A Survey of the Ontogeny of Tool Use: From Sensorimotor Experience to Planning
— Source link

Frank Guerin, Norbert Krüger, Dirk Kraft

Institutions: University of Aberdeen, University of Southern Denmark

Topics: Developmental robotics

Related papers:

- Intrinsic Motivation Systems for Autonomous Mental Development

- The Ecological Approach to Visual Perception

- The Origins of Intelligence in Children

- Behavior-Grounded Representation of Tool Affordances

- Learning Object Affordances: From Sensory--Motor Coordination to Imitation

A Survey of the Ontogeny of Tool Use: From Sensorimotor Experience to Planning

Guerin, Frank; Krüger, Norbert; Kraft, Dirk

Go to publication entry in University of Southern Denmark's Research Portal

A Survey of the Ontogeny of Tool Use: From Sensorimotor Experience to Planning

Frank Guerin, Norbert Krüger, and Dirk Kraft

Abstract—In this paper, we review current knowledge on tool use development in infants in order to provide relevant information to cognitive developmental roboticists seeking to design artificial systems that develop tool use abilities. This information covers: 1) sketching developmental pathways leading to tool use competences; 2) the characterization of learning and test situations; 3) the crystallization of seven mechanisms underlying the developmental process; and 4) the formulation of a number of challenges and recommendations for designing artificial systems that exhibit tool use abilities in complex contexts.

Index Terms—Developmental psychology, developmental robotics, infant behavior, tool use.

I. INTRODUCTION

THIS survey paper is targeted at researchers in artificial intelligence[1] (AI) interested in pursuing a developmental approach[2] to achieving robust object manipulation competence and basic tool use in their systems. The paper presents relevant research from studies in developmental psychology (mostly of human infants). In addition to reporting individual results, the paper identifies core mechanisms we believe to be in operation during the development of tool use in infants. Based on these mechanisms we then present general recommendations, which may be useful for those who wish to build artificial systems that exhibit a similar development.

From a roboticist's perspective, a central issue is how to reach a point in development where planning complex actions with objects is possible. Planning ability relies on a knowledge of planning operators, which describe preconditions and post-conditions of actions. How the knowledge of such operators de-

velops is a key issue for this paper. In the developmental psychology literature, there are some constructs reasonably close to planning operators, for example the *sensorimotor schema* [2] and [6].[3] A sensorimotor schema gathers together the perceptions and associated actions involved in the performance of a habitual behavior. The schema represents knowledge generalized from all the experiences of that behavior. It also includes knowledge about the context in which the behavior was performed as well as expectations about the effects. We will later refer to the parts of a schema as *context*, *action*, and *effect*. This paper will trace the development of sensorimotor schemas from their origins in the first months through to the point where they represent sufficiently abstract knowledge that they can be recruited for planning operations.

We take the *perception–action perspective* on tool use development [5], which sees a continuous trajectory of development from early exploratory interactions with objects and surfaces through to more advanced manipulations, including tool use. From this perspective, developments in perceptual and motor skills are potentially very relevant to understanding the abilities that precede tool use and might serve as foundations for it. Therefore, this survey will also devote some attention to these "precursors" to tool use. This is in line with the trend in modern AI to shift the focus away from "high-level" cognitive skills and more towards "lower-level" control of actions in the world [7]. This is not a denial of the existence of high-level skills, but rather a realization of the need to build them from low-level sensorimotor skills.

Tool use is an interesting phenomenon because it is a very obvious demonstration of intelligence, and it is relatively easy to analyze due to its external manifestation. In addition, its ontogeny is particularly interesting because it shows a development from simple sensorimotor behaviors to behaviors exhibiting the hallmarks of advanced cognition. In its simplest forms, tool use requires no more than simple context-specific sensorimotor knowledge, such as an expectation about the effect of an action. In its advanced forms, tool use requires knowledge of objects, distances, forces, and their interactions, and the ability to manipulate some form of internal representations of these. Furthermore, it is surmised that this knowledge of the physical world may be crucial as a foundation for more advanced cognition [8]. Therefore, a study of the development of this foundational knowledge, through tool use, would be an important step towards understanding advanced cognition.

[1]This paper will use the term artificial intelligence (AI) in the broad sense of the discipline concerned with any type of intelligent information processing in artificial systems, including as subdisciplines: cognitive robotics, computer vision, and all areas within computational intelligence.

[2]We will not attempt to argue for the developmental approach here, as that has been done elsewhere (e.g., [1]).

[3]The term "sensorimotor schema" comes from Piaget [2], but similar ideas like "sensorimotor process" [3], "skill" [4], or "perception-action routine" [5] are used by other psychologists. We use the term in a broad sense to capture the general idea shared by these works.
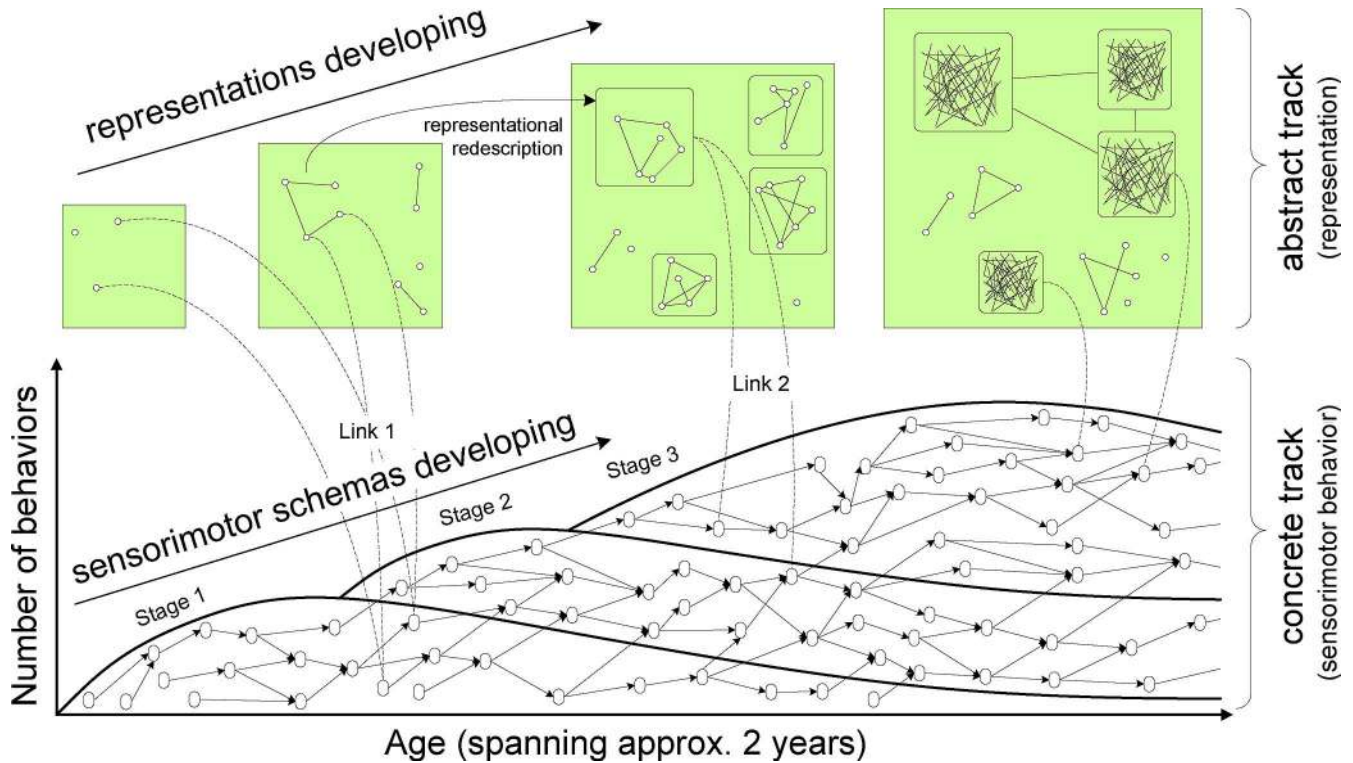
Fig. 1. Conceptual diagram providing an overview of infant developments leading to tool use.

Developmental psychology does not yet have a complete theory of cognitive development.[4] Therefore, roboticists do not have an abstract mathematical "theory of development" they could apply to any task (contrast this with, e.g., Shannon's mathematical theory of communication [10]). In the absence of such a theory, roboticists who wish to build systems that develop in a similar way to human infants may need to follow the same concrete tasks that human infants do. If a roboticist builds a robot to follow an infant's developmental sequence of increasingly sophisticated tasks, then this increases the chance that the roboticist may discover similar mechanisms of development.

The contributions of this paper are summarized as follows.[5]

1) It sketches the **development paths** along which several examples of simple tool use may be acquired. This means we describe a sequence of increasingly sophisticated behaviors that lead to some examples of tool use. This is valuable as it often reveals pathways that simplify the acquisition of a behavior and intermediate competences one might not have considered.[6]

2) It gives **concrete examples** of simple tool use and their precursors. These simple tasks are good candidate tasks

for experimentation with artificial systems because they help to avoid the danger of attempting an overly advanced task (which might force a solution to be coded in a nondevelopmental way).[7]

3) It gives some insight into how **knowledge** (of actions and object relationships) may be **represented**; this comes from our analysis of the sensorimotor basis for such knowledge, which begins with subjective sensorimotor experiences, and gradually becomes generalized to capture more objective knowledge about the effects of certain operations when applied in certain situations.

4) It gives some insight into **general mechanisms of sensorimotor development**, and how they apply in the development of tool use (see the list of mechanisms towards the end of this section).

We capture our view of the ontogeny of tool use in humans with the conceptual diagram of Fig. 1. The diagram shows two parallel tracks of development. At the bottom is the "concrete" track, which shows the development of sensorimotor schemas observable in infant behavior. At the top is the "abstract" track, which shows the parallel development of the underlying representations used by the infant. In this paper, we will mainly describe the "concrete track." For the abstract track, we stress that we know very little; psychology is mostly the study of behavior, and conjectures of cognitive models are quite limited at the present time.[8]

The lower (concrete) track shows a directed acyclic graph; each node represents a newly created sensorimotor schema, (which corresponds to a new observable behavior arising at

[4]There exist interesting considerations of design principles connected to developmental processes (see, e.g., [9]) that can guide design processes. However, a complete theory should account for developments in all domains; it should explain how the same mechanism can develop different types of knowledge depending on the environment it interacts with, and it should be detailed enough to allow for the computer modeling of a complete longitudinal developmental sequence.

[5]The first two points are purely from behavior, but the last two are delving into the internal mechanism and, hence, are more speculative.

[6]This is similar to the way in which the fossil record can reveal pathways along which complex organs developed, whereas in the absence of the fossil evidence, the evolutionary development seems difficult to explain.

[7]*Nondevelopmental* here means making use of hand-coded knowledge from a human, as opposed to having the system learn for itself.

[8]For example, limited to isolated episodes or aspects development [11].

this time) .[9] The directed edges of the graph have the meaning "is a necessary precursor"; i.e., the later behavior builds on the previous one(s). Acquisitions are (mostly) accumulative, e.g., babies suck things less as they get older, but in general they do not forget. Most behaviors we deal with can still be elicited in older infants. For the concrete track, we can categorize the behaviors as belonging to three consecutive and overlapping stages (indicated by the overlapping curves), described as follows.[10]

A. Behaviors Without Objects

This stage starts with the development of a small number of innate behavior patterns, which are in general not linked to any object. The vision and motor system become calibrated, leading to the ability to grasp seen objects, which facilitates the transition to the next stage. In parallel with this, on the (upper) abstract track, we have initially isolated fragmentary representations, which function in limited environmental contexts; these develop and become gradually connected (e.g., across different senses), allowing for transfer of knowledge (see, e.g., Section V-B5).

B. Behaviors With Single Objects

In this stage, accidental events start a linking process between action and object perception. This leads to specification and branching of sensorimotor schemas concerned with single objects; their effects become increasingly predictable. In parallel, extensive training data is generated on concrete object–action experiences. This allows the abstract track to find connections between similar experiences, and so to generalize across them; this constitutes the beginnings of the construction of more general object and affordance representations, which become increasingly task independent (by virtue of the fact that they generalize across multiple tasks). We begin to have a generic and powerful world model, which constitutes itself as an *independent* entity shared by the different sensorimotor schemes. Unfortunately, this internal world model is only indirectly deducible from the observation of behavior and constitutes a big challenge to roboticists (see [13]).

