactions. Unlike prior opponent modeling approaches, ours is fundamentally game theoretic and takes advantage of recent algorithms for automated abstraction and equilibrium computation rather than relying on domain-specific prior distributions, historical data, or a hand-crafted set of features. Experiments show that our algorithm leads to significantly higher win rates (than an approximate-equilibrium strategy) against several opponents in limit Texas Hold'em – the most studied imperfect-information game in computer science – including competitors from recent AAAI computer poker competitions.

**False-name Bidding in First-price Combinatorial Auctions with Incomplete Information**
*Atsushi Iwasaki, Atsushi Katsuragi, Makoto Yokoo*
(Poster G 59)

False-name bids are bids submitted by a single agent under multiple fictitious names such as multiple e-mail addresses. False-name bidding can be a serious fraud in Internet auctions since identifying each participant is virtually impossible. It is shown that even the theoretically well-founded Vickrey-Clarke-Groves auction (VCG) is vulnerable to falsename bidding. Thus, several auction mechanisms that cannot be manipulated by false-name bids, i.e., false-name-proof mechanisms, have been developed. This paper investigates a slightly different question, i.e., how do they affect (perfect) Bayesian Nash equilibria of first-price combinatorial auctions? The importance of this question is that first-price combinatorial auctions are by far widely used in practice than VCG, and can be used as a benchmark for evaluating alternate mechanisms. In an environment where false-name bidding are possible, analytically investigating bidders' behaviors is very complicated, since nobody knows the number of real bidders. As a first step, we consider a kind of minimal settings where falsename bids become effective, i.e., an auction with two goods where one naive bidder competes with one shill bidder who may pretend to be two distinct bidders. We model this auction as a simple dynamic game and examine approximate Bayesian Nash equilibria by utilizing a numerical technique. Our analysis revealed that false-name bidding significantly affects the first-price auctions. Furthermore, the shill bidder has a clear advantage against the naive bidder.

## Session C4 – Teamwork

*Room: 101C*
*Chair: Juan A. Rodriguez*
*10:30 – 12:10*

**Metastrategies in the Colored Trails Game**
*Steven de Jong, Daniel Hennes, Karl Tuyls, Ya'akov (Kobi) Gal*
(Poster R 57)

This paper presents a novel method to describe and analyze strategic interactions in settings that include multiple actors, many possible actions and relationships among goals, tasks and resources. It shows how to reduce these large interactions to a set of bilateral normal-form games in which the strategy space is significantly smaller than the original setting, while still preserving many of its strategic characteristics. We demonstrate this technique on the Colored Trails (CT) framework, which encompasses a broad family of games defining multi-agent interactions and has been used in many past studies. We define a set of representative heuristics in a three-player CT setting. Choosing players' strategies from this set, the original CT setting is analytically decomposed into canonical bilateral social dilemmas, i.e., Prisoners' Dilemma, Stag Hunt and Ultimatum games. We present a set of criteria for generating strategically interesting CT games and empirically show that they indeed decompose into bilateral social dilemmas if players play according to the heuristics. Our results have significance for multi-agent systems researchers in mapping large multi-player task settings to well-known bilateral normal-form games in a way that facilitates the analysis of the original setting.

**Computing Stable Outcomes in Hedonic Games with Voting-Based Deviations**
*Martin Gairing, Rahul Savani*
(Poster R 58)

We study the computational complexity of finding stable outcomes in hedonic games, which are a class of coalition formation games. We restrict our attention to a nontrivial subclass of such games, which are guaranteed to possess stable outcomes, i.e., the set of symmetric additively-separable hedonic games. These

games are specified by an undirected edge-weighted graph: nodes are players, an outcome of the game is a partition of the nodes into coalitions, and the utility of a node is the sum of incident edge weights in the same coalition. We consider several stability requirements defined in the literature. These are based on restricting feasible player deviations, for example, by giving existing coalition members veto power. We extend these restrictions by considering more general forms of preference aggregation for coalition members. In particular, we consider voting schemes to decide if coalition members will allow a player to enter or leave their coalition. For all of the stability requirements we consider, the existence of a stable outcome is guaranteed by a potential function argument, and local improvements will converge to a stable outcome. We provide an almost complete characterization of these games in terms of the tractability of computing such stable outcomes. Our findings comprise positive results in the form of polynomialtime algorithms, and negative (PLS-completeness) results. The negative results extend to more general hedonic games.

**Empirical Evaluation of Ad Hoc Teamwork in the Pursuit Domain**
*Samuel Barrett, Peter Stone, Sarit Kraus*
(Poster B 62)

The concept of creating autonomous agents capable of exhibiting ad hoc teamwork was recently introduced as a challenge to the AI, and specifically to the multiagent systems community. An agent capable of ad hoc teamwork is one that can effectively cooperate with multiple potential teammates on a set of collaborative tasks. Previous research has investigated theoretically optimal ad hoc teamwork strategies in restrictive settings. This paper presents the first empirical study of ad hoc teamwork in a more open, complex teamwork domain. Specifically, we evaluate a range of effective algorithms for on-line behavior generation on the part of a single ad hoc team agent that must collaborate with a range of possible teammates in the pursuit domain.

**Decision Theoretic Behavior Composition**
*Nitin Yadav, Sebastian Sardina*
(Poster B 63)

The behavior composition problem involves realizing a virtual target behavior (i.e., the desired module) by suitably coordinating the execution of a set of partially controllable available components (e.g., agents, devices, processes, etc.) running in a shared partially predictable environment. All existing approaches to such problem have been framed within *strict* uncertainty settings. In this work, we propose a framework for automatic behavior composition which allows the seamless integration of classical behavior composition with decision-theoretic reasoning. Specifically, we consider the problem of *maximizing the "expected realizability"* of the target behavior in settings where the uncertainty can be quantified. Unlike previous proposals, the approach developed here is able to (better) deal with instances that do not accept "exact" solutions, thus yielding a more practical account for real domains. Moreover, it is provably strictly more general than the classical composition framework. Besides formally defining the problem and what counts as a solution, we show how a decision-theoretic composition problem can be solved by reducing it to the problem of finding an optimal policy in a Markov decision process.

**Solving Election Manipulation Using Integer Partitioning Problems**
*Andrew Lin*
(Poster G 60)
An interesting problem of multi-agent systems is that of voting, in which the preferences of autonomous agents are to be combined. Applications of voting include modeling social structures, search engine ranking, and choosing a leader among computational agents. In the setting of voting, it is very important that each agent presents truthful information about his or her preferences, and not manipulate. The choice of election system may encourage or discourage voters from manipulating. Because manipulation often results in undesirable consequences, making the determination of such intractable is an important goal. An interesting metric on the robustness of an election system concerns the frequency in which opportunities of manipulations occur in a given election system. Previous work by Walsh has evaluated the frequency of manipulation in the context of very specific election systems, particularly veto, when the number of candidates is limited to at most three, by

showing that manipulation problems in these systems can be directly viewed as problems of (Two-Way) Partition, and then using the best known heuristics of Partition. Walsh also claimed similar results hold for k-candidate veto election by way of problems involving multi-way partitions. We show that the results for k-candidate veto elections do not follow directly from common versions of partition problems and require non-trivial modifications to Multi-Way Partition. With these modifications, we confirm Walsh's claim that these elections are also vulnerable to manipulation. Our new computational problems also allow one to evaluate manipulation in the general case of k-candidate scoring protocols. We investigate the complexity of manipulating scoring protocols using new algorithms we derive by extending the known algorithms of Multi-Way Partition. It is our conclusion that the problems of manipulation in more general scoring protocols of four or more candidates are not vulnerable to manipulation using extensions of the current known algorithms of Multi-Way Partition. This may be due to weaknesses in these algorithms or complexity in manipulating general scoring protocols.

## Session A5 – Learning Agents

*Room: 101A*
*Chair: Paul Scerri*
*13:30 – 15:10*

### Using Iterated Reasoning to Predict Opponent Strategies

*Michael Wunder, John Robert Yaros, Michael Littman, Michael Kaisers*
(Poster G 61)

The field of multiagent decision making is extending its tools from classical game theory by embracing reinforcement learning, statistical analysis, and opponent modeling. For example, behavioral economists conclude from experimental results that people act according to levels of reasoning that form a "cognitive hierarchy" of strategies, rather than merely following the hyper-rational Nash equilibrium solution concept. This paper expands this model of the iterative reasoning process by widening the notion of a level within the hierarchy from one single strategy to a distribution over strategies, leading to a more general framework of multiagent decision making. It provides a measure of sophistication for strategies and can serve as a guide for designing good strategies for multiagent games, drawing it's main strength from predicting opponent strategies. We apply these lessons to the recently introduced Lemonade-stand Game, a simple setting that includes both collaborative and competitive elements, where an agent's score is critically dependent on its responsiveness to opponent behavior. The opening moves are significant to the end result and simple heuristics have achieved faster cooperation than intricate learning schemes. Using results from the past two real-world tournaments, we show how the submitted entries fit naturally into our model and explain why the top agents were successful.

### Cognitive Policy Learner: Biasing Winning or Losing Strategies

*Dominik Dahlem, Jim Dowling, William Harrison*
(Poster R 59)

In continuous learning settings stochastic stable policies are often necessary to ensure that agents continuously adapt to dynamic environments. The choice of the decentralised learning system and the employed policy plays an important role in the optimisation task. For example, a policy that exhibits fluctuations may also introduce non-linear effects which other agents in the environment may not be able to

cope with and even amplify these effects. In dynamic and unpredictable multiagent environments these oscillations may introduce instabilities. In this paper, we take inspiration from the limbic system to introduce an extension to the weighted policy learner, where agents evaluate rewards as either positive or negative feedback, depending on how they deviate from average expected rewards. Agents have positive and negative biases, where a bias either magnifies or depresses a positive or negative feedback signal. To contain the non-linear effects of biased rewards, we incorporate a decaying memory of past positive and negative feedback signals to provide a smoother gradient update on the probability simplex, spreading out the effect of the feedback signal over time. By splitting the feedback signal, more leverage on the win or learn fast (WoLF) principle is possible. The cognitive policy learner is evaluated using a small queueing network and compared with the fair action and weighted policy learner. Emphasis is placed on analysing the dynamics of the learning algorithms with respect to the stability of the queueing network and the overall queueing performance.

**Agent-Mediated Multi-Step Optimization for Resource Allocation in Distributed Sensor Networks**
*Bo An, Victor Lesser, David Westbrook, Michael Zink*
(Poster R 60)

Distributed collaborative adaptive sensing (DCAS) of the atmosphere is a new paradigm for detecting and predicting hazardous weather using a large dense network of short-range, low-powered radars to sense the lowest few kilometers of the earths atmosphere. In DCAS, radars are controlled by a collection of Meteorological Command and Control (MC&C) agents that instruct where to scan based on emerging weather conditions. Within this context, this work concentrates on designing efficient approaches for allocating sensing resources to cope with restricted real-time requirements and limited computational resources. We have developed a new approach based on explicit goals that can span multiple system heartbeats. This allows us to reason ahead about sensor allocations based on expected requirements of goals as they project forward in time. Each goal explicitly specifies end-users' preferences as well as a prediction of how a phenomena will move. We use a genetic algorithm to generate scanning strategies of each single MC&C and a distributed negotiation model to coordinate multiple MC&Cs' scanning strategies over multiple heartbeats. Simulation results show that as compared to simpler variants of our approach, the proposed distributed model achieved the highest social welfare. Our approach also has exhibited similarly very good performance in an operational radar testbed that is deployed in Oklahoma to observe severe weather events.

**Integrating Reinforcement Learning with Human Demonstrations of Varying Ability**
*Matthew E. Taylor, Halit Bener Suay, Sonia Chernova*
(Poster R 61)

This work introduces Human-Agent Transfer (HAT), an algorithm that combines transfer learning, learning from demonstration and reinforcement learning to achieve rapid learning and high performance in complex domains. Using experiments in a simulated robot soccer domain, we show that human demonstrations transferred into a baseline policy for an agent and refined using reinforcement learning significantly improve both learning time and policy performance. Our evaluation compares three algorithmic approaches to incorporating demonstration rule summaries into transfer learning, and studies the impact of demonstration quality and quantity, as well as the effect of combining demonstrations from multiple teachers. Our results show that all three transfer methods lead to statistically significant improvement in performance over learning without demonstration. The best performance was achieved by combining the best demonstrations from two teachers.

## Session B5 – Auction and Incentive Design

*Room: 101B*
*Chair: Rajiv Maheswaran*
*13:30 – 15:10*

**Incentive Design for Adaptive Agents**
*Yiling Chen, Jerry Kung, David C. Parkes, Ariel D. Procaccia, Haoqi Zhang*
(Poster R 62)

We consider a setting in which a principal seeks to induce an adaptive agent to select a target action by providing incentives on one or more actions. The agent maintains a belief about the value for each action—which may update based on experience—and selects at each time step the action with the maximal sum of value and associated incentive. The principal observes the agent's selection, but has no information about the agent's current beliefs or belief update process. For inducing the target action as soon as possible, or as often as possible over a fixed time period, it is optimal for a principal with a per-period budget to assign the budget to the target action and wait for the agent to want to make that choice. But with an across-period budget, no algorithm can provide good performance on all instances without knowledge of the agent's update process, except in the particular case in which the goal is to induce the agent to select the target action once. We demonstrate ways to overcome this strong negative result with knowledge about the agent's beliefs, by providing a tractable algorithm for solving the offline problem when the principal has perfect knowledge, and an analytical solution for an instance of the problem in which partial knowledge is available.

**A Truth Serum for Sharing Rewards**
*Arthur Carvalho, Kate Larson*
(Poster R 63)

We study a problem where a group of agents has to decide how a joint reward should be shared among them. We focus on settings where the share that each agent receives depends on the subjective opinions of its peers concerning that agent's contribution to the group. To this end, we introduce a mechanism to elicit and aggregate subjective opinions as well as for determining agents' shares. The intuition behind the proposed mechanism is that each agent who believes that the others are telling the truth has its expected share maximized to the extent that it is well-evaluated by its peers and that it is truthfully reporting its opinions. Under the assumptions that agents are Bayesian decision-makers and that the underlying population is sufficiently large, we show that our mechanism is incentive-compatible, budget-balanced, and tractable. We also present strategies to make this mechanism individually rational and fair.

**Capability-Aligned Matching: Improving Quality of Games with a Purpose**
*Che-Liang Chiou, Jane Yung-Jen Hsu*
(Poster B 64)

So far computer cannot satisfyingly solve many tasks that are extremely easy for human, such as image recognition or common sense reasoning. A partial solution is to delegate algorithmically difficult computation task to human, called human computation. The Game with a Purpose (GWAP), in which computational task is transformed into a game, is perhaps the most popular form of human computation. A simplified adverse selection model for output-agreement / simultaneous-verification GWAP was built, using the ESP Game as example. The experiment results favored an adverse selection model over an moral hazard model. We were particularly interested in output quality of a GWAP affected by how players are matched with each other, and proposed capability-aligned matching (CAM) versus commonly-used random matching. The analysis showed that when compared with random mathcing, the CAM improved output quality. The experiment confirmed conclusions drawed from the analysis, and further pointed out that task-human matching scheme was as important as human-human matching scheme studied in this paper. The main contribution of this paper is the analysis and empirical evaluation of human-human matching scheme, showing that capability-aligned matching can improve quality of GWAP.

**False-name-proof Mechanism Design without Money**
*Taiki Todo, Atsushi Iwasaki, Makoto Yokoo*
(Poster G 62)

Mechanism design studies how to design mechanisms that result in good outcomes even when agents strategically report their preferences. In traditional settings, it is assumed that a mechanism can enforce payments to give an incentive for agents to act honestly. However, in many Internet application domains, introducing monetary transfers is impossible or undesir-

able. Also, in such highly anonymous settings as the Internet, declaring preferences dishonestly is not the only way to manipulate the mechanism. Often, it is possible for an agent to pretend to be multiple agents and submit multiple reports under different identifiers, e.g., by creating different e-mail addresses. The effect of such false-name manipulations can be more serious in a mechanism without monetary transfers, since submitting multiple reports would have no risk. In this paper, we present a case study in false-name-proof mechanism design without money. In our basic setting, agents are located on a real line, and the mechanism must select the location of a public facility; the cost of an agent is its distance to the facility. This setting is called the facility location problem and can represent various situations where an agent's preference is single-peaked. First, we fully characterize the deterministic false-name-proof facility location mechanisms in this basic setting. By utilizing this characterization, we show the tight bounds of the approximation ratios for two objective functions: social cost and maximum cost. We then extend the results in two natural directions: a domain where a mechanism can be randomized and a domain where agents are located in a tree. Furthermore, we clarify the connections between false-name-proofness and other related properties.

**Majority-Rule-Based Preference Aggregation on Multi-Attribute Domains with CP-Nets**
*Minyi Li, Quoc Bao Vo, Ryszard Kowalczyk*
(Poster R 64)

This paper studies the problem of majority-rule-based collective decision-making where the agents' preferences are represented by CP-nets (Conditional Preference Networks). As there are exponentially many alternatives, it is impractical to reason about the individual full rankings over the alternative space and apply majority rule directly. Most existing works either do not consider computational requirements, or depend on a strong assumption that the agents have acyclic CP-nets that are compatible with a common order on the variables. To this end, this paper proposes an efficient SAT-based approach, called MajCP (Majority-rule-based collective decision-making with CP-nets), to compute the majority winning alternatives. Our proposed approach only requires that each agent submit a CP-net; the CP-net can be cyclic, and it does not need to be any common structures among the agents' CP-nets. The experimental results presented in this paper demonstrate that the proposed approach is computationally efficient. It offers several orders of magnitude improvement in performance over a Brute-force algorithm for large numbers of variables.

## Session C5 – Simulation and Emergence

*Room: 101C*
*Chair: Elisabeth Sklaar*
*13:30 – 15:10*

**Emerging Cooperation on Complex Networks**
*Norman Salazar, Juan Antonio Rodriguez-Aguilar, Josep Lluís Arcos, Ana Peleteiro, Juan C. Burguillo-Rial*
(Poster B 65)

The dynamic formation of coalitions is a well-known area of interest in multi-agent systems (MAS). Coalitions can help self-interested agents to successfully cooperate and coordinate in a mutually beneficial manner. Moreover, the organization provided by coalitions is particularly helpful for largescale MAS. In this paper we present a distributed approach for coalition emergence in large-scale MAS. In particular, we focus on MAS with agents interacting over complex networks since they provide a realistic model of the nowadays interconnected world (e.g. social networks). Our experiments show the effectiveness of our coalition emergence approach in achieving full cooperation over different complex networks. Furthermore, they provide a clear picture of the strong influence the topology has on coalition emergence.

**An Investigation of the Vulnerabilities of Scale Invariant Dynamics in Large Teams**
*Robin Glinton, Paul Scerri, Katia Sycara*
(Poster G 63)

Large heterogeneous teams in a variety of applications must make joint decisions using large volumes of noisy and uncertain data. Often not all team members have access to a sensor, relying instead on information shared by peers to make decisions. These sensors can become permanently corrupted through hardware failure or as a result of the actions of a malicious adversary. Previous work showed that when the trust between agents was tuned to a specific value the resulting dynamics of the system had a property called scale invariance which led to agents reaching highly accurate conclusion with little communication. In this paper we show that these dynamics also leave the

system vulnerable to most agents coming to incorrect conclusions as a result of small amounts of anomalous information maliciously injected in the system. We conduct an analysis that shows that the efficiency of scale invariant dynamics is due to the fact that large number of agents can come to correct conclusions when the difference between the percentage of agents holding conflicting opinions is relatively small. Although this allows the system to come to correct conclusions quickly, it also means that it would be easy for an attacker with specific knowledge to tip the balance. We explore different methods for selecting which agents are Byzantine and when attacks are launched informed by the analysis. Our study reveals global system properties that can be used to predict when and where in the network the system is most vulnerable to attack. We use the results of this study to design an algorithm used by agents to effectively attack the network, informed by local estimates of the global properties revealed by our investigation.

**The Evolution of Cooperation in Self-Interested Agent Societies: A Critical Study**
*Lisa-Maria Hofmann, Nilanjan Chakraborty, Katia Sycara*
(Poster B 66)

We study the phenomenon of evolution of cooperation in a society of self-interested agents using repeated games in graphs. A repeated game in a graph is a multiple round game, where, in each round, an agent gains payoff by playing a game with its neighbors and updates its action (state) by using the actions and/or payoffs of its neighbors. The interaction model between the agents is a two-player, two-action (cooperate and defect) Prisoner's Dilemma (PD) game (a prototypical model for interaction between self-interested agents). The conventional wisdom is that the presence of network structure enhances cooperation and current models use multiagent simulation to show evolution of cooperation. However, these results are based on particular combination of interaction game, network model and state update rules (e.g., PD game on a grid with imitate your best neighbor rule leads to evolution of cooperation). The state-of-the-art lacks a comprehensive picture of the dependence of the emergence of cooperation on model parameters like network topology, interaction game, state update rules and initial fraction of cooperators. We perform a thorough study of the phenomenon of evolution of cooperation using (a) a set of popular categories of networks, namely, grid, random networks, scale-free networks, and small-world networks and (b) a set of cognitively motivated update rules. Our simulation results show that the evolution of cooperation in networked systems is quite nuanced and depends on the combination of network type, update rules and the initial fraction of cooperating agents. We also provide an analysis to support our simulation results.

