09341 Abstracts Collection
Cognition, Control and Learning for Robot Manipulation in Human Environments
— Dagstuhl Seminar —

Michael Beetz[1], Oliver Brock[2], Gordon Cheng[3] and Jan Peters[4]

[1] TU München, D
beetzm@in.tum.de
[2] Univ. of Massachusetts - Amherst, USA
oliver.brock@tu-berlin.de
[3] ATR - Kyoto, J
gordon.cheng@ieee.org
[4] MPI für biologische Kybernetik - Tübingen, D
jan.peters@tuebingen.mpg.de

Abstract. From 16.08. to 21.08.2009, the Dagstuhl Seminar 09341 "Cognition, Control and Learning for Robot Manipulation in Human Environments " 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. Mobile manipulation, cognition, control, learning, humanoid robot, unstructured environments

09341 Summary – Cognition, Control and Learning for Robot Manipulation in Human Environments

High performance robot arms are faster, more accurate, and stronger than humans.

Yet many manipulation tasks that are easily performed by humans as part of their daily life are well beyond the capabilities of such robots. The main reason for this superiority is that humans can rely upon neural information processing and control mechanisms which are tailored for performing complex motor skills, adapting to uncertain environments and to not imposing a danger to surrounding humans. As we are working towards autonomous service robots operating and performing manipulation in the presence of humans and in human living and working environments, the robots must exhibit similar levels of flexibility, compliance, and adaptivity.

The goal of this Dagstuhl seminar is to make a big step towards pushing robot manipulation forward such that robot assisted living can become a concrete vision for the future.

In order to achieve this goal, the computational aspects of everyday manipulation tasks need to be well-understood, and requires the thorough study of the interaction of perceptual, learning, reasoning, planning, and control mechanisms.

The challenges to be met include cooperation with humans, uncertainty in both task and environments, real-time action requirements, and the use of tools. The challenges cannot be met by merely improving the software engineering and programming techniques. Rather the systems need built-in capabilities to deal with these challenges. Looking at natural intelligent systems, the most promising approach for handling them is to equip the systems with more powerful cognitive mechanisms.

The potential impact of bringing cognition, control and learning methods together for robotic manipulation can be enormous. This urge for such concerted approaches is reflected by a large number of national and international research initiatives including the DARPA cognitive systems initiative of the Information Processing Technoloy Office, various integrated projects funded by the European Community, the British Foresight program for cognitive systems, huge Japanese research efforts, to name only a few.

As a result, many researchers all over the world engage in cognitive system research and there is need for and value in discussion. These discussions become particularly promising because of the growing readiness of researchers of different disciplines to talk to each other.

Early results of such interdisciplinary crossfertilization can already be observed and we only intend to give a few examples: Cognitive psychologists have presented empirical evidence for the use of Bayesian estimation and discovered the cost functions possibly underlying human motor control. Neuroscientists have shown that reinforcement learning algorithms can be used to explain the role of Dopamine in the human basal ganglia as well as the functioning of the bea brain. Computer scientists and engineers have shown that the understanding of brain mechanisms can result into realiable learning algorithms as well as control setups. Insights from artificial intelligence such as Bayesian networks and the associated reasoning and learning mechanisms have inspired research in cognitive psychology, in particular the formation of causal theory in young children.

These examples suggest that (1) successful computational mechanisms in artificial cognitive systems tend to have counterparts with similar functionality in natural cognitive systems; and (2) new consolidated findings about the structure and functional organization of perception and motion control in natural cognitive systems indicate in a number of cases much better ways of organizing and specifying computational tasks in artificial cognitive systems.