C. Object–Object Behaviors

Further branching and refinement of the schemas continues in this stage, but the new element is that the sensorimotor schemas are extended to deal with relationships among objects, which deliver the basic units for tool use. Because schemas now necessarily deal with relationships among objects, the representation of spatial locations and transforms within space begins to be constructed (abstract track). Object representations become elaborated to integrate parts of objects and different perspectives, as well as physical properties influencing their interaction. In the abstract track we also have some connected fragments of representations, which may be reformulated to form a new more general representation subsuming the old versions

[9]This graph has a strong similarity with Fischer and Hencke [12, Fig. 2].

[10]Our choice to group things in *three* stages is somewhat arbitrary, as it suits the observations we want to describe; Piaget uses six stages [2], and Fischer uses four [4], for the same period.

(the process of *representational redescription* [14]). In addition, at this stage, simple examples of planning can already be observed. The schemas are now usable in a wider variety of contexts, and their effects increasingly predictable; therefore, it is possible to plan a sequence of actions while still maintaining a high degree of predictability of the effects of such action sequences. These developments (on both tracks) are ongoing and do not stop where our figure stops.

Perhaps the greatest mystery in cognitive development is understanding how practice with concrete tasks in specific situations leads to the development of more abstract general knowledge of the physical world, such as improving representations of objects and space. We can only give some small insights into this in this paper, and to this end Fig. 1 also illustrates (with dashed curves) links between the abstract and concrete tracks; these links are bidirectional. To avoid clutter, only a few links are shown, but in reality all representational fragments will be linked to sensorimotor schemas.

In one direction (see Fig. 1, Link 1), representations may be built up (or existing fragmentary representations linked up) from the action of sensorimotor schemas; when a schema acts in a variety of contexts, it discovers sensory abstractions, which predict its success, and these abstractions are preliminary representations (e.g., representations of shape for grasping, see Section III-A). Such representations can immediately link to actions, which can manipulate the represented object or spatial relation. In the other direction (see Fig. 1, Link 2), more advanced schemas make use of the newly formed representations, for example in their description of the context in which a behavior may be performed, or its effects, or the control policy followed during execution of the schema.

This development process (see Fig. 1) allows us to deduce guidelines for how to set up an artificial developing system. In particular, the developments on the concrete track are reasonably well studied and observable, allowing us to deduce some of the mechanisms underlying them, as follows.

M1 (Repetition): Each schema seeks to repeat itself opportunistically (this explains *play*, see Section II-B), leading to its own refinement, and also to the accidental discovery of new effects in new situations (such accidental discoveries can subsequently be intentionally exploited).

M2 (Variation and Selection): Actions are performed with high variability in order to discover new results, and understand the effects of parameter variations, and later those results that give desired outcomes are selected.

M3 (Differentiation): Schemas are differentiated when an unexpected result is sufficiently interesting to warrant its own specialized schema. Motor differentiation changes the schema's action (motor control program). Sensor differentiation changes the schema's context and effect.

M4 (Decomposition): A single schema may be broken into a number of sequential chunks so that refinements of the individual parts can take place, as well as flexible reassembly (see, e.g., Sections V-B1 and V-D); this can increase the predictability and maturity of the schema.

M5 (Composition): Schemas can be composed to form simple composite sequences or higher order schemas, which control relationships among lower schemas.

M6 (Modularization): Composite schemas may be initially crudely connected sequences, but can then be refined by repetition, variation, and selection, to produce a "smooth atom" [15], which could then be put under control of another (further composition).

We will give concrete descriptions of how these mechanisms are exemplified through different behaviors in the course of the paper. These mechanisms are crucial for the development of planning competences. M1 and M2 are important for increasing the predictive power of the schemas, and by means of M3–M6, new schemas (planning operators) can be generated.

In this paper, we are not addressing the mechanisms underlying the development of the abstract track, but there is a need for a mechanism that synchronizes the development of both tracks:

M7 (Representational Redescription): When similarities are noted among a set of sensorimotor schemas, a new more abstract representation can be created, which can reformulate the knowledge captured in the former schemas within a more generic framework.

The largest part of this paper is Section V, which goes through examples of tool use and precursors. However, before this, it is necessary to first cover some more general preliminaries. Section II sets forth our perspective on the problem of tool use and how competence develops. Section III gives an overview of sensorimotor schemas (which is the unit of knowledge we will use in analysis in this paper). Section IV provides an overview of various perspectives on cognitive development in order to explain the different psychological approaches and to put the later results in context. Section V is the main part of the paper and presents the evidence from various behavioral studies. Section VI briefly looks at developments on the abstract track. Section VII reflects on these results and draws conclusions relevant to cognitive roboticists.

II. THE PROBLEM OF TOOL USE

Tool use is an example of problem solving. It involves selecting the right tool or tools, spatially arranging the right relationships between tools and target objects, and performing the appropriate manipulations to solve the problem. The problem may be solved by simple trial-and-error in the world, or by advanced planning. (Thus, we can consider planning as a special case of problem solving). This section defines the problem and outlines the techniques human infants (and some animals) seem to apply to simplify the search for solutions.

We distinguish between 1) general problem-solving abilities (such as planning and search techniques), and 2) domain-specific knowledge of specific actions' preconditions and effects in different situations (i.e., AI planning operators, which we call sensorimotor schemas.) This section focuses on general abilities while we will look at the development of specific abilities in terms of sensorimotor schemas in Section V.

Section II-A looks at the size of the problem space and how it can be reduced by various techniques. Sections II-B and II-C look at research on infants' competences in general problem solving and planning, and how these develop. Section II-D looks at the role of social learning in problem solving. Section II-E

sums up how the development of infant problem solving could be important to cognitive roboticists.

A. Managing the Problem Space

A mathematical formulation of the complete problem in a tool using scenario can consider all the degrees of freedom of the actor, and the objects involved (typically a tool and a target object, but possibly other objects as well). The spatial–temporal relation between objects, which are to interact, is of prime importance; this can be described by a set of relative parameters [16]. First, to determine if an object is suitable as a tool for operating on another object to attain a certain goal, the actor needs to monitor, and possibly react to, changes in these relative spatial parameters [16]; this implies consideration of the objects' shapes and possible spatial relationships. Second, in order to use one object as a tool on another, the values of these relative parameters must be appropriately controlled. One must also consider whether these relative parameters [17] 1) must be produced or maintained sequentially or concurrently (sequentially is easier); 2) require active monitoring for their maintenance; and 3) are managed by direct contact or through the intermediate action of an object. Furthermore, a consideration of required forces and velocities is necessary (which are also parameters).

From this perspective, many tool use problems have a very large problem space. In practice, there are a number of ways in which the total degrees of freedom are greatly reduced, by dealing with the problem space via smaller manageable subspaces (as exemplified in the next paragraph). This happens because infants (and other animals) tend not to tackle the whole problem space to find a solution, but rather their search is constrained by prior experience, habits, and knowledge. The space reduction methods include 1) *sequencing*; 2) *stereotypic behavior*; and 3) *sensory abstraction*.

First, tasks are usually solved by a *sequence of actions*, where each step need only consider (and control) a limited number of degrees of freedom. Consider a capuchin monkey who cracks a nut by first transporting it to a large "anvil" stone, next retrieving a suitable "hammer" stone, and then raising the hammer high to strike the nut [17]. In total there are three objects being put in a relationship, and the number of degrees of freedom is large, but the sequencing of actions leads to a series of smaller problems. It is not necessary to contemplate the relative positions of all three objects; instead one may consider only a pair at a time. To enact a sequence discovered perhaps by chance, coordination of components does not require that the whole problem space be considered at one time (see Section I, M5).

Second, within any step, the motor actions and the sensory elements considered do not include all those available to the animal, but are constrained by existing sensorimotor "units," which often manifest themselves as *stereotypic behaviors*. An easily identifiable sensorimotor unit is the knowledge related to any habitual behavior, for example the banging action of an infant. This is what we called a *sensorimotor schema* (see Section I and also Section III, associating perception and action). Infants (and other animals) tend to constrain the actions they try out on objects to a limited repertoire of stereotypical behaviors [18], even though their motor apparatus has a far greater range of possibilities. This restricts their motion, but

they can perform their behaviors with high variation, so there is a distribution of actions associated to each stereotypical behavior. With the mechanisms of repetition (M1) as well as variation and selection (M2), these schemas lead to a reasonably constrained exploration of the problem space, while still allowing for less constrained exploration when desired [to provoke new results, and lead to differentiation (M3)].

Third, *sensory abstractions* constrain the space; for example, five-month-old infants use only depth and motion to determine object boundaries (and not color, for example), probably because these have higher ecological validity [19, pp. 149]. This perceptual simplification means that such infants face a "smaller" problem than adults in many scenarios, because it is surmised that objects that have lost their depth boundaries are not seen as objects [20], [21]. In terms of development, these units of knowledge may be partially predetermined by genetics, and/or composed by the organism from other units and fragments (see, e.g., Section VI). This paper will sketch this development where possible.

We have sketched the above *strategies* to facilitate the search in the problem space. The fact that humans (or animals) typically do not consider the whole solution space means that they will often arrive at suboptimal solutions, and this is to be expected in a developmental approach. When a solution is first assembled, the animal will tend to perform component parts in a habitual way (i.e., the way in which those components had been performed before they were recruited to solve the current problem), but over time these may be refined to be more efficient for the task at hand. However, an engineer considering all degrees of freedom available to the animal may be able to find a more efficient solution, which would not occur to the animal. The cost of the animal's approach is suboptimality, but the benefit is tractability, because the search for a solution may be intractable if all degrees of freedom are considered.[11]

From a robotics perspective, the infant system seems to start with mechanisms that ensure rather simple state and action spaces, which are then extended over time. Providing a rather simple initial state space seems also to be reasonable in an artificial developing system, and this might even be crucial to making learning and development possible. This idea of constraining the space can be implemented by imposing constraints [23], [24] or by allowing them to arise naturally as the consequence of a simple learning algorithm [25].

B. Planning and Playing

In looking at infant behavior and development from a "zoomed out" perspective, we could see three types of behavior 1) reflex; 2) play; and 3) problem-solving (these are also overlapping waves). Reflexes seem to happen in early stages of development and serve to bootstrap the development process. Play could be described as an affordance-based activity, where affordances of objects in the environment suggest certain behaviors. There is a close relationship between problem solving and free play with infants seamlessly switching between the two. Bruner states, "In play, ends are altered to suit means,

rather than means being altered to achieve an end held constant, as in problem-solving [26]." Infants sometimes lack the capacity to hold ends in mind, and so the means may take over in some problem-solving attempts (e.g., lifting the barrier in Section V-C5). Free play is an effective way for infants to learn about the effects of actions and means–end relationships, so it is an important part of the development of schemas. Play is explained simply by the mechanisms of repetition (M1) and variation and selection (M2), but given a large set of possible activities, an important issue is to decide which actions to do and which actions not to do; this is an area of growing interest in the AI community under the heading of "intrinsic motivation" [27]. In this section, we focus more on problem solving, but also highlight some connections with play (see Section V-C8 for more on play).

Problem solving and planning are very evident in the second year of development, and to a lesser extent before. These are complex activities requiring task analysis, monitoring of the solution, memory to retain goals and subgoals, organization of successive attempts, and the use of discovered information to guide further attempts [28]. This search of the space of actions can happen via (at least) three mechanisms[12]: 1) simple forward search; 2) forward search with heuristics; or 3) means–end subgoaling (discussed below). The first step in problem solving is to choose a goal that is not immediately attainable. The simplest case is where the goal is seen, such as an out-of-reach object. Alternatively, it may be recently seen, such as an object that has just been hidden. Sometimes the goal is unseen, but it may be triggered by the sight of something that is often used as an element in a procedure leading to the achievement of that goal; for example, an infant sees a coat, which triggers the desire to be taken outside. Finally, the goal may be internally triggered by a physical need, such as hunger.

1) Forward Search: This is where actions are tried out in the real world, in an effort to achieve the goal. This does not require any mental simulation of future states or actions. It is required that the infant have the ability to pursue a goal, and in order to avoid exhaustive search of all possibilities, the infant should have some knowledge of how to "use information about the difference between what was achieved and what was intended to guide subsequent activity" [28]. According to studies cited by Willatts, newborns have these abilities, with evidence of some goal directed search with hand to mouth in a limited region. With external objects, goal-directed behavior may appear as early as three months; this was shown for shaking a mobile hanging above a crib, furthermore these 3-month-olds were able to hold a relatively complex goal in memory (the achievement of a certain amount of shaking in a mobile) (see [28, and the studies cited therein]).