**A Model of Norm Emergence and Innovation in Language Change**
*Samarth Swarup, Andrea Apolloni, Zsuzsanna Fagyal*
(Poster G 64)

We analyze and extend a recently proposed model of linguistic diffusion in social networks, to analytically derive time to convergence, and to account for the innovation phase of lexical dynamics in networks. Our new model, the degree-biased voter model with innovation, shows that the probability of existence of a norm is inversely related to innovation probability. When the innovation rate in the population is low, variants that become norms are due to a peripheral member with high probability. As the innovation rate increases, the fraction of time that the norm is a peripheral-introduced variant and the total time for which a norm exists at all in the population decrease. These results align with historical observations of rapid increase and generalization of slang words, technical terms, and new common expressions at times of cultural change in some languages.

**Dynamic Level of Detail for Large Scale Agent-Based Urban Simulations**
*Laurent Navarro, Fabien Flacher, Vincent Corruble*
(Poster B 67)

Large scale agent-based simulations typically face a trade-off between the level of detail in the representation of each agent and the scalability seen as the number of agents that can be simulated with the computing resources available. In this paper, we aim at bypassing this trade-off by considering that the level of detail is itself a parameter that can be adapted automatically and dynamically during the simulation, taking into account elements such as user focus, or specific events. We introduce a framework for such a methodology, and detail its deployment within an existing simulator dedicated to the simulation of urban infrastructures. We evaluate the approach experimentally along two criteria: (1) the impact of our methodology on the resources (CPU use), and

(2) an estimate of the dissimilarity between the two modes of simulation, i.e. with and without applying our methodology. Initial experiments show that a major gain in CPU time can be obtained for a very limited loss of consistency.

## Session D5 – Logic-Based Approaches II

*Room: 101D*
*Chair: Munindar Singh*
*13:30 – 15:10*

### Reasoning About Local Properties in Modal Logic

*Hans van Ditmarsch, Wiebe van der Hoek, Barteld Kooi*
(Poster R 65)

In modal logic, when adding a syntactic property to an axiomatisation, this property will semantically become true in all models, in all situations, under all circumstances. For instance, adding a property like $K_a p \rightarrow K_b p$ (agent $b$ knows at least what agent $a$ knows) to an axiomatisation of some epistemic logic has as an effect that such a property becomes globally true, i.e., it will hold in all states, at all time points (in a temporal setting), after every action (in a dynamic setting) and after any communication (in an update setting), and every agent will know that it holds, it will even be common knowledge. We propose a way to express that a property like the above only needs to hold locally: it may hold in the actual state, but not in all states, and not all agents may know that it holds. We can achieve this by adding relational atoms to the language that represent (implicitly) quantification over all formulas, as in $\forall p(K_a p \rightarrow K_b p)$. We show how this can be done for a rich class of modal logics and a variety of syntactic properties.

### Knowledge and Control

*Wiebe van der Hoek, Nicolas Troquard, Michael Wooldridge*
(Poster R 66)

Logics of propositional control, such as van der Hoek and Wooldridge's CL-PC, were introduced in order to represent and reason about scenarios in which each agent within a system is able to exercise unique control over some set of system variables. Our aim in the present paper is to extend the study of logics of propositional control to settings in which these agents have incomplete information about the society they occupy. We consider two possible sources of incomplete information. First, we consider the possibility that an agent is only able to "read" a subset of the overall system variables, and so in any given system state, will have partial information about the state of the system. Second, we consider the possibility that an agent has incomplete information about which agent controls which variables. For both cases, we introduce a logic combining epistemic modalities with the operators of CL-PC, investigate its axiomatization, and discuss its properties.

### Strategic Games and Truly Playable Effectivity Functions

*Valentin Goranko, Wojciech Jamroga, Paolo Turrini*
(Poster R 67)

A well known (and often used) result by Marc Pauly states that for every playable effectivity function E there exists a strategic game that assigns to coalitions exactly the same power as E, and vice versa. While the latter direction of the correspondence is correct, we show that the former does not always hold in the case of infinite game models. We point out where the proof of correspondence goes wrong, and we present examples of playable effectivity functions in infinite models for which no equivalent strategic game exists. Then, we characterize the class of truly playable effectivity functions, that does correspond to strategic games. Moreover, we discuss a construction that transforms any playable effectivity function into a truly playable one while preserving the power of most (but not all) coalitions. We also show that Coalition Logic is not expressive enough to distinguish between playable and truly playable effectivity functions, and we extend it to a logic that can make this distinction while enjoying finite axiomatization and finite model property.

### Scientia Potentia Est

*Thomas Ågotnes, Wiebe van der Hoek, Michael Wooldridge*
(Poster R 68)

In epistemic logic, Kripke structures are used to model the distribution of information in a multi-agent system. In this paper, we present an approach to quantifying how much information each particular agent in a system has, or how important the agent is, with respect to some fact represented as a goal formula. It is typically the case that the goal formula is distributed knowledge in the system, but that no individual agent alone knows it. It might

be that several different groups of agents can get to know the goal formula together by combining their individual knowledge. By using power indices developed in voting theory, such as the Banzhaf index, we get a measure of how important an agent is in such groups. We analyse the properties of this notion of information-based power in detail, and characterise the corresponding class of voting games. Although we mainly focus on distributed knowledge, we also look at variants of this analysis using other notions of group knowledge. An advantage of our framework is that power indices and other power properties can be expressed in standard epistemic logic. This allows, e.g., standard model checkers to be used to quantitatively analyse the distribution of information in a given Kripke structure.

**Tractable Model Checking for Fragments of Higher-Order Coalition Logic**
*Patrick Doherty, Barbara Dunin-Kęplicz, Andrzej Szałas*
(Poster B 68)

A number of popular logical formalisms for representing and reasoning about the abilities of teams or coalitions of agents have been proposed beginning with the Coalition Logic (CL) of Pauly. Agotnes et al. introduced a means of succinctly expressing quantification over coalitions without compromising the computational complexity of model checking in CL by introducing Quantified Coalition Logic (QCL). QCL introduces a separate logical language for characterizing coalitions in the modal operators used in QCL. Boella et al., increased the representational expressibility of such formalisms by introducing Higher-Order Coalition Logic (HCL), a monadic second-order logic with special set grouping operators. Tractable fragments of HCL suitable for efficient model checking have yet to be identified. In this paper, we relax the monadic restriction used in HCL and restrict ourselves to the diamond operator. We show how formulas using the diamond operator are logically equivalent to second-order formulas. This permits us to isolate and define well-behaved expressive fragments of second-order logic amenable to model-checking in PTIME. To do this, we appeal to techniques used in deductive databases and quantifier elimination. In addition, we take advantage of the monotonicity of the effectivity function resulting in exponentially more succinct representation of models. The net result is identification of highly expressible fragments of a generalized HCL where model checking can be done efficiently in PTIME.

# Friday

## Session A6 – Robotics and Learning

*Room: 101A*
*Chair: Gal Kaminka*
*10:30 – 12:10*

**Active Markov Information-Theoretic Path Planning for Robotic Environmental Sensing**
*Kian Hsiang Low, John M. Dolan, Pradeep Khosla*
(Poster R 69)

Recent research in multi-robot exploration and mapping has focused on sampling environmental fields, which are typically modeled using the Gaussian process (GP). Existing information-theoretic exploration strategies for learning GP-based environmental field maps adopt the non-Markovian problem structure and consequently scale poorly with the length of history of observations. Hence, it becomes computationally impractical to use these strategies for in situ, real-time active sampling. To ease this computational burden, this paper presents a Markov-based approach to efficient information-theoretic path planning for active sampling of GP-based fields. We analyze the time complexity of solving the Markov-based path planning problem, and demonstrate analytically that it scales better than that of deriving the non-Markovian strategies with increasing length of planning horizon. For a class of exploration tasks called the transect sampling task, we provide theoretical guarantees on the active sampling performance of our Markov-based policy, from which ideal environmental field conditions and sampling task settings can be established to limit its performance degradation due to violation of the Markov assumption. Empirical evaluation on real-world temperature and plankton density field data shows that our Markov-based policy can generally achieve active sampling performance comparable to that of the widely-used non-Markovian greedy policies under less favorable realistic field conditions and task settings while enjoying significant computational gain over them.

**Horde: A Scalable Real-time Architecture for Learning Knowledge from Unsupervised Sensorimotor Interaction**
*Richard S. Sutton, Joseph Modayil, Michael Delp, Thomas Degris, Patrick M. Pilarski, Adam White, Doina Precup*
(Poster R 70)

Maintaining accurate world knowledge in a complex and changing environment is a perennial problem for robots and other artificial intelligence systems. Our architecture for addressing this problem, called Horde, consists of a large number of independent reinforcement learning sub-agents, or demons. Each demon is responsible for answering a single predictive or goal-oriented question about the world, thereby contributing in a factored, modular way to the system's overall knowledge. The questions are in the form of a value function, but each demon has its own policy, reward function, termination function, and terminal-reward function unrelated to those of the base problem. Learning proceeds in parallel by all demons simultaneously so as to extract the maximal training information from whatever actions are taken by the system as a whole. Gradient-based temporal-difference learning methods are used to learn efficiently and reliably with function approximation in this off-policy setting. Horde runs in constant time and memory per time step, and is thus suitable for learning online in real-time applications such as robotics. We present results using Horde on a multi-sensored mobile robot to successfully learn goal-oriented behaviors and long-term predictions from off-policy experience. Horde is a significant incremental step towards a real-time architecture for efficient learning of general knowledge from unsupervised sensorimotor interaction.

**On Optimizing Interdependent Skills: A Case Study in Simulated 3D Humanoid Robot Soccer**
*Daniel Urieli, Patrick MacAlpine, Shivaram Kalyanakrishnan, Yinon Bentor, Peter Stone*
(Poster B 69)

In several realistic domains an agent's behavior is composed of multiple interdependent skills. For example, consider a humanoid robot that must play soccer, as is the focus of this paper. In order to succeed, it is clear that the robot needs to walk quickly, turn sharply, and kick the ball far. However, these individual skills are ineffective if the robot falls down when switching from walking to turning, or if it cannot position itself behind the ball for a kick. This paper presents a learning architecture for a humanoid robot soccer agent that has been fully deployed and tested within the RoboCup 3D simulation environment. First, we demonstrate that individual skills such as walking and turning can be parameterized and optimized to match the best performance statistics reported in the literature. These results are achieved through effective use of the CMA-ES optimization algorithm. Next, we describe a framework for optimizing skills in conjunction with one another, a little-understood problem with substantial practical significance. Over several phases of learning, a total of roughly 100-150 parameters are optimized. Detailed experiments show that an agent thus optimized performs comparably with the top teams from the RoboCup 2010 competitions, while taking relatively few man-hours for development.

**Metric Learning for Reinforcement Learning Agents**
*Matthew E. Taylor, Brian Kulis, Fei Sha*
(Poster R 71)

A key component of any reinforcement learning algorithm is the underlying representation used by the agent. While reinforcement learning (RL) agents have typically relied on hand-coded state representations, there has been a growing interest in learning this representation. While inputs to an agent are typically fixed (i.e., state variables represent sensors on a robot), it is desirable to automatically determine the optimal relative scaling of such inputs, as well as to diminish the impact of irrelevant features. This work introduces HOLLER, a novel distance metric learning algorithm, and combines it with an existing instance-based RL algorithm to achieve precisely these goals. The algorithms' success is highlighted via empirical measurements on a set of six tasks within the mountain car domain.

## Session B6 – Energy Applications

*Room: 101B*
*Chair: Janusz Marecki*
*10:30 – 12:10*

### Cooperatives of Distributed Energy Resources for Efficient Virtual Power Plants

*Georgios Chalkiadakis, Valentin Robu, Ramachandra Kota, Alex Rogers, Nicholas R. Jennings*
(Poster R 72)

The creation of Virtual Power Plants (VPPs) has been suggested in recent years as the means for achieving the cost-efficient integration of the many distributed energy resources (DERs) that are starting to emerge in the electricity network. In this work, we contribute to the development of VPPs by offering a game-theoretic perspective to the problem. Specifically, we design cooperatives (or "cooperative VPPs"—CVPPs) of rational autonomous DER-agents representing small-to-medium size renewable electricity producers, which coalesce to profitably sell their energy to the electricity grid. By so doing, we help to counter the fact that individual DERs are often excluded from the wholesale energy market due to their perceived inefficiency and unreliability. We discuss the issues surrounding the emergence of such cooperatives, and propose a pricing mechanism with certain desirable properties. Specifically, our mechanism guarantees that CVPPs have the incentive to truthfully report to the grid accurate estimates of their electricity production, and that larger rather than smaller CVPPs form; this promotes CVPP efficiency and reliability. In addition, we propose a scheme to allocate payments within the cooperative, and show that, given this scheme and the pricing mechanism, the allocation is in the core and, as such, no subset of members has a financial incentive to break away from the CVPP. Moreover, we develop an analytical tool for quantifying the uncertainty about DER production estimates, and distinguishing among different types of errors regarding such estimates. We then utilize this tool to devise protocols to manage CVPP membership. Finally, we demonstrate these ideas through a simulation that uses real-world data.

### How Agents Can Help Curbing Fuel Combustion – a Performance Study of Intersection Control for Fuel-Operated Vehicles

*Natalja Pulter, Heiko Schepperle, Klemens Böhm*
(Poster R 73)

Traffic causes pollution and demands fuel. When it comes to vehicle traffic, intersections tend to be a main bottleneck. Traditional approaches to control traffic at intersections have not been designed to optimize any environmental criterion. Our objective is to design mechanisms for intersection control which minimize fuel consumption.

This is difficult because it requires a specialized infrastructure: It must allow vehicles and intersections to communicate, e.g., vehicles send their dynamic characteristics (position, speed etc.) to the intersection more or less continuously so that it can estimate the fuel consumption. In this context, the use of software agents supports the driver by reducing the necessary degree of direct interaction with the intersection.

In this paper, we quantify the fuel consumption with existing agent-based approaches for intersection control. Further, we propose a new, agent-based mechanism for intersection control, with minimization of fuel consumption as an explicit design objective. It reduces fuel consumption by up to 26% and waiting time by up to 98%, compared to traffic lights. Thus, agent-based mechanisms for intersection control may reduce fuel consumption in a way that is substantial.

### Decentralized Coordination Of Plug-in Hybrid Vehicles For Imbalance Reduction In A Smart Grid

*Stijn Vandael, Klaas De Craemer, Nelis Boucké, Tom Holvoet, Geert Deconinck*
(Poster G 65)

Intelligent electricity grids, or 'Smart Grids', are being introduced at a rapid pace. Smart grids allow the management of new distributed power generators such as solar panels and wind turbines, and innovative power consumers such as plug-in hybrid vehicles. One challenge in Smart Grids is to fulfill consumer demands while avoiding infrastructure overloads. Another challenge is to reduce imbalance costs: after ahead scheduling of production and consumption (the so-called 'load schedule'), unpredictable changes in production and consumption yield a cost for repairing this balance. To cope with these risks and costs, we propose a decentralized, multi-agent system solution for coordinated charging of PHEVs in a Smart Grid.

Essentially, the MAS utilizes an "intention graph" for expressing the flexibility of a fleet of PHEVs. Based on this flexibility, charging of PHEVs can be rescheduled in real-time to reduce imbalances. We discuss and evaluate two scheduling strategies for reducing imbalance costs: reactive scheduling and proactive scheduling. Simulations show that reactive scheduling is able to reduce imbalance costs by 14%, while proactive scheduling yields the highest imbalance cost reduction of 44%.

**Online Mechanism Design for Electric Vehicle Charging**
*Enrico H. Gerding, Valentin Robu, Sebastian Stein, David C. Parkes, Alex Rogers, Nicholas R. Jennings*
(Poster B 70)

Plug-in hybrid electric vehicles are expected to place a considerable strain on local electricity distribution networks, requiring charging to be coordinated in order to accommodate capacity constraints. We design a novel online auction protocol for this problem, wherein vehicle owners use agents to bid for power and also state time windows in which a vehicle is available for charging. This is a multi-dimensional mechanism design domain, with owners having non-increasing marginal valuations for each subsequent unit of electricity. In our design, we couple a greedy allocation algorithm with the occasional "burning" of allocated power, leaving it unallocated, in order to adjust an allocation and achieve monotonicity and thus truthfulness. We consider two variations: burning at each time step or on-departure. Both mechanisms are evaluated in depth, using data from a real-world trial of electric vehicles in the UK to simulate system dynamics and valuations. The mechanisms provide higher allocative efficiency than a fixed price system, are almost competitive with a standard scheduling heuristic which assumes non-strategic agents, and can sustain a substantially larger number of vehicles at the same per-owner fuel cost saving than a simple random scheme.

## Session C6 – Voting Protocols

*Room: 101C*
*Chair: Takayuki Ito*
*10:30 – 12:10*

**Homogeneity and Monotonicity of Distance-Rationalizable Voting Rules**
*Edith Elkind, Piotr Faliszewski, Arkadii Slinko*
(Poster R 74)

Distance rationalizability is a framework for classifying voting rules by interpreting them in terms of distances and consensus classes. It can also be used to design new voting rules with desired properties. A particularly natural and versatile class of distances that can be used for this purpose is that of votewise distances, which "lift" distances over individual votes to distances over entire elections using a suitable norm. In this paper, we continue the investigation of the properties of votewise distance-rationalizable rules initiated in Elkind et al. We describe a number of general conditions on distances and consensus classes that ensure that the resulting voting rule is homogeneous or monotone. This complements the results of Elkind et al., where the authors focus on anonymity, neutrality and consistency. We also introduce a new class of voting rules, that can be viewed as "majority variants" of classic scoring rules, and have a natural interpretation in the context of distance rationalizability.

**Possible Winners When New Alternatives Join: New Results Coming Up!**
*Lirong Xia, Jérôme Lang, Jérôme Monnot*
(Poster R 75)

In a voting system, sometimes multiple new alternatives will join the election after the voters' preferences over the initial alternatives have been revealed. Computing whether a given alternative can be a co-winner when multiple new alternatives join the election is called the possible co-winner with new alternatives (PcWNA) problem and was introduced by (Chevaleyre et al., 2010). In this paper, we show that the PcWNA problems are NP-complete for the Bucklin, Copeland$_0$, and maximin (a.k.a. Simpson) rule, even when the number of new alternatives is no more than a constant. We also show that the PcWNA problem can be solved in polynomial time for plurality with runoff. For the approval rule, we examine three different ways to extend a linear order with new alternatives, and characterize the computational complexity

of the PcWNA problem for each of them.

**The Complexity of Voter Partition in Bucklin and Fallback Voting: Solving Three Open Problems**
*Gábor Erdélyi, Lena Piras, Jörg Rothe*
(Poster B 71)

Electoral control models ways of changing the outcome of an election via such actions as adding/deleting/partitioning either candidates or voters. These actions modify an election's participation structure and aim at either making a favorite candidate win ("constructive control") or prevent a despised candidate from winning ("destructive control"). To protect elections from such control attempts, computational complexity has been used to show that electoral control, though not impossible, is computationally prohibitive. Recently, Erdélyi and Rothe proved that Brams and Sanver's fallback voting, a hybrid voting system that combines Bucklin with approval voting, is resistant to each of the standard types of control except five types of voter control. They proved that fallback voting is vulnerable to two of those control types, leaving the other three cases open. We solve these three open problems, thus showing that fallback voting is resistant to all standard types of control by partition of voters–which is a particularly important and well-motivated control type, as it models "two-district gerrymandering." Hence, fallback voting is not only fully resistant to candidate control but also fully resistant to constructive control, and it displays the broadest resistance to control currently known to hold among natural voting systems with a polynomial-time winner problem. We also show that Bucklin voting behaves almost as good in terms of control resistance. Each resistance for Bucklin voting strengthens the corresponding control resistance for fallback voting.

**An Algorithm for the Coalitional Manipulation Problem under Maximin**
*Michael Zuckerman, Omer Lev, Jeffrey S. Rosenschein*
(Poster B 72)

We introduce a new algorithm for the Unweighted Coalitional Manipulation problem under the Maximin voting rule. We prove that the algorithm gives an approximation ratio of $1\frac{2}{3}$ to the corresponding optimization problem. This is an improvement over the previously known algorithm that gave a 2-approximation. We also prove that its approximation ratio is no better than $1\frac{1}{2}$, i.e., there are instances on which a $1\frac{1}{2}$ -approximation is the best the algorithm can achieve. Finally, we prove that no algorithm can approximate the problem better than to the factor of $1\frac{1}{2}$, unless $\mathcal{P} = \mathcal{NP}$.

