Report from Dagstuhl Seminar 11141

Plan Recognition

Edited by

Robert P. Goldman[1], Christopher W. Geib[2], Henry Kautz[3], and Tamim Asfour[4]

1 **SIFT, LLC – Minneapolis, US,** rpgoldman@sift.net
2 **University of Edinburgh, GB,** cgeib@inf.ed.ac.uk
3 **University of Rochester, US,** kautz@cs.rochester.edu
4 **KIT – Karlsruher Institut für Technologie, DE,** asfour@kit.edu

Abstract

This Dagstuhl seminar brought together researchers with a wide range of interests and backgrounds related to plan and activity recognition. It featured a substantial set of longer tutorials on aspects of plan and activity recognition, and related topics and useful methods, as a way of establishing a common vocabulary and shared basis of understanding. Building on this shared understanding, individual researchers presented talks about their work in the area. There were also panel discussions which addressed questions about how to best foster progress in the field — specifically how to improve our ability to compare different plan and activity recognition algorithms — and address the question of whether to assume rationality in the modeled agents (a question that is of great concern in many fields at this time). This report presents a summary of the talks and discussions at the seminar.

Seminar 03.–08. April, 2011 – www.dagstuhl.de/11141
1998 ACM Subject Classification I.2.m Artificial Intelligence, Miscellaneous
Keywords and phrases Artificial intelligence, plan recognition, intent recognition, activity recognition.
Digital Object Identifier 10.4230/DagRep.1.4.1

1 Executive Summary

Robert P. Goldman
Christopher W. Geib
Tamim Asfour
Henry Kautz

Plan recognition, activity recognition, and intent recognition all involve making inferences about other actors from observations of their behavior, i.e., their interaction with the environment and with each other. The observed actors may be software agents, robots, or humans. This synergistic area of research combines and unifies techniques from user modeling, machine vision, intelligent user interfaces, human/computer interaction, autonomous and multi-agent systems, natural language understanding, and machine learning. It plays a crucial role in a wide variety of applications including:

- assistive technology
- software assistants

- computer and network security
- behavior recognition
- coordination in robots and software agents
- e-commerce and collaborative filtering

This Dagstuhl seminar brought together researchers with a wide range of interests and backgrounds related to plan and activity recognition. It featured a substantial set of longer tutorials on aspects of plan and activity recognition, and related topics and useful methods, as a way of establishing a common vocabulary and shared basis of understanding. These were:

- Plan recognition and discourse;
- Plan recognition and psychology;
- Probabilistic methods;
- Plan recognition and learning;
- Grammatical methods and
- Planning and plan recognition.

The common ground constructed by these tutorials provided a basis that individual researchers could build upon when sharing their specific interests and developments.

One challenge to progress in plan recognition is that there has not been a shared agreement about what constitutes plan recognition: what are its inputs and outputs, and what constitutes a good answer. In particular, this has inhibited progress because it is difficult to clearly compare new work in plan recognition with preceding work (quantitative comparisons are almost impossible), there is a paucity of shared data sets, etc. Coming into the seminar, the organizing committee proposed that the field might be improved by the introduction of a plan recognition competition, modeled on competitions in AI planning (the International Planning Competition), SAT solving, etc. Discussions at the seminar concluded that it would be premature to introduce such a competition at this time. Participants felt that a more productive use of community resources would be to develop a shared repository of plan and activity recognition data sets. A number of participants volunteered to provide their data sets, and there has been movement towards establishing a common public repository.

Plan Recognition: background

The earliest work in plan recognition was rule-based; researchers attempted to come up with inference rules that would capture the nature of plan recognition. However without an underlying formal model these rule sets are difficult to maintain and do not scale well.

In 1986, Kautz and Allen (K&A) published an article, "Generalized Plan Recognition" [7] that framed much of the work in plan recognition to date. K&A defined the problem of plan recognition as the problem of identifying a minimal set of *top-level actions* sufficient to explain the set of observed actions. Plans were represented in a plan graph, with top-level actions as root nodes and expansions of these actions into unordered sets of child actions representing plan decomposition. To a first approximation, the problem of plan recognition was then a problem of graph covering. K&A formalized this view of plan recognition in terms of McCarthy's circumscription. Kautz [6] presented an approximate implementation of this approach that recast the problem as one of computing vertex covers of the plan graph.

A number of early plan recognition systems used techniques such as rule-based systems [9], vertex covering, etc. Such techniques are not able to take into account differences in the *a priori* likelihood of different goals. Observing an agent going to the airport, this algorithm views "air travel," and "terrorist attack" as equally likely explanations, since they explain (cover) the observations equally well.

To the best of our knowledge, Charniak was the first to argue that plan recognition was best understood as a specific case of the general problem of *abduction*, or reasoning to the best explanation [3, e.g.,]. Charniak and Goldman (C&G) [2] argued that, viewing plan recognition as abduction, it could best be done as Bayesian (probabilistic) inference. Bayesian inference supports the preference for minimal explanations, in the case of equally likely hypotheses, but also correctly handles explanations of the same complexity but different likelihoods. For example, if a set of observations could be equally well explained by two hypotheses, theft and bragging being one, and theft alone being the other, simple probability theory (with some minor assumptions), will tell us that the simpler hypothesis is the more likely one. On the other hand, if as above, the two hypotheses were "air travel" and "terrorist attack," and each explained the observations equally well, then the prior probabilities will dominate, and air travel will be seen to be the most likely explanation. There have been many similar approaches to the problem, based on cost minimization, etc.