Keywords: Mobile manipulation, cognition, control, learning, humanoid robot, unstructured environments

Joint work of: Beetz, Michael; Brock, Oliver; Cheng, Gordon; Peters, Jan

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2364

Decisional issues for Human-Robot Interaction

Rachid Alami (LAAS - Toulouse, FR)

If the robot really intends to act in human environment and to not be switched off very quickly by unsatisfied, unhappy or offended users, it has really to make drastic efforts to exhibit a behaviour that is not only safe but also legible and acceptable. It has to avoid incongruous or intrusive conduct. We are doing our best to help.

Keywords: Decisional issues for Human-Robot Interaction, Task planning

Investigation of comprehensive and compact control programs for autonomous robots

Michael Beetz (TU München, DE)

I consider one of the key challenges for robot manipulation in human environments to be the investigation of comprehensive and compact control programs for autonomous robots that enable them to perform everyday manipulation tasks and activities flexibly, reliably, and skillfully. Even the seemingly simple task of picking up an object requires a robot to make many informed decisions regarding to where to stand, which hand to use, how to reach for the object, which grasp to apply, where to grasp, how much force to apply, how to lift, where and how to hold the object, etc. This complex decision making gets even more complicated because all decisions might have to be taken regarding the context. Typically a robot should grasp a bottle differently depending on whether it intends to fill a glass or put the bottle away. Sophisticated decision making is also necessary to successfully performing incompletely specified tasks such as "set the table" (without spelling out what has to be put on the table) and to perform task variants such as "setting the table for a children birthday party".

I believe that robot control programs with such capabilities can be realized when we do not try to explicitly code the decision making but rather enable the robot to infer the right decisions from learned models. Achieving this goal requires in my view a new generation of AI methods. The new generation of methods must be action-centric: the purpose of reasoning is to act appropriately. Concepts should be defined in terms of their roles in actions, experience and knowledge should be acquired to a large degree through observation and performance of activities. Realizing this new generation of methods requires new research efforts in many AI fields and the application of these methods to robot manipulation from the very beginning.

Keywords: AI-based autonomous robotics, integration of perception, knowledge processing, learning, reasoning, planning, and manipulation, Cognition for technical systems

For dexterous manipulation look at the human example

Sven Behnke (Universität Bonn, DE)

Humans are not so good at many things, but they excel in their manipulation skills. When constructing dexterous robots for mobile manipulation, it might be a good idea to look at the human example, understand key principles, and transfer these to technical systems.

One of the most striking features of the human manipulation system is hierarchy. Hierarchy can be found in all of its aspects, including perception, planning, and also action control. Another important feature is reactivity. The human manipulation system does not seem to be very deliberative, but plans consist of few steps on each level of abstraction only. The human manipulation system is able to adapt to changing conditions, and it can learn manipulation strategies. Furthermore, human limbs are lightweight and compliant. This contributes to safety.

Above features are in sharp contrast to classical industrial manipulators. Hence, classical performance metrics might not be meaningful any more. Fortunately, there are competitions such as the RoboCup@Home league that allow for a direct comparison of different approaches to robot construction, perception, and behavior control.

Keywords: Biologically inspired information processing, humanoid robots, computer vision, speech processing, and machine learning

Progress in manipulation

Oliver Brock (TU Berlin, DE)

Fundamental progress in manipulation will not be achieved through a specific technological or scientific advancement but instead by developing an understanding of how to suitably compose perception, control, planning, mechanisms, etc.

Keywords: Mobile manipulation, perception, grasping, motion generation, architectures of factorization, computational biology

Understanding through creating

Gordon Cheng (ATR - Kyoto, JP)

I take the view of 'understanding through creating', over the past 10 or so years, my focus has been in realising humanoid robotic systems that are closer to humans. I believe this view shall provide benefits to science, engineering and society. This approach shall also extend to our understanding and the creation of robotic systems for better manipulation in human environments.

Keywords: Humanoid robotics, cognitive systems, brain- machine interfaces, biomimetic of human vision, computational neuroscience of vision, action understanding, human-robot interaction, active vision, mobile robot navigation

Visual perception and control of motion.