2) Forward Search With Heuristics: Heuristics can give an estimate of how likely a potential candidate action is to lead to the goal. An example heuristic could be the reduction in distance between a desired object and the infant; actions that the infant expects will reduce this distance will be chosen in preference to others. If forward search is used with appropriate heuristics, it can lead to a very sophisticated problem-solving behavior,

---

[11]Apart from tool use, the same strategy of degree of freedom reduction is seen in pure motor control problems (see [22, Sec. 3.4]).

[12]These are well known from AI [29, pp. 375–416]

and could explain a great deal of the problem solving observed in infants, even those behaviors thought to be the result of a high-level representation involving mental simulation [28]. For example, the use of a long stick to retrieve an out-of-reach object could be accounted for by a trial-and-error search, with appropriate heuristics.

Heuristics may also aid in the search for the appropriate parameters for an action (for example, the force to be exerted on an object). After varying a parameter, the infant may understand if the variation is "going in the right direction." Evidence that such relationships could be deduced comes from [30] who showed that by 15 months (and not before) infants can predict object weight from size (evident from results on grasp development); given that they only have experience of having lifted a certain finite set of objects, they must interpolate for previously unseen objects; this suggests that in general if they have sufficient experience with the effects of some values of a particular parameter or dimension (in perception or motor control), they could interpolate for unseen cases. The kind of knowledge acquired from such interpolation can help greatly in constraining forward search for problem solving, where an action with variable parameters is being used as a means; given a few trials, the relationship between parameters of the action being applied, and its effects, could be recognized, and so the range of search can be narrowed considerably. This is an example of the mechanism of variation and selection (M2), which allows the relationships between initial conditions and effects to be studied.

3) Means–End Behavior: Here, one starts with the goal and searches for a means to achieve it. The simplest form is where a single means action makes the goal action possible. An example is pulling a cloth in order to bring an object resting on it within reach (where grabbing this object is the goal). Simple means–end behavior has been described by Piaget [2] as emerging at about 8 months. Willatts [28] showed a transition from accidental retrieval to intentional retrieval from 6 to 8 months (see Section V-C3). Furthermore, by 9 months, it was shown that infants can adjust the means action (cloth pull) as appropriate to the goal, in situations where the goal may be far or near. Willatts argues that the basic ability to perform means–end behavior is present in the first 6 months, but only appears for manual tasks between 6–8 months because the infant has just acquired new manual skills and is learning about their effects. In terms of sensorimotor schema, their contexts (preconditions) and effects (postconditions) are becoming refined via this practice.

What is special about means–end behavior here is that it generally involves a composition (M5) of schemas, one acting on the means object and one on the goal, and so the composition implicitly captures a relationship among these objects, and through practice the infant learns this relationship (i.e., learns situations where the composition works or does not); the pattern here is one of fortuitous success, followed later by understanding (see Section V-C3). This type of accidental discovery of relationships among objects could explain the emergence of relational play (i.e., using object–object relationships, see Section V-C3) shortly thereafter. While means–end planning is straightforward in robotics (given well-defined planning operators), we do not find many examples where the accidental discovery of means–end sequences leads to the development of the individual planning operators themselves (see [31] for an exception).

More complex means–end behavior involves working backwards from the goal to find a series of subgoals, which will lead to eventual solution. This requires mental representation of intermediate states (whereas forward search can in principle be done without such representation). Forward search can also be done with mental representation where courses of action are tried out mentally before being tried in the world. As noted above, forward search can also make use of sophisticated heuristics to guide the search, and in observing an infant solving a problem by a sequence of actions it may not be possible to determine if forward search or means–end analysis is being used during that particular episode; Willatts [28] states that he does not know of any empirical way of distinguishing the two alternatives.

4) Affordance-Based Activity Within Planning: In addition to simulating forwards (by heuristic search) or backwards (in a hierarchically directed means–end fashion) there is also evidence to suggest that some fragments of solutions in the middle of a possible sequence may be so compelling that children feel obliged to use them. A study of older children (average age of 32 months) by Cox *et al.* [16] required them to move a disk (with a duck, a swan, a frog, or a fish painted on top) from the center of a circular table towards the boundary (which had a trough painted blue for water). In order to move the disk the children were required to use a cane with a hooked shape at the end. They were presented with the cane in a variety of different starting orientations. Despite the fact that the children could have easily swept the disk to the edge of the table in a single motion, on 79% of the trials the children chose to enclose the disk in the hook, and children almost always chose to move the object closer to them. This suggests that the fact that the hook fitted very well with the shape of the disk triggered a fragment of an action sequence that was too compelling for the children to ignore, even though it did not lead to the most efficient solution to the problem. This is an example of where problem solving and affordance-based play are not separated. Affordance-based play has also been implemented in robots in robotics [32]–[37] (see Section III-C). Learned affordances are used in planning [33], but the affordance-based play is not mixed with planning.

In summary, for complex problems, it seems plausible that a child may see the scene and trigger the simulation of many fragmentary sequences of actions. These may be actions from the current state forwards, or from the goal backwards, or parts in the middle of a sequence; these will then be assembled in some sequence, which is expected to achieve the goal. This may happen at runtime or in advance if sufficient knowledge of actions and effects is available. The important message for cognitive roboticists is that there is evidence that basic planning mechanisms are applied at very early stages of development (and hence are likely to a substantial degree innate) but that the library of planning operators is very limited in the beginning, and this partly explains why not much planning is observed in younger infants.

C. Domain-General Abilities in Problem Solving

The above described the main strategies infants seem to use in planning, but did not address how general planning abilities develop over time. Willatts holds that there are no major discontinuous changes in strategy, but he does describe some of the developments in underlying generic cognitive abilities, which would lead to improved search in older infants [28]. These include the following.

1) *Memory of what has already been tried:* This has been tested with search tasks, to see that the infant does not return to a location that has already been searched. This shows improvement from 14–16 months of age [28].

2) *Backtracking if there has been failure:* This has been tested in a task involving nesting cups. Children from 18–30 months were unlikely to backtrack to a previous configuration, but from 30–42 months there was a significant increase in this behavior [28].

3) *Memory of goals:* Younger children may get distracted and forget the goal, the potential depth of their search is therefore limited [28].

4) *Organization of search:* If the order of search is organized systematically it means there is no need to remember what has been tried. For example, infants seem to try easy actions first, and harder ones later. The evidence for this organized strategy increases from 12–24 months, but it may well be innate, and not very apparent in 12-month-olds simply because they only have simple actions in their repertoire [28].

5) *Inhibition of errors* (for example, the tendency to repeat a previously successful action, in the wrong situation [38]): This does seem to improve throughout infancy, but task-specific effects are very strong.

It is not clear if the delay in these developments simply reflects limitations of the maturing brain, or if it is important to ensure that a high variation in testing of sensorimotor schemas happens before more sophisticated planning is attempted. In any case, cognitive roboticists might be reminded not to try to trigger complex planning with sensorimotor schemas at too early a stage of development but to focus on the grounding of these schemas.

D. A Brief Note on Social Learning

Most examples of tool use we cover rely on social learning, either directly or indirectly (for learning precursors to the behavior). In most socially learned examples of a skilled behavior, there is an element of imitation and an element of self-exploration. For example, in infant learning of self-feeding with the spoon, the infant initially imitates the behavior demonstrated by the adult, but the result is quite crude, and the infant shows little understanding of the various components in the sequence of behavior. Over time (several months), in addition to observing adults, the infant experiments with the constituent parts and refines the behavior, eventually producing an effective behavior. There is also an ongoing interaction between social learning and exploratory learning; there is a limit to how much a learner can advance through learning socially from a master demonstrating a skill, and when the learner does further self-exploratory learning, he or she subsequently can profit more from the same demonstration (because he or she now has a greater understanding of the relationships among relevant parameters in the constituent parts).

In this paper, we focus on the self-exploration part of the problem rather than the social learning. A great deal of literature exists on social learning in infancy [39]–[41] and would warrant a survey of its own. In a robotics scenario it is relatively easy to provide input from a teacher. For example, a human can take hold of the robot's arm and perform an action with it, or the robot can be given a handcrafted example of the correct motion to solve a certain tool use problem in a specific situation. After this, the challenge is to make the robot adapt this appropriately in new situations. This requires the robot to develop an understanding of the relevant parameters and relationships among them in the task, as well as appropriate representations, which is what we focus on in this survey.

E. Conclusion for Cognitive Roboticists

From the preceding subsections we can draw some important conclusions. First, Section II-A suggests that constraints on the complexity of state and action spaces can be designed at the beginning of development, and can help to bootstrap learning and development (and avoid the posing of unreasonably large learning tasks). Second, for domain-general aspects, we have sketched in Section II-B the main strategies (e.g., forward search, means–end), which appear to be in use in early infant planning. It seems that the basic (domain-general) infrastructure for planning is in place relatively early on, but that the library of (domain specific) schemas (planning operators) is relatively empty, so that not much planning will be observed; this library becomes filled during ongoing development by increasingly accurate sensorimotor schemas. Finally, the domain-general aspects of development include the gradually developing abilities listed in Section II-C (for example, memory); it is possible that following such a schedule of development is advantageous so that the younger infant is presented with simpler more manageable problems (much like degrees of freedom are constrained in early tool use; see Section II-A).

Based on this, planning would seem to be relatively easy to emulate in a robot as it matches techniques, which are already mature in AI. However, a proper grounding of schemas in sensorimotor experience seems to require a long development with a great deal of experiences of the causes and effects of sensorimotor schemas, which must be learned by testing them in many different contexts (by playing). This poses a real challenge to roboticists because it requires *a huge amount* of *meaningful* experiences; this is still difficult to achieve by means of real robots nowadays due to unstable hardware and limited sensors (see Section V-B5), and inadequate representational structures for interpreting and assimilating the data. The remainder of this paper will mainly focus on the development of these (domain specific) schemas (i.e., Fig. 1, lower track).

III. BACKGROUND ON SENSORIMOTOR SCHEMAS

The previous section discussed problem solving and planning abilities; this section will focus on the planning operators themselves (here called sensorimotor schemas, see Section I). We look at schemas in psychology and neuroscience, and closely

related constructs, and also the efforts to formalize them computationally.

A. Piaget's Sensorimotor Schemas

The sensorimotor schema has its origins in Piaget's work [2].[13] It includes knowledge about the context in which an action can be performed as well as expectations about the effects. The schema can also encompass some higher level planning knowledge because from Piaget's stage IV onwards (see Section IV-B) there exist higher level schemas, which coordinate relationships among lower level schemas (for example in coordinating means–end behavior). There have been a number of AI works explicitly modelling Piagetian sensorimotor schemas [31], [42]–[44]. These all have three-part schemas of the form context/action/result. These models especially focus on allowing schemas to be learned autonomously from experience, and also to facilitate the composition of schemas (M5), and the construction of higher order schemas from lower order ones [31]. These examples mostly work in simplified simulations; it is likely that additional techniques would be required if one wished to scale them up to the higher dimensional state spaces of robotics scenarios (for example, new sensory abstractions would be needed).

B. Planning Operators in Artificial Intelligence

Since the early days of AI, STRIPS-like planning operators [45] have been used to do planning in closed and deterministic worlds. Every action has clearly defined binary pre- and post-conditions; everything is assumed to be observable and only the agent's actions can change the state space in which operations take place. STRIPS-like planning operators (and later extensions [46]) have been used successfully in restricted domains but the major problem is the reliance on human programmers to predefine all operators. Neither the actual action execution by the robot (which can highly vary with the scene context) nor the pre- and post- conditions are subject to any learning. As a consequence, they require a completely designed world; the limits of this approach have been accurately pinpointed by Brooks [7] and Sutton [47]. There are, however, some more recent works in AI planning that can learn planning operators [48], [49]. These effectively learn pre- and post-conditions, but lack the mechanism of motor differentiation (M3) to generate new actions.

C. Affordances

There is a close relationship between sensorimotor schemas and the Gibsonian notion of affordances [50]. The visual perception of a handle of a cup, for example, can be associated with the action of grasping it and the effect of having a stable grasp. AI implementations inspired by Gibsonian affordances

[13]The use of sensorimotor schemas in analysis does not necessarily mean that we need to bring along all the Piagetian baggage, which many have criticized; e.g., Fischer [4] says Piaget's schema is too general, because according to Piaget's theory two skills using the same schema should happen at the same time, and this cannot explain the phenomenon of horizontal décalage (for example, that conservation of liquid is acquired long before conservation of mass). This can be circumvented simply by making schemas less general. We may have context specific schemas for each skill, which need not arise together (although some processes such as Representational Redescription [14] may later recognize the similarity).

arrive at knowledge structures very similar to the implementations of Piagetian schemas [32], [33], coding the effect of executing a particular behavior on an object. In other robotics work, Modayil and Kuipers showed in [51] how the effects of a physical robot's actions on its perception can be learned. Fitzpatrick and Metta [35] learned affordances for pushing objects and made initial steps towards learning categories (e.g., roll-able). Stoytchev [36] was able to learn tool affordances for a set of tools by correlating the tool used, its movement, and the effects on other objects in the scene. Affordances for grasping have been learned through "play" and associated with object models [52]. Hart and Grupen showed how intrinsic reward can be given for the discovery of affordances [37]; in this work generalizable control programs can be learned and can evolve from one task to the next. These generalizable control programs could be understood as the action core of a sensorimotor schema.

