Report from Dagstuhl Seminar 22391

Cognitive Robotics

Fredrik Heintz*[1], Gerhard Lakemeyer*[2], and Sheila McIlraith*[3]

1 **Linköping University, SE.** `fredrik.heintz@liu.se`
2 **RWTH Aachen University, DE.** `gerhard@kbsg.rwth-aachen.de`
3 **University of Toronto, CA.** `sheila@cs.toronto.edu`

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 22391 on the topic of "Cognitive Robotics". Cognitive Robotics is concerned with endowing robots or software agents with higher level cognitive functions that involve reasoning, for example, about goals, perception, actions, the mental states of other agents, and collaborative task execution. The seminar is the latest event in a series of events on this topic that were initiated in 1998. With its roots in knowledge representation and reasoning, the program for this seminar was influenced by transformative advances in machine learning and deep learning, by recent advances in human-robot interactions, and by issues that arise in the development of trustworthy cognitive robotic systems. Reflective of this, the seminar featured sessions devoted to the following four themes: cognitive robotics and KR, verification of cognitive robots, human-robot interaction and robot ethics, and planning and learning. Each theme consisted of plenary talks, plenary discussions and working groups resulting in a research road map for the coming years. There was also a poster session where new or published results could be presented by the participants.

The seminar was very successful and well received by the participants thanks to the excellent environment for exchanging ideas provided by Schloss Dagstuhl.

Seminar September 25–30, 2022 – http://www.dagstuhl.de/22391

2012 ACM Subject Classification Computing methodologies → Cognitive robotics; Computing methodologies → Planning and scheduling; Computing methodologies → Machine learning; Human-centered computing → Human computer interaction (HCI); Computer systems organization → Robotics

Keywords and phrases Artificial Intelligence, Knowledge Representation and Reasoning, Cognitive Robotics, Verification, Human-robot Interaction, Robot Ethics, Machine Learning, Planning

Digital Object Identifier 10.4230/DagRep.12.9.200

1 Executive Summary

Fredrik Heintz (Linköping University, SE, `fredrik.heintz@liu.se`*)*
Gerhard Lakemeyer (RWTH Aachen University, DE, `gerhard@kbsg.rwth-aachen.de`*)*
Sheila McIlraith (University of Toronto, CA, `sheila@cs.toronto.edu`*)*

Cognitive Robotics is concerned with endowing robots or software agents with higher level cognitive functions that involve reasoning, for example, about goals, perception, actions, the mental states of other agents, collaborative task execution, etc. This research agenda has historically been pursued by describing, in a language suitable for automated reasoning, enough of the properties of the robot, its abilities, and its environment, to permit it to make high-level decisions about how to act. Such properties were typically encoded by a human, but with recent advances in machine learning, many of these properties, and the

* Editor / Organizer

determination of how to act, can be learned or adapted through experience. This in turn raises the question of how we can ensure that robots, or other intelligent agents, can be constructed in a manner that is compatible with human values and modes of interactions.

The Cognitive Robotics workshop series has been running since 1998 and includes a Dagstuhl Seminar held in 2010. While progress in Cognitive Robotics has undoubtedly been made over the past twenty years, it is fair to say that we are still far away from creating truly cognitive robots. In particular, the years since the previous Dagstuhl Seminar have seen tremendous progress in many areas that touch on the realisation of cognitive robots such as advances in human-robot interaction and machine learning.

This seminar featured sessions devoted to the following four themes:

Cognitive Robotics and KR: While knowledge representation and reasoning (KR) has played a role in robotic systems for many years, for example, by incorporating domain knowledge in the form of description logic-based ontologies or using automated planning systems for high-level robot control, obstacles remain, which prevent today's robots from benefiting from the true potential of KR. In this session we re-visited the state of the art of how KR is used in robotics and discussed challenges and possible benchmark problems that would demonstrate the need and benefit of KR techniques for cognitive robots. The session was organized by Michael Beetz, University of Bremen.

Verification of Cognitive Robots: Verification has been an active research area in formal methods for many years. It is also an important topic when it comes to cognitive robots, especially when it comes to achieving trustworthiness. However, the sheer complexity of the interplay between a robot's hard- and software components makes verification particularly challenging. In this session we discussed where we currently stand in terms of verifying cognitive robots and what challenges lie ahead. The session was organized by Michael Fisher, University of Manchester.