**Computational Complexity of Two Variants of the Possible Winner Problem**
*Dorothea Baumeister, Magnus Roos, Jörg Rothe*
(Poster G 66)

A possible winner of an election is a candidate that has, in some kind of incomplete-information election, the possibility to win in a complete extension of the election. The first type of problem we study is the Possible co-Winner with respect to the Addition of New Candidates (PcWNA) problem, which asks, given an election with strict preferences over the candidates, is it possible to make a designated candidate win the election by adding a limited number of new candidates to the election? In the case of unweighted voters we show NP-completeness of PcWNA for a broad class of pure scoring rules. We will also briefly study the case of weighted voters. The second type of possible winner problem we study is Possible Winner/co-Winner under Uncertain Voting System (PWUVS and PcWUVS). Here, uncertainty is present not in the votes but in the election rule itself. For example, PcWUVS is the problem of whether, given a set $C$ of candidates, a list of votes over $C$, a distinguished candidate $c \in C$, and a class of election rules, there is at least one election rule from this class under which $c$ wins the election. We study these two problems for a class of systems based on approval voting, the family of Copeland$^\alpha$ elections, and a certain class of scoring rules. Our main result is that it is NP-complete to determine whether there is a scoring vector that makes c win the election, if we restrict the set of possible scoring vectors for an m-candidate election to those of the form $(\alpha_1, \ldots, \alpha_{m-4}, x_1, x_2, x_3, 0)$, with $x_i = 1$ for at least one $i \in \{1, 2, 3\}$.

## Session D6 – Trust and Organisational Structure

*Room: 101D*
*Chair: Jordi Sabatier*
*13:00 – 12:10*

### Trust as Dependence: A Logical Approach

*Munindar P. Singh*
(Poster G 67)

We propose that the trust an agent places in another agent declaratively captures an architectural *connector* between the two agents. We formulate trust as a generic modality expressing a relationship between a truster and a trustee. Specifically, trust here is *definitionally independent* of, albeit constrained by, other relevant modalities such as commitments and beliefs. Trust applies to a variety of attributes of the relationship between truster and trustee. For example, an agent may trust someone to possess an important capability, exercise good judgment, or to intend to help it. Although such varieties of trust are hugely different, they respect common logical patterns. We present a logic of trust that expresses such patterns as reasoning postulates concerning the static representation of trust, its dynamics, and its relationships with teamwork and other agent interactions. In this manner, the proposed logic illustrates the general properties of trust that reflect natural intuitions, and can facilitate the engineering of multiagent systems.

### Multi-Layer Cognitive Filtering by Behavioral Modeling

*Zeinab Noorian, Stephen Marsh, Michael Fleming*
(Poster G 68)

In the absence of legal enforcement procedures for the participants of an open e-marketplace, trust and reputation systems are central for resisting against threats from malicious agents. Such systems provide mechanisms for identifying the participants who disseminate unfair ratings. However, it is possible that some of the honest participants are also victimized as a consequence of the poor judgement of these systems. In this paper, we propose a two-layer filtering algorithm that cognitively elicits the behavioral characteristics of the participating agents in an e-marketplace. We argue that the notion of unfairness does not exclusively refer to deception but can also imply differences in dispositions. The proposed filtering approach aims to go beyond the inflexible judgements on the quality of participants and instead allows the human dispositions that we call optimism, pessimism and realism to be incorporated into our trustworthiness evaluations. Our proposed filtering algorithm consists of two layers. In the first layer, a consumer agent measures the competency of its neighbors for being a potentially helpful adviser. Thus, it automatically disqualifies the deceptive agents and/or the newcomers that lack the required experience. Afterwards, the second layer measures the credibility of the surviving agents of the previous layer on the basis of their behavioral models. This tangible view of trustworthiness evaluation boosts the confidence of human users in using a web-based agent-oriented e-commerce application.

### Argumentation-Based Reasoning in Agents with Varying Degrees of Trust

*Simon Parsons, Yuqing Tang, Elizabeth Sklar, Peter McBurney, Kai Cai*
(Poster B 73)

In any group of agents, trust plays an important role. The degree to which agents trust one another will inform what they believe, and, as a result the reasoning that they perform and the conclusions that they come to when that involves information from other agents. In this paper we consider a group of agents with varying degrees of trust of each other, and examine the combinations of trust with the argumentation-based reasoning that they can carry out. The question we seek to answer is "What is the relationship between the trust one agent has in another and the conclusions that it can draw using information from that agent?", and show that there are a range of answers depending upon the way that the agents deal with trust.

### A Particle Filter for Bid Estimation in Ad Auctions with Periodic Ranking Observations

*David Pardoe, Peter Stone*
(Poster G 69)

Keyword auctions are becoming increasingly important in today's electronic marketplaces. One of their most challenging aspects is the limited amount of information revealed about other advertisers. In this paper, we present a particle filter that can be used to estimate the bids of other advertisers given a periodic ranking of their bids. This particle filter makes use of models of the bidding behavior of other advertisers, and so we also show how such models can be learned from past bidding data. In experiments in the Ad Auction scenario of the Trading Agent Competition, the combination of this particle filter and

bidder modeling outperforms all other bid estimation methods tested.

**Conviviality Measures**
*Patrice Caire, Baptiste Alcalde, Leendert van der Torre, Chattrakul Sombattheera*
(Poster R 76)

Conviviality has been introduced as a social science concept for multiagent systems to highlight soft qualitative requirements like user friendliness of systems. In this paper we introduce formal conviviality measures for dependence networks using a coalitional game theoretic framework, which we contrast with more traditional efficiency and stability measures. Roughly, more opportunities to work with other people increases the conviviality, whereas larger coalitions may decrease the efficiency or stability of these involved coalitions. We first introduce assumptions and requirements, then we introduce a classification, and finally we introduce the conviviality measures. We use a running example from robotics to illustrate the measures.

## Session A7 – Argumentation and Negotiation

*Room: 101A*
*Chair: Simon Parsons*
*13:00 – 14:40*

**Choosing Persuasive Arguments for Action**
*Elizabeth Black, Katie Atkinson*
(Poster B 74)

We present a dialogue system that allows agents to exchange arguments in order to come to an agreement on how to act. When selecting arguments to assert, an agent uses a model of what is important to the recipient agent. The system lets the agents agree to an action that each finds acceptable, but does not necessarily demand that they resolve their differing preferences. We present an analysis of the behaviour of our system and develop a mechanism with which an agent can develop a model of another's preferences.

**Argumentation Strategies for Plan Resourcing**
*Chukwuemeka D. Emele, Timothy J. Norman, Simon Parsons*
(Poster G 70)

What do I need to say to convince you to do something? This is an important question for an autonomous agent deciding whom to approach for a resource or for an action to be done. Were similar requests granted from similar agents in similar circumstances? What arguments were most persuasive? What are the costs involved in putting certain arguments forward? In this paper we present an agent decision-making mechanism where models of other agents are refined through evidence from past dialogues, and where these models are used to guide future argumentation strategy. We empirically evaluate our approach to demonstrate that decision-theoretic and machine learning techniques can both significantly improve the cumulative utility of dialogical outcomes, and help to reduce communication overhead.

**Multi-Criteria Argument Selection In Persuasion Dialogues**
*Tom L. van der Weide, Frank Dignum, John-Jules Ch. Meyer, H. Prakken, Gerard A.W. Vreeswijk*
(Poster G 71)

The main goal of a persuasion dialogue is to persuade, but agents may have a number of additional goals concerning the dialogue duration, how much and what information is shared or how aggressive the agent is. Several criteria have been proposed in the literature covering different aspects of what may matter to an agent, but it is not clear how to combine these criteria that are often incommensurable and partial. This paper is inspired by multi-attribute decision theory and considers argument selection as decision-making where multiple criteria matter. A meta-level argumentation system is proposed to argue about what argument an agent should select in a given persuasion dialogue. The criteria and sub-criteria that matter to an agent are structured hierarchically into a value tree and meta-level argument schemes are formalized that use a value tree to justify what argument the agent should select. In this way, incommensurable and partial criteria can be combined.

**Analyzing Intra-Team Strategies for Agent-Based Negotiation Teams**
*Víctor Sánchez-Anguix, Vicente Julián, Vicente Botti, Ana García-Fornes*
(Poster R 77)

An agent-based negotiation team is a group of two or more agents with their own and possibly conflicting preferences who join together as a single negotiating party because they share a common goal which is related to the negotiation. Scenarios involving negotiation teams require coordination among party mem-

bers in order to reach a good agreement for all of the party members. An intra-team strategy defines what decisions are taken by the negotiation team and when and how these decisions are taken. Thus, they are tightly linked with the results obtained by the team in a negotiation process. Environmental conditions affect the performance of the different intra-team strategies in different ways. Thus, team members need to analyze their environment in order to select the most appropriate strategy according to the current conditions. In this paper, we analyze how environmental conditions affect different intra-team strategies in order to provide teams with the knowledge necessary to select the proper intra-team strategy.

**The Effect of Expression of Anger and Happiness in Computer Agents on Negotiations with Humans**
*Celso M. de Melo, Peter Carnevale, Jonathan Gratch*
(Poster R 78)

There is now considerable evidence in social psychology, economics, and related disciplines that emotion plays an important role in negotiation. For example, humans make greater concessions in negotiation to an opposing human who expresses anger, and they make fewer concessions to an opponent who expresses happiness, compared to a no-emotion-expression control. However, in AI, despite the wide interest in negotiation as a means to resolve differences between agents and humans, emotion has been largely ignored. This paper explores whether expression of anger or happiness by computer agents, in a multi-issue negotiation task, can produce effects that resemble effects seen in human-human negotiation. The paper presents an experiment where participants play with agents that express emotions (anger vs. happiness vs. control) through different modalities (text vs. facial displays). An important distinction in our experiment is that participants are aware that they negotiate with computer agents. The data indicate that the emotion effects observed in past work with humans also occur in agent-human negotiation, and occur independently of modality of expression. The implications of these results are discussed for the fields of automated negotiation, intelligent virtual agents and artificial intelligence.

## Session B7 – Planning

*Room: 101B*
*Chair: Sven Koenig*
*13:00 – 14:40*

**Toward Error-Bounded Algorithms for Infinite-Horizon DEC-POMDPs**
*Jilles S. Dibangoye, Abdel-Illah Mouaddib, Brahim Chaib-draa*
(Poster B 75)

Over the past few years, attempts to scale up infinite-horizon DEC-POMDPs are mainly due to approximate algorithms, but without the theoretical guarantees of their exact counterparts. In contrast, $\epsilon$-optimal methods have only theoretical significance but are not efficient in practice. In this paper, we introduce an algorithmic framework ($\beta$-PI) that exploits the scalability of the former while preserving the theoretical properties of the latter. We build upon $\beta$-PI a family of approximate algorithms that can find (provably) errorbounded solutions in reasonable time. Among this family, H-PI uses a branch-and-bound search method that computes a near-optimal solution over distributions over histories experienced by the agents. These distributions often lie near structured, low-dimensional subspace embedded in the high-dimensional sufficient statistic. By planning only on this subspace, H-PI successfully solves all tested benchmarks, outperforming standard algorithms, both in solution time and policy quality.

**Distributed Model Shaping for Scaling to Decentralized POMDPs with Hundreds of Agents**
*Prasanna Velagapudi, Pradeep Varakantham, Katia Sycara, Paul Scerri*
(Poster B 76)

The use of distributed POMDPs for cooperative teams has been severely limited by the incredibly large joint policyspace that results from combining the policy-spaces of the individual agents. However, much of the computational cost of exploring the entire joint policy space can be avoided by observing that in many domains important interactions between agents occur in a relatively small set of scenarios, previously defined as *coordination locales* (CLs). Moreover, even when numerous interactions *might* occur, given a set of individual policies there are relatively few *actual* interactions. Exploiting this observation and building on an existing model shaping algorithm, this paper

presents D-TREMOR, an algorithm in which cooperative agents iteratively generate individual policies, identify and communicate possible interactions between their policies, shape their models based on this information and generate new policies. D-TREMOR has three properties that jointly distinguish it from previous DEC-POMDP work: (1) it is completely distributed; (2) it is scalable (allowing 100 agents to compute a "good" joint policy in under 6 hours) and (3) it has low communication overhead. D-TREMOR complements these traits with the following key contributions, which ensure improved scalability and solution quality: (a) techniques to ensure convergence; (b) faster approaches to detect and evaluate CLs; (c) heuristics to capture dependencies between CLs; and (d) novel shaping heuristics to aggregate effects of CLs. While the resulting policies are not globally optimal, empirical results show that agents have policies that effectively manage uncertainty and the joint policy is better than policies generated by independent solvers.

**Efficient Planning in R-max**
*Marek Grześ, Jesse Hoey*
(Poster G 72)

PAC-MDP algorithms are particularly efficient in terms of the number of samples obtained from the environment which are needed by the learning agents in order to achieve a near optimal performance. These algorithms however execute a time consuming planning step after each new state-action pair becomes known to the agent, that is, the pair has been sampled sufficiently many times to be considered as known by the algorithm. This fact is a serious limitation on broader applications of these kind of algorithms. This paper examines the planning problem in PAC-MDP learning. Value iteration, prioritized sweeping, and backward value iteration are investigated. Through the exploitation of the specific nature of the planning problem in the considered reinforcement learning algorithms, we show how these planning algorithms can be improved. Our extensions yield significant improvements in all evaluated algorithms, and standard value iteration in particular. The theoretical justification to all contributions is provided and all approaches are further evaluated empirically. With our extensions, we managed to solve problems of sizes which have never been approached by PAC-MDP learning in the existing literature.

**Multiagent Argumentation for Cooperative Planning in DeLP-POP**
*Pere Pardo, Sergio Pajares, Eva Onaindia, Pilar Dellunde, Lluís Godo*
(Poster G 73)

This contribution proposes a model for argumentation-based multi-agent planning, with a focus on cooperative scenarios. It consists in a multiagent extension of DeLP-POP, partial order planning on top of argumentation-based defeasible logic programming. In DeLP-POP, actions and arguments (combinations of rules and facts) may be used to enforce some goal, if their conditions (are known to) apply and arguments are not defeated by other arguments applying. In a cooperative planning problem a team of agents share a set of goals but have diverse abilities and beliefs. In order to plan for these goals, agents start a stepwise dialogue consisting of exchanges of plan proposals, plus arguments against them. Since these dialogues instantiate an A search algorithm, these agents will find a solution if some solution exists, and moreover, it will be provably optimal (according to their knowledge).

## Session C7 – Game Theory II

*Room: 101C*
*Chair: Milind Tambe*
*13:00 – 14:40*

**Computing a Self-Confirming Equilibrium in Two-Player Extensive-Form Games**
*Nicola Gatti, Fabio Panozzo, Sofia Ceppi*
(Poster R 79)

The Nash equilibrium is the most commonly adopted solution concept for non-cooperative interaction situations. However, it underlays on the assumption of common information that is hardly verified in many practical situations. When information is not common, the appropriate game theoretic solution concept is the self-confirming equilibrium. It requires that every agent plays the best response to her beliefs and that the beliefs are correct on the equilibrium path. We present, to the best of our knowledge, the first study on the computation of a self-confirming equilibrium for two-player extensive-form games. We provide algorithms, we analyze the computational complexity, and we experimentally evaluate the performance of our algorithms in terms of computational time.

**Computing Time-Dependent Policies for Patrolling Games with Mobile Targets**
*Branislav Bošanský, Viliam Lisý, Michal Jakob, Michal Pěchouček*
(Poster B 77)

We study how a mobile defender should patrol an area to protect multiple valuable targets from being attacked by an attacker. In contrast to existing approaches, which assume stationary targets, we allow the targets to move through the area according to an a priori known, deterministic movement schedules. We represent the patrol area by a graph of arbitrary topology and do not put any restrictions on the movement schedules. We assume the attacker can observe the defender and has full knowledge of the strategy the defender employs. We construct a game-theoretic formulation and seek defender's optimal randomized strategy in a Stackelberg equilibrium of the game. We formulate the computation of the strategy as a mathematical program whose solution corresponds to an optimal time-dependent Markov policy for the defender. We also consider a simplified formulation allowing only stationary defender's policies which are generally less effective but are computationally significantly cheaper to obtain. We provide experimental evaluation examining this trade-off on a set of test problems covering various topologies of the patrol area and various movement schedules of the targets.

**Quality-bounded Solutions for Finite Bayesian Stackelberg Games: Scaling up**
*Manish Jain, Christopher Kiekintveld, Milind Tambe*
(Poster G 74)

The fastest known algorithm for solving General Bayesian Stackelberg games with a finite set of follower (adversary) types have seen direct practical use at the LAX airport for over 3 years; and currently, an (albeit non-Bayesian) algorithm for solving these games is also being used for scheduling air marshals on limited sectors of international flights by the US Federal Air Marshals Service. These algorithms find optimal randomized security schedules to allocate limited security resources to protect targets. As we scale up to larger domains, including the full set of flights covered by the Federal Air Marshals, it is critical to develop newer algorithms that scale-up significantly beyond the limits of the current state-of-theart of Bayesian Stackelberg solvers. In this paper, we present a novel technique based on a hierarchical decomposition and branch and bound search over the follower type space, which may be applied to different Stackelberg game solvers. We have applied this technique to different solvers, resulting in: (i) A new exact algorithm called HBGS that is orders of magnitude faster than the best known previous Bayesian solver for general Stackelberg games; (ii) A new exact algorithm called HBSA which extends the fastest known previous security game solver towards the Bayesian case; and (iii) Approximation versions of HBGS and HBSA that show significant improvements over these newer algorithms with only 12% sacrifice in the practical solution quality.

**Approximation Methods for Infinite Bayesian Stackelberg Games: Modeling Distributional Payoff Uncertainty**
*Christopher Kiekintveld, Janusz Marecki, Milind Tambe*
(Poster G 75)

Game theory is fast becoming a vital tool for reasoning about complex real-world security problems, including critical infrastructure protection. The game models for these applications are constructed using expert analysis and historical data to estimate the values of key parameters, including the preferences and capabilities of terrorists. In many cases, it would be natural to represent uncertainty over these parameters using continuous distributions (such as uniform intervals or Gaussians). However, existing solution algorithms are limited to considering a small, finite number of possible attacker types with different payoffs. We introduce a general model of infinite Bayesian Stackelberg security games that allows payoffs to be represented using continuous payoff distributions. We then develop several techniques for finding approximate solutions for this class of games, and show empirically that our methods offer dramatic improvements over the current state of the art, providing new ways to improve the robustness of security game models.

**Solving Stackelberg Games with Uncertain Observability**
*Dmytro Korzhyk, Vincent Conitzer, Ronald Parr*
(Poster G 76)

Recent applications of game theory in security domains use algorithms to solve a Stackelberg model, in which one player (the leader) first commits to a mixed strategy and then the other player (the follower) observes that strategy and best-responds to

it. However, in real-world applications, it is hard to determine whether the follower is actually able to observe the leader's mixed strategy before acting. In this paper, we model the uncertainty about whether the follower is able to observe the leader's strategy as part of the game (as proposed in the extended version of Yin et al. [17]). We describe an iterative algorithm for solving these games. This algorithm alternates between calling a Nash equilibrium solver and a Stackelberg solver as subroutines. We prove that the algorithm finds a solution in a finite number of steps and show empirically that it runs fast on games of reasonable size. We also discuss other properties of this methodology based on the experiments.

## Session D7 – Virtual Agents II

*Room: 101D*
*Chair: Von Wun Sun*
*13:00 – 14:40*

**A Style Controller for Generating Virtual Human Behaviors**
*Chung-Cheng Chiu, Stacy Marsella*
(Poster G 77)

Creating a virtual character that exhibits realistic physical behaviors requires a rich set of animations. To mimic the variety as well as the subtlety of human behavior, we may need to animate not only a wide range of behaviors but also variations of the same type of behavior influenced by the environment and the state of the character, including the emotional and physiological state. A general approach to this challenge is to gather a set of animations produced by artists or motion capture. However, this approach can be extremely costly in time and effort. In this work, we propose a model that can learn styled motion generation and an algorithm that produce new styles of motions via style interpolation. The model takes a set of styled motions as training samples, and can create new motions that are the generalization among given styles of motions. Our style interpolation algorithm can blend together motions with distinct styles, and it also helps improve the performance of previous work. We verify our algorithm using walking motions of different styles, and the experimental results show that our method is significantly better than previous work.

