ICT Call 7
ROBOHOW.COG
FP7-ICT-288533

Deliverable D1.4:

Knowledge enabled decision making and inference in human everyday manipulation activities

SEVENTH FRAMEWORK PROGRAMME

January 31st, 2015

Project acronym:	ROBOHOW.COG
Project full title:	Web-enabled and Experience-based Cognitive Robots that Learn Complex Everyday Manipulation Tasks
Work Package:	WP 1
Document number:	D1.4
Document title:	Knowledge enabled decision making and inference in human everyday manipulation activities
Version:	1.0
Delivery date:	January 31st, 2015
Nature:	Report
Dissemination level:	Restricted (RE)
Authors:	Roy de Kleijn (UL) George Kachergis (UL) Bernhard Hommel (UL)

The research leading to these results has received funding from the European Union Seventh Framework Programme FP7/2007-2013 under grant agreement n^{O}288533 ROBOHOW.COG.

Contents

Summary

In this deliverable, we identify key issues in human everyday action and how they relate to robotics. Action control as used for such everyday manipulation tasks are remarkably different than robotic action control for industrial robotics and human action control in psychology labs.

The deliverable consists of two chapters. The first chapter is based on a journal paper, and discusses how psychological research into everyday action can inform roboticists.

The second chapter is based on a book chapter. This chapter focuses on the relationship between the field of cognitive robotics and cognitive psychology. The field of cognitive robotics successfully uses several concepts from cognitive psychology, and is an example of how psychological research can be valuable in robotics.

Chapter 1

Everyday human action, and what roboticists should know

This chapter focuses on the issues that arise when generalizing observations made in experimental psychology labs to the real world, and what this implies for robotics. Whereas simple responses to unambiguous stimuli have been the subject of study for more than a century, more complex actions have remained elusive for cognitive scientists. However, we believe that research in human action control has learned enough to inform (cognitive) robotics, and reach out with four research areas that may prove to be valuable.

First, the gap between symbolic planning and subsymbolic sensorimotor information needs to be bridged. Everyday, sequential, action in humans does not only consist of determining the optimal order of subsequences, but also shows sensorimotor context effects that optimize execution of the complete sequence.

Second, feedforward and feedback mechanisms for motor control need to be integrated to make the most efficient use of available information. In human action, unnecessary parameters are omitted from the feedforward plan to be retrieved from the environment online. This ensures speedy action onset and provides for adaptive behavior.

Third, everyday action can often be viewed as hierarchical in nature. For example, in the pancake making example, the subaction *flip the pancake* could consist of subactions *grab the spatula*, *move the spatula under the pancake*, and *flip spatula*. However, the representation of such hierarchical information in the human brain is unclear, and the subject of active investigation. We have made progress in unraveling this problem in the form of a cognitive model based on the LEABRA model.

Last, human action is highly contextualized. The same subaction can consist of different subactions based on its position in the sequence or purpose. The presence and activation of goal representation may be necessary to produce optimal action plans.

This chapter is based on the following publication that has been published in Frontiers in Neurorobotics:

> de Kleijn, R., Kachergis, G., and Hommel, B. (2014). Everyday robotic action: Lessons from human action control. *Frontiers in Neurorobotics, 8:13.*

frontiers in
NEUROROBOTICS

REVIEW ARTICLE
published: 17 March 2014
doi: 10.3389/fnbot.2014.00013

Everyday robotic action: lessons from human action control

Roy de Kleijn[1,2] *, George Kachergis[1,2] and Bernhard Hommel[1,2]

[1] Cognitive Psychology Unit, Institute for Psychological Research, Leiden University, Leiden, Netherlands
[2] Leiden Institute for Brain and Cognition, Leiden University, Leiden, Netherlands

Edited by:
Michail Maniadakis, Foundation for Research and Technology – Hellas, Greece

Reviewed by:
José F. Fontanari, University of São Paulo, Brazil
Maryam Abbapour Babaei, University Putra Malaysia, Malaysia

***Correspondence:**
Roy de Kleijn, Cognitive Psychology Unit, Institute for Psychological Research, Leiden University, Wassenaarseweg 52, 2333 AK Leiden, Netherlands
e-mail: kleijnrde@fsw.leidenuniv.nl

Robots are increasingly capable of performing everyday human activities such as cooking, cleaning, and doing the laundry. This requires the real-time planning and execution of complex, temporally extended sequential actions under high degrees of uncertainty, which provides many challenges to traditional approaches to robot action control. We argue that important lessons in this respect can be learned from research on human action control. We provide a brief overview of available psychological insights into this issue and focus on four principles that we think could be particularly beneficial for robot control: the integration of symbolic and subsymbolic planning of action sequences, the integration of feedforward and feedback control, the clustering of complex actions into subcomponents, and the contextualization of action-control structures through goal representations.

Keywords: complex action, action control, action sequencing, naturalistic action, goal-directed behavior

INTRODUCTION

In a relatively short time span, the discipline of robotics has advanced from producing industrial non-autonomous, repetitive machines to semi-autonomous agents that should be able to function in a dynamic, human-driven world. Simple examples include automatic vacuum cleaners such as Roombas, but more flexible and autonomous humanoid robots are currently under development (e.g., the RoboHow.Cog project: www.robohow.eu). As robots perform more and more everyday human activities such as household chores, interacting with humans, and thereby almost becoming citizens in our societies, we believe that psychologists can provide relevant knowledge about human behavior that is generalizable to robots.

Like early approaches to artificial intelligence (AI), traditional cognitive psychology considers behavior (of biological or artificial agents) to emerge from discrete series of cognitive operations that take information from the environment (registered by sensory organs or artificial sensors), process this information in more or less complex ways, and eventually manipulate something in the environment as a result of this processing. In psychology, this *discrete, serial processing model* of cognition has been successful in explaining various psychological phenomena, but for one reason or another most research has focused on the early and middle stages of this process, leaving action and motor control far behind. Indeed, psychology as an autonomous science has historically shown an impressive neglect of the study of action and motor control, to the extent that it has even been called the "Cinderella of psychology" (Rosenbaum, 2005).

Fortunately, however, more recent approaches have emphasized the role of action not only as an output function but as a precondition and basic ingredient of human cognition (e.g., Clark, 1997; Hommel et al., 2001; O'Regan and Noe, 2001). These recent approaches have criticized the traditional sequential-stage account of human behavior for analyzing action as a consequence of stimuli. They argue that action is more aptly characterized as people's means to produce stimuli (desired outcomes), rather than as a means to respond to stimuli (Hommel, 2009). Moreover, actions are more than mere ballistic outputs: they are events that unfold in time and that must be structured in such a way that their outcome satisfies current needs and goals. Consider, for example, the act of tea-making, which consists of a number of components: (1) boiling water, (2) putting a tea bag in a teapot, (3) pouring the boiling water in the teapot, and (4) pouring the tea in one or more cups. Executing these different components in such a way that the intended goal is eventually achieved requires planning. In the following, we will provide a brief overview of available psychological insights into how this planning works in humans, and how these insights might inform the creation of robotic everyday action systems. At the moment, although robot actions mimic human action, the control systems are in fact quite different. We will confine our discussion to four principles that we think could be particularly beneficial for robot control: the integration of symbolic and subsymbolic planning of action sequences, the integration of feedforward and feedback control, the clustering of complex actions into subcomponents, and the contextualization of action-control structures through goal representations.

INTEGRATING SYMBOLIC AND SUBSYMBOLIC PLANNING

In contrast to the ballistic, single-step actions that participants in laboratory experiments often carry out, everyday action commonly consists of multiple components, as in the tea-making example. In AI and robotics, multi-component actions are commonly planned at a symbolic level, with each action component being represented by an arbitrary symbol or function. The STRIPS (Stanford Research Institute Problem Solver) planner (Fikes and Nilsson, 1971) is a famous example: it serves to translate an initial state into an intended goal state by determining the subset of actions (defined as a symbolically described relation between sets of pre- and post-conditions) needed to do so. The format of all representations involved is symbolic allowing all goals and actions to be represented in basically the same way, although they can be arbitrarily linked to subsymbolic trigger states. This uniformity allows for a very efficient planning process, as action components can be easily manipulated and exchanged until the entire plan is optimal.

Symbolic action planning of this sort is consistent with early models of human action planning, which typically connected underspecified symbolic action representations with subsymbolic trigger states that took care of timing. For instance, Margaret Washburn considered that later action components might be triggered by the perception of the execution of the previous one: "If the necessary stimulus for pronouncing the last syllable of a series were the muscular contractions produced in pronouncing the next to the last syllable, then the proper sequence of movements would be insured" (Washburn, 1916, p. 9). Along the same lines, James (1890) suggested a *serial chaining model*, according to which each action component is triggered by the perception of the sensory feedback produced by the previous component. Accordingly, learners will create associations linking the motor patterns and their sensory consequences in a chain-like fashion.

As more studies were conducted, however, it was found that chaining accounts of sequential behavior cannot account for several empirical observations. In a seminal paper, the neurophysiologist Lashley (1951) pointed out that the serial chaining models of the time were not adequate, because: (1) movements can still be executed if sensory feedback is impaired; (2) some movements are executed too quickly to have time to process feedback from preceding actions, and (3) errors in behavior suggest the presence of predetermined action plans (Rosenbaum et al., 2007). Rosenbaum et al. (2007) added further arguments against a chaining account of sequential action. For example, the time needed to initiate an action is a function of its complexity (Henry and Rogers, 1960; Klapp, 1977; Rosenbaum, 1987), suggesting that the agent anticipates later action components before beginning to execute the first.

Along the same lines, Cohen and Rosenbaum (2004); [for another good example see Van der Wel and Rosenbaum (2007)] had participants grasp a vertical cylinder placed on a platform and move it to another platform that was either higher or lower than the initial location. The researchers determined the vertical location of the grasp, and found that the grasp location was dependent on the expected end state. More specifically, subjects tended to choose a lower grasp location when bringing the cylinder to a higher position, and vice versa. Likewise, when subjects were asked to move the cylinder back to its starting position, they tended to grasp it in the location where they grasped it before. This *end-state comfort effect* suggests that people anticipate the position that they will assume after the action has been completed.