Human-robot Interaction and Robot Ethics: For cognitive robots to be useful in human environments, effective human-robot interaction (HRI) plays a crucial role. Besides the technological challenges such as multi-modal communication, ethical considerations have become more and more important. These range from robots observing norms and conventions to humans viewing robots as moral agents. In this session we discussed the many facets of robot ethics in the context of HRI and identified a number of future challenges and open problems. The session was organized by Matthias Scheutz, Tufts University.

Planning and Learning: While planning and learning have traditionally been separate research tracks in cognitive robotics, recent work has shown how action primitives that form the basis of planning can be learned from data without background knowledge, thus avoiding the need for hand-crafted solutions. In this sessions this work and related proposals were discussed and a roadmap with short- and long-term challenges was drawn up. The session was organized by Hector Geffner, ICREA and Universitat Pompeu Fabra, Spain. The format of the sessions varied and consisted of one or more plenary talks, plenary discussions and/or working groups. Working groups for all four themes discussed challenges and roadmaps for the future, and one representative of each group presented their findings on the last day of the seminar. Besides talks and discussions that centered around the four themes, the seminar also featured two invited talks by Luis Lamb, Universidade Federal Do Rio Grande Do Sul, on neurosymbolic AI and by Jan Peters, TU Darmstadt, on robot learning. In addition, a number of participants gave poster presentations on their research.

The organizers of the seminar wish to thank Schloss Dagstuhl for providing such an excellent environment for exchanging ideas on how to move the field of cognitive robotics forward.

2 Table of Contents

3 Overview of Talks

3.1 Knowledge Representation and Reasoning for Cognition-enabled Robot Manipulation

Michael Beetz (Universität Bremen, DE)

Robotic agents that can accomplish manipulation tasks with the competence of humans have been the holy grail for AI and robotics research for more than 50 years. However, while the fields made huge progress over the years, this ultimate goal is still out of reach. I believe that this is the case because the knowledge representation and reasoning methods that have been proposed in AI so far are necessary but still too abstract. In this talk I propose to endow robots with the capability to mentally "reason with their eyes and hands," that is to internally emulate and simulate their perception-action loops based on photo-realistic images and faithful physics simulations, which are made machine-understandable by casting them as virtual symbolic knowledge bases. These capabilities allow robots to generate huge collections of machine-understandable manipulation experiences, which they can then generalize into commonsense and intuitive physics knowledge applicable to open manipulation task domains. The combination of learning, representation, and reasoning will equip robots with an understanding of the relation between their motions and the physical effects they cause at an unprecedented level of realism, depth, and breadth, and enable them to master human-scale manipulation tasks. This breakthrough will be achievable by combining simulation and visual rendering technologies with mechanisms to semantically interpret internal simulation data structures and processes.

3.2 Online Replanning with Human-in-The-Loop for Non-Prehensile Manipulation in Clutter – A Trajectory Optimization based Approach

Tony Cohn (University of Leeds, GB)

Joint work of RafaelPapalla, Anthony G. Cohn, Mehmet R. Dogar

We are interested in the problem where a number of robots, in parallel, are trying to solve reaching through clutter problems in a simulated warehouse setting. In such a setting, we investigate the performance increase that can be achieved by using a human-in-the-loop providing guidance to robot planners. These manipulation problems are challenging for autonomous planners as they have to search for a solution in a high- dimensional space. In addition, physics simulators suffer from the uncertainty problem where a valid trajectory in simulation can be invalid when executing the trajectory in the real-world. To tackle these problems, we propose an online-replanning method with a human-in-the-loop. This system enables a robot to plan and execute a trajectory autonomously, but also to seek high-level suggestions from a human operator if required at any point during execution. This method aims to minimize the human effort required, thereby increasing the number of robots that can be guided in parallel by a single human operator. We performed experiments in simulation and on a real robot, using an experienced and a novice operator. Our results show a significant increase in performance when using our approach in a simulated warehouse scenario and six robots.

3.3 Joint Perceptual Learning and Natural Language Acquisition for Autonomous Robots

Tony Cohn (University of Leeds, GB)

Joint work of Muhannad Alomari, Fangjun Li, David C. Hogg, Anthony G. Cohn

In this work, the problem of bootstrapping knowledge in language and vision for autonomous robots is addressed through novel techniques in grammar induction and word grounding to the perceptual world. In particular, we demonstrate a system, called OLAV, which is able, for the first time, to (1) learn to form discrete concepts from sensory data; (2) ground language (n-grams) to these concepts; (3) induce a grammar for the language being used to describe the perceptual world; and moreover to do all this incrementally, without storing all previous data. The learning is achieved in a loosely-supervised manner from raw linguistic and visual data. Moreover, the learnt model is transparent, rather than a black-box model and is thus open to human inspection. The visual data is collected using three different robotic platforms deployed in real-world and simulated environments and equipped with different sensing modalities, while the linguistic data is collected using online crowdsourcing tools and volunteers. The analysis performed on these robots demonstrates the effectiveness of the framework in learning visual concepts, language groundings and grammatical structure in these three online settings.