**The Face of Emotions: A Logical Formalization of Expressive Speech Acts**
*Nadine Guiraud, Dominique Longin, Emiliano Lorini, Sylvie Pesty, Jérémy Rivière*
(Poster G 78)

In this paper, we merge speech act theory, emotion theory, and logic. We propose a modal logic that integrates the concepts of belief, goal, ideal and responsibility and that allows to describe what a given agent expresses in the context of a conversation with another agent. We use the logic in order to provide a systematic analysis of expressive speech acts, that is, speech acts that are aimed at expressing a given emotion (e.g. to apologize, to thank, to reproach, etc.).

**I've Been Here Before! Location and Appraisal in Memory Retrieval**
*Paulo F. Gomes, Carlos Martinho, Ana Paiva*
(Poster B 78)

The objective of our current work was to create a model for agent memory retrieval of emotionally relevant episodes. We analyzed agent architectures that support memory retrieval realizing that none fulfilled all of our requirements. We designed an episodic memory retrieval model consisting of two main steps: *location ecphory*, in which the agent's current location is matched against stored memories associated locations; and *recollective experience*, in which memories that had a positive match are re-appraised. We implemented our model and used it to drive the behavior of characters in a game application. We recorded the application running and used the videos to create a non-interactive evaluation. The evaluation's results are consistent with our hypothesis that agents with memory retrieval of emotionally relevant episodes would be perceived as more believable than similar agents without it.

**From Body Space to Interaction Space - Modeling Spatial Cooperation for Virtual Humans**
*Nhung Nguyen, Ipke Wachsmuth*
(Poster B 79)

This paper introduces a model which connects representations of the space surrounding a virtual humanoid's body with the space it shares with several interaction partners. This work intends to support virtual humans (or humanoid robots) in near space interaction and is inspired by studies from cognitive neurosciences on the one hand and social interaction studies on the other hand. We present our work on

learning the body structure of an articulated virtual human by using data from virtual touch and proprioception sensors. The results are utilized for a representation of its reaching space, the so-called peripersonal space. In interpersonal interaction involving several partners, their peripersonal spaces may overlap and establish a shared reaching space. We define it as their interaction space, where cooperation takes place and where actions to claim or release spatial areas have to be adapted, to avoid obstructions of the other's movements. Our model of interaction space is developed as an extension of Kendon's F-formation system, a foundational theory of how humans orient themselves in space when communicating. Thus, interaction space allows for analyzing the spatial arrangement (i.e., body posture and orientation) between multiple interaction partners and the extent of space they share. Peripersonal and interaction space are modeled as potential fields to control the virtual human's behavior strategy. As an example we show how the virtual human can relocate object positions toward or away from locations reachable for all partners, and thus influencing the degree of cooperation in an interaction task.

**Effect of Time Delays on Agents' Interaction Dynamics**
*Ken Prepin, Catherine Pelachaud*
(Poster G 79)

While speaking about social interaction, psychology claims as crucial the temporal correlations between interactants' behaviors: to give to their partners a feeling of natural interaction, interactants, be human, robotic or virtual, must be able to react on appropriate time. Recent approaches consider autonomous agents as dynamical systems and the interaction as a coupling between these systems. These approaches solve the issue of time handling and enable to model synchronization and turn-taking as phenomenon emerging with the coupling. But when complex computations are added to their architecture, such as processing of video and audio signals, delays appear within the interaction loop and disrupt this coupling. We model here a dyad of agents where processing delays are controlled. These agents, driven by oscillators, synchronize and take turns when there is no delay. We describe the methodology enabling to evaluate the synchrony and turn-taking emergence. We test oscillators coupling properties when there is no delay: coupling occurs if coupling strength is inferior to the parameter controlling oscillators natural period and if the ratio between oscillators periods is inferior to 1/2. We quantify the maximal delays between agents which do not disrupt the interaction: the maximal delay tolerated by agents is proportional to the natural period of the coupled system and to the strength of the coupling. These results are put in perspective with the different time constraints of human-human and human-agent interactions.

# Main Program – Extended Abstracts

## Thursday

### Red Session

*Room: 201*
*17:00 – 18:00*

**A Computational Model of Achievement Motivation for Artificial Agents**
*Kathryn E. Merrick*
(Poster R 1)

Computational models of motivation are tools that artificial agents can use to autonomously identify, prioritize, and select the goals they will pursue. Previous research has focused on developing computational models of arousal-based theories of motivation, including novelty, curiosity and interest. However, arousal-based theories represent only one aspect of motivation. In humans, for example, curiosity is tempered by other motivations such as the need for health, safety, competence, a sense of belonging, esteem from others or influence over others. To create artificial agents that can identify and prioritize their goals according to this broader range of needs, new kinds of computational models of motivation are required. This paper expands our 'motivation toolbox' with a new computational model of achievement motivation for artificial agents. The model uses sigmoid curves to model approach of success and avoidance of failure. An experiment from human psychology is simulated to test the new model in virtual agents. The results are compared to human results and existing theoretical and computational models. Results show that virtual agents using our model exhibit statistically similar goal-selection characteristics to humans with corresponding motive profiles. In addition, our model outperforms existing models of achievement motivation in this respect.

**Incremental DCOP Search Algorithms for Solving Dynamic DCOPs**
*William Yeoh, Pradeep Varakantham, Xiaoxun Sun, Sven Koenig*
(Poster R 2)

Distributed constraint optimization problems (DCOPs) are well-suited for modeling multi-agent coordination problems. However, most research has focused on developing algorithms for solving static DCOPs. In this paper, we model dynamic DCOPs as sequences of (static) DCOPs with changes from one DCOP to the next one in the sequence. We introduce the ReuseBounds procedure, which can be used by any-space ADOPT and any-space BnB-ADOPT to find cost-minimal solutions for all DCOPs in the sequence faster than by solving each DCOP individually. This procedure allows those agents that are guaranteed to remain unaffected by a change to reuse their lower and upper bounds from the previous DCOP when solving the next one in the sequence. Our experimental results show that the speedup gained from this procedure increases with the amount of memory the agents have available.

**MetaTrust: Discriminant Analysis of Local Information for Global Trust Assessment**
*Liu Xin, Gilles Tredan, Anwitaman Datta*
(Poster R 3)

A traditional approach to reasoning about the trustworthiness of a transaction is to determine the trustworthiness of the specific agent involved, based on its past behavior. As a departure from such traditional trust models, we propose a transaction centered trust model (MetaTrust) where an agent uses its previous transactions to assess the trustworthiness of a potential transaction based on associated meta-information, which is capable of distinguishing successful transactions from unsuccessful ones. This meta information is harnessed using a machine learning algorithm (namely, discriminant analysis) to extract relationships between the potential transaction and previous transactions.

**Efficient Penalty Scoring Functions for Group Decision-making with TCP-nets**
*Minyi Li, Quoc Bao Vo, Ryszard Kowalczyk*
(Poster R 4)

This paper studies the problem of collective decision-making in combinatorial domain where the agents' preferences are represented by qualitative models with TCP-nets (Tradeoffs-enhanced Conditional Preference Network). The features of TCP-nets enable us to easily encode human preferences and the relative importance between the decision variables; however, many group decision-making methods re-

quire numerical measures of degrees of desirability of alternative outcomes. To permit a natural way for preference elicitation while providing quantitative comparisons between outcomes, we present a computationally efficient approach that compiles individual TCP-nets into ordinal penalty scoring functions. After the individual penalty scores are computed, we further define a collective penalty scoring function to aggregate multiple agents' preferences.

**A Curious Agent for Network Anomaly Detection**
*Kamran Shafi, Kathryn E. Merrick*
(Poster R 5)

This paper presents a novel approach to intrusion detection using curious agents to detect anomalies in network data. Curious agents use computational models of novelty-seeking behavior and interest, based on human curiosity, to reason about their experiences in their environment. They are online, single-pass agents that respond to the similarity, frequency and recentness of their experiences. As such, they combine a number of important characteristics required for intrusion detection. This paper presents a generic, curious reflex agent model for network intrusion detection and the results of experiments with a number of variants of this model. Specifically, five different models of curiosity are compared for their ability to detect first instances of attacks in the KDD Cup data set. Results show that our curious agents can achieve high detection rates for intrusions, with moderate false-positive rates.

**Agents, Pheromones, and Mean-Field Models**
*H. Van Dyke Parunak*
(Poster R 6)

Some agent-based models use digital analogs of insect pheromones for coordination. Such models are intermediate between classical agent-based models and equation-based "mean field" models. Their position in this range can be adjusted by pheromone parameters (notably, the propagation factor).

**Basis Function Discovery using Spectral Clustering and Bisimulation Metrics**
*Gheorghe Comanici, Doina Precup*
(Poster R 7)

We study the problem of automatically generating features for function approximation in reinforcement learning. We build on the work of Mahadevan and his colleagues, who pioneered the use of spectral clustering methods for basis function construction. Their methods work on top of a graph that captures state adjacency. Instead, we use bisimulation metrics in order to provide state distances for spectral clustering. The advantage of these metrics is that they incorporate reward information in a natural way, in addition to the state transition information. We provide theoretical bounds on the quality of the obtained approximation, which justify the importance of incorporating reward information. We also demonstrate empirically that the approximation quality improves when bisimulation metrics are used instead of the state adjacency graph in the basis function construction process.

**Incentive Compatible Influence Maximization in Social Networks and Application to Viral Marketing**
*Mayur Mohite, Y. Narahari*
(Poster R 8)

Information diffusion and influence maximization are important and extensively studied problems in social networks. Various models and algorithms have been proposed in the literature in the context of the influence maximization problem. A crucial assumption in all these studies is that the influence probabilities are known to the social planner. This assumption is unrealistic since the influence probabilities are usually private information of the individual agents and strategic agents may not reveal them truthfully. Moreover, the influence probabilities could vary significantly with the type of the information flowing in the network and the time at which the information is propagating in the network. In this paper, we use a mechanism design approach to elicit influence probabilities truthfully from the agents. Our main contribution is to design a scoring rule based mechanism in the context of the influencer-influencee model. In particular, we show the incentive compatibility of the mechanisms and propose a reverse weighted scoring rule based mechanism as an appropriate mechanism to use.

**On Optimal Agendas for Package Deal Negotiation**
*Shaheen Fatima, Michael Wooldridge, Nicholas R. Jennings*
(Poster R 9)

This paper analyzes bilateral multi-issue negotiation where the issues are indivisible, there are time constraints in the form of deadlines and discount factors. The issues are negotiated using the package deal procedure. The set of issues to be negotiated is called the negotiation agenda. The agenda is crucial since the outcome of negotiation depends on the agenda. This paper therefore looks at the decision making involved in choosing a negotiation agenda. The scenario we look at is as follows. There are $m > 2$ issues available for negotiation. But from these, an agent must choose $g < m$ issues and negotiate on them. Thus the problem for an agent is to choose an agenda (i.e, a subset of $g$ issues). Clearly, from all possible agendas (i.e., all possible combinations of $g$ issues), an agent must choose the one that maximizes its expected utility and is therefore its optimal agenda. To this end, this paper presents polynomial time methods for choosing an agent's optimal agenda.

**An Abstract Framework for Reasoning About Trust**
*Elisabetta Erriquez, Wiebe van der Hoek, Michael Wooldridge*
(Poster R 10)

We present an abstract framework that allows agents to form coalitions with agents that they believe to be trustworthy. In contrast to many other models, we take the notion of distrust to be our key social concept. We use a graph theoretic model to capture the distrust relations within a society, and use this model to formulate several notions of mutually trusting coalitions. We then investigate principled techniques for how the information present in our distrust model can be aggregated to produce individual measures of how trustworthy an agent is considered to be by a society.

**Message-Passing Algorithms for Large Structured Decentralized POMDPs**
*Akshat Kumar, Shlomo Zilberstein*
(Poster R 11)

Decentralized POMDPs provide a rigorous framework for multi-agent decision-theoretic planning. However, their high complexity has limited scalability. In this work, we present a promising new class of algorithms based on probabilistic inference for infinite-horizon ND-POMDPs—a restricted Dec-POMDP model. We first transform the policy optimization problem to that of likelihood maximization in a mixture of dynamic Bayes nets (DBNs). We then develop the Expectation-Maximization (EM) algorithm for maximizing the likelihood in this representation. The EM algorithm for ND-POMDPs lends itself naturally to a simple messagepassing paradigm guided by the agent interaction graph. It is thus highly scalable w.r.t. the number of agents, can be easily parallelized, and produces good quality solutions.

**Jogger: Models for Context-Sensitive Reminding**
*Ece Kamar, Eric Horvitz*
(Poster R 12)

We describe research on principles of context-sensitive reminding that show promise for serving in systems that work to jog peoples' memories about information that they may forget. The methods center on the construction and use of a set of distinct probabilistic models that predict (1) items that may be forgotten, (2) the expected relevance of the items in a situation, and (3) the cost of interruption associated with alerting about a reminder. We describe the use of this set of models in the Jogger prototype that employs predictions and decision-theoretic optimization to compute the value of reminders about meetings.

**Spatio-Temporal A* Algorithms for Offline Multiple Mobile Robot Path Planning**
*Wenjie Wang, Wooi Boon Goh*
(Poster R 13)

This paper presents an offline collision-free path planning algorithm for multiple mobile robots using a 2D spatial-time map. In this decoupled approach, a centralized planner uses a Spatio-Temporal A* algorithm to find the lowest time cost path for each robot in a sequentially order based on its assigned priority. Improvements in viable path solutions using wait time insertion and adaptive priority reassignment strategies are discussed.

**Influence of Head Orientation in Perception of Personality Traits in Virtual Agents**
*Diana Arellano, Nikolaus Bee, Kathrin Janowski, Elisabeth André, Javier Varona, Francisco J. Perales*
(Poster R 14)

The aim of this research is to explore the influence of static visual cues on the perception of a character's personality traits: extraversion, agreeableness and emotional stability. To measure how users perceived personality, we conducted a web-based study with 133 subjects who rated 54 images of a virtual character with varying head orientations and gaze.

**Conflict Resolution with Argumentation Dialogues**
*Xiuyi Fan, Francesca Toni*
(Poster R 15)

Conflicts exist in multi-agent systems for a number of reasons: agents have different interests and desires; agents hold different beliefs; agents make different assumptions. To resolve conflicts, agents need to better convey information to each other and facilitate fair negotiations yielding jointly agreeable outcomes. We present a two-agent, dialogical conflict resolution scheme developed with the Assumption-Based Argumentation (ABA) framework.

**Reasoning Patterns in Bayesian Games**
*Dimitrios Antos, Avi Pfeffer*
(Poster R 16)

Bayesian games have been traditionally employed to describe and analyze situations in which players have private information or are uncertain about the game being played. However, computing Bayes-Nash equilibria can be costly, and becomes even more so if the common prior assumption (CPA) has to be abandoned, which is sometimes necessary for a faithful representation of real-world systems. We propose using the theory of reasoning patterns in Bayesian games to circumvent some of these difficulties. The theory has been used successfully in common knowledge (non-Bayesian) games, both to reduce the computational cost of finding an equilibrium and to aid human decision-makers in complex decisions. In this paper, we first show that reasoning patterns exist for every decision of every Bayesian game, in which the acting agent has a reason to deliberate. This implies that reasoning patterns are a complete characterization of the types of reasons an agent might have for making a decision. Second, we illustrate practical applications of reasoning patterns in Bayesian games, which allow us to answer questions that would otherwise not be easy in traditional analyses, or would be extremely costly. We thus show that the reasoning patterns can be a useful framework in analyzing complex social interactions.

**Using Coalitions of Wind Generators and Electric Vehicles for Effective Energy Market Participation**
*Matteo Vasirani, Ramachandra Kota, Renato L.G. Cavalcante, Sascha Ossowski, Nicholas R. Jennings*
(Poster R 17)

Wind power is becoming a significant source of electricity in many countries. However, the inherent uncertainty of wind generators does not allow them to participate in the forward electricity markets. In this paper, we foster a tighter integration of wind power into electricity markets by using a multi-agent coalition formation approach to form virtual power plants of wind generators and electric vehicles. We identify the four different phases in the life-cycle of a VPP, each characterised by its own challenges that need to be addressed.

**Negotiation Over Decommitment Penalty**
*Bo An, Victor Lesser*
(Poster R 18)

We consider the role of negotiation in deciding decommitment penalties. In our model, agents simultaneously negotiate over both the contract price and decommitment penalty in the contracting game and then decide whether to decommit from contracts in the decommitment game. Experimental results show that setting penalties through negotiation achieved higher social welfare than other exogenous penalty setting mechanisms.

**Ship Patrol: Multiagent Patrol under Complex Environmental Conditions**
*Noa Agmon, Daniel Urieli, Peter Stone*
(Poster R 19)

In the problem of multiagent patrol, a team of agents is required to repeatedly visit a target area in order to monitor possible changes in state. The growing popularity of this problem comes mainly from its immediate applicability to a wide variety of domains. In this paper we concentrate on frequency-based patrol, in which the agents' goal is to optimize a frequency criterion, namely, minimizing the time between visits to a set of interest points. In situations with varying environmental conditions, the influence of changes in the conditions on the cost of travel may be immense. For example, in marine environments, the travel time of ships depends on parameters such as

wind, water currents, and waves. Such environments raise the need to consider a new multiagent patrol strategy which divides the given area into regions in which more than one agent is active, for improving frequency. We prove that in general graphs this problem is intractable, therefore we focus on simplified (yet realistic) cyclic graphs with possible inner edges. Although the problem remains generally intractable in such graphs, we provide a heuristic algorithm that is shown to significantly improve point-visit frequency compared to other patrol strategies.

**Empirical and Theoretical Support for Lenient Learning**
*Daan Bloembergen, Michael Kaisers, Karl Tuyls*
(Poster R 20)

Recently, an evolutionary model of Lenient Q-learning (LQ) has been proposed, providing theoretical guarantees of convergence to the global optimum in cooperative multi-agent learning. However, experiments reveal discrepancies between the predicted dynamics of the evolutionary model and the actual learning behavior of the Lenient Q-learning algorithm, which undermines its theoretical foundation. Moreover it turns out that the predicted behavior of the model is more desirable than the observed behavior of the algorithm. We propose the variant Lenient Frequency Adjusted Q-learning (LFAQ) which inherits the theoretical guarantees and resolves this issue. The advantages of LFAQ are demonstrated by comparing the evolutionary dynamics of lenient vs non-lenient Frequency Adjusted Q-learning. In addition, we analyze the behavior, convergence properties and performance of these two learning algorithms empirically in the Battle of the Sexes (BoS) and the Stag Hunt (SH). Significant deviations arise from the introduction of leniency, leading to profound performance gains in coordination games against both lenient and non-lenient learners.

**A Formal Framework for Reasoning about Goal Interactions**
*Michael Winikoff*
(Poster R 21)

A defining characteristic of intelligent software agents is their ability to flexibly and reliably pursue goals, and many modern agent platforms provide some form of goal construct. However, these platforms are surprisingly naive in their handling of interactions between goals. Whilst previous work has provided mechanisms to identify and react appropriately to various sorts of interactions, it has not provided a framework for reasoning about goal interactions that is generic, extensible, formally described, and that covers a range of interaction types.

**On-line Reasoning for Institutionally-Situated BDI agents**
*Tina Balke, Marina De Vos, Julian Padget, Dimitris Traskas*
(Poster R 22)

Institutions offer the promise of a means to govern open systems, in particular, open multi-agent systems. Research in logics and their derived tools now support the specification, verification and enactment of institutions (or organizations, depending on the terminology of the tool). Most effort to date has tended to focus on the static properties of institutions, such as whether a particular state of affairs is reachable or not from a given set of initial conditions. Such models are useful in forcing the designer to state their intentions precisely, and for testing (static) properties. We call this off-line reasoning. We identify two problems in the direct utilization of off-line models in the governance of live systems: (i) static model artefacts that are typically aspects of agent behaviour in the dynamic model (ii) over-specification of constraints on actions, leading to undue limitation of agent autonomy. Agents need to be able to query an institution for (dynamic) properties. We call this on-line reasoning. In this paper we present a methodology to extract the on-line specification from an off-line one and use it to support BDI agents to realize a norm-governed multi-agent system.

**Strategy Purification**
*Sam Ganzfried, Tuomas Sandholm, Kevin Waugh*
(Poster R 23)

There has been significant recent interest in computing good strategies for large games. Most prior work involves computing an approximate equilibrium strategy in a smaller abstract game, then playing this strategy in the full game. In this paper, we present a modification of this approach that works by constructing a deterministic strategy in the full game from the solution to the abstract game; we refer to this procedure as purification. We show that purification, and its generalization which we call thresholding, lead to significantly stronger play than the standard approach in a wide variety of experimental domains. One can view these approaches as ways of

achieving robustness against one's own lossy abstraction.