Ernst Dickmanns (Universität der Bundeswehr München, DE)

Expectation-based, Multi-focal, Saccadic (EMS-) vision in a closed perception - action loop with one of the knowledge base components encompassing "dynamic maneuvers" (control time histories) for transitions from state A to state B (skills) is considered essential. Both a simultaneous wide field of view (low resolution) and a pointable high-resolution imaging capability have to be available in parallel serving a common instance for interpretation.

Keywords: Visual perception and control of motion

Action-Related Place-Based Mobile Manipulation

Andreas Fedrizzi (TU München, DE)

In mobile manipulation, the position to which the robot navigates has a large influence on the ease with which a subsequent manipulation action can be performed. Whether a manipulation action succeeds depends on many factors, such as the robot's hardware configuration, the controllers the robot uses to achieve navigation and manipulation, the task context, and uncertainties in state estimation.

In this paper, we present 'ARPlace', an action-related place which takes these factors, and the context in which the actions are performed into account. Through experience-based learning, the robot first learns a so-called generalized success model, which discerns between positions from which manipulation succeeds or fails. On-line, this model is used to compute a ARPlace, a probability distribution that maps positions to a predicted probability of successful manipulation, and takes the uncertainty in the robot and object's position into account. In an empirical evaluation, we demonstrate that using ARPlaces for least-commitment navigation improves the success rate of subsequent manipulation tasks substantially.

Keywords: Mobile manipulation, learning, model acquisition, robotics

Joint work of: Fedrizzi, Andreas; Stulp, Freek; Beetz, Michael

See also: Freek Stulp, Andreas Fedrizzi, and Michael Beetz. Action-Related Place-Based Mobile Manipulation. In Proceedings of the International Conference on Intelligent Robots and Systems (IROS). 2009.

Advances in Robotics through Open Source Software

Brian Gerkey (Willow Garage, Inc. - Menlo Park, US)

The key to advancing the state of the art in mobile manipulation (and robotics generally) is to build Open, shareable systems. It must be possible to reproduce results that support claims in scientific papers, by running the authors' own code. To paraphrase Buckheit and Donoho (writing about wavelet analysis): An article or a video describing an experimental robotic system is not the scholarship itself, but merely advertising for the scholarship. The actual scholarship is in the complete software environment that was used to perform the experiment.

Keywords: Robot programming and simulation, path-planning and control, multi-robot coordination

Important issues for building smart robots

Joachim Hertzberg (Universität Osnabrück, DE)

To begin with, I don't think that autonomous service (household) robots and mobile manipulation in a workspace shared between humans and robots is the real issue. It certainly isn't for me. The problems that I am interested in would also help build – no, they would be REQUIRED for building smart Mars, or Alpha Centauri, or city sewer robots, where neither washing dishes nor cooperating in a harmless way with humans is a problem. It is nice to use the dishwashing household robot as the Drosophila in research (and in such a Dagstuhl Seminar, for that matter), but just to state the point: It is neither the issue in scientific results, nor a relevant focus in terms of applications. (Seriously, do you dream of a household robot? I don't.)

Then what IS the issue? It is closed-loop goal-directed perception-and-reasoning-and-action (the three are inseparable) in real world and with bounded resources, given incomplete and inaccurate knowledge and information. I know this is old hats. I also know that each and every part of this picture and some of their combinations have been investigated deeply and with many good local results. Yet we still fail to understand the combination and integration of all the sub-issues. A household robot is a reasonable Drosophila here, because, if taken seriously, it doesn't allow you to cheat and leave out your least favorite subset of the sub-issues.

Yet, having the full picture in mind does not mean you know where to start drawing it and building a household robot (or a Mars, or Alpha Centauri, or city sewer one). Some people say that we don't have the right architecture idea. May be it's that, but to me this makes it sound more technical than it really is – I think our lack of knowledge or imagination is even wider than that, namely, on a level of concepts and ideas. (That is why I think looking at biological cognition from time to time may provide insight, although I am not primarily interested in human, or mouse, or drosophila cognition.)