Another broad area of attack on the problem of plan recognition has been to reformulate it as a parsing problem [10, e.g.,]. Parsing-based approaches to plan recognition promise greater efficiency than other approaches, but at the cost of making strong assumptions about the ordering of plan steps. The major problem with parsing as a model of plan recognition is that it does not treat partially-ordered plans or interleaved plans well. Approaches that use statisical parsing [11, e.g.,] combine parsing and Bayesian approaches.

Finally, there has been a large amount of very promising work done using variations of Hidden Markov Models (HMMs) [1], techniques that came to prominence in signal processing applications, including speech recognition. These approaches offer many of the efficiency advantages of parsing approaches, but with the additional advantages of incorporating likelihood information and of supporting machine learning to automatically acquire their plan models. Standard HMMs seem to be insufficiently expressive to capture planful behavior, but a number of researchers have extended them to hierarchical formulations, that capture more complicated intentions. Conditional Random Fields [8], dynamic Bayes nets, and other probabilistic models have also been used.

Much of this latter work has been done under the rubric of *activity recognition*. The early work in this area very carefully chose the term *activity* or *behavior recognition* to distinguish it from plan recognition. The distinction to be made between activity recognition and plan recognition is the difference between recognizing a single (possibly complex) activity and recognizing the relationships between a set of such activities that result in a complete plan. Much of the work on activity recognition can be seen as discretizing a sequence of possibly noisy and intermittent low-level sensor readings into coherent actions that could be treated as inputs to a plan recognition system.

Several researchers have been interested in using plan recognition to improve team coordination [4, 5]. That is, if agents in a team can recognize what their teammates are doing, then they can better cooperate and coordinate. They may also be able to learn something about their shared environment. For example, a member of a military squad who sees a teammate ducking for cover may infer that there is a threat, so that it also takes precautions.

References

1 H. Bui and S. Venkatesh and G. A. W. West. Policy recognition in the abstract Hidden Markov Model. In *Journal of Artificial Intelligence Research*, vol. 17, pages 451–499, 2002.

2 Eugene Charniak and Robert P. Goldman. A Bayesian model of plan recognition. *Artificial Intelligence*, 64(1):53–79, November 1993.
3 Eugene Charniak and Drew McDermott. *Introduction to Artificial Intelligence.* Addison Wesley, Reading, MA, 1985.
4 Marcus J. Huber, Edmund H. Durfee and Michael P. Wellman. The Automated mapping of plans for plan recognition. In *Proceedings of the National Conference on Artificial Intelligence*, pages 344–351, 1994.
5 Gal A. Kaminka, D.V. Pynadath, and M. Tambe. Monitoring teams by overhearing: a multi-agent plan-recognition approach. In *Journal of Artificial Intelligence Research*, vol. 17, pages 83–135, 2002.
6 Henry Kautz. *A Formal Theory of Plan Recognition.* Technical Report, Department of Computer Science, University of Rochester, May 1987.
7 Henry Kautz and James F. Allen. Generalized plan recognition. In *Proceedings of the National Conference on Artificial Intelligence*, pages 32–38, 1986.
8 Lin Liao, Dieter Fox and Henry Kautz. Learning and inferring transportation routines. In *Proceedings of the National Conference on Artificial Intelligence*, 2004.
9 C.F. Schmidt and N.S. Sridharan and J.L. Goodson. The plan recognition problem: an intersection of psychology and artificial intelligence. In *Artificial Intelligence*, vol. 11, pages 45–83, 1978.
10 Marc Vilain. Getting serious about parsing plans: a grammatical analysis of plan recognition. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-90)*, pages 190–197, 1990.
11 David V. Pynadath and Michael P. Wellman. Accounting for context in plan recognition with application to traffic monitoring. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 472–481, 1995.

2 Table of Contents

3 Overview of Talks

3.1 From Motion to Text and Back for Humanoid Robots

Tamim Asfour (KIT – Karlsruhe Institute of Technology, DE)

Joint work of Asfour, Tamim; Dillmann, Ruediger
URL http://his.anthropomatik.kit.edu/english/65.php

Semantic representations are a prerequisite for the development of cognitive capabilities and understanding in robots as well as for cooperation, interaction and communication with humans. Building such representations from sensorimotor experience rely on organizing the system's sensorimotor experience to provide data structures which can be used at different levels of the systems hierarchy and breaks through the gap between sensorimotor level and symbolic level.

In this talk, we present our recent work on building humanoid robots able to act, interact and autonomously acquire knowledge in the real world. Results are presented towards the implementation of integrated 24/7 humanoid robots able to 1) perform complex grasping and manipulation tasks in a kitchen environment 2) autonomously acquire object knowledge through active visual and haptic exploration and 3) learn actions from human observation and imitate them in goal-directed manner.

Further, we discuss how a motion library can be built from observation of human demonstration and how the elements of such library can be represented and enriched by additional constraints such as objects involved in an action, forces applied to an object and agents involved in interaction and cooperation tasks.

The resulting data structures, together methods of natural language processing will facilitate the link between sensorimotor experience and linguistic representations.