The same conclusion is suggested by studies on context effects in speech production. For example, people round their lips before pronouncing the *t* in the word *tulip*, in anticipation of pronouncing the *u* later in the sequence (Daniloff and Moll, 1968; Bell-Berti and Harris, 1979; Fowler, 1980; Rosenbaum, 1991). This does not seem to be a purely epiphenomenal property of human action; one can easily see how this produces more efficient, smoother speech, and a more careful use of the human speech-production "hardware." An analogous action blending effect occurs when people reach for objects: people adaptively flex their fingers while moving the hand toward an object (Jeannerod et al., 1995), and has been observed to develop when sequentially moving a cursor through a learned series of stimuli (Kachergis et al., under review). Compared to typical step-wise robotic motion, this action blending seems to be more efficient, using predictive motion to minimize the time and energy required to achieve the goal.

Further insights into human sequential action planning come from Gentner et al. (1980), who conducted a photographic study of a skilled typist. Using high-speed photography, they analyzed the hand movements of a 90-wpm typist, and found that the typist's hands were moving continuously, with fingers starting to move toward a destination before several preceding characters were to be typed. In fact, for 96% of all keystrokes, movement was initiated on average 137 ms before the preceding keystroke was completed, and for 21% the movement was initiated before the preceding keystroke was initiated. Larochelle (1984) presents a similar but more extensive study, analyzing the typing of four professional typists while they typed either words or non-words, of which half were typed with one hand, and the other half with two hands. In more than half of the trials the movement was initiated before completion of the previous keystroke for two-handed trials.

These interactions between early and later sequence elements cast doubt on a simple chaining theory of sequential action. Rosenbaum et al. (2007) interpreted these findings as evidence that sensory feedback is not a necessary component for action sequencing, in keeping with the conclusion of Lashley (1951). They argued that "the state of the nervous system can predispose the actor to behave in particular ways in the future," (p. 526), or, there are *action plans* for some behaviors. And yet, studies on spontaneous speech repair (e.g., Nakatani and Hirschberg, 1994) also show that people are very fast in fixing errors in early components of a word or sentence, much too fast to assume that action outcomes are evaluated only after entire sequences are completed. This means that action planning cannot be exclusively feedforward, as Lashley (1951) seemed to suggest, but must include several layers of processing, with lower levels continuously checking whether the current action component proceeds as expected. In other words, action planning must be a temporally extended process in which abstract representations to some extent provide abstract goal descriptions, which must be integrated with lower-level subsymbolic representations controlling sensorimotor loops. The existence of subsymbolic sensorimotor representations would account for context and anticipation effects, as described above. In the more general field of knowledge representation, some authors even take it one step further, positing that subsymbolic, sensorimotor representations are *necessary* for higher-level symbolic cognition. For example, Barsalou's (1993,1999) perceptual symbol systems theory defines cognition as embedded in the world, stating that agents form grounded models via perception and interaction with their environments. With these models, the representation of abstract concepts can be implemented using grounded perceptual symbols. The empirical support for theories like these motivate the notion that both symbolic and subsymbolic representations can (and should) work together to account for human cognition.

A good example for an action planning model that includes one symbolic and one subsymbolic level is the typewriting model suggested by Rumelhart and Norman (1982). To control typing the word "WORD," say, the model would assume that the symbolic/"semantic" representation WORD would activate motor units controlling the finger movements required to type "W," "O," "R," and "D" in parallel. This parallel activation allows for crosstalk between the different units, which would account for context effects and anticipations. At the same time, the activated units are prevented from firing prematurely by means of a forward-inhibition structure. That is, each unit is inhibiting all following units in the sequence (so that the "W" unit inhibits the "O," "R," and "D" units, the "O" unit the "R" and "D" units, and the "R" the "D" unit) and release that inhibition only once they are executed. The dynamics of these inhibition and release processes automatically produce the necessary sequence. It is thought that such activation and inhibition processes play a role even in young infants (Verschoor et al., unpublished). Immediate feedback, though not explicitly addressed by Rumelhart and Norman (1982), could serve to repair the actions controlled by particular units, but the feedback would not be needed to produce the sequence – a major advantage over chaining models. For an overview of similar models and other action domains, see Logan and Crump (2011).

The main lesson for robotic everyday action control is that purely symbolic planning may be too crude and context-insensitive to allow for smooth and efficient multi-component actions. Introducing multiple levels of action planning and action control may complicate the engineering considerably, but it is also likely to make robot action more flexible and robust – and less "robotic" to the eye of the user.

INTEGRATING FEEDFORWARD AND FEEDBACK MECHANISMS

In perfectly predictable environments such as industrial construction halls, there is hardly any need for feedback mechanisms. Indeed, early industrial robots, such as Unimate, could rely on fully preprogrammed feedforward control for repetitive multi-component actions such as picking up and manipulating objects (Hägele et al., 2008). However, real-life environments are much too unpredictable to allow for purely feedforward control. Considering that purely feedback-based control is often much too slow to allow for real-life human action, it is unsurprising that human action control seeks for an optimal integration of feedforward and feedback mechanisms.

One of the earliest studies into feedforward planning is Henry and Rogers (1960), which compared reaction times of participants performing a simple finger movement to reaction times of a moderately complex arm movement (reaching and grasping) in response to a stimulus. The authors found that participants performing the more complex movement showed a 20% increase in reaction time, with as much as a 25% increase for even more complex movement. This suggests the existence of feedforward action planning prior to action execution.

Linguistic studies have shown a similar effect. Eriksen et al. (1970) had participants read aloud two-digit numbers consisting of a varying number of syllables. Longer numbers were shown to have a longer onset delay. In order to account for the possibility that factors other than motor planning play a role, participants were given the same task with a delay between stimulus onset and vocalization. Here, the effect disappeared, again providing evidence for pre-execution action plan formation.

However, while it may be tempting to conclude that an action plan is formed completely before action onset, incremental approaches to sequential action posit that this is not the case. Palmer and Pfordresher (2003) argued that it is unlikely for actors to have access to all elements in a long sequence, as this would place unnecessarily large demands on memory – just think of a conductor starting to conduct a 4-h Wagner opera. Instead, planning and execution co-occur in time, limiting access to sequence elements that appeared much earlier or that lie far in the future. Evidence for this was indeed found by Sternberg et al. (1988), in which six participants prepared and produced sequences of mono- or tri-syllabic words. In addition to the length effect discussed above, preparation times were found to increase with length of the word sequence until approaching asymptote (which was 10.3 ± 0.6 words for sequences of mono-syllabic words and 6.4 ± 0.9 words for tri-syllabic words). This suggests that plan formation and execution occur simultaneously, at least for longer sequences of actions, with a limited capacity.

However, feedforward mechanisms alone cannot account for such complex action as our tea-making example. A complete feedforward program would need to incorporate numerous unknown parameters, such as the exact location and physical properties (e.g., weight) of all necessary objects. The prior unavailability of such parameters is not the only reason feedback mechanisms might be helpful. Some parameters might be possible to include in a feedforward program, but would simply be more efficient or optimal if filled in online, such as grip strength. Even if all this information were available, an actor still needs to be able to correct possible – sometimes inevitable – perturbations in action execution.

Indeed, it seems that the presence of uncertainty (i.e., unavailability of necessary parameters) increases the importance of feedback mechanisms. Saunders and Vijayakumar (2011) fitted participants with a prosthetic hand that could provide vibrotactile feedback. Using this prosthetic hand, they were asked to manipulate objects of different weights. Manipulating both feedforward uncertainty by adding an unpredictable delay in the prosthetic hand and feedback information by manipulating vibrotactile feedback, they found that performance decreased when feedback was removed in situations with feedforward uncertainty. This illustrates that human action emerges from the interaction of feedforward and feedback mechanisms.

Integrating feedforward and feedback mechanisms holds the promise to get the best from two worlds. Feedforward mechanisms are likely to determine the necessary action components and pre-load at least some of them before initiating the action (Henry and Rogers, 1960), and to selectively tune attention to stimuli and stimulus dimensions that are relevant for the task

(Hommel, 2010). Feedback processes, in turn, provide excellent accuracy – often at the cost of speed (Seidler et al., 2004). These strengths and weaknesses have motivated hybrid models claiming that feedforward mechanisms provide the skeleton of action plans which leave open slots for parameters provided by feedback processes (Schmidt, 1975; Glover, 2004; Hommel, 2010). A particularly good example of this kind of interaction is provided by the observations of Goodale et al. (1986). In a clever experiment, participants were asked to rest their hand on a platform and point to a visual target presented at a random location on an imaginary line in their right visual field. The participants were not told that in half of the trials the target changed location during the first saccade. The authors found that participants would successfully point to the target on these trials without even being aware of the location change, and without additional delay. As feedforward programming is thought to take time, a fast and online feedback mechanism of which participants are unaware has to be responsible for this finding. After this study showing online adaptation of hand velocity, Prablanc and Martin (1992) found that these results generalize to two dimensions. Using stimuli presented on a screen, it was found that both the velocity and trajectory of the hand were adjusted online. This demonstrates that action is the result of a preprogrammed action plan (the initial movement of the hand) combined with online adaptation to reach goal requirements. Interestingly, such a division of labor fits well with the architecture of the human brain, which includes both a slow, cognitively penetrated ventral route from perception to action and a fast dorsal sensorimotor loop (for a broader overview, see Milner and Goodale, 1995).

It is clear that both feedforward and feedback mechanisms are responsible for producing complex action, but there remain a number of unanswered questions. Are feedforward processes always responsible for certain actions? How are these plans learned, and how do people know when to apply them? How does feedback on a lower level result in action re-planning on a higher level, and does this require conscious intervention? What is the division of labor between feedback and feedforward mechanisms? How fluid is it – how hierarchical?

We know that with practice, the roles of feedback and feedforward processes change. In a standard rapid aimed limb movement paradigm, participants are asked to perform a manual action in order to reach a target. During such tasks, the response can be regarded as having two elements: (1) a ballistic primary movement, thought to be controlled by a feedforward mechanism, and (2) a secondary, corrective movement, thought to be caused by a feedback mechanism. Pratt and Abrams (1996) used such a paradigm to investigate the effect of practice on the weight of primary and secondary movements. Participants were asked to repeatedly move a visual cursor to a target location using wrist rotation. With more practice, the percentage of time spent in the first movement increased, while time spent in the second movement decreased. As the first movement is feedforward-controlled, this suggests that practice reduces the need of feedback control, as the feedforward process becomes more accurate. But will this learning generalize to new situations with similar action requirements, and is it long-lasting?

To investigate the relationship between practice and feedback control, Proteau et al. (1987) had participants practice an aiming task on either 200 or 2000 trials and found that, when visual feedback was taken away, participants who had more practice were more impaired by the removal of feedback. This is not what one would expect if practice simply shifts control to feedforward processes. Subsequent research has shown that, with practice, higher peak velocities are reached in the early phase of movement, thereby leaving more time for corrective submovements based on feedback. Thus, instead of a shift from feedback control to feedforward control, feedback processes seem to be optimized as a result of practice (Proteau et al., 1987; Khan et al., 1998; Elliott et al., 2010).