Yet, waiting for the Grand Theory about closed-loop goal-directed etc. robotics to emerge before continuing to work is clearly no option, so we are probably all working busily on sub-parts of the big picture. I am currently most interested in the question how, in the closed-loop perception-and-reasoning-and-action process, the results of symbolic reasoning influence the perception, i.e., the sensor data acquisition and the data interpretation. To give an idea for a household robot, consider the two rules: "If I have perceived a chair, I should be prepared for sensing a table." and "If I am perceiving a tiled floor, then I am not in the living room, and if my localization says so, then it is mistaken."

Results of symbolic reasoning should influence physical action for taking sensor data (e.g., go to some place to disambiguate some data), the process of knowledge base maintenance (e.g., initiate knowledge revision), the plan-based control of action (e.g., reacting to the fact that some piece of knowledge turns certain or uncertain, respectively) – basically everything else. Just to make sure: That is not to say that "reasoning comes first". In terms of architecture, if you call it so, reasoning is the job of some threads that run in parallel to others (e.g., reactive control), where some of the threads implement or contribute to slow control cycles (reasoning), others fast ones (reaction). The interesting question for me is how the slow control modules doing reasoning would influence the fast reactive ones (and in particular those that have to do with perception), contributing to a goal-directed overall behavior without enslaving or hindering reaction.

Keywords: Artificial Intelligence, autonomous mobile robots, sensor data interpretation, semantic mapping, action planning, plan- based robot control

Embodiment in Robotics

Koh Hosoda (Osaka University, JP)

Embodiment plays a crucial role for realizing adaptive behavior of an agent. For this purpose, therefore, we have to design its morphology vary carefully.

Keywords: Legged Locomotion, Infant Locomotion, Anthropomorphic hand and arm; Adaptive behavior based on Embodiment

Robots as effective collaborators for human endeavors

Odest Chadwicke Jenkins (Brown University - Providence, US)

The guiding vision of my research is to realize robots as effective collaborators for human endeavors. Towards this goal, my research pursues the use of learning and data-driven approaches to robot programming, communication, and manipulation. Manipulation, in particular, is a highly crucial application of our work that aims to enable robots to achieve desired effects in the physical world.

Learning from demonstration (LfD) is a central theme of our efforts for human users to teach autonomous robots to perform manipulation tasks through natural modes of instruction.

Robotics is poised to revolutionize the way society performs physical task similar to how personal computing revolutionized how we manage information. As many have hypothesized, the continuing emergence of robotics complements and analogizes well with the expansive growth of computing technology. Robotics can be seen as a major extension of the Internet beyond virtual/digital environments and into the real/physical world in which we live. Recent trends are indicative of increasingly stronger, faster, safer, and cheaper robot platforms that will eventually become ubiquitous systems and commercial-off-the-shelf products.

Unlike personal computing, two primary issues make robotics a distinct challenging: uncertainty and purpose. The growth of personal computing has been due in large part to the "write local, run global" approach to software development. That is, a program written by a developer (write local) will reliably perform the same way when distributed to users across the world (run global) as for the original developer.

Unfortunately, robotics lacks this property due to the degrees of ambiguity inherent in current robot perception and actuation. Further, the variability in environments, objectives, and expectations across users presents new levels of uncertainty as robots.

Keywords: Robot learning, human-robot interaction, human motion tracking, robotic manipulation

Extracting Planar Kinematic Models Using Interactive Perception

Dov Katz (Univ. of Massachusetts - Amherst, US)

To accomplish a manipulation task, a robot must have some knowledge about the various objects in its environment and how its actions can affect them. With this knowledge, the robot can determine a sequence of actions for achieving its manipulation objective. I believe that the major challenge for developing autonomous manipulation is acquiring this necessary knowledge. It follows that a prerequisite for solving autonomous manipulation is the development of a set of skills for acquiring information about the state of the world, the objects within it, and the effect of one's own actions on the state of the world. Assessing the state of the world, however, can be quite challenging. The robot's environments may contain many objects, each having many properties. Some of these objects may be either partially or completely unobservable due to obstruction or lighting. Moreover, objects can move, disappear or reappear frequently and unexpectedly.