D. Schemas in Neuroscience

Neuroscientific evidence from monkeys shows that object shape is coded in a motor area of the brain, which is involved mostly in the control of hand movements [53]; the authors of this study concluded that "every time an object is presented, its visual features are automatically (regardless of any intention to move) 'translated' into a potential motor action. This potential action describes the pragmatic physical properties of the objects." This lends credibility to the idea of a sensorimotor schema as a unit of knowledge, which links a particular perception with an appropriate action. This aspect of brain architecture has also been modeled computationally [54].

E. Limitations of the Sensorimotor Schema

The sensorimotor schema always combines issues of perceptual, motor, and cognitive development, and in this, it is not always compatible with contemporary views. For example, in Piaget's view, perception tends to be built up by experience with acting in various contexts. While this has clearly been shown to be the case in an experiment with cat locomotion, for example in [55], there is less evidence to support all the cases in which Piaget held that the same process occurs. A classic example is in the case of the means–end action of retrieving a hidden object. Piaget held that it was through experiences with acting on objects in relationships such as "in front" that the perceptual and representational competence to understand about hidden objects was constructed. Contemporary views hold that many perceptual competences may be more independent from action competence, and that in many cases perceptual competence might come first [19, pp. 247 and pp. 260]. Nevertheless, even if perceptual competence does lead, there may still be a place for action to help with the interpretation of that perception [56, pp. 176]. It seems that the idea of sensorimotor schemas is not completely invalidated by later results, but the schema is not the only unit of knowledge and may need to be complemented with pure perceptual or motor competences, which may mature according to some internal developmental processes [for example the onset of stereoscopic depth perception [19, pp. 96], or the arrival of stereotyped behaviors (see Section V-A)]; once they do become available, it seems plausible that sensorimotor schemas may again come into play to integrate them [57, pp. 148]. Fig. 1

(lower track) shows some nodes that do not have precursors, and these correspond to sensorimotor schemas that integrate newly matured perceptual or motor competences.

It is not clear if the sensorimotor schema can also account for high-level representations of objects, for example 3D models that can be manipulated mentally. Knowledge of objects, which is abstracted from any particular action, would seem to be a prerequisite for more advanced types of planning. In particular, the fact that older children and adults do have excellent generalization abilities across domains indicates that representations, which are relatively independent of particular tasks, must arise during development. However, the developmental evidence suggests this is a long and protracted process. This is illustrated in the upper (abstract) track of Fig. 1.

F. Challenges in Formalizing Schemas

The major challenge in formalizing sensorimotor schemas is probably to capture the dynamic properties of its actual acquisition process. This acquisition requires us to 1) ground the sensorimotor schemas in sensorimotor experience (meaning to learn suitable pre- and post-conditions as well as associated success likelihoods) by means of the mechanisms M1 and M2; 2) define processes that lead to the creation of new sensorimotor schemas by means of the mechanisms M3–M6; and 3) integrate the sensorimotor schemas into the dynamically changing abstract world knowledge by means of representational redescription (M7). This final point means that at advanced stages of infancy, schemas must be referring to relationships among objects in their contexts and results. The process of schema acquisition requires the performance of a huge amount of meaningful actions providing the required experience to feed the developmental process in a complex cognitive architecture. This architecture is likely to be equipped with a not insignificant degree of innate structure such as innate behaviors and a rudimentary planning machinery (see Section IV-A) based on which the still to be acquired schemas develop from initial sensorimotor schemas. In this context, we recently introduced the concept of so called "object-action complexes" (OACs) [58] as one possibility for formalizing the complex acquisition process associated with sensorimotor schemas. As for the process of representational redescription, perhaps the closest analog in robotics is Bratko's use of ILP [59] to find "insights," leading to a "better representation language which facilitates a more compact representation of the current knowledge."

In this paper, we give evidence that the development and enrichment of sensorimotor schemas can be directly linked to the development of tool use, and hence provides a means to tackling the same problem in robotics. Insight into the development of sensorimotor schemas in infants hence can also guide the design of artificial agents that learn to use tools.

IV. Perspectives on Development

This section provides an overview of the different psychological approaches to understanding infant development. Section IV-A looks at maturation and learning during the process of development. Section IV-B reviews Piaget's theory, which is probably the most well-known theory of cognitive development. Section IV-C looks at Siegler's overlapping waves theory (from which we borrowed our three "stages" as seen in Fig. 1). Section IV-D looks at the more recently popular dynamic systems perspective on development. Finally, we look at how a consistent picture of development could be found in these theories.

A. Development: Maturation and Learning

Development includes both maturation and learning. Maturation is a change, due to biological growth (or aging) in the organism, without the need for environmental influences [56, pp. 3], e.g., growth of certain centers in the brain. Learning is a change due to information processing; for example, a change in the organism's competence resulting from the processing of information from the environment. An extreme nativist viewpoint would posit that all development is due to maturation, with new brain structures unfolding according to a preset plan hard-wired by the genome. An extreme empiricist viewpoint would posit that all development is due to learning, with new mental structures being constructed due to the processing of new information from the environment (i.e., the software is changing but not the hardware). Contemporary viewpoints lie between the extremes. The evidence from the literature suggests a very complex bidirectional interaction between physical changes (in brain and body) and mental changes due to learning [60]. The impact of changes in the body has been studied in the development of locomotion [60], and body changes must also pose problems for infants learning to use tools, but we know of no studies addressing this. In this paper, we cannot address brain or body changes in any detail since we are mainly concerned with the description of the observable development of sensorimotor schemas. The issue of interaction between innate structure and learning is only indirectly observable (as with the interaction of the concrete and abstract track indicated in Fig. 1). This problem is also fundamental for roboticists since they need to determine the prior structures of the systems they design.

There is quite a bit of research on the innateness of certain kinds of knowledge in neurophysiology (see, e.g., [61] and [62]) and developmental psychology (see, e.g., [19]), which allows for postulates on reasonable innate structures in robot systems (see also [52]). In Section V-A, we particularly point to a number of innate behaviors that are used to bootstrap the developmental process (see Section V-B); we also gave evidence for a certain degree of innate machinery for planning in Section II-B.

B. Piagetian Schema Development

Piaget's theory is called constructivism and is based on the idea of the infant gradually building up knowledge structures as he/she interacts with the environment. On the nature/nurture spectrum, it is closer to nurture (i.e., the empiricist viewpoint in Section IV-A). Piaget uses the *sensorimotor schema* (see Section I) as the unit of knowledge. Piaget defined six sequential stages during sensorimotor development (from approximately 0–2 years) [2], [20], with a qualitative difference between the sensorimotor schemas in use in each stage. Piaget's early stages roughly map to the three stages we outlined in Section I.

Piaget's Stages I-II roughly correspond to our Stage 1 (see Section V-A), including reflexes such as sucking and grasping,

as well as integrating different modalities (auditory, tactile, visual), and mastering reaching to grasp. Each behavior is associated with its own global schema, which generalizes from experiences where the action happens, and recognizes the situations where the action is triggered, and the expectation of what sensory impressions arise while the action is in progress.

Piaget's Stage III roughly corresponds to our Stage 2 (see Section V-B) and involves repeating results fortuitously discovered with objects in the environment, such as shaking a rattle. During this stage there is a rapid growth in the number of schemas in the infant's repertoire, as new schemas are differentiated (M3) from previous ones in order to repeat interesting discoveries (e.g., squeezing, shaking, striking, scraping, rubbing, and pulling).

Piaget's Stages IV-V roughly correspond to our Stage 3 (see Section V-C); means–end sequences of actions are performed. For example, the infant will intentionally displace an obstacle in order to retrieve a desirable object that is visible behind it. This requires two distinct sensorimotor schemas–one for the means action (displace the obstacle), and one for the end (grab the desired object). This implies that the sensorimotor schemas must now incorporate relatively advanced knowledge of the world; they must capture the effect of an action on the relationships between objects (for example the relationship "in front"). Schemas are also intentionally varied (Piaget's Stage V) so that the relationships between initial conditions and effects can be studied. Piaget's Stage VI (roughly from 18–24 months) involves internal representation of objects, actions, and effects; this gives rise to covert planning (though Willatts is sceptical and sees a more continuous development in planning abilities [28]).

Through all this progression, there is a gradual increase in the abstractness and objectiveness of the knowledge captured by sensorimotor schemas; earlier schemas capture subjective knowledge locked in particular contexts, while later schemas abstract away from these contexts and capture knowledge about relationships between objects and actions in the world.

Piaget's stages do not have crisp boundaries between them, and some behaviors are intermediate. The stages also have significant overlap, so that a child who acquires his first Stage V behaviors will also be spending a significant proportion of his time engaging in behaviors belonging to earlier stages. However, the sequential ordering is strict, i.e., a child who exhibits behavior from stage n must have previously exhibited some behaviors from stage $n - 1$.

In Piaget's theory, development happens either through the modification of individual schemas or the relationships between them. Within each of his six sensorimotor stages, schemas individually develop; by being executed in varied situations, they refine the motor action of the schema and also refine their knowledge of the various effects produced in various contexts. Transitions between stages are explained through coordinations among schemas. For example, means–end behavior emerges in his fourth sensorimotor stage, and this development is explained as a process of coordination between the schemas of previous stages. For example, the schema of hitting an obstacle (means) could be coordinated with the schema of grabbing an object (end) in order to remove an obstacle to prehension. Computational modellers of Piagetian development (see [31]

and [42]–[44]) have captured hallmarks of his theory; e.g., the acquisition of higher order knowledge based on basic schemas from a previous stage. However, these are only demonstrated in "toy" domains; we believe it is necessary to work in more realistic domains (i.e., closer to what an infant experiences) in order to get close to modelling the mechanisms of development that an infant needs to use. The history of AI shows that techniques that work in toy domains do not necessarily shed any light on the techniques needed for realistic worlds [7].[14]

In the last few decades, a great deal has been written about where Piaget was right and wrong; Siegler gives a good brief account in [63]. In summary, sometimes Piaget overestimated infants' abilities, and sometimes he underestimated them. However, possibly the biggest problem from a computational point of view is simply the vagueness of his theory; it gives a rough sketch of how the development happens, but leaves the mechanism of development very underspecified. Despite all the criticisms, Piaget's theory remains one of the few attempts to explain the whole of development, and quite probably is a reasonable sketch of the outline of how the mechanism of development works.

C. Overlapping Waves Theory

Siegler's "overlapping waves" theory of development [64, pp.7] holds that at a particular age a child will have a number of different strategies for tackling a problem (for example, this could be ways of approaching a particular tool use problem); these different ways of thinking are all active at the same time and may give rise to different conclusions, thus explaining how a child may approach the same problem in different ways on successive days. The different ways of thinking continue to compete with each other over long time scales (e.g., several months); with development, there are gradual changes whereby more successful ways of thinking become used more frequently, and others are used less frequently.

We are not aware of any computational work that emulates these overlapping waves; computational approaches tend to seek to learn the "correct" strategy and then to stick with it.[15] In contrast, "overlapping waves" seems to be a "sloppy" way of thinking; for example, if one strategy is clearly leading to failure and another to success, it would seem logical to abandon the first; however, children tend to continue using "wrong" strategies (albeit less frequently) for some time. It is possible that this approach leads to increased robustness, because typical interactions in real-world situations have many uncontrolled variables and do not give such clear cut results as a science experiment would. In such situations, it may make sense not to abandon any alternative for quite some time, so that there are always alternative strategies to fall back on if the one that first appeared promising eventually proves not to be so. Furthermore, it has been shown that children who exhibit

---

[14]We do not see the need to work with real robots; simulation may be adequate provided the world is rich enough to allow typical infant tasks to be attempted. Conversely, real robots may be inadequate if the tasks are oversimplified.