3.4 Verifying Autonomous Systems

Michael Fisher (University of Manchester, GB)

Autonomy represents a step-change in systems development and requires new approaches to system architectures, to systems analysis and to effective usage.

In this presentation, I describe an approach that utilises the modularity and heterogeneity of (robotic) software architectures to provide a hybrid agent architecture. Then, a range of verification techniques can be applied to the different components, from formal verification applied to the core autonomous decision-making through to varieties of testing used in other parts of the system.

Finally, an important component is the use of runtime verification (or runtime monitoring) to check for anomolies and violations. Together, these mechanisms provide a basis for more relaible, transparent, trustworthy and verifiable autonomous systems.

3.5 Top-down Representation Learning for Acting and Planning

Hector Geffner (ICREA and Universitat Pompeu Fabra, ES)

Recent breakthroughs in AI have shown the remarkable power of deep learning and deep reinforcement learning. These developments, however, have been tied to specific tasks, and progress in out-of-distribution generalization has been limited. While it is assumed that these limitations can be overcome by incorporating suitable inductive biases in neural nets, this is left vague and informal, and does not provide meaningful guidance. In this talk, I articulate a different learning approach where representations are learned over domain-independent target languages whose structure and semantics yield a meaningful and strongly biased hypothesis space. The learned representations do not emerge then from biases in a low level architecture but from a general preference for the simplest hypothesis that explain the data. I illustrate this general idea by considering three learning problems in AI planning: learning general actions models, learning general policies, and learning general subgoal structures ("intrinsic rewards"). In all these cases, learning is formulated and solved as a combinatorial optimization problem although nothing prevents the use of deep learning techniques instead. Indeed, learning representations over domain-independent languages with a known structure and semantics provides an account of what is to be learned, while learning representations with neural nets provides a complementary account of how representations can be learned. The challenge and the opportunity is to bring the two approaches together.

3.6 Better Autonomy Through Uncertainty

Nick Hawes (Oxford University, GB)

Due to the challenges of perception and action, and inevitable inaccuracies in world modelling, the results of a robot's interactions with its environment are inherently stochastic. To successfully complete extended missions under such conditions it is therefore essential that autonomous robots use techniques from decision-making under uncertainty to plan goal-directed behaviour. In this talk I will give an overview of our recent work on planning under uncertainty for autonomous robots, drawing examples from mobile service robots, underwater vehicles, and quadrupeds.

3.7 Cognitive Robotics – A KR Perspective

Gerhard Lakemeyer (RWTH Aachen University, DE)

In this overview talk I address some of the main representation and reasoning techniques that have been used in robotic systems. On the representation side, these include simple databases (logical literals), description logics, and geometric or topological maps with semantic annotations. On the reasoning side, we find methods for temporal, spatial, and uncertainty reasoning as well as automated planning techniques. I also touch upon the need for execution monitoring and failure diagnosis. At the end of my talk I briefly introduce the RoboCup Logistics League, where robots interact with machines in a production logistics scenario and which can serve as a benchmark for applying KR in robotics, both in simulation and on real robots.

3.8 Learning and Reasoning in Neurosymbolic AI

Luis Lamb (Universidade Federal Do Rio Grande Do Sul, BR))

Neurosymbolic AI aims to bring together the statistical nature of machine learning and the logical essence of reasoning in AI systems. Recently, leading technology companies and research groups have put forward agendas for the development of the field, as modern AI systems require sound reasoning and improved explainability. In this talk, we highlight Neurosymbolic AI research results that led to applications and novel developments towards building richer AI systems. We summarize how the field evolved over the years and how it can potentially contribute to improved explainability and the effective integration of learning and reasoning in robust AI.

3.9 Learning Grounded Language for Human Interaction

Cynthia Matuszek (University of Maryland, Baltimore County, US)

Robots deployed today largely perform a predefined set of tasks in limited, controlled environments. In order to handle the complexity of human-centric spaces, it is necessary to learn about the world and tasks from human end users, and natural language is a key modality for such learning. Two high level approaches to understanding and learning from such language are, first, learning probabilistic grammars describing the perceptual state of the world and, second, learning directly from speech, without any textual intermediary. This talk describes work on using a combination of language and perceptual data to learn about how people describe objects in the world, with the long-term goal of understanding tasks and instructions presented in natural language by non-specialist end users. The importance of using speech directly is discussed, and the effectiveness of using featurized speech is compared to ASR-based approaches. Using speech not only improves performance on the language grounding task, but also reduces performance differences among different demographic groups, leading to more immediately deployable robotic systems.

