Report from Dagstuhl Seminar 13072

Mechanisms of Ongoing Development in Cognitive Robotics

Edited by

Jacqueline Fagard[1], Roderic A. Grupen[2], Frank Guerin[3], and Norbert Krüger[4]

1 **Université Paris Descartes, FR,** `jacqueline.fagard@parisdescartes.fr`
2 **University of Massachusets – Amherst, US,** `grupen@cs.umass.edu`
3 **University of Aberdeen, GB,** `f.guerin@abdn.ac.uk`
4 **University of Southern Denmark – Odense, DK,** `norbert@mmmi.sdu.dk`

Abstract

In cognitive robotics "ongoing development" refers to the ability to continuously build on what the system already knows, in an ongoing process, which acquires new skills and knowledge, and achieves more sophisticated levels of behaviour. Human infants are possibly the best known demonstrators of this ability; developmental psychology has many results documenting what infants can and cannot do at various ages, however we know very little about the mechanisms underlying the development. On the robotics side, creating a computational system which displays ongoing development is still an unsolved problem. There are major unsolved questions regarding the mechanisms of ongoing development, in both biological and artificial systems; for example: how to transfer existing skills to a new context, how to build on existing skills, and how to represent knowledge (or skills). The primary aim of the seminar was to bring together researchers from two communities (developmental robotics and infant developmental psychology) in order to spawn new collaborative research projects which will advance our scientific understanding of the mechanisms underlying ongoing development (whether in infants or robots). We especially focused on perception, understanding and manipulation skills relating to physical objects in the world, and the skills which infants acquire in approximately the 4-24 months period. The main outcomes of the seminar were ideas about how the communities could work together to advance their respective goals. This requires psychologists to become computer scientists to some degree, and computer scientists to become psychologists. In addition each may need to be willing to help to solve some challenge problems posed by the other community in order to have their challenges tackled in turn.

Seminar 11.–15. February, 2013 – www.dagstuhl.de/13072

1998 ACM Subject Classification I.2 Artificial Intelligence, I.2.0 General: Cognitive simulation, Philosophical foundations, I.2.6 Learning, I.2.9 Robotics, Manipulators, I.2.10 Vision and Scene Understanding

Keywords and phrases Developmental psychology, Infancy, Motor skill development, Perceptual development, Origins of concepts, Developmental robotics, Affordances, Intrinsic motivation, Transfer of skills/knowledge

Digital Object Identifier 10.4230/DagRep.3.2.55

Mechanisms of Ongoing Development in Cognitive Robotics, *Dagstuhl Reports*, Vol. 3, Issue 2, pp. 55–91
Editors: Jacqueline Fagard, Roderic A. Grupen, Frank Guerin, and Norbert Krüger
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Executive Summary

Jacqueline Fagard
Roderic A. Grupen
Frank Guerin
Norbert Krüger

In cognitive robotics "ongoing development" refers to the ability to continuously build on what the system already knows, in an ongoing process, which acquires new skills and knowledge, and achieves more sophisticated levels of behaviour. Human infants are possibly the best known demonstrators of this ability; developmental psychology has many results documenting what infants can and cannot do at various ages, however we know very little about the mechanisms underlying the development. On the robotics side, creating a computational system which displays ongoing development is still an unsolved problem. There are major unsolved questions regarding the mechanisms of ongoing development, in both biological and artificial systems; for example: how to transfer existing skills to a new context, how to build on existing skills, and how to represent knowledge (or skills).

The primary aim of the seminar was to bring together researchers from two communities (developmental robotics and infant developmental psychology) in order to spawn new collaborative research projects which will advance our scientific understanding of the mechanisms underlying ongoing development (whether in infants or robots). We especially focused on perception, understanding and manipulation skills relating to physical objects in the world, and the skills which infants acquire in approximately the 4-24 months period.

Working groups were formed in the areas of (i) transfer of means/skills; (ii) motor skills/manipulation; (iii) concepts/representations; (iv) motivation; (v) visual perception. These discussed gaps between what infants and robots can do and what research might close the gap. In discussion groups the most significant issue that was raised (and discussed at length) was how to get psychologists and roboticists talking together and doing research together, as there seems to exist a wide gap between the communities. It was concluded that there was a need for psychologists to become computer scientists and computer scientists to become psychologists; i.e. that the meeting of the two fields would not happen simply by people getting together in a room, but that the meeting must happen inside individual heads. Furthermore challenge problems were posed by each of the two respective communities; challenges which they would like the other community to work on.

2 Table of Contents

Working Groups

Kevin’s Game

3 Desired Seminar Outcomes and Progress

3.1 Desired Tangible Outcomes

It is the aim that the following outcomes will be pursued after the conclusion of the seminar:

- A "Roadmap Paper" to understand ongoing development by creating a working model.
  - This idea is being taken forward.
  - The idea is not to state times (e.g. achieve in 5, 10 yrs), but milestones, and parallel/serial work.
  - The roadmap should tell one what to do in what order, where to focus, and where to go next.
- Journal special issue
  - TAMD special issue should have deadline coming in October 2013
  - A special section for a psychology journal is also considered (Jeff Lockman)
- An edited book along the lines of the book "Stone Knapping: The Necessary Conditions for a Uniquely Hominin Behaviour (McDonald Institute Monographs) [Hardcover] V. Roux (Author, Editor), B. Bril (Editor)". The idea would be to present a similar volume which addresses the question of what is missing from robots and present in humans which permits ongoing development.
  - Gianluca Baldassarre said this might be difficult on the CS/robotics side, as books are not valued here, hence it is difficult to get people to contribute. In any event this is a long term plan (5-10) years and another meeting of potential contributors would take place first. It was agreed that a sequel to this seminar would be appropriate in three years' time.

3.2 Social Outcomes

It was an aim of the seminar to get new pairs of people working to identify work they want to undertake together: possible papers, projects, psychology experiments, coordination/integration of computational work in different labs.

4 Overview of Talks

4.1 What are intrinsic motivations? A biological and computational perspective

Gianluca Baldassarre (ISTC-CNR – Rome, IT)

Main reference G. Baldassarre, "What are intrinsic motivations? A biological perspective," in Proc. of the 2011 IEEE Int'l Conf. onDevelopment and Learning (ICDL'11), Vol.2, pp. 1–8, IEEE, 2011.
URL http://dx.doi.org/10.1109/DEVLRN.2011.6037367

The concept of "intrinsic motivation", initially proposed and developed within psychology, is gaining an increasing attention within cognitive sciences for its potential to produce open-ended learning machines and robots. However, a clear definition of the phenomenon is not yet available. This presentation aims to clarify what intrinsic motivations are from a biological perspective and from a computational perspective. To this purpose, it first shows how intrinsic motivations can be defined contrasting them to extrinsic motivations from an evolutionary (and engineering) perspective: whereas extrinsic motivations guide learning of behaviours that directly increase fitness (or satisfy the user/designer purposes), intrinsic motivations drive the acquisition of knowledge and skills that contribute to produce behaviours that increase fitness (or user satisfaction) only in a later stage. Given this key difference, extrinsic motivations generate learning signals on the basis of events involving body homeostatic regulations (accomplishment of user purposes), whereas intrinsic motivations generate transient learning signals mainly based on events taking place within the brain itself (or within the controller of the robot/intelligent machine). These ideas are supported by presenting (preliminary) taxonomies and examples of biological mechanisms underlying the two types of motivations, and also by linking them to some of the most commonly used mechanisms proposed by the literature to implement intrinsic motivations in robots and machines.

4.2 What babies do that might be hard for robots to do

Emily W. Bushnell (Tufts University, US)

Joint work of Bushnell, Emily W.; Brugger, Amy; Lariviere, Leslie. A.; Mumme, Donna. L., Sidman, Jason; Yang, Dahe J.
Main reference D.J. Yang, J. Sidman, E.W. Bushnell, "Beyond the information given: Infants' transfer of actions learned through imitation," Journal of Experimental Child Psychology, Vol. 106, Issue 1, pp. 62–81, 2010.
URL http://dx.doi.org/10.1016/j.jecp.2009.12.005

Learning by imitation is a very efficient, prominent, and productive learning mechanism during human infancy. If robots are to learn as infants do, they will have to be built to imitate. However, infant imitation is not a simple, straight- forward process. In this presentation, I discuss some aspects of infant imitation that may be difficult to program into robots. First, infant imitation is "optional" – if babies see a sequence of actions, they may subsequently imitate only some of these behaviors and not others. Research from my lab indicates that infants are more likely to imitate an action when it is causally relevant to achieving a goal than when it is unnecessary to the goal. They are also more likely to imitate an action followed by an effect than an unadorned action, but less so if they already know another way to create that effect. Infant imitation also interacts with their mind-reading abilities; infants will imitate irrelevant and inefficient actions if they perceive social cues that the demonstrator means for them to do so, and they will perform behaviors they perceive as intended by a demonstrator in preference to the behaviors actually observed. Furthermore, the weighting of these various parameters affecting imitation – causal relevance, efficacy, efficiency, social cueing, etc. – is not fixed. A given parameter may override another in one context, whereas in a different context their influence may be reversed. Capturing this flexibility within a robotic learning system may be a challenge.

Infant imitation is also "generative". Work in my lab shows that by 15 months of age, infants robustly transfer actions learned by imitation to new object contexts which have not been demonstrated for them. Such transfer is a developmental and also a phylogenetic achievement; very young infants and non-human primates do not likewise transfer learned responses across object contexts so readily. Transfer from imitation hinges on extracting an action from the observed action-object-effect context, so the action becomes a distinct entity (representation) that may be combined in a grammar-like way with other objects to potentially produce new effects. Thus transfer enhances the value of learning by imitation considerably, as acquired actions serve to guide infants' subsequent exploration so that is more focused and non-random. However, the capacity to transfer also requires some constraints to limit instances of "negative transfer" or overgeneralization. Identifying the biases that both propel and restrain infants' transfer from imitation is a goal for further developmental research, and likewise programming such priors into intelligent machines is a task for roboticists.

References

1 Brugger, A., Lariviere, L. A., Mumme, D. L., and Bushnell, E. W. (2007). Doing the right thing: Infants' selection of actions to imitate from observed event sequences. Child Development, 78, 806 -824.