3.2 Bayesian Theory of Mind: Modeling Joint Belief-Desire Inference

Chris L. Baker (MIT - Cambridge, US)

Joint work of Baker, Chris L.; Saxe, Rebecca R.; Tenenbaum, Joshua B.
Main reference Baker, C.L., Saxe, R.R., Tenenbaum, J.B., "Bayesian Theory of Mind," Proceedings of the Thirty-Second Annual Conference of the Cognitive Science Society, to appear.

We present a computational framework for understanding Theory of Mind (ToM): the human capacity to make joint inferences about the beliefs and desires underlying the observed actions of other agents. Bayesian ToM (BToM) formalizes the concept of intentional agency at the heart of ToM as a partially observable Markov decision process (POMDP), and performs Bayesian inference over this structured model to reconstruct an agent's joint belief state and reward function given observations of its behavior in some environmental context. We test the BToM framework by collecting people's joint inferences of agents' desires and beliefs about unobserved aspects of the environment in response to stimuli of agents moving in simple spatial scenarios. BToM performs substantially better than two simpler variants: one in which desires are inferred without reference to an agents' beliefs, and another in which beliefs are inferred without reference to the agent's dynamic observations in the environment.

3.3 Eliciting Plan Recognition Cues by Provoking Opponents in RTS Games

Francis Bisson (Université de Sherbrooke, CA)

Joint work of Bisson, Francis; Kabanza, Froduald; Benaskeur, Abder Rezak; Irandoust, Hengameh
Main reference Bisson, F., Kabanza, F., Benaskeur, A. and Irandoust, H., "Provoking Opponents to Facilitate the Recognition of their Intentions," Proceedings of the AAAI Student Abstract and Poster Program, 2011.
URL http://planiart.usherbrooke.ca/ bisson/papers/aaai2011-poster.pdf

For agents evolving in adversarial environments such as RTS games, it is necessary to be able to recognize the goals of their opponents in the environment. However, most adversarial plan recognizers rely on a passive observation of the opponents, gathering and analyzing cues related to their goals. In contrast, in this talk I will present preliminary results for a plan recognition approach that can provoke the opponents in order to observe their reactions, and use the resulting cues to disambiguate the current set of hypotheses on their goals.

3.4 Knowledge-rich Plans – How they Enable Explanation, Recognition, and Repair

Susanne Biundo (Universität Ulm, DE)

Hybrid planning combines the traditional planning paradigms of hierarchical task network (HTN) and partial-order causal-link (POCL) planning. The resulting systems are able to use predefined standard solutions like in pure HTN planning, but can also develop (parts of) a plan from scratch or modify a default solution in cases where the initial state deviates from the presumed standard. This flexibility makes hybrid planning particularly well suited for real-world applications.

Based on a completely declarative description of actions, tasks, and solution methods, hybrid planning smoothly integrates reasoning about procedural knowledge and causalities and allows for the generation of knowledge-rich plans of action. The information those plans comprise includes causal dependencies between actions on both abstract and primitive levels as well as information about their hierarchical and temporal relationships. By making use of this information, as well as of the underlying declarative domain models, capabilities like the generation of courses of action on various abstraction levels, the stable repair of failed plans, plan recognition, and the explanation of different solutions for a given planning problem can be implemented by advanced automated reasoning techniques.

3.5 Thinking about Evaluation and Corpora for Plan Recognition

Nate Blaylock (IHMC – Pensacola, US)

Joint work of Blaylock, Nate; Allen, James

Lately, many fields of AI have benefited from labeled corpora as common resources for training and evaluating performance. I will discuss some of the issues of creating corpora for plan recognition and argue that this would be a worthwhile investment for our community. I will also discuss a range of metrics for evaluating the performance of plan recognizers.

3.6 "The gist of the matter": On plan understanding, behaviour prediction, and referring expressions

Michael Brenner (Universität Freiburg, DE)

The talk discussed the problem of understanding the purpose of a plan. I argued that this is more than recognising the goal state it tries to achieve, but rather a characterisation of the plan in relation to "deficits" in (and other constraints on) the initial state. Such a characterisation (called a "gist") can be described indepently of the specific initial state by making use of referring expressions, similarly to their use in natural-language processing.

The talk then discussed cost measures for gists and some initial ideas for recognising them given an observed plan.

3.7 AbRA: An Abductive, Rationalizing Agent for Plan Recognition

Will Bridewell (Stanford University, US)

Joint work of Bridewell, Will; Langley, Pat
Main reference Bridewell, W., & Langley, P., "A computational account of everyday abductive inference," Proceedings of the 33rd Annual Meeting of the Cognitive Science Society, 2011.

Plan recognition is a naturally abductive task. That is, looking at the actions of an agent, one must make assumptions about its underlying plan. AbRA is a novel, logic-based system for carrying out socially aware, abductive inference.

This system emphasizes cyclic, online operation, incremental extension of explanations, a shifting focus of attention, and a data-driven inference mechanism. Guided by local coherence, AbRA constructs an explanation/plan that ties the observations into a plausible, although not necessarily correct or even optimally rational, story. Here I provide an intuitive description of the system, report preliminary results on a complex plan recognition domain, and plot our current research trajectory.