[15]One could consider reinforcement learning (RL) with options to be similar, because an option may still be selected even if it has not been yielding good rewards recently. However, Siegler's "overlapping waves" are, generally speaking, about higher levels of abstraction than those found in RL options.

more varied ways of thinking learn more from training [65], and more generally, variability in psychological development may play the same critical role that it plays in evolutionary development. In Fig. 1, the overlapping waves drawn on the lower part apply equally well to the representational redescriptions in the upper part of the diagram; i.e., older (more context specific) representations will not be immediately retired when newer (more generic) representations come online; the alternative representations will continue to operate in parallel for some time, with one or the other being used depending on the task.

Siegler broadly agrees with Piaget's constructivist theory, but he also highlights the importance of aspects that might be neglected by an excessive focus on constructivism, for example, the acquisition of associative knowledge (learned in specific contexts) or more generally, the issue of knowledge retrieval processes [65]. Siegler points out that it may not make sense to ask "whether children 'have' a concept or strategy or theory at a given age;" instead it may make more sense to investigate "the set of conceptualizations and strategies and theories that children know and the mechanisms by which they choose among them" [65].

D. Dynamic Systems Approach

Piaget gives the impression of a rational infant that will take sensible actions if he/she has the relevant knowledge. The dynamic systems view [60] attempts to explain behavior at a lower level, via the activation and interaction of various low level processes such as perceiving, moving, and remembering; the eventual behavior observed is explained in terms of these processes and may not always appear rational from a more global perspective. This can lead to different conclusions being drawn from behavioral studies. For example, according to Piaget's view, if an infant knows where a hidden object is, then the infant can be expected to attempt to retrieve it from there; however, in the dynamic systems perspective, an infant may reach toward the wrong location because of an inability to suppress a response performed earlier [38], or simply because some alternative action has a higher activation (even though the infant might at the same time have an expectation of perceiving the object in the new location). The folk psychology concept of "knowledge" is at too coarse a grain for dynamic systems explanations, so that the question of whether or not an infant really "knows" something (e.g., the location of an object) is not meaningful; the behavior emerging from the infant's lower level processes may seem to demonstrate knowledge under some circumstances and not others (i.e., context dependent). This relatively new approach to understanding development helps to explain earlier observations that often noted infants' considerable difficulty in inhibiting "obvious" actions, or actions in progress [15]. It also could explain some of Siegler's observations of the context specificity of knowledge, and the switching between different strategies in different circumstances. Dynamic systems have been employed in robotics for motor control [37], [66], and some higher level aspects of cognition, e.g., for scene understanding [67]. However, these computational works, to date, only capture a very small part of the scope envisaged by the

psychologists: e.g., it is envisaged that dynamic systems could be employed to account for the development of a grounded understanding of human concepts; Thelen and Smith [60] outline how a concept such as force could be developed by generalizing from walking, reaching, crawling, and pushing.

E. Conclusion on Developmental Theories

The theories sketched above are not entirely consistent on all details, but it is possible to find a consistent theory, which incorporates their major aspects, with some adjustments. From Piaget's theory, we can take the notion of sensorimotor schemas, and a mechanism of development, which builds new schemas by operations such as differentiation or composition of old schemas (see M1–M6). The overlapping waves theory can be accommodated by ensuring that older behaviors will not be replaced at once when newer more sophisticated behaviors develop; instead both will continue in parallel and may be elicited in different contexts. The dynamic systems approach impresses on us the necessity to model at a fairly low level, so that sensorimotor schemas may be quite context specific, and triggered in certain situations, without a global overview ensuring consistency and rationality in behavior. This means that developments to new "stages" do not happen all at once but include a protracted phase of intermediate behaviors where behaviors in some domains are more advanced than others. Abstract domain-general knowledge and representations may be very slow to arise. The brief review above also supports some of the mechanisms we identified in Section I as underlying development. For example, the mechanism of variation and selection (M2) is very evident and is believed to be a primary mechanism in development both at low levels such as learning motor synergies [68] and also at higher levels in selecting which strategies to use [65].

V. THE BEHAVIORAL STUDIES (CONCRETE TRACK)

This section looks at behavioral developments in the first two years of life, which we believe to be relevant to the development of tool use. The developmental descriptions in this section (and Section V-C specifically) are not intended to comprehensively describe how particular examples of tool use develop, but instead to be suggestive of the ways in which aspects of tool use development unfold in infancy. We stress the importance of sensorimotor learning as a precursor to tool use, and we explain how the six mechanisms (distilled in Section I) are in operation in the examples to follow. To give a complete developmental trajectory for a particular tool use example would require extensive longitudinal studies, which have not yet been carried out.

This section is organized roughly in order of developments, which build on each other, including the supposed precursors to tool use, and simple examples of tool use of increasing complexity. These behaviors are summarized in Fig. 1 where there are three overlapping waves for the three types of behavior:1) behaviors without objects (see Section V-A); 2) behaviors with single objects (see Section V-B); and 3) object-object behaviors (see Section V-C). These subsections cover the precursors to tool use, and then Section V-D takes an in depth look at one particular example of tool use, which is common in the second
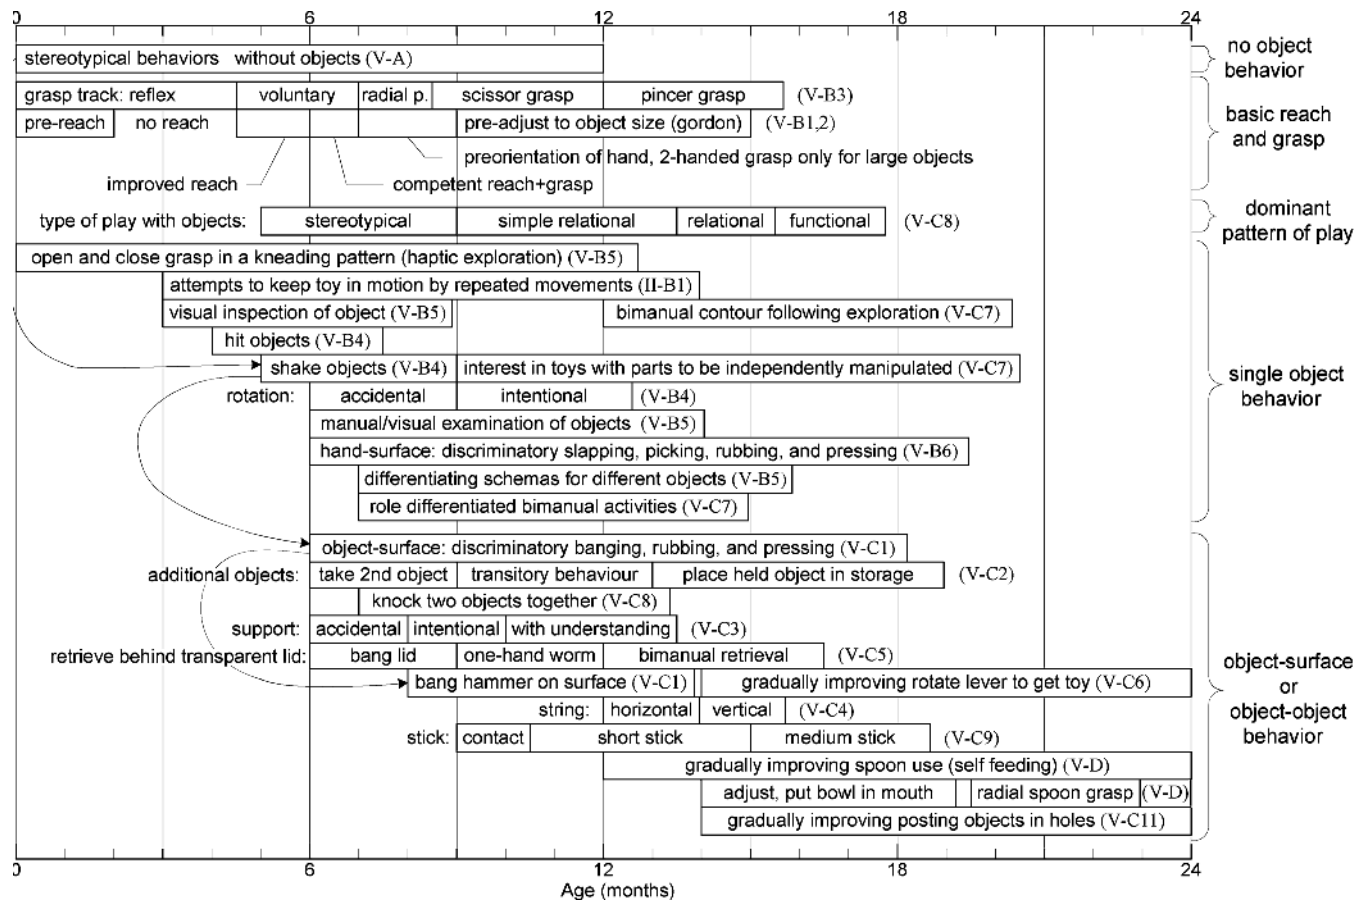
Fig. 2. Behaviors at various ages. The left side of each rectangle indicates the age at which the behavior appears; the right side is meaningless. Cross references to the section of this paper that describes the behavior are given in parentheses. The arrows show an example of a chain of supposed precursor behaviors. Some of the ages here come from Uzgiris and Hunt [69] with the warning that they should not be interpreted as typical for infants in general. They are provided to indicate an order of progression, rather than precise timings.

year of infancy: self-feeding with a spoon. Fig. 2 graphs the individual behavioral developments covered in this section.

A. Behaviors Without Objects

Here, we analyze some typical behavior patterns of infants, which do not require manual contact with objects or surfaces. First, there are a number of "reflexes" such as reaching or rooting for the breast. However, von Hofsten [70] cites evidence showing that these and other examples of supposed reflexes do not in fact share the expected properties of reflexes (e.g., elicited and automatic) and, in fact, turn out to be under voluntary control. He states that, as with other mammals, it should not be surprising to find sophisticated prestructured actions in human neonates. Second, there are "rhythmical stereotypical behaviors," which Thelen describes as being more complex than reflexes, but less variable and flexible than full voluntary behavior [18]. This lack of variability may be an example of the strategy of initially reducing the degrees of freedom of the motor control problem (see Section II-A). Roboticists have also implemented a similar strategy to reduce the variability in early actions, using rhythmic movements [71] or "goal babbling" [72].

Thelen [73] observed infants longitudinally during the first year, and recorded all rhythmical stereotypical behaviors; this

meant any movement that was repeated at least three times at regular short intervals of about a second or less. Forty seven distinct behaviors were observed, appearing at different times. We will describe eight of these, involving arms, hands, and fingers, which seem most relevant as precursors to tool use (rather than leg movements, or whole body movements, etc.). The numbers in parentheses describe the percentage of sampled infants who exhibited the behavior.

- Arm wave (100%)–a rapid flapping of the arm vertically from the shoulder . This leads to surface slapping behavior (see Section V-B6), and also waving of objects and banging them on surfaces (see Section V-C1).
- Finger flex (100%)–lexion and extension of all four fingers simultaneously, and often the thumb. This probably leads to exploratory behaviors with objects (see Section V-B5).
- Hand rotate (90%) and flex (80%)–a rhythmic rotation, bending, and extending of the wrist. This is subsequently performed with objects; possibly it is used in object exploration (see Section V-B5).
- Clap hands together (75%) (referred to by Thelen as pat-a-cake)–this later leads to stereotyped banging objects together (see Section V-C8).
- Arm fly together (20%) and plucking (15%)–these were similar to clapping, but the hands were not extended to

slap palms together; hands were brought together and then thrown apart; these may also be precursors to banging objects together (see Section V-C8).

- Finger rotate (15%)–similar to "the movement used in turning a large dial, where the fingers are rotated slightly outward" [73]; this may lead to rotation of lids/dials.

Other stereotypical behaviors relevant to tool use are only performed with objects, and so are covered in Section V-C1.

The percentage of the infants' time spent engaging in these movements rose during the period from about 1 month through to 6–7 months [18], after which it plateaued and then fell off towards the end of the first year. The average time spent engaging in the movements at the peak was approximately 9%. Although the overall frequency of stereotyped behaviors declined in the second half of the year, the number of different types of behaviors rose, because new behaviors were added without the loss of older ones (unfortunately, the study did not report on which specific behaviors appeared at which times). Behaviors tended to be more variable around the time of their first appearance than later [which may be the mechanism of variation followed by selection, (M2)].

