10081 Abstracts Collection
Cognitive Robotics
— Dagstuhl Seminar —

Gerhard Lakemeyer[1], Hector Levesque[2] and Fiora Pirri[3]

[1] RWTH Aachen, DE
gerhard@cs.rwth-aachen.de
[2] University of Toronto, CA
hector@cs.toronto.edu
[3] University of Rome "La Sapienza", IT
pirri@dis.uniroma1.it

Abstract. From 21.02. to 26.02.2010, the Dagstuhl Seminar 10081 "Cognitive Robotics " was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Cognitive roboticsm, Knowledge representation and reasoning, Machine learning, Cognitive science, Cognitive vision

10081 Summary – Cognitive Robotics

Research in robotics has traditionally emphasized low-level sensing and control tasks including sensory processing, path planning, and manipulator design and control. Research in Cognitive Robotics, on the other hand, emphasizes those cognitive functions that enable robots and software agents to reason, act and perceive in changing, incompletely known, and unpredictable environments. Such robots must, for example, be able to reason about goals, to choose actions and to focus on patterns, objects and events according to the task execution and the cognitive states of other agents, by taking into account time, resources, and the consequences of their decisions. In short, cognitive robotics is concerned with integrating reasoning, perception, and action within a uniform theoretical and implementation framework.

The term cognitive robotics and the vision that knowledge representation and reasoning plays a fundamental role in the design of cognitive robots was first laid out by the late Ray Reiter in his lecture on receiving the Research Excellence Award by the International Joint Conference on Artificial Intelligence (IJCAI) in 1993. Since 1998, biannual Cognitive Robotics workshop with Dagstuhl being the seventh in this series.

While the earlier workshops were largely a forum for presenting state-of-the-art research results, the purpose of the Dagstuhl event was to broaden the view and bring together people from various disciplines to shed new light on the issues in cognitive robotics. In this respect we were very fortunate to have participants from areas such as robotics, machine learning, cognitive vision, computational neuroscience, and knowledge representation and reasoning.

Given the diversity of the group, we spent the first day with tutorial-style presentation, starting out with an overview of Cognitive Robotics in the sense of Ray Reiter's vision by Gerhard Lakemeyer. This was followed by presentations on Computational Neuroscience by Laurent Itti, Planning and Execution Monitoring by Brian Williams, Probabilistic Reasoning by Eyal Amir, Cognitive Vision by Jim Little, and Human-Robot Interaction by Geert-Jan Kruijff. The rest of the Workshop consisted of research presentations, a panel, and three breakout discussion groups on the following topics: the nature of perception, symbolic and numerical uncertainty, and the role of automated reasoning. The Proceedings collected here represent submissions from all but 2 of the participants at the Workshop in the order they were given. In some cases we have short abstracts of the talks, and in others, full research papers. We thank the participants and the copyright holders for permission to use these papers.

Keywords: Cognitive roboticsm, Knowledge representation and reasoning, Machine learning, Cognitive science, Cognitive vision

Combining Planning and Motion Planning

Eyal Amir (University of Illinois, USA)

Robotic manipulation is important for real, physical world applications. General Purpose manipulation with a robot (eg. delivering dishes, opening doors with a key, etc.) is demanding. It is hard because (1) objects are constrained in position and orientation, (2) many non-spatial constraints interact (or interfere) with each other, and (3) robots may have multidegree of freedoms (DOF). In this paper we solve the problem of general purpose robotic manipulation using a novel combination of planning and motion planning. Our approach integrates motions of a robot with other (non-physical or external-to-robot) actions to achieve a goal while manipulating objects. It differs from previous, hierarchical approaches in that (a) it considers kinematic constraints in configuration space (C-space) together with constraints over object manipulations; (b) it automatically generates high-level (logical) actions from a C-space based motion planning algorithm; and (c) it decomposes a planning problem into small segments, thus reducing the complexity of planning.

Joint work of: Choi, Jaesik ; Amir, Eyal

Keywords: Motion Planning, Factored Planning, Robotic arm

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2629

Challenges for domestic service robots

Sven Behnke (Universität Bonn, DE)

Benchmarking robotic systems is difficult. Robot competitions, such as RoboCup and the DARPA Grand and Urban Challenges provide a standardized test environment and allow for the direct comparison of different systems. For domestic service robots, in recent years, the RoboCup@Home league has been established. Domestic service tasks require three main skills from autonomous robots: robust navigation, mobile manipulation, and intuitive communication with the users. My team NimbRo developed the robot Dynamaid for these competitions. For robust navigation, Dynamaid has a base with four individually steerable differential wheel pairs, which allow omnidirectional motion. For mobile manipulation, Dynamaid is additionally equipped with two anthropomorphic arms that include a gripper, and with a trunk that can be lifted as well as twisted. For intuitive multimodal communication, the robot has a microphone, stereo cameras, and a movable head. Dynamaid can perceive persons and objects in its environment, recognize and synthesize speech. Together with our communication robot Robotinho, it won the Innovation award at RoboCup 2009.