2 Yang, D.J., Sidman, J., and Bushnell, E. W. (2010) Beyond the information given: Infants' transfer of actions learned through imitation. Journal of Experimental Child Psychology, 106, 62 -81.

3 Yang, D. J., Bushnell, E.W., Buchanan, D. W., & Sobel, D. M. (in press). Infants' use of contextual cues in the generalization of causal actions. Journal of Experimental Child Psychology.

4.3 Learning Language to Describe the Activities in Videos

Paul R. Cohen (University of Arizona – Tucson, US)

Developmental robotics deals with learning fundamental cognitive structures and processes by interacting with the environment over long time frames. I am particularly interested in language learning and learning in service of vision. As specific examples I would describe our work in DARPA's Mind's Eye initiative, where the task is to generate natural language descriptions of surveillance videos; our work on learning deep semantics for spatial language; and a new project called the Bayesian Blackboard, an architecture for integrating top-down and bottom-up processes in a probabilistically sound way.

http://w3.sista.arizona.edu/~cohen/

4.4 Observational learning of tool use: Understanding the goal of the experimenter can enhance infants' learning of a use of a novel tool

Rana Esseily (Université Paris Ouest Nanterre, FR)

Joint work of Esseily, Rana; Rat-Fischer, Lauriane; O'Regan, Kevin; Fagard, Jacqueline
Main reference R. Esseily, L. Rast-Fischer, K. O'Regan, J. Fagard, "Understanding the experimenter's intention improves 16-month-olds' observational learning of the use of a novel tool," Cognitive Development, Vol. 28, Issue 1, pp. 1–9, 2013.
URL http://dx.doi.org/10.1016/j.cogdev.2012.10.001

In the beginning of the second year of life, infants become highly capable at learning by observation new means end actions such as opening a box with one hand to retrieve an object with the other hand (Esseily et al., 2010). However tool use studies show that before the end of the second year, infants fail to learn by observation how to use a tool to retrieve an out of reach object (Fagard et al., 2011). The aim of our studies was to investigate why do 16-month-old infants who have already developed some observational learning capacities, fail to learn by observation a tool use action. We claim that in order to learn by observation a new target action, infants have to understand the goal of that action. Thus, if infants do not understand the goal of using the tool, they will not be able to predict and anticipate the demonstrator's actions and thus to relate those actions (the experimenter pulling the tool) with their consequences (the toy coming within reach). We tested this hypothesis by showing 16-month-old infants an explicit demonstration of the goal of the experimenter before demonstrating the target action. We tested 65 16-month-old infants on a tool use action consisting in grasping a rake-like object to retrieve an out of reach toy. Infants were randomly assigned to one of 5 groups: spontaneous group (spontaneous manipulation of tool use), classic demonstration group (observation of a model performing directly the demonstration of the target action), intention prior to demonstration group (observation of a model showing her goal by stretching her hand toward the toy before performing the demonstration of the target action), and two additional groups to control for local and stimulus enhancement. The results show that infants in the intention prior to demonstration group performed significantly better than infants in all other groups. However the results also show that infants' performance was not perfect and even though infants made a connection between the toy and the tool, the toy was not always successfully retrieved. One of the reasons learning was not perfect can be that the experimenter's goal was not sufficiently enhanced. Thus, in another ongoing study, we aim at making the goal of the experimenter even more salient by providing infants a situation where the goal is incongruous with the action performed, thus attracting their attention to that goal (the experimenter throws away the toy as soon as she retrieves it using the tool). The preliminary results show that the incongruity makes the situation humoristic for some infants and it is precisely those infants who laugh at the demonstration, who learn perfectly the target action; whereas infants who do not laugh, do not learn by observation how to retrieve the toy using the tool. Hypotheses regarding the underlying mechanisms responsible for these results will be discussed.

4.5 Infants' failure to retrieve an out-of-reach toy with a rake: what is lacking until 18 months?

Jacqueline Fagard (Université Paris Descartes, FR)

Main reference L. Rat-Fischer, J.K. O'Regan, J. Fagard, "The emergence of tool use during the second year of life," Journal of Experimental Child Phsychology, Vol. 113, Issue 3, pp. 440-446, 2012.
URL http://dx.doi.org/10.1016/j.jecp.2012.06.001

Both robotics and developmental psychology explore how an organism becomes autonomous, learns new abilities, and builds on these abilities. In other words both investigate the emergence of higher cognitive functions through learning and development, from perception and action. We choose the emergence of tool use in infants to tackle this question. Tools allow one to overcome the limits of one's body in interacting with the environment. In everyday life infants can use a toothbrush or a spoon not long after their first year. At the same age they may even be able to use a rake-like tool to retrieve a toy in an experimental situation if the toy is placed inside the tool, thus if no spatial gap lies between rake and toy and the toy may come by simple contingency as soon as the rake is moved (Bates, Carlsonluden, & Bretherton, 1980; Brown 1990; van Leeuwen, Smitsman, & van Leeuwen, 1994). It is thus amazing to observe that it is not until 18 months that, in normal conditions, infants succeed at using a rake to retrieve a toy when the latter is placed at distance and to the side of the rake, an observation that we did in our longitudinal (6 infants) and cross-sectional studies (60 infants). In both studies infants failed spontaneously but also after demonstrations from an adult of how to use the rake (Rat-Fischer, O'Regan, & Fagard, 2012; Fagard, Rat-Fisher, & O'Regan, 2012). This late success raises the question of what does it take to an infant to learn to use a new tool. What do infants need to use the rake? We see at least five components of success.

1/ Being able to grasp and move the rake: they obviously can do that (they can grasp the rake, bang on the table, throw it away, etc. at 12 months and even earlier).

2/ Being willing to retrieve the toy: even though it looks sometimes that the toy itself is less interesting that the raking of the toy (quite often they give the toy back to the experimenter, like a dog with a ball), they always indicate that they want the toy since pointing toward the toy in a begging gesture is the most frequent behaviour before success. However we have examples of high motivation being efficient to increase the rate of success (food as a toy) and of too high motivation leading to regression (object too much desired leading to crying and fussing)

3/ Knowing that the rake would allow bringing the toy closer (functionality of the rake): if it was the only problem, they would succeed after the first demonstration from the adult, which is not the case.

4/ Knowing where should be positioned the rake (behind the toy).

5/ Being able to precisely position the rake behind the toy: if it was the only problem, they would try hard and fail. This behaviour of near-success is observed very late, and usually gives way to success within the same session.

Among these five components, some are more on the manual control side (1, which is obviously not a limitation at the age tested, and 5), some on the motivational side (2), and some more on the cognitive side (3, 4). To explore further which one of these components is the most limiting constraint, we first showed an infant repeated demonstrations of using a rake to bring an object closer between 9 and 12 months. He never had the opportunity to manipulate the rake himself (pure visual familiarization). We then followed him longitudinally from 12 to 18 months, in the same conditions as for the six infants of our longitudinal study. This was one way of testing whether understanding the functionality of the rake would help the infant succeed before 18 months, despite the lack of manual practice. Results showed that this infant was able to succeed at using the rake to retrieve a toy placed to the side of the rake much earlier than 18 months (some near-success at 12 months, a few successes at 13 and 14 months, many successes at 16 months). Most importantly, as opposed to all infants tested so far, this infant almost never rejected the rake. These results show that repeatedly observing the functionality of the rake helps succeed earlier. This indicates how component 3 is an important limiting constraint for tool use. However, from the observation of this infant, it was also clear that even when he tried to use the rake to bring the toy closer, it was extremely difficult for him to put the rake precisely behind the toy. Thus, components 5 (and may be 4) were also a limiting constraint. The most likely hypothesis, thus, is that it is the combination of all these components which are needed for success. According to Bruner, skill emerges from the addition of sub-routines which are slowly integrated into a successful behaviour (Bruner, 1970). Besides, a negative influence of the cognitive load of a task on the quality of the infant's movement and a negative influence of the motor load of a task on the infant's understanding of the task has been shown and explained by some motor-cognitive trade-off (Boudreau & Bushnell, 1996). This might be understood as a limitation of processing capacities or attentional resources when one or the other components required for success is made more difficult. In the case of tool-use and our experiment of visual familiarization, it is conceivable that helping understand the functionality of the tool (component 3) frees the infant's mind to try hard on where to place the rake (component 4) and how to do it (component 5). To confirm this preliminary result and to compare the impact of the cognitive and manual-control components on success at tool use, we are now comparing two groups of infants, one group with only manual familiarization with the rake alone, and one group with only visual familiarization with the action of retrieving an out-of-reach object with the rake. Both groups are familiarized during five sessions before being tested at 16 months in the same conditions as in our previous studies.

4.6 Mechanisms for Development of Sensory Abstraction

Severin Fichtl (University of Aberdeen, GB)

Joint work of Fichtl, Severin; Alexander, John; Mustafa, Wail; Kraft, Dirk; Jorgensen, Jimmy; Krüger, Norbert; Guerin, Frank

We are currently interested in three areas, all related to ongoing developmental learning in robotics:

1) Sensor differentiation: Sensor differentiation: Infants extend their repertoire of behaviours from initially simple behaviours with single objects to complex behaviours dealing with spatial relationships among objects. We are interested in the mechanisms underlying this development in order to achieve similar development in artificial systems. One mechanism is sensorimotor differentiation, which allows one behaviour to become altered in order to achieve a different result; the old behaviour is not forgotten, so differentiation increases the number of available behaviours. Differentiation requires the learning of both, sensory abstractions and motor programs for the new behaviour; here[1] we focus only on one sensory aspect: learning to recognise situations in which the new behaviour succeeds. We experimented with learning these situations in a realistic physical simulation of a robotic manipulator interacting with various objects, where the sensor space includes the robot arm position data and a kinect based vision system. The mechanism for learning sensory abstractions for a new behaviour is a component in the larger enterprise of building systems which emulate the mechanisms of infant development.