Since the behaviors described above have no obvious precursors, and given that older children and adults do not perform these behaviors, it is not possible that they were imitated, and it is surmised that they are innate. Furthermore, Thelen states that "the onset of particular stereotypes is largely dependent upon events intrinsic to the infant" [73]. Thelen suggests that the behaviors are not much affected by the environment, but rather "internally guided" [73]. Thelen suspects that the behaviors may emerge as by-products of the normal maturation of motor control circuits, but that they may be opportunistically used by infants for the purpose of bootstrapping further development, for example, by encouraging actions, which will at some point lead to interesting results. Evidence that spontaneous behavior such as kicking can be transformed into an instrumental behavior comes from Rovee and Rovee [74] (as early as 10 weeks) and also Piaget's work (his second sensorimotor stage, which is at approximately 4 months). This follows the pattern of accidental discovery followed by intentional exploitation (already seen in Section II-B, for means–end behavior, and which we will see arise in many later behaviors). Furthermore, Thelen [18] speculates that these stereotyped movements may be incorporated in hierarchically structured advanced skills [composition (M5), and means–end in Section II-B3)], where the stereotyped movements form low-level subunits. Recent AI work shows how this could fit in the framework of hierarchical reinforcement learning [71].

Apart from stereotypical movements, calibration of the vision, motor, and proprioceptive systems seems to also take place in the first months of life, which is related to various calibration tasks connected to vision-based robotics [75]–[77]. Infants may regard their hands moving during the second month, but the vision does not guide the hands [2, pp. 102]. Subsequent to this, vision augments the activity of the hand. Infants engage in extensive self-exploration by 2–3 months [78]. Young infants, when viewing their own movements, are sensitive to visual-proprioceptive contingency; infants that are 3 months are able to discriminate between direct and delayed views of self-produced

leg movements. [16] Rochat [78] concluded that there is evidence for a perceptual-based *body schema* at this time. It would be interesting to know how this body schema can predict interaction with seen objects. Bower [79, pp. 123] has shown that there seems to be an innate knowledge[17] of when a primitive reach should contact a seen object because infants were distressed by a "virtual object," which they could see and reach for, but which produced no tactile sensation. However, this expectation of contact is likely to be very specific to reaching for a target; we are not aware of any experiments, which determine when the infant displays knowledge of expected collisions with objects that are not the targets of reaching.

In summary, these behaviors without objects give some of the initial sensorimotor schemas, which form the beginning of the developmental story illustrated in Fig. 1; together with the basic mechanisms M1–M6, they start the developmental process, and become differentiated (M3) and composed (M5) to produce new schemas. This happens because these initial behavioral patterns cause interesting events to occur (touch, sound etc.), leading to differentiation of schemas (M3); subsequent variation and selection (M2) helps to fine-tune these new schemas so that they become increasingly predictable (and potentially intentional). Additionally, some of these innate patterns seem to trigger multisensory expectations indicating an innate multi-sensory experience space [80]. These behavioral patterns play an important role in learning a body schema.

B. Behaviors With Single Objects

1) Learning to Reach and Grasp: Reaching is an obvious precursor to dealing with objects, but it is probably particularly relevant to simple tool use because similar problems seem to be involved. For example, in learning to control a stick there is a similarity between learning to bring the hand in contact with a seen object, and learning to bring the stick (which could be viewed as an extension of the hand, see, e.g., [81]) in contact with an object. One needs to solve a degrees of freedom problem to control joints, and a scaling problem to apply the correct force to each element; furthermore there may be visual feedback (servoing) to control the movement to the target. It is probable that evolution has developed innate routines to bootstrap the development of reaching, and it is plausible that some of that innate machinery may be reused for tool use.

Neonates seem to have a very premature reach and grasp mechanism, which reduces in frequency over the first two weeks of life, and is hard to elicit in the period from 4–20 weeks [79], [82, Ch. 6]. The neonate primitive reach motion is visually elicited, and ballistic, with no visual feedback to correct the reach motion while it is in progress [56, pp. 38] and it has a relatively low success rate (9–40% [19, pp. 250], [79]). Bruner [83] suggests that this type of innate response is coded in by evolution to serve as a "launching stock" from which skilled actions can then be constructed. The mechanism of variation and selection (M2) seems to be important in building on such innate patterns to develop mature reaching, as shown by Thelen *et al.*'s

---

[16]Leg movements are also included in Thelen's stereotypical behaviors, but we have omitted them here because they seem less relevant to tool use.

[17]His youngest experimental subject was 4 days old

[84] detailed study. One striking result of their study was that there were dramatic differences in how four different infants first reached for toys, and how they developed; each seemed to have their own developmental pathway.[18] There is little evidence that the capability to grasp accurately is present in the first month of life, although some elements such as hand opening are present [56]. The hand opening disappears at about two to three months, and the behavior becomes more of a swiping with a closed fist, which is then replaced by an open-handed reaching, but with no grasp [85]. The reduction in reaching at about 7 weeks does not appear to be due to any loss of interest in nearby objects, because visual fixation increases; it is possible that the reduction in interest is because of the excitement of looking inhibiting the motor response, but also there must be some reorganization of the motor control to explain how the hand opening becomes separated from reaching [86].

It is not until about the fourth month that proper visually elicited grasping will commence and that the infant will bring the hand into view to grasp seen objects even when the hand is not initially in view [69, pp.110]. The more mature reach and grasp, which appears at about five months has about an 80% chance of contacting the target, with the possibility of visual feedback to correct the reach in progress, although the grasp coordination appears to have regressed [79]. It seems that the primitive reach-grasp was undifferentiated (i.e., the reach and grasp are coupled), and by 20 weeks two separate motions have replaced it (M4), which can now be executed independently or coordinated. Gordon [87] also notes that this type of regression followed by advancement is common: reorganizations can initially result in a decline of the motor skills, before improvement. After this period, the coordinated reach and grasp develops rapidly; Bower reports 100% accuracy on visually elicited reaching [79, pp. 174] at 6 months.

Studies with prisms, which distort the infant's vision, show that the infants can use the visual feedback to correct their reaches [88]. However, the visual feedback does not seem to be necessary for normal successful reaching; studies of reaching for glowing objects in the dark [89] conclusively show that sight of the infant's limb is not necessary either at the onset of successful reaching or in succeeding weeks; proprioceptive information must therefore be employed. Nevertheless, vision does seem to be important in normal infants' development of a common mapping system for auditory, proprioceptive, and visual information, because reaching in blind infants is substantially delayed [89].

A computational model of the development of reaching appears in [90] and is also shown to model recalibration in a condition analogous to prismatic adaptation studies with infants [91]. The robotics community has developed multiple methods for integrating visual feedback into reaching movements and they are called *visual servoing* (for an overview, see [92], [93]). A more developmental inspired approach to learning to grasp is shown in [94]. Before being able to apply this method to the object that should be grasped, the object needs to be determined in the

scene. One method that segments an area (the potential object) in a biologically motivated fashion is [95]. The method is able to do this segmentation procedure based on a starting point on the object. This requires a visual attention mechanism to focus on the possible interesting objects in the scene (see [96] for an overview from a cognitive robotics perspective).

2) Refining the Reach and Grasp Movement Parameters: Bruner [83] sees a commonality between the development of reaching to grasp, and the development of other goal-directed skilled actions. In each case the behavior starts out with a series of component acts (for reaching these include raising of arms, ballistic flinging, and closing the hand), but the sequence is not correct and the acts are crude; once the sequence becomes correct each component is "drastically altered to fit task requirements." With more practice the whole sequence becomes energy efficient, which suggests that there may be a feedback in the system based on efficiency. Eventually, the whole sequence becomes modularized (M6) so that it can appear as a component in new higher order sequences.

There is a developmental progression in the infant's use of information about the size of objects for grasping; infants as young as eight weeks make more reaches to a graspable ball than to one that is too large [56, pp. 43]); 5-month-olds tend to reach with two hands regardless of size, whereas 7- to 8-month-olds use two hands for large objects more often than for small ones, and at 11 to 12 months reaching closely reflects the object's diameter [97]. A similar pattern appears for the thumb–index finger angle opening during the reach, which increases after 7 to 8 months, as well as the adjustment of the angle to the object diameter, and the proportion of the object within the hand opening at touch [97]. In studies of 5-, 9-, and 13-month-old infants [56, pp. 44], it was found that all infants began hand closure before contact, but younger infants began closure later.

For orientation, it was found that 9-month-olds rotated their hand to adjust to the orientation of a stick before grasping, whereas 5-month-olds did most of the adjustment after contact. A detailed longitudinal study [85] has shown a qualitative change in grasp preorientation occurring between 5 and 7 months, which is roughly in line with other results [98]. Investigations of the importance of vision to preorientation [99] have shown that 7-month-olds could preorient correctly to grasp a glowing rod in the dark, showing that visual monitoring of the hand's orientation was not necessary. For 9-month-olds, it was shown that they could orient correctly and grasp the rod if shown the rod only before reaching, but the rod remained darkened during reach onset and grasp. Again (as in the previous section), proprioception must be used here, and the authors suggest that vision is used to calibrate the proprioceptive information from the limbs, and thereafter vision is not necessary.

These results on size and orientation show that fragments of a practical object representation are present at this early stage (see also Section VI); by "practical" we mean that it may be encoded via the motor action that can grasp it. Such representations allow for the distinction of object size and orientation, while the strategies to adapt grasping appropriately evolve over time. We surmise that these early fragments of representation form the basis for later more complete and action-independent representations (see Fig. 1, upper track).

---

[18]Such development pathways give useful material, which helps in the discovery of the underlying mechanisms of development. Also, they illustrate a degree of robustness and generality in these mechanisms, in that they can cope with a variety of different configurations.
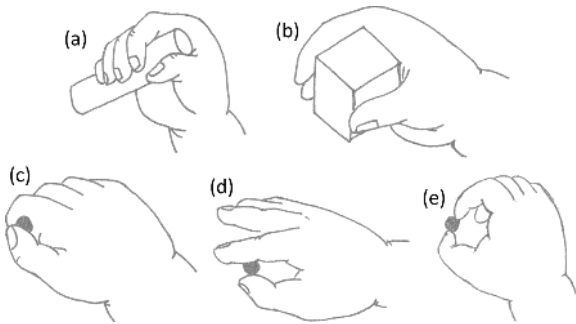
Fig. 3. Developing grasps. (a) Voluntary palmar grasp. (b) Radial palmar grasp. (c) Scissors grasp. (d) Inferior pincer grasp. (e) Pincer grasp.

3) Development of Grasping Competences: Infant grasping abilities develop throughout the first year of life. There is an initial reflex, called the "neonatal palmar grasp," which is present from birth, and where the fingers close on stimulation of the palm; there are conflicting results [100] about whether this disappears gradually over the first year, or disappears after two months, or gradually merges into voluntary grasping.

With regard to voluntary grasping, Touwen [100], in his longitudinal observations of 27 male infants, noted a development through five different phases of grasping. The *voluntary palmar grasp* [see Fig. 3(a)] uses the whole hand (appearing on average at about 4 1/2 months). The *radial palmar grasp* [see Fig. 3(b)] is performed mainly with the area of the palm between the thumb and the second and third finger, i.e., with the thumb opposed (appears on average at about 7 months). The *scissors grasp* uses the volar sides (the volar is the lower surface of the hand, i.e., which includes the palm) of the extended thumb and index finger (appears on average at about 8 1/2 months), or alternatively the side of the curled index finger [101], as shown in Fig. 3(c). There follow some intermediate stages before the proper pincer grasp that include the *scissor-pincer grasp*, which uses the tip of the index finger and the volar side of the extended thumb (appears on average at about 11 months), an alternative *inferior pincer*, which uses the tip of the thumb and volar surface of index finger [101] [see Fig. 3(d)] and the *pincer grasp* [see (Fig. 3(e)], which uses the volar surfaces of the tips of the thumb and index finger almost exclusively (appears on average at about 12 months). Development is by no means finished here; the second year will see an improvement in the use of appropriate forces, and the grasp will not approximate adult performance until 6 to 8 years, with further subtle improvements in adjustments of force and anticipatory control continuing until adolescence [87]. This gives potentially valuable indications for roboticists on how the control of artificial dexterous hands can be learned in some hierarchical scheme starting from a coarse to a fine grained representation. We are not aware of any works in robotics that have already addressed this problem directly. Instead, works that transfer grasps known for a specific object to new situations reactively (e.g., [102]) or that transfer grasps between similar objects (from the same category) [103] have been shown.