**Agent-Based Container Terminal Optimisation**
*Stephen Cranefield, Roger Jarquin, Guannan Li, Brent Martin, Rainer Unland, Hanno-Felix Wagner, Michael Winikoff, Thomas Young*
(Poster R 24)

Container terminals play a critical role in international shipping and are under pressure to cope with increasing container traffic. The problem of managing container terminals effectively has a number of characteristics that suggest the use of agent technology would be beneficial. This paper describes a joint industry-university project which has explored the applicability of agent technology to the domain of container terminal management.

**Solving Delayed Coordination Problems in MAS**
*Yann-Michaël De Hauwere, Peter Vrancx, Ann Nowé*
(Poster R 25)

Recent research has demonstrated that considering local interactions among agents in specific parts of the state space, is a successful way of simplifying the multi-agent learning process. By taking into account other agents only when a conflict is possible, an agent can significantly reduce the state-action space in which it learns. Current approaches, however, consider only the immediate rewards for detecting conflicts. This restriction is not suitable for realistic systems, where rewards can be delayed and often conflicts between agents become apparent only several time-steps after an action has been taken. In this paper, we contribute a reinforcement learning algorithm that learns where a strategic interaction among agents is needed, several time-steps before the conflict is reflected by the (immediate) reward signal.

**Human-like Memory Retrieval Mechanisms for Social Companions**
*Mei Yii Lim, Ruth Aylett, Patricia A. Vargas, Wan Ching Ho, João Dias*
(Poster R 26)

This paper demonstrates a biologically- and psychologically- inspired human – like computational memory focusing on the retrieval mechanisms – Spreading Activation and Compound Cue for a companion agent's episodic memory (EM) that helps the agent to manage it's memory more efficiently and enable it to have a more natural interaction with the user.

**Forgetting Through Generalisation - A Companion with Selective Memory**
*Mei Yii Lim, Ruth Aylett, Patricia A. Vargas, Sibylle Enz, Wan Ching Ho*
(Poster R 27)

This research investigates event generalisation in computational episodic memory for artificial companions. Two studies indicated a preference of a biologically-inspired selective memory over an absolute memory companion. Consequently, we present a preliminary implementation of a forgetting mechanism that enables the companion to create "generalised event representations" from its experiences allowing the companion to learn from past encounters.

**Representation of Coalitional Games with Algebraic Decision Diagrams**
*Karthik .V. Aadithya, Tomasz P. Michalak, Nicholas R. Jennings*
(Poster R 28)

With the advent of algorithmic coalitional game theory, it is important to design coalitional game representation schemes that are both compact and efficient with respect to solution concept computation. To this end, we propose a new representation for coalitional games, which is based on Algebraic Decision Diagrams (ADDs). Our representation is fully expressive, compact for many games of practical interest, and enables polynomial time Banzhaf Index, Shapley Value and core computation.

**Game Theoretical Adaptation Model for Intrusion Detection System**
*Martin Rehak, Michal Pěchouček, Martin Grill, Jan Stiborek, Karel Bartos*
(Poster R 29)

We present a self-adaptation mechanism for Network Intrusion Detection System which uses a game-theoretical mechanism to increase system robustness against targeted attacks on IDS adaptation. We model the adaptation process as a strategy selection in sequence of single stage, two player games. The key innovation of our approach is a secure runtime game definition and numerical solution and real-time use of game solutions for dynamic system reconfiguration. Our approach is suited for realistic environ-

ments where we typically lack any ground truth information regarding traffic legitimacy/maliciousness and where the significant portion of system inputs may be shaped by the attacker in order to render the system ineffective. Therefore, we rely on the concept of challenge insertion: we inject a small sample of simulated attacks into the unknown traffic and use the system response to these attacks to define the game structure and utility functions. This approach is also advantageous from the security perspective, as the manipulation of the adaptive process by the attacker is far more difficult. Our experimental results suggest that the use of game-theoretical mechanism comes with little or no penalty when compared to traditional self-adaptation methods.

**Solving Strategic Bargaining with Arbitrary One-Sided Uncertainty**
*Sofia Ceppi, Nicola Gatti, Claudio Iuliano*
(Poster R 30)

Bilateral bargaining has received a lot of attention in the multi-agent literature and has been studied with different approaches. According to the strategic approach, bargaining is modeled as a non-cooperative game with uncertain information and infinite actions. Its resolution is a long-standing open problem and no algorithm addressing uncertainty over multiple parameters is known. In this paper, we provide an algorithm to solve bargaining with any kind of one-sided uncertainty. Our algorithm reduces a bargaining problem to a finite game, solves this last game, and then maps its strategies with the original continuous game. We prove that with multiple types the problem is hard and only small settings can be solved in exact way. In the other cases, we need to resort to concepts of approximate equilibrium and to abstractions for reducing the size of the game tree.

**Manipulation in Group Argument Evaluation**
*Martin Caminada, Gabriella Pigozzi, Mikołaj Podlaszewski*
(Poster R 31)

Given an argumentation framework and a group of agents, the individuals may have divergent opinions on the status of the arguments. If the group needs to reach a common position on the argumentation framework, the question is how the individual evaluations can be mapped into a collective one. This problem has been recently investigated. In this paper, we study under which conditions these operators are Pareto optimal and whether they are manipulable.

**Abstraction for Model Checking Modular Interpreted Systems over ATL**
*Michael Köster, Peter Lohmann*
(Poster R 32)

We propose an abstraction technique for model checking multi-agent systems given as modular interpreted systems (MIS) which allow for succinct representations of compositional systems. Specifications are given as arbitrary ATL formulae, i.e., we can reason about strategic abilities of groups of agents. Our technique is based on collapsing each agent's local state space with hand-crafted equivalence relations, one per strategic modality. We develop a model checking algorithm and prove its soundness. This makes it possible to perform model checking on abstractions (which are much smaller in size) rather than on the concrete system which is usually too complex, thereby saving space and time.

**VIXEE an Innovative Communication Infrastructure for Virtual Institutions**
*Tomas Trescak, Marc Esteva, Inmaculada Rodriguez*
(Poster R 33)

Virtual Institutions (VI) provide many interesting possibilities for social virtual environments, collaborative spaces and simulation environments. VIs combine Electronic Institutions and 3D Virtual Worlds. While Electronic Institutions are used to establish the regulations which structure participants interactions, Virtual Worlds are used to facilitate human participation. In this paper we propose Virtual Institution Execution Environment (VIXEE) as an innovative communication infrastructure for Virtual Institutions. Main features of the infrastructure are i) the causal connection between Virtual World and Electronic Institutions layers, ii) the automatic generation and update of VIs 3D visualization and iii) the simultaneous participation of users from different Virtual World platforms.

**Smart Walkers! Enhancing the Mobility of the Elderly**
*Mathieu Sinn, Pascal Poupart*
(Poster R 34)

The idea of Smart Walkers is to equip customary rolling walkers with sensors in order to assist users, caregivers and clinicians. The integral part of the

Smart Walkers is an autonomous agent which monitors the activity of the user, assesses his physical conditions, and detects potential risks of falls. In this paper, we study methods which enable the agent to recognize the user activity from the sensor measurements. The proposed methods use Conditional Random Fields with features based on discriminant rules. A special case are features which, in order to distinguish between two activities, compare the sensor measurements to thresholds learned by a linear classifier. Experiments with real user data show that the methods achieve a good accuracy; the best results are obtained using "smooth" thresholds based on sigmoid functions.

**Modeling Empathy for a Virtual Human: How, When and to What Extent?**
*Hana Boukricha, Ipke Wachsmuth*
(Poster R 35)

Going along the questions of how, when and to what extent does empathy arise in humans, we propose an approach to model empathy for EMMA - an Empathic MultiModal Agent - based on three processing steps: First, the Empathy Mechanism by which an empathic emotion is produced. Second, the Empathy Modulation by which the empathic emotion is modulated. Third, the Expression of Empathy by which EMMA's modulated empathic emotion is expressed through her multiple modalities. The proposed model is integrated in a conversational agent scenario involving the virtual humans MAX and EMMA.

**Multi-Agent Abductive Reasoning with Confidentiality**
*Jiefei Ma, Alessandra Russo, Krysia Broda, Emil Lupu*
(Poster R 36)

In the context of multi-agent hypothetical reasoning, agents typically have partial knowledge about their environments, and the union of such knowledge is still incomplete to represent the whole world. Thus, given a global query they need to collaborate with each other to make correct inferences and hypothesis, whilst maintaining global constraints. There are many real world applications in which the confidentiality of agent knowledge is of primary concern, and hence the agents may not share or communicate all their information during the collaboration. This extra constraint gives a new challenge to multi-agent reasoning. This paper shows how this dichotomy between "open communication" in collaborative reasoning and protection of confidentiality can be accommodated, by extending a general-purpose distributed abductive logic programming system for multi-agent hypothetical reasoning with confidentiality. Specifically, the system computes consistent conditional answers for a query over a set of distributed normal logic programs with possibly unbound domains and arithmetic constraints, preserving the private information within the logic programs.

**Reasoning About Preferences in BDI Agent Systems**
*Simeon Visser, John Thangarajah, James Harland*
(Poster R 37)

BDI agents often have to make decisions about which plan is used to achieve a goal, and in which order goals are to be achieved. In this paper we describe how to incorporate preferences (based on the *LPP* language) into the BDI execution model.

## Blue Session

*Room: 201*
*18:00 – 19:00*

**Probabilistic Hierarchical Planning over MDPs**
*Yuqing Tang, Felipe Meneguzzi, Katia Sycara, Simon Parsons*
(Poster B 1)

In this paper, we propose a new approach to using probabilistic hierarchical task networks (HTNs) as an effective method for agents to plan in conditions in which their problem-solving knowledge is uncertain, and the environment is non-deterministic. In such situations it is natural to model the environment as a Markov decision process (MDP). We show that using Earley graphs, it is possible to bridge the gap between HTNs and MDPs. We prove that the size of the Earley graph created for given HTNs is bounded by the total number of tasks in the HTNs and show that from the Earley graph we can then construct a plan for a given task that has the maximum expected value when it is executed in an MDP environment.

**Can Trust Increase the Efficiency of Cake Cutting Algorithms?**
*Roie Zivan*
(Poster B 2)
Fair division methods offer guarantees to agents of the proportional size or quality of their share in a division

of a resource (cake). These guarantees come with a price. Standard fair division methods (or "cake cutting" algorithms) do not find efficient allocations (not Pareto optimal). The lack of efficiency of these methods makes them less attractive for solving multi-agent resource and task allocation. Previous attempts to increase the efficiency of cake cutting algorithms for two agents resulted in asymmetric methods that were limited in their ability to find allocations in which both agents receive more than their proportional share.

Trust can be the foundation on which agents exchange information and enable the exploration of allocations that are beneficial for both sides. On the other hand, the willingness of agents to put themselves in a vulnerable position due to their trust in others, results in loss of the fairness guarantees that motivate the design of fair division methods.

In this work we extend the study on fair and efficient cake cutting algorithms by proposing a new notion of trust-based efficiency, which formulates a relation between the level of trust between agents and the efficiency of the allocation. Furthermore, we propose a method for finding trust-based efficiency. The proposed method offers a balance between the guarantees that fair division methods offer to agents and the efficiency that can be achieved by exposing themselves to the actions of other agents. When the level of trust is the highest, the allocation produced by the method is globally optimal (social welfare).

**Decentralized Decision Support for an Agent Population in Dynamic and Uncertain Domains**

*Pradeep Varakantham, Shih-Fen Cheng, Nguyen Thi Duong*

(Poster B 3)

This research is motivated by problems in urban transportation and labor mobility, where the agent flow is dynamic, non-deterministic and on a large scale. In such domains, even though the individual agents do not have an identity of their own and do not explicitly impact other agents, they have implicit interactions with other agents. While there has been much research in handling such implicit effects, it has primarily assumed controlled movements of agents in static environments. We address the issue of decision support for individual agents having involuntary movements in dynamic environments. For instance, in a taxi fleet serving a city: (i) Movements of a taxi are uncontrolled when it is hired by a customer. (ii) Depending on movements of other taxis in the fleet, the environment and hence the movement model for the current taxi changes. Towards addressing this problem, we make three key contributions: (a) A framework to represent the decision problem for individuals in a dynamic population, where there is uncertainty in movements; (b) A novel heuristic technique called Iterative Sampled OPtimization (ISOP) and greedy heuristics to solve large scale problems in domains of interest; and (c) Analyze the solutions provided by our techniques on problems inspired from a real world data set of a taxi fleet operator in Singapore. As shown in the experimental results, our techniques are able to provide strategies that outperform "driver" strategies with respect to: (i) overall availability of taxis; and (ii) the revenue obtained by the taxi drivers.

**Adaptive Decision Support for Structured Organizations: A Case for OrgPOMDPs**

*Pradeep Varakantham, Nathan Schurr, Alan Carlin, Christopher Amato*

(Poster B 4)

In today's world, organizations are faced with increasingly large and complex problems that require decision-making under uncertainty. Current methods for optimizing such decisions fall short of handling the problem scale and time constraints. We argue that this is due to existing methods not exploiting the inherent structure of the organizations which solve these problems. We propose a new model called the OrgPOMDP (Organizational POMDP), which is based on the partially observable Markov decision process (POMDP). This new model combines two powerful representations for modeling large scale problems: hierarchical modeling and factored representations. In this paper we make three key contributions: (a) Introduce the OrgPOMDP model; (b) Present an algorithm to solve OrgPOMDP problems efficiently; and (c) Apply OrgPOMDPs to scenarios in an existing large organization, the Air and Space Operation Center (AOC). We conduct experiments and show that our OrgPOMDP approach results in greater scalability and greatly reduced runtime. In fact, as the size of the problem increases, we soon reach a point at which the OrgPOMDP approach continues to provide solutions while traditional POMDP methods cannot. We also provide an empirical evaluation to highlight the benefits of an organization implementing an OrgPOMDP policy.

**iCLUB: An Integrated Clustering-Based Approach to Improve the Robustness of Reputation Systems**
*Siyuan Liu, Jie Zhang, Chunyan Miao, Yin-Leng Theng, Alex C. Kot*
(Poster B 5)

The problem of unfair testimonies has to be addressed effectively to improve the robustness of reputation systems. We propose an integrated **CLU**stering-**B**ased approach called **iCLUB** to filter unfair testimonies for reputation systems using multi-nominal testimonies, in multiagent-based electronic commerce. It adopts clustering and considers buying agents' local and global knowledge about selling agents. Experimental evaluation demonstrates promising results of our approach in filtering various types of unfair testimonies.

**Effective Variants of Max-Sum Algorithm to Radar Coordination and Scheduling**
*Yoonheui Kim, Michael Krainin, Victor Lesser*
(Poster B 6)

This work proposes new techniques for saving communication and computational resources when solving distributed constraint optimization problems in an environment where system hardware resources are clustered. Using a pre-computed policy and two phase propagation on Max-Sum algorithm, the system performance on Radar scheduling problem improves in terms of communication and computation.

**Improved Computational Models of Human Behavior in Security Games**
*Rong Yang, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe, Richard John*
(Poster B 7)

It becomes critical to address human adversaries' bounded rationality in security games as the real-world deployment of such games spreads. To that end, the key contributions of this paper include: (i) new efficient algorithms for computing optimal strategic solutions using Prospect Theory and Quantal Response Equilibrium; (ii) the most comprehensive experiment to date studying the effectiveness of different models against human subjects for security games. Our new techniques outperform the leading contender for modelling human behavior in security games in experiment with human subjects.

**Agent-Based Resource Allocation in Dynamically Formed CubeSat Constellations**
*Chris HolmesParker, Adrian Agogino*
(Poster B 8)

In the near future, there is potential for a tremendous expansion in the number of Earth-orbiting CubeSats, due to reduced cost associated with platform standardization, availability of standardized parts for CubeSats, and reduced launching costs due to improved packaging methods and lower cost launchers. However, software algorithms capable of efficiently coordinating CubeSats have not kept up with their hardware gains, making it likely that these CubSats will be severely underutilized. Fortunately, these coordination issues can be addressed with multiagent algorithms. In this paper, we show how a multiagent system can be used to address the particular problem of how a third party should bid for use of existing Earth-observing CubeSats so that it can achieve optical coverage over a key geographic region of interest. In this model, an agent is assigned to every CubeSat from which observations may be purchased, and agents must decide how much to offer for these services. We address this problem by having agents use reinforcement learning algorithms with agent-specific shaped rewards. The results show an eight fold improvement over a simple strawman allocation algorithm and a two fold improvement over a multiagent system using standard reward functions.

**A Simple Curious Agent to Help People be Curious**
*Han Yu, Zhiqi Shen, Chunyan Miao, Ah-Hwee Tan*
(Poster B 9)

Curiosity is an innately rewarding state of mind that, over the millennia, has driven the human race to explore and discover. Many researches in pedagogical science have confirmed the importance of being curious to the students' cognitive development. However, in the newly popular virtual world-based learning environments (VLEs), there is currently a lack of attention being paid to enhancing the learning experience by stimulating the learners' curiosity. In this paper, we propose a simple model for curious agents (CAs) which can be used to stimulate learners' curiosity in VLEs. Potential future research directions will be discussed.

**Social Instruments for Convention Emergence**
*Daniel Villatoro, Jordi Sabater-Mir, Sandip Sen*
(Poster B 10)

In this paper we present the notion of Social Instruments as a set of mechanisms that facilitate the emergence of norms from repeated interactions between members of a society. Specifically, we focus on two social instruments: rewiring and observation. Our main goal is to provide agents with tools that allow them to leverage their social network of interactions when effectively addressing coordination and learning problems, paying special attention to dissolving metastable subconventions. Finally, we present a more sophisticated social instrument (observation + rewiring) for robust resolution of subconventions, which works dissolving SelfReinforcing Substructures (SRS) in the social network.

**Learning By Demonstration in Repeated Stochastic Games**
*Jacob W. Crandall, Malek H. Altakrori, Yomna M. Hassan*
(Poster B 11)

Despite much research in recent years, newly created multiagent learning (MAL) algorithms continue to have one or more fatal weaknesses. These weaknesses include slow learning rates, failure to learn nonmyopic solutions, and inability to scale up to domains with many actions, states, and associates. To overcome these weaknesses, we argue that fundamentally different approaches to MAL should be developed. One possibility is to develop methods that allow people to teach learning agents. To begin to determine the usefulness of this approach, we explore the effectiveness of learning by demonstration (LbD) in repeated stochastic games.

**Maximizing Revenue in Symmetric Resource Allocation Systems When User Utilities Exhibit Diminishing Returns**
*Roie Zivan, Miroslav Dudík, Praveen Paruchuri, Katia Sycara*
(Poster B 12)

Consumers of resources in realistic applications (e.g., web, multimedia) typically derive diminishing-return utilities from the amount of resource they receive. A resource provider who is deriving an equal amount of revenue from each satisfied user (e.g., by online advertising), can maximize the number of users by identifying a satisfaction threshold for each user, i.e., the minimal amount of resource the user requires in order to use the service (rather than drop out). A straightforward approach is to ask users to submit their minimal demands (direct revelation). Unfortunately, self-interested users may try to manipulate the system by submitting untruthful requirements.

We propose an incentive-compatible mechanism for maximizing revenue in a resource allocation system where users are ex-ante symmetric (same amount of revenue for any satisfied user) and have diminishing-return utility functions. Users are encouraged by the mechanism to submit their true requirements and the system aims to satisfy as many users as possible. Unlike previous solutions, our mechanism does not require monetary payments from users or downgrading of service. Our mechanism satisfies the number of users within a constant factor of the optimum. Our empirical evaluation demonstrates that in practice, our mechanism can be significantly closer to the optimum than implied by the worst-case analysis. Our mechanism can be generalized to settings when revenue from each user can differ. Also, under some assumptions and adjustments, our mechanism can be used to allocate resource periodically over time.

**Collaborative Diagnosis of Exceptions to Contracts**
*Özgür Kafalı, Francesca Toni, Paolo Torroni*
(Poster B 13)

Exceptions constitute a great deal of autonomous process execution. In order to resolve an exception, several participants should collaborate and exchange knowledge. We believe that argumentation technologies lend themselves very well to be used in this context, both for elaborating on possible causes of exceptions, and for exchanging the result of such elaboration. We propose an open and modular multi-agent framework for handling exceptions using agent dialogues and assumption-based argumentation as the underlying logic.

**Genetic Algorithm Aided Optimization of Hierarchical Multiagent System Organization**
*Ling Yu, Zhiqi Shen, Chunyan Miao, Victor Lesser*
(Poster B 14)

In this paper, we propose a genetic algorithm aided optimization scheme for designing the organization of hierarchical multiagent systems. We introduce the hierarchical genetic algorithm, in which hierarchical crossover with a repair strategy and mutation

of small perturbation are used. The phenotypic hierarchical structure space is translated to the genome-like array representation space, which makes the algorithm genetic-operator-literate. Our experiments show that competitive structures can be found by the proposed algorithm. Compared with traditional operators, the new operators produced better organizations of higher utility more consistently. The proposed algorithm extends the search processes of the state-of-the-art multiagent organization design methodologies, and is more computationally efficient in a large search space.