Attending to Motion: an object-based approach

Anna Belardinelli (Universität Bielefeld, DE)

Visual attention is the biological mechanism allowing to turn mere sensing into conscious perception. In this process, object-based modulation of attention provides a further layer between low-level space/feature-based region selection and full object recognition. In this context, motion is a very powerful feature, naturally attracting our gaze and yielding rapid and effective shape distinction.

Moving from a pixel-based account of attention to the definition of proto-objects as perceptual units labelled with a single saliency value, we present a framework for the selection of moving objects within cluttered scenes. Through segmentation of motion energy features, the system extracts coherently moving proto-objects defining them as consistently moving blobs and produces an object saliency map, by evaluating bottom-up distinctiveness of each object candidate with respect to its surroundings, in a center-surround fashion.

Keywords: Visual attention model, motion selection, saliency map

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2628

Robust and Efficient Visual SLAM for Spatial Cognition

Andrew Calway (University of Bristol, GB)

It has long been understood that autonomous exploration of previously unseen environments to gain spatial awareness and understanding will be a fundamental capability of future robotic systems.

The development of techniques for simultaneous localisation and mapping (SLAM) - determining world structure and topology whilst simultaneously estimating position - has therefore been at the centre of robotics research for many years. This has been based on a variety of different sensors, including lasers, vision, GPS, odometry and ultrasonics. Of these, vision based techniques have perhaps the greatest potential given the low cost and small form factor of the sensor and the richness of the data.

Recent years have seen significant advances in visual SLAM, driven in the main by the possibility of a highly portable position sensing device, especially in applications such as wearable computing. Several systems now exist which are capable of highly robust and efficient localisation and mapping, operating at rates in excess of standard video frame rates, giving genuine real-time capability. As such, they provide huge potential for advancing sensing capability in robotic systems and further as a catalyst for advancing spatial cognition, especially when coupled with other advances in computer vision, such as object and activity recognition.

In this talk I will summarise some of these advances in visual SLAM, focusing on work carried out at Bristol on robustness to erratic motion, relocalisation in previously built maps, extracting higher order map structure and mechanisms for efficient map representation.

Learning functional object categories and event classes from video

Anthony G. Cohn (University of Leeds, GB)

In this talk I present ongoing work at Leeds, in collaboration with Krishna Sridhar and David Hogg on inducing functional object categories from (qualitative) spatio-temporal descriptions of the participating objects. First we mine event classes from activity descriptions expressed using qualitative spatio-temporal graphs. We presume that events may overlap temporally, and have shared participants. I present two techniques for doing this: a bottup up level-wise technique, and a top down technique which presupposes an event generation model, and where the learning task consist of finding the best explanation of the observed activity under various assumptions. Good explanations are characterised by having few event classes, where each event class is compact (in terms of how similar its instance are) and by a notion we call "interactivity", which measures the extent to which an event consists of highly interacting objects, without "bystanders". Having formed these event classes, object categories can be formed by clustering those objects which take the same roles in a particular event. We have experimented with these techniques in two domains: a kitchen scenario, and aircraft turnovers. I also briefly mention a techniuque for robustly obtaining qualitative spatial relations from noisy video data using a trained HMM.

Keywords: Cognitive Vision, Events, Unsupervised Learning, qualitative spatio-temporal representation

Full Paper:
http://www.comp.leeds.ac.uk/qsr/pub/ecai08.pdf

See also: Learning Functional Object-Categories from a Relational Spatio-Temporal Representation, M. Sridhar, A. G. Cohn and D. C. Hogg ECAI, edited by M. Ghallab, C. D. Spyropoulos, N. Fakotakis and N. M. Avouris, Frontiers in Artificial Intelligence and Applications, 178 , pp 606-610, IOS Press, (2008).

Research with Collaborative Unmanned Aircraft Systems

Patrick Doherty (Linköping University, SE)

We provide an overview of ongoing research which targets development of a principled framework for mixed-initiative interaction with unmanned aircraft systems (UAS). UASs are now becoming technologically mature enough to be integrated into civil society. Principled interaction between UASs and human resources is an essential component in their future uses in complex emergency services or bluelight scenarios. In our current research, we have targeted a triad of fundamental, interdependent conceptual issues: delegation, mixed- initiative interaction and adjustable autonomy, that is being used as a basis for developing a principled and well-de

fined framework for interaction. This can be used to clarify, validate and verify different types of interaction between human operators and UAS systems both theoretically and practically in UAS experimentation with our deployed platforms.