3.8 Plan Recognition for User-Adaptive Interaction

Cristina Conati (University of British Columbia – Vancouver, CA)

I will first give some examples of how we use plan/goal/activity recognition in user-adaptive interactive systems. I will then introduce two directions we are exploring to improve the accuracy and usability of user-adaptive interaction: (i) using eye-gaze information to inform plan recognition; (ii) Explaining to the user aspects of the system's reasoning to increase user trust in the system's adaptive interventions

3.9 Activity Recognition for a Knowledge Worker Assistant

Thomas Dietterich (Oregon State University, US)

Joint work of Dietterich, Thomas; Shen, Jianqiang; Bao, Xinlong; Keiser, Victoria; Bui, Hung
Main reference Dietterich, T. G., Bao, X., Keiser, V., Shen, J., "Machine Learning Methods for High Level Cyber Situation Awareness," In Jajodia, S., Liu, P., Swarup, V. and Wang, C. (Eds.) , *"Cyber Situation Awareness"*. Springer, 2009, pp. 227–247.
URL http://web.engr.oregonstate.edu/ tgd/publications/csa-dietterich-bao-keiser-shen.pdf

Knowledge workers execute hundreds of simple digital workflows in a typical work week. We will describe three forms of activity recognition that seek to assist knowledge workers with these workflows. The first is the TaskTracer Project Predictor, which attempts to infer which project the user is working on based on observed desktop activity. The second is the TaskTracer Folder Predictor, which predicts which folder the user wishes to access when opening or saving a file. Experimental studies show that Folder Predictor reduces by 50desired folder. The third is a method for discovering and recognizing workflow executions as part of an effort to provide proactive assistance to desktop knowledge workers.

3.10 Tutorial: Plan Recognition via Inverse Reinforcement Learning

Thomas Dietterich (Oregon State University, US)

Reinforcement learning methods seek an optimal policy for an unknown Markov Decision Process by interacting with that process to learn the transition and reward functions. Inverse Reinforcement learning is given the transition function and the optimal (or expert) policy and seeks to find the reward function. More generally, Inverse RL can be viewed as attempting to infer the goals underlying observed behavior. A closely-related task is to infer a expert's policy from demonstrations.

The components of an MDP (reward function, policy, value function, state visitation probabilities) are inter-related, and Inverse RL methods can be categorized based on the primary component that they attempt to learn. Given observed behavior, there are multiple reward functions and value functions consistent with it (even asymptotically), which makes direct attempts to learn these components ill-defined. In contrast, methods that seek to directly learn the policy or the state visitation probabilities appear to be more successful, because these are uniquely specified by observed behavior (at least asymptotically). This tutorial surveys several methods for learning reward functions, state visitation probabilities, and policies and concludes that learning state visitation probabilities is the most promising approach.

3.11 Plan Recognition and Collaborative Assistants

George Ferguson (University of Rochester, US)

This talk discusses the roles of plan recognition in the design and implementation of collaborative assistants—intelligent systems that interact naturally to help people solve problems. Two key roles are identified: (1) the disambiguation of natural language input to support mixed-initiative interaction and learning from demonstration, and (2) tracking user performance during execution to support both mixed-initiative interaction, task- and context-sensitive help, and overt instruction or teaching. These roles are illustrated with examples from systems we have implemented in the past. We also describe a new research thrust based on combining natural language description with low- level sensor data for learning models of real-world tasks performed by humans.

3.12 Planning and Plan Recognition

Hector Geffner (Universidad Pompeu Fabra – Barcelona, ES)

Joint work of Geffner, Hector; Ramirez, Miquel
Main reference M. Ramirez and H. Geffner, "Probabilistic Plan Recognition using off-the-shelf Classical Planners," Proc. AAAI-2010.
URL http://www.dtic.upf.edu/~hgeffner

Plan recognition is like planning in reverse: while in planning the goal is given and a plan is sought; in plan recognition, part of a plan is given, and the goal and complete plan are sought. Until recently, however, plan recognition has been addressed using methods which are not related to planning such as parsing algorithms, Bayesian network procedures, and specialized methods.

In almost all cases, the space of possible plans or activities to be recognized is assumed to be given by a suitable library or set of policies.

Recently, an approach that does not require the use of a plan library and which uses planning technology, as been developed by Baker, Saxe, and Tenenbaum, on the one hand, and by Ramirez and Geffner, on the other. In this approach, the plan recognition problem is mapped into a collection of planning problems that can be solved with off-shelf-planners. The posterior distribution over the possible goals given the observation is inferred from basic probability lows and costs derived from the use of a planner. The approach has been used to perform planning recognition over classical planning models, Markov Decision Processes (MDPs), and Partially Observable MDPs (POMDPs).

In this invited talk, I review the relevant ideas from AI Planning and their use for formulating and solving the plan recognition problem.

3.13 Grammatical Methods for Plan Recognition

Christopher W. Geib (University of Edinburgh, GB)

This talk is an overview of prior and current work on viewing plan recognition as a parsing task given a formal grammar and a sequence of observations. It covers work in using: regular, context free, probabilistic state-dependent, plan tree, tree adjoining, and combinatorial categorial grammars.

3.14 Plan recognition: a historical survey, part I

Robert P. Goldman (SIFT – Minneapolis, US)

The seminar opened with a historical survey of plan recognition. We present a taxonomy of plan recognition problems, including the conventional distinction between keyhole, intended, and adversarial plan recognition, but touching on other dimensions such as fallible versus ideal agents, complete versus partial observability, open versus closed worlds, static versus evolving sets of intentions, and expressiveness of plan representation. After outlining the dimensions of plan representation, we proceeded to review methods used in the early history of plan recognition. We began with early techniques, based on rule-based systems, then moved on to discuss the formalization of the field, based on Kautz and Allen's generalized theory of plan recognition, and Vilain's parsing-based complexity analysis of the theory. We discussed systems inspired by this work, including techniques based on parsing and on minimal graph cover.