This inherent complexity of real-world environment raises two fundamental questions for modeling real-world environments: what properties of the environment to measure, and how to interpret information that is only partially or indirectly observable. The first question is about choosing what properties of the environment to measure. The state space required to fully specify most unstructured environments is very high-dimensional. This complexity is due to the many objects contained in a typical environment; each with multiple visual, auditory, and tactile properties. It is impossible to acquire all of this information. Robots, therefore, must limit their measurements to aspects of the environment that best approximate the true state of the world, without affecting their ability to accomplish the given task.

The second challenge is about understanding this information. A robot may only be able to obtain partial and noisy observations. There may be some properties that it cannot directly measure. Furthermore, the robot may be missing some necessary context or prior experience to correctly interpret its observations. This challenge raises questions such as what set of pixels constitutes an object? What information can be inferred by considering the relationship between observable objects? What interaction and with which object will reveal new, previously unobservable, information? And how will the robot's actions affect the state of the world?

To answer these two questions, what to measure and how to interpret our measurements, I believe that a new conceptual and algorithmic framework is required. This framework will provide robots with concepts such as temporal relationship, cause and effect, deductive and inductive logic, and the ability to leverage past experience. Finally, it will provide researchers with important insights about the right way to represent manipulation knowledge and skills.

Keywords: Autonomous mobile manipulation in real-world environments

Empiricism in Robotics

Charles C. Kemp (Georgia Institute of Technology, US)

The more time I've spent in robotics, the more I've been drawn to empiricism. The Wikipedia article on empiricism has this nice text, "It is a fundamental part of the scientific method that all hypotheses and theories must be tested against observations of the natural world, rather than resting solely on a priori reasoning, intuition, or revelation." I'm very interested in what the natural world will teach us as we build autonomous mobile manipulators and test them in human environments. I am often surprised by what works and what doesn't. It's not yet clear to me what is hard and what is simple.

Keywords: Autonomous mobile manipulation, human-robot interaction, assistive robots, healthcare robots, bio-inspired approaches, AI

Motion Planning for Constrained Manipulation Tasks

James Kuffner (Carnegie Mellon University - Pittsburgh, US)

One of the grand challenges in artificial intelligence is to create truly "general-purpose" autonomous robots for home, hospital, and office environments. This talk will discuss some of the challenges of motion autonomy and present an overview of some practical automatic motion planning methods for object grasping and manipulation under obstacle and other generalized task constraints. Experimental results on several humanoid platforms around the world will be shown, along with some new efforts in mobile manipulation. Finally, the long-term prospects for the future development of robot autonomy as it relates to search-based AI will be discussed.

Keywords: Motion planning, obstacle avoidance, task constraints, humanoid robots

Using prediction as a key component for robotic manipulation

Alexis Maldonado (TU München, DE)

I am interested in using prediction as a key component for robotic manipulation. An expectation of the sensory stream during actions can not only improve performance (like in feed-forward control), but also be used to detect abnormal conditions. The basic idea is to separate the endogenous and exogenous components of the sensory data. Force/Torque sensing in manipulators makes this finally possible for robotic actuators, so I also look forward to finding force/torque sensors in new robot designs, along with appropriate interfaces that allow fast close-loop control on the velocity and force levels. Having experience with several robotics systems, I have also come to appreciate the open platforms (software and hardware) and their importance for robotics development and benchmarking.