**Complexity of Multiagent BDI Logics with Restricted Modal Context**
*Marcin Dziubiński*
(Poster B 15)

In this paper we present and discuss a novel language restriction for modal logics for multiagent systems that can reduce the complexity of the satisfiability problem from EXPTIME-hard to NPTIME-complete. In the discussion we focus on a particular BDI logic, called TeamLog, which is a logic for modelling cooperating groups of agents and which possesses some of the characteristics typical to other BDI logics.

**Extension of MC-net-based Coalition Structure Generation: Handling Negative Rules and Externalities**
*Ryo Ichimura, Takato Hasegawa, Suguru Ueda, Atsushi Iwasaki, Makoto Yokoo*
(Poster B 16)

Forming effective coalitions is a major research challenge in AI and multi-agent systems. A Coalition Structure Generation (CSG) problem involves partitioning a set of agents into coalitions so that the social surplus is maximized. Ohta *et al.* introduce an innovative direction for solving CSG, i.e., by representing a characteristic function as a set of rules, a CSG problem can be formalized as the problem of finding a subset of rules that maximizes the sum of rule values under certain constraints. This paper considers two significant extensions of the formalization/algorithm of Ohta *et al.*, i.e., (i) handling negative value rules and (ii) handling externalities among coalitions.

**Diagnosing Commitments: Delegation Revisited**
*Özgür Kafalı, Paolo Torroni*
(Poster B 17)

The success of contract-based multiagent systems relies on agents complying with their commitments. When something goes wrong, the key to diagnosis lies within the commitments' mutual relations as well as their individual states. Accordingly, we explore how commitments are related through the three-agent commitment delegation operation. We then propose exception diagnosis based on such a relation.

**ADAPT: Abstraction Hierarchies to Succinctly Model Teamwork**
*Meirav Hadad, Avi Rosenfeld*
(Poster B 18)

In this paper we present a lightweight teamwork implementation through use of abstraction hierarchies. The basis of this implementation is ADAPT, which supports **A**utonomous **D**ynamic **A**gent **P**lanning for **T**eamwork. ADAPT's novelty stems from how it succinctly decomposes teamwork problems into two separate planners: a **task** network for the set of activities to be performed by a specific agent and a separate **group** network for addressing team organization factors. Because abstract search techniques are the basis for creating these two components, ADAPT agents are able to effectively address teamwork in dynamic environments without explicitly enumerating the entire set of possible team states. During run-time, ADAPT agents then expand the teamwork states that are necessary for task completion through an association algorithm to dynamically link its task and group planners. As a result, ADAPT uses far fewer team states than existing teamwork models. We describe how ADAPT was implemented within a commercial training and simulation application, and present evidence detailing its success in concisely and effectively modeling teamwork.

**Rip-off: Playing the Cooperative Negotiation Game**
*Yoram Bachrach, Pushmeet Kohli, Thore Graepel*
(Poster B 19)

We propose "Rip-off", a new multi-player bargaining game based on the well-studied weighted voting game (WVG) model from cooperative game theory. Many different solution concepts, such as the Core and the Shapley value have been proposed to analyze models

such as WVGs. However, there is little work on analyzing how humans actually play in such settings. We conducted several experiments where we let humans play "Rip-off". Our analysis reveals that although solutions of games played by humans do suffer from certain biases, a player's average payoff over several games is roughly reflected by the Shapley value.

**Interfacing a Cognitive Agent Platform with a Virtual World: a Case Study using Second Life**
*Surangika Ranathunga, Stephen Cranefield, Martin Purvis*
(Poster B 20)

Online virtual worlds provide a rich platform for remote human interaction, and are increasingly being used as a simulation platform for multi-agent systems and as a way for software agents to interact with humans. It would therefore be beneficial to provide techniques allowing high-level agent development tools, especially cognitive agent platforms such as belief-desire-intention (BDI) programming frameworks, to be interfaced with virtual worlds. This is not a trivial task as it involves mapping potentially unreliable sensor readings from complex virtual environments to a domain-specific abstract logical model of observed properties and/or events. This paper investigates this problem in the context of agent interactions in a multi-agent system simulated in Second Life. We present a framework which facilitates the connection of any multi-agent platform with Second Life, and demonstrate it in conjunction with the Jason BDI interpreter.

**Message-Generated Kripke Semantics**
*Jan van Eijck, Floor Sietsma*
(Poster B 21)

We show how to generate multi-agent Kripke models from message exchanges. With these models we can analyze the epistemic consequences of a message exchange. One novelty in this approach is that we include the messages in our logical language. This allows us to model messages that mention other messages and agents that reason about messages. Our framework can be used to model a wide range of different communication scenarios.

**Substantiating Quality Goals with Field Data for Socially-Oriented Requirements Engineering**
*Sonja Pedell, Tim Miller, Leon Sterling, Frank Vetere, Steve Howard, Jeni Paay*
(Poster B 22)

We propose a method for using ethnographic field data to substantiate agent-based models for socially-oriented systems. We investigate in-situ use of domestic technologies created to encourage fun between grandparents and grandchildren separated by distance. The field data added an understanding of what *intergenerational fun* means when imbued with concrete activities. Our contribution is twofold. First, we extend the understanding of agent-oriented concepts by applying them to household interactions. Second, we establish a new method for informing quality goals with field data to enable development of novel applications in the domestic domain.

**Normative Programs and Normative Mechanism Design**
*Nils Bulling, Mehdi Dastani*
(Poster B 23)

The environment is an essential component of multi-agent systems, which is often used to coordinate the behaviour of individual agents. Recently many programming languages have been proposed to facilitate the implementation of such environments. This extended abstract is motivated by the emerging programming languages that are designed to implement environments in terms of normative concepts such as norms and sanctions. We propose a formal analysis of normative environment programs from a mechanism design perspective. By doing this we aim at relating normative environment programs to mechanism design, setting the stage for studying formal properties of these programs such as whether a set of norms implements a specific social choice function in a specific equilibria.

**Privacy-Intimacy Tradeoff in Self-disclosure**
*Jose M. Such, Agustin Espinosa, Ana García-Fornes, Carles Sierra*
(Poster B 24)

In this paper, we introduce a self-disclosure decision-making mechanism based on information-theoretic measures. This decision-making mechanism uses an intimacy measure between agents and the privacy loss that a particular disclosure may cause.

**Reasoning About Norm Compliance**
*Natalia Criado, Estefania Argente, Vicente Botti, Pablo Noriega*
(Poster B 25)

This paper proposes a reasoning process to allow agents to decide when and how norms should be violated or obeyed. The coherence-based reasoning mechanism proposed in this paper, allows *norm aware* agents to confront the norm compliance dilemma and build alternatives for such normative decisions.

**Emergence of Norms for Social Efficiency in Partially Iterative Non-Coordinated Games**
*Toshiharu Sugawara*
(Poster B 26)
We discuss the emergence of social norms for efficiently resolving conflict situations through reinforcement learning and investigate the features of the emergent norms, where conflict situations can be expressed by non-cooperative payoff matrix and will remain if they fail to resolve the conflicts.

**On the Construction of Joint Plans through Argumentation Schemes**
*Oscar Sapena, Alejandro Torreño, Eva Onaindia*
(Poster B 27)

The term Multi-Agent Planning (MAP) refers to any kind of planning in domains in which several independent agents plan and act together. In this paper, we present a novel argumentation-based approach for multiple agents that plan cooperatively while having different abilities, different (and possibly conflicting) views of the world, and different rationalities. The argumentation-based partial-order planning model allows agents to solve MAP problems by proposing partial solutions, giving out opinions on the adequacy of these proposals and modifying them to the benefit of the overall process. We propose the use of a presumptive argumentation model based on the instantiation of argument schemes and associated critical questions to a MAP context.

**Team Coverage Games**
*Yoram Bachrach, Pushmeet Kohli, Vladimir Kolmogorov*
(Poster B 28)

*Team Coverage Games* (TCGs) are a representation of cooperative games, where the value a coalition generates depends on both individual contributions of its members and synergies between them. The synergies are expressed in terms of the importance of the agents in various teams. TCGs model the synergy as a reduction in utility that occurs when team members are missing, causing the team not to achieve its full potential. We focus on the case where the utility reduction incured is a concave function of the importance of the missing team members and analyze the domain from a computational game theoretic perspective.

**Agent-based Inter-Company Transport Optimization**
*Klaus Dorer, Ingo Schindler, Dominic Greenwood*
(Poster B 29)

In previous work we and other authors have shown that agent-based systems are successful in optimizing delivery plans of single logistics companies and are meanwhile successfully productive in industry. In this paper we show that agent-based systems are particularly useful to also optimize transport across logistics companies. In inter-company optimization, privacy is of major importance between the otherwise competing companies. Some data has to be treated strictly private like the cost model or the constraint model. Other data like order information has to be shared. However, typically the amount of orders released to other companies has also to be limited. We show that our agent-based approach can be easily fine tuned to trade off privacy against the benefit of cooperation.

**Belief/Goal Sharing BDI Modules**
*Michal Cap, Mehdi Dastani, Maaike Harbers*
(Poster B 30)

This paper proposes a modularisation framework for BDI based agent programming languages developed from a software engineering perspective. Like other proposals, BDI modules are seen as encapsulations of cognitive components. However, unlike other approaches, modules are here instantiated and manipulated in a similar fashion as objects in object orientation. In particular, an agent's mental state is formed dynamically by instantiating and activating BDI modules. The agent deliberates on its active module instances, which interact by sharing their beliefs and goals.

**Neural Symbolic Architecture for Normative Agents**
*Guido Boella, Silvano Colombo Tosatto, Artur d'Avila Garcez, Valerio Genovese, Dino Ienco, Leendert van der Torre*
(Poster B 31)

In this paper we propose a neural-symbolic architecture to represent and reason with norms in multi-agent systems. On the one hand, the architecture contains a symbolic knowledge base to represent norms and on the other hand it contains a neural network to reason with norms. The interaction between the symbolic knowledge and the neural network is used to learn norms. We describe how to handle normative reasoning issues like contrary to duties, dilemmas and exceptions by using a priority-based ordering between the norms in a neural-symbolic architecture.

**No Smoking Here: Compliance Differences Between Legal and Social Norms**
*Francien Dechesne, Virginia Dignum*
(Poster B 32)

The values shared within a society influence the (social) behaviour of the agents in that society. This connection goes through implicit and explicit norms. In the case of a norm conflict, an agent will decide to comply with one or more of the applicable norms, while violating others. Our interest is how the *type* of the norms may play a role in such decision, and take the chosen behaviour of an agent to depend on a personal preference order on the norm types. We distinguish three different types of norms: legal norms, social norms and private norms. We illustrate our approach using the introduction of the law prohibiting smoking in cafes: we present a simulation of this situation involving agents' preferences over different norm types. The results of this simulation are used for an explorative model for normative reasoning based on norm types. We discuss a possible connection between the composition of a society in terms of these profiles and its culture and the relevance of the model with respect to *value sensitive design* of socio-technological systems.

**Agents That Speak: Modelling Communicative Plans and Information Sources in a Logic of Announcements**
*Philippe Balbiani, Nadine Guiraud, Andreas Herzig, Emiliano Lorini*
(Poster B 33)

We present a modal logic of belief and announcements in a multi-agent setting. This logic allows to express not only that $\psi$ holds after the announcement of $\varphi$ as in standard public announcement logic (PAL), but also that the announcement of $\varphi$ occurs. We use the logic to provide a formal analysis of several concepts that are relevant for multi-agent systems (MAS) theory and applications: the notions of communicative action (an agent informs another agent about something) and communicative intention (an agent has the intention to inform another agent about something), and the notion of information source.

**Procedural Fairness in Stable Marriage Problems**
*Mirco Gelain, Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, Toby Walsh*
(Poster B 34)

The stable marriage problem is a well-known problem of matching men to women so that no man and woman, who are not married to each other, both prefer each other. It has a wide variety of practical applications, ranging from matching resident doctors to hospitals, to matching students to schools, or more generally to any two-sided market. Given a stable marriage problem, it is possible to find a male-optimal (resp., female-optimal) stable marriage in polynomial time. However, it is sometimes desirable to find stable marriages without favoring one group at the expenses of the other one. To achieve this goal, we consider a local search approach to find stable marriages with the aim of exploiting the nondeterminism of local search to give a fair procedure. We test our algorithm on classes of stable marriage problems, showing both its efficiency and its sampling capability over the set of all stable marriages, and we compare it to a Markov chain approach.

**Tag-Based Cooperation in N-Player Dilemmas**
*Enda Howley, Jim Duggan*
(Poster B 35)

This paper studies the emergence of cooperation in the N-Player Prisoner's Dilemma (NPD) using a tag-mediated interaction model. Tags have been widely used to bias agent pairwise interactions which facilitates the emergence of cooperation. This paper shows some of the key parameters that influence the emergence of cooperation in an evolutionary setting. The aim of this paper is to demonstrate the most vital factors that are commonly ignored in many existing NPD studies.

**Heuristic Multiagent Planning with Self-Interested Agents**
*Matt Crosby, Michael Rovatsos*
(Poster B 36)

The focus of multiagent planning research has recently turned towards domains with self-interested agents leading to the definition of Coalition-Planning Games (CoPGs). In this paper, we investigate algorithms for solving a restricted class of "safe" CoPGs, in which no agent can benefit from making another agent's plan invalid. We introduce a novel, generalised solution concept, and show how problems can be translated so that they can be solved by standard single-agent planners. However, standard planners cannot solve problems like this efficiently. We then introduce a new multiagent planning algorithm and the benefits of our approach are illustrated empirically in an example logistics domain.

**Mining Qualitative Context Models from Multiagent Interactions**
*Emilio Serrano, Michael Rovatsos, Juan Botia*
(Poster B 37)

We present a novel method for analysing the behaviour of multiagent systems on the basis of the semantically rich information provided by agent communication languages and interaction protocols. Contrary to analysis methods that rely on observing more low-level patterns of behaviour, our method is based on exploiting the semantics. These languages and protocols which can be used to extract *qualitative* properties of observed interactions. This can be achieved by interpreting the logical constraints associated with protocol execution paths or individual messages as models of the *context* of an observed interaction, and using them as features of learning samples.

**Partially Observable Stochastic Game-based Multi-Agent Prediction Markets**
*Janyl Jumadinova, Prithviraj Dasgupta*
(Poster B 38)

We present a novel representation of the prediction market using a partially observable stochastic game with information (POSGI), that can be used by each trading agent to precisely calculate the state of the market. We then propose that a correlated equilibrium (CE) strategy can be used by the agents to dynamically calculate the prices at which they should trade securities in the prediction market. Simulation results comparing the CE strategy within our POSGI model with five other strategies commonly used in similar markets show that the CE strategy results in improved price predictions and higher utilities to the agents as compared to other strategies.

## Green Session

*Room: 201*
*19:00 – 20:00*

**A Cost-Based Transition Approach for Multiagent Systems Reorganization**
*Juan M. Alberola, Vicente Julián, Ana García-Fornes*
(Poster G 1)

In this paper we present an organization transition model that is based on costs along with an associated organization transition mechanism. This mechanism calculates how a current instance of an organization can evolve to a future instance and how costly this evolution is.

**Towards an Agent-Based Proxemic Model for Pedestrian and Group Dynamics: Motivations and First Experiments**
*Sara Manzoni, Giuseppe Vizzari, Kazumichi Ohtsuka, Kenichiro Shimura*
(Poster G 2)

This paper introduces the first experiments of an innovative approach to the modeling and simulation of crowds of pedestrians considering the presence of groups as a crucial element influencing overall system dynamics. *In-silico* experimental results are discussed in relation to *in-vitro* experiments (experimental observations on the movement of pedestrians and groups).

**Batch Reservations in Autonomous Intersection Management**
*Neda Shahidi, Tsz-Chiu Au, Peter Stone*
(Poster G 3)

The recent robot car competitions and demonstrations have convincingly shown that fully autonomous vehicles are feasible with current or near-future intelligent vehicle technology. Looking ahead to the time when such autonomous cars will be common, Dresner and Stone proposed a new intersection control protocol called *Autonomous Intersection Management* (AIM) and showed that by leveraging the capacities of autonomous vehicles we can devise a reservation-based intersection control protocol that is much more efficient than traffic signals and stop signs. Their pro-

posed protocol, however, handles reservation requests one at a time and does not prioritize reservations according to their relative importance and vehicles' waiting times, causing potentially large inequalities in granting reservations. For example, at an intersection between a main street and an alley, vehicles from the alley can take a very long time to get reservations to enter the intersection. In this research, we introduce a prioritization scheme to prevent uneven reservation assignments in unbalanced traffic. Our experimental results show that our prioritizing scheme outperforms previous intersection control protocols in unbalanced traffic.

**Multi-Agent, Reward Shaping for RoboCup KeepAway**
*Sam Devlin, Marek Grześ, Daniel Kudenko*
(Poster G 4)

This paper investigates the impact of reward shaping in multi-agent reinforcement learning as a way to incorporate domain knowledge about good strategies. In theory, potential-based reward shaping does not alter the Nash Equilibria of a stochastic game, only the exploration of the shaped agent. We demonstrate empirically the performance of statebased and state-action-based reward shaping in RoboCup KeepAway. The results illustrate that reward shaping can alter both the learning time required to reach a stable joint policy and the final group performance for better or worse.

**Approximating Behavioral Equivalence of Models Using Top-K Policy Paths**
*Yifeng Zeng, Yingke Chen, Prashant Doshi*
(Poster G 5)

Decision making and game play in multiagent settings must often contend with behavioral models of other agents in order to predict their actions. One approach that reduces the complexity of the unconstrained model space is to group models that tend to be behaviorally equivalent. In this paper, we seek to further compress the model space by introducing an approximate measure of behavioral equivalence and using it to group models.

**Reflection about Capabilities for Role Enactment**
*M. Birna van Riemsdijk, Virginia Dignum, Catholijn M. Jonker, Huib Aldewereld*
(Poster G 6)

An organizational modeling language can be used to specify an agent organization in terms of its roles, organizational structure, norms, etc. Using such an organizational specification to organize a multi-agent system should make the agents more effective in attaining their purpose, or prevent certain undesired behavior from occurring. Agents who want to enter and play roles in an organization are expected to understand and reason about the organizational specification. An important aspect that such organization-aware agents should be able to reason about is role enactment. In particular, agents should be able to reflect on whether they have the capabilities to play a role in an organization. In future work it needs to be made precise when an agent can be said to have a certain capability, and how an agent can reflect on its capabilities. This is necessary for programming role enactment in organization-aware agents.

**Prognostic Normative Reasoning in Coalition Planning**
*Jean Oh, Felipe Meneguzzi, Katia Sycara, Timothy J. Norman*
(Poster G 7)

In this paper we describe a software assistant agent that can proactively assist human users situated in a time-constrained coalition environment. The cognitive workload is significantly increased when the user must not only cope with a complex environment, but also with a set of unaccustomed rules that prescribe how the coalition planning process must be carried out. In this context, we introduce the notion of *prognostic norm reasoning* to predict the user's likely normative violations, allowing the assistant agent to plan and take remedial actions before the violations actually occur. To the best of our knowledge, our approach is the first that manages norms in a proactive and autonomous manner.

**Virtual Agent Perception in Large Scale Multi-Agent Based Simulation Systems**
*Dane Kuiper, Rym Z. Wenkstern*
(Poster G 8)

In this paper we discuss virtual agent perception in large scale open environment based MABS.

**A Formal Analysis of the Outcomes of Argumentation-based Negotiations**
*Leila Amgoud, Srdjan Vesic*
(Poster G 9)

This paper tackles the problem of exchanging arguments in negotiation dialogues, and provides first characterizations of the outcomes of such rich dialogues.

**Modeling the Emergence of Norms**
*Logan Brooks, Wayne Iba, Sandip Sen*
(Poster G 10)

Norms or conventions can be used as external correlating signals to promote coordination between rational agents and hence have merited in-depth study of the evolution and economics of norms both in the social sciences and in multi-agent systems. While agent simulations can be used to gain a cursory idea of when and what norms can evolve, the estimations obtained by running simulations can be costly to obtain, provide no guarantees about the behavior of the system, and may overlook some rare occurrences. We use a theoretical approach to analyze a system of agents playing a convergence game and develop models that predict (a) how the system's behavior will change over time, (b) how much time it will take for it to converge to a stable state, and (c) how often the system will converge to a particular norm.

**Introducing Homophily to Improve Semantic Service Search in a Self-adaptive System**
*E. del Val, M. Rebollo, Vicente Botti*
(Poster G 11)

Humans create efficient social structures in a self-organized way. People tend to join groups with other people with similar characteristics. This is called homophily. This pa- per proposes how homophily can be introduced in Service- Oriented Multiagent Systems (SOMAS) to create efficient self-organized structures.