While the first generation of robots and other intelligent systems had a strong preference for feedforward control, not in the least because of the rather predictable environments they were implemented in, some modern systems rely heavily on feedback control to perform actions – especially humanoid systems operating in real-world scenarios. This is likely to work as long as action production in such robots is slower than the feedback loops informing them (Plooij et al., 2013), but progress in action mechanics is likely to make hybrid feedforward/feedback systems an attractive alternative in the near future.

HIERARCHICAL ACTION REPRESENTATION

Human actions can often be described in a hierarchical fashion: "Going on vacation" implies action such as "packing my bags," "getting the car," "loading it," "driving down to city X," and so forth and so on. Many authors have taken that to imply that action control is hierarchical as well. According to Lashley (1951), only a hierarchical organization of actions and action plans can provide the opportunity to have the same motor acts acquire different meanings, depending on the context in which the motor act is performed. In Miller et al. (1960) seminal book, action plans are even hierarchical by definition: "A Plan is any hierarchical process in the organism that can control the order in which a sequence of operations is to be performed" (p. 16). And yet, while it is certainly uncontroversial that it is possible to *describe* actions as hierarchical, this need not have any implication for the cognitive organization of actions. As Badre (2008) argues, "the fact that a task can be represented hierarchically does not require that the action system itself consist of structurally distinct processing levels" (p. 193; see also Klein, 1983). Moreover, it is not always clear what authors mean if they say that actions are organized in a hierarchical fashion.

Uithol et al. (2012) noted that there are at least two ways to look at hierarchical action. These two ways differ in what are considered to be the different levels in such a hierarchy. One way to look at action hierarchies is the view of part-whole relations. In this account, each level in the hierarchy exists solely as the sum of lower-level units. In other words, an action unit such as "get a pan for pancake making" consists of the subunits "open the cupboard," "take pan from cupboard," "place pan on counter," and "close the cupboard." It should be clear that when all subordinate units are present, the superordinate unit "get a pan" is also present, as it is identical to the sum of its parts. Uithol et al. (2012) argues that this kind of hierarchy does not provide an explanation of the complex action; it merely provides a thorough description of the to-be-explained action, in which higher levels are more complex than lower levels. It also does not give information about the causal relationship between the different levels in the hierarchy, as you cannot consider an element to be the cause of its own parts. Another restriction of this type of hierarchy is that it can only accommodate levels that are of a similar nature. That is, actions can only be divided into sub-actions, not into objects or world states.

Another way to view hierarchies is to see the different levels as representing causal relations between the levels. In this approach, units on a higher level causally influence units on a lower level. In this type of hierarchy, lower-level units can be modulated by higher-level units. In contrast with the part-whole hierarchy, lower levels are not necessarily less complex than higher levels. Goals that are formulated as simple and propositional states can be the cause of more complex elements. Using this hierarchical approach also opens up the possibility of states or objects being the cause of an action, as it does not have the limitation of requiring action-type goals.

Uithol et al. (2012) proposed a new model, in which the fundamental foundation for the hierarchical structure is not cause-and-effect (i.e., goals cause motor acts), or complexity (i.e., complex motor acts such as grabbing a pan consist of simpler acts such as flexing fingers and grasping the handle), but temporal stability. In this view, stable representations can be considered goal-related, while more temporary representations reflect motor acts on different levels, not unlike the more enduring conceptual representations and the less enduring motor units of Rumelhart and Norman's (1982) model discussed above. However, this representation proposal does not include a model of how the hierarchies within a task are abstracted and learned from experience, nor of how they may be shared across tasks despite requiring different parameterizations.

Botvinick and Plaut (2004) tackled some of these issues, pointing out that not only is it unclear how existing hierarchical models learn hierarchies from experience, but also that most theoretical accounts lead to a circular reference: acquiring sequence knowledge relies on the ability to identify event boundaries, which in turn requires sequence knowledge. A further problem is sequencing in hierarchical structures; many models (e.g., Rumelhart and Norman, 1982; Houghton, 1990) solve that by means of forward inhibition, but this only works on units at the lowest level of a hierarchy. Botvinick and Plaut (2004) offered a recurrent connectionist network model that helps avoiding these problems. Using computer simulations they showed that such a network, which contains no inherent hierarchical structure, can learn a range of sequential actions that many consider hierarchical. The hierarchy, they argued, emerges from the system as a whole. The network they used is a three-layer recurrent network, with an input layer representing held objects and fixated objects, an output layer representing actions to be taken, and a hidden layer (with recurrent connections) for the internal representation. Having trained this network on a routine complex task (making coffee or tea), they showed that it can perform complex action that can be considered hierarchical in nature (e.g., varying orders of subactions leading to the same outcome) without relying on a hierarchical system architecture. The network also showed slips of action when the internal representation layer was degraded, as well as other action errors found in empirical studies, although Cooper and Shallice (2006) suggest that the relative frequency and types of errors shown by the recurrent model do not match human subjects.

We believe that architectures offering such hierarchical behavior, without necessarily being hierarchically structured, can provide robots with the needed flexibility to function in a dynamic, human-driven world. Botvinick and Plaut's (2004) model seems to be able to account for some aspects of flexible behavior, but more complex and biologically inspired models such as LEABRA (O'Reilly, 1996; Kachergis et al., under review) promise to generalize to other tasks, as well as being able to learn relatively fast, two aspects of human behavior we consider essential to emulate in robot behavior.

CONTEXTUALIZING ACTION CONTROL

As pointed out above, one of the reasons why Lashley (1951) considered action representations to be necessarily hierarchically organized was the fact that the meaning and purpose of action components vary with the goal that they serve to accomplish: while making a kicking movement with your right leg can easily be replaced by moving your head sideways when trying to score a goal in a soccer game, that would not be a particularly good idea when performing a group can-can on stage during a performance of Orpheus in the Underworld. In other words, goals are needed to contextualize action components. In AI, robotics, and some information-processing approaches in psychology, the main function of goal representation is to guide the selection of task components, including stimulus and response representations or perception-action rules. In traditional processing models, like ACT-R or Soar (Laird et al., 1987; Anderson, 1993), goal representations limit the number of production rules considered for a task, which reduces the search space and makes task preparation more efficient (Cooper and Shallice, 2006). Moreover, goals commonly serve as a reference in evaluating an action, when comparing the current state of the environment with the desired state (Miller et al., 1960).

This practice was challenged by Botvinick and Plaut (2004), who pointed out at least two problems with goal representations in cognitive models. First, goals themselves may be context-dependent. The goal of cleaning the house may have rather different implications depending on whether it serves to satisfy the expectations of one's partner or to prepare for a visit of one's mother-in-law. Likewise, the goal of stirring will produce somewhat different behavior depending on whether one is stirring egg yolks or cement. Most models that postulate the existence of goals do not allow for such context dependence. Second, it is argued that many everyday activities do not seem to have definable, or at least not invariant goals; just think of playing a musical instrument or taking a walk. The authors demonstrated that goal-directed behavior can be achieved without the explicit representation of goals. In the previously mentioned simulation studies with recurrent neural networks, they were able to simulate goal-directed actions that operate very much like Miller et al.'s (1960) TOTE units, without any need to represent the goal explicitly. Obviating the need for representing goals, such a model could be applied to behavior with non-obvious goals, such as taking a walk as a consequence of feeling restless or having the thought of fresh air (Botvinick and Plaut, 2004).

Cooper and Shallice (2006) took issue with this non-representationalist account of goals, giving at least two reasons why goals *should* be implemented in cognitive models. First, goals allow for the distinction between critical and supporting actions. When making pancakes, the subaction of adding egg to the mixture consists of picking up an egg, breaking it (above the bowl), and discarding the empty shell (not above the bowl). It should be clear that the breaking of the egg is the most important action in this sequence. Dissociating important actions from non-important actions can account for skipping unnecessary steps. When applying butter to two slices of toast, it is not necessary to execute the supporting actions "discard knife" and "pick up knife" between the two executions of the "butter toast" action program. Second, the implementation of goals would allow for subactions that serve the same purpose to be interchanged. For example, flipping a pancake by flipping it in the air or flipping it using a spatula would both be perfectly good methods for pancake flipping, and the shared goal allows these actions to be interchanged. Models without goal representation can only show this behavior if they are explicitly trained on all the alternative actions that can be taken. To make the realization that a set of actions are equivalent for achieving a goal, a model would in essence have to contain a representation of that goal.

Interestingly, however, goal representations (whether explicit or implicit) can play an important role in contextualizing cognitive representations. Most representational accounts assume that representations of stimulus and action events are invariant. The need to contextualize representations – i.e., to tailor them to the particular situation and task at hand – thus seems to put the entire burden on the goal, so that the explicit representation of the goal seems to be a necessary precondition for adaptive behavior. But, from a grounded cognition perspective, it seems that alternative scenarios are possible. In a grounded cognition framework, the representation of objects and object categories takes an embodied form, using modal features from at least the visual, motor, and auditory modalities (Prinz and Barsalou, 2000). For example, the concept of apple would be represented by a network of visual codes representing <green> and <round>, but also the auditory <crunchy sound> of biting into it. The embodied cognition framework has already been successfully implemented in robot platforms such as iCub, and shows stimulus compatibility effects similar to those that can be observed in humans (Macura et al., 2009; Pezzulo et al., 2011).

According to the Theory of Event Coding (Hommel et al., 2001), events are represented – like objects – in a feature-based, distributed fashion. This will mean that the aforementioned apple would be represented by a network of codes representing not only the apple's perceptual features such as being <greenish> and <round>, but also its properties such as being <edible>, <graspable>, <carryable>, <throwable>, and so forth. In this view, one of the main roles of goals is to emphasize (i.e., increase the weight of) those features that in the present task are of particular importance. This means that when hungry, the feature of being <edible> will be primed in advance and become more activated when facing an apple, while <throwability> will become more important when being in danger and trying to defend oneself. Several studies have provided evidence that goals are indeed biasing attentional settings toward action-relevant feature dimensions (e.g., Fagioli et al., 2007; Wykowska et al., 2009; Kühn et al., 2011), suggesting that the impact of goals goes beyond the selection of production rules and outcome evaluation. Interestingly, this kind of "intentional weighting" function (Memelink and Hommel, 2013) can be considered to represent the current goal without requiring any explicit representation – very much along the lines of Botvinick and Plaut (2004).