4) Developing the Stereotypical Behaviors With Objects: Once reaching and grasping are mastered (at about 5 months),

the stereotypical movements discussed in Section V-A (without objects) can easily be performed when an object is grabbed, and this is exactly what happens [due to the mechanism of repetition (M1)]. Thus, arm waving becomes waving with an object held in the hand, such as a rattle. The opening and closing of the fingers can be done with an object, either to catch (and release) it or scratch it. These behaviors lead to new interesting effects, and so are reinforced (corresponding to Piaget's third stage which occurs from about 4 to 8 months). This is quite a clear example of the process of differentiation of schemas (M3), which can be followed by variation and selection (M2) to refine the newly differentiated schema. Other behaviors appearing include rubbing an object with the hand, or squeezing an object, or crushing (e.g., paper). The use of stereotypical behaviors on objects, to learn the effects, matches closely robotics work on affordances [33]; however, we are not aware of robotics work, which also does differentiation (M3) to adjust the motor action to better achieve new effects, and eventually to branch out to a series of new behaviors, as depicted in moving from left to right in Fig. 1 (concrete track).

Some developments with a superficially similar appearance to tool use can also appear at this time, for example, an infant can pull strings to shake something tied to the end, or if the infant is in possession of a stick, and it accidentally hits an object, then this action can be repeated. This "accidental to intentional" pattern is a feature of the repetition mechanism M1. However, the combination is discovered by accident and the relationship between the objects is not understood; this is shown by the fact that strings that are close, but not connected to the object, will also be pulled [2], and in the stick example, there is no ability to control the direction in which the stick is pointing.

Infants may also rotate an object during random exploration at 6 months [2], although it is doubtful that it is done with the intention of seeing the other side, as they are incapable of fully turning objects to find a hidden desired side until about 9 months [69, pp. 120]. They may, however, intentionally half-rotate an object to bring a seen desired part to the mouth [2].

5) Multi-Sensory Object Exploration: There is an obvious progression in the sophistication of the infant's object manipulation abilities, but in parallel with this there is the less obvious development in the infant's perceptual abilities. It is surmised that, in some instances, the behavioral developments help the perceptual development [104], [105]; this has been studied in the case of haptic perception. Before reaching to grasp is mastered, objects are manipulated, and information about them is gathered haptically. Infants haptically perceive object size probably during the first months of life [104]. Newborns have been shown to discriminate objects by haptic exploration (e.g., cylinder versus prism) based on shape [80]. More specifically, 2-month-olds can gather partial knowledge of shape from clues such as points, curves, or the presence or absence of a hole [80]. We saw in relation to grasping (see Section V-B3) that a visually based object representation was forming, now we see here that a haptic based representation is forming in parallel. Strategies and representations for haptic exploration in robotic systems have been developed (e.g., [106], [107]). Most of these approaches still suffer from the poor quality of available tactile sensors (compared to the tactile sensors of humans).

There is also evidence of cross-sensorial transfer in neonates [80]; i.e., infants who were habituated by haptically exploring an object were subsequently able to visually discriminate between that object and another novel object. This transfer from touch to vision has also been shown to be present at 2 months, but transfer in the other direction is absent (i.e., infants do not haptically discriminate between two objects if they have just seen one of them) [80]. Interestingly, at 5 months (when reaching and grasping are coordinated) the reverse has been shown: infants transfer from vision to touch, but no longer from touch to vision [80]. It has been surmised that the two senses may have their own representations developing at different rates, and at certain times the level of representation in each one might not facilitate transfer; the haptic system seems to be present very early and to mature slowly, whereas the visual system appears later but develops more rapidly [80].

Infants can perceive hardness or compliance by 6 months and possibly earlier; the development of such perception may be facilitated by the main action performed with objects for about the first three months, which is to grasp in a fist and open and close in a kneading pattern. Infants can perceive tactile texture[19] by 6 months, but not earlier [104], [108]; this may be because it relies on practice with exploratory rubbing movements. In normal contexts infants perceive weight at probably about 9 months, which notably comes after they have experience with waving objects, although there are exceptions [108]. In darkness, when infants are seated upright on the parent's lap, 3-month-olds can perceive weight; the darkness removes visual stimuli that may be consuming the infants' attention, and the infants' posture means that as soon as they have possession of the objects they have to support their weight (as opposed to fingering them on a table) [109]. There is also some evidence that properties of temperature, texture, and compliance may be perceived by 3 months if in the dark [109]. These advanced results in darkness highlight the difficulty of determining the competences of infants; failure to perform a task might not reflect a lack of ability, but rather it may be simply because a competing stimulus was more exciting.

The above examples show how behavior may facilitate a perception. In the other direction, infants' behaviors with objects are affected by their haptic perception. Lockman presented infants at ages 6, 8, and 10 months with hard and soft objects [110]. Infants at all ages squeezed the soft object significantly more than the hard object, but squeezing of the soft object increased significantly with age.

To model similar behaviors in artificial systems, sensors comparable to the human system are required. While powerful visual sensors are readily available nowadays the selection of available complex tactile sensors is very limited. These sensors show a significantly worse performance compared to human abilities in some or all of the following dimensions: spatial resolution, sensitivity, disparity between capabilities of hands and capabilities of sensors, long term stability, and system integration [111]. In this sense, further progress on tactile sensor development needs

---

[19]Texture here means the fine-grained property distinguishing same-shaped blocks if they were covered with, e.g., a towel, rubber, smooth cloth, or plain wood. Two objects appearing the same to the eye may have different textures, so the term is not synonymous with its use in computer vision.

to me made before similar behaviors can be replicated in artificial systems.

6) Hand-to-Surface Interactions: Hand-to-surface interactions tend to occur after manipulation of objects because the infant usually needs to be seated (with some assistance at about 6 months [104]) in order to access surfaces. Interactions between the hand and a surface can be considered to be quite similar to the hand with an object, and again the rhythmic flexion and extension of fingers (described above [73]) was also often performed to scratch a surface. Lockman specifically studied surface interactions; he presented infants at ages 6, 8, and 10 months with surfaces that were liquid, discontinuous (net), flexible (sponge), or rigid [110]. He recorded actions of slapping, picking, rubbing, and pressing; these actions may themselves be derived from stereotypical behaviors (e.g., slapping from waving) or recently acquired grasping actions (e.g., picking).

He found that infants discriminate. For example, they pressed a flexible surface more than the other three and rubbing was more prevalent across liquid; furthermore, the discrimination develops with age, becoming more pronounced. Overall, we see that once infants are grasping and acting on objects (or surfaces), they begin to discriminate the properties of those objects, and this must link to developing object representations, which are beginning to be formed. By now the infant understands something of the properties of individual objects (as a result of differentiation, M3). This means that the sensorimotor schemas (e.g., for banging and pressing) include sensory abstractions, which discriminate between different objects and surfaces (in order to predict different consequences for the action being executed on each one). This discriminating knowledge forms a substrate, which will allow the infant to progress to learning about the effects of actions involving relationships among these objects and surfaces. Recent robotics work has shown that a robot equipped with different sensors (vibrotactile sensors [112], accelerometers [113] or strain gauges and Polyvinylidene Fluoride sensors [114]), and performing exploratory movements, can learn to discriminate different surfaces; this should likewise be useful for building multi-sensorial object representations.

C. Object–Object Behaviors

This section covers the early object–object interactions that involve controlling relationships among objects. Through these interactions, knowledge about the relationships between objects is acquired.

1) Object-to-Surface Interactions: Two of Thelen's stereotypical movements [73] were performed only as an object–surface interaction. One movement consisted of an infant holding an object and rubbing it (horizontally) against the surface of a table or floor, with movement from the shoulder. The second was a push–pull movement from the elbow (flex and extend) with the arm parallel to the floor or table. This was typically done for an object too heavy to be lifted, so instead it was pushed back and over on the surface; it arises due to the interaction between an innate motor behavior (push–pull of the elbow) and the constraints of the physical world. A plausible explanation is that differentiation (M3) adjusts the innate motor

action, producing a new version, which can subsequently be used even to push light objects across a surface.

The stereotypical waving action has been performed with objects, as noted in Section V-B6; now when a hard surface is present, this can lead to the behavior of banging on the surface to produce an interesting sound. Lockman studied how this banging action becomes discriminatory (a result of sensor differentiation, M3). This means that infants begin to learn what object-surface combinations will produce the sound (i.e., learning contexts and effects of the new schema), and perform motor differentiation (M3) to increase the sound. Again, Lockman presented infants at ages 6, 8, and 10 months with surfaces that were liquid, discontinuous (net), flexible (sponge), or rigid, and also with hard and soft blocks [110]. In this case, there was an inability of the younger infants to discriminate some of the relationships. All infants banged differentially, banging more frequently with the hard block on the net and rigid surfaces, but only the 10-month-olds banged the hard block more than the soft one. Also, only the 10-month-olds differentially rubbed across surfaces, rubbing over the rigid one for longer. This again suggests that infants develop by first focusing on the exploration of individual objects and only at a later stage focusing on object relations. As a consequence, by 10 months, there is a general improvement in the ability to handle relationships between a grasped object and a surface it acts on. Possibly the limitation of younger infants is important because it prolongs the period of dealing with simple relationships, so that they can be learned thoroughly.

A further study [115, Ch. 21] tested infants, from 8 to 10 months, on a monthly basis, with hammers for banging surfaces. The hammer heads were hard, soft, or half hard/half soft. The surfaces were hard or soft. Infants at this age were able to hold a hammer by its handle, to use as a tool. It was found that all ages banged the hard hammer more than the soft hammer, and on the hard not soft surface. Furthermore, there were more hits with the hard side of the mixed hammer.

This action then is almost tool use, however, orienting a hammer to a surface is easier than directing it at a specific object. When an infant is presented with two surfaces on a table, side by side, and is able to selectively bang on one surface, then the infant shows awareness that these two surfaces are distinct. This is very close to selectively banging against another object, and forms a possible bridge to banging a held object against a stationary one.

Overall, the path we have traced shows how the original stereotypical movements could help to bootstrap the development of object-object actions primarily via the mechanism of differentiation (M3), which is itself triggered by accidental discovery following repetition (M1).

2) Taking an Additional Object: As soon as infants can grasp one object, they will inevitably face situations where they want to grab another even though they are already holding one. Bruner [15] examined the way infants respond to being handed multiple toys, one after the other. Infants from five different age groups were tested.

At 4–5 months, some infants could not even get the first toy, and some could not hold it for long. Some infants succeeded in taking the second toy, but only because they inadvertently dropped the first before taking the second. In general, infants tended not to grab the second toy if the first toy was already in the process of being taken to the mouth.

At 6–8 months, good command of the grasp was attained. On being presented with the second toy infants at this age often transferred the first toy to the other hand, to free the preferred hand for reaching. The development of this behavior came from taking the first object to the midline in order to hold it with the two hands, and then reaching out with the nearest hand; and this then evolved into an anticipatory handover. This is a composition (M5) of "handing over" and "reaching for the next object." Sometimes, instead of the transfer, the infant would reach across with the empty hand.

At 9–11 months, one fifth of the trials successfully dealt with three to four objects. The strategy employed was to put one in reserve storage (in the lap or beside the infant) to free the hand for the second, although this often (50% of the time) triggers another grasp attempt immediately, i.e., the infant forgets why he put down the first object and/or cannot inhibit the action of retrieving it again immediately (see again Section II-B on affordance-based to goal-based action). Overcoming this difficulty requires a capacity to delay the retrieval response and to maintain the intention for grasping the new object. These abilities are obviously also important for more complex problem solving requiring planning (see Section II-B).

At 12–14 months, the storage strategy was well-developed, and furthermore, the infants can place an object in storage before a third or fourth object is handed to them. It is not clear how the storage strategy develops from the handover strategy. At 15–17 months, the mean number of objects the infant can take possession of has gone from 3.0 to 3.7, and objects are stored in one way consistently. Overall, from 6–17 months, there is a gradual increase in leaving it there, rather than a sudden step change. The development shows a process of integration of the constituent acts (pull to self, place in storage, reach for new object) into a successful behavior (see modularization, M6) [15]. We can also see that knowledge of space and special locations is necessary for dealing with more than one object. In the earlier interactions with a single object the reach and grasp were triggered, and then the infant manipulated the object. Location was implicitly coded in the reach behavior, but the infant was not forced to be aware of this. However, when two objects are being handled, the infant is forced to become aware of locations apart from the location implicit in a reach. We have mentioned the beginnings of object representation before; we see here the beginnings of spatial representation, which becomes more abstract in a similar way (see Fig. 1, abstract track).