Keywords: Mobile manipulation, robot reaching and grasping, robotic middleware and integration, free (as in freedom) software for robotics

Generality and Simple Hands (preprint of paper to appear, ISRR 2009)

Matthew T. Mason (Carnegie Mellon University - Pittsburgh, US)

While complex hands seem to offer generality, simple hands are more practical for most robotic and telerobotic manipulation tasks, and will remain so for the foreseeable future.

This raises the question: how do generality and simplicity trade off in the design of robot hands? This paper explores the tension between simplicity in hand design and generality in hand function. It raises arguments both for and against simple hands; it considers several familiar examples; and it proposes a concept for a simple hand design with associated strategies for grasping and object localization. The central idea is to use knowledge of stable grasp poses as a cue for object localization. This leads to some novel design criteria, such as a desire to have only a few stable grasp poses. We explore some of the design implications for a bin-picking task, and then examine some experimental results to see how this approach might be applied in an assistive object retrieval task.

Keywords: Robot hands, generality, clutter, grasp, bin picking, object retrieval

Joint work of: Mason, Matthew T.; Srinivasa, Siddhartha; Vazquez, Andres

See also: Proceedings, Int Symp Robotics Research. 2009.

Predictive ability for human motion

Takamitsu Matsubara (Nara Institute of Science and Technology, JP)

Predictive ability for human motion is required for robot systems in human environments!

Keywords: Policy learning, human motion prediction, reinforcement learning

From biology to robots: the RobotCub project

Giorgio Metta (Italian Institute of Technology - Genova, IT)

Until recently, the study of cognition and the neurophysiological basis of human behavior was the subject of quite separate disciplines such as psychology, neurophysiology, cognitive science, computer science, and philosophy, among others. Mental processes were mainly studied in the framework of abstract theories, mathematical models, and disembodied artificial intelligence. It has now become clear that mental processes are strongly entwined with the physical structure of the body and its interaction with the environment.

As such, the study of cognition and intelligence is more and more dependent on the use of physical bodies, and, ultimately, on the use of humanoid robots. The RobotCub humanoid platform represents an opportunity to move this research agenda forward: through open collaboration, on the common theme of embodied cognition, enabled by a shared humanoid platform, and supported by an institute capable of delivering on this vision in the long run.

I will present results on modeling computational processes of the brain in the field of action recognition and speech perception and connect them to the development of the RobotCub project.

I will also show various preliminary porting of these ideas on the RobotCub humanoid robot.

Keywords: Neuroscience, robotics, manipulation

Motor control in humans and robots

Michael Mistry (Disney Research - Pittsburgh, US)

What makes humans capable of elegant and dexterous motion? How can we control our robots to achieve similar performance? Based on observations and studies of motor control in humans, I believe in the following underlying principles: 1) the motor control system learns and uses predictive internal models of dynamics, 2) motion is planned in a task oriented space, utilizing available redundancy to increase robustness, and 3) stiffness can be adjusted to trade off the requirements of accuracy, compliance, and effort. I believe that understanding these principles of control and movement in both humans and robots will be critical if we are ever to achieve robust and safe robot manipulation in real-life human environments. To test our models of these control principles, we use high degree-of-freedom, force-controllable robots, such as the Sarcos Humanoid platform. Force control allows us to apply feed-forward torque commands to compensate for dynamics, as well as to adjust joint stiffness levels to trade off compliance and robustness to modeling errors. We attempt to emulate human-like control and performance on the Sarcos humanoid, by designing task-space, model-based controllers, using inverse kinematics and inverse dynamics algorithms, which take into account full-body dynamics and redundancy, including the under-actuated floating-base.

Keywords: Dexterous motion in humans and humanoid robots, internal models, model-based control, task-space control, redundancy resolution

Manipulation, Vision, and Planning for integrated humanoid system

Kei Okada (University of Tokyo, JP)

Integrated humanoid system of manipulation, vision, planning is the key to realize human assistive tasks in daily environments. Specially, the system capable of recognize failure and plan recovery actions are important. Thus, we are focusing on developing humanoid system with task-relevant knowledge that integrate motion planning, visual recognition and high-level task planning. The system was evaluated on real humanoid in real environments and a subsequent version is build upon the verified system. This developmental approach helps to push forward our researches.