Another potential role of goals is related to temporal order. In chaining models, the dimension of time was unnecessary because the completion of each component automatically "ignites" the next component. The same holds for current planners in cognitive robotics, which commonly fix the order of action subcomponents (e.g., CRAM: Beetz et al., 2010). But action plans may follow a more abstract syntax instead, much like how syntactic constraints of natural languages allow for various possible sequences. For instance, consider the process of making tea. With the possible exception of true connoisseurs, it doesn't make any difference for most tea drinkers whether one puts the tea or the water into the cup first; i.e., the order of these two subactions is interchangeable. A truly flexible system would thus allow for any of these orders, depending on whether water or tea is immediately at hand. While a chaining model would not allow for changing the original order, a more syntactic action plan would merely define possible slots for particular subcomponents (e.g., Rosenbaum et al., 1986), so that the actual order of execution would be an emerging property of the interaction of the syntactic plan and the situational availability of the necessary ingredients.

These considerations suggest that robotic systems need to incorporate at least some rudimentary aspects of time and temporal order to get on par with humans. Along these lines, Maniadakis and Trahanias (2011) have propagated the idea that robotic systems should be equipped with some kind of temporal cognition, be it by incorporating temporal logic or event calculus. Indeed, recent robotic knowledge representation systems, such as KnowRob (Tenorth and Beetz, 2012), do possess the ability to do spatiotemporal reasoning about the changing locations of objects, such as predicting when and where objects can be found.

CONCLUSION

We have discussed how conceptions of robotic action planning can benefit from insights into human action planning. Indeed, we believe that constructing truly flexible and autonomous robots requires inspiration from human cognition. We focused on four basic principles that characterize human action planning, and we have argued that taking these principles on board will help to make artificial cognition more human-like.

First, we have discussed evidence that human action planning emerges from the integration of a rather abstract, perhaps symbolic representational level and concurrent planning at a lower, more concrete representational level. It is certainly true that multilevel planning can create difficult coordination problems. Using grounded cognition approaches in robotics is potentially a good method to ground such higher-level symbolic representation in lower-level sensorimotor representations, which may allow robot action to become more flexible and efficient.

Second, we have argued that human action planning emerges from the interplay of feedforward and feedback mechanisms. Again, purely feedforward or purely feedback architectures are likely to be more transparent and easier to control. However, fast, real-time robotic action in uncertain environments will require a hybrid approach that distributes labor much like the human brain does by combining slow and highly optimized feedforward control with fast sensorimotor loops that continuously update the available environmental information. A major challenge for the near future will be to combine such hybrid systems with error-monitoring and error-correcting mechanisms. When preparing pancake dough, accidentally pouring some milk outside the bowl would need to trigger a fast correction mechanism informed by low-level sensory feedback but not necessarily the re-planning of (or crying over) the entire action. However, if for some reason the entire milk carton is emptied by this accident, leaving the agent without the necessary ingredient, feedback would have to propagate to higher, more abstract or more comprehensive planning levels to decide whether the plan needs to be aborted. How this works in detail and how decisions are made as to which level is to be informed is not well understood, but progress is being made. Research into feedback processes has yielded information about the optimal speed of sensorimotor loops (Joshi and Maass, 2005), and we find it reasonable to expect that models using such fast feedback loops combined with accurate feedforward planning can ultimately produce human-like motor performance in robots.

Third, we have argued that while descriptions of human actions may refer to a hierarchy, it is not yet clear whether the cognitive – *in vivo* or *in silico* – representations of such actions need to be explicitly hierarchical as well. Equally unclear is whether representations that differ in hierarchical level would necessarily need to differ in format. However, it is clear that representations that are considered to be "higher in hierarchy" are more comprehensive. The concept of "making a pancake," say, is necessarily richer and more abstract than the associated lower-level actions of "reaching for egg" and "grabbing a pan," suggesting that the latter two are more directly grounded in sensorimotor activity (Kraft et al., 2008). Future research will need to investigate how representations at different planning levels (or different levels of description) interact or relate to each other.

The nature of goals and their role in action control is also a matter of ongoing research. The two different viewpoints – i.e., that goals require explicit representation or not – seem to reflect different preferences in conceptualization and modeling techniques, and it may well turn out that an explicit representation of goals in the preferred modeling language translates to a more implicit representation of goals in the actual functional or neural architecture. In robotics, most modern plan languages use a form of explicit goal-related action control that defines a goal as a required world state on which constraints can be imposed. Such a structure is flexible enough to allow equifinality, but it is unclear how knowledge about the various means to produce a result is acquired. Ultimately, we believe that subsymbolic programming approaches may allow for more adaptive, "human" representational architectures – though likely more difficult to engineer and define provably safe operating conditions for.

To conclude, we believe that the construction of robots that are up to real-life, everyday actions in environments that are as uncertain as human environments requires the consideration of cognitive principles like the four principles we have discussed in this article. The benefit of doing so will be twofold. For one, it will strongly increase the flexibility of robots. For another, it will make robots more human-like in the eyes of the human user, which will help us understand and cooperate with our future robotic colleagues.

ACKNOWLEDGMENT

The preparation of this work was supported by the European Commission (EU Cognitive Systems project ROBOHOW.COG; FP7-ICT-2011).

REFERENCES

Anderson, J. R. (1993). *Rules of the Mind*. Hillsdale, NJ: Erlbaum.

Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends Cogn. Sci.* 12, 193–200. doi: 10.1016/j.tics.2008.02.004

Barsalou, L. W. (1993). "Flexibility, structure, and linguistic vagary in concepts: manifestations of a compositional system of perceptual symbols," in *Theories of Memory*, eds A. C. Collins, S. E. Gathercole, and M. A. Conway (London: Lawrence Erlbaum Associates), 29–101.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behav. Brain Sci.* 22, 577–660.

Beetz, M., Mösenlechner, L., and Tenorth, M. (2010). "CRAM: a cognitive robot abstract machine for everyday manipulation in human environments," in *IEEE/RSJ International Conference on Intelligent RObots and Systems*, Taipei.

Bell-Berti, F., and Harris, K. S. (1979). Anticipatory coarticulation: some implications from a study of lip rounding. *J. Acoust. Soc. Am.* 65, 1268–1270. doi: 10.1121/1.382794

Botvinick, M., and Plaut, D. C. (2004). Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychol. Rev.* 111, 395–429. doi: 10.1037/0033-295X.111.2.395

Clark, A. (1997). *Being there: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.

Cohen, R. G., and Rosenbaum, D. A. (2004). Where grasps are made reveals how grasps are planned: generation and recall of motor plans. *Exp. Brain Res.* 157, 486–495. doi: 10.1007/s00221-004-1862-9

Cooper, R. P., and Shallice, T. (2006). Hierarchical schemas and goals in the control of sequential behavior. *Psychol. Rev.* 113, 887–916. doi: 10.1037/0033-295X.113.4.887

Daniloff, R., and Moll, K. (1968). Coarticulation of lip rounding. *J. Speech Hear. Res.* 11, 707–721. doi: 10.1044/jshr.1104.707

Elliott, D., Hansen, S., Grierson, L. E., Lyons, J., Bennett, S. J., and Hayes, S. J. (2010). Goal-directed aiming: two components but multiple processes. *Psychol. Bull.* 136, 1023–1044. doi: 10.1037/a0020958

Eriksen, C. W., Pollack, M. D., and Montague, W. E. (1970). Implicit speech: mechanism in perceptual encoding? *J. Exp. Psychol.* 84, 502–507. doi: 10.1037/h0029274

Fagioli, S., Hommel, B., and Schubotz, R. I. (2007). Intentional control of attention: action planning primes action-related stimulus dimensions. *Psychol. Res.* 71, 22–29. doi: 10.1007/s00426-005-0033-3

Fikes, R., and Nilsson, N. (1971). STRIPS: a new approach to the application of theorem proving to problem solving. *Artific. Intell.* 2, 189–208. doi: 10.1016/0004-3702(71)90010-5

Fowler, C. A. (1980). Coarticulation and theories of extrinsic timing control. *J. Phonet.* 8, 113–133.

Gentner, D. R., Grudin, J., and Conway, E. (1980). *Skilled Finger Movements in Typing (Technical Report 8001)*. San Diego, CA: Center for Human Information Processing.

Glover, S. (2004). Separate visual representations in the planning and control of action. *Behav. Brain Sci.* 27, 3–24.

Goodale, M. A., Pélisson, D., and Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature* 320, 748–750. doi: 10.1038/320748a0

Hägele, M., Nilsson, K., and Norberto Pires, J. (2008). "Industrial robotics," in *Springer Handbook of Robotics*, eds B. Siciliano and O. Khatib (Berlin: Springer), 963–986.

Henry, F. M., and Rogers, D. E. (1960). Increased response latency for complicated movements and a "memory drum" theory of neuromotor reaction. *Res. Q.* 31, 448–458.

Hommel, B. (2009). Action control according to TEC (theory of event coding). *Psychol. Res.* 73, 512–526. doi: 10.1007/s00426-009-0234-2

Hommel, B. (2010). "Grounding attention in action control: the intentional control of selection," in *Effortless Attention: A New Perspective in the Cognitive Science of Attention and Action*, ed. B. J. Bruya (Cambridge, MA: MIT Press), 121–140. doi: 10.7551/mitpress/9780262013840.003.0006

Hommel, B., Müsseler, J., Aschersleben, G., and Prinz, W. (2001). The theory of event coding (TEC): a framework for perception and action planning. *Behav. Brain Sci.* 24, 849–878. doi: 10.1017/S0140525X01000103

Houghton, G. (1990). "The problem of serial order: a neural network model of sequence learning and recall," in *Current Research in Natural Language Generation*, eds R. Dale, C. Mellish, and M. Zock (San Diego, CA: Academic Press), 287–319.