**Adaptive Regulation of Open MAS: an Incentive Mechanism based on Modifications of the Environment**
*Roberto Centeno, Holger Billhardt*
(Poster G 12)

The global objective of open multiagent systems might be in conflict with individual preferences of rational agents participating in such systems. Addressing this problem, we propose a mechanism able to attach incentives to agent actions such that the global utility of the system is improved. Such incentives are dynamically adjusted to each agent's preferences by using institutional agents called *incentivators*.

**Allocating Spatially Distributed Tasks in Large, Dynamic Robot Teams**
*Steven Okamoto, Nathan Brooks, Sean Owens, Katia Sycara, Paul Scerri*
(Poster G 13)

For an interesting class of emerging applications, a large robot team will need to distributedly allocate many more tasks than there are robots, with dynamically appearing tasks and a limited ability to communicate. The LA-DCOP algorithm can conceptually handle both large-scale problems and multiple tasks per robot, but has key limitations when allocating spatially distributed tasks. In this paper, we extend LA-DCOP with several alternative acceptance rules for robots to determine whether to take on an additional task, given the interaction with the tasks it has already committed to. We show that these acceptance rules dramatically outperform a naive LA-DCOP implementation. In addition, we developed a technique that lets the robots use completely local knowledge to adjust their task acceptance criteria to get the best possible performance at a given communication bandwidth level.

**Bounded Optimal Team Coordination with Temporal Constraints and Delay Penalties**
*G. Ayorkor Korsah, Anthony Stentz, M. Bernardine Dias*
(Poster G 14)

We address the problem of optimally assigning spatially distributed tasks to a team of heterogeneous mobile agents in domains with inter-task temporal constraints, such as precedence constraints. Due to delay penalties, satisfying the temporal constraints impacts the overall team cost. We present a mathematical model of the problem, a benchmark anytime bounded optimal solution process, and an analysis of the impact of delay penalties on problem difficulty.

**A Perception Framework for Intelligent Characters in Serious Games**
*Joost van Oijen, Frank Dignum*
(Poster G 15)

In this paper we present a framework for perceptual attention for BDI-agents embodied in a virtual environment. To abstract from different data represen-

tations present in virtual environments, we introduce an ontological information model to define environment semantics. Further, the framework proposes an approach for goal-directed perception based on the specification of different interests, which define selected information from the environment an agent is interested in. These interests filter and extract specific information from a continuous incoming flow of sensory information. In this way we can find a balance between stimulus-driven control and goal based control over perception. Thus preventing an agent to become flooded with irrelevant information on the one hand or blinded for unexpected events that need to be reacted to on the other hand.

**SR-APL: A Model for a Programming Language for Rational BDI Agents with Prioritized Goals**
*Shakil M. Khan, Yves Lespérance*
(Poster G 16)

To provide efficiency, current BDI agent programming languages with declarative goals only support a limited form of rationality – they ignore other concurrent intentions of the agent when selecting plans, and as a consequence, the selected plans may be inconsistent with these intentions. In this paper, we develop logical foundations for a rational BDI agent programming framework with prioritized declarative goals that addresses this deficiency. We ensure that the agent's chosen declarative goals and adopted plans are consistent with each other and with the agent's knowledge. We show how agents specified in our language satisfy some key rationality requirements.

**Designing Petri Net Supervisors for Multi-Agent Systems from LTL Specifications**
*Bruno Lacerda, Pedro U. Lima*
(Poster G 17)

In this paper, we use LTL to specify acceptable/desirable behaviours for a system modelled as a Petri net, and create a Petri net realization of a supervisor that is guaranteed to enforce them, by appropriately restricting the uncontrolled behaviour of the system.We illustrate the method with an application to the specification of coordination requirements between the members of a team of simulated soccer robots.

**Friend or Foe? Detecting an Opponent's Attitude in Normal Form Games**
*Steven Damer, Maria Gini*
(Poster G 18)

We study the problem of achieving cooperation between two self-interested agents that play a sequence of different randomly generated normal form games. The agent learns how much the opponent is willing to cooperate and reciprocates. We present empirical results that show that both agents benefit from cooperation and that a small number of games is sufficient to learn the cooperation level of the opponent.

**The BDI Driver in a Service City**
*Marco Lützenberger, Nils Masuch, Benjamin Hirsch, Sebastian Ahrndt, Axel Heßler, Sahin Albayrak*
(Poster G 19)

Most traffic simulation frameworks move vehicles from some location A to some location B as the result of different equations of motion or fluid dynamics. As it is, reality is much more complex because what actually happens on the road is not only determined by physics of motion, but also by the perception and attitudes of the drivers. In this work, we introduce an approach which considers a driver's state of mind within large scale traffic simulations. For this purpose we describe a BDI based conceptualisation of a driver and extend common simulation topologies with service oriented concepts.

**Identifying and Exploiting Weak-Information Inducing Actions in Solving POMDPs**
*Ekhlas Sonu, Prashant Doshi*
(Poster G 20)

We present a method for identifying actions that lead to observations which are only weakly informative in the context of partially observable Markov decision processes (POMDP). We call such actions as weak-(inclusive of zero-)information inducing. Policy subtrees rooted at these actions may be computed more efficiently. While zero-information inducing actions may be exploited without error, we show that the error due to the quicker backup for weak but non-zero information inducing actions is bounded. We demonstrate the substantial computational savings that exploiting such actions may bring to exact and approximate solutions of POMDPs.

**Teamwork in Distributed POMDPs: Execution-time Coordination Under Model Uncertainty**
*Jun-Young Kwak, Rong Yang, Zhengyu Yin, Matthew E. Taylor, Milind Tambe*
(Poster G 21)

Despite their worst-case NEXP-complete planning complexity, DEC-POMDPs remain a popular framework for multiagent teamwork. This paper introduces effective teamwork under model uncertainty (i.e., potentially inaccurate transition and observation functions) as a novel challenge for DEC-POMDPs and presents MODERN, the first execution-centric framework for DEC-POMDPs explicitly motivated by addressing such model uncertainty. MODERN's shift of coordination reasoning from planning-time to execution-time avoids the high cost of computing optimal plans whose promised quality may not be realized in practice. There are three key ideas in MODERN: (i) it maintains an exponentially smaller model of other agents' beliefs and actions than in previous work and then further reduces the computation-time and space expense of this model via bounded pruning; (ii) it reduces execution-time computation by exploiting BDI theories of teamwork, and limits communication to key trigger points; and (iii) it limits its decision-theoretic reasoning about communication to trigger points and uses a systematic markup to encourage extra communication at these points — thus reducing uncertainty among team members at trigger points.

**Escaping Local Optima in POMDP Planning as Inference**
*Pascal Poupart, Tobias Lang, Marc Toussaint*
(Poster G 22)

Planning as inference recently emerged as a versatile approach to decision-theoretic planning and reinforcement learning for single and multi-agent systems in fully and partially observable domains with discrete and continuous variables. Since planning as inference essentially tackles a non-convex optimization problem when the states are partially observable, there is a need to develop techniques that can robustly escape local optima. We propose two algorithms: the first one adds nodes to the controller according to an increasingly deep forward search, while the second one splits nodes in a greedy fashion to improve reward likelihood.

**Toward Human Interaction with Bio-Inspired Teams**
*Michael A. Goodrich, P. B. Sujit, Jacob W. Crandall*
(Poster G 23)

Although much work has been done on designing autonomy and user interfaces for managing small teams of independent robots, much less is known about managing large-scale bio-inspired robot (BIRT) teams. In this paper, we explore human interaction with BIRT teams in an information foraging task. We summarize results from two small experiments that use two types of BIRT teams in a foraging task. The results illustrate differences in BIRT performance for different types of human interaction, and illustrate how performance robustness can vary as a function of interaction type.

**Escaping Heuristic Depressions in Real-Time Heuristic Search**
*Carlos Hernández, Jorge A. Baier*
(Poster G 24)

Heuristic depressions are local minima of heuristic functions. While visiting one them, real-time (RT) search algorithms like LRTA will update the heuristic value for most of their states several times before escaping, resulting in costly solutions. Existing RT search algorithm tackle this problem by doing more search and/or lookahead but do not guide search towards leaving depressions. We present eLSS-LRTA, a new RT search algorithm based on LSSLRTA that actively guides search towards exiting regions with heuristic depressions. We show that our algorithm produces better quality solutions than LSS-LRTA for equal values of lookahead in standard RT benchmarks.

**Pseudo-tree-based Algorithm for Approximate Distributed Constraint Optimization with Quality Bounds**
*Tenda Okimoto, Yongjoon Joe, Atsushi Iwasaki, Makoto Yokoo*
(Poster G 25)

Most incomplete DCOP algorithms generally do not provide any guarantees on the quality of the solutions. In this paper, we introduce a new incomplete DCOP algorithm that can provide the upper bounds of the absolute/relative errors of the solution, which can be obtained a priori/a posteriori, respectively. The evaluation results illustrate that this algorithm can obtain better quality solutions and bounds compared to ex-

isting bounded incomplete DCOP algorithms, while the run time of this algorithm is much shorter.

**Concise Characteristic Function Representations in Coalitional Games Based on Agent Types**
*Suguru Ueda, Makoto Kitaki, Atsushi Iwasaki, Makoto Yokoo*
(Poster G 26)

Forming effective coalitions is a major research challenge in AI and multi-agent systems. Thus, coalitional games, including coalition structure generation, have been attracting considerable attention from the AI research community. Traditionally, the input of a coalitional game is a black-box function called a characteristic function. In this paper, we develop a new concise representation scheme for a characteristic function, which is based on the idea of agent types. This representation can be exponentially more concise than existing concise representation schemes. Furthermore, this idea can be used in conjunction with existing schemes to further reduce the representation size.

**Iterative Game-theoretic Route Selection for Hostile Area Transit and Patrolling**
*Ondřej Vaněk, Michal Jakob, Viliam Lisý, Branislav Bošanský, Michal Pěchouček*
(Poster G 27)

A number of real-world security scenarios can be cast as a problem of transiting an area patrolled by a mobile adversary, where the transiting agent aims to choose its route so as to minimize the probability of encountering the patrolling agent, and vice versa. We model this problem as a twoplayer zero-sum game on a graph, termed the transit game. In contrast to the existing models of area transit, where one of the players is stationary, we assume both players are mobile. We also explicitly model the limited endurance of the patroller and the notion of a base to which the patroller has to repeatedly return. Noting the prohibitive size of the strategy spaces of both players, we employ iterative oracle-based algorithms including a newly proposed accelerated scheme, to obtain optimum route selection strategies for both players. We evaluate the developed approach on a range of transit game instances inspired by real-world security problems in the urban and naval security domains.

**Abduction Guided Query Relaxation**
*Samy Sá, João Alcântara*
(Poster G 28)

We investigate how to improve cooperative communication between agents by representing knowledge bases as logic programs extended with abduction. In this proposal, agents try to provide explanations whenever they fail to answer a question. Query Relaxation is then employed to search for answers related to the query, characterizing cooperative behavior. Our contributions bring insightful improvements to relaxation attempts and the quality of related answers. We introduce rational explanations and use them to efficiently guide the search for related answers in a relaxation tree.

**A Message Passing Approach To Multiagent Gaussian Inference for Dynamic Processes**
*Stefano Ermon, Carla Gomes, Bart Selman*
(Poster G 29)

In [1], we introduced a novel distributed inference algorithm for the multiagent Gaussian inference problem, based on the framework of graphical models and message passing algorithms. We compare it to current state of the art techniques and we demonstrate that it is the most efficient one in terms of communication resources used. Moreover, we show experimentally that it outperforms the other methods in terms of estimation error on a general class of problems, even in presence of data loss.

**Multiagent Environment Design in Human Computation**
*Chien-Ju Ho, Yen-Ling Kuo, Jane Yung-Jen Hsu*
(Poster G 30)

This research explores human computation systems as a multi-agent platform for coordinating human brain power. We are interested in the problem of designing human computation systems. In particular, we focus on the situation where developers cannot redesign the whole platform but can only make limited changes to the environments. We extend the framework of environment design to multiple agent in the context of human computation. To incorporate the collective information from multiple agents, we propose two approaches, agent type elicitation and collaborative filtering, under different assumptions. The formulation and algorithms provide solutions for developers in human computation systems to find the environment settings maximizing their goal functions.

**Social Distance Games**
*Simina Brânzei, Kate Larson*
(Poster G 31)

In this paper we introduce and analyze social distance games, a family of non-transferable utility coalitional games where an agent's utility is a measure of closeness to the other members of the coalition. We study both social welfare maximisation and stability in these games from a graph theoretic perspective. We investigate the welfare of stable coalition structures, and propose two new solution concepts with improved welfare guarantees. We argue that social distance games are both interesting in themselves, as well as in the context of social networks.

**Agent Sensing with Stateful Resources**
*Adam Eck, Leen-Kiat Soh*
(Poster G 32)

In the application of multi-agent systems to real-world problems, agents often suffer from bounded rationality where agent reasoning is limited by 1) a lack of knowledge about choices, and 2) a lack of resources required for reasoning. To overcome the former, the agent uses sensing to refine its knowledge. However, sensing can also require limited resources, leading to inaccurate environment modeling and poor decision making. In this paper, we consider a novel and difficult class of this problem where agents must use stateful resources during sensing, which we define as resources whose state-dependent behavior changes over time based on usage. Specifically, such sensing changes the state of a resource, and thus its behavior, producing a phenomenon where the sensing activity can and will distort its own outcome. We term this the Observer Effect after the similar phenomenon in the physical sciences. Given this effect, the agent faces a strategic tradeoff between satisfying the need for 1) knowledge refinement, and 2) avoiding corruption of knowledge due to distorted sensing outcomes. To address this tradeoff, we use active perception to select sensing activities and model activity selection as a Markov decision process (MDP) solved through reinforcement learning where an agent optimizes knowledge refinement while considering the state of the resource used during sensing.

**Modeling Bounded Rationality of Agents During Interactions**
*Qing Guo, Piotr Gmytrasiewicz*
(Poster G 33)

In this paper, we propose that bounded rationality of another agent be modeled as errors the agent is making while deciding on its action. We are motivated by the work on quantal response equilibria in behavioral game theory which uses Nash equilibria as the solution concept. In contrast, we use decision-theoretic maximization of expected utility. Quantal response assumes that a decision maker is approximately rational, i.e., is maximizing its expected utility but with an error rate characterized by a single error parameter. Another agent's error rate may be unknown and needs to be estimated during an interaction. We show that this error rate can be estimated using Bayesian update of a suitable conjugate prior, and that it has a sufficient statistic of fixed dimension under strong simplifying assumptions. However, if the simplifying assumptions are relaxed, the quantal response does not admit a finite dimensional sufficient statistic, and a more complex update is needed.

**Comparing Action-Query Strategies in Semi-Autonomous Agents**
*Robert Cohn, Edmund H. Durfee, Satinder Singh*
(Poster G 34)

We consider semi-autonomous agents that have uncertain knowledge about their environment, but can ask what action the operator would prefer taking in the current or in a potential future state. Asking queries can help improve behavior, but if queries come at a cost (e.g., due to limited operator attention), the number of queries needs to be minimized. We develop a new algorithm for selecting action queries by adapting the recently proposed Expected Myopic Gain (EMG) from its prior use in settings with reward or transition probability queries to our setting of action queries, and empirically compare it to the current state of the art.

**A Multimodal End-of-Turn Prediction Model: Learning from Parasocial Consensus Sampling**
*Lixing Huang, Louis-Philippe Morency, Jonathan Gratch*
(Poster G 35)

Virtual human, with realistic behaviors and social skills, evoke in users a range of social behaviors normally only seen in human face-to-face interactions.

One of the key challenges in creating such virtual humans is to give them human-like conversational skills, such as turn-taking skill. In this paper, we propose a multimodal end-of-turn prediction model. Instead of recording face-to-face conversation data, we collect the turn-taking data using Parasocial Consensus Sampling (PCS) framework. Then we analyze the relationship between verbal and nonverbal features and turn-taking behaviors based on the consensus data and show how these features influence the time people use to take turns. Finally, we present a probabilistic multimodal end-ofturn prediction model, which enables virtual humans to make real-time turn-taking predictions. The result shows that our model achieves a higher accuracy than previous methods did.

**Scalable Adaptive Serious Games using Agent Organizations**
*Joost Westra, Frank Dignum, Virginia Dignum*
(Poster G 36)

Serious games and other training applications have the requirement that they should be suitable for trainees with different skill levels. Current approaches either use human experts or a completely centralized approach for this adaptation. These centralized approaches become very impractical and will not scale if the complexity of the game increases. Agents can be used in serious game implementations as a means to reduce complexity and increase believability but without some centralized coordination it becomes practically impossible to follow the intended storyline of the game and select suitable difficulties for the trainee. In this paper we show that using agent organizations to coordinate the agents is scalable and allows adaptation in very complex scenarios while making sure the storyline is preserved the right difficulty level for the trainee is preserved.

**Integrating power and reserve trade in electricity networks**
*Nicolas Höning, Han Noot, Han La Poutré*
(Poster G 37)

In power markets, the trade of reserve energy will become more important as more intermittent generation is traded. In this work, we propose a novel bidding mechanism for the integration of power and reserve markets. It adds expressivity to reserve bids and facilitates planning.

# Demonstrations

**BDI Agent model Based Evacuation Simulation**
*Masaru Okaya, Tomoichi Takahashi*
(Demo 1)

In this paper, we present BDI agent model that social relationships affect agents' evacuation behaviors. The model causes evacuation grouping and breaking in a crowd and the simulation results shows that it generates emergent behaviors in a crowd evacuation as a result of each behavior. Our evacuation simulation provides an analysis tool to assess the safety of occupants in building evacuation.

**An Interactive Tool for Creating Multi-Agent Systems and Interactive Agent-based Games**
*Henrik Hautop Lund, Luigi Pagliarini*
(Demo 2)

Utilizing principles from parallel and distributed processing combined with inspiration from modular robotics, we developed the modular interactive tiles. As an educational tool, the modular interactive tiles facilitate the learning of multi-agent systems and interactive agent-based games. The modular and physical property of the tiles provides students with hands-on experience in exploring the theoretical aspects underlying multi-agent systems which often appear as challenging to students. By changing the representation of the cognitive challenging aspects of multi-agent systems education to a physical (hands-on) one, the challenge may become much easier and fun to face for the students.

**Towards Robot Incremental Learning Constraints from Comparative Demonstration**
*Rong Zhang, Shangfei Wang, Xiaoping Chen, Dong Yin, Shijia Chen, Min Cheng, Yanpeng Lv, Jianmin Ji, Dejian Wang, Peijia Shen*
(Demo 3)

This paper presents an attempt on incremental robot learning from demonstration. Based on previously learnt knowledge about a task in simpler situations, a robot learns to fulfill the same task properly in a more complicated situation through analyzing comparative demonstrations and extracting new knowledge, especially the constraints that the task in the new situation imposes on the robot's behaviors.

**Teleworkbench: Validating Robot Programs from Simulation to Prototyping with Minirobots**
*A. Tanoto, F. Werner, U. Rückert, H. Li*
(Demo 4)

This paper describes a Demo showing the role of the Teleworkbench in the validation process of a multi-agent system, e.g., a traffic management system. In the Demo, we show the capability of the Teleworkbench in seamlessly bridging the simulation and experimentation with real robots. During experiments, important information is logged for analysis purpose. Additionally, a graphical user interface enables geographically distributed users to perform some levels of interactivity, e.g., watch the video or command the robots.

**A MAS Decision Support Tool for Water-Right Markets**
*Adriana Giret, Antonio Garrido, Juan A. Gimeno, Vicente Botti, Pablo Noriega*
(Demo 5)

We present a MAS decision support tool, as an open and regulated virtual organization, that uses intelligent agents to manage a flexible water-right market. The application goal of this tool is to be used as a simulator to assist in decision-taking processes for policy makers. The simulator focuses on demands and, in particular, on the type of regulatory (in terms of norms selection and agents behaviour), and market mechanisms that foster an efficient use of water while also trying to prevent conflicts among parties. Technically, it contributes with a testbed to explore policy-simulation alternatives under an agreement-technology perspective, thus promoting agreements fulfillment.

**An Implementation of Basic Argumentation Components**
*Mikołaj Podlaszewski, Martin Caminada, Gabriella Pigozzi*
(Demo 6)

The current implementation provides a demonstration of a number of basic argumentation components that can be applied in the context of multi-agent systems. These components include algorithms for calculating argumentation semantics, as well as for deter-

mining the justification status of the arguments and providing explanation in the form of formal discussion games. Furthermore, the current demonstrator also includes the first implementation we know of regarding argument-based judgment aggregation theory.