3.10 Reward Machines: Formal Languages and Automata for Reinforcement Learning

Sheila McIlraith (University of Toronto, CA)

Reinforcement Learning (RL) is proving to be a powerful technique for building sequential decision-making systems in cases where the complexity of the underlying environment is difficult to model. Two challenges that face RL are reward specification and sample complexity. Specification of a reward function – a mapping from state to numeric value – can be challenging, particularly when reward-worthy behaviour is complex and temporally extended. Further, when reward is sparse, it can require millions of exploratory episodes for an RL agent to converge to a reasonable quality policy. In this talk I'll show how formal languages and automata can be used to represent complex non-Markovian reward functions. I'll present the notion of a Reward Machine, an automata-based structure that provides a normal form representation for reward functions, exposing function structure in a manner that greatly expedites learning. Finally, I'll also show how these machines can be generated via symbolic planning or learned from data, solving (deep) RL problems that otherwise could not be solved.

3.11 Model Learning for Plannning

Christian Muise (Queens University – Kingston, CA)

Model learning can primarily be characterized across three dimensions: (1) the input data format; (2) the output model components; and (3) the priors/partial models that we start with. Here, we explore two settings where model learning for planning has been studied.

First, we detail the Model Acquisition Toolkit (MACQ): a library dedicated to learning action theories from state traces of various forms. Each technique in the library comes with its own priors, but collectively the library provides the most comprehensive treatment to date of extracting action theories from discrete time series data.

The second work explores how strong priors influenced by planning concepts can aid in learning planning models from image pairs alone. By embedding strong notions of action representation into the learning architecture itself, we are able to learn action theories and state representations that can be given to off-the-shelf planners.

These are but two modern examples of how model learning is being explored in the context of planning.

3.12 Hardware Acceleration: Why, What, How, Use Cases?

Bernhard Nebel (Universität Freiburg, DE)

This talk does not report on research results, but rather on perspectives of how hardware acceleration can be exploited for automatic planning. Focusing on RPG-style heuristics, it is sketched how such heuristics estimators can be compiled into sequential circuits for moderately large planning tasks, which opens up the possibility to implement that on standard FPGAs. Since 80-90% of the compute time in planning systems is spent on computing heuristic estimates, this could result in a speedup of one order of magnitude.

3.13 Robot Learning: Quo Vadis?

Jan Peters (TU Darmstadt, DE)

Autonomous robots that can assist humans in situations of daily life have been a long standing vision of robotics, artificial intelligence, and cognitive sciences. A first step towards this goal is to create robots that can learn tasks triggered by environmental context or higher level instruction. However, learning techniques have yet to live up to this promise as only few methods manage to scale to high-dimensional manipulator or humanoid robots. In this talk, we investigate the challenges for robot learning from both the symbolic and subsymbolic perspective! We show how symbols can arise in a robot learning system and can used to further the general application of robot learning. We also discuss how classically disjunct approaches from first order insight can be used as inductive biases for faster learning using the simulation based approach. We describe the work in various robotic scenarios ranging from tactile manipulation to robot juggling.

3.14 HRI and Robot Ethics

Matthias Scheutz (Tufts University – Medford, US)

Robot ethics is no different from bioethics, information ethics, environmental ethics, etc. in that as a technology it has impact on humans human societies. It is different from all other technologies in that AI enables the development and deployment of autonomous systems that perceive their environment and determine their actions without human ado. AI/robot ethics thus raises the question of whether these systems can operate in human societies and interact with humans in a way that is ethical and acceptable to humans, not causing any harm. For this, robots need to be able to learn human norms from observations and instructions and follow them. When norm conflicts arise, they need to be able to determine the best course of action and justify their choices by appealing to principles used for their decisions. How to build a robotic architecture capable of all of this is the main challenge of ethical HRI!

3.15 Active Learning in Risky Environments: Exploring Deep-Sea Volcanoes and Ocean Worlds

Brian Williams (MIT – Cambridge, US)

Machine learning is a powerhouse in information rich environments. However, machine learning remains challenging when data is sparse, is costly to collect, and is dangerous and complex to acquire. As two examples, ocean exploration and subsea inspection use autonomous vehicles to perform information gathering, to answer questions about the environment. In these applications, communication is limited, vehicles need to be autonomous, environments are risky, and resources are constrained.