James, W. (1890). *Principles of Psychology*, Vol. 1. New York: Holt. doi: 10.1037/10538-000

Jeannerod, M., Arbib, M. A., Rizzolatti, G., and Sakata, H. (1995). Grasping objects: the cortical mechanisms of visuomotor transformation. *Trends Neurosci.* 18, 314–320. doi: 10.1016/0166-2236(95)93921-J

Joshi, P., and Maass, W. (2005). Movement generation with circuits of spiking neurons. *Neural Comput.* 17, 1715–1738. doi: 10.1162/0899766054026684

Khan, M. A., Franks, I. M., and Goodman, D. (1998). The effect of practice on the control of rapid aiming movements: evidence for an interdependency between programming and feedback processing. *Q. J. Exp. Psychol. Hum. Exp. Psychol.* 51A, 425–444. doi: 10.1080/713755756

Klapp, S. T. (1977). Reaction time analysis of programmed control. *Exerc. Sport Sci. Rev.* 5, 231–253. doi: 10.1249/00003677-197700050-00008

Klein, R. (1983). Nonhierarchical control of rapid movement sequences. *J. Exp. Psychol. Hum. Percept. Perform.* 9, 834–836. doi: 10.1037/0096-1523.9.5.834

Kraft, D., Baseski, E., Popovic, M., Batog, A. M., Kjær-Nielsen, A., Krüger, N., et al. (2008). Exploration and planning in a three level cognitive architecture. *Proceedings of the International Conference on Cognitive Systems (CogSys 2008)*, Karlsruhe.

Kühn, S., Keizer, A., Rombouts, S. A. R. B., and Hommel, B. (2011). The functional and neural mechanism of action preparation: roles of EBA and FFA in voluntary action control. *J. Cogn. Neurosci.* 23, 214–220. doi: 10.1162/jocn.2010.21418

Laird, J. E., Newell, A., and Rosenbloom, P. S. (1987). SOAR: an architecture for general intelligence. *Artific. Intell.* 33, 1–64. doi: 10.1016/0004-3702(87)90050-6

Larochelle, S. (1984). "Some aspects of movements in skilled typewriting," in *Attention and Performance. Control of Language Processes*, Vol. 10, eds H. Bouma and D. G. Bouwhuis (Hillsdale, NJ.: Erlbaum), 43–54.

Lashley, K. S. (1951). "The problem of serial order in behavior," in *Cerebral Mechanisms in Behavior*, ed. L. A. Jeffress (New York: Wiley), 112–131.

Logan, G. D., and Crump, M. J. C. (2011). Response to M. Ullsperger and Danielmeier's E-Letter. *Sci. E Lett.* (February 9, 2011).

Macura, Z., Cangelosi, A., Ellis, R., Bugmann, D., Fischer, M. H., and Myachykov, A. (2009). A cognitive robotic model of grasping. *Proceedings of the Ninth International Conference on Epigenetic Robotics*, Venice.

Maniadakis, M., and Trahanias, P. (2011). Temporal cognition: a key ingredient of intelligent systems. *Front. Neurorobot.* 5:2. doi: 10.3389/fnbot.2011.00002

Memelink, J., and Hommel, B. (2013). Intentional weighting: a basic principle in cognitive control. *Psychol. Res.* 77, 249–259. doi: 10.1007/s00426-012-0435-y

Miller, G. A., Galanter, E., and Pribram, K. H. (1960). *Plans and the Structure of Behavior*. New York: Holt, Rinehart & Winston. doi: 10.1037/10039-000

Milner, A. D., and Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.

Nakatani, C. H., and Hirschberg, J. (1994). A corpus-based study of repair cues in spontaneous speech. *J. Acoust. Soc. Am.* 95, 1603–1616. doi: 10.1121/1.408547

O'Regan, J. K., and Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behav. Brain Sci.* 24, 939–1031. doi: 10.1017/S0140525X01000115

O'Reilly, R. C. (1996). *The LEABRA Model of Neural Interactions and Learning in the Neocortex*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA.

Palmer, C., and Pfordresher, P. Q. (2003). Incremental planning in sequence production. *Psychol. Rev.* 110, 683–712. doi: 10.1037/0033-295X.110.4.683

Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. H., McRae, K., and Spivey, M. J. (2011). The mechanics of embodiment: A dialog on embodiment and computational modeling. *Front. Psychol.* 2:5. doi: 10.3389/fpsyg.2011.00005

Plooij, M., de Vries, M., Wolfslag, W., and Wisse, M. (2013). *Optimization of Feedforward Controllers to Minimize Sensitivity to Model Inaccuracies*. Paper submitted to and accepted for IROS 2013, Tokyo.

Prablanc, C., and Martin, O. (1992). Automatic control during hand reaching at undetected two-dimensional target displacements. *J. Neurophysiol.* 67, 455–469.

Pratt, J., and Abrams, R. A. (1996). Practice and component submovements: the roles of programming and feedback in rapid aimed limb movements. *J. Mot. Behav.* 28, 149–156. doi: 10.1080/00222895.1996.9941741

Prinz, J. J., and Barsalou, L. W. (2000). "Steering a course for embodied representation," in *Cognitive Dynamics: Conceptual Change in Humans and Machines*, eds E. Dietrich and A. Markman (Cambridge, MA: MIT Press), 51–77.

Proteau, L., Marteniuk, R. G., Girouard, Y., and Dugas, C. (1987). On the type of information used to control and learn an aiming movement after moderate and extensive practice. *Hum. Mov. Sci.* 6, 181–199. doi: 10.1016/0167-9457(87)90011-X

Rosenbaum, D. A. (1987). Successive approximations to a model of human motor programming. *Psychol. Learn. Motivat.* 21, 153–182. doi: 10.1016/S0079-7421(08)60028-6

Rosenbaum, D. A. (1991). *Human Motor Control*. New York, NY: Academic.

Rosenbaum, D. A. (2005). The Cinderella of psychology: the neglect of motor control in the science of mental life and behavior. *Am. Psychol.* 60, 308–317. doi: 10.1037/0003-066X.60.4.308

Rosenbaum, D. A., Cohen, R. G., Jax, S. A., Weiss, D. J., and van der Wel, R. (2007). The problem of serial order in behavior: Lashley's legacy. *Hum. Mov. Sci.* 26, 525–554. doi: 10.1016/j.humov.2007.04.001

Rosenbaum, D. A., Weber, R. J., Hazelett, W. M., and Hindorff, V. (1986). The parameter remapping effect in human performance: evidence from tongue twisters and finger fumblers. *J. Mem. Lang.* 25, 710–725. doi: 10.1016/0749-596X(86)90045-8

Rumelhart, D. E., and Norman, D. A. (1982). Simulating a skilled typist: a study of skilled cognitive-motor performance. *Cogn. Sci.* 6, 1–36. doi: 10.1207/s15516709cog0601_1

Saunders, I., and Vijayakumar, S. (2011). The role of feed-forward and feedback processes for closed-loop prosthesis control. *J. NeuroEng. Rehabilit.* 8, 60. doi: 10.1186/1743-0003-8-60

Schmidt, R. A. (1975). A schema theory of discrete motor skill learning. *Psychol. Rev.* 82, 225–260. doi: 10.1037/h0076770

Seidler, R. D., Noll, D. C., and Thiers, G. (2004). Feedforward and feedback processes in motor control. *Neuroimage* 22, 1775–1783. doi: 10.1016/j.neuroimage.2004.05.003

Sternberg, S., Knoll, R. L., Monsell, S., and Wright, C. E. (1988). Motor programs and hierarchical organization in the control of rapid speech. *Phonetica* 45, 175–197. doi: 10.1159/000261825

Tenorth, M., and Beetz, M. (2012). Knowledge processing for autonomous robot control. *AAAI Spring Symposium on Designing Intelligent Robots: Reintegrating AI*, Stanford.

Uithol, S., van Rooij, I., Bekkering, H., and Haselager, W. F. G. (2012). Hierarchies in action and motor control. *J. Cogn. Neurosci.* 24, 1077–1086. doi: 10.1162/jocn_a_00204

Van der Wel, R. P. R. D., and Rosenbaum, D. A. (2007). Coordination of locomotion and prehension. *Exp. Brain Res.* 176, 281–287. doi: 10.1007/s00221-006-0618-0

Washburn, M. F. (1916). *Movement and Mental Imagery*. Boston: Houghton Mifflin.

Wykowska, A., Schubö, A., and Hommel, B. (2009). How you move is what you see: action planning biases selection in visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 1755–1769. doi: 10.1037/a0016798

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 01 December 2013; accepted: 25 February 2014; published online: 17 March 2014.
Citation: de Kleijn R, Kachergis G and Hommel B (2014) Everyday robotic action: lessons from human action control. Front. Neurorobot. ***8****:13. doi: 10.3389/fnbot.2014.00013*
This article was submitted to the journal Frontiers in Neurorobotics.

Chapter 2

Crossroads of cognitive robotics and psychology

In the previous chapter we discussed the merits of psychological insights for robotics, and argued that the field of robotics can use the human brain as inspiration for designing robotic systems. One of the areas that has taken this advice to heart is *cognitive robotics*.

In this chapter, we focus on the intersection of cognitive psychology and cognitive robotics. Many skills needed for everyday action, such as affordance inference and action-effect learning have successfully–although limited–been demonstrated in cognitive robots.

This chapter is based on the following book chapter, part of the book Cognitive Robotics by Hooman Samani:

> de Kleijn, R., Kachergis, G., and Hommel, B. (2015). Robotic action control: On the crossroads of cognitive psychology and robotics. In H. Samani (Ed.), *Cognitive robotics.* Taylor & Francis.

CHAPTER 9

Robotic Action Control: On the Crossroads of Cognitive Psychology and Cognitive Robotics

Roy de Kleijn

Leiden Institute for Brain and Cognition,
Leiden University, The Netherlands.

George Kachergis

Leiden Institute for Brain and Cognition,
Leiden University, The Netherlands.

Bernhard Hommel

Leiden Institute for Brain and Cognition,
Leiden University, The Netherlands.

CONTENTS

The field of robotics is shifting from building industrial robots that can perform repetitive tasks accurately and predictably in constrained settings, to more autonomous robots that should be able to perform a wider range of tasks, including everyday household activities. To build systems that can handle the uncertainty of the real world, it is important for roboticists to look at how humans are able to perform in such a wide range of situations and contexts–a domain that is traditionally the purview of cognitive psychology. Cognitive scientists have been rather successful in bringing computational systems closer to human performance. Examples include image and speech recognition and general knowledge representation using parallel distributed processing (e.g., modern deep learning models).

Similarly, cognitive psychologists can use robotics to complement their research. Robotic implementations of cognitive systems can act as a "computational proving ground", allowing accurate and repeatable real-world testing of model predictions. All too often, theoretical predictions–and even carefully-conducted model simulations–do not scale up or even correspond well to the complexity of the real world. Psychology should always seek to push theory out of the nest of the laboratory and see if it can take flight. Finally, cognitive psychologists have an opportunity to conduct experiments that will both inform roboticists as they seek to make more capable cognitive robots, and increase our knowledge of how humans perform adaptively in a complex, dynamic world. In this chapter, we will give a broad but brief overview of the fields of cognitive psychology and robotics, with an eye to how they have come together to inform us about how (artificial and natural) actions are controlled.