2) Intrinsic Motivation: In order to deal with the realistic and high dimensional environments which we encounter in our Sensor differentiation research we have to apply some strategy in order to render the complex learning problems feasible. A standard approach to decrease complexity and increase convergence speed is dimensionality reduction, which transforms the state space by projecting it to a lower dimensional feature space. In our work, we have developed a variation of intrinsic motivation called Certainty Based Curiosity (CBC)[2] in order to efficiently explore the space to facilitate quick learning. The idea behind CBC is to label samples that are likely to add most information to the model. This is achieved by labelling the sample which the current model is most unsure about how to classify. To label a sample means to perform an action in a given environment and the different samples equate to different actions that are available to the agent. In contrast to other Intrinsic Motivation algorithms, like Intelligent Adaptive Curiosity, it actively reduces the amount of training needed to improve classifiers and predictors.

3) Learning Spatial Object Relations which determine the Outcome of Actions In order

to construct complex plans and to achieve elaborate tasks it is essential for an agent to understand the qualitative structure and spatial relations of the objects in its environment. Our agents' vision system uses kinect or stereo cameras to generate a 3D point cloud of its environment and from this extracts a texlet based representation[3] of the scene. From this texlet representation we extract relevant information about the spatial relation between objects and store this information in form of 2D relation histograms. This information is extracted by calculating certain relations between object texlets. In this work we use two different distance relations to learn spatial relations. First we calculate the absolute distance of two texlets in the $X - Y$ plane, neglecting the difference in height. The other distance we calculate is the difference in height with respect to the texlet of object 2. From labelled histograms we train Random Forest models to recognise spatial relations. Preliminary experiments suggest that this is a valid approach to learning Spatial relations in 3D environments.

References

1 Fichtl, S., Alexander, J., Kraft, D., Jorgensen, J. A., Krüger, N., Guerin, F. (2012) Rapidly learning preconditions for means-ends behaviour using active learning, ICDL

2 Fichtl, S., Alexander, J., Kraft, D., Jorgensen, J. A., Krüger, N., Guerin, F. (2013) Learning object relationships which determine the outcome of actions, Paladyn (Special Issue on Advances in Developmental Robotics), http://dx.doi.org/10.2478/ s13230-013-0104-x

3 Pugeault, N., Wörgötter, F., Krüger, N. (2010) Visual primitives: Local, condensed, and semantically rich visual descriptors and their applications in robotics, International Journal of Humanoid Robotics (Special Issue on Cognitive Humanoid Vision)

4.7 Model-Based Belief Dynamics for Manipulation Planning

Roderic A. Grupen (University of Massachusets – Amherst, US)

This presentation proposes a data-driven computational approach that accumulates both skills and experience. Skills and partial models of the world are the focus of an intrinsically motivated exploration driven by the difference between expectations and observations[Hart]. Examples are presented of a skill hierarchy accumulated by Dexter (the UMass bimanual humanoid) over the course of approximately four days of training using this approach. These skills include:

1. a policy (searchtrack) for searching for and then tracking visual features
2. a policy (reachgrasp) for reaching to and grasping an object
3. a policy (pick-and-place) for putting one object in contact with a second object
4. a few simple assembly policies (stacks of objects)

We introduce a Bayes filter for representing objects in terms of probabilistic models of how these actions cause effects and then formulate plans that optimize the information gain of the learning system[Sen]. The presentation concludes with new demonstrations of this framework configured to discriminate between objects by composing informative sequences of manual and visual actions.

4.8 What can we learn from infants' reaches to out-of-reach objects?

Beata Joanna Grzyb (Universitat Jaume I – Castellon de la Plana, ES)

The knowledge of one's own action capabilities and bodily characteristics plays a crucial role in perceptuo-motor behavior and hence needs to be incorporated, very early in life, in a bodily frame of reference for action. In general, the bodily frame of reference has to be updated throughout life to properly accommodate changes in perceptual, action or cognitive abilities. We investigated how infants' knowledge of their reachable space changes as their capabilities change over a relatively short developmental timescale. Reaching action provides a good measure of infants' body (and space) awareness, since to successfully reach for an object infants need to know not only the distance to the object, but also how far they can reach and lean without losing balance.

Five experiments compared 9- and 12-month olds in reaching tasks to targets at varying distances – manipulating the salience of the objects, the novelty of the motor act via added wrist weights, and the ordering of the target distances (random, near to far, far to near). The results show that older infants, 12-month-olds do not honor in their attempted reaches a boundary between targets at reachable and not reachable distances but reach to targets at patently unreachable distances. For the infants in our empirical studies, it is likely that few of the 9-month olds were walking or "cruising" upright while holding on to a support, but it is highly likely that many of the 12 month olds were walking or spending time in some form of pre-walking activity in an upright posture. Thus, the developmental change in the alignment between attempted and successful reaching distances could be related to the transition to walking.

We extended our Experiment 1 to include infants with different walking abilities: non-walkers, walkers with help, and independent walkers. The results of our extended Experiment show that walkers (with or without help) constantly reached for the nonreachable target, whereas non-walkers reached less showing better alignment of their reaching attempts to the distances they can reach. An examination of reaches to far distances as a function of trial block reveals that all infants reached with high probability the first time the object was presented. The reaches of non-walkers, however, decreased over trial blocks showing a clear adjustment of reaching behavior at the "near boundary" distances in the task. Walkers in contrast persistently reached to far distance regardless of the trial block showing little adjustment of their behavior with failures to make contact at the far distances .

The decision whether to reach or not for an object depends on many cognitive, motivational, social, perceptual and motor factors. Developmental changes in any or several of these components could be central to the present findings. With all these possibilities in mind, we offer three hypotheses as starting points for understanding why 12-month-old infants with more walking experience reach to targets at nonreachable distances. These hypotheses are: (i) the decreased ability to learn from negative outcome while reaching makes infants fine-tune their walking skill, (ii) the processes responsible for integration of different visual depth cues reorganize themselves at the onset of walking so as to incorporate information from self-motion-based depth cues, (iii) the representation of space changes with the onset of walking; near and far space are being integrated with the reaching and walking actions to constitute a coherent space representation. These hypotheses have been modeled and their plausibility subsequently tested in a robotic setup. The results of robot experiments showed that these hypotheses are not mutually exclusive and overlap in underlying mechanisms, providing further evidences that goal directed reaching is a complicated skill with a long and protracted developmental course.

We advocate that new impetus to robotics can be given from these studies aiming at improving the efficacy of contemporary robotic systems. From a pragmatic point of view, a robot should be able to purposefully and consistently interact with its environment, by grounding its skills on the integration of different stimuli. Such skills could be based on building a representation of its nearby environment, representation which can be exploited for more precise and complex interactions with the environment components. The representation of space should be plastic, and change with the acquisition of new motor skills to properly reflect current robot action abilities. A robotic system should be provided with the ability to autonomously build a coherent representation of the environment for purposeful exploration and actuation in both peripersonal and extrapersonal space, through the active interaction with the environment in a similar way as infants do. Such joint studies should advance robotics, and give some insights for further understanding of human cognitive development, and the nature of embodied intelligence more generally.

4.9 The Structure of Knowledge in Development

Frank Guerin (University of Aberdeen, GB)

Joint work of Guerin, Frank; Alexander, John; Fichtl, Severin
Main reference F. Guerin, N. Krüger, D. Kraft, "A Survey of the Ontogeny of Tool Use: from Sensorimotor Experience to Planning," IEEE Transactions on Autonomous Mental Development, Vol. 5, Issue 1, pp. 18–45, 2012.
URL http://dx.doi.org/10.1109/TAMD.2012.2209879

Finding the appropriate representation for knowledge is long-standing difficulty in Artificial Intelligence (AI). Studying infants to get some idea of the types of representations they might be using is one possible way to attack the AI problem. Presumably infant representations are simpler and fewer than adult ones, and may provide a point of entry to understand adult ones. In previous work we analysed the development of infant behaviours and the parallel development of representations; two tracks which seem to bootstrap each other. Because of the close linkage of these two tracks, all representations, or fragments of representations are associated with behaviours, so that the infant knows what can be done with each concept or fragment thereof. The present talk focuses more on the representational track, and in particular identifies the need for structure in these representations so that concepts are constructed from components which (i) allows an infant to focus on facets of a complex concept, and to know the behavioural possibilities which facilitates planning); (ii) can be re-used by other concepts; (iii) can facilitate analogical reasoning via components shared with other concepts; (iv) can facilitate the construction of advanced concepts from components.

4.10 Sensorimotor Loop Simulations for Tool-Use

Verena V. Hafner (HU Berlin, DE)

Joint work of Hafner, Verena V.; Schillaci, Guido; Lara, Bruno
Main reference G. Schillaci, B. Lara, V.V. Hafner, V.V. "Internal Simulations for Behaviour Selection and Recognition," in Proc. of the 3rd Int'l Workshop on Human Behaviour Understanding (HBU'12), LNCS, Vol. 7559, pp. 148–160, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-34014-7_13

In order to chose and perform appropriate actions, one can internally simulate an action and its predicted outcome. We implemented internal models based on pairs of inverse and forward models on a humanoid robot. The models were learned during body babbling. In the specific experiment, two different models were learned: one for the robot reaching an object with its arm, and one for the robot reaching an object with a tool, a stick attached to the robot's arm serving as an elongated end-effector. The robot could thus internally simulate the desired action for a given reaching position, and make a decision of whether to use the tool or not [2]. The same mechanism of internal models is used to recognise actions of others and even to distinguish between self and others [1, 6]. We are currently investigating the use of internal models to recognise human behaviour, e.g. in throwing [3]. The work is related to our previous work on body maps [5] and intrinsic motivation for exploratory learning [4].

References

1 Schillaci, G., Lara, B. and Hafner, V.V. (2012), Internal Simulations for Behaviour Selection and Recognition, in Human Behaviour Understanding 2012, A.A. Salah et al. (Eds.), Lecture Notes in Computer Science, Volume 7559, pp. 148-160.
2 Schillaci, G., Hafner, V. V., Lara, B. (2012), Coupled Inverse-Forward Models for Action Execution Leading to Tool-Use in a Humanoid Robot, Proceedings of the 7th ACM/IEEE International Conference on Human- Robot Interaction (HRI 2012), pp. 231-232, Boston, USA.
3 Frömer, R., Hafner, V.V. and Sommer, W. (2012), Aiming for the bull's eye: throwing investigated with event related brain potentials, Psychophysiology, Volume 49, Issue 3, pages 335-344, Wiley New York.
4 Oudeyer, P.-Y., Kaplan, F., Hafner, V.V. (2007), Intrinsic Motivation Systems for Autonomous Mental Development, IEEE Transactions on Evolutionary Computation, Special Issue on Autonomous Mental Development, Volume: 11, Issue: 2, pp. 265-286
5 Hafner, V.V. and Kaplan, F. (2008), Interpersonal Maps: How to Map Affordances for Interaction Behaviour, In: E. Rome et al. (Eds.): Affordance-Based Robot Control, LNAI 4760, pp. 1-15, Springer-Verlag Berlin Heidelberg
6 Schillaci, G., Hafner, V.V., Lara, B. and Grosjean, M. (2013), Is That Me? Sensorimotor Learning and Self-Other Distinction in Robotics, in Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI 2013), Tokyo, Japan.