Keywords: Multi-agent systems, robotics, human-robot interaction, delegation

Joint work of: Doherty, Patrick; Kvarnström, Jonas; Heintz, Fredrik; Landen, D.; Olsson, P.-M.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2630

golog.lua: Towards a Non-Prolog Implementation of Golog for Embedded Systems

Alexander Ferrein (University of Cape Town, ZA)

Spatial problems come in a variety of forms: as physical tasks as in the piano movers' problem; or more abstractly as in natural language or in a formal geometric specification; or in some form in between as in a diagram that exhibits aspects of physical space and aspects of abstractions. Similarly, we have a variety of options for solving spatial problems ranging from concrete to abstract. In my talk, I will present different ways in which we can solve spatial problems and I will discuss some of the features of the different approaches.

I will emphasize cognitive aspects of problem solving and I propose selecting reasoning methods with regards to the contexts of their respective task environments. From a cognitive perspective, a particularly relevant method is to use spatial structures for solving spatial problems. I will call this approach Spatial Computing. Spatial Computing enables us solving problems of a certain type in a particularly efficient way. I will address the question whether the principles underlying Spatial Computing can be exploited for the development of special spatial computers; and if so, whether these computers will be restricted to solving spatial problems or whether they may be of more general use. I will suggest that humans employ principles of Spatial Computing for solving problems that go far beyond spatial problems.

Keywords: Spatial computing, right brain hemispere

Attentive Monitoring and Adaptive Control in Cognitive Robotics

Alberto Finzi (University of Napoli, IT)

In this work, we present an attentional system for a robotic agent capable of adapting its emergent behavior to the surrounding environment and to its internal state. In this framework, the agent is endowed with simple attentional mechanisms regulating the frequencies of sensory readings and behavior activations. The process of changing the frequency of sensory readings is interpreted as an increase or decrease of attention towards relevant behaviors and particular aspects of the external environment. In this paper, we present our framework discussing several case studies considering incrementally complex behaviors and tasks.

Keywords: Attention, behavior-based control, robotics

Joint work of: Burattini, E. ; Finzi, Alberto ; Rossi, S. ; Staffa, Maria Carla

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2632

Spatial computing - or how to design a right brain hemisphere

Christian Freksa (Universität Bremen, DE)

Spatial problems come in a variety of forms: as physical tasks as in the piano movers' problem; or more abstractly as in natural language or in a formal geometric specification; or in some form in between as in a diagram that exhibits aspects of physical space and aspects of abstractions.

Similarly, we have a variety of options for solving spatial problems ranging from concrete to abstract. In my talk, I will present different ways in which we can solve spatial problems and I will discuss some of the features of the different approaches. I will emphasize cognitive aspects of problem solving and I propose selecting reasoning methods with regards to the contexts of their respective task environments. From a cognitive perspective, a particularly relevant method is to use spatial structures for solving spatial problems. I will call this approach Spatial Computing. Spatial Computing enables us solving problems of a certain type in a particularly efficient way. I will address the question whether the principles underlying Spatial Computing can be exploited for the development of special spatial computers; and if so, whether these computers will be restricted to solving spatial problems or whether they may be of more general use. I will suggest that humans employ principles of Spatial Computing for solving problems that go far beyond spatial problems.

Keywords: Spatial computing, right brain hemispere

Modeling the observed behavior of a robot through Machine Learning

Malik Ghallab (LAAS - Toulouse, FR)

Artificial systems are becoming more and more complex, almost as complex in some cases as natural systems. Up to now, the typical engineering question was *"how do I design my system to behave according to some specifications"*. However, the incremental design process is leading to so complex artifacts that engineers are more and more addressing a quite different issue of *"how do I model the observed behavior of my system"*. Engineers are faced with the same problem as scientists studying natural phenomena. It may sound strange for an engineer to engage in observing and modeling what a system is doing, since this should be inferable from the models used in the system's design stage. However, a modular design of a complex artifact develops only local models that are combined on the basis of some composition principle of these models; it seldom provides global behavior models.

These general remarks hold in computer sciences throughout several examples of complex systems, ranging from multi-core processors to internet networks. This talk will illustrate the global approach of observation and modeling on the problem of understanding and predicting the behavior of a mobile robot.

Robots are becoming very complex, with a large number of sensory-motor functions combining dozens of actuators and sensors, offering the capabilities of many navigation and manipulation skills, and allowing the execution of sophisticated tasks. The design of these robots usually relies on some reasonable assumptions about the environment and does not model explicitly changing, open-ended environments with human interaction. Hence, a precise observation model of a given robot behavior in a varying and open environment can be essential for understanding how the robot operates within that environment, for predicting its behavior and for improving it.