We concluded the first part of this talk with a discussion of techniques based on Bayes networks.

3.15 Behavior Recognition and Demonstration for Human/Robot Cooperation

Tetsunari Inamura (NII – Tokyo, JP)

Main reference Tetsunari Inamura, Keisuke Okuno, "Robotic Motion Coach: Effect of Motion Emphasis and Verbal Expression for Imitation Learning," Proc. 3rd International Conference on Cognitive Neurodynamics, p.186, 2011.

Behavior Recognition and Demonstration for Human-Robot Cooperation In this talk, development of a robotic coaching system is discussed. recognition and reproduction of human's whole body motion patterns are focused towards establishment of natural human-robot cooperation through demonstration of motion and speech conversation. A robotic coaching system should be a target application because the robot should recognize user's motion, demonstrate modified motion according to user's level of skill, and generate advice with verbal expression. Abstract of motion pattern using HMMs and a phase space are proposed. Using the phase space, motion emphasis and generation of verbal expression are integrated. A robotic simulator platform is also introduced towards a basis of evaluation of plan recognition for human-robot interaction application.

References

1 Tetsunari Inamura and Keusuke Okuno. *Robotic Motion Coach: Effect of Motion Emphasis and Verbal Expression for Imitation Learning*. The 3rd International Conference on Cognitive Neurodynamics, p.186, 2011.

3.16 Plan recognition challenges in real-time strategy games

Froduald Kabanza (Université de Sherbrooke, CA)

Main reference Kabanza, F., Bellefeuille, P., Bisson, F., Benaskeur, A., and Irandoust, H., "Opponent Behaviour Recognition for Real-Time Strategy Games," Proc. of AAAI Workshop on Plan, Activity and Intent Recognition (PAIR), 2010.
URL http://planiart.usherbrooke.ca/kabanza/publications/10/pair10-opponent.pdf

In real-time strategy (RTS) games, players recruit and manoeuvre army units in order to defeat their opponents. The victory condition may vary from one game or scenario to another, but it usually involves destroying some or all of the opponent's assets. A key component of the player's situation awareness in this context is the recognition of his opponent's intent and plans. This presentation covers some of the main challenges posed by the intent and plan recognition problems in RTS games and sketch the main building blocks of a conceptual plan recognition method geared towards addressing these challenges. The method is still a concept in the early development stage, and the presentation will be aimed at stimulating a discussion and encouraging the audience to comment on it rather than demonstrating its effectiveness. The RTS domain is used for concrete scenarios, but the fundamental intent and plan recognition problems that we are addressing remain relevant to other adversarial domains.

3.17 Survey of Probabilistic Activity and Plan Recognition

Henry A. Kautz (University of Rochester, US)

We provide an overview of probabilistic plan recognition methods. These include:
- HMM
- Layered HMM
- Dynamic Bayesian Networks
- Stochastic Grammars
- Conditional Random Fields
- Relational Markov Model

- Markov Logic
- Bayesian Inverse Planning
- Inverse Reinforcement Learning
- N-Gram Models

3.18 Mobile Intention Recognition And Spatially Constrained Grammars

Peter Kiefer (Universität Bamberg, DE)

Main reference Peter Kiefer, "The Mobile Intention Recognition Problem And An Approach Based On Spatially-Constrained Grammars," PhD Thesis, to appear 2011.

Mobile intention recognition differs from the general plan and intention recognition problem by the availability of spatial context information for each input behavior. This talk proposes to use the specific properties of spatial context, such as continuity and hierarchies, for the disambiguation of mobile behavior sequences.

Most current approaches for interpreting mobile behavior focus on activity recognition, not on (high-level) intentions. This talk argues that formal grammars, enhanced with spatial information, are well-suited for representing high-level intentions. Formal grammars make expressiveness properties explicit, are cognitively comprehensible, and allow for easy geographic portability - a requirement crucial in mobile assistance domains.

Many mobile assistance domains require us to represent behaviors and intentions with formal grammars of higher expressiveness than context-free grammars. This talk proposes to enhance the mildly context-sensitive Tree-Adjoining Grammar formalism, well-known in natural language processing, with spatial constraints, yielding in a spatial grammar specifically useful to express the Visit-/Revisit-pattern frequently occuring in mobile assistance.

3.19 Probabilistic Plan Recognition

Kathryn B. Laskey (George Mason University – Fairfax, US)

Plan recognition is naturally viewed as a problem in inference under uncertainty. From observations of an agent's actions (or effects actions), a plan recognition system attempts to infer the agent's goal and explain the actions in terms of a plan for achieving the goal. Typically, there are multiple explanations for any sequence of actions. Probability is a natural approach to weighing the relative plausibility of alternative explanations.

Attractive features of probability include its strong theoretical foundation, its unified view of inference and learning, and its practical success in a growing body of applications. On the other hand, probabilistic inference and learning are NP-hard, and achieving sufficiently expressive yet tractable representations is a major challenge. This talk provides an overview of major probabilistic representations used for plan recognition, describes common exact and approximate inference methods, and identifies research challenges.