9.1 EARLY HISTORY OF THE FIELDS

9.1.1 History of cognitive psychology

Before cognitive psychology and robotics blended into the approach now known as cognitive robotics, both fields already had a rich history. Cognitive psychology as we now know it has had a rocky past (as have most psychological disciplines, for that matter). Breaking away from philosophy, after briefly attempting to use introspection to observe the workings of the mind, the field of psychology found it more reliable to rely on empirical evidence.

Although making rapid strides using this empirical evidence, for example in the form of Donders' now classic reaction time experiments which proposed stages of processing extending from perception to action, early cognitive psychology came to be dominated by a particular approach, *behaviorism*. This position, popularized by Watson [54] and pushed further by Skinner [46], held that the path for psychology to establish itself as a natural science on par with physics and chemistry would be to restrict itself to observable entities such as stimuli and responses. In this sense, behaviorists such as Skinner were strongly anti-representational, i.e., against the assumption of internal knowledge and states in the explanation of behavioral observations. On the other hand, the focus on observable data brought further rigor into the field, and many interesting effects were described and explained.

The behaviorist approach dominated the field of psychology during the first half of the 20th century. In the 1950s, seeming limitations of behaviorism fueled what some scholars would call the *neocognitive revolution*. Starting with Chomsky's scathing 1951 review of Skinner's book that tried to explain how infants learn language by simple association, many researchers were convinced that behaviorism could not explain fundamental cognitive processes such as learning (especially language) and memory. The foundations of the field of artificial intelligence were also nascent, and pursuing explanations of high-level, uniquely human aptitudes–e.g., analytical thought, reasoning, logic, strategic decision-making–grew in popularity.

9.1.2 The computer analogy

Another factor contributing to the neocognitive revolution was the emergence of a new way to describe human cognition as similar to electronic computer systems. The basic mechanism operating computers was (and still is, in a fundamental way) gathering input, processing it, and outputting the processed information, not unlike the basic cognitive model of stimulus detection, storage and transformation of stimuli, and response production.

Clearly, this processing of information requires some representational states which are unaccounted for (and dismissed as unnecessary) by behaviorists. This new way to look at human cognition as an information processing system not only excited psychologists as a way of understanding the brain, but the analogy also raised hopes for building intelligent machines. The idea was that if computer systems could use the same rules and mechanisms as the human brain, they could also *act* like humans. Perhaps the most well-known proponent of this optimistic vision was Turing [51], who suggested that it wouldn't be long before machine communication would be indistinguishable from human communication. Maybe the secret of cognition lies in the way the brain transforms and stores data, it was thought.

Alas, the optimists would be disappointed. It soon became clear that computers and humans have very different strengths and weaknesses. Computers can calculate pi to twenty significant digits within mere milliseconds. Humans can read terrible handwriting. Clearly, humans are not so comparable to basic input-output systems after all. It would take another 25 years for cognitive psychology and artificial intelligence to begin their romance once again, in the form of the *parallel distributed processing* (PDP) approach [40].

9.1.3 Early cognitive robots

With this idea of smart computer systems in mind, it seemed almost straightforward to add embodiment to build intelligent agents. The first cognitive robots were quite simple machines. The *Machina Speculatrix* [53] consisted of a mobile platform, two sensors, actuators and 'nerve cells'. Understandably, these robots were designed to mimic behavior of simple animals, and could move safely around a room and recharge themselves using relatively simple approach and avoidance rules.

Due to their simplicity, it was questionable exactly how *cognitive* these robots were–they are more related to cybernetics and control theory (e.g., [5])–but soon enough complexity made its way into cognitive robotics.

From the 1960s, robots would be able to represent knowledge and plan sequences of operations using algorithms such as *STRIPS* [17], that would now be considered essential knowledge for every AI student. The STRIPS planner, which represents goal states and preconditions and attempts to derive the action sequences that would achieve them before carrying them out, is quite slow to execute. Moreover, this type of planning suffers from its closed world assumption (i.e., that the environment and all relevant states are known–by programming–and will not change), and the massive complexity of the real world, leading to intractable computations. Yet the general approach taken by STRIPS–of modeling the environment, possible actions and state transformations, and goal states via predicate logic, and operating robots via a sense-plan-act loop–has dominated cognitive robotics for quite some time, and is still a strong thread today.

Various behavior-based robotics architectures and algorithms–taking some inspiration from biological organisms–have been developed in the past few decades. An early, influential example is Rodney Brooks' subsumption architecture [9], which eschews planning entirely–"planning is just a way of avoiding to figure out what to do next", using a defined library of basic behaviors arranged hierarchically to generate behavior based on incoming stimuli. Although fast and often generating surprisingly complex behavior from simple rules (see also [8]), the subsumption architecture and many other behavior-based robotics algorithms do not yet incorporate much from the lessons to be learned from psychological studies in humans.

9.2 ACTION CONTROL

9.2.1 Introduction

One of the other areas that shows considerable overlap between robots and humans is motor/action control. Two types of control systems govern motor action: *feedforward* and *feedback* control systems.

A feedforward motor control system sends a signal from the (human or robotic) motor planning component to the relevant motor component using predetermined parameters, executing said action. Information from the environment can be considered only before execution begins, which makes feedforward control suitable for predictable environments.

In contrast, a feedback motor control system incorporates information from itself or the environment (feedback) more or less continuously to modulate the control signal. In this way, the system can dynamically alter its behavior in response to a changing environment.

9.2.2 Feedforward and feedback control in humans

For many years, psychology and related disciplines have approached action control from rather isolated perspectives. As the probably first systematic study on movement control by Woodworth [55] had provided strong evidence for the contribution of environmental information, many authors have tried to develop closed-loop models of action control that rely on a continuous feedback loop (e.g., [1]). At the same time, there was strong evidence from animal and lesion studies [31, 49] and from theoretical considerations [34] that various movements can be considered in the absence of sensorimotor feedback loops, which has motivated the development of feedforward models (e.g., [22]).

Schmidt [43] was one of the first who argued that human action control consists of both feedforward and feedback components. According to his reasoning, human agents prepare a movement schema that specifies the relevant attributes of the intended movement but leave open parameter slots that are specified by using online environmental information. Neuroscientific evidence has provided strong support for such a hybrid control model, suggesting that off-line action planning along a ventral cortical route is integrated with online sensorimotor specification along a dorsal route [19, 18].

In particular, feedforward mechanisms seem to determine the necessary action components off-line and pre-load at least some of them before initiating the action [22], and to selectively tune attention to stimuli and stimulus dimensions that are relevant for the task [24]. Feedback processes, in turn, provide excellent accuracy–often at the cost of speed [44]. These strengths and weaknesses have motivated hybrid models claiming that feedforward mechanisms provide the skeleton of action plans which leave open slots for parameters provided by feedback processes [43, 18, 24].

A particularly good example of this kind of interaction is provided by the observations of Goodale and colleagues [20]. In a clever experiment, participants were asked to rest their hand on a platform and point to a visual target presented at a random location on an imaginary line in their right visual field. The participants were not told that in half of the trials the target changed location during the first saccade. The authors found that participants would successfully point to the target on these trials without even being aware of the location change, and without additional delay. As feedforward programming is thought to take time, a fast and online feedback mechanism of which participants are unaware has to be responsible for this finding.

On a higher level, interaction between feedforward and feedback systems must exist for goal-directed action to be carried out. Higher level, goal-directed action planning, such as planning to make pancakes would be impossible to plan in a completely feedforward fashion: it would require all motor parameters to be specified a priori, and thus would require exact knowledge of the position and properties of all necessary equipment and ingredients, such as weight, friction coefficients, etc.

Instead, many of these parameters can be filled in online by using information from the environment. It is not necessary to know the exact weight of a pan, because you can determine that easily by picking it up: you increase the exerted force until the pan leaves the surface of the kitchen counter. Although, you likely also learn a distribution of probable pan weights (e.g., more than 50 g and less than 10 kg) from your experience of other pans–or even just similarly-sized objects.

Interaction between feedforward and feedback becomes even more apparent on a higher level when a planned action fails to be executed. When a necessary ingredient is missing, replanning (or cancellation) of a preprogrammed action sequence may be necessary: if there is no butter, can I use oil to grease up the pan? Somehow, this information gathered by feedback processes must be communicated to the higher level action planner.

9.2.3 Feedforward and feedback control in robots

The theorizing on action control in robotic systems must be considered rather ideological, sometimes driven by the specifics of particular robots and/or tasks considered and sometimes by broadly generalized anti-representationalist attitudes. Many early robots only had a handful of sensors and responded in a fixed pattern of behavior given a particular set of stimuli. Some robots were even purely feedforward, performing the same action or action sequence, with no sensory input whatsoever [37]. Feedforward or simple reactive control architectures make for very brittle behavior: even complex, carefully-crafted sequences of actions and reactions will appear clumsy if the environment suddenly presents an even slightly novel situation.

More complex architectures have been proposed, often with some analogy to biology or behavior, giving birth to the field of *behavior-based robotics.* The *subsumption architecture* [9] was a response to the traditional GOFAI, and posited that complex behavior need not necessarily require a complex control system. Different behaviors are represented as layers that can be inhibited by other layers. For example, a simple robot could be provided with the behaviors *wandering*, *avoiding*, *pickup*, and *homing*. These behaviors are hierarchically structured, with each behavior inhibiting its preceding behavior [4].

This hierarchy of inhibition between behavior is (although somewhat more complex) also visible in humans. For example, if your pants are (accidentally) set on fire while doing the dishes, few people would finish the dishes before stopping, dropping, and rolling. In other words, some behaviors take precedence over others. An approach similar to the subsumption architecture has been proposed by [3]. The *motor schema* approach also uses different, parallel layers of behavior, but does not have the hierarchical coordination as the subsumption approach does. Instead, each behavior contributes to the robot's overall response.

On a higher level, as noted in the previous section, other problems arise. When a planned action fails to succeed, for example because a robot can't find a pan to make pancakes in, replanning is necessary. The earliest AI planners such as GPS would simply backtrack to the previous choice point and try an alternative subaction. However, this does not guarantee the eventual successful completion of the action. Other planners, such as ABSTRIPS [41], use a hierarchy of representational levels. When it fails to complete a subaction, it could return to a more abstract level.