Machine learning techniques are developed for acquiring the behavior models we are seeking. Three different approaches will be illustrated. In the first approach we learn from experience very robust ways of performing a high-level task such as *"navigate to"*. The designer specifies a collection of *skills* represented as *hierarchical tasks networks*, whose primitives are sensorymotor functions. The skills provide different ways of combining these sensory-motor functions to achieve the desired task. The specified skills are assumed to be complementary and to cover different situations. The relationship between control states, defined through a set of taskdependent features, and the appropriate skills for pursuing the task is learned as a finite observable Markov decision process. This MDP provides a general policy for the task; it is independent of the environment and characterizes the abilities of the robot for the task.

In the second and third approaches, we learn from observations and we model as stochastic automata the behavior of the robot in performing a given task. We use two different techniques:

1. Hidden Markov models, where part of the learning problems are how to acquire the finite observation space and the finite state space;
2. Dynamic Bayes networks, that can be less readable from a user's point of view, but that are used to improve online the robot behavior.

The talk will survey these approach, the tradeoffs, advantages and complexity of each approach, how the robotics experiments have been carried out, and the obtained results. The details of this research pursued jointly with several colleagues and PhD students can be found out in particular in the following publications.

M.FOX, M.GHALLAB, G.INFANTES, D.LONG. Robot introspection through learned hidden Markov models. Artificial Intelligence, 170(2): 59-113, Feb. 2006

B.MORISSET, M.GHALLAB, Learning how to Combine Sensory-Motor Functions into a Robust Behavior. Artificial Intelligence, 172(4-5): 392-412, March 2008

G. INFANTES, F. INGRAND, M. GHALLAB, Learning Behaviors Models for Robot Execution Control. Proc. 17th European Conference on Artificial Intelligence ECAI 2006, Aug. 2006

G.INFANTES , F.INGRAND , M.GHALLAB. Learning behavior models for robot execution control. 16th International Conference on Automated Planning and Scheduling (ICAPS), Anableside (GB), 6-10 June 2006, pp.394-397

Keywords: Robotics, Machine Learning

Stream-Based Reasoning in DyKnow

Fredrik Heintz (Linköping University, SE)

The information available to modern autonomous systems is often in the form of streams. As the number of sensors and other stream sources increases there is a growing need for incremental reasoning about the incomplete content of sets of streams in order to draw relevant conclusions and react to new situations as quickly as possible.

To act rationally, autonomous agents often depend on high level reasoning components that require crisp, symbolic knowledge about the environment. Extensive processing at many levels of abstraction is required to generate such knowledge from noisy, incomplete and quantitative sensor data. We define knowledge processing middleware as a systematic approach to integrating and organizing such processing, and argue that connecting processing components with streams provides essential support for steady and timely flows of information.

DyKnow is a concrete and implemented instantiation of such middleware, providing support for stream reasoning at several levels. First, the formal KPL language allows the specification of streams connecting knowledge processes and the required properties of such streams. Second, chronicle recognition incrementally detects complex events from streams of more primitive events. Third, complex metric temporal formulas can be incrementally evaluated over streams of states.

DyKnow and the stream reasoning techniques are described and motivated in the context of a UAV traffic monitoring application.

Keywords: Knowledge representation, autonomous systems, stream-based reasoning

Joint work of: Heintz, Fredrik; Kvarnström, Jonas; Doherty, Patrick

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2627

Detecting humans activities in video and still images

Vaclav Hlavac (Czech Technical University, CZ)

The method for recognizing human actions based on pose primitives will be presented. In learning mode, the parameters representing poses and activities are estimated from videos. In run mode, the method can be used both for videos or still images. For recognizing pose primitives, we extend a histogram of Oriented Gradient (HOG) based descriptor to better cope with articulated poses and cluttered background. Action classes are represented by histograms of poses primitives. For sequences, we incorporate the local temporal context by means of n-gram expressions. Action recognition is based on a simple histogram comparison. Unlike the mainstream video surveillance approaches, the proposed method does not rely on background subtraction or dynamic features and thus allows for action recognition in still images.

Keywords: Videoanalytics, human activity in video

Joint work of: Hlavac, Vaclav; Thurau, Christian

Fast Replanning

Sven Koenig (USC - Los Angeles, US)

Planning systems for mobile robots often need to operate in domains that are only incompletely known or change dynamically. In this case, they need to replan quickly as their knowledge changes. Replanning from scratch is often very time consuming. In this talk, I will discuss ongoing research by us and others on incremental heuristic search. Incremental search methods reuse information from previous searches to find solutions to series of similar search tasks potentially much faster than is possible by solving each search task from scratch, while heuristic search methods use heuristic knowledge in form of approximations of the goal distances to focus the search and solve search problems potentially much faster than uninformed search methods. I will discuss the theory behind incremental heuristic search and several of its applications. This is joint work with my students and ex-students C. Bauer, D. Furcy, M. Likhachev, Y. Liu, A. Ranganathan and X. Sun.