3.20 Plan Recognition Using Multi-Entity Bayesian Networks and PR-OWL

Kathryn B. Laskey (George Mason University – Fairfax, US)

Joint work of Carvalho, Rommel N.; Costa, Paulo C. G.; Laskey, Kathryn B.; and Chang, KuoChu
Main reference Carvalho, Rommel N.; Costa, Paulo C. G.; Laskey, Kathryn B.; Chang, KuoChu, "PROGNOS: Predictive Situational Awareness with Probabilistic Ontologies," Proceedings of the Thirteenth International Conference of the Society of Information Fusion (FUSION 2010).
URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5711970

Increasingly expressive languages are emerging for representing and reasoning with probability. Multi-Entity Bayesian Networks (MEBN) is a first-order language for specifying probabilistic knowledge bases as parameterized fragments of Bayesian networks. MEBN fragments (MFrags) can be instantiated and combined to form arbitrarily complex graphical probability models. An MFrag represents probabilistic relationships among a conceptually meaningful group of uncertain hypotheses. The PR-OWL probabilistic ontology language, based on MEBN, extends OWL to allow expression of uncertainty about attributes and relations. An example is given of a MEBN theory for maritime domain awareness and its implementation as a PR-OWL probabilistic ontology.

3.21 Plan Recognition in/with Agent Programming Languages

Yves Lespérance (York University – Toronto, CA)

Joint work of Lesperance, Yves; Goultiaeva, Alexandra
Main reference Goultiaeva, A. and Lespérance, Y., "Incremental Plan Recognition in an Agent Programming Framework," In Working Notes of the AAAI 2007 Workshop on Plan, Activity, and Intent Recognition (PAIR'07), Vancouver, BC, July, 2007.
URL http://www.cse.yorku.ca/~lesperan/papers/PAIR07.pdf

In the talk, I discuss how agent programming languages can be used for specifying plans for plan recognition, and also how plan recognition capabilities could be usefully added to such languages. I focus on the ConGolog agent programming language, based on the situation calculus. I review an account of plan recognition for this setting [1], where ConGolog plan libraries are used. This provides a very expressive language for specifying plans.

It supports several forms of nondeterminism and allows sketchy plan templates to be specified. Also it is closed under union, intersection, and complementation, so one can specify the set of runs that are part of the plan in a completely compositional way. I discuss how the account can be simplified by restricting attention to situation-determined programs, where the remaining program after a partial execution is uniquely determined [2].

References

1 Goultiaeva, A. and Lespérance, Y., "Incremental Plan Recognition in an Agent Programming Framework." In *Working Notes of the AAAI 2007 Workshop on Plan, Activity, and Intent Recognition (PAIR'07),* Vancouver, BC, July, 2007.
2 De Giacomo, G., Lespérance, Y., and Muise, C., "Agent Supervision in Situation-Determined ConGolog." To appear in *Working Notes of the 9th International Workshop on Nonmonotonic Reasoning, Action and Change (NRAC-2011),* Barcelona, Spain, July, 2011.

3.22 Yappr: From LL parsing to plan recognition

John Maraist (SIFT – Minneapolis, US)

Joint work of Geib, Christopher W.; Goldman, Robert P.; Maraist, John
Main reference Christopher W. Geib, John Maraist, Robert P. Goldman, "A New Probabilistic Plan Recognition Algorithm Based on String Rewriting," Proc. of the 18th International Conference on Automated Planning and Scheduling (ICAPS-2008), Sydney, Australia, 2008.

We present the probabilistic HTN plan recognition algorithm Yappr by evolution from its motivating classical parsing algorithm. We begin with a simple stack-based automaton for nondeterministic LL parsing, and identify the three refinements which produce Yappr: precompilation of plans for efficient, deterministic retrieval; replacement of the parser stack with a graph allowing multiple application points; and maintenance of multiple explanations rather than a single parse. We conclude with a look forward at the advantages and disadvantages of moving from an LL-based to an LR- based approach.

3.23 No More Plan Libraries – The Case for a Structureless World

David Pattison (The University of Strathclyde – Glasgow, GB)

Joint work of Pattison, David; Long, Derek

Plan Libraries have always gone hand-in-hand with Plan Recognition. Having a concise set of possible plans to map to an agent's observed actions allows for reliable goal recognition, next-action prediction, intermediate states and further prediction and analysis.

The problem is that these libraries don't exist in the real world. Construction of a plan/goal library by hand is a labour-intensive process with a highly bespoke output. In the past 10 years work has moved towards the generation of these libraries at runtime, but computer-generated plans can never be truly perfect, with invalid or unwanted entries an unavoidable side-effect. In this talk I will discuss my own work on Goal Recognition as a Planning problem as an example of such a model, and the need to move away from a library-based standard in recognition.

3.24 Modeling Theory of Mind as Plan Recognition

David Pynadath (University of Southern California – Marina del Rey, US)

Human social interaction relies on our ability to model each other as (mostly) rational actors. Despite the uncertainty we may have of another's intentions and subjective beliefs, our theory of mind provides valuable leverage that we exploit whenever possible. Thus, multiagent modeling of social situations can benefit from a computational implementation of theory of mind. I present one such implementation where an agent reuses its own decision-theoretic planning to generate expectations about the behavior of others. By combining its uncertain beliefs about the possible mental states of others with a planning mechanism, it becomes straightforward to recast this problem as plan recognition. Inverting the planning process generates abductive reasoning that an agent can use to update its beliefs about others as it observes their behavior. While such recursive beliefs can become prohibitively complex as the number of agents increases, I also show that the decision-theoretic context gives each agent a utility-based metric for deciding what it can safely ignore about everyone else.