4.11 Body schema in humans and animals and how to learn and model it in robots

Matej Hoffmann (Universität Zürich, CH)

Main reference M. Hoffmann, H. Marques, A. Hernandez Arieta, H. Sumioka, M. Lungarella, R. Pfeifer, "Body schema in robotics: a review," IEEE Transactions on Autonomous Mental Development, Vol. 2, Issue 4, pp. 304–324, 2010.
URL http://dx.doi.org/10.1109/TAMD.2010.2086454

The mechanisms that underlie body representations are co-responsible for many of the admiring capabilities of humans: combining information from multiple sensory modalities, controlling complex bodies, adapting to growth, failures, or using tools. These features are also desirable in robots. We review the concept of body schema in robotics. First, we briefly examine application-oriented research: being able to automatically synthesize, extend, or adapt a model of its body gives more autonomy and resilience to a robot. Second, we summarize the research area in which robots are used as tools to verify hypotheses on the mechanisms underlying biological body representations.

Finally, we present a case study, in which we performed a quantitative analysis of sensorimotor flows in a running quadruped robot using tools from information theory (transfer entropy). Starting from very little prior knowledge, through systematic variation of control signals and environment, we show how the agent can discover the structure of its sensorimotor space, identify proprioceptive and exteroceptive sensory modalities, and acquire a primitive body schema.

References

1 M. Hoffmann, H. Marques, A. Hernandez Arieta, H. Sumioka, M. Lungarella, and R. Pfeifer, "Body schema in robotics: a review," *IEEE Trans. Auton. Mental Develop.*, vol. 2 (4), pp. 304–324, 2010.
2 N. Schmidt, M. Hoffmann, K. Nakajima, and R. Pfeifer, "Bootstrapping perception using information theory: Case studies in a quadruped robot running on different grounds," *Advances in Complex Systems J.*, vol. 16, no. 6, 2012.

4.12 Thinking Like A Child: The Role of Surface Similarities in Stimulating Creativity

Bipin Indurkhya (IIIT – Hyderabad, IN)

Main reference B. Indurkhya, "Thinking like a child: the role of surface similarities in stimulating creativity," in Proc. of the AAAI-2013 Spring Symposium Series: Creativity and (Early) Cognitive Development, Stanford University, Palo Alto, California (USA), 2013.
URL http://www.aaai.org/ocs/index.php/SSS/SSS13/paper/view/5725/5924

An oft-touted mantra for creativity is: think like a child. We focus on one particular aspect of child-like thinking here, namely surface similarities. Developmental psychology has convincingly demonstrated, time and again, that younger children use surface similarities for categorization and related tasks; only as they grow older they start to consider functional and structural similarities. We consider examples of puzzles, research on creative problem solving, and two of our recent empirical studies to demonstrate how surface similarities can stimulate creative thinking. We examine the implications of this approach for designing creativity-support systems.

4.13 Affordances, Verbs, Nouns and Adjectives

Sinan Kalkan (Middle East Technical University – Ankara, TR)

Joint work of Kalkan, Sinan; Sahin, Erol; Yuruten, Onur; Uyanık, Kadir Fırat; Çalişkan, Yiğit; Bozcuoğlu, Asil Kaan
Main reference O. Yürüten, K.F. Uyanık, Y. Çalişkan, A. Kaan Bozcuoğlu, Erol Şahin, Sinan Kalkan, "Learning Adjectives and Nouns from Affordances on the iCub Humanoid Robot," in Proc. of the 12th Int'l Conf. on Adaptive Behavior (SAB'12), LNCS, Vol. 7426, pp. 330–340, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-33093-3_33

Learning and conceptualizing word categories in language such as verbs, nouns and adjectives are very important for seamless communication with robots. Along these lines, we linked the notion of affordance proposed by Gibson to (i) conceptualize verbs, nouns and adjectives, and (ii) demonstrate how a robot can use them for several important tasks in Robotics. For verbs, we compare different conceptualization views proposed by Psychologists over the years. Moreover, we show that there is an important underlying distinction between adjectives and nouns, as supported by recent findings and theories in Psychology, Language and Neuroscience.

4.14 Robots, Skills and Symbols

George Konidaris (MIT – Cambridge, US)

Joint work of Konidaris, George; Kuindersma, Scott; Barto, Andrew; Grupen, Roderic; Kaelbling, Leslie; Lozano-Perez, Tomas

My presentation approaches the problem of designing hierarchical control structures for robots that enable high-level symbolic reasoning, while ultimately remaining grounded in low-level sensorimotor control. The central theme of my talk is that the way to build such hierarchies is around learning sensorimotor skills. I first briefly cover my existing work on autonomous robot skill acquisition, which demonstrates that we are beginning to understand how to build robots that can discover skills through solving one task, and transfer them to more effectively solve future tasks. I will then consider the problem of symbolic planning using acquired skills—in particular, the question of which symbols are required to express and evaluate plans composed of sequences of skills. My (preliminary) work in this area shows that symbolic predicates corresponding to the preconditions and effects of the agent's skills are sufficient for task-level planning in any problem, and necessary in some. The immediate implication of this is that a robot's skills, environment and goal directly and completely specify the symbolic representation that it should use for planning. Since this representation is grounded and amenable to learning, a robot can acquire a symbolic representation appropriate for planning from its own experience.

4.15 Remarks to Frank Guerin's talk

Norbert Krueger (University of Southern Denmark – Odense, DK)

As a reply to Frank Guerin's talk, I dwell on four problems connected to developmental robotics:

1. Suitable hardware with enough dexterity and stability;
2. Defining meaningful initial behaviours;
3. Interaction of the behavioral track and representational track;
4. The definition of required prior knowledge.

4.16 Constructing the Foundations of Commonsense Knowledge

Benjamin Kuipers (University of Michigan, US)

Joint work of Kuipers, Benjamin; Pierce, David; Modayil, Joseph; Mugan, Jonathan; Xu, Changhai
URL http://web.eecs.umich.edu/~kuipers/research/whats-new.html

An embodied agent experiences the physical world through low-level sensory and motor interfaces (the "pixel level"). However, in order to function intelligently, it must be able to describe its world in terms of higher-level concepts such as places, paths, objects, actions, goals, plans, and so on (the "object level"). How can higher-level concepts such as these, that make up the foundation of commonsense knowledge, be learned from unguided experience at the pixel level? I will describe progress on providing a positive answer to this question.

This question is important in practical terms: As robots are developed with increasingly complex sensory and motor systems, and are expected to function over extended periods of time, it becomes impractical for human engineers to implement their high-level concepts and define how those concepts are grounded in sensorimotor interaction. The same question is also important in theory: Must the knowledge of an AI system necessarily be programmed in by a human being, or can the concepts at the foundation of commonsense knowledge be learned from unguided experience?

4.17 Building Tool Use from Object Manipulation

Jeffrey J. Lockman (Tulane University, US)

Main reference J.J. Lockman, "A perception-action perspective on tool use development," Child Development, 71(1), pp. 137–144, 2000.
URL http://www.ncbi.nlm.nih.gov/pubmed/10836567

Tool use has long been considered a cognitive advance. In contrast, in our work we suggest that tool use should be considered a problem of perceptuomotor adaptation in which individuals learn over an extended period of time how a tool changes the action possibilities or affordances of the hand.

Specifically, we have been studying the development of object manipulation in infants and how the behaviors involved in object manipulation transition to tool use. Our work indicates that infants adapt to changes in the properties of their hands when holding objects – a key component of tool use. In the second half year, they combine objects and surfaces together selectively, varying the actions that they perform based on the properties of the object in hand and the type of surface that the object contacts. Equally important, they do so when they hold handled objects: infants relate objects located at the end of the handle to surfaces appropriately, even though they are holding the handle and not the object directly.

Likewise, at a motor level, there is continuity in the behaviors that support the emergence of tool use. Our work employing motion tracking technology and kinematics indicates that in the second half year, infants naturally adapt the percussive up-down movements involved in banging in ways that make these actions ideally suited for instrumental tool use. We suggest that through spontaneous and repeated performance of banging behaviors, infants become skilled in controlling these behaviors, easing the transition toward incorporating these behaviors into such instrumental forms of tool use as hammering.

More broadly, we maintain that there is considerable utility in framing the problem of the emergence of tool use as an ongoing process of perceptuomotor adaptation. Such a process-oriented approach not only offers a way of linking the manual behaviors of infants to the tool use behaviors of older children, but also provides a way of viewing tool use as a product of more general perception- action processes that characterize the functioning of all organisms. This approach, in turn, may offer clues for promoting flexibility and learning in artificial agents that are designed for tool use.

4.18 Constructing Space

J. Kevin O'Regan (Université Paris Descartes, FR)

Joint work of O'Regan, J. Kevin; Laflaquière, Alban; Terekhov, Alexander
URL http://www.kevin-oregan.net

Space seems to be given to us a priori, as a container which contains "stuff" like "objects" that can "move". Among the objects are our "bodies", which we can use to "act upon" the objects. These actions obey certain mathematical constraints dictated by the fact that space is three-dimensional and more or less Euclidean. But for our brains such goings-on are only nerve firings, and nerve firings can occur without there being such a thing as space outside the body. So how can the nerve firings lead to space? Evolution may have built our brains to create space, but how can this have come about? What patterns of nerve firing enable this to be done?