Collaborative Activity and Human-Robot Interaction

Geert-Jan M. Kruijff (DFKI Saarbrücken, DE)

The tutorial will highlight issues in processing situated dialogue in human-robot interaction. We see such interaction as part of a larger collaborative activity. The human and the robot are engaged in a form of joint activity, and dialogue as (inter)action is part of that. This makes utterance meaning more than "just content." A robot needs to figure out what the utterance refers to, what new information it provides – but also, why somebody said it, and how the robot is supposed to act upon it.

Processing and managing dialogue therefore needs to be integrated with modeling multi-agent beliefs and -intentions. We discuss the implications of environment uncertainty and -dynamics on how such multi-agent models need to be modeled and maintained. This leads to more general implications for cognitive architecture design. Particularly, we focus on the grounding of multi-agent models in various types of short- and long-term memories, and how that grounding can make it possible to use these models as interface between different types of deliberative cognitive processes like planning, dialogue, and reasoning. The tutorial illustrates the discussed approaches on several robot systems.

Keywords: Human-robot interaction, situated spoken dialogue

Full Paper:
http://talkingrobots.dfki.de

Action Languages – Knowledge Representation and Reasoning for Cognitive Robots

Gerhard Lakemeyer (RWTH Aachen, DE)

In this tutorial I will give an overview of how action languages, exemplified by the situation calculus, are capable of capturing the knowledge needed by a cognitive agent to arrive at reasoned decisions about which actions to take. In many ways, the talk is about Ray Reiter's vision who first coined the phrase Cognitive Robotics in 1993 and it is partly based on a recent survey article on Cognitive Robotics in the Handbook of Knowledge Representation by Hector Levesque and myself. The presentation has two parts: in the first part I introduce the situation calculus, emphasizing aspects which are of particular importance to robot control; in the second part I discuss the action programming language Golog, which is based on the situation calculus and has been used in practice to control a variety of robots.

Keywords: Knowledge representation, reasoning about action, cognitive robotics

Cognitive Robotics

Gerhard Lakemeyer (RWTH Aachen, DE)

This chapter is dedicated to the memory of Ray Reiter. It is also an overview of cognitive robotics, as we understand it to have been envisaged by him.1 Of course, nobody can control the use of a term or the direction of research. We apologize in advance to those who feel that other approaches to cognitive robotics and related problems are inadequately represented here.

Joint work of: Lakemeyer, Gerhard; Levesque, Hector, J.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2633

Improving the Performance of Complex Agent Plans Through Reinforcement Learning

Matteo Leonetti (University of Rome "La Sapienza", IT)

Agent programming in complex, partially observable, and stochastic domains usually requires a great deal of understanding of both the domain and the task in order to provide the agent with the knowledge necessary to act effectively. While symbolic methods allow the designer to specify declarative knowledge about the domain, the resulting plan can be brittle since it is difficult to supply a symbolic model that is accurate enough to foresee all possible events in complex environments, especially in the case of partial observability. Reinforcement

Learning (RL) techniques, on the other hand, can learn a policy and make use of a learned model, but it is difficult to reduce and shape the scope of the learning algorithm by exploiting a priori information. We propose a methodology for writing complex agent programs that can be effectively improved through experience.We show how to derive a stochastic process from a partial specification of the plan, so that the latter's perfomance can be improved solving a RL problem much smaller than classical RL formulations. Finally, we demonstrate our approach in the context of Keepaway Soccer, a common RL benchmark based on a RoboCup Soccer 2D simulator.

Joint work of: Leonetti, Matteo; Iocchi, Luca

Keywords: Agent programming, planning, reinforcement learning, semi non-Markov decision process

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2634

Plan recognition in a situation calculus-based agent programming framework

Yves Lesperance (York University, CA)

Plan recognition is capability that can be useful in many contexts, such as producing helpful computer interfaces, monitoring and assisting cognitively impaired people, and detecting suspicious behavior. In this talk, I will present a formal model of plan recognition for inclusion in a cognitive agent programming framework. The model is based on the Situation Calculus and the ConGolog agent programming language. This provides a very rich plan specification language. Our account supports incremental recognition, where the set of hypotheses about matching plans is progressively filtered as more actions are observed. This is specified using a transition system account. The model also supports hierarchically structured plans and recognizes subplan relationships. We will also discuss work in progress on developing a version of the framework that uses a Bayesian approach to compute the probability of matching plan hypotheses. We will present an application of this to building "smarter" non-player characters in video games.