Our vision is to create systems that answer information queries by performing active learning in risky environments. These systems 1) generate information gathering plans that bound risk, while maximizing information with respect to a set of questions being asked, 2) continuously adapt plans based on what is observed and what remains unanswered and 3) incorporates informative measures and risk within operational plans, at multiple levels of abstraction.

The talk introduces a series of model-based agent programming paradigms that support this process of active learning in risky environments, starting with state and decision-theoretic programming. The talk then focuses on planning and learning methods that are needed to support two new programming paradigms – information theoretic and risk-aware programming. These approaches are demonstrated in the context of a 2019 ocean campaign, to explore the Columbo volcano in the Mediterranean Sea.

4 Poster Presentations

All participants provided a brief oral introduction and an overview of their research at the outset of the 5-day seminar. All participants were also given the opportunity to present their work in a poster session. The following is a list of participants who presented their research as posters.

Mohamed Behery and Gerhard Lakemeyer

Poster: *Assistive Robot Teleoperation Using Phase Switching Behavior Trees*
Authors: Mohamed Behery, Minh Trinh, Christian Brecher, Gerhard Lakemeyer
Related Publications: Not published yet.

Anthony G Cohn

Poster: *A framework for categorising AI evaluation instruments*
Authors: A G Cohn, José Hernández-Orallo, Julius Sechang Mboli, Yael Moros-Daval, Zhiliang Xiang, Lexin Zhou
Related Publications: `https://ceur-ws.org/Vol-3169/paper3.pdf`

Jasmin Grosinger

Poster: *Proactivity*
Author: Jasmin Grosinger
Related Publications: Not published yet.

Till Hofmann and Gerhard Lakemeyer

Poster: *Controlling Golog Programs against MTL Constraints*
Authors: Till Hofmann, Stefan Schupp, Gerhard Lakemeyer
Related Publications: Not published yet.

Mikhail Khodak

Poster: *Learning Algorithms and Learning Algorithms*
Authors: Mikhail Khodak
Related Publications: listed at the bottom of the poster.

Sven Koenig

Poster: *Multi-Agent Path Finding (MAPF) and Its Applications*
Authors: Many, as listed on poster
Related Publications: `http://idm-lab.org/project-p.html`

Yves Lespérance

Poster: *Plan Recognition in a High Level Belief-Based Programming Language*
Authors: Yves Lespérance, Alistair Scheuhammer, Yu Chen, and Petros Faloutsos

Setareh Maghsudi

Poster: *Multi-Agent Reinforcement Learning*
Authors: Setareh Maghsudi
Related Publications: based on several publications as listed in the poster,

Sheila McIlraith

Poster: *LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning*
Authors: Alberto Camacho, Rodrigo Toro Icarte, Toryn Q. Klassen, Richard Valenzano, Sheila A. McIlraith
Related Publications: based on several publications as listed in the poster.

Bernhard Nebel

Poster: *The Complexity of MAPF on Directed Graphs & The Small Solution Hypothesis*
Authors: Bernhard Nebel
Related Publications: *The Small Solution Hypothesis for MAPF on Strongly Connected Directed Graphs is True*, arXiv:2210.04590.

Maayan Shvo

Poster: *Proactive Robotic Assistance via Theory of Mind*
Authors: Maayan Shvo, Ruthrash Hari, Ziggy O'Reilly, Sophia Abolore, Nina Wang, Sheila A. McIlraith
Related Publications: *Proactive Robotic Assistance via Theory of Mind*, IROS 2022.

5 Working groups

The seminar focused on four themes central to cognitive robotics, with one expert among the participants organizing a session around each theme: cognitive robotics and KR (Michael Beetz), verification of cognitive robots (Michael Fisher), HRI and robot ethics (Matthias Scheutz), and planning and learning (Hector Geffner). The format of the sessions varied and consisted of one or more plenary talks, plenary discussions and/or working groups. Working groups for all four themes discussed challenges and roadmaps for the future, and one representative of each group presented their findings on the last day of the seminar: Gerhard Lakemeyer (cognitive robotics and KR), Fredrik Heintz (verification of cognitive robots), Cynthia Matuszek (HRI and robot ethics), Christian Muise (planning and learning). Here is a summary.