The problem is complicated by the fact that sensory receptors do not signal spatial properties directly. For example in vision, distance is confounded with size; position is confounded with eye and body posture. In hearing, distance must be deduced from a combination of intensity and inter-aural time differences. Another problem is that in order to deduce spatial properties of the environment, the brain needs to know something about the body's own spatial structure. And this is signalled by proprioceptive receptors whose outputs are also ambiguous. Finally, some a priori knowledge of body structure would seem to be necessary. So how can space arise from such a magma of neural firings?

When we think carefully about what space really is, we realize that we cannot hope to find space as a feature of the environment that is directly perceived. Space is a construction that allows us to describe our worlds more conveniently. It is a collection of invariants linking neural output to neural input.

Extracting such invariants must allow the brain to define concepts like “body”, “environment”, “action”, “object”, “position”, “movement”, “distance”. Underlying such concepts are further facts like Separability: What I do here is generally not affected by what I do there; Relativity: Objects can be placed in the same spatial relation here as there; Impenetrability: Generally two objects cannot simultaneously occupy the same position; Group structure: some actions done on objects obey certain combinatorial rules independently of what the objects are. All of these notions are a few of many that are aspects of what we call space, but not all may be necessary for animals to function properly. Even humans’ notion of space may not rigorously encompass all these notions.

To understand better what are the basic concepts underlying the notion of space, a way to proceed is to build artificial agents of different degrees of complexity and see what notions of space they require in order to function. In my talk, I will present different agents illustrating different aspects of space, and will speculate how the underlying invariants could be learnt. I will show a naive agent that understands space as a set of “viewpoints from which things can be observed”. I will show how this agent can determine the dimension of this space and acquire its metric properties.

4.19 Learning from multiple motives. A reservoir computing approach

Mohamed Oubbati (Universität Ulm, DE)

Joint work of Oubbati, Mohamed; Palm, Günther
URL http://www.uni-ulm.de/in/neuroinformatik/forschung/neurobotik.html

Intrinsic-Extrinsic motivation can be viewed as another version of the mind-body dualism, such that intrinsic motives (e.g. curiosity) are those of the mind, while extrinsic motives (e.g. Hunger) are those of the body. The pressure exerted by such motives will keep a situated agent on the track to learn how to make trade-off between them in order to maintain its internal equilibrium. We are interested in studying how several motives influences the decision making process of the agent. We propose to integrate the concept of Reservoir Computing within the frame of Adaptive Dynamic Programming so that the agent learns to act and adapt in presence of several sources of reward. A single reservoir maybe trained to estimate several value functions simultaneously. This would be possible, because recurrent networks are able to learn from heterogeneous data, i.e. memory is in the recurrent activation, not only in the synaptic weights. In this way, a single reservoir could be able to cope with the conflicting demands imposed by different rewards.

4.20 Developmental Mechanisms for Autonomous Life-Long Skill Learning in Robots and Humans

Pierre-Yves Oudeyer (INRIA – Bordeaux, FR)

Joint work of Oudeyer, Pierre-Yves; Kaplan, Frédéric; Baranes, Adrien; Hafner, V.; Nguyen, Mai; Stulp, Freek; Lopes, Manuel
Main reference P.-Y. Oudeyer, A. Baranes, F. Kaplan "Intrinsically Motivated Learning of Real-World Sensorimotor Skills with Developmental Constraints," in G. Baldassarre, M. Mirolli, (eds.)., Intrinsically Motivated Learning in Natural and Artificial Systems, Springer, 2013.
URL http://dx.doi.org/10.1007/978-3-642-32375-1_13
URL http://www.pyoudeyer.com/OudeyerBaranesKaplan13.pdf

Developmental robotics studies and experiments mechanisms for autonomous life-long learning of skills in robots and humans. One of the crucial challenges is due to the sharp contrast between the high-dimensionality of their sensorimotor space and the limited number of physical experiments they can make within their life-time. This also includes the capability to adapt skills to changing environments or to novel tasks. To achieve efficient life-long learning in such complex spaces, humans benefit from various interacting developmental mechanisms which generally structure exploration from simple learning situations to more complex ones. I will present recent research in developmental robotics that has studied several ways to transpose these developmental learning mechanisms to robots [4], and which allowed to generate original hypothesis for mechanisms of infant development [2, 5, 7]. In particular, I will present and discuss computational mechanisms of intrinsically motivated active learning, which automatically select training examples of increasing complexity [6, 5, 2], or tasks through goal babbling [1], and their interaction with imitation learning [3], as well as maturation and body growth where the number of sensori and motor degrees-of-freedom evolve through phases of freezing and freeing [4, 7]. I will discuss them both from the point of view of modeling sensorimotor and cognitive development in infants and from the point of view of technology, i.e. how to build robots capable to learn efficiently in high-dimensional sensorimotor spaces.

References

1. Baranes, A., Oudeyer, P-Y. (2013) Active Learning of Inverse Models with Intrinsically Motivated Goal Exploration in Robots, Robotics and Autonomous Systems, 61(1), pp. 49–73. http://dx.doi.org/10.1016/j.robot.2012.05.008.
2. Kaplan F. and Oudeyer P-Y. (2007) In search of the neural circuits of intrinsic motivation, Frontiers in Neuroscience, 1(1), pp. 225–236.
3. Nguyen M., Baranes A. and P-Y. Oudeyer (2011) Bootstrapping intrinsically motivated learning with human demonstrations, in proceedings of the IEEE International Conference on Development and Learning, Frankfurt, Germany.
4. Oudeyer P-Y., Baranes A., Kaplan F. (2013) Intrinsically Motivated Learning of Real-World Sensorimotor Skills with Developmental Constraints, in Intrinsically Motivated Learning in Natural and Artificial Systems, eds. Baldassarre G. and Mirolli M., Springer.
5. Oudeyer P-Y. Kaplan F. and V. Hafner (2007) Intrinsic motivation systems for autonomous mental development, IEEE Transactions on Evolutionary Computation, 11(2), pp. 265–286.
6. Schmidhuber, J. (1991) Curious model-building control systems, in: Proc. Int. Joint Conf. Neural Netw., volume 2, pp. 1458–1463.
7. Stulp F., Oudeyer P-Y. (2012) Emergent Proximo-Distal Maturation with Adaptive Exploration, in Proceedings of IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-Epirob), San Diego, USA.

4.21 Learning Much From Little Experience

Justus Piater (Universität Innsbruck, AT)

Joint work of Piater, Justus; Szedmak, Sandor

Learning about objects and actions upon them should take advantage of previously-acquired knowledge of objects and actions. We introduce a framework for propagating object-action knowledge to new objects via action-specific similarity functions and action parameter transformations that are learned simultaneously from limited experience. The framework is based on a generalized regression algorithm capable of simultaneously learning object-object, object-action and action-action relations. These relations can be quite general and can represent notions such as similarities, parameter transformations or success probabilities.

4.22 What do infants perceive from the spatial relations between objects? Data from 6- to 20-months-old infants

Lauriane Rat-Fischer (Université Paris Descartes, FR)

Joint work of Rat-Fischer, Lauriane; Florean, Cecilia; O'Regan, J. Kevin; Fagard, Jacqueline

From birth, infants have to coordinate vision and action to explore their environment. Around 10 months of age, they start reaching for out-of-reach objects by pulling a string attached to them, or a cloth on which the objects stand. This type of behavior, called means-end behavior, involves a key concept: the notion of spatial connectedness. Psychologists have shown that the presence/absence of a spatial gap between objects influences the performance of infants in such means-end behaviours. Infants are able to identify composite objects as a unique object when both are contiguous and move in a similar way. However, little is known about what infants perceive and understand from the spatial relationship between unmoving objects. When do infants start to consider the spatial connection as a relevant information to identify composite objects? And then, as soon as they understand that two contiguous objects are connected to each other, how do they apply these informations to solve problems involving the retrieval of out-of-reach objects? Are these informations sufficient in situations with more complex spatial relationships? Two behavioral studies [one involving an eyetracker] on infants aged 6 to 20 months gives us more informations on infants' expectations of composite objects, and their perception of spatial connectedness.

4.23 What Robot(ic)s might learn from Children

Helge Ritter (Neuroinformatics Group Faculty of Technology and Cluster of Excellence Cognitive Interaction Technology (CITEC), Bielefeld University, DE)

The world of current robots is very different (and very far) from the world of children: there is a strong bias to "solve tasks", to carry out "useful activities" and to deal with artificial, mostly rigid objects. In comparison, children primarily play, or engage in social behavior.

Their actions often exhibit a low degree of precision, but high variability and the ability to deal with soft and deformable objects.

Therefore, a first strong message from these differences concerns representational biases: we might question our bias for representing actions for geometrically well-defined, mostly rigid objects, and shaping behavior with a strong emphasis on well-defined goals and constraints. One challenge would be to come closer to abilities of children manifested in coping with deformable objects, such as clothes, food, or toys and materials such as plasticine.

A second major aspect are interfaces. Despite a high appreciation for robustness and flexibility, we still deal with rather rigid interfaces in robotics that are akin to "clockworks". This is not only a matter of fact for most mechanical parts of our robots (such as rather rigid arms and grippers), but also the way we use and combine data structures which we can only accomplish by employing a very high degree of precision for their specification. In contrast, the interfaces that we see at work in children appear extremely exible: highly tactile hands, the emergence of language instead of precise codes, the ability to match and compare objects and situations based on patterns such as "Gestalts" and the formation of capabilities such as an effectively useable "theory of mind" to efficiently approximate the complex inner states of agents.

We believe that recent developments open up fruitful directions to come better to grip at least with some of these challenges. Our own work has been focused on a better understanding of how to replicate some of the "interface flexibility" of the human hand in robots, combining compliant movement grasping strategies and approaches for realizing touch and tactile behavior in a range of contexts. We have developed algorithms for dextrous manipulation, such as bimanual unscrewing actions, or folding of paper under real-time visual feedback. Another line of research has been the study of neural network approaches for Gestalt perception with the goal of action coordination based on Gestalt principles. This work is connected with further research lines within CITEC but outside of our group that emphasizes the role of social interaction for learning.