**AgentC: Agent-based System for Securing Maritime Transit**
*Michal Jakob, Ondřej Vaněk, Branislav Bošanský, Ondřej Hrstka, Michal Pěchouček*
(Demo 7)

Recent rise in maritime piracy prompts the search for novel techniques for addressing the problem. We therefore developed AgentC, a prototype system that demonstrates how agent-based traffic management techniques can be used to improve the security of transit through piracy-affected areas. Combining agent-based modeling and simulation of maritime traffic and novel route planning and vessel scheduling techniques, the system shows the promising potential of agent-based methods for increasing maritime security.

**Bee-Inspired Foraging In An Embodied Swarm**
*Sjriek Alers, Daan Bloembergen, Daniel Hennes, Steven de Jong, Michael Kaisers, Nyree Lemmens, Karl Tuyls, Gerhard Weiss*
(Demo 8)

We show the emergence of Swarm Intelligence in physical robots. We transfer an optimization algorithm which is based on beeforaging behavior to a robotic swarm. In simulation this algorithm has already been shown to be more effective, scalable and adaptive than algorithms inspired by ant foraging. In addition to this advantage, bee-inspired foraging does not require (de-)centralized simulation of environmental parameters (e.g. pheromones).

**The Social Ultimatum Game and Adaptive Agents**
*Yu-Han Chang, Rajiv Maheswaran*
(Demo 9)

The Ultimatum Game is a key exemplar that shows how human play often deviates from "rational" strategies suggested by game-theoretic analysis. One explanation is that humans cannot put aside the assumption of being in a multi-player multi-round environment that they are accustomed to in the real world. We introduce the Social Ultimatum Game (SUG), where players can choose their partner among a society of agents, and engage in repeated interactions of the Ultimatum Game. We develop mathematical models of human play that include "irrational" concepts such as fairness and adaptation to the expectations of the society. We will display a system where people can play SUG against a mixed system of other humans and autonomous agents based on our mathematical models.

**DipTools: Experimental Data Visualization Tool for the DipGame Testbed**
*Angela Fabregues, David López-Paz, Carles Sierra*
(Demo 10)

DipGame is a testbed for negotiation. It permits to test negotiation algorithms, even if enriched with argumentation, trust or reputation techniques. It is very appropriate to run experiments that mix humans and agents. In this demonstration we introduce a tool to visualise data obtained from DipGame experiments.

**TALOS: A Tool for Designing Security Applications with Mobile Patrolling Robots**
*Nicola Basilico, Nicola Gatti, Pietro Testa*
(Demo 11)

TALOS is a software tool for supporting the development of security applications with mobile patrolling robots. Exploiting TALOS's functionalities, a user can easily compose a patrolling setting and apply recent algorithms presented in the multi-agent literature to find optimal patrolling strategies. Results can be evaluated and compared with intuitive graphical representations and an interacting game can be played by the user in a simulated patrolling scenario.

**Vision-Based Obstacle Run for Teams of Humanoid Robots**
*Jacky Baltes, Chi Tai Cheng, Jonathan Bagot*
(Demo 12)

This demonstration shows a team of small humanoid robots traverse an environment through a set of obstacles. The robots' brain are implemented using mobile phones for vision, balance, and processing. The robots use particle filters to localize themselves and to map the environment. A frontier-based exploration algorithm is used to direct the robots to overcome obstacles and to explore all regions of the environment.

**Evolutionary Design of Agent-based Simulation Experiments**
*James Decraene, Yew Ti Lee, Fanchao Zeng, Mahinthan Chandramohan, Yong Yong Cheng, Malcolm Yoke Hean Low*
(Demo 13)

We present CASE (complex adaptive systems evolver), a framework devised to conduct the design of agent-based simulation experiments using evolutionary computation techniques. This framework enables one to optimize complex agent-based systems, to exhibit pre-specified behavior of interest, through the use of multi-objective evolutionary algorithms and cloud computing facilities.

**Interactive Storytelling with Temporal Planning**
*Julie Porteous, Jonathan Teutenberg, Fred Charles, Marc Cavazza*
(Demo 14)

Narrative time has an important role to play in Interactive Storytelling (IS) systems. In contrast to prevailing IS approaches which use implicit models of time, in our work we have used an explicit model of narrative time. The goal of the demonstration IS system is to show how this explicit temporal representation and reasoning can help overcome certain problems experienced in IS systems such as the coordination of virtual agents and system inflexibility with respect to the staging of virtual agent actions. The fully implemented system features virtual agents and situations inspired by Shakespeare's play *The Merchant of Venice*.

**Agent-based Network Security Simulation**
*Dennis Grunewald, Marco Lützenberger, Joël Chinnow, Rainer Bye, Karsten Bsufka, Sahin Albayrak*
(Demo 15)

We present *NeSSi*$^2$ , the Network Security Simulator, a simulation environment that is based on the service-centric agent platform JIAC. It focuses on network security-related scenarios such as attack analysis and evaluation of countermeasures. We introduce the main *NeSSi*$^2$ concepts and discuss the motivation for realizing them with agent technology. Then, we present the individual components and examples where *NeSSi*$^2$ has been successfully applied.

**Experimental Evaluation of Teamwork in Many-Robot Systems**
*Andrea D'Agostini, Daniele Calisi, Alberto Leo, Francesco Fedi, Luca Iocchi, Daniele Nardi*
(Demo 16)

This paper describes a Multi-robot System, its features and a set of experiments that show its performance under different conditions (e.g., network and sensor unrealiability, number of robots, etc). The system has been fully deployed in a simulated environment; in addition, experiments with a hybrid team of cooperating simulated and real robots will be presented.

# Doctoral Consortium Abstracts

### Reasoning About Norms Within Uncertain Environments

*Natalia Criado*

The main aim of my thesis is the development of agents capable of reasoning about norms given that they are situated in an uncertain environment. The n-BDI agent architecture developed in my thesis is aimed at allowing agents to determine which and how norms will be obeyed and supporting agents when facing with norm violations.

### Privacy and Self-disclosure in Multiagent Systems

*Jose M. Such*

Agents usually encapsulate their principals' personal data attributes, which can be disclosed to other agents during agent interactions, producing a potential loss of privacy. We propose self-disclosure decision-making mechanisms for agents to decide whether disclosing personal data attributes to other agents is acceptable or not. Moreover, we also propose secure agent infrastructures to protect the information that agents decide to disclose from undesired accesses.

### Policies for Role Based Agents in Environments with Changing Ontologies

*Fatih Tekbacak, Tugkan Tuglular, Oguz Dikenelli*

Software agents try to achieve the goals of roles that they have in an environment. It is supposed that the dynamic structure of role based agents can be connected with updatable domain ontologies of the environment. Ontology evolution can cause the update of agent behaviors or access restrictions to ontological elements. So regulation for the agent behaviors may be needed. Our motivation is to create a suitable policy model for agents, environments and organizations when ontologies in the environment can change.

### Human Factors in Computer Decision-Making (PhD Thesis Extended Abstract)

*Dimitrios Antos*

This thesis investigates whether incorporating ideas from human decision-making in computer algorithms may help improve agents' decision-making performance, as either independent actors or in collaboration with humans. For independent actors, psychological cognitive appraisal theories of emotion are used to develop a lightweight algorithm that dynamically re-prioritizes their goals to direct their attention. In experiments in quickly changing and highly uncertain domains these agents are shown to perform as well as agents that compute expensive optimal solutions, and exhibit robustness with respect to the parameters of the environment. For agents interacting with humans, it is investigated whether expressing emotions has the ability to convey traits like trustworthiness and skill, and whether the appropriate emotional expression can help forge mutually beneficial relationships with the human. Finally, the theory of reasoning patterns is leveraged to analyze games and make it possible to answer questions about a system's strategic behavior without having to compute an expensive, precise solution. This theory is also employed to the generate advice for human decision-makers in complex games. This advice has been experimentally shown to improve their decision-making performance.

### Security in the Context of Multi-Agent Systems

*Gideon D. Bibu*

Security of systems and information has always been a challenge to organisations and industries. Many technical solutions including firewalls, encryption and anti-virus software have been used, yet security still remains a problem. These security solutions failures are largely due to the fact that as systems become more complex, a lot of interaction is involved between various actors. Some of these interactions usually leave room for security vulnerabilities which are simply not accounted for by the technical security solutions: there are just too many possibilities. My research is focused on this aspect of organisational security. The proposed approach to this involves the monitoring of events for traces of behaviours that may eventually circumvent the security regulations of the organisation. The methodology includes organisational modeling and simulation of self monitoring agents using a normative framework.

### Agent Dialogues and Argumentation

*Xiuyi Fan*

Agents have different interests and desires. Agents also hold different beliefs and assumptions. To accomplish tasks jointly, agents need to better convey information between each other and facilitate fair negotiations. In this thesis, we investigate agent dialogue systems developed with the Assumption-Based Argumentation (ABA) framework. Agents represent their beliefs and desires in ABA. Knowledge is exchanged via ABA arguments through dialogues. Main contributions include (1) understanding the connection between dialogues and argumentation frameworks and (2) applying argumentation dialogues in various agent applications.

### Massively Multi-Agent Pathfinding made Tractable, Efficient, and with Completeness Guarantees

*Ko-Hsin Cindy Wang*

Multi-agent path planning is a challenging problem with numerous real-life applications, including robotics, logistics, military operations planning, disaster rescue, and computer games. We look at navigating large numbers of mobile units to their targets on navigation graphs such as grid maps. The size of problems examined is significantly larger than can be handled using optimal multi-agent pathfinding algorithms in practice.

We introduced MAPP, a *tractable* algorithm for multi-agent path planning on undirected graphs. MAPP and its extended versions are complete on well specified and tractably testable classes of problems. They have low-polynomial worst-case upper bounds for the running time, the memory requirements, and the length of solutions.

Experiments on realistic game grid maps, with uniformly randomly generated start and target locations for each unit, show MAPP as a state-of-the-art multi-agent pathfinding algorithm in terms of scalability and success ratio (i.e., percentage of solved units). Even on challenging scenarios with 2000 units, MAPP solves 92% to 99.7% of units. FAR and WHCA, two fast but incomplete algorithms that were previously state-of-the-art in terms of scalability, solve as few as 17.5% and 12.3% of these problems. The quality of MAPP's solutions is empirically analyzed using multiple quality criteria: total travel distance, makespan, and sum of actions (including move and wait actions). MAPP is competitive in terms of solution quality and speed with FAR and WHCA*. MAPP further provides the formal characterizations that FAR and WHCA* lack, on problems it can solve as well as low-polynomial upper bounds on the resources required.

As optimal algorithms have limited scalability, we evaluated the solution quality of suboptimal algorithms using lower bounds of optimal values. We showed that MAPP's solutions have a reasonable quality. For example, MAPP's total travel distance is on average 19% longer than a lower bound on the optimal value.

### Securing Networks Using Game Theory: Algorithms and Applications

*Manish Jain*

Extensive transportation networks have become the economic backbone of the modern age. Thus, securing these networks against the increasing threat of terrorism is of vital importance. However, protecting critical infrastructure using limited security resources against intelligent adversaries in the presence of the uncertainty and complexities of the real-world is a major challenge. While game-theoretic approaches have been proposed for security domains, traditional methods cannot scale to realistic problem sizes (up to billions of action combinations), even in the absence of uncertainty.

My thesis proposes new models and algorithms that have not only advanced the state of the art in game-theory, but have actually been successfully deployed in the real-world. For instance, IRIS has been in use by the Federal Air Marshal Service for scheduling officers on some international flights since October 2009. My thesis contributes to a very new area that uses insights from large-scale optimization for game-theoretic problems. It represents a successful transition from game-theoretic advancements to real-world applications that are already in use, and it has opened exciting new avenues to greatly expand the reach of game theory.

### Decentralized Semantic Service Discovery based on Homophily for Self-Adaptive Service-Oriented MAS

*E. del Val*

The aim of my PhD thesis is to propose a decentralized system for service management based on the social concept of homophily. The system provides self-organizing features, and it is established and maintained without supervision. Each agent manages autonomously events such as searching services, joining or leaving the system, which reduces the service management and the structure maintenance cost. Agents,

considering only local information, carry out all these tasks.

**A Cost-Oriented Reorganization Reasoning for Multiagent Systems Organization Transitions**

*Juan M. Alberola*

Our main aim for this work is to provide a reorganization model which take into consideration the costs associated to the reorganization process. The reorganization model must provide the measurament of costs from the agent perspective (what does it cost the agent to play a new/other role) and from the organization perspective (what does it cost the organization to have an specific agent playing an specific role and how does it benefit from that). Furthermore, this measurement should be defined for static costs and also for dynamic costs. This reorganization model will allow us to reason about both reorganization dimensions: before and after reorganization. The first dimension is focused on measuring the effectiveness of the organization in the future and analyzing whether the organization will be able to cope with some changed circumstances. The second dimension is focused on measuring the impact of the problems appeared during the reorganization process in the cost of change.

**Graphical Multiagent Models**

*Quang Duong*

I introduce a graphical representation for modeling multi-agent systems based on different kinds of reasoning about agent behavior. I seek to investigate this graphical model's predictive and representative capabilities across various domains, and examine methods for learning the graphical structure from agent interaction data. I also propose to explore the framework's scalability in large real-world scenarios, such as social networks, and evaluate its prediction performance with existing network behavior models.

**Extended Abstract of Elisabetta Erriquez Thesis**

*Elisabetta Erriquez*

Agents in Multi-Agent Systems depend on interactions with others to achieve their goals. Often, goals of agents conflict with each other, and agents can be unreliable or deceitful. Therefore, trust and reputation are key issues in this domain. As in human societies, software agents must interact with other agents in settings where there is the possibility that they can be exploited. This suggests the need for computational models of trust and reputation that can be used by software agents, and accordingly, much research has investigated this issue over the past decade. This thesis concentrates on two important questions, therefore it is divided in two parts. The first question is what sources agents can use to build their trust of others upon. The second question is how agents can use trust and reputation concepts to form stable coalitions.

**Improving Game-tree Search by Incorporating Error Propagation and Social Orientations**

*Brandon Wilson*

Game-tree search algorithms, such as the two-player Minimax algorithm and its multi-player counterpart, Max-n, are a fundamental component in the development of computer programs for playing extensive-form games. The success of these algorithms is limited by the underlying assumptions on which they are built. For example, it is traditionally assumed that deeper search always produces better decisions and also that search procedures can assume all players are selfish and ignore social orientations. Deviations from these assumptions can occur in real games and can affect the success of a traditional search algorithms. The goal of my thesis is to determine when such deviations occur and modify the search procedure to correct the errors that are introduced.

**Negotiation Teams in Multiagent Systems**

*Víctor Sánchez-Anguix*

In this paper, I present my ongoing research on agent-based negotiation teams. An agent-based negotiation team is a group of two or more agents with their own and possibly conflicting goals that join together as a single negotiation party because they share a common goal that is related to the negotiation. Our research goal is to provide agent-based solutions for problems which may need negotiation teams.

**Real-World Security Games: Toward Addressing Human Decision-Making Uncertainty**

*James Pita*

Game theory is a useful tool for reasoning about interactions between agents and in turn aiding in the decisions of those agents. In fact, Stackelberg games are natural models for many important applications such as oligopolistic markets and security domains. Indeed, Stackelberg games are at the heart of three deployed systems, ARMOR; IRIS; and GUARDS, for aiding security officials in making critical resource allocation decisions. In Stackelberg games, one player,

the leader, commits to a strategy and follower makes her decision with knowledge of the leader's commitment. Existing algorithms for Stackelberg games efficiently find optimal solutions (leader strategy), however, they critically assume that the follower plays optimally. Unfortunately, in many applications, agents face human followers (adversaries) who - because of their bounded rationality and possibly limited information of the leader strategy - may deviate from their expected optimal response. Not considering these likely deviations when dealing with human adversaries may cause an unacceptable degradation in the leader's reward, particularly in security applications where these algorithms have seen deployment. To that end, I explore robust algorithms for agent interactions with human adversaries in security applications. I have developed a number of robust algorithms for a class of games known as "Security Games" and am working toward enhancing these approaches for a richer models of these games that I developed known as "Security Circumvention Games".

**A Multi-Agent System for Predicting Future Event Outcomes**
*Janyl Jumadinova*

My research focuses on understanding and analyzing prediction markets using multi-agent system and game theory-based tools. Prediction markets, in which trading agents trade outcomes of events (like "Barack Obama will win the presidential election"), are similar to stock markets in that they aggregate the beliefs of many agents with different information about an event, and produce a single number, a market price. The main challenge in designing, deploying and operating prediction markets is the lack of understanding about what makes a successful prediction market, i.e.(1) how information affects traders and prediction market price, (2) how to incentivize traders to reveal their socially useful information, and (3) under what conditions do prediction markets perform the best. I address the first question by developing a multi-agent system that is used to analyze the effect of information on the prediction market performance and show that our model provides a better understanding and novel insights into the behavior of prediction markets and its participants. For the second question I propose a correlated equilibrium strategy for the trading agents within a partially observable stochastic game-based model that incentivizes traders to reveal their true beliefs and show that the market's participants following my strategy achieve higher profit by 35-127%. And for the third question I apply Boolean Network techniques to study the dynamics of prediction markets under various conditions and show under what parameter values the prediction market stabilizes and how its dynamics change with the introduction of noise and increased number of traders.

**A Study of Computational and Human Strategies in Revelation Games**
*Peled Noam*

This thesis focuses on the design of autonomous agents which can negotiate with people using argumentation strategies. Argumentation is the ability to argue and to persuade another party to accept a desired agreement, to acquire or give information, to coordinate goals and actions and to find and verify evidence. Argumentation is endemic to human interaction. It facilitates knowledge about people's positions, and may improve the final outcome of negotiation. Despite the importance of argumentation within the general framework of negotiations, work on argumentation over the last few years has focused almost exclusively on the context of rational interactions between self-interested, automated agents.

**Thesis Research Abstract: Modeling Crowd Behavior Based on Social Comparison Theory**
*Natalie Fridman*

Modeling crowd behavior is an important challenge for agent-based simulation. My overall goal is to provide a single computational cognitive mechanism that, when executed by individual agents, would give rise to different crowd behaviors, depending on the perceptions and actions available to each individual. I propose a novel model of crowd behavior, based on Social Comparison Theory (SCT), a popular social psychology theory that has been continuously evolving. I am pursuing a concrete algorithmic framework for SCT and evaluating it on different social behaviors. Moreover, I have begun to explore the use of qualitative reasoning techniques to model global (macro-level) social phenomena in demonstrations. I believe that this is the first use of QR techniques for such purposes.

**Cooperation between Self-Interested Agents in Normal Form Games**
*Steven Damer*

We study how to achieve cooperation between two self-interested agents that play repeated randomly generated normal form games. We take inspiration

from a model originally designed to identify cooperative actions by humans who play a game, but we use the model in a prescriptive rather than descriptive manner. To identify cooperative intent, agents use a particle filter to learn the parameters of the model.

**Group Decision Making in Multiagent Systems with Abduction**
*Samy Sá*

In Multiagent Systems (MAS), various activities are related to decisions involving a group of agents such as negotiation, auctions and social choice. Group Decision Making (GDM) specializes in the situations where a group of agents need to pick one of possibly many options from a set and commit to it. We intend to provide a novel GDM framework in which the agents are able to employ abductive reasoning and discuss the options towards consensus.

**Security Games with Mobile Patrollers**
*Ondřej Vaněk*

To optimally secure large and complex infrastructures against crime activities, a scalable model for optimal defender allocation is needed. Game theory is successfully used to formalize the problem as a two-player game between an attacker and a defender. We consider both player to be mobile and we focus on proper path intersection modeling and we observe the trade-off between fidelity and computational complexity. We search for the a Nash Equlibirium of the game using oracle based algorithms and we evaluate the robustness of the solution in a multi-agent simulation where some assumptions made do not strictly hold.

**Self-Organization in Decentralized Agent Societies through Social Norms**
*Daniel Villatoro*

My research on normative systems is focused on the spreading and acceptance of social norms, defined as *norm support*. Our understanding of norm support deals with the problem of which norm is established as the dominant. Specifically, we deal with two different branches of the research on normative sytems: conventional norms and essential norms. On the one hand conventional norms fix one norm amongst a set of norms that are equally efficient as long as every member of the population uses the same (e.g. communication protocols, greetings, driving side of the road), and on the other hand, essential norms solve or ease collective action problems, where there is a conflict between the individual and the collective interests. The scientific question of this research is how to accelerate the establishment of a common norm in virtual societies: in the case of conventional norms, by dissolving the subconventions; and in the case of essential norms, by studying the effects of punishment and norm internalization.

**A Trust Model for Supply Chain Management**
*Yasaman Haghpanah*

My thesis will contribute to the field of multi-agent systems by proposing a novel and formal trust-based decision model for supply chain management.

# Author Index