We have implemented these algorithms within a social simulation framework, PsychSim, that has supported simulations of various scenarios, including bilateral negotiation, language and cultural training, and urban stabilization operations.

3.25 Intentions in Collaboration: Insights from Meaning

Matthew Stone (Rutgers University – Piscataway, US)

Conversation is one of many cases where we want to attribute intentions to agents exhibiting improvised, fluid, expert strategic behavior. Understanding an utterance, on the received view, is just recognizing the speaker's communicative intention. But does this make sense? Any action could implicitly prepare for an open-ended array of contingencies, reflecting an open-ended array of expectations its agent brings to the situation. And agents may well choose those actions through heuristic problem solving and learned strategies—processes that fit poorly with the intuitive notions of deliberation and commitment used by intention theorists. How can we actually carve out recognizable intentions from such a complex ensemble of factors? Think of recognizing the intention of the Roshambo player you see throw Rock, who might specifically intend to make a throw chosen at random, to play best response against you, or both.

In this talk, I will try to clarify both the received view of meaning as intention and the place of intention recognition in collaborative activity. We have surprisingly strong judgments about what people can and cannot mean with individual utterances, and about how those meanings fit together over the course of a conversation. These judgments motivate a specific kind of collaborative intention: a system of public categories of action, coordinated across agent and teammates, classifying each action based on the content of the mental representation that immediately underpins its performance. Playing Rock fits in such a system, playing at random or playing best response do not. I close by sketching how an implemented agent—one that might not actually meet traditional standards for having individual or shared intentions—could use such categories to pursue its utterances in a meaningful and collaborative way.

3.26 Assuming the Human's Cognitive State as Basis for Assistant System Initiative

Ruben Strenzke (Unibw – München, DE)

Joint work of Strenzke, Ruben; Schulte, Axel
Main reference Strenzke, Ruben; Schulte, Axel, "Modeling the Human Operator's Cognitive Process to Enable Assistant System Decisions," Proc. of Goal, Activity and Plan Recognition (GAPRec) Workshop in conjunction with International Conference on Automated Planning and Scheduling (ICAPS) 2011.
URL http://icaps11.icaps-conference.org/proceedings/gaprec/strenzke.pdf

In the Manned-Unmanned Teaming application a transport helicopter commander is also controlling multiple reconnaissance unmanned aerial vehicles (UAVs) that shall reduce the risk of the transport mission. This human operator is therefore responsible of planning and re-planning a multi-aircraft mission under situation-dependent time pressure. To lower the workload generated thereby, he/she shall be supported by a cognitive assistant system that is designed with respect to the Cooperative Automation and mixed-initiative planning approaches.

In order to decide, whether, when, and in which way to take initiative, the assistant system has to know about the human operator's goal, assume his/her plan, and evaluate his/her activity. The latter is also necessary in order to estimate the current workload situation.

In our implementation the assistant system assumes the human plan by taking into account the mission order (goal constraints) and the partial or complete plan entered by the human into the system (plan constraints). The assistant system then checks, if this plan is feasible, complete, and compare this plan with what itself has planned automatically. On this basis the assistant system can decide to take initiative to urge the operator to improve plan quality and enforce timely plan execution.

For another task of this human operator (identification of vehicles during route reconnaissance by UAVs) there have been workload estimation experiments based on operator activity, i.e. his/her manual and visual interaction with the system.

The workload estimation can be accomplished by HMMs per operator task and per workload level. In case of workload higher than normal, so-called self-adaptive strategies are observable, which alter the human behavior.

The great challenge remains the combination of methods like used in the two approaches mentioned. This would allow to assume the operator's cognitive state in a broader context, which makes sense because in the cognitive planning process the goals lead to the plan, the plan leads to choosing action, action leads to behavior, and workload leads to errors induced in the different steps of this process.

3.27 Coupling Plan Recognition with Plan Repair for Real-Time Opponent Modeling

Gita Reese Sukthankar (University of Central Florida – Orlando, US)

Joint work of Sukthankar, Gita Reese; Laviers, Kennard
Main reference Kennard Laviers, Gita Sukthankar, "A Real-time Opponent Modeling System for Rush Football," Proceedings of International Joint Conference on Artificial Intelligence, July 2011.

One drawback with using plan recognition in adversarial games is that often players must commit to a plan before it is possible to infer the opponent's intentions. In such cases, it is valuable to couple plan recognition with plan repair, particularly in multi-agent domains where complete replanning is not computationally feasible. This paper presents a method for learning plan repair policies in real-time using Upper Confidence Bounds for Trees (UCT). We demonstrate how these policies can be coupled with plan recognition in an American football game (Rush 2008) to create an autonomous offensive team capable of responding to unexpected changes in defensive strategy. Our real-time version of UCT learns play modifications that result in a significantly higher average yardage and fewer interceptions than either the baseline game or domain-specific heuristics.