5.1 Cognitive Robotics and KR

Knowledge Representation and Reasoning (KR) has been a concern in cognitive robotics for many years, beginning with the robot *Shakey* developed at SRI in the late sixties. While ontological knowledge, formalized using description logics, and automated planning systems, among other things, can be found in many robotic applications, KR has yet to play a central role in building cognitive robots. In this working group, we discussed and collected some of the challenges that remain in order to leverage the true potential of KR for cognitive robotics. The following lists the main findings and recommendations.

5.1.1 Challenges

- How does a robot know when system 2 is needed (meta cognition)?
- Finding suitable open-ended robotic tasks that demonstrate the need for KR.
- Industrial use cases, where humans and robots collaborate during production (issues in planning, HRI).
- Addressing problems with long-tail phenomena, which are best solved with commonsense.
- How to acquire commonsense for specific tasks.
- Standardization of KR formalisms would help with the uptake (as has happened with OWL).
- Creating a NELL (lifelong learned KB) for robots.
- How to control the complexity of a task? Compilation techniques?
- How can a robot be taught like a human or, how to transfer conceptual representations of a human to a robot?
- How to build a system that can perform a task after watching a video that shows how to do it. How to do it with tools different from those in the video.

- How to build systems that can introspect on their own actions and explain what they are doing.
- How to build robots with a theory of mind (going beyond traditional BDI, which does not consider action, perception, failures, uncertainty).

5.1.2 Reasons why KR is not yet central to robotics and possible ways to overcome this

- When working with robots, 90% of the time is spent on things other than KR. For roboticist, KR issues are often an afterthought, while KR people cannot grapple with the complexity of robots.
- KR for robotics is lacking a "playground" such as benchmarks suitable for testing/evaluating implemented systems. (Attempts like RoboCup Logistics in simulation were not taken up by the planning people because of the complexity, see also Multi-Agent Programming)
- Appropriate environments need to be developed (RoboCup?)
- Those need to be spread and advertized via tutorials at the KR and ICAPS conferences.
- Similarly, KR tools need to be created for use by roboticists.

5.1.3 Roadmap (5–10 years)

- Principled approaches to abstraction of perception.
- Goal reasoning for robots.
- Rationalizing existing implemented KR systems like KnowRob.
- A theory of explainable behavior and its realization in cognitive robots.
- A theory of mind for robots.

5.2 Verification of Cognitive Robots

Verification and validation of complex cognitive robots is very challenging and existing methods, mainly from formal verification, can only be applied to relatively simple cases. This section summarizes the challenges, directions for future research and provides a roadmap towards verification of cognitive robots.

5.2.1 Challenges

- Correct-by-design
- End-to-end verification
- Composing verified components into verified systems
- Combining partial/abstract offline verification with complete/detailed online verification
- Minimum assumption verification, combined with a risk model to assess the risk involved in the assumptions, combined with runtime verification of the assumption to get the minimum risk system
- Systematic combination of partial verification and testing, verify those parts that can be verified, and then systematically test the rest
- Verify models that are used by for example solvers
- Understand the limitations of what is verifiable
- How to build systems that can be verified? What architectures enables verification?

- Verifying learning systems
- Verifying systems that interact with people
- Continuous (online) verification of learning and interacting systems

5.2.2 General Direction of Development

- From static deterministic simple environments to dynamic non-deterministic complex adversarial environments
- From one-off large-scale efforts for verifying particular components to systematic methods for verifying components to tools that automate the verification of components
- From simple components to complex components to simple static systems to complex dynamic systems of components to open, dynamic and learning systems-of-systems

5.2.3 Roadmap

- 5-years
  - Develop verified plan verifiers that can verify plan instances
  - Verified solvers, such as planners, which are guaranteed to generate verified solutions
  - Verified skills under (potentially strong) assumptions about sensors and external behaviors
  - Principled combination of testing, off-line verification and on-line verification of static systems
  - Early involvement with regulators to jointly agree on what to verify, also related to translating high-level abstract properties into things that can be quantified and (probabilistically) verified
- 10-years
  - A formal understanding of what can be verified
  - Methods for formally verifying solvers
  - Methods for verifying skills (and other robot behaviors) and reducing the assumptions under which these are guaranteed to work
  - Verified (simple) cognitive robots using (verified) solvers and (verified) skills to achieve non-trivial goals
  - Principled (off-line/on-line) verification of (simple) cognitive robots that improve their behavior over time (learning)

5.3 HRI and Robot Ethics

Human-robot interaction (HRI) with its many facets and interdisciplinary nature is of key importance for cognitive robotics, with ethical concerns playing an important role as well. In this working group challenges for HRI and robot ethics were discussed and collected along several dimensions: humans modeling robots and vice versa, norms, communication and information flow, and proactive behavior. In the following, we summarize our main findings.