We are also pursuing sensor-ware motion planning, deformable object manipulation and learning tool manipulation from observation as next challenges developmental approach that integrated all We have been working tow along with maintaining the integrated system for HRP2 humanoid robot based on developmental approach.

Keywords: Humanoid, 3D Vision, Robot System

Human-Robot Interaction

Angelika Peer (TU München, DE)

Can knowledge obtained in the analysis of physical human-human interaction transferred to improve human-robot interaction?

Keywords: Robotics, Control, Teleoperation, Haptics, Haptic Interaction: mechatronics, redundant manipulators, telepresence and teleaction, multi-user teleoperation, shared control, human-machine systems, human-robot interaction, human-human interaction, virtual reality, multi-modal human-system interfaces, design of haptic input devices, behavior modelling, intention recognition

Learning approaches for Autonomous Robots

Jan Peters (MPI für biologische Kybernetik - Tübingen, DE)

The only way to create autonomous robots that can deal with the uncertainty inherent to human-inhabited environments is through appropriate robot learning approaches. On the other hand, new machine learning methods will result from this domain as well as an improved the understanding of human motor control.

Keywords: Machine learning, motor skills, robot control, manipulation, locomotion, human motor learning

Coupling Sensing and Acting

Martin Riedmiller (Universität Freiburg, DE)

To get to the next generation of intelligent systems, we have to couple the modules of sensing and the modules of acting in a much tighter way as it currently is.

Keywords: Machine learning, robotics, neural networks, reinforcement learning, autonomous mobile robots, vision based acting

Equipping personal robots with powerful perception systems

Radu B. Rusu (TU München, DE)

My mission is to equip personal robots with powerful perception systems that enable them to perform everyday manipulation in human environments. To this end, I believe that 3D range sensing and semantic map modeling are key aspects for accomplishing this goal.

Throughout the past few years, I have created original algorithms for the processing of 3D data and have tackled difficult problems such as: data registration, surface reconstruction, 3D point feature estimation, robust model fitting, and learning classes of objects from sensed data, to name a few. These contributions have allowed me to investigate and answer questions such as:

- how to enable 3D perception based on dense point clouds to reconstruct objects from partial views in ways that are sufficient for grasping;
- how to enable the basic data driven point cloud processing mechanisms to have on the fly semantic interpretations;
- how can we basically deeply incorporate object detection and categorization mechanisms into the point cloud processing structures;
- how can we build useful functional object based models on the environment that include true meaningful object categorizations useful for personal robotics.

My most current work includes higher level interpretation mechanisms for 3D perception systems that allow robots to get human-like views of the world. The environment and objects present in it are no longer polygonal soups or collection of points, but they become rather meaningful important parts for a robot in the task of fulfilling a given goal objective. This way a kitchen becomes an entity containing tables, and cupboards, and kitchen appliances as static objects.

Cupboards and drawers are considered cuboid structures, with a planar frontal door and a fixture (handle or knob) used to open or close them, and contain movable objects of interest. Tables support objects that are to be manipulable and graspable in a robot's task. The movable objects themselves are created from collections of primitive geometric shapes, and can be decomposed into parts which simplify the problem of picking them up or placing them in a certain way. The fusion of 3D geometry-based methods with artificial intelligence (machine learning) techniques has been one of the key successes of our mapping approach. Tackling all these problems is an extremely complex task, and though it might require a few more years to get robots to see things the way humans do, my personal belief is that we are on the way to having the necessary tools to solve this.