However, truly intelligent systems should be more flexible in handling such unforeseen events. If a robot cannot make me a pizza with ham, maybe it should make me one with bacon? Generalizing and substituting appropriate remain an elusive ability for robots, although vector space models of semantics (e.g., BEAGLE; [28]) offer a step in the right direction. Like neural networks, these models represent items (e.g., words) in a distributed fashion, using many-featured vectors with initially low similarity between random items. As the model learns–say, by reading documents, item representations are updated to make them more similar (on a continuous scale) to contextually similar items. These continually-updated representations can be used to extract semantic as well as syntagmatic (e.g., part-of-speech) relationships between items. Beyond text learning, vector space models may ultimately be used to learn generalizable representations for physical properties and manipulations of objects and environments.

9.2.4 Robotic action planning

It is understood that reaching movements in humans have an initial ballistic, feedforward component, followed by a slower, feedback-driven component that corrects for error in the initial movement. As people become more adept at reaching to targets at particular distances, a greater portion of their movement is devoted to the initial feedforward component, and less time is spent in the feedback component, thus speeding response times. Understanding how this happens should enable roboticists to make more adaptive, human-like motor planning systems for robots.

In this line of research, Kachergis et al. [29] studied sequence learning using mouse movements. Inspired by earlier work of Nissen and Bullemer [38], subsequences of longer sequences were acquired by human participants during a learning phase. The participants seem to implicitly extract the subsequences from longer sequences by showing faster response times and context effects.

These findings cast doubt on a simple chaining theory of sequential action. Rosenbaum et al. [39] interpreted these findings as evidence that sensory feedback is not a necessary component for action sequencing, in keeping with the conclusion of Lashley [34]. They argued that "the state of the nervous system can predispose the actor to behave in particular ways in the future," (p. 526), or, there are action plans for some behaviors. And yet, studies on spontaneous speech repair (e.g., [36]) also show that people are very fast in fixing errors in early components of a word or sentence, much too fast to assume that action outcomes are evaluated only after entire sequences are completed. This means that action planning cannot be exclusively feedforward, as Lashley [34] seemed to suggest, but must include several layers of processing, with lower levels continuously checking whether the current action component proceeds as expected. In other words, action planning must be a temporally extended process in which abstract representations to some extent provide abstract goal descriptions, which must be integrated with lower-level subsymbolic representations controlling sensorimotor loops. The existence of subsymbolic sensorimotor representations would account for context and anticipation effects, as described above.

The main lesson for robotic motor planning is that purely symbolic planning may be too crude and context-insensitive to allow for smooth and efficient multi-component actions. Introducing multiple levels of action planning and action control may complicate the engineering considerably, but it is also likely to make robot action more flexible and robust–and less "robotic" to the eye of the user.

9.3 ACQUISITION OF ACTION CONTROL

9.3.1 Introduction

In order for humans or robots to be able to achieve their goals, it is necessary for them to know what effect an action would have on their environment. Or, reasoning back, what actions are required to produce a certain effect in the environment. Learning relevant action-effect bindings as an infant is a fundamental part of development and likely bootstraps later acquisition of general knowledge.

In humans, learned action-effects seem to be stored bidirectionally. Following Lotze [35] and Harless [21], James [27] noted that intentionally creating a desired effect requires knowledge about, and thus the previous acquisition of action-effect contingencies. The *Theory of Event Coding* (TEC; [25]) is a comprehensive empirically well-supported (for recent reviews, see [23, 45]) theoretical framework explaining the acquisition and use of such action-effect bindings for goal-directed action. TEC states that actions and their expected effects share a common neural representation. Therefore, performing an action activates the expectation of relevant effects and thinking of (i.e., intending or anticipating) an action's effects activates motor neurons responsible for achieving those effects.

9.3.2 Human action-effect learning

9.3.2.1 Traditional action-effect learning research

In traditional cognitive psychology experiments, action-effect bindings are acquired by having humans repetitively perform an action (such as pressing a specific button on a keyboard), after which an effect (such as a sound or a visual stimulus) is presented. After a certain amount of exposure to this combination of action and effect, evidence suggests that a bidirectional binding has been formed. When primed with a previously learned effect, people respond faster with the associated action [15]. This action-effect learning is quite robust but sensitive to action-effect contingency and contiguity [16].

9.3.2.2 Motor babbling

Of course, action-effect learning does not only happen in artificial environments such as psychology labs. In fact, action-effect learning in humans starts almost instantly after birth [52] and some would argue even before. Young infants perform uncoordinated movements known as *body* or *motor babbling*. Most of these movements will turn out to be useless, however, some of them will have an effect that provides the infant with positive feedback. For example, a baby could accidentally push down with its right arm while lying on its belly, resulting in rolling on its back and seeing all sorts of interesting things. Over time, the infant will build up action-effect associations for actions it deems useful, and can perform motor acts by imagining their intended effects.

Having mastered the intricacies of controlling the own body, higher level action-effects can be learned in a manner similar to motor babbling. Eenshuistra et al. [14] give the example of driving a spacecraft that you are trying to slow down. If nobody ever instructed you on how to do that, your best option would probably be pressing random buttons until the desired effect is reached (be careful with that self-destruct button!). Once you have learned this action-effect binding, performance in a similar situation in the future will be much better.

9.3.3 Robotic action-effect learning

The possibility that cognition can be grounded in sensorimotor experience and represented by automatically created action-effect bindings has attracted some interest of cognitive roboticists already. For instance, Kraft et al. [32] have suggested a three-level cognitive architecture that relies on object-action complexes, that is, sensorimotor units on which higher-level cognition is based. Indeed, action-effect learning might provide the cognitive machinery to generate action-guiding predictions and the off-line, feedforward component of action control. This component might specify the invariant aspects of an action, that is, those characteristics that need to be given for an action to reach its goal, to create its intended effect while an online component might provide fresh environmental information to specify the less goal-relevant parameters, such as the speed of a reaching movement when taking a sip of water from a bottle [24]. Arguably, such a system would have the benefit of allowing for more interesting cognitive achievements than the purely online, feedback-driven systems that are motivated by the situated-cognition approach [10]. At the same time, it would be more flexible than systems that rely entirely on the use of internal forward models [13]. Thus, instead of programmers trying to imagine all possible scenarios and enumerate reasonable responses, it might be easier to create robots that can learn action-effect associations appropriate to their environment and combine them with online information.

In robots as well as in humans, knowledge about one's own body is required to acquire knowledge about the external world. Learning how to control your limbs–first separately and then jointly (e.g., walking)–clearly takes more than even the first few years of life: after learning to roll over, crawl, and then walk, we are still clumsy at running and sport for several years (if, indeed, we ever become very proficient). Motor babbling helps develop tactile and proprioception–as well as visual and even auditory cues–of what our body in motion feels like. Knowing these basic actions and their effects on ourselves (e.g., what hurts) lays the foundation for learning how our actions can affect our environments.

In perhaps the first ever study of motor babbling in a (virtual) robot, Kuperstein [33] showed how random movement execution can form associations between a perceived object-in-hand position and the corresponding arm posture. This association is bidirectional, and as such is in line with ideomotor (or TEC) theory. We (and others, e.g., [11]) believe that such bidirectional bindings can help robots overcome traditional problems, such as inverse model inference from a forward model.

More recent investigations in robotic motor babbling have extended and optimized the method to include behavior that we would consider *curiosity* in humans. For example, Saegusa et al. [42] robotically implemented a sensorimotor learning algorithm that organized learning in two phases: *exploration* and *learning*. In the exploration stage, random movements are produced, while in the learning stage the action-effect bindings (or, more specifically, mapping functions) are optimized. The robot can then decide to learn bindings that have not yet been learned well.

9.4 DIRECTIONS FOR THE FUTURE

9.4.1 Introduction

Many questions remain with respect to the acquisition and skillful performance of not only well-specified, simple actions (e.g., reaching to a target) but of complex actions consisting of various components and involving various effectors. Indeed, how can we create a learning algorithm that can go from basic motor babbling to both successful goal-directed reaching, grasping, and manipulations of objects? To accomplish this obviously difficult goal, it will likely be beneficial for psychologists to study infants' development of these abilities and beneficial for cognitive roboticists to learn more from human capabilities.

9.4.2 Affordance learning

Object manipulation and use is an indispensable activity for robots working in human environments. Perceiving object affordances–i.e., what a tool can do for you or how you can use an object–seems to be a quick, effortless judgment for humans, in many cases. For example, when walking around and seeing a door, you automatically pull the handle to open it.

One of the ways robots can perform object affordance learning is by motor babbling using simple objects as manipulators (e.g., [47]). In a so-called *behavioral babbling stage* a robot applies randomly chosen behaviors to a tool and observes their effects on an object in the environment. Over time, knowledge about the functionality of a tool is acquired, and can be used to manipulate a novel object with the tool.

As impressive as this may sound, this approach does not allow for easy generalization, and the robot cannot use this knowledge to manipulate objects using another, similar, tool. More recent approaches, such as demonstrated by Jain and Inamura [26] infer functional features from objects to generalize affordances to unknown objects. These functional features are supposed to be object invariant within a tool category.

In humans, an approach that seems successful in explaining affordance inference is based on Biederman's *recognition-by-components theory* [6]. This theory allows for object recognition by segmenting an encountered object in elementary geometric parts called *geons*. These are simple geometric shapes such as cones, cylinders and blocks. By reducing objects to a combination of more elementary units invariance is increased, simplifying object classification. Biederman recognized 36 independent geons, having a (restricted) generative power of 154 million three-geon objects.

In addition to being useful for object classification, geons can also be used to infer affordances. For example, a spoon is suitable for scooping because its truncated hollow sphere at the end of its long cylinder allows for containing things, and an elongated cylinder attached to an object can be used to pick it up.

One very promising example of the use of geons in affordance inference is demonstrated by Tenorth and Beetz [50]. This technique matches perceived objects to three-dimensional CAD models from a public database such as Google Warehouse. These models are then segmented into geons, which makes affordance inference possible.

However, the affordances that geons give us need to be learned in some way. Teaching robots how to infer what a tool can be capable of remains difficult. Ultimately, we want affordances to develop naturally during learning: be it from watching others, from verbal instruction, or from embodied experimentation. Task context is also an important aspect of affordance learning: depending on the situation, a hammer can be used as a lever, a paperweight, a missile, or well, a hammer. To understand how context affects action planning, studying naturalistic scenes and human activities jointly seems essential (cf. [2]).