Joint work of: Goultieva, Alexandra; Lesperance, Yves; McClymont, Joshua

Mapping and meaning

Jim Little (University of British Columbia, CA)

of which are suficient to enable robots to understand where they are. Many tasks such as manipulation and search simply require relative localization. The next challenge is to situate the activities and tasks of mobile robots in maps labeled with semantic information.

The labels can be learned from annotated image databases. Spatial relations among objects provide clues towards identifying the kinds of activities that associate with sites in living and working spaces. Visual recognition is becoming more capable of recognizing generic objects so a robot can determine that it is in, for example, a kitchen and therefore will expect food preparation and cooking activities in that locale. Places supply prior information about the events and actions that occur there. I will discuss how place categorization proceeds from object recognition and how it enables robots to act in cooperation with people at home and at work.

Context and place categorization for assistive robotics

Jim Little (University of British Columbia, CA)

Online annotated image databases are increasingly common. We develop a spatial-semantic modeling system that learns object-place relations from such databases, and then apply these relations to real-world tasks. In the home such a system can label novel scenes with place information, as we demonstrate on test scenes drawn from the same source as our training set. We have designed our system for future enhancement of a robot platform that performs state-of- the-art object recognition and creates object maps of realistic environments. In this context, we demonstrate the use of spatial-semantic information to perform clustering and place labeling of object maps obtained from real homes and offices. Then place information feeds back into the robot system to inform an object search planner about likely locations of a query object.

On First-Order Definability and Computability of Progression for Local-Effect Actions and Beyond

Yongmei Liu (Sun Yat-sen University, CN)

In a seminal paper, Lin and Reiter introduced the notion of progression for basic action theories in the situation calculus. Unfortunately, progression is not first-order definable in general. Recently, Vassos, Lakemeyer, and Levesque showed that in case actions have only local effects, progression is firstorder representable. However, they could show computability of the first-order representation only for a restricted class. Also, their proofs were quite involved. In this paper, we present a result stronger than theirs that for local-effect actions, progression is always first-order definable and computable. We give a very simple proof for this via the concept of forgetting. We also show first-order definability and computability results for a class of knowledge bases and actions with non-local effects. Moreover, for a certain class of local-effect actions and knowledge bases for representing disjunctive information, we show that progression is not only firstorder definable but also efficiently computable.

Keywords: Action and change, knowledge representation

Joint work of: Lakemeyer, Gerhard; Yongmei, Liu

Coming up With Good Excuses: What to do When no Plan Can be Found

Bernhard Nebel (Universität Freiburg, DE)

Plan recognition is capability that can be useful in many contexts, such as producing helpful computer interfaces, monitoring and assisting cognitively impaired people, and detecting suspicious behavior. In this talk, I will present a formal model of plan recognition for inclusion in a cognitive agent programming framework. The model is based on the Situation Calculus and the ConGolog agent programming language. This provides a very rich plan specification language. Our account supports incremental recognition, where the set of hypotheses about matching plans is progressively filtered as more actions are observed. This is specified using a transition system account. The model also supports hierarchically structured plans and recognizes subplan relationships. We will also discuss work in progress on developing a version of the framework that uses a Bayesian approach to compute the probability of matching plan hypotheses. We will present an application of this to building "smarter" non-player characters in video games.

Keywords: Planning, knowledge representation

Joint work of: Göbeldecker, Moritz; Keller, Thomas; Eyerich, Patrick; Brenner, Michael; Nebel, Bernhard

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2773

A Constraint-Based Approach for Plan Management in Intelligent Environments

Federico Pecora (University of Örebro, SE)

State of the art robotic and sensor systems can be leveraged to achieve intelligent and proactive service-providing functionalities that are useful in a number of domains, such as assistive workplaces or domotic environments. Two key cognitive capabilities are often essential for achieving such functionalities, namely (1) human activity recognition and (2) planning for controlling robotic actuation devices. In this talk, I argue that in order to be effective, these two cognitive processes must be seamlessly integrated, both at the modeling level and at the reasoning level. Existing solutions are mostly well suited for dealing with one or the other problem is isolation (e.g., Hidden Markov Models and Hierarchical Task Network planning). I will present our current work on an integrated approach to activity recognition, task planning and execution. The approach employs a uniform formalism based on Allen's interval algebra to represent both the criteria for activity recognition and a planning domain for proactive service enactment. Constraint-based temporal reasoning techniques are used in a closed loop with deployed sensors and actuators to seamlessly interleave context deduction and plan generation/execution. This talk will focus on the achievements and expose the open issues entailed by the solutions we have explored so far. These issues include: efficient temporal reasoning algorithms for on-line inference; how to deal with multiple hypotheses in context recognition; combining data-driven approaches with model-driven approaches to obtain a more flexible and adaptable system; and dealing with sensory uncertainty.