Keywords: 3D Semantic Perception for Mobile Manipulation, Machine Learning

Physical grounding of Robots

Juergen Sturm (Universität Freiburg, DE)

In order to enable robots to robustly perform everyday manipulation tasks, I believe that robots need to be physically better grounded in the real world than they are right now. I think that robots need to be able to continuously learn about the (changing) environment and perceive the outcomes of their own actions from self-observation to become more versatile. Further, I think that many meaningful manipulation challenges that can be used to guide our research can be found in domestic environments.

Keywords: Kinematic model learning, structure selection, graphical models, body schema learning, self-observation in robotics, sensor-motor learning, learning by imitation, tactile sensing

Statistical Relational Learning techniques for autonomous robots

Martijn van Otterlo (K.U. Leuven, BE)

I think it is only natural to conceive of a robotic system as an intelligent machine that perceives the world as consisting of objects and relations among them. In manipulation domains, I think it is very effective and efficient to plan, reason and learn at the level of objects, which can be things that must be manipulated, but also humans, parts of the environment and even aspects of the robot itself. In addition, a robotic system obviously needs a number of other, more lower level, representational layers and operational mechanisms and I think two of the most prominent questions here are i) how information is propagated from low-level sensor data to high-level, cognitive representations and mechanisms (aka, the symbol grounding and anchoring problems) and ii) which parts of a robotic (manipulation) task are especially suitable for planning and learning at the object level.

In order to facilitate answering these questions, as well as to investigate and build systems that can work with higher level representational devices consisting of objects and relations, I think a fruitful synergy could emerge between the field of robotics and that of statistical relational learning. The former shows a recent interest in learning and using powerful representations and, based on these, compact control policies expressed over generic representations of tasks and environments. The latter field has developed a complete gamut of techniques for machine learning in probabilistic, relational domains, including techniques for (un)supervised learning, reinforcement learning and decision-theoretic planning. It is the robotic (manipulation) domain where demands of the former field can be met by techniques of the latter. In particular, the robotic domain poses an excellent set of challenges concerning grounding and anchoring in the context of noisy and low-level sensor data, and concerning how to use inferential mechanisms on learned knowledge in real, probabilistic domains. Older so-called GOFAI approaches that focused on higher-level representations too often did not suffice for such rich environments because they lacked robustness to noise and stochasticity, but I am sure that through the use of state-of-the-art probabilistic logical machine learning techniques, we can move to something like MOFAI (modern old-fashioned artificial intelligence) to tackle interesting and important problems.

Keywords: Decision-theoretic learning, planning and reasoning; (relational) knowledge representation; reinforcement learning, probabilistic logical learning; action formalisms and cognitive architectures; grounding and anchoring

Human biological components for building robots

Patrick van der Smagt (DLR Oberpfaffenhofen, DE)

"To follow nature as my teacher" - that is Patrick van der Smagt's motto, taken from Stravinsky's opera The Rake's Progress. Van der Smagt leads the bionics group of the Institute of Robotics and Dynamics of the German Aerospace Center (DLR). His group focuses on understanding nature, in particular how the human extremities work, the output of which targets DLRŠs planned integrated human hand-arm system, which resembles its human counterpart as closely as possible in terms of dynamics and kinematics.

But why follow a new path in robotics? Although mechanical arms and hands, including the highly-integrated DLR Four-Finger Hand, already have a wide range of applications, they are far from reaching the dexterity and applicability of their human ancestors: they are not as flexible, agile, light-weight and robust, and cannot yet be seen as a reasonable alternative. This is true for actuation, sensing, as well as control.

To match the dexterity of the human upper extremities, van der Smagt's group analyses various aspects of the human biological components (using magnetic resonance tomography, sonography, electrode implantation, biophysics, and other tools) in order to copy their behaviour. When integrated in the DLR hand-arm system, this system can perform in a human-like fashion, allowing for a perfect integration with the human body, but also as a copying robotic companion. Where will robotics go?

Keywords: Biomimetic robotics; bionic learning