An overarching exciting aspect of "what children can do for robotics" is that they allow us to observe how the capability of cognitive interaction emerges, and how it does so within ressource limits of processing and time that appear parsimonious in comparison to current large scale systems. How this parsimony is achieved is a major open question. An overarching long term goal in CITEC is to bring together several of the above-mentioned research strands into an architecture that allows us to bootstrap cognitive interaction from resource-parsimonious, guided growth and adaptivity, with only parsimonious blueprinting of initial "scaffolds" that direct the development of the system. As it appears, children are a strong existence proof and a great encouragement for successful solutions along such an approach.

4.24 Meta-Morphogenesis theory as background to Cognitive Robotics and Developmental Cognitive Science

Aaron Sloman (University of Birmingham, GB)

How could our minds and the rest of life have come from a cloud of dust?

Since its beginnings, we have made a lot of progress in AI and Cognitive Science in some areas, and done abysmally in others. That's because there are some very deep problems about animal intelligence that have not been solved, and some have not even been noticed by most researchers. These include problems connected with human mathematical competences (e.g. in geometry) and problem solving competences in other animals. There are also some allegedly hard problems that are actually not so hard – for people who have understood advances in virtual machinery, e.g. problems about the evolution, and implementation of qualia and various ways of being self conscious. I'll suggest a (Turing-inspired) strategy for trying to clarify the problems and gain ideas about the solutions produced by evolution that so far surpass anything we have in AI/Robotics. The strategy is to attempt to identify and explain transitions in the evolutionary history of biological information-processing – since microbes – since it is possible that current animals, including humans, still make important use of solutions to old problems, in ways that would not occur to us starting with computers. I call this the Meta-Morphogenesis project – partly inspired by reading Turing's paper on Morphogenesis (1952), and partly because mechanisms that produce new mechanisms, can also produce new mechanisms for producing new mechanisms! I have to select and over-simplify because the topic is far too large for a single talk. It's a project, not results, and this is just a sample. The focus is not physical morphology, or behaviour, but information-processing. My slides expand on these points.

Revised versions of the slides uploaded will be available at
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk107

4.25 Think Like a Child: Creativity, Perceptual Similarity, Analogy and How to Make Adults Think like a Child

Georgi Stojanov (The American University of Paris, FR)

Joint work of Stojanov, Georgi; Indurkhya, Bipin; Roda, Claudia; Kianfar, Dana

We put forward an assumption that creative thinking and creative behavior are an integral part of typical human cognitive development. Therefore, by looking into the early stages of this development, we can learn more about creativity. Conversely: by exploring creative behavior we might be able to learn something about early cognitive development. In addition, we believe that analogy is a core part of the creativity and developmental mechanisms. During the evolution, we have acquired enough innate knowledge which is crucial for bootstrapping the cognitive development in newborns, and continuously extending it mainly via analogical reasoning and behavior.

Some researchers of creativity make a distinction between historical-creativity (H-creativity) and psychological-creativity (P-creativity), which is about small creative deeds, probably new only to the individual performing them. According to our basic assumption, we also hypothesize that they share the same basic cognitive mechanisms, and that creative perception (in viewing an artifact) involves the same mechanisms that are responsible for generating creative artifacts. Moreover, these mechanisms can also be observed during cognitive development: a constant re-conceptualization of one' s understanding of their environment in the process of agent-environment interaction, maturation, and education. If this hypothesis is accepted, then it suggests that by exercising and stimulating creative perception, we can also strengthen the ability to generate creative ideas and artifacts in the individual.

We have re-casted Piaget's theory of cognitive development by describing assimilation and accommodation as progressive reasoning by analogy starting from early analogizing in terms of bodily sensory-motor schemas, to analogies in mature cognitive agents who have developed object representations and language.

One of the consequences of the above would be that if we would able to induce child-like behavior in adults, this would result in increased ability for creative behavior and creative problem solving.

For example it is well known that young children have much shorter attention spans, or that they tend to focus on surface similarities in categorization tasks (see Bipin Indurkhya's abstract). Another property unique for human infants is the pretend play. During a typical episode of pretend play, children detach themselves from the immediate here and now, and pretend that some objects actually represent other objects, depending on their scenario. For example, they may pretend that they are parents, dolls are their children, and that the banana is a phone which they may use to call their children's school. This is an example of analogy at work: seeing given object/situation as something else.

In this context (trying to make adults think like a child), we have designed a methodology where adult subjects (18 to 22 years old) are given a standard creativity test (in our case that was the ideational fluency test); while they were performing the test we were interrupting them with simple tasks irrelevant for the primary one. Our hypothesis was that this would lead to better performance as a result of widening and defocusing their attention. Although our first results were not encouraging — participants who were being interrupted did not show an increase in their ideational fluency, we believe that this methodology (making adults to "think like a child") is a promising one and we are currently working on different experimental scenarios. One explanation may be that interruptions lead to stress? Another may be related to motivation: the subject didn't receive any compensation for participating in the experiment. The literature on the subject shows both positive and negative results in experiments similar to ours.

4.26 What Infants Can Teach Us About The Way We Program Robots

Alexander Stoytchev (Iowa State University, US)

Joint work of Stoytchev, Alexander; Sinapov; Jivko; Sukhoy; Vladimir; Sriffith; Shane
Main reference A. Stoytchev, "Baby Gym For Robots: A New Platform For Testing Developmental Learning Algorithms," In Proc. of the 2011 AAAI Workshop on Lifelong Learning from Sensorimotor Experience, held at the 25th National Conf. on Artificial Intelligence (AAAI), pp. 63–64, 2011.
URL http://www.aaai.org/ocs/index.php/WS/AAAIW11/paper/view/3895

This talk will focus on recent research results that show how a robot can solve multiple tasks based on what it learns during a developmental period similar to a child's play. During this period the robot actively tries to grasp, lift, shake, touch, scratch, tap, push, drop, and crush objects. At the end of this period the robot knows what different objects sound like when they are dropped, feel like when they are squeezed, etc. Because these properties are grounded in the robot's sensorimotor repertoire the robot can autonomously learn, test, and verify its own representations without human intervention. The talk will demonstrate how the robot can use this information to recognize objects, separate objects into functional categories, and even find the odd-one-out in a set of objects. The talk will also demonstrate how the robot can use sensorimotor interactions to bootstrap the development of its visual system in the context of a button-pressing task. Experiments on learning the properties of cup using water will also be presented.

4.27 Unsupervised Discovery of Actions and Action Possibilities

Emre Ugur (ATR – Kyoto, JP)

Joint work of Ugur, Emre; Oztop, Erhan; Sahin Erol
Main reference E. Ugur, E. Oztop, E. Sahin, "Goal emulation and planning in perceptual space using learned affordances," Robotics and Autonomous Systems, 59:7–8, pp. 580–595, 2011.
URL http://dx.doi.org/10.1016/j.robot.2011.04.005

In our framework that is inspired from Developmental Psychology, the robot's discovery of action possibilities is realized in two sequential phases. In the first phase, the robot that initially possesses a basic action and reflex discovers new behavior primitives by exercising the action and by monitoring the changes created in its initially crude perception system. In the second phase, the robot explores a more complicated environment by executing the discovered behavior primitives and using more advanced perception to learn further action possibilities, aka affordances. After learning affordances through self-interaction and self-observation, the robot can make plans to achieve desired goals, emulate end states of demonstrated actions, monitor the plan execution and take corrective actions using the perceptual structures employed or discovered during learning.

This research was partially supported by a contract with the Ministry of Internal Affairs and Communications entitled, 'Novel and innovative R&D making use of brain structures'.

4.28 Do We Need Models to Develop Robot Vision?

Markus Vincze (TU Wien, AT)

Joint work of Wohlkinger, Walter; Aldoma, Aitor; Rusu, Radu Bogdan; Tombari, Frederico; di Stefano, Luigi
Main reference W. Wohlkinger, A. Aldoma, R.B. Rusu, M. Vincze, "3DNet: Large-Scale Object Class Recognition from CAD Models," in Proc. of the 2012 IEEE Int'l Conf. on Robotics and Automation (ICRA'12), pp. 5384–5391, IEEE, 2012.
URL http://dx.doi.org/10.1109/ICRA.2012.6225116
Main reference A. Aldoma, F. Tombari, L. di Stefano, M. Vincze, "A Global Hypotheses Verification Method for 3D Object Recognition," in Proc. of the 12th European Conf. on Computer Vision (ECCV'12), LNCS, Vol. 7574, pp. 511–524, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-33712-3_37

The robots we wish to build work in a human environment. It is know (e.g., studies by M. Land) that humans strongly use models to cope with the complexity of their environment. Hence, it is argued that models play a strong role in vision. Consequently, the work presents attempts to learn models of objects and the environment, understand when models are complete,and then use models to detect target items given a robotics task. It can be shown that the use of models is highly beneficial to improve the robustness of object and object class detection as well as scene segmentation and object tracking.

4.29 Piaget for Robots: Implementing Accommodation and Assimilation in a Machine

Florentin Wörgötter (Universität Göttingen, DE)

Using generative models it is possible to implement the two piagetian mechanisms of Accommodation and Assimilation in a robot. By viewing a human the machine extracts the observed action and compares it to its (earlier acquired) action knowledge. If the action to entirely new it stores it as a whole (accommodation), if it is similar to a known action only the novels sub-aspects are memorized (assimilation). This is possible because of a new type of action representation – the Semantic Event Chain (SEC) – by which the “essence” of an action is extracted. This allows the machine to distinguish between known and unknown actions. Altogether this leads to a substantial speed-up of learning and supplements other learning mechanisms (e.g. learning by exploration) in an efficient way.

5 Working Groups

Working groups met on the Tuesday.

5.1 Group on transfer of means/skills

(Konidaris, Bushnell, Liu, Alexander, Ugur)

We were mostly in agreement that the our focus should be on policy transfer, where an agent learns a policy in one task that can be redeployed in other tasks. We felt the most interesting directions involved learning policy libraries and then using the agent’s context to determine which might be applicable to new tasks. Challenges here included determining which of the many features available to the agent could be useful for this prediction, whether we could infer that features used to predict the usefulness of some actions might also be likely to predict the usefulness of others, how to include structured bias into this process (in the form of both motion primitives and prior knowledge about the way the world works) and how to avoid negative transfer. This discussion was partly motivated by Prof. Bushnell’s example of babies both under- and over-generalizing learned skills.