Joint work of: Pecora, Federico; Cirillo, Marcello; Brenner, Michael

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2635

Combining Qualitative Modelling and Trial-and-Error Learning for Skill Acquisition

Claude Sammut (Univ. of New South Wales, AU)

Pure reinforcement learning does not scale well to domains with many degrees of freedom and particularly to continuous domains. We introduce a hybrid method in which a symbolic planner constructs an approximate solution to a control problem. Subsequently, a numerical optimisation algorithm is used to refine the qualitative plan into an operational policy. The method is demonstrated on the problem of learning a stable walking gait for a bipedal robot.

Keywords: Planning, Qualitative Reasoning, Reinforcement Learning

Self-Maintenance for Autonomous Robots in the Situation Calculus

Stefan Schiffer (RWTH Aachen, DE)

In order to make a robot execute a given (high-level) program more robustly we want to enable it to take care of its self-maintenance requirements during online execution of this program. This requires the robot to know about the (internal) states of its components, constraints that restrict execution of certain actions and possibly also how to recover from faulty situations. The general idea is to implement a transformation process on the program to be performed based on explicit (temporal) constraints. Afterwards, a 'guarded' execution of the transformed program should result in more robust behavior.

Keywords: Domestic mobile robotics, self-maintenance, robustness

Joint work of: Schiffer, Stefan; Wortmann, Andreas; Lakemeyer, Gerhard

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2627

The GLAIR Cognitive Architecture

Stuart C. Shapiro (SUNY - Buffalo, US)

GLAIR (Grounded Layered Architecture with Integrated Reasoning) is a multi-layered cognitive architecture for cognitive robots operating in real, virtual, or simulated environments.ă The highest layer of the GLAIR Architecture, the Knowledge Layer (KL), is implemented in the SNePS knowledge representation and reasoning system, contains the beliefs of the robot, and is the layer in which conscious reasoning, planning, and act selection is performed.ă The lowest layer of the GLAIR Architecture, the Sensori-Actuator Layer (SAL), contains the controllers of the sensors and effectors of the hardware or software robot.ă Between the KL and the SAL is the Perceptuo-Motor Layer (PML),which grounds the KL symbols in perceptual structures and subconscious actions, contains deictic and modality registers for providing the agent's sense of situatedness in the environment, and handles translation and communication between the KL and the SAL.ă The presentation will concentrate on connecting the reasoning mind with the acting/sensing body, the implementation of embodiment and situatedness, and the grounding of SNePS terms in PML perceptual structures,and registers.

Keywords: Cognitive Robotics, Cognitive Architectures, Embodiment, Situatedness, Symbol Grounding

Joint work of: Shapiro, Stuart C. ; Bona, Jonathan P.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2772

Cognitive robotics, embodied cognition and human-robot interaction

J. Gregory Trafton, Naval Research Lab., US)

In the past few years, we have been building embodied cognitive models. We use embodied representations in both the online and the offline sense (Wilson, 200Z). We use online cognition for "here and now" interactions (Brooks, 1990) while we use offline cognition by using body based representations for thinking even when decoupled from the environment (e.g., Tucker and Ellis, 2001). We build our models to be as cognitively plausible as possible, using cognitive representations, strategies, and procedures with the belief that this makes human-robot interaction easier and stronger. We have a number of online models, including gaze-following (Trafton, Harrison, Fransen, and Bugajska, 2009) and level 1 visual perspective taking (Trafton and Harrison, under review). We also have a model of offline cognition based on Tucker and Ellis' 2001 data (Harrison and Trafton, in press). We believe that our process models of human cognition not only help us under- stand how people perform different tasks but also provide a different approach to human-robot interaction.

Harrison, A. M., and Trafton, J. G. (in press). Cognition for action: an architectural account for grounded interaction. Proceedings of the Cognitive Science Society, 2010.

Trafton and Harrison (under review). Embodied Spatial Cognition. Trafton, J.G., Harrison, A.M., Fransen, B.R., and Bugajska, M. (2009) An embodied model of infant gaze- following. In A. Howes, D. Peebles, R. Cooper (Eds.), 9th International conference on cognitive modeling - ICCM2009, Manch- ester, UK.

Tucker, M., and Ellis, R. (2001) The potentiation of grasp types during visual object categorization. V isual cognition, 8, 769-800.

Wilson, M. (2002). Six views of embodied cognition. Psychonomic Bulletin and Review.