5.3.1 Humans modeling robots

- How can we build systems where it is possible for people to have an accurate model of the robot's capabilities and internal state?
- Possibly we will always interpret its behavior or lack thereof as if it were a human.
  - Does it matter if it is human-shaped?

  - We will update our mental model over time to more accurately capture the robot's state.
  - How can we design a robot such that people's model of it is more accurate?
- Maybe a principle of robot design should be to work with the model of the robot that people have, rather than trying to affect that model.
- Transparency – where does the data come from for learning?

5.3.2 Robots modeling humans

- Understanding/demonstrating social behavior:
  - Depends in part on reasoning about plans, beliefs, goals
  - Timing, dialog, cutting in, ...
  - What can we learn from cognitive science interests, e.g., human-human interaction studies?
- Some characteristics can be learned from data, but not all.
  - What are the features that such a representation would need to learn that model humans?
  - It is different if you are learning ethical principles.
- Need to consider roles and role-switching to handle such learning.
  - Speeds up planning in a collaborative setting to have an understanding of acceptable behavior/social norms.

5.3.3 Norms: representing, learning, following them

- What is the best formalism to express norms/ethical principles?
- Need dialog/some capability of learning from being "told."
- Need more general reasoning and more commonsense/general knowledge
  - It depends on how expensive plan changes are, how long the planning horizon is, etc.
- How to learn norms? From observations, instructions, ...?
  - Norms vary in importance, consequences.
  - We learn norms from a variety of mechanisms:
    important things are written down, less important things are told, some things are just learned from demonstrations.
- Important for norm learners: must be able to learn norms online.
  - Can't do a single model and then be done with it.
  - Online learning and online adaptation.
- Do we have to learn norms in context?
  - General vs. specific vs. culturally-modulated norms.
- Challenge: doing online learning, but not trying stupid things that violate social requirements.
  - But children push boundaries to improve understanding.
  - Learn in simulation?
- Concept of risk, balancing information gain with possible seriousness of a transgression.
  - Four choices: be extremely conservative to try to minimize norm violations; watch and see; ask; or try it and see what happens
- What can we do in simulation?
- Norms can be complex/contradictory/overlapping.
  - Learning sufficiently to act appropriately is difficult.
- Need to watch for signals and adjust norms over time.

  - How do you know how to adjust behaviors?
  - How do you recognize signals that you have violated a norm?
- There is a gradation from benign to strong social norms (mild vs. serious), long-term vs short-term – is there a spatio-temporal hierarchy?
  - People violate norms all the time.
- Challenge: what do we start with?
- Challenge: a lot has to happen in parallel; there is a control problem of making the layers of the robot architecture work together with timing.

5.3.4 Communication and information flow

- For HRI, humans and robots need to communicate.
  - Many modalities of interagent communication.
  - Language, legible behavior, ...
- Need some model of information flow that is deliberate on the part of at least one actor.
  - There exists work on recognition and activity/plan recognition–what else is there that robots can learn from passive observation?
- How can we communicate by inferring from behavior?
  - If you act to make your model clear via inference, you are communicating.
  - Some things are also best conveyed via being told, e.g., driving regulations.
- Just conveying information is not enough.

5.3.5 Proactive behavior

- Desirable for robots to be not purely reactive, but
- More of a problem for the robot to get things wrong when assisting than to do nothing.
  - Do people have a charitable view of a robot if the robot meant well but messed up?
  - Apologizing helps.
  - Depends partly on horizon – for how long will it be bad at something before it becomes good/helpful?

5.3.6 Grand and small challenges

- Supermarket:
  - Sub-problem: socially aware spill detector;
  - Sub-problem: getting something from the shelf for someone.
- Polite restaurant server:
  - When to interrupt, how long to leave the table alone, ...
- Shared manipulation/physical HRI:
  - Joint manipulation (putting all the dishes on the trolley);
  - Joint cooking.
- Seeing-eye Spot robot
  - Intelligent disobedience;
  - Epistemic reasoning about human's beliefs, intentions.

5.4 Planning and Learning

Planning and learning have traditionally been two separate research tracks within cognitive robots. Lately, several research groups have started to study the combination and integration of planning and learning. For example learning symbols or primitives from observations. To achieve this, it is important to use the right inductive biases in learning to ground the AI system in the world. The key to complex behavior is being able to compose these into more complex plans or composite behaviors, thus planning based on these learned primitives clearly adds a significant value. This section provides a roadmap to achieve this in the form of three short-term challenges and four longer term challenges.