5.2 Group on motor skills/manipulation

(J. Lockman, B. Kuipers, A. Sloman, R. Grupen, J. Fagard, E. Ugur, L. Rat-Fischer)

Interesting time points for infant development

1. Fetal period – birth
2. Reaching – grasping (5-7 mo)
3. Manipulation (6-12 months)
4. Means-end, multiple-step actions (7-24 months) (No time to deal with this one)

5.2.1 Fetal period

Infant development: To know motor primitives, it is important to start from long before birth: fetuses start moving at 8 weeks, then show general movement, then isolated movements at 10 weeks, then hand-head touching a week later, then soon thumb sucking, and finally thumb sucking with anticipation of mouth opening, all develops within a few weeks. Thus sensorimotor contingencies detection starts several months before birth and we should try to understand the very beginning. (why thumb sucking in particular? The tip of the thumb and the mouth have a particularly high density of sensors; prepare the infant for being breastfed; the space is limited in utero; target toward body always easier than target toward the outside; etc.)

Robotic: Few fetuses studies beside Kuniyoshi's simulation.

Propose more studies with less priors in the robot?

5.2.2 Birth

Infant development: Huge change in the environment: from liquid to aerial environment (in addition to better vision of the hand and of the environment). Does the learning in utero help the neonate scaffold his motor repertoire or does he have to relearn from scratch? Would it be more difficult to learn the 'general landscape' of this gesture without having experienced this in utero?

If the effect of the environment is minimized, as in the condition of freed motor control ("motricite liberee", Grenier), with the neck supported by the adult, then the neonate is capable of pre-reaching to a bright object.

Robotic: No robot was ever programmed to go from liquid to aerial environment! One idea: have robots go from kinematic position to dynamic. There has been work that exploit kinematic-then-dynamic strategies (Rosenstein's weight lifter robot, for example), and policy iteration techniques that start from kinematic seeds (Schaal, Kuindersma).

In general, infants spend a lot of time discovering the dynamics of their body. There are numerical techniques that automatically identify the dynamics of a limb or of the limb and an object that is grasped. These techniques in robotics can acquire the forward dynamics of the limb up to a set of equivalent parameter settings and with more training get better. Moreover, having identified the dynamics of the limb, these algorithms can also identify the inertial parameters of grasped objects while undergoing generic movements. These results have not yet been generalized to "whole-body" dynamics.

5.2.3 5–7 months: emergence of grasping

Infant development: In order to reach and grasp objects, infants need to know how to compensate gravity, how to slow down their movement before touching the object, anticipate the shape, size, etc. of the object so that the hand arrives around it, ready to grasp. He does it by freezing the degrees of freedom of the arm.

There are several studies focused on corrective movements instead of integrated movements (Berthier), and reach-touch, reach-grasp skill learning (Grupen). However, these studies do not match the human infant's broad attention to motor contingencies. One of the appropriate foci for study involves frameworks for attention and exploration in robots aimed specifically in competency and situation assessment in unstructured situations

Robotic: How can robotic help solve understand the mechanisms underlying changes? For instance robotics has solved the forward-inverse problem but we still don't understand how infants solve the problem.

Several studies have explored "freezing and thawing" degrees of freedom with inconclusive results. For example, Luc Berthouze has analyzed jolly jumpers and swinging motions using robots that first engage hip motions and then knees. The simple proximal degree of freedom followed by proximal plus distal degree of freedom has not yet proven to lead to optimal 2 DOF strategies. The proximal-distal proposition has, therefore, not yet been demonstrated conclusively.

5.2.4 6–12 months: emergence of manipulation

Infant development: as soon as the infant is able to grasp objects, he manipulates them. A lot of mouthing but also many actions adapted to the object's characteristics (banging hard object on hard surface for instance, etc. Ruff, Lockman). They often transfer the object from hand to hand, use role-differentiated bimanual actions.
Robotic: Need for bimanual coordination? (is it difficult for a robot to transfer object from hand to hand? This is a very common behaviour in infants starting to explore objects)

Tasks involving two hands can easily be decomposed into left and right hand roles in robot control systems, however, to optimize the behaviour of the bimanual system, a subsequent optimization period is required. Several analytical approaches are applicable (dynamic motion primitives, policy iteration techniques), but I am not aware that this has been demonstrated in bimanual tasks in a compelling fashion.
Conclusion: should we treat robots as a different species?! Humans develop, grow physically whereas robots do not ? does it change the learning processes? Infants & robot learn in an opposite way: – infants start with a bunch of explorations, – robots start with "what is my task?"

5.2.5 Comments from Rod Grupen

In general, skills and abilities in infants and robots are still acquired in quite different way. Infants build layer upon layer of support skills by exploration that seems to be independent of any other purpose than to acquire comprehensive mastery of increasingly sophisticated relationships to the world. No task is required. The state of the art in robotics, however, typically starts with a target task and is reduced into pieces that are described algorithmically. Typically, a designer anticipates all the events and intermediate states and therefore, the robot is unsupported by the same breadth of contingencies that the infant spends all of its time constructing during the sensorimotor stage of development. This is an opportunity for both fields. If the artifacts of development in human infants can be transformed into theories of the processes of development, then these can be verified on robot platforms to create better robots as well as to codify theories of development in animals.

5.2.6 Comments from Aaron Sloman

A general point is that I don't believe we have an adequate ontology for formulating theories about the information processing going on at different stages in an infant, nor the changes that can occur (possibly along different trajectories in different individuals, including changes in architectures, forms of representation, ontologies, data-structures, information actually stored, transformations of information, control strategies – e.g. selecting capabilities that exist but may or may not be used at various stages).

Compare the differences between the Aristotelian attempt to explain behaviours of physical matter (in terms of something like the goals of different kinds of matter and what they do to achieve their goals), and Newton's explanations in terms of inertial mass, velocities, accelerations and forces.

The changes we need are probably much more complex than the transition from Aristotle to Newton.

Jackie Chappell and I wrote an invited paper for ”The International Journal of Unconventional Computing” trying to bring out the different relations between genome and development at different stages of development, suggesting that some of the genomic influences can only operate on results of previous learning that are environment-dependent. This undermines many evo-devo debates. The paper is online here:
http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609

One aspect that changes during development in ways that current theories in AI and psychology do not seem to address is the role of mechanisms for representation of possibilities (in the world, in the child) and the control of use of those mechanisms. The developing understanding of sets of possibilities and constraints on those possibilities appears to be the basis of at least some mathematical competences, including those that led our ancestors to the discovery/creation of Euclidean geometry.

I also wonder how many differences there are between infants in different cultures and different physical environments (e.g. cave dwellers, tent dwelling nomads, infants with and without planar surfaces in their environment...

5.3 Group on concepts/representations

(O'Regan, Woergoertter, Stoychev, Ugur, Stojanov)

The discussion in this group was on a much more general level, which was probably to be expected given the breadth of the topic. Just to give an idea: the discussion went from philosophical theories of concepts (feature, rule, prototype, or example based) to the conceptual framework of children raised in non-typical environments (raised by animals, in confinement, and so on). We started by criticizing current mainstream research in developmental psychology, agreeing that if Piaget were to submit a paper to Jean Piaget Society annual conference, it would probably be rejected.

During the discussion two opposing views emerged:

1. Concepts emerge gradually, starting with the motor babbling, which gives rise to repetitive sensory-motor experience, especially before the time infants are capable of walking. These similar sensory-motor trajectories are then clustered producing a primary categorization tool for objects that can be physically manipulated. The process continues with the development of language, when these primary categories start being labeled. The syntax of the language then allows for hypothetical (never-seen) word constructs, novel categories to be expressed. Perception of objects/situations, in this view, would involve bottom up activation of these clusters and spread activation to close ones.
2. The opposing view, articulated by a practicing robotics researcher, was that this cannot be the whole story, for the way it has been presented, the previous theory could not give a satisfactory account of the emergence of abstract concepts like 'containment'. Also, perception of an object as a member of some category happens virtually instantly in human perception. But, so far, we do not know of such fast algorithms that will start from the pixels and come out with the object category name instantly. A suggestion was made that we probably need some grammar-like structures, perhaps innate (the name of Jerry Fodor was mentioned), to account for the 'immediacy of perception'.

We certainly did not come to a point to suggest some experimental designs for psychologists, but did discuss possibilities of exploring perception in congenitally blind people, or people born with other types of sensory-motor deficiencies. We noted that none of us were aware of psychological research that would suggest what kinds of data structures/algorithms were better suited for modeling concepts in machines. Currently, one might say that the bottom-up ('clustering') approach is more popular at least among researchers in developmental robotics.

5.4 Group on motivation (e.g. what to explore, and what is not interesting)

Outcomes of the meeting of the group.

- Focus on intrinsic motivations (IMs)
- Focus on IMs in infants, but also a bit older children and adults
- Focus on psychological problems (we considered robots as models to study children, although being aware that this will be also useful for technology in the future)

Why is studying IMs important?

- They are fundamental for individual learning, when there are no EMs or social pressures. E.g., consider a child playing. A large amount of knowledge and skills are acquired based IMs (we played a game: formulating the subjective estimation of the percent of knowledge and skills acquired by a 1 year old child during 1 month based on IMs. The results was: Verena, Gianluca, Beata, 90
- IMs are fundamental for education, but the basic mechanisms and principles behind them are not well understood: if they will be better understood we could improve the education system
- Because they are implicitly exploited in most developmental psychology experiments (where are take for granted), so if you understand them you can control such experiments better
- ...but IMs are not studied much per se: they should, given their importance.

Why are IMs not studied much in developmental psychology?

- Because the focus of whole "scientific" psychology is on cognition rather than on motivations/emotions (treated in psychology only for therapy, etc.).
- Because developmental psychology rarely studies mechanisms; e.g., it very often studies when different cognitive capacities emerge (50-70
- So, importance of:
  - collaboration with modelers for developing experiments on mechanisms
  - finding new experimental paradigms suitable to study IMs

Methodological problems and solutions:

- Problem: How to study IMs for ongoing development in the limited time of experiments? Solution: Longitudinal studies can also help a lot to do this Solutions: developmental experiments are always exploiting IMs, so:
  - We can look at existing research to have info
  - They show we can study IMs mechanisms, e.g. creating suitable set ups (novel objects, agency).