Acquiring scene models through human-robot interaction

Sven Wachsmuth (Universität Bielefeld, DE)

The success of human-robot interaction critically depends on aligned representations of the mutual intentions as well as the surrounding scene. This cannot be computed from purely bottom-up processing. Even if the general goal is clear for a human interaction partner, the variation of utterances how they would instruct a robot is extremely large. In terms of the visual perception of the environment we have to deal with previously unknown objects, places and complex scene structures including tables, chairs, shelfs, etc Even if the general goal is clear for a human interaction partner, the variation of utterances how they would instruct a robot is extremely large. In terms of the visual perception of the environment we have to deal with previously unknown objects, places and complex scene structures including tables, chairs, shelfs, etc.

Thus, the goal of a cognitive robot is to learn as much as possible from the direct interaction with humans in a structured way. Most work in this field concentrates on the speech interpretation and gesture recognition side assuming that a propositional scene representation is available. Less work has been dedicated to the extraction of relevant scene structures that underlies these propositions. Psycholinguistic studies have shown that people mostly use clear hierarchical structures in the verbalization of spatial scene description. Objects are typically put in spatial relation if they have a common supporting structure or are in a direct supporting relation. This information can be exploited to hypothesize the underlying scene structure from verbal descriptions. We generate 3D plane hypotheses from scene objects referenced by the speaker. These define grouping relations between 3D planar patches that are extracted from depth images recorded by a Time-of-Flight camera (Swadzba et al. 2009). Using this scheme scene structures can be acquired through human-robot interaction on a semantic and functionally defined level. Without assuming any pre-known model of the specific room, we show that the system aligns its sensor-based room representation to a semantically meaningful representation typically used by the human descriptor.

Another aspect of such teaching scenarios is that untrained users have an insuficient expectation about successful human-robot interactions. Having the task to teach the robot, there is a large variability in verbalization behavior of laypersons. Thus, human-robot interaction is significantly improved if the system provides dialog structure and engages the human in an explanatory teaching scenario. We specifically target untrained users, who are supported by mixed-initiative interaction using verbal and non-verbal modalities. The principles of dialog structuring are based on a novelty detection and curiousity behavior of the robot on the one side and user clari

cation dialogs on the other side. In the scenario a person needs to teach object names and manipulation grips to the robot. System development is followed and interleaved with an interactive evaluation approach. The system is based on an extensible, event-based interaction architecture. It was shown in a video study that users benfit from the provided dialog structure resulting in a predictable and successful human-robot interaction (Ltkebohle et al. 2009). References

1. Ltkebohle, I., Peltason, J., Schillingmann, L., Elbrechter, C.,Wrede, B., Wachsmuth, S., Haschke, R. , "The Curious Robot - Structuring Interactive Robot Learning", International Conference on Robotics and Automation, Kobe, Japan, IEEE, 14/05/2009.
2. Swadzba, A., Vorwerg, C., Wachsmuth, S., Rickheit, G. , "Computational Model for the Alignment of Hierarchical Scene Representations in Human-Robot Interaction", International Joint Conference on Artificial Intelligence, Pasadena, CA, USA, AAAI Press, pp. 1857-1863, 14/07/2009.

Keywords: Human Robot Interaction, scene understanding, computer vision

Exploiting Spatial and Temporal Flexibility for Exploiting Spatial and Temporal Flexibility for Plan Execution of Hybrid, Under-actuated Systems

Brian C. Williams (MIT Cambridge, US)

Robotic devices, such as rovers and autonomous spacecraft, have been successfully controlled by plan execution systems that use plans with temporal flexibility to dynamically adapt to temporal disturbances. To date these execution systems apply to discrete systems that abstract away the detailed dynamic constraints of the controlled device. To control dynamic, under-actuated devices, such as agile bipedal walking machines, we extend this execution paradigm to incorporate detailed dynamic constraints. Building upon prior work on dispatchable plan execution, we introduce a novel approach to flexible plan execution of hybrid under-actuated systems that achieves robustness by exploiting spatial as well as temporal plan flexibility. To accomplish this, we first transform the high-dimensional system into a set of low dimensional, weakly coupled systems. Second, to coordinate these systems such that they achieve the plan in real-time, we compile a plan into a concurrent timed flow tube description. This description represents all feasible control trajectories and their temporal coordination constraints, such that each trajectory satisfies all plan and dynamic constraints. Finally, the problem of runtime plan dispatching is reduced to maintaining state trajectories in their associated flow tubes, while satisfying the coordination constraints. This is accomplished through an efficient local search algorithm that adjusts a small number of control parameters in real-time. The first step has been published previously; this paper focuses on the last two steps. The approach is validated on the execution of a set of bipedal walking plans, using a high fidelity simulation of a biped.

Joint work of: Hofmann, Andreas G. ; Williams, Brian C.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2774