5.4.1 Roadmap

- Short-term Challenge 1: Bootstrap the knowledge – preliminary information, partial models
  - Examples of success: Given a partial PDDL for any planning model learning approach, and observing an improved performance in acquiring the rest of the model. Easy for aspects like SAT/ASP-based learning, but not so clear for deep learning methods that try to acquire things in an end-to-end way.
  - ETA 4 years
- Short-term Challenge 2: Life-long learning – models will drift, change, no longer be valid, etc
  - Example of success: Reliably able to detect when the model is no longer valid and how it has changed. Detect when new object types or new instances of an object are introduced. Being able to modify existing knowledge/model so that action which failed after world change now succeeds. Taking advantage of past experiences to quickly adapt to new environments.
  - ETA: 4 years
- Short-term Challenge 3: Leveraging our model specifications / formal languages to help traditional learning, e.g. interpretability/explainability: "why did my model do this?"; robustness: ensuring a DL system performs as expected; fairness: detecting biases, establishing and verifying fairness criteria
  - Example of success: Reasoning-based approach to verify / validate the concepts learned by traditional DL systems (e.g., interrogating LLM's for consistent reasoning)
  - ETA: 4 years
- Long-term Challenge 1: Integration of learned dynamics and hand-crafted models. Understanding the aspects of the models learned by agents – aligning / grounding the symbols specified & learned (including grounding language).
  - This challenge includes generating high level plans to deal with all kinds of complex environment which could include those with non-rigid objects (e.g. bed sheets) but also cluttered environments, uncertain environments, environments with other agents who change the environment dynamically.
  - Example of success: Creating plans that involve learning dynamics e.g., folding a bed sheet.
  - ETA: 10+ years
- Long-term challenge 2: HRI-style Model Acquisition. How to ground symbols interactively (with human users) to iteratively build a planning model (including objects, fluents, actions, etc). Aligning agent's internal language to the one used by the human. Extension – ability to align to multiple humans, using different concepts and languages/phrasing.

  - Example of success: Robot being capable of interactively receiving instructions (while clarifying) and performing the task given by human operators. ALFRED may be an initial starting point (language is all pre-known, as are the goals).
  - ETA: 10+ years
- Long-term challenge 3: Exploration based learning – i..e the robot actively exploring the world and trying to perform experiments to learn more about the world, and its capabilities and how actions affect the world.
  - Example of success: Simulated environment to place an egocentric agent in – success measured in properly acquiring a correct (or correct enough) planning model
  - ETA: 10+ years
- Long-term challenge 4: Multi-agent/human collaboration – learning how to collaborate with another agent to perform a task
  - Example of success: learning how to hand over an object, or jointly moving some large object, or collaborating to build some object (one agent holding the work-piece to resist forces such as sawing or drilling being applied by a second agent).
  - ETA: 5+ years

Participants

- Michael Beetz
Universität Bremen, DE
- Mohamed Behery
RWTH Aachen, DE
- Jens Claßen
Roskilde University, DK
- Anthony Cohn
University of Leeds, GB
- Frank Dignum
University of Umeå, SE
- Alexander Ferrein
Fachhochschule Aachen, DE
- Michael Fisher
University of Manchester, GB
- Hector Geffner
UPF – Barcelona, ES
- Jasmin Grosinger
University of Örebro, SE
- Nick Hawes
University of Oxford, GB
- Fredrik Heintz
Linköping University, SE
- Till Hofmann
RWTH Aachen, DE
- Mikhail Khodak
Carnegie Mellon University – Pittsburgh, US
- Sven Koenig
USC – Los Angeles, US
- Gerhard Lakemeyer
RWTH Aachen, DE
- Yves Lesperance
York University – Toronto, CA
- Setareh Maghsudi
Universität Tübingen, DE
- Cynthia Matuszek
University of Maryland, Baltimore County, US
- Sheila McIlraith
University of Toronto, CA
- Christian Muise
Queen's University – Kingston, CA
- Bernhard Nebel
Universität Freiburg, DE
- Tim Niemueller
Intrinsic Innovation – München, DE
- Ron Petrick
Heriot-Watt University – Edinburgh, GB
- Sebastian Sardiña
RMIT University – Melbourne, AU
- Matthias Scheutz
Tufts University – Medford, US
- Stefan Schiffer
RWTH Aachen University, DE
- Maayan Shvo
University of Toronto, CA
- Gerald Steinbauer
TU Graz, AT
- Brian C. Williams
MIT – Cambridge, US