Lab Seminar
Foundations of Artificial Intelligence

14. Oktober 2009

12:20 – 12:50 Jens Witkowski

Eliciting Non-Common Prior Beliefs in Peer Prediction

Reputation systems rely on truthful feedback from their users. The "Peer Prediction Method" that was developed by Miller-et-al and extended by Jurca and Faltings is arguably the most prominent feedback mechanism in the literature. An obstacle with regard to its application are the strong common knowledge assumptions. Especially the commonly held prior belief, although prevailing in economic theory, is too strict for many multiagent settings. Two issues stand out here, in particular: first, the agents' beliefs are often unknown to the mechanism and, second, different agents often hold different beliefs. We therefore study procedures that not only elicit the agents' feedback reports but also their prior beliefs. While the simultaneous announcement of the prior belief and the signal is not feasible, we show that truthfulness can be obtained through a temporal separation of the prior announcement and the reception of the signal. Based on this finding, we design a mechanism that is both truthful and budget-optimal.

14:30 – 15:30 Malte Helmert

Planen als heuristische Suche: Theoretische Grenzen und praktische Algorithmen

In diesem Übersichtsvortrag stelle ich meine Forschungsarbeiten der letzten Jahre über heuristische Suchverfahren für Probleme der klassischen Handlungsplanung und andere Anwendungen vor. Im theoretischen Teil sehen wir, dass der Ansatz der heuristischen Suche eigentlich hoffnungslos ist; im algorithmischen Teil sehen wir dann, dass er trotzdem sehr gut funktioniert.

15:30 – 16:00 Gabi Röger

Combining Heuristic Estimators for Satisficing Planning

Joint work with Malte Helmert.

The problem of effectively combining multiple heuristic estimators has been studied extensively in the context of optimal planning, but not in the context of satisficing planning. To narrow this gap, we empirically examine several ways of exploiting the information of multiple heuristics in a satisficing best-first search algorithm, comparing their performance in terms of coverage, plan quality and runtime. Our empirical results indicate that using multiple heuristics for satisficing search is indeed useful and that the best results are not obtained by the most obvious combination methods.

16:40 – 17:10 Dapeng Zhang

A Life-Long Learning Method for Conditional Random Fields

Convergence is one of the important properties of a learning algorithm. It guarantees the results of the learning, but confines the abilities of the learned algorithm. Conditional Random Fields (CRFs) can be regarded as a fully connected stochastic state automaton. Typically, the states of CRFs are predefined and their parameters can be estimated [Lafferty et al. 2001]. We develop a method to increase the number of the states in CRFs during the learning process. Each time after the number of the states has been increased, their parameters are estimated until the convergence. Consequently, the resulting learning has a much bigger improvement space than normal CRFs. The system based on this method can thus acquire a "life-long learning" ability. The recent experiments on the table soccer robot will be discussed.

17:10 – 17:40 Christian Dornhege

Integration of symbolic and geometric planning

I will present initial results of a hybrid planning system linking symbolic and geometric planning. The symbolic planning system has been enhanced using "Semantic Attachments" to allow the integration of low-level reasoners as geometric planning. The system will be applied to construct a manipulation planner by using a modified blocks-world domain that integrates a probabilistic roadmap planner into Temporal Fast Downward.

17:40 – 18:10 Patrick Eyerich, Thomas Keller

Task Planning for Autonomous Robots

With the increasing number of abilities of autonomous robots, it becomes more and more important to plan their behavior online because the preparation of scripts that can handle every possible situation is highly error-prone. We present some issues that arise when adapting a PDDL-based planning system to real-world problems, and show ways how to tackle those challenges. In particular, we cover issues like the generation of PDDL-tasks based on information distributed among loosely coupled robot subsystems, repeated planning, execution, and monitoring, or how to include subplanning processes like path and motion planning into a symbolic planning system.

15. Oktober 2009

9:30 – 10:00 Michael Brenner, Moritz Göbelbecker

Continual Planning in the CogX Project

In this talk, we will present the EU-funded project CogX ("Cognitive Systems that Self-Understand and Self-Extend") and our own work in this project towards robots that are able to work in open ended, challenging environments, dealing with novelty, uncertainty and change. In such environments, the robots should understand (the limits of) their own abilities and knowledge (self-understanding), and be able to extend these based on their experiences and goals (self-extension). Our specific task in Freiburg is to study representations and methods that enable the robot to plan its physcial actions as well as its interactions with humans and its internal processes. Furthermore, we want to enable the robot to decide which specific planning technique to use for a given task, which information to use for planning and how to learn from previous planning episodes. In the talk, we will present the state of our work two weeks before the first major review of the project and discuss where we want to go from there.

10:00 – 10:30 Alexander Kleiner, Dali Sun

On Pirates Supporting Business Tasks

Although a remarkably high degree of automation has been reached in production and intra-logistics nowadays, handcarts and forklifts manually steered by humans are still indispensable in many situations, which is clearly undesirable due to high costs and risk of failure. Alternatively, the problem is solved in automated warehouses with fixed installed conveyors either overhead- or floor-based. However, those solutions have several drawbacks, for example, when the business model of the company changes existing installations have to be re-designed. In this talk, we introduce a novel approach to coordinate decentralized teams of autonomous robots performing intra-logistics tasks. Key challenges in this scenario are to maintain transportation tasks within a peer-to-peer network, keeping the network structure intact, and most importantly, to facilitate a fair distribution of robots among loading stations. The contribution of this talk is a solution to all of these problems at the same time by adopting a variant of the successful concept behind file sharing systems on the Internet, namely weighted decentralized hash tables (WDHTs). Experimental results presented show, that our method reaches a fair distribution of robots over loading stations.

11:10 – 11:40 Robert Mattmüller

Component-based Abstraction Refinement for Timed Controller Synthesis

We present a novel technique for synthesizing controllers for distributed real-time environments with safety requirements. Our approach is an abstraction refinement extension to the on-the-fly algorithm by Cassez et al. Based on partial compositions of some environment components, each refinement cycle constructs a sound abstraction that can be used to obtain under- and over-approximations of all valid controller implementations. This enables (1) early termination if an implementation does not exist in the over-approximation, or, if one does exist in the under-approximation, and (2) pruning of irrelevant moves in subsequent refinement cycles. In our refinement loop, the precision of the abstractions incrementally increases and converges to all specification-critical components.

We implemented our approach in a prototype synthesis tool and evaluated it on an industrial benchmark. In comparison with the timed game solver UPPAAL-TiGa, our technique outperforms the nonincremental approach by an order of magnitude.

11:40 – 12:10 Sebastian Kupferschmid

Interference Pruning

Directed model checking is a well-established technique to efficiently tackle the state explosion problem when the aim is to find error states in large systems. We propose a novel technique for traversing the state space based on the notion of interference. Our approach is orthogonal to the model checking process and can be applied to a wide range of search methods without losing completeness. We have implemented our method and empirically evaluated its potential on a range of non-trivial case studies. Compared to standard model checking techniques, we are able to detect subtle bugs with shorter error traces, consuming less memory and time.

12:10 – 12:40 Jan-Georg Smaus

Finding Errors of Hybrid Systems by Optimising an Abstraction-Based Quality Estimate

Joint work with Stefan Ratschan

We present an algorithm for falsifying safety properties of hybrid systems, i.e., for finding a trajectory to an unsafe state. The approach is to approximate how close a point is to being an initial point of an error trajectory using a real-valued quality function, and then to use numerical optimisation to search for an optimum of this function. The function is computed by running simulations, where information coming from abstractions computed by a verification algorithm is exploited to determine whether a simulation looks promising and should be continued or cancelled. This information becomes more reliable as the abstraction becomes more refined. We thus interleave falsification and verification attempts.