3.28 Efficient Hybrid Algorithms for Plan Recognition and Detection of Suspicious and Anomalous Behavior

Dorit Zilberbrand (Givat Shmuel, IL)

Plan recognition is the process of inferring other agents' plans and goals based on their observable actions. Modern applications of plan recognition, in particular in surveillance and security raise several challenges. First, a number of key capabilities are missing from all but a handful of plan recognizers: (a) handling complex multi-featured observations; (b) dealing with plan execution duration constraints; (c) handling lossy observations (where an observation is intermittently lost); and (d) handling interleaved plans. Second, essentially all previous work in plan recognition has focused on recognition accuracy 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 work we present set of efficient plan recognition algorithms that are capable of handling the variety of features required of realistic recognition tasks. We also 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).We demonstrate the efficacy of the techniques described above, by applying them to the problem of detecting anomalous and suspicious behavior. The system contains the symbolic plan recognition algorithm, which detects anomalous behavior, and the utility-based plan recognizer which reasons about the expected cost of hypotheses. These two components form a highly efficient hybrid plan recognizer capable of recognizing abnormal and potentially dangerous activities. We evaluate the system with extensive experiments, using real-world and simulated activity data, from a variety of sources.

4 Panel Discussions

4.1 A Plan recognition competition?

Christopher W. Geib (University of Edinburgh – Edinburgh, UK)

Little of the research published in plan recognition reports on results that can be directly compared to previous research. The community has yet to agreed on standard data sets or benchmark problems that all systems are expected to be evaluated against. In an effort to address this same problem the AI planning research community established the International Planning Competition (IPC). At Dagstuhl, we had a panel discussion to consider if a similar competition would benefit the plan recognition community. The pannel members were Christopher Geib, Hector Geffner, Jerry Hobbs, and Froduald Kabanza. There was lively debate both pro and con, and while there were strong arguments in favor it was not universally agreed that an IPC-style competition would be in the best interests of the plan recognition community. There was significant concern that such a competition could fragment the small and still growing plan recognition community and might unintentionally limit future research directions. That said, it was generally agreed that more efforts should be made, especially by leaders in the community, to share data sets and report work in a way that enabled more directly comparable results. A number of participants in the seminar agreed to make their data sets freely available.

4.2 Rational versus fallible agents

Matthew Stone (Rutgers University – Piscataway, US)

The panel on rational versus fallible agents primarily addressed the issues involved in using automated plan recognition approaches to understand human activity. There are several respects in which humans may not be perfect decision makers. For example, as Professor Conati observed, student learners solving problems may apply incorrect approaches or use erroneous facts. She found it crucial to model these mistakes in her systems for plan recognition for intelligent tutoring. Similarly, Professor Stone pointed out that many examples of indirection in conversation seem to rely on the heuristics and biases of human decision making, and cited a number of likely cases from the work of Steven Pinker. His dialogue systems increasingly assume very constrained reasoning on the part of interlocutors. At the same time, however, Chris Baker emphasized that there are many domains where people do exhibit expert behavior which systems need to understand. Assumptions of rationality can be very effective in these domains in making good predictions with simple models and minimal training. Meanwhile, Professor Kautz observed that many learning approaches to plan recognition put the focus on finding reliable patterns of activity, and make few assumptions one way or the other about the rationality of target agents. Looking forward, the panel recommended that researchers aim to factor out assumptions about agents from their algorithms wherever possible, so that the community can focus on techniques that generalize across diverse agent populations and tasks. It is also important to evaluate this dimension of plan recognition systems, to understand where and when assumptions of rationality are violated, and what effects such cases have both on the performance of plan recognition algorithms and the contribution of plan recognition to broader measures of system performance.

5 Invited Talks

Invited talks have their abstracts (where available) in the main body of the report.

- Plan recognition and discourse, Jerry Hobbs
- Plan recognition and psychology, Chris Baker
- Probabilistic methods, Kathryn Blackmond Laskey
- Plan recognition and learning, Tom Dietterich
- Grammatical methods, Christopher W. Geib
- Planning and plan recognition, Hector Geffner

Participants

- Tamim Asfour
KIT – Karlsruhe Institute of Technology, DE
- Chris L. Baker
MIT – Cambridge, US
- Francis Bisson
Université de Sherbrooke, CA
- Susanne Biundo
Universität Ulm, DE
- Nate Blaylock
IHMC – Pensacola, US
- Michael Brenner
Universität Freiburg, DE
- Will Bridewell
Stanford University, US
- Cristina Conati
University of British Columbia – Vancouver, CA
- Thomas Dietterich
Oregon State University, US
- George Ferguson
University of Rochester, US
- Hector Geffner
Universidad Pompeu Fabra – Barcelona, ES
- Christopher W. Geib
University of Edinburgh, GB
- Robert P. Goldman
SIFT – Minneapolis, US
- Jerry Hobbs
University of Southern California – Marina del Rey, US
- Tetsunari Inamura
NII – Tokyo, JP
- Froduald Kabanza
Université de Sherbrooke, CA
- Henry A. Kautz
University of Rochester, US
- Peter Kiefer
Universität Bamberg, DE
- Kathryn B. Laskey
George Mason University – Fairfax, US
- Yves Lesperance
York University – Toronto, CA
- John Maraist
SIFT – Minneapolis, US
- David Pattison
The University of Strathclyde – Glasgow, GB
- David Pynadath
University of Southern California – Marina del Rey, US
- Matthew Stone
Rutgers Univ. – Piscataway, US
- Ruben Strenzke
Unibw – München, DE
- Gita Reese Sukthankar
University of Central Florida – Orlando, US
- Dorit Zilberbrand
Givat Shmuel, IL