3) Pulling the Supporting Object: Willatts analyzed the task of pulling a towel to retrieve a supported toy at from 6 to 8 months [116]. He recorded not only success or failure on the task, but also monitored the infant's gaze, in order to have an objective measure that could discriminate between accidental success or intentional success. Younger infants (about 6 months) tend to give up on the toy and play with the towel instead, but in doing so they often accidentally bring the toy into reach. Willatts was able to monitor the infant's gaze and to show that there was a transition. Whereas the younger infants (6 months) gave up on looking at the toy, as they got older, there were more glances to-

wards the toy (8 months), suggesting that pulls of the towel were intentional in order to retrieve the toy. These kinds of accidental discoveries lead the child to understand the effects of various actions on object-object relationships, and lead to the development of a repertoire of "means" actions that can be employed to achieve goals [2], [20]. This is an example of the mechanism of repetition (M1) leading to accidental discovery followed by composition (M5) of the means–end behavior, which then allows intentional exploitation and differentiation (M3) of the motor action to more effectively pull the towel along the table.

Note that in the case of support, the necessary relationship (on top of) is not understood at 8 months, and up until 10 months or later the infant will still pull a support even if the desired object is held above it and not touching it [69, pp.111], or resting on an object close to the support [2, pp. 283]. The acquisition of this knowledge requires sensor differentiation (M3) in the context of the pull-support schema.

4) The String: A string is tied to an object and must be pulled in order to bring the object within reach. The string behavior is particularly easy because of its unbreakable contact with the object [117]; it is hard to go wrong (in contrast to the stick, Section V-C9), if the string is shaken wildly the object will still not be lost from the end of it, and is quite likely to be brought closer. The behavior is learned by differentiation (M3) following accidental success (i.e., as in the support, initially pulling the string without being aware that it will bring the object closer). One difference from the support is that if the string is not straight, it may require several iterations of reaching, pulling, releasing, and reaching. It also involves composition (M5) because the full behavior composes pulling the string with reaching for and taking possession of the object. Two entities are comprehended in a spatial relationship by this composition (object and string). Uzgiris and Hunt [69] tested two different string situations. The easier situation was on a horizontal surface; the more difficult string behavior was when the object must be raised vertically. The horizontal strategy fails because the object falls if the string is released; success requires bimanual control, typically with one hand pulling, and then passing control to the second hand, which prevents the object from falling, while the first hand stretches again. Uzgiris and Hunt observed the horizontal string task at 12 months, and the vertical string task at 13 months [69, pp. 111].

5) Obstacle Removal/Avoidance: This behavior is a step towards tool use because the relationship between two objects must be acknowledged (obstacle and desired object), and one must either remove the obstacle or detour around it. Learning the means–end coordination to remove an obstacle in the way of grasping is one of the first means–end behaviors described by Piaget [2, pp. 217], which he places at 7 1/2 months; as with the support, it may be learned by an accidental discovery followed by intentional exploitation. This can be learned by decomposition (M4) of the waving motion (see Section V-C1) such that only part is performed, thus taking the object out of the way. Again, it also involves composition (M5), and two objects are thereby comprehended in a spatial relationship. This is a difficult problem for infants at this age because they are not used to dealing with two objects, and so it is hard for them to execute an action on an object (obstacle) that is not the current goal; Piaget speculates that other two-object behaviors such as placing one object aside in order to take another (see Section V-C2) may derive from obstacle removal [2, pp. 217].

Bruner [15] looked at the task of retrieving a toy from behind a transparent lid. The lid could be easily lifted, but fell down if not held open. The behaviors observed in infants were very much in line with Siegler's multiple strategies in overlapping waves (see Section IV-C); infants used various strategies each of which peaked at certain ages and decreased only gradually. The youngest (6- to 8-month-olds) went directly for the toy and then engaged in banging of the (closed) lid, which may have become an end in itself; this behavior gradually decreased with age, but still appeared in some trials for the oldest infants. Infants of 9 to 11 months predominantly used two different strategies:1) raising and closing the lid, which also seemed to become an end in itself (an example of play taking over, Section II-B); and 2) raising the lid with one hand and carefully "worming" the same hand into the opening so that the hand (and arm) prevent the lid from closing. This is a differentiation (M3) of the reach schema, to keep the lid open, and during the execution both lid and object must be monitored. The fourth strategy was two-handed: the lid is opened with one or two hands, followed by a reach with one or two hands, but the lid is not held open long enough for efficient retrieval. This behavior had some presence in all groups, gradually increasing and peaking for the oldest (15 to 17 months). The final strategy involved holding the box open with one hand while the other hand retrieved the toy; this increased sharply after 12 months. Thereafter, there was no new strategy, but this strategy became less effortful and quicker. Progress to the final strategies is likely not a result of accidentally happening to use two hands, and having success; rather, it probably depends on general advances in bimanual control (see Section V-C7), and so this is an example of where progress to a node (i.e., behavior) in Fig. 1 (concrete track) may have to wait for all its necessary precursors to be ready.

6) Rotate a Lever: This task involves a 42-inch lever that can rotate about its center on a table. One side of the bar is within reach of the child, but the far side is inaccessible. An attractive toy is tied to the far side. The child must rotate the bar in order to bring the toy around to the reachable area. This is difficult because the child must take the unusual action of pushing the bar away in order to bring the toy closer.

Koslowski and Bruner [118] tested children of three age groups (12–14, 14–16, and 16–24 months) and categorized the strategies they used. Strategy 1 was categorized as *Linear*: reaching directly for the toy, trying to push the bar in a straight line towards or away, pulling the table. Strategy 2 was referred to as *Oscillation*: pushing the bar back and forth, but never rotating more than 45° from the midline, and tending to return it to midline after rotation. Here rotation is differentiated (M3) from pulling by variation and selection (M2); this (motor) differentiation is an example of "shaping" by the environment (modifying the dynamics of the situation [119]); the child is presented with what looks like a free stick, and the child pulls it as though it were, but because its motion is constrained by the fulcrum it only moves in one way, and the child gradually discovers this motion. Strategy 3 is *Partial Rotation*: rotating

45° and then stopping to consider the new position, but not making a concerted effort to reach for the toy. Considering the states is likely an example of sensor differentiation (M3) for the new rotation schema. In strategy 4, the children are *Absorbed* in the rotation activity, often rotating the toy within reach, but ignoring the toy [see comments on affordance-based play in Section II-B, which is the mechanism of repetition (M1)]. Strategy 5 was referred to as *Rotate and Capture*.

There was a progression towards more advanced strategies with age. Younger children found it difficult to suppress the linear strategy, and this explains the oscillation strategy. After rotating a bit, they resort to pulling the bar straight towards them (hence returning it to midline). Repeated failure with this almost forces the child to consider unidirectional rotation. Thereafter the child can pay attention to two aspects of the apparatus: either the relation between movement of the bar and the position of the goal, or the way in which the movement of the bar can be effected. The authors suspect that both cannot be attended to simultaneously due to information processing capacity limits (see Section II-C); therefore they must be first modularized (M6). While focusing on looking at the toy, little progress is made in unidirectional rotation, on the other hand focusing on the rotation leads to strategy 4. Eventually the fact that the goal is within reach is noted; this sort of accidental discovery bears some similarity to the discovery of the support (see Section V-C3). The change in strategies used was inline with Siegler's overlapping waves theory (see Section IV-C); there was a marked increase in the use of strategy 3 by the 14- to 16-month-olds, and the 16- to 24-month-olds had the largest number of children using strategies 4 and 5, but older strategies had not died away completely. This is also inline with the idea of schemas being the unit of behavior: the novice child has a well-developed schema for pulling in a straight line, but is only developing the schema for rotation; the child must ignore the goal in order to focus on developing the rotation schema further, so that it can be later used as a means action (see Section II-B). The child's modular approach to the problem has a major benefit. "Not only is the problem solved, but it is solved for a wide variety of circumstances and forms in which it is likely to be encountered, wherever the lever may point, whatever its shape, and so forth. Transfer, so to speak, is built into the solution" [118].

7) Advancing Bimanual Control and Object Manipulation: Bimanual control is required in many tool use scenarios, and it is reckoned to be an important component in explaining why human tool use capabilities exceed those of other animals [115, Ch. 24]. Behaviors such as holding an object in one hand, and striking or stroking it with the other are the beginnings of "role differentiated bimanual activity," and appear as early as 7 months [120]. Some toys are more likely to elicit bimanual activities than others at particular ages, but overall the frequency of bimanual activities increases linearly with age [120]. Infants of 7 months were as likely to execute bimanual activities on toys with no movable parts, as those with moving parts, but from 9 months onwards, toys with no movable parts elicited few responses, and infants seemed more interested in toys with parts to be independently manipulated [120]. It is surmised that these developments require a combination of neural developments (i.e., maturation), as well as having the appropriate

objects, and also understanding the properties of those objects (which can arise through sensor differentiation (M3) [120]. More sophisticated bimanual actions appear towards the end of the first year. For example "contour-following" can be observed at 12 months, which involves holding the toy in one hand and maneuvering it, while the fingertips of the other hand are moved smoothly over its edges [121], or employing a single finger or pincer action for manipulation with the second hand at 11 or 13 months [120]. Bimanual control has been implemented in the iCub using the Passive Motion Paradigm, to perform two-handed grasping and moving of objects [122]; role differentiated bimanual activity has been shown recently for opening of a screw-top jar [123].

8) Relational Play: Apart from intentional problem solving, there is also a natural progression towards object-object interactions in infants' free play. An overview of infant free play is shown in Fig. 2 (just below the grasping track). When infants of various ages were presented with a wide variety of toys, three categories of play were observed [124]: 1) stereotypical play was dealing with a single object (mouthing, fingering, waving, banging) and dominated at 9 1/2 months; 2) relational play dealt with associations of two or more objects and dominated at 13 1/2 months; and 3) functional play was using a toy in a manner deemed appropriate to an adult, such as using a comb to comb a doll's hair; it dominated at 15 1/2 months. This study shows that by 13 1/2 months most infants prefer to explore relationships among objects, rather than exploring objects individually. A further study [125] gave more insight into the precursors to full relational play; at 7 months, the very simple relational action of banging two objects together was common, and by 9 months infants could do very simple relational acts such as touching a spoon to the base of a pot; it was between 9 and 13 months that most infants made the transition from these simple relational acts to "accommodative" relational acts such as putting a lid on a pot or a spoon in a cup. At around 10–11 months, infants begin to establish the links between particular objects and their "canonical actions," e.g., a hammer is for banging, a brush is for sweeping, and also the spatial relationships that must be established between tool and target object; e.g., the relations "in" (key or screwdriver in a slot), on (one block on another), and under (put a spatula under a pancake) [108].

These developments are what some might call "stage change" because there is a qualitative difference in the behavior. We are not aware of any computational model that can capture this. It seems to require advancement in the infant's knowledge of object relationships and spatial relationships. For example, beginning with a very practical knowledge of a relationship between two concrete objects, discovered during play, the infant generalizes to other similar objects, and thus it becomes a more abstract knowledge (increasing abstraction is shown by moving to the right on Fig. 1, abstract track), which may involve processes of representational redescription (M7).

9) The Stick: The behavior of the stick entails using a stick to move an out-of-reach object and bring it within reach of the hand. There is in fact a spectrum of behaviors under the broad umbrella of "the stick." The simpler end of the spectrum consists of using a short stick to retrieve a barely-out-of-reach object with a single sweep of the arm; the arm is initially extended

so the stick reaches beyond the object, and then it is brought towards the body (which may happen towards the end of the first year). The more complex end of the spectrum includes behaviors using a long object (e.g., a long stick or mop) to knock an object from side to side until it can be reached (which may be placed at about 3 years [126]). Uzgiris and Hunt tested a medium-length stick (18 inches long) [69, pp. 150]. They place the behavior at 15–18 months [69, pp. 111]. It is a relatively difficult behavior when compared to strings and supports because the tool is not given in the appropriate relationship from the outset; the infant must create the appropriate relationship. This is a complex example of differentiation (M3) and composition (M5). Its relatively late appearance suggests that it requires certain precursors to be present, for example, an understanding of spatial relationships to understand the necessity of getting behind the object without knocking it further away.

Brown [126] tested children ranging from 17 to 36 months old for transfer ability in retrieval tasks with a variety of stick-like tools, some of which had a hook or rake at the head. Tools varied in length, rigidity, color pattern, and type of head. The children never selected nonrigid tools. Overall, the children seemed to understand quite well the properties required of an effective tool for the task. Brown makes a strong case for the ability to transfer being very much domain specific, and related to the child's understanding of causality in the particular task. The stick behavior has been tackled in robotics by considering the tool as an extension of the robot's body schema (see survey [127]). Observing video of infants (e.g., at 11 months) struggling to use a stick suggests that they do not solve the problem in this way; in particular, it is notable how they struggle to control the stick, and so cannot move its end as they move their hand. The body-schema approach does not capture the understanding of causality that infants seem to achieve, i.e., the understanding that an independent object (stick) can effect motion in another; the