- Other problem: adults, but also children (and even babies!) feel a lot of social pressure and implicit requests for tasks.
- Other problem: it is in general difficult to study motivation in the lab as you cannot control motivation.

IMs and other motivations:

- Very often you have multiple motivations:
  - multiple IMs
  - multiple EMs
  - multiple social motivations and they together drive behaviour
- One interesting problem is how these different sources of motivations are arbitrated

Relations between IMs and social motivations:

- Problem of the relation between IMs and social motivations, for example imitation.
- Imitation might be an innate drive, so not related to IMs.
- Or imitation might be (at least in part) the consequence of IMs as they drive the child to engage in experiences that maximise learning rate.

5.5 Visual Perception

What is missing in robotics as opposed to children?

- Segmentation: developmental studies assume objects as entities, however, this is one of the big open visual perception questions.
- Perception of gestalts: It is unclear how to group percepts and how to weigh different Gestalt principles.
- Stable object and object class recognition, similarly for navigation and localisation
- Small parts and small objects: difficult to perceive (Kinect is blind to small objects) and difficult to find reliably.
- Object permanence: knowing this is the backside of the object seen before.
- Object models: taking into account changes in objects such as cutting them or other non-trivial deformations.
- Hardware is missing: hands, tactile sensors, robustness of HW is lacking, 6-7DOF arms suitable for mobile robots, costs are far too high.

5.6 Miscellaneous points spanning above subareas

Rod Grupen made the point that some problems might be too difficult in a single modality. The group in general identified a lack of psychology research on mechanisms of development (most research is about abilities at timepoints). What new experimental paradigms could be proposed to address this deficiency?

6 Kevin's Game

This was a game proposed by Kevin O'Regan.

6.1 Part A

The Rules:

1. The psychologists pose a concrete question about a particular behaviour
2. Roboticist volunteers give 6 minute explanations suggesting why their particular theoretical constructs can explain the behaviour
3. The psychologists comment and evaluate on the responses

6.1.1 Question from Psychologist Emily Bushnell

There is a library of elementary skills. How does this battery of skills increase in number, and how do they recombine and get modified over developmental time?

6.1.2 Ben Kuipers

Using QLAP, and starting from raw sensory data, the agent builds a hierarchy of actions aimed towards making the combinations coherent for a given task. One problem is how to find meaningful qualitative states from the continuous data. Seek "contexts" in which observed contingencies between events become more reliable To attain "goals", find conditions where appropriate contexts exist that allow wanted actions to occur Use reinforcement learning to search the space of qualitative values that correspond to those contexts. Method allows simple skills to develop into higher level skills. The search space is reduced by this progressive method. "NEVER SOLVE HARD PROBLEMS". Solve easy problems first.

Comments from Psychologists: An interesting approach.

6.1.3 Gianluca Baldassarre

There are three parts of the model: (i) a skill learning part: makes a map between state of the world and motor behaviour; (ii) Goal creating part; (iii) Motivation part: (intrinsic, extrinsic, social).

Goals can be set either through motivation or through a change in the Environment.

Comments from Psychologists: The most translatable into psychology among the robotic competitors.

6.1.4 Rod Grupen

Build General Motor schemes which correspond to major "routes" to solve problems, fit them to smaller problems.

Comments from Psychologists: Good metaphors, good timing; good tripping, sometimes lost en route though.

6.1.5 Florentin Wörgötter

List the things you can do with your hand; three types: grasp, Take down, put on top. Describe the essential characteristics of the different things you can do. Learning problem: two components: 1. Acquire the "essence"; 2. Find out how to do it nicely.

Comments from Psychologists: Very Good list.

6.1.6 Alex Stoytchev

No answers, only analogies... Skills are “replicators” which evolve. Where did the skills come from? How do they evolve? Is imprinting fundamentally different from the mechanisms that make skills evolve?

Comments from Psychologists: Extra credit for being brief. Good to be clumsy at the beginning.

6.1.7 Aaron Sloman

Vive l’architecture!

Comments from Psychologists: Vive l’architecture!

Some robotics/AI researchers have noticed that human and animal competences have different “layers”, providing very different capabilities.

Reactive architectures can provide rich and versatile behavioural competences, with little understanding of how they work what their limitations are, what might have gone wrong in a successful action, ... Insects and perhaps the majority of successful behaving organisms have only that kind of intelligence.

There’s another architectural layer that seems to have evolved much later and in fewer species which involves not just the ability to perform successfully, but the ability to consider unrealised possibilities, to speculate about unknown structures and processes (e.g. in the distance, out of sight, in the past in the future, or what might have happened but did not) and to understand some of the constraints on those possibilities, which can be used in planning, predicting, explaining, designing new machines or buildings or changing actions to improve them, without depending on trial and error.

Yet another collection of competences which seems to be even rarer and probably develops more slowly in humans involves meta-semantic competences: being able to represent and reason about things that represent and reason (including having goals, plans, beliefs, preferences, puzzles, etc.)

I conjecture that in children the three types of architecture develop in parallel, with increasing roles for the second and third layers over time, and that the patterns of interaction are so diverse (across individuals, and across cultures) that human developmental trajectories are far more variable than in other species., a fact that can be missed by some researchers. For more on this see
http://www.cs.bham.ac.uk/research/projects/cogaff/#overview

To be continued.

6.2 Part B

The Rules:
Roboticists ask psychologists what precise experiments they would like psychologists to do.

6.2.1 Ben Kuipers:

Analogously to Hilbert’s list of important problems in mathematics, create a small and concrete set of puzzles that the roboticists could model, and that psychologists think would be productive. Find ways to visualize individual results with their intrinsic variation.

6.2.2 Markus Vincze & Justus Piater

Create experiments that could separate maturational learning from (adult-type) learning. Deteriming that something happens at Month X does not mean anything useful to roboticists. It is better to specify the necessary capacities that precede each next developing capacity. Individual trajectories are useful because they help specify sequences.

6.2.3 Gianluca Baldassarre

How do children set their goals when there is no social pressure?

6.2.4 Frank Guerin

More longitudinal observations of individual infants' day by day evolution (e.g. as Ester Thelen did analysing reaching). Make detailed video databases of individual evolution and make them available.

6.2.5 Pierre-Yves Oudeyer

Experiments to show whether development of language vocalization use the same mechanisms as skill learning?

6.2.6 Alex Stoytchev

It would be desirable if psychologists could publish unaveraged results, to show what each individual child does, and their individual trajectory of development. Averaging loses a lot of important data.

7 Challenge Problems

Ben Kuipers initiated a discussion about "challenge problems" along the lines of Hilbert's problems. This came to be known as "Ben's Ten" within the group.

The following is the proposal for Ben's Ten concrete psychological problems to be solved by roboticists:

1. To model the Rovee Collier experiments showing the effect of a change of context on an infant skill. Can the reinforcement learning framework explain this progressive specialisation, and effect of context, perhaps through generalization and transfer of the RL skills. (Testing basic reinforcement-type learning on infant phenomena.)
2. To model the infant's intrinsic motivation and the developmental progression in their spontaneous exploration of objects of different complexity(moderate discrepancy theory) (cf. Emily Bushnell). Perhaps modelled on Pierre-Yves Oudeyer's experiment to test intrinsic motivation
3. To model the process of symbol generation/abstraction/object concept formation, e.g. Sinan Kalkan's noun/adjective distinction; chunking.
4. To model child abilities in sequential attention/planning/sequential behaviour (perhaps modelled with hierarchical reinforcement learning policies).
6. To model convincingly the A not B task.
7. To model infant abilities in Combining elements which are related (spontaneously sorting objects).

8. To model the following phenomena described by Piaget: the child's difficulty with problems like: rose is a flower and a flower is a plant.
9. To model the problem described by Jacqueline Fagard of lifting a rod from slot, where the rod has board stuck on top; infants struggle to do this, why? It seems to be related to retrieving an object from on top of a support, where the object loses one boundary.
10. Pavolovian conditioning has 47 phenomena which have not been explained mechanistically (cf. Stoytchev) (cf behaviourism: discriminative learning; generalization.).

Aaron Sloman had given the example of Richard Young's modeling with production system of Piaget's serial sorting task, as a particularly clear example of how AI techniques could model some psychological phenomena.

Participants

- John Alexander
University of Aberdeen, GB
- Gianluca Baldassare
ISTC-CNR – Rome, IT
- Emily W. Bushnell
Tufts University, US
- Paul R. Cohen
Univ. of Arizona – Tucson, US
- Rana Esseily
Univ. Paris Ouest Nanterre, FR
- Jacqueline Fagard
Université Paris Descartes, FR
- Severin Fichtl
University of Aberdeen, GB
- Roderic A. Grupen
University of Massachusets – Amherst, US
- Beata Joanna Grzyb
Universitat Jaume I – Castellon de la Plana, ES
- Frank Guerin
University of Aberdeen, GB
- Verena V. Hafner
HU Berlin, DE
- Matej Hoffmann
Universität Zürich, CH
- Bipin Indurkhya
AGH University of Science & Technology – Krakow
- Sinan Kalkan
Middle East Technical University – Ankara, TR
- George Konidaris
MIT, US
- Norbert Krüger
University of Southern Denmark – Odense, DK
- Benjamin Kuipers
University of Michigan, US
- Ales Leonardis
University of Birmingham, GB
- Honghai Liu
University of Portsmouth, GB
- Jeffrey J. Lockman
Tulane University, US
- Bärbel Mertsching
Universität Paderborn, DE
- J. Kevin O'Regan
Université Paris Descartes, FR
- Mohamed Oubbati
Universität Ulm, DE
- Pierre-Yves Oudeyer
INRIA – Bordeaux, FR
- Justus Piater
Universität Innsbruck, AT
- Lauriane Rat-Fischer
Université Paris Descartes, FR
- Helge Ritter
Universität Bielefeld, DE
- Aaron Sloman
University of Birmingham, GB
- Georgi Stojanov
The American University of Paris, FR
- Alexander Stoytchev
Iowa State University, US
- Emre Ugur
ATR – Kyoto, JP
- Markus Vincze
TU Wien, AT
- Florentin Wörgötter
Universität Göttingen, DE