Learning geon affordances that can be generalized to object affordances seems a fruitful approach to automating affordance learning in robots, although it is early to say whether this or other recent approaches will fare better. For example, deep neural networks use their multiple hidden layers along with techniques to avoid overfitting to learn high-level perceptual features for discriminating objects. The representations learned by such networks are somewhat more biologically-plausible than geon decompositions, and thus may be more suitable for generalization (although cf. [48] for generalization problems with deep neural networks).

9.4.3 Everyday action planning

A major obstacle in the way of robots performing everyday actions is the translation of high-level, symbolic task descriptions into sensorimotor action plans. In order to make such translations, one method would be to learn the other way around: by observing sensorimotor actions, segment and classify the input.

Everyday action is characterized by sequential, hierarchical action subsequences. Coffee- and tea-making tasks, for example, have shared subsequences such as adding milk or sugar. Moreover, the goal of adding sugar might be accomplished in one of several ways: e.g., tearing open and adding from a packet, or spooning from a bowl or box. Also, these subsequences do not necessarily have to performed in the same order every time (with some constraints, of course). It is this flexibility and ability to improvise that makes everyday action so natural for humans, yet so hard for robots.

Cognitive models that represent hierarchical information have been proposed (e.g. [12], [7]), but differ in the way they represent these hierarchies. One approach explicitly represents action hierarchies by hard-coding them into the model–hardly something we can do for a general autonomous robot–whereas the latter models hierarchy as an emergent property of the recurrent neural network. More recently, the model put forth by Kachergis et al. [30], uses a neural network with biologically plausible learning rules to extract hierarchies from observed sequences, needing far fewer exemplars than previous models.

9.5 CONCLUSION

In this chapter, we have discussed several concepts that are shared between cognitive robotics and cognitive psychology in order to argue that the creation of flexible, truly autonomous robots depends on the implementation of algorithms that are designed to mimic human learning and planning. Thus, there are many relevant lessons from cognitive psychology for aspiring cognitive roboticists.

Ideomotor theory and its implementations such as TEC provide elegant solutions to action-effect learning. Robotic motor learning algorithms that use motor babbling to bootstrap higher-order learning seem to be promising, and require little *a priori* knowledge given by the programmer, ultimately leading to more flexible robots.

Generalization of action plans is still a very difficult problem. Inferring hierarchical structure of observed or learned action sequences seems to be a promising approach, although the structure of everyday action seems to be nearly as nuanced and intricate to untangle as the structure of human natural language–and less well-studied, at this point. Again, we believe that biologically inspired learning models such as LeabraTI can play a role in making robotic action more human-like.

The overlapping interests of cognitive robotics and cognitive psychology has proven fruitful so far. Mechanisms like motor babbling and affordance inference, which are extensively studied in humans, can provide robots with techniques to make their behavior more flexible and human-like. We believe human inspiration for robots can be found at an even lower level by incorporating biologically-inspired neural models for learning in robots.

9.6 Bibliography

[1] J. A. Adams. A closed-loop of motor learning. *Journal of Motor Behavior*, 3:111–150, 1971.

[2] E. E. Aksoy, B. Dellen, M. Tamosiunaite, and F. Worgotter. Execution of a dual-object (pushing) action with semantic event chains. In *11th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2011.

[3] R. C. Arkin. Motor schema-based mobile robot navigation. *International Journal of Robotics Research*, 8:92–112, 1989.

[4] R. C. Arkin. Behavior-based robotics. MIT Press: Cambridge, MA, 1998.

[5] W. R. Ashby. *An introduction to cybernetics*. London: Chapman & Hall, 1956.

[6] I. Biederman. Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94:115–147, 1987.

[7] M. Botvinick and D. C. Plaut. Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychological Review*, 111(2):395–429, Apr 2004.

[8] V. Braitenberg. *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press, 1984.

[9] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14–23, 1986.

[10] R. Brooks. Intelligence without reason. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, volume 1, pages 569–595, 1991.

[11] D. Caligiore, D. Parisi, N. Accornero, M. Capozza, and G. Baldassarre. Using motor babbling and Hebb rules for modeling the development of reaching with obstacles and grasping. In *Proceedings of the 2008 International Conference on Cognitive Systems*, 2008.

[12] R. P. Cooper and T. Shallice. Hierarchical schemas and goals in the control of sequential behavior. *Psychological Review*, 113(4):887–916; discussion 917–31, Oct 2006.

[13] Y. Demiris and A. Dearden. From motor babbling to hierarchical learning by imitation: a robot developmental pathway. In *International Workshop on Epigenetic Robotics*, pages 31–37, 2005.

[14] R.M. Eenshuistra, M.A. Weidema, and B. Hommel. Development of the acquisition and control of action-effect associations. *Acta Psychologica*, 115:185–209, 2004.

[15] B. Elsner and B. Hommel. Effect anticipation and action control. *Journal of Experimental Psychology: Human Perception and Performance*, 27:229–240, 2001.

[16] B. Elsner and B. Hommel. Contiguity and contingency in the acquisition of action effects. *Psychological Research*, 68:138–154, 2004.

[17] R. Fikes and N. Nilsson. Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.

[18] S. Glover. Separate visual representations in the planning and control of action. *Behavioral and Brain Sciences*, 27:3–24, 2004.

[19] M. A. Goodale and A. D. Milner. Separate visual pathways for perception and action. *Trends in Neurosciences*, 15:20–25, 1992.

[20] M. A. Goodale, D. Pelisson, and C. Prablanc. Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, 320:748–750, 04 1986.

[21] E. Harless. Der Apparat des Willens. *Zeitschrift für Philosophie und philosophische Kritik*, 38:50–73, 1861.

[22] F. M. Henry and D. E. Rogers. Increased response latency for complicated movements and a "memory drum" theory of neuromotor reaction. *Research Quarterly*, 31:448–458, 1960.

[23] B. Hommel. Action control according to TEC (theory of event coding). *Psychological Research*, 73:512–526, 2009.

[24] B. Hommel. *Effortless attention: A new perspective in the cognitive science of attention and action*, chapter Grounding attention in action control: The intentional control of selection, pages 121–140. Cambridge, MA: MIT Press, 2010.

[25] B. Hommel, J. Müsseler, G. Aschersleben, and W. Prinz. The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24:849–937, 2001.

[26] R. Jain and T. Inamura. Bayesian learning of tool affordances based on generalization of functional feature to estimate effects of unseen tools. *Artificial Life and Robotics*, 18:95–103, 2013.

[27] W. James. *Principles of psychology*, volume 1. New York: Holt, 1890.

[28] M. N. Jones and D. J. K. Mewhort. Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114:1–37, 2007.

[29] G. Kachergis, F. Berends, R. de Kleijn, and B. Hommel. Trajectory effects in a novel serial reaction time task. In *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 2014.

[30] G. Kachergis, D. Wyatte, R. C. O'Reilly, R. de Kleijn, and B. Hommel. A continuous time neural model for sequential action. *Philosophical Transactions of the Royal Society B*, 369:20130623, 2014.

[31] H. D. Knapp, E. Taub, and A. J. Berman. Movements in monkeys with deafferented forelimbs. *Experimental Neurology*, 7:305–315, 1963.

[32] D. Kraft, E. Baseski, M. Popovic, A.M. Batog, A. Kjær-Nielsen, N. Krüger, R. Petrick, C. Geib, N. Pugeault, M. Steedman, T. Asfour, R. Dillmann, S. Kalkan, F. Wörgötter, B. Hommel, R. Detry, and J. Piater. Exploration and planning in a three level cognitive architecture. In *Proceedings of the International Conference on Cognitive Systems (CogSys 2008),*, Karlsruhe, 2008.

[33] M. Kuperstein. Neural model of adaptive hand-eye coordination for single postures. *Science*, 239:1308–1311, 1988.

[34] K. S. Lashley. *Cerebral mechanisms in behavior*, chapter The problem of serial order in behavior, pages 112–131. New York: Wiley, 1951.

[35] R. H. Lotze. *Medicinische Psychologie oder Physiologie der Seele*. Weidmann, 1852.

[36] C. Nakatani and J. Hirschberg. A corpus-based study of repair cues in spontaneous speech. *Journal of the Acoustical Society of America*, 95:1603–1616, 1994.

[37] S. B. Niku. *Introduction to robotics: analysis, control, applications*. Wiley, 2010.

[38] M. J. Nissen and P. Bullemer. Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19:1–32, 1987.

[39] D. A. Rosenbaum, R. G. Cohen, S. A. Jax, D. J. Weiss, and R. van der Wel. The problem of serial order in behavior: Lashley's legacy. *Human Movement Science*, 26:525–554, 2007.

[40] D. E. Rumelhart and J. L. McClelland. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume I.* Cambridge, MA: MIT Press, 1986.

[41] E. D. Sacerdoti. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence*, 5:115–135, 1974.

[42] R. Saegusa, G. Metta, G. Sandini, and S. Sakka. Active motor babbling for sensory-motor learning. In *IEEE International Conference on Robotics and Biomimetics*, pages 794–799, 2008.

[43] R. A. Schmidt. A schema theory of discrete motor skill learning. *Psychological Review*, 82:225–260, 1975.

[44] R. D. Seidler, D. C. Noll, and G. Thiers. Feedforward and feedback processes in motor control. *Neuroimage*, 22:1775–1783, 2004.

[45] Y. K. Shin, R. W. Proctor, and E. J. Capaldi. A review of contemporary ideomotor theory. *Psychological Bulletin*, 136(943–974), 2010.

[46] B. F. Skinner. *The behavior of organisms: An experimental analysis*. Cambridge, Massachusetts: B.F. Skinner Foundation, 1938.

[47] A. Stoytchev. Autonomous learning of tool affordances by a robot. In *AAAI 2005*, 2005.

[48] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *CoRR*, abs/1312.6199, 2013.

[49] E. Taub, I. A. Goldberg, and P. Taub. Deafferentation in monkeys: Pointing at a target without visual feedback. *Experimental Neurology*, 46:178–186, 1975.

[50] M. Tenorth and M. Beetz. KnowRob: A knowledge processing infrastructure for cognition-enabled robots. part 1: The KnowRob system. *International Journal of Robotics Research*, 2013.

[51] A. M. Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.

[52] S. A. Verschoor, M. Spapé, S. Biro, and B. Hommel. From outcome prediction to action selection: Developmental change in the role of action-effect bindings. *Developmental Science*, 16:801–814, 2013.

[53] W. G. Walter. *The living brain*. London: Duckworth, 1953.

[54] J. B. Watson. Psychology as the behaviourist views it. *Psychological Review*, 20:158, 1913.

[55] R. S. Woodworth. The accuracy of voluntary movement. *Psychological Review*, 3:1–119, 1899.