



BISFAI 2007


9th Bar-Ilan Symposium on the Foundations of AI
in conjunction with
ISCOL, Israeli Seminar on Computational Linguistics
Robotics, supported by the Ministry of Science

June 20-22, 2007
Bar-Ilan University
Ramat Gan, Israel

<http://www.cs.biu.ac.il/bisfai07>

At Gonda Brain Research Center

Conference Booklet



Introduction to BISFAI

BISFAI, the Bar-Ilan Symposium on Foundations of Artificial Intelligence, was established in 1989 by Prof. Martin Golumbic. Since then, it has taken place every other year (besides 2003), and is regularly held at Bar-Ilan University by the Department of Computer Science, with the exception of 2005, when it was hosted by the Caesarea Rothschild Institute (CRI) at University of Haifa.

The biennial symposium focuses on the foundations of AI and, over the years, has often had additional specialized themes. It boasts excellent talks in different subfields of AI. Two books and several special issues of journals have been published from selected papers presented at BISFAI. World renowned AI researchers regularly visit and lecture at BISFAI, which gives an opportunity for the Israeli AI community, especially the students, to meet their peers from abroad. To many, BISFAI has become the premier event of AI happening in Israel.

BISFAI 2007 took place at Bar-Ilan, June 20-22, 2007. The symposium chair was Sarit Kraus, the program chairs were Gal Kaminka and Shlomo Argamon, and the organization chair was Ariel Frank. The event was held in conjunction with the Israeli Seminar on Computational Linguistics (ISCOL), organized by Ari Rappoport and Ido Dagan, and a special Robotics track, organized by Gal Kaminka and Eli Kolberg. Martin Golumbic was the honorary chair. The keynote speakers in BISFAI 2007 (in alphabetic order) were:

Yaacov Choueka (Bar-Ilan U.)
Dana Nau (U. of Maryland)
Milind Tambe (U. of Southern California)
Manuela Veloso (Carnegie Mellon U.)
Michael Wooldridge (U. of Liverpool)

The following is a list of the keynote speakers from previous BISFAI conferences:

BISFAI 2005

Richard Korf (UCLA)
John McCarthy (Stanford U.)
Helmut Prendinger (Japanese National Institute of Informatics)
Bart Selman (Cornell U.)
Yoav Shoham (Stanford U.)
Shlomo Zilberstein (U. of Massachusetts)

BISFAI 2001

Yaacov Choueka (Bar-Ilan U.)
Jieh Hsiang (National Taiwan U.)
Daphne Koller (Stanford U.)
Oliviero Stock (ITC-irst, Trento)
Moshe Vardi (Rice U.)

BISFAI 1999

Leo Joskowicz (Hebrew U.)
Stanley Rosenschein (Stanford U.)
Reid Simmons (Carnegie Mellon U.)

BISFAI 1997

Craig Boutilier (U. of British Columbia)
Michael Kearns (AT&T)
Yoav Shoham (Stanford)
Wolfgang Wahlster (German Research Center for AI, DFKI GmbH)

BISFAI 1995

Robert Berwick (MIT)
Aravind Joshi (U. of Pennsylvania)
Hans Kamp (Stuttgart U.)
Jack Minker (U. of Maryland)
Sergei Nirenburg (New Mexico State U.)
Naftaly Tishby (Hebrew U.)

BISFAI 1993

Robert Aumann (Hebrew U.)
Barbara Grosz (Harvard U.)
Vladimir Lifschitz (U. of Texas at Austin)
Jeffrey Rosenschein (Hebrew U.)

BISFAI 1991

Hans Berliner (Carnegie Mellon U.)
Ernest Davis (Courant Institute, NYU)
Monique Pavel (U. of Paris)
Stanley Peters (Stanford U.)

BISFAI 1989

Joe Halpern (IBM Research)
Johann A. Makowsky (Technion)
John McCarthy (Stanford U.)
Judea Pearl (UCLA)
Ron Rivest (MIT)

BISFAI has been an ongoing success due to the dedication of the Israeli AI community members, and especially thanks to the previous Chairs who have diligently served (some quite a few times ☺) over the years (in alphabetical order): Nachum Dershowitz, Ronen Feldman, Ariel Frank, Claudia Goldman, Martin Charles Golumbic, Leo Joskowicz, Moshe Koppel, Sarit Kraus, Daniel Lehmann, Leora Morgenstern, Mori Rimon, Uri Schild, Eli Shamir, and Solomon Eyal Shimony.

Conference Program

רביעי, ד' תמוז, תשס"ז Wednesday, June 20, 2007

09:00-09:25 Registration

09:25-10:50 Session ISCOL 1 (ISCOL opening) Session chair: Ari Rappoport
Roy Bar-Haim, Ido Dagan, Iddo Greental and Eyal Shnarch. <i>Semantic Inference at the Lexical-Syntactic Level</i>
Ben Sandbank, Shimon Edelman and Eytan Ruppín. <i>From ConText to Grammar: A Step towards Practical Probabilistic Context Free Grammar Inference</i>
Dana Dannells and Louise Deleger. <i>Multilingual Generation of Medical Information</i>
Alon Lavie and Shuly Wintner. <i>Rapid Prototyping of a Transfer-based Hebrew-to-English Machine Translation System</i>

09:45-10:50 Session Robotics 1 (Robotics opening) Session chair: Israel Wagner
Eli Kolberg, Yoram Reich and Ilya Levin. <i>Design Methodology for Mobile Robots</i>
Yehuda Elmaliah, Noa Agmon and Gal Kaminka. <i>Multi-Robot Area Patrol under Frequency Constraints</i>
Gal Kaminka, Ari Yakir, Dan Erusalimchik and Nirom Cohen-Nov. <i>Towards Collaborative Task and Team Maintenance</i>

10:50-11:20 Coffee break

11:20-13:00 Session ISCOL 2 Session chair: Ido Dagan
Roi Reichart and Ari Rappoport. <i>An Ensemble Method for Selection of High Quality Parses</i>
Yoav Goldberg and Michael Elhadad. <i>Toward Better Understanding of Hebrew NP Chunks</i>
Reut Tsarfaty and Khalil Sima'an. <i>Dimensions of Parameterization for Modern Hebrew Statistical Parsing</i>
Saib Mansour, Khalil Sima'an and Yoad Winter. <i>Smoothing a Lexicon-based POS Tagger for Arabic and Hebrew</i>
Joshua Waxman. <i>Transliteration Bitext Corpus Generation and Cross-Language Information Retrieval Using Minimum Edit Distance and Character Cluster Equivalence Classes</i>

11:20-13:00 Session Robotics 2 Session chair: Eli Kolberg
Yan Virin, Guy Shani and Solomon Eyal Shimony. <i>Scaling Up: Solving POMDPs through Value Based Clustering</i>
Eliyahu Osherovich, Vladimir Yanovski, Israel Wagner and Freddy Bruckstein. <i>Robust and Efficient Covering of Unknown Continuous Domains with Simple, Ant-Like Agents</i>
Ami Berler and Solomon Eyal Shimony. <i>Mapping AWOL Using a Concrete Domain</i>
Ariel Felner, Roni Stern, Zahi Benaya and Tal Fadida. <i>Problem Solving by Moving Agents in Physical Graphs</i>

13:00-14:15 Lunch break

14:15-15:30 Keynote 1 (BISFAI Opening)
Session chair: Sarit Kraus
BISFAI Greetings
Keynote: Yaacov Choueka. <i>You've come a long way, Baby! Responsa meet Full-text and CL, one-thousand-year manuscripts meet AI</i>

15:30-16:10 Session ISCOL 3
Session chair: Michael Elhadad
Dmitry Davidov, Ari Rappoport and Moshe Koppel. <i>Fully Unsupervised Discovery of Concept-Specific Relationships by Web Mining</i>
Idan Szpektor and Ido Dagan. <i>Learning Canonical Forms of Entailment Rules</i>

15:30-16:10 Session AI 1 (Clustering)
Session chair: Amnon Meisels
Maxim Binshtok, Ronen Brafman, Solomon Eyal Shimony, Ajay Mani and Craig Boutilier. <i>Computing Optimal Subsets</i>
Rachel Ben-Eliyahu-Zohary, Ran Giladi, Philip Hendrix and Stuart M. Shieber. <i>Clustering Ad-Hoc Networks: Experiments in Local Search</i>

16:10-16:40 Coffee break

16:40-18:00 Session ISCOL 4
Session chair: Alon Itai
Roi Reichart and Ari Rappoport. <i>Self-Training for Enhancement and Domain Adaptation of Statistical Parsers Trained on Small Datasets</i>
Ido Dagan, Oren Glickman, Alfio Massimiliano Gliozzo, Efrat Hershkovic Marmorshtein and Carlo Strapparava. <i>Direct Word Sense Matching for Lexical Substitution</i>
Yael Netzer, Meni Adler, David Gabay and Michael Elhadad. <i>Can You Tag the Modal? You Should</i>
Judit Bar-Ilan, Kevin Keenoy, Mark Levene and Eti Yaari. <i>Ranking Preferences of Search Results – A User Study</i>

16:40-18:00 Session AI 2 (Search)
Session chair: Carmel Domshlak
Roie Zivan, Moshe Zazone and Amnon Meisels. <i>Min-Domain retroactive ordering for Asynchronous Backtracking</i>
Ariel Felner and Nir Ofek. <i>Combining Perimeter Search and Pattern Database Abstractions</i>
Dragan Bosnacki, Edith Elkind, Blaise Genest and Doron Peled. <i>On Commutativity Based Edge Lean Search</i>
Uzi Zahavi and Ariel Felner. <i>Inconsistent Heuristics</i>

תמישי, ה' תמוז, תשס"ז Thursday, June 21, 2007

08:30-09:00 Registration

09:00-10:45 Session AI 3 (Machine Learning) Session chair: Shaul Markovitch
Ariel Procaccia, Aviv Zohar, Yoni Peleg and Jeffrey Rosenschein. <i>Learning Voting Trees</i>
Ron Katz and Sarit Kraus. <i>Gender-Sensitive Automated Negotiators</i>
Roy Fox and Moshe Tennenholtz. <i>A Reinforcement Learning Algorithm with Polynomial Interaction Complexity for Only-Costly-Observable MDPs</i>
Larry Manevitz and Hananel Hazan. <i>History-Dependent Neurons and Identification of Temporal Sequences</i>
Alon Altman, Avivit Bercovici-Boden and Moshe Tennenholtz. <i>Learning in one-shot strategic form games</i>

10:45-11:15 Coffee break Demonstration: Avi Pfeffer, <i>IBAL: Turning Probabilistic Reasoning into Programming.</i>
--

11:15-12:15 Keynote 2 Session chair: Gal Kaminka
Keynote: Manuela Veloso. <i>Selective Use of Multiple Sources of Robot Sensory Information</i>

12:15-13:00 Session AI 4 (Game Complexity) Session chair: Moshe Tennenholtz
Yoram Bachrach and Jeffrey Rosenschein. <i>Computing the Banzhaf Power Index in Network Flow Games</i>
Edith Elkind, Leslie Ann Goldberg, Paul Goldberg and Michael Wooldridge. <i>Computational Complexity of Weighted Threshold Games</i>

13:00-14:00 Lunch break

14:00-15:45 Session AI 5 (Planning and Plan Recognition) Session chair: Solomon Eyal Shimoni
Michael Katz and Carmel Domshlak. <i>Structural Patterns of Tractable Sequentially-Optimal Planning</i>
Carmel Domshlak and Vitaly Mirkis. <i>Cost-Sharing Approximations for h^+</i>
Hannaneh Hajishirzi and Eyal Amir. <i>Stochastic Filtering in a Probabilistic Action Model</i>
Dorit Avrahami-Zilberbrand and Gal Kaminka. <i>Incorporating Observer Biases in Keyhole Plan Recognition (Efficiently!)</i>
Nirom Cohen-Nov Slapak and Gal Kaminka. <i>Computational Load and Performance in Integrated Multi-Agent Intention Recognition</i>

15:45-16:15 Coffee break Demonstration: Avi Pfeffer, <i>IBAL: Turning Probabilistic Reasoning into Programming.</i>
--

16:15-17:15 Keynote 3 Session chair: Ronen Brafman
Keynote: Dana Nau. <i>May All Your Plans Succeed! (Or Have a High Expected Utility)</i>

17:15-18:15 Session AI 6 (Modelling) Session chair: Meir Kalech
Shlomo Berkovsky, Ariel I. Gorfinkel, Tsvi Kuflik and Larry M. Manevitz. <i>Case-Based to Content-Based User Model Mediation and its Effectiveness</i>
Guy Shani. <i>A Stereotypes-Based Hybrid Recommender System for Media Items</i>
Orna Peleg, Zohar Eviatar, Larry M. Manevitz and Hananel Hazan. <i>Differences and Interactions between Cerebral Hemispheres When Processing Ambiguous Homographs</i>

שישי, ו' תמוז, תשס"ז Friday, June 22, 2007

08:30-09:00 Registration

09:00-10:45 Session AI 7 (Multi-Agent Systems I) Session chair: Ariel Felner
Avi Rosenfeld, Sarit Kraus and Charlie Ortiz. <i>Quantifying the Expected Utility of Information in Multi-Agent Scheduling Tasks</i>
Meir Kalech, Gal Kaminka and Michael Lindner. <i>Matrix-Based Representation for Coordination Fault Detection: A Formal Approach</i>
Alex Rogers, Esther David, Terry Payne and Nicholas R. Jennings. <i>An Advanced Bidding Agent for Advertisement Selection on Public Displays</i>
Ido Levy and Claudia Goldman. <i>Distributed On-line Learning Cooperative Multi-Agent System</i>
Avi Rosenfeld, Claudia Goldman, Sarit Kraus and Gal Kaminka. <i>An Agent Architecture for Hybrid P2P Free-Text Search</i>

10:45-11:15 Coffee break

11:15-12:15 Keynote 4 Session chair: Martin Golumbic
Keynote: Milind Tambe. <i>Multiagent and Agent-human Teamwork: Hybrid Approaches</i>

12:15-13:00 Session AI 8 (Game Theory) Session chair: Jeff Rosenschein
Aviv Zohar, Hagay Levin and Michael Schapira. <i>Incentive-Compatible Distributed Routing</i>
Michal Penn, Maria Polukarov and Moshe Tennenholtz. <i>Congestion Games with Load-Dependent Failures: Identical Resources</i>

13:00-14:00 Lunch break

14:00-14:40 Session AI 9 (Multi-Agent Systems II) Session chair: Esther David
Natalie Fridman and Gal Kaminka. <i>Towards a Cognitive Model of Crowd Behavior Based on Social Comparison Theory</i>
Noam Hazon, Paul E. Dunne, Sarit Kraus and Michael Wooldridge. <i>How to Rig an Election</i>

14:40-15:40 Keynote 5 (BISFAI closing) Session chair: Eyal Amir
Keynote: Michael Wooldridge. <i>Logic for Automated Mechanism Design – A Progress Report</i>

Conference Abstracts

רביעי, ד' תמוז, תשס"ז Wednesday, June 20, 2007

09:25-10:50 Session ISCOL 1

Roy Bar-Haim, Ido Dagan, Iddo Greental and Eyal Shnarch. *Semantic Inference at the Lexical-Syntactic Level*

Semantic inference is an important component in many natural language understanding applications. Classical approaches to semantic inference rely on complex logical representations. However, practical applications usually adopt shallower lexical or lexical-syntactic representations, but lack a principled inference framework. We propose a generic semantic inference framework that operates directly on syntactic trees. New trees are inferred by applying entailment rules, which provide a unified representation for varying types of inferences. We describe the construction of new rule base that covers inferences related to generic linguistic structures, as well as treatment of negation and modality. We further create entailment rules for specific lexical-based inferences, using both automatic acquisition methods and existing manually-created resources. Finally, we present a novel evaluation methodology for inference systems, based on relation extraction from a large corpus. This methodology is particularly suitable for robust component-wise evaluation of knowledge-based inference systems. Initial empirical evaluation supports the validity of our inference approach.

Ben Sandbank, Shimon Edelman and Eytan Ruppim. *From ConText to Grammar: A Step towards Practical Probabilistic Context Free Grammar Inference*

Unsupervised probabilistic context free grammar inference commonly proceeds towards ever more generative grammars during the training process. Consequently, an over-generalization error made during an early phase in the process is perpetuated even though it is potentially possible to detect it from the training data. Furthermore, existing algorithms usually employ a constituency criterion that is decoupled from their substitutability criterion. Here we present a new algorithm that works without a separate constituency criterion, and defines constituency only when it aids generalization. We also introduce an over-generalization control scheme which undoes errors during training. The resulting algorithm's "behavior" is analyzed and its performance is shown to surpass that of state of the art grammar inference algorithms.

Dana Dannells and Louise Deleger. *Multilingual Generation of Medical Information*

Multilingual generation systems aim to produce understandable texts in multiple languages from one knowledge representation. We adapted an existing prototype multilingual generator that presents simulated breast cancer EHRs in English to French and Swedish. The purpose of this work was to test how much effort it would require to modify a limited domain, template-based English generator to enable it to generate in French and Swedish. We describe the adaptation to both languages, viewing the grammatical aspects involved and explaining the modifications performed. The results prove that relatively little effort is required to produce practical output in the additional languages. This work illustrates how the same underlying knowledge representation can be used to generate output texts in multiple languages with only minor linguistic modifications.

Alon Lavie and Shuly Wintner. *Rapid Prototyping of a Transfer-based Hebrew-to-English Machine Translation System*

We describe the rapid development of a preliminary Hebrew-to-English Machine Translation system under a transfer-based framework specifically designed for rapid MT prototyping for languages with limited linguistic resources. The task is particularly challenging due to two main reasons: the high lexical and morphological ambiguity of Hebrew and the dearth of available resources for the language. Existing, publicly available resources were adapted in novel ways to support the MT task. The methodology behind the system combines two separate modules: a transfer engine which produces a lattice of possible translation segments, and a decoder which searches and selects the best scoring combination of translation segments according to an English language model and several additional features. We demonstrate that a small manually crafted set of transfer rules suffices to produce legible translations. Performance results are evaluated using state of the art measures and are shown to be encouraging.

09:45-10:50 Session Robotics 1

Eli Kolberg, Yoram Reich and Ilya Levin. *Design Methodology for Mobile Robots*

We present a design methodology for mobile robots. We designed this methodology in the context of a robotics course in high schools. The motivation for designing this new methodology was improving the robots' robustness and reliability and preparing students for becoming better designers. The new methodology proved to be highly successful in designing top quality robots. In the methodology design, we explored and adapted design methods to the specific designers, the nature of the product, the environment, the product needs, and the design context goals. At the end of this comprehensive design, we selected a synergetic integration of six methods to compose the methodology for this product context: conceptual design, fault tolerant design, atomic requirements, using fuzzy logic for the control of robotics systems, creative thinking method, and microprogramming design.

Yehuda Elmaliah, Noa Agmon, Gal Kaminka. *Multi-Robot Area Patrol under Frequency Constraints*

This paper discusses the problem of generating patrol paths for a team of mobile robots inside a designated target area. Patrolling requires an area to be visited repeatedly by the robot(s) in order to monitor its current state. First, we present frequency optimization criteria used for evaluation of patrol algorithms. We then present a patrol algorithm that guarantees maximal uniform frequency, i.e., each point in the target area is covered at the same optimal frequency. This solution is based on finding a circular path that visits all points in the area, while taking into account terrain directionality and velocity constraints. Robots are positioned uniformly along this path, using a second algorithm. Moreover, the solution is guaranteed to be robust in the sense that uniform frequency of the patrol is achieved as long as at least one robot works properly.

Gal Kaminka, Ari Yakir, Dan Erusalimchik and Nirrom Cohen-Nov. *Towards Collaborative Task and Team Maintenance*

There is significant interest in modeling teamwork in agents. In recent years, it has become widely accepted that it is possible to separate team work from taskwork, providing support for domain-independent teamwork at an architectural level, using teamwork models. However, existing teamwork models (both in theory and practice) focus almost exclusively on achievement goals, and ignore maintenance goals, where the value of a proposition is to be maintained over time. Such maintenance goals exist both in taskwork (i.e., agents take actions to maintain a condition while a task is executing), as well as in teamwork (i.e., agents take actions to maintain the team). This paper presents mechanisms for collaborative maintenance in both taskwork and teamwork, allowing for flexible selection of the maintenance protocol. The mechanism is integrated and evaluated in two teamwork architectures for situated agent teams: DIESEL, an implemented teamwork and taskwork architecture, built on top of Soar, and BITE, an architecture for physical behavior-based robots. We provide details of these implementations, and the results from experiments demonstrating the benefits of support for collaborative maintenance processes, in several dynamic rich domains. We show that the use of collaborative maintenance leads to significant improvement in task performance in all domains.

11:20-13:00 Session ISCOL 2

Roi Reichart and Ari Rappoport. *An Ensemble Method for Selection of High Quality Parses*

While the average performance of statistical parsers gradually improves, they still attach to many sentences annotations of rather low quality. The number of such sentences grows when the training and test data are taken from different domains, which is the case for major web applications such as information retrieval and question answering. In this paper we present a Sample Ensemble Parse Assessment (SEPA) algorithm for detecting parse quality. We use a function of the agreement among several copies of a parser, each of which trained on a different sample from the training data, to assess parse quality. We experimented with both generative and re-ranking parsers (Collins, Charniak and Johnson respectively). We show superior results over several baselines, both when the training and test data are from the same domain and when they are from different domains. For a test setting used by previous work, we show an error reduction of 31% as opposed to their 20%.

Yoav Goldberg and Michael Elhadad. *Toward Better Understanding of Hebrew NP Chunks*

In (Goldberg and Elhadad, 2007), we presented two techniques (SVM Model Tampering and Anchored Learning) for investigating the SVM learning process and resulting models. These techniques were applied to the task of SVM based Hebrew NP Chunking. The results were better understanding of SVM based chunking, of the role lexical features play in the learned models, of the Hebrew NP chunking

task definition, and identification of several deficiencies in the NP Chunking corpus. This paper focuses on the Hebrew specific issues raised by that work, which are presented with more detail.

Reut Tsarfaty and Khalil Sima'an. *Dimensions of Parameterization for Modern Hebrew Statistical Parsing*

We evaluate empirically the performance of enriched unlexicalized treebank Probabilistic Context Free Grammars (PCFGs) for parsing Modern Hebrew (MH). We show that contrary to experience with parsing the WSJ Penn Treebank, the Immediate-head unlexicalized variety does not necessarily outperform PCFGs for Semitic languages that differ in structure and characteristics from English. We demonstrate that enriching unlexicalized PCFGs with morphological features percolated up the tree outperforms plain PCFGs on the newly available MH Treebank (v2.0). We further show that an (unlexicalized) Immediate-head variety enriched with morphologically marked agreement features achieves even better performance. We conclude that morphologically-rich languages introduce an additional dimension of parameterization that is orthogonal to the horizontal/vertical dimensions used by current unlexicalized parsers (Johnson 1998, Klein and Manning 2003) and its contribution is essential and complementary. Our best model uses three dimensions of parameterization and our best results are on a par with those of a fully lexicalized parser for Modern Standard Arabic trained on a much larger treebank.

Saib Mansour, Khalil Sima'an and Yoad Winter. *Smoothing a Lexicon-based POS Tagger for Arabic and Hebrew*

We propose an enhanced Part-of-Speech (POS) tagger of Semitic languages that treats Modern Standard Arabic (henceforth Arabic) and Modern Hebrew (henceforth Hebrew) using the same probabilistic model and architectural setting. We start out by porting an existing Hidden Markov Model POS tagger for Hebrew to Arabic by exchanging a morphological analyzer for Hebrew with Buckwalter's (2002) morphological analyzer for Arabic. This gives state-of-the-art accuracy (96.12%), comparable to Habash and Rambow's (2005) analyzer-based POS tagger on the same Arabic datasets. However, further improvement of such analyzer-based tagging methods is hindered by the incomplete coverage of standard morphological analyzer (Bar Haim et al., 2005). To overcome this coverage problem we supplement the output of Buckwalter's analyzer with synthetically constructed analyses that are proposed by a model which uses character information (Diab et al., 2004) in a way that is similar to Nakagawa's (2004) system for Chinese and Japanese. A version of this extended model that (unlike Nakagawa) incorporates synthetically constructed analyses also for known words achieves 96.28% accuracy on the standard Arabic test set.

Joshua Waxman. *Transliteration Bitext Corpus Generation and Cross-Language Information Retrieval Using Minimum Edit Distance and Character Cluster Equivalence Classes*

We use a modified Levenshtein minimum edit distance algorithm for cross-language fuzzy string matching in order to automatically pair transliterations of Hebrew and Aramaic phrases from the Babylonian Talmud with the original Hebrew and Aramaic phrases. Thus, we create a bitext transliteration corpus of Hebrew and Aramaic which can be used to test the accuracy of cross-language information retrieval, and for statistical training of transliteration models. We customize the algorithm for the specific type of data and search at hand. Finally, we propose a modification of the standard minimum edit distance algorithm to allow mappings of characters in alphabet A to clusters of characters in alphabet B, thus allowing for such mappings as the Hebrew letter "shin" to "sh" (for Hebrew/Aramaic) or "sch," the Hebrew letter "yud" to "ei" (for Hebrew/Aramaic), and English "j" to "daled zayin shin" (for Yiddish), all while maintaining the dynamic programming character of the standard implementation of minimum edit distance.

11:20-13:00 Session Robotics 2

Yan Virin, Guy Shani and Solomon Eyal Shimony. *Scaling Up: Solving POMDPs through Value Based Clustering*

Partially Observable Markov Decision Processes (POMDPs) provide a rich model for modeling agents operating under partial knowledge of the environment. Since finding an optimal POMDP policy is intractable, approximation techniques have been a main focus of research, among them point-based algorithms, which scale up relatively well - up to thousands of states. An important decision in a point-based algorithm is the order of backup operations over belief states. Prioritization techniques for ordering the sequence of backup operations reduce the number of needed backups considerably but involve significant overhead. In this paper we suggest a new way to order backups, based on a soft

clustering of the belief space. Our novel soft clustering method relies on the solution of the underlying MDP. Empirical evaluation verifies that our method rapidly computes a good order of backups, showing orders of magnitude improvement in runtime over a number of benchmarks.

Eliyahu Osherovich, Vladimir Yanovski, Israel Wagner and Freddy Bruckstein. *Robust and Efficient Covering of Unknown Continuous Domains with Simple, Ant-Like A(gents)*

We propose a new "Mark-Ant-Walk" algorithm for robust and efficient covering of continuous domains by ant-like robots with very limited capabilities. The robots can mark places visited with pheromone marks and sense the level of the pheromone in their local neighborhood. In case of multiple robots these pheromone marks can be sensed by all robots and provide the only way of (indirect) communication between the robots. The robots are assumed to be memory-less, and to have no information besides that mentioned above. Despite the robots' simplicity, we show that they are able, by running a very simple rule of behavior, to ensure efficient covering of arbitrary connected domains, including non-planar and multidimensional ones. The novelty of our algorithm lies in the fact that, unlike previously proposed methods, our algorithm works on continuous domains without relying on some "induced" underlying graph that effectively reduces the problem to a discrete case of graph covering. In addition we demonstrate various benefits of our successive visits of the robot at the same location that makes it suitable for patrolling. Finally we provide a new theoretical bound on covering time of a wide family of such mark and walk algorithms.

Ami Berler and Solomon Eyal Shimony. *Mapping AWOL Using a Concrete Domain*

Collaboration of several intelligent agents on a shared task is a complex research issue. Even when each agent seeks to optimize global team utility, the environment may force decisions to be decentralized due to limited communication and uncertainty about the environment. As our domains are stochastic, and not completely observable, we use standard methods based on Markov assumptions in order to model these problems. A typical scheme is modeling the decision problem as a Decentralized Markov Decision Process (DEC-MDP), to which finding an optimal policy is extremely intractable. Making special assumptions about the domain, e.g. independence in transitions or observations, is one typical scheme to overcome the problem. Such assumptions, however, frequently do not conform to realistic environments. Another approach is to reduce the state-space and action space by using a plan-space abstraction, solve the abstract problem, in order to achieve a (hopefully) near-optimal policy. The type of abstraction we use to reduce the state space the domain was proposed in early papers: the Abstract World for Opportunistic Local decisions (AWOL) framework. In order to evaluate the model empirically, we show how the model is used in the context of a simulated environment: a discrete, two dimensional, simplified version of the Unreal Tournament (TM) computer game. As a comparative experiment, we are implementing four different methodologies to solve the multi-agent decision problem. We compare the results of these methods w.r.t. computation time for the solution, and expected utility.

Ariel Felner, Roni Stern, Zahi Benaya and Tal Fadida. *Problem Solving by Moving Agents in Physical Graphs*

There are many known algorithms to graph problems. We examine a general framework for implementing these algorithms on physical graphs where a traveling agent needs to physically explore relevant portions of the graph. We then discuss how this framework can be generalized to the case where a number of agents cooperate in solving these problems. We examine theoretically and experimentally several communication paradigms and several overall cost functions and suggest different way to handle these aspects based on our general framework.

14:15-15:30 Keynote 1

Keynote: Yaacov Choueka. *You've come a long way, Baby! Responsa meet Full-text and CL, one-thousand-year manuscripts meet AI*

The Responsa Project, a full-text retrieval system for the Responsa and other major corpora of Rabbinical literature, was recently awarded the prestigious Israel Price for 2007. Yaacov Choueka, Professor Emeritus of Computer Science and one of the two main leaders and developers of the project, will present interesting highlights from the project history, including research problems confronted by the developers well before they were studied by the research community at large, and will conclude with some hints about AI problems related to a new project on Computers and Jewish Studies - The Friedberg Genizah Project - he is now directing.

15:30-16:10 Session ISCOL 3

Dmitry Davidov, Ari Rappoport and Moshe Koppel. *Fully Unsupervised Discovery of Concept-Specific Relationships by Web Mining*

We present a web mining method for discovering and enhancing relationships in which a specified concept participates. We discover a whole range of relationships focused on the given concept, rather than generic known relationships as most previous work. Our method is based on clustering patterns that contain concept words and other words related to them. We evaluate the method on three different rich concepts and find that in each case the method generates a broad variety of relationships with good precision.

Idan Szpektor and Ido Dagan. *Learning Canonical Forms of Entailment Rules*

We propose a modular approach to paraphrase and entailment-rule learning that addresses the morpho-syntactic variability of lexical-syntactic templates. Using an entailment module that captures generic morpho-syntactic regularities, we transform every identified template into a canonical form. This way, statistics from different template variations are accumulated for a single template form. Additionally, morpho-syntactic redundant rules are not acquired. This scheme also yields more informative evaluation for acquisition quality, since bias towards rules with many frequent variations is avoided.

15:30-16:10 Session AI 1 (Clustering)

Maxim Binshtok, Ronen Brafman, Solomon Eyal Shimony, Ajay Mani and Craig Boutilier. *Computing Optimal Subsets*

Various tasks in decision making and decision support require selecting a preferred subset of items from a given set of feasible items. Recent work in this area considered methods for specifying such preferences based on the attribute values of individual elements within the set. Of these, the approach of Brafman et. al. 06 appears to be the most general. In this paper, we consider the problem of computing an optimal subset given such a specification. The problem is shown to be NP-hard in the general case, necessitating heuristic search methods. We consider two algorithm classes for this problem: direct set construction, and implicit enumeration as solutions to appropriate CSPs. New algorithms are presented in each class and compared empirically against previous results. The paper was accepted for AAAI-07 conference both as oral and poster presentation.

Rachel Ben-Eliyahu-Zohary, Ran Giladi, Philip Hendrix and Stuart M. Shieber. *Clustering Ad-Hoc Networks: Experiments in Local Search*

Partitioning a wireless ad-hoc network into clusters is a known method for achieving reduction in resource consumption. In this paper we use local search in order to divide the network into k -diameter clusters - clusters in which the distance between any given two nodes is at most k . This problem has application also in multi-agent systems, where we have some kind of a communication network among the agents and we want to divide them into groups in a way that reduces communication cost. Two methods are used in order to evaluate the quality of the local search. The first method compares the results of the local search with the results of a distributed algorithm having a guaranteed upper bound of the optimal clustering. The second method tests the local search on graphs for which the optimal clustering is known in advance. Both tests show that the local search approach is superior to the distributed algorithm for a large family of graphs.

16:40-18:00 Session ISCOL 4

Roi Reichart and Ari Rappoport. *Self-Training for Enhancement and Domain Adaptation of Statistical Parsers Trained on Small Datasets*

Creating large amounts of annotated data to train statistical PCFG parsers is expensive, and the performance of such parsers declines when training and test data are taken from different domains. In this paper we use self-training in order to improve the quality of a parser and to adapt it to a different domain, using only small amounts of manually annotated seed data. We report significant improvement both when the seed and test data are in the same domain and in the out-of-domain adaptation scenario. In particular, we achieve 50% reduction in annotation cost for the in-domain case, yielding an improvement of 66% over previous work, and a 20-33% reduction for the domain adaptation case. This is the first time that self-training with small labeled datasets is applied successfully to these tasks. We were also able to formulate a characterization of when self-training is valuable.

Ido Dagan, Oren Glickman, Alfio Massimiliano Gliozzo, Efrat Hershkoviz Marmorshtein and Carlo Strapparava. *Direct Word Sense Matching for Lexical Substitution*

This paper investigates conceptually and empirically the novel sense matching task, which requires to recognize whether the senses of two synonymous words match in context. We suggest direct approaches to the problem, which avoid the intermediate step of explicit word sense disambiguation, and demonstrate their appealing advantages and stimulating potential for future research.

Yael Netzer, Meni Adler, David Gabay and Michael Elhadad. *Can You Tag the Modal? You Should*

Computational linguistics methods are typically first developed and tested in English. When applied to other languages, s from English data are often applied to the target language. One of the most common such assumption is that a “standard” part-of-speech (POS) tagset can be used across languages with only slight variations. We discuss in this paper a specific issue related to the definition of a POS tagset for Modern Hebrew, as an example to clarify the method through which such variations can be defined. It is widely assumed that Hebrew has no syntactic category of modals. There is, however, an identified class of words which are modal-like in their semantics, and can be characterized through distinct syntactic and morphologic criteria.

We have found wide disagreement among traditional dictionaries on the POS tag attributed to such words. We describe three main approaches when deciding how to tag such words in Hebrew. We illustrate the impact of selecting each of these approaches on agreement among human taggers, and on the accuracy of automatic POS taggers induced for each method. We finally recommend the use of a "modal" tag in Hebrew and provide detailed guidelines for this tag. Our overall conclusion is that tagset definition is a complex task which deserves appropriate methodology.

Judit Bar-Ilan, Kevin Keenoy, Mark Levene and Eti Yaari. *Ranking Preferences of Search Results – A User Study*

We describe the results of an experiment designed to study user preferences for different orderings of search results from three major search engines. Users were asked to choose the best ordering from two different orderings of the same search results: each pair consisted of the search engine’s original top-ten ordering and a synthetic ordering created from the same top-ten results retrieved by the search engine. This process was repeated for different queries and different synthetic orderings. The results show that there is a slight overall preference for the search engines’ original orderings, but the preference is rarely significant. Users choice of the “best” result from each of the different orderings, seems to indicate that placement on the page (i.e. whether the result appears near the top) is the most important factor used in determining the quality of the result, and not the actual content of the displayed snippet.

16:40-18:00 Session AI 2 (Search)

Roie Zivan, Moshe Zazone and Amnon Meisels. *Min-Domain retroactive ordering for Asynchronous Backtracking*

A new type of ordering heuristics for dynamic ordering asynchronous backtracking (ABT_DO) on DisCSPs is presented. Agents can be moved to a position that is higher than that of the target of the backtrack (culprit). This new type of heuristics does not follow the restrictions on the heuristics of previous versions of dynamically ordered ABT.

The degree of the retroactivity of heuristics is dependent upon the size of additional Nogood storage that agents are allowed to keep. This size is defined by a parameter k which limits the size of stored Nogoods.

The performance of the retroactive ordered ABT is found to be worse when larger Nogood storage is used. The best performing version is the one that uses a specific min-domain heuristic with no additional storage of Nogoods. The best version of retroactive ordered ABT is faster by a factor of 2 than the best heuristic of dynamically ordered ABT.

Ariel Felner and Nir Ofek. *Combining Perimeter Search and Pattern Database Abstractions*

A pattern database abstraction (PDB) is a heuristic function in a form of a lookup table. A PDB stores the cost of optimal solutions for instances of abstract problems (sub-problems). These costs are used as admissible heuristics for the original problem. Perimeter search (PS) is a form of bidirectional search. First, a breadth-first search is performed backwards from the goal state. Then, a forward search is executed towards the nodes of the perimeter.

In this paper we study the effect of combining these two techniques. We describe two methods for

doing this. The simplified method uses a regular PDB (towards a single goal state) but uses the perimeter to correct heuristics of nodes outside the perimeter. The second, more advanced method is to build a PDB that stores the cost of reaching any node of the perimeter from a given pattern. Although one might see great potential for speedup in the advanced method, we theoretically show that surprisingly most of the benefit of combining perimeter and PDBs is already exploited by the first method. We also provide experimental results that confirm our findings. We then study the behavior of our new approach when combined with methods for using multiple PDBs such as maxing and adding.

Dragan Bosnacki, Edith Elkind, Blaise Genest and Doron Peled. *On Commutativity Based Edge Lean Search*

The problem of **state space search** is fundamental to artificial intelligence. Often, the state space is huge, so optimizing the search may be crucial. We consider the problem of visiting all states in a graph where edges are generated by actions and the (reachable) states are not known in advance. Some of the actions may commute, i.e., they result in the same state for every order in which they are taken. We show how to use commutativity to achieve full coverage of the states while traversing considerably fewer edges.

Uzi Zahavi and Ariel Felner. *Inconsistent Heuristics*

In the field of heuristic search it is well-known that improving the quality of an admissible heuristic can significantly decrease the search effort required to find an optimal solution. Existing literature often assumes that admissible heuristics are consistent, implying that consistency is a desirable attribute. To the contrary, this paper shows that an inconsistent heuristic can be preferable to a consistent heuristic. Theoretical and empirical results show that, in many cases, inconsistency can be used to achieve large performance improvements. The conventional wisdom about inconsistent heuristics is wrong.

תמישי, ה' תמוז, תשס"ז Thursday, June 21, 2007

09:00-10:45 Session AI 3 (Machine Learning)

Ariel Procaccia, Aviv Zohar, Yoni Peleg and Jeffrey Rosenschein. *Learning Voting Trees*

Binary voting trees provide a succinct representation for a large and prominent class of voting rules. In this paper, we investigate the PAC-learnability of this class of rules. We show that, while in general a learning algorithm would require an exponential number of samples, if the number of leaves is polynomial in the size of the set of alternatives then a polynomial training set suffices. We apply these results in an emerging theory: automated design of voting rules by learning.

Ron Katz and Sarit Kraus. *Gender-Sensitive Automated Negotiators*

This paper introduces an innovative approach for automated negotiating using the gender of human opponents. Our approach segments the information acquired from previous opponents, stores it in two databases, and models the typical behavior of males and of females. The two models are used in order to match an optimal strategy to each of the two sub-populations. In addition to the basic separation, we propose a learning algorithm which supplies an online indicator for the gender separability-level of the population, which tunes the level of separation the algorithm activates. The algorithm we present can be generally applied in different environments with no need for configuration of parameters. Experiments in 4 different one-shot domains, comparing the performance of the gender based separation approach with a basic approach which is not gender sensitive, revealed higher layoffs of the former in almost all the domains. Moreover, using the proposed learning algorithm further improved the results.

Roy Fox and Moshe Tennenholtz. *A Reinforcement Learning Algorithm with Polynomial Interaction Complexity for Only-Costly-Observable MDPs*

An Unobservable MDP (UMDP) is a POMDP in which there are no observations. An Only-Costly-Observable MDP (OCOMDP) is a POMDP which extends an UMDP by allowing a particular costly action which completely observes the state. We introduce UR-MAX, a reinforcement learning algorithm with polynomial interaction complexity for unknown OCOMDPs.

Larry Manevitz and Hananel Hazan. *History-Dependent Neurons and Identification of Temporal Sequences*

We describe a new kind of abstract artificial neuron, called here MM-neurons, based on the relatively recent discovery that refractory periods and other constants in biological neurons are not constants but are history dependent. We then show that, in principle, such neurons can be used for temporal detection and classification; and that in a natural way it leads to the transfer of rate codes to temporal ones. We examine this in the most simple situations (McCullough Pitts with history dependent refractory periods) where we see that the system forms clusters of periodic input, and these clusters are robust to the architecture of the network. We also indicate future directions of research where on the one hand this phenomenon can be combined with the "usual" discriminatory ability of standard neural networks, and on the other hand can be combined with more temporal networks (like pulsed neurons).

Alon Altman, Avivit Bercovici-Boden and Moshe Tennenholtz. *Learning in one-shot strategic form games*

We propose a machine learning approach to action prediction in one-shot games. In contrast to the huge literature on learning in games where an agent's model is deduced from its previous actions in a multi-stage game, we propose the idea of inferring correlations between agents' actions in different one-shot games in order to predict an agent's action in a game which she did not play yet. We define the approach and show, using real data obtained in experiments with human subjects, the feasibility of this approach. Furthermore, we demonstrate that this method can be used to increase payoffs of an adequately informed agent. This is, to the best of our knowledge, the first proposed and tested approach for learning in one-shot games, which is the most basic form of multiagent interaction.

11:15-12:15 Keynote 2

Keynote: Manuela M. Veloso. *Selective Use of Multiple Sources of Robot Sensory Information*

An autonomous robot needs to assess the state of the environment, make decisions towards achieving its goals, and execute the selected actions. In teams of autonomous robots, robots have individually limited perception, but can communicate state information to each other creating therefore multiple perceptual inputs. We present an algorithm for selectively merging and using the own perceptual data combined with the communicated data from teammate robots. In general, robots face the challenge of combining multiple sources of sensory information. We illustrate different concrete instances of this problem and discuss a prioritized approach to effectively merge multi-modal information. The talk will be organized as an explanation of the algorithms underlying a series of different robot videos, including robot soccer players, humanoid robot soccer commentators, and machine visual object recognition.

12:15-13:00 Session AI 4 (Game Complexity)

Yoram Bachrach & Jeffrey Rosenschein. *Computing the Banzhaf Power Index in Network Flow Games*

Preference aggregation is used in a variety of multiagent applications, and as a result, voting theory has become an important topic in multiagent system research. However, power indices (which reflect how much "real power" a voter has in a weighted voting system) have received relatively little attention, although they have long been studied in political science and economics. The Banzhaf power index is one of the most popular; it is also well-defined for any simple coalitional game.

In this paper, we examine the computational complexity of calculating the Banzhaf power index within a particular multiagent domain, a network flow game. Agents control the edges of a graph; a coalition wins if it can send a flow of a given size from a source vertex to a target vertex. The relative power of each edge/agent reflects its significance in enabling such a flow, and in real-world networks could be used, for example, to allocate resources for maintaining parts of the network. We show that calculating the Banzhaf power index of each agent in this network flow domain is $\#P$ -complete. We also show that for some restricted network flow domains there exists a polynomial algorithm to calculate agents' Banzhaf power indices.

Edith Elkind, Leslie Ann Goldberg, Paul Goldberg and Michael Wooldridge. *Computational Complexity of Weighted Threshold Games*

Weighted threshold games are coalitional games in which each player has a weight (intuitively corresponding to its voting power), and a coalition is successful if the sum of its weights exceeds a given threshold. Key questions in coalitional games include finding coalitions that are stable (in the sense that no member of the coalition has any rational incentive to leave it), and finding a division of payoffs to coalition members (an imputation) that is fair. We investigate the computational complexity of such questions for weighted threshold games. We study the core, the least core, and the nucleolus, distinguishing those problems that are polynomial-time computable from those that are NP-hard, and providing pseudo-polynomial and approximation algorithms for the NP-hard problems.

14:00-15:45 Session AI 5 (Planning and Plan Recognition)

Michael Katz and Carmel Domshlak. *Structural Patterns of Tractable Sequentially-Optimal Planning*

We study the complexity of sequentially-optimal classical planning, and discover new problem classes for whose such optimization is tractable. The results are based on exploiting numerous structural characteristics of planning problems, and a constructive proof technique that connects between certain tools from planning and tractable constraint optimization. In particular, we believe that structure-based tractability results of this kind may help devising new admissible search heuristics. We discuss the prospects of this direction along a principled extension of pattern-database heuristics to "structural patterns" of unlimited dimensionality.

Carmel Domshlak and Vitaly Mirkis. *Cost-Sharing Approximations for h^+*

Relaxations based on (either complete or partial) ignoring delete effects of the actions provide the basis for some seminal classical planning heuristics. However, the palette of the conceptual tools exploited by these heuristics remains rather limited. We study a framework for approximating the optimal cost solutions for problems with no delete effects that bridge between certain works on heuristic search for probabilistic reasoning and classical planning. In particular, this framework generalizes some previously known, as well as suggests some novel, tools for heuristic estimates for Strips planning.

Hannaneh Hajishirzi and Eyal Amir. *Stochastic Filtering in a Probabilistic Action Model*

Stochastic filtering is the problem of estimating the state of a dynamic system after time passes and given partial observations. It is fundamental to automatic tracking, planning, and control of real-world stochastic systems such as robots, programs, and autonomous agents. This paper presents a novel sampling-based filtering algorithm. Its expected error is smaller than sequential Monte Carlo sampling techniques given a fixed number of samples, as we prove and show empirically. It does so by sampling deterministic action sequences and then performing exact filtering on those sequences. These results are promising for applications in stochastic planning, natural language processing, and robot control.

Dorit Avrahami-Zilberbrand and Gal Kaminka. *Incorporating Observer Biases in Keyhole Plan Recognition (Efficiently!)*

Plan recognition is the process of inferring other agents' plans and goals based on their observable actions. Essentially all previous work in plan recognition has focused on the recognition process itself, with no regard to the use of the information in the recognizing agent. As a result, low-likelihood recognition hypotheses that may imply significant meaning to the observer, are ignored in existing work. In this paper, we present novel efficient algorithms that allow the observer to incorporate her own biases and preferences—in the form of a utility function—into the plan recognition process. This allows choosing recognition hypotheses based on their expected utility to the observer. We call this Utility-based Plan Recognition (UPR). While reasoning about such expected utilities is intractable in the general case, we present a hybrid symbolic/decision-theoretic plan recognizer, whose complexity is $O(NDT)$, where N is the plan library size, D is the depth of the library and T is the number of observations. We demonstrate the efficacy of this approach with experimental results in several challenging recognition tasks.

Nirom Cohen-Nov Slapak and Gal Kaminka. *Computational Load and Performance in Integrated Multi-Agent Intention Recognition*

Intention recognition is the ability to reason and infer information about others based on observations of their behavior. Unfortunately, to date, there have been only a handful of investigations into the integration of intention recognition into agent architectures. In particular, there are open questions as to the effect of intention recognition on the computational resources available to the agent. The issue is of particular importance in modern virtual environments, where the agent may be interacting with multiple other agents. This paper tackles this question empirically, taking first steps to identify heuristics that may be used to allow a recognition process to keep track of several agents in parallel, while minimizing the computational costs of such tracking.

16:15-17:15 Keynote 3

Keynote: Dana S. Nau. *May All Your Plans Succeed! (Or Have a High Expected Utility)*

Automated planning technology has become mature enough to be useful in applications that range from game-playing to control of space vehicles. In this talk, Dr. Nau will discuss where automated-planning research has been, where it is likely to go, where he thinks it **should** go, and some major challenges. The presentation is an updated version of Dr. Nau's invited talk at AAAI-05.

17:15-18:15 Session AI 6 (Modelling)

Shlomo Berkovsky, Ariel I. Gorfinkel, Tsvi Kuflik and Larry M. Manevitz. *Case-Based to Content-Based User Model Mediation and its Effectiveness*

Systems providing personalized services to users have a need to build and maintain a User Model (UM). One important user modeling task is initialization; moreover, for some systems this task is crucial, yet no user data from prior interactions is available. An example is museum systems, in which both the repeat use-age by a user is rare, and the total interaction on a single use is relatively limited. Due to lack of standards in the representation of UMs, commercial competition, and privacy issues, distinct personalized service-providing systems build their own specific models and store their information in incompatible manners. Thus, despite the fact that plenty of data on a specific user might exist in other systems, it is typically unavailable for use in the initial phase of the given system. This work puts forward the design of a user model mediation idea. This is demonstrated in an implementation of a specific system (Museum Visitors' guide system) under the PIL project, where the UM is Content-based and the initial information is imported from an external trip planning system that

uses a Case-based UM representation. This work also puts forward an evaluation method for the effectiveness of user modeling adaptation and initialization techniques. This evaluation method is based on simulation. By suggesting a way to define a "Gold standard" it allows comparing the effectiveness of different techniques, specifically comparing between the mediation technique and other UM initialization techniques.

Guy Shani. *A Stereotypes-Based Hybrid Recommender System for Media Items*

Many Recommender Systems use either Collaborative Filtering (CF) or Content-Based (CB) techniques to receive recommendations for products. Both approaches have advantages and weaknesses. Combining the two approaches together can overcome most weaknesses. However, most hybrid systems combine the two methods in an ad-hoc manner.

In this paper we present a hybrid approach for recommendations, where a user profile is a weighted combination of user stereotypes, created automatically through a clustering process. Each stereotype is defined by an ontology of item attributes. Our approach provides good recommendations for items that were rated in the past and is also able to handle new items that were never observed by the system. Our algorithm is implemented in a commercial system for recommending media items. The system is envisioned to function as personalized media (audio, video, print) service within mobile phones, online media portals, sling boxes, etc. It is currently under development within Deutsche Telekom Laboratories - Innovations of Integrated Communication projects.

Orna Peleg, Zohar Eviatar, Larry M. Manevitz and Hananel Hazan. *Differences and Interactions between Cerebral Hemispheres When Processing Ambiguous Homographs*

It is well known that the brain (especially the cortex) is structurally separable into two Hemispheres. Many neuropsychological studies show that the process of ambiguity resolution requires the intact functioning of both cerebral hemispheres. Moreover, these studies suggest that while the LH quickly selects one alternative, the RH maintains alternate meanings. However, these hemi-spheres are connected through the corpus callosum and presumably the exchange of information is useful. In addition, many works show that the Left Hemisphere (LH) is more influenced by the phonological aspect of written words whereas lexical processing in the Right Hemisphere (RH) is more sensitive to visual form. This distinction suggests that the interconnections between the hemispheres may be used to strengthen or correct incorrect interpretations by one hemisphere. We test this hypothesis by (i) postulating that in the Left Hemisphere (LH) orthography, phonology and semantics are interconnected while (ii) the Right Hemisphere (RH), phonology is not connected directly to orthography and hence its influence must be mitigated by semantical processing (iii) seeing if corrections in ambiguous word processing can be aided by information in the other hemisphere. We investigate this by complementary human psychophysical experiments and by dual (one RH, one LH) computational neural network model architecturally modified from [14] Kawamoto's [1993] model to follow our hypothesis. Since the different models have different rates of convergence, we test (iii) by halting processing, and using an analogue to priming to compare the rate of convergence to a corrected semantics in the LH working alone and working with information obtained from the RH at the same point in processing. In this paper we present results of the computational model and show that (i) results obtained from the two hemispheres separately are analogous to the human experiments and (ii) use of the RH information does in-deed help such corrections.

שישי, ו' תמוז, תשס"ז Friday, June 22, 2007

09:00-10:45 Session AI 7 (Multi-Agent Systems I)

Avi Rosenfeld, Sarit Kraus and Charlie Ortiz. *Quantifying the Expected Utility of Information in Multi-Agent Scheduling Tasks*

In this paper we investigate methods for analyzing the expected value of adding information in distributed task scheduling problems. As scheduling problems are NP-complete, no polynomial algorithms exist for evaluating the impact a certain constraint, or relaxing the same constraint, will have on the global problem. We present a general approach where local agents can estimate their problem tightness, or how constrained their local sub-problem is. This allows these agents to immediately identify many problems which are not constrained, and will not benefit from sending or receiving further information. Next, agents use traditional machine learning methods based on their specific local problem attributes to attempt to identify which of the constrained problems will most benefit from human attention. We evaluated this approach within a distributed cTAEMS scheduling domain and found this approach was overall quite effective.

Meir Kalech, Gal Kaminka and Michael Lindner. *Matrix-Based Representation for Coordination Fault Detection: A Formal Approach*

Teamwork requires that team members coordinate their actions. The representation of the coordination is a key requirement since it influences the complexity and flexibility of reasoning team-members. One aspect of this requirement is detecting coordination faults as a result of intermittent failures of sensors, communication failures, etc. Detection of such failures, based on observations of the behavior of agents, is of prime importance. Though different solutions have been presented thus far, none has presented a comprehensive and formal resolution to this problem.

This paper presents a formal approach to representing multi-agent coordination, and multi-agent observations, using matrix structures. This representation facilitates easy representation of coordination requirements, modularity, flexibility and reuse of existing systems. Based on this representation we present a novel solution for fault-detection that is both generic and efficient for large-scale teams.

Alex Rogers, Esther David, Terry Payne and Nicholas R. Jennings. *An Advanced Bidding Agent for Advertisement Selection on Public Displays*

In this paper we present an advanced bidding agent that participates in first-price sealed bid auctions to allocate advertising space on BluScreen – an experimental public advertisement system that detects users through the presence of their Bluetooth enabled devices. Our bidding agent is able to build probabilistic models of both the behavior of users who view the adverts, and the auctions that it participates within. It then uses these models to maximize the exposure that its adverts receive. We evaluate the effectiveness of this bidding agent through simulation against a range of alternative selection mechanisms including a simple bidding strategy, random allocation, and a centralized optimal allocation with perfect foresight. Our bidding agent significantly outperforms both the simple bidding strategy and the random allocation, and in a mixed population of agents it is able to expose its adverts to 25% more users than the simple bidding strategy. Moreover, its performance is within 7.5% of that of the centralized optimal allocation despite the highly uncertain environment in which it must operate.

Ido Levy and Claudia Goldman. *Distributed On-line Learning Cooperative Multi-Agent System*

This paper presents a distributed version of the Q-learning algorithm that combines communication efficiently to speed up the learning process of a group of agents acting together. We have modeled and simulated a distributed system with on-line learning cooperative agents that have the ability to communicate in order to help each other. We discuss the trade off between communication cost and agent coordination and show that in most cases, depending on the communication cost, we will have better system utilization.

We focus on how to shorten the length of the distributed learning process by specifying the agent's decision process that can be integrated with any on-line learning procedure. We define the agent-to-agent communication and information exchange model in a collaborative learning environment. While there is no proof of convergence for Q-learning in distributed environments, we suggest a practical method that tries to improve the utilization of such systems by speeding up the learning process.

Avi Rosenfeld, Claudia Goldman, Sarit Kraus and Gal Kaminka. *An Agent Architecture for Hybrid P2P Free-Text Search*

Recent advances in peer to peer (P2P) search algorithms have presented viable structured and unstructured approaches for full-text search. We posit that these existing approaches are each best suited for different types of queries. We present PHIRST, the first system to facilitate effective full-text search within P2P networks. PHIRST works by effectively leveraging between the relative strengths of these approaches. Similar to structured approaches, agents first publish terms within their stored documents. However, frequent terms are quickly identified and not exhaustively stored, resulting in a significantly reduction in the system's storage requirements. During query lookup, agents use unstructured searches to compensate for the lack of fully published terms. Additionally, they explicitly weigh between the costs involved with structured and unstructured approaches, allowing for a significant reduction in query costs. We evaluated the effectiveness of our approach using both real-world and artificial queries. We found that in most situations our approach yields near perfect recall. We discuss the limitations of our system, as well as possible compensatory strategies.

11:15-12:15 Keynote 4

Keynote: Milind Tambe. *Multiagent and Agent-human Teamwork: Hybrid Approaches*

How do we build multiagent systems? Today, within the agents and multiagent systems community, we see four competing approaches: logic-based belief-desire-intention (BDI), decision-theory and its incarnation in distributed Markov decision problems (distributed MDPs or POMDPs), distributed constraint optimization and finally, auctions or game-theoretic approaches. While there is exciting progress within each approach, there is a lack of cross-cutting research.

This talk will outline a case for hybrid approaches as we scale to complex multiagent domains. In particular, for the past decade, the TEAMCORE research group has focused on building agent teams in complex, dynamic domains. While our early work was inspired by BDI, we will present an overview of recent research that uses DCOPs and distributed POMDPs in building agent teams. While hybrids of DCOP and distributed POMDPs have helped us address problems of efficiency and expressiveness, game-theoretic considerations have helped build a new family of DCOP algorithms. Finally, in our BDI-POMDP hybrid approaches, BDI team plans are exploited to improve POMDP tractability, and POMDPs improve BDI team plan performance. We present some recent results from applying these hybrids to complex domains such as teams of software personal assistants for office environments and agent teams for disaster rescue simulations.

12:15-13:00 Session AI 8 (Game Theory)

Aviv Zohar, Hagay Levin and Michael Schapira. *Incentive-Compatible Distributed Routing*

We present a novel game-theoretic model that captures many of the intricacies of inter-domain routing in today's Internet. In this model, the strategic agents are source nodes located on a network, who aim to send traffic to a unique destination node. The interaction between the agents is dynamic and complex – asynchronous, sequential, and based on partial information. We consider a simple distributed protocol in which all agents are instructed to behave myopically. We show that in realistic and well-studied settings, this protocol is immune to rational manipulations. I.e., not only does myopic behavior of all players converge to a "stable" routing outcome, but no player has motivation to unilaterally deviate from the protocol. Moreover, we show that even coalitions of players of any size cannot improve their routing outcomes by collaborating. The only protocol used for inter-domain routing, namely the Border Gateway Protocol (BGP), is based on such myopic interaction. Hence, from a practical perspective, our results imply a strategic justification for BGP by showing that it possesses extremely strong game-theoretic properties.

From a theoretical perspective, the strengths of our results lie in the following facts: Firstly, unlike the vast majority of works in mechanism design, our results do not require any monetary transfers (to or by the agents). Secondly, we allow significantly more complex preferences over routes than previous works. Finally, our results are a first step in the exploration of myopic behavior in Algorithmic Mechanism Design.

Michal Penn, Maria Polukarov and Moshe Tennenholtz. *Congestion Games with Load-Dependent Failures: Identical Resources*

We define a new class of games, congestion games with load-dependent failures (CGLFs), which generalizes the well-known class of congestion games, by incorporating the issue of resource failures into congestion games. In a CGLF, agents share a common set of resources, where each resource has a cost and a probability of failure. Each agent chooses a subset of the resources for the execution of his task, in order to maximize his own utility. The utility of an agent is the difference between his benefit from successful task completion and the sum of the costs over the resources he uses. CGLFs possess two novel features. It is the first model to incorporate failures into congestion settings, which results in a strict generalization of congestion games. In addition, it is the first model to consider load-dependent failures in such framework, where the failure probability of each resource depends on the number of agents selecting this resource. Although, as we show, CGLFs do not admit a potential function, and in general do not have a pure strategy Nash equilibrium, our main theorem proves the existence of a pure strategy Nash equilibrium in every CGLF with identical resources and non-decreasing costfunctions.

14:00-14:40 Session AI 9 (Multi-Agent Systems II)

Natalie Fridman and Gal Kaminka. *Towards a Cognitive Model of Crowd Behavior Based on Social Comparison Theory*

Models of crowd behavior facilitate analysis and prediction of human group behavior, where people are affected by each other's presence. Unfortunately, existing models leave many open challenges. In particular, psychology models often offer only qualitative description, while computer science models are often simplistic, and are not reusable from one simulated phenomenon to the next. We propose a novel model of crowd behavior, based on Festinger's Social Comparison Theory (SCT). We propose a concrete algorithmic framework for SCT, and evaluate its implementation in several crowd behavior scenarios. Results from task measures and human judges evaluation shows that the SCT model produces improved results compared to base models from the literature.

Noam Hazon, Paul E. Dunne, Sarit Kraus and Michael Wooldridge. *How to Rig an Election*

We investigate the extent to which it is possible to rig the agenda of an election or competition so as to favor a particular candidate in the presence of imperfect information about the preferences of the electorate. We assume that what is known about an electorate is the probability that any given candidate will beat another. As well as presenting some analytical results relating to the complexity of finding and verifying agenda, we develop heuristics for agenda rigging, and investigate the performance of these heuristics for both randomly generated data and real-world data from tennis and basketball competitions.

14:40-15:40 Keynote 5

Keynote: Michael Wooldridge. *Logic for Automated Mechanism Design – A Progress Report*

Over the past few years, strategy logics such as Coalition Logic and ATL have proved to be a powerful formal tool for representing and reasoning about the properties of game-like mechanisms such as social choice procedures. These logics share the same technical heritage as CTL-like temporal logics, which have proved to be amenable to automation, and have been shown to be extremely useful in the analysis and verification of reactive systems via model checking. We believe that in exactly the same way, model checking tools for strategic logics will prove to be of similar value in the automated analysis and verification of social choice procedures. In this talk, we present a survey of our work in this area. We begin by introducing ATL-like logics, and demonstrating that they form a natural tool for the specification of social choice procedures. We show how ATL model checkers can be used to verify economic properties of social choice mechanisms. We then discuss the main issues required to turn this vision into a reality, focusing on issues such as the succinct representation of social choice rules, the complexity of reasoning with such representations, and the handling of preferences.

This talk will report joint work with Thomas Agotnes (Bergen), Wiebe van der Hoek (Liverpool), Marc Pauly (Stanford), and Paul E. Dunne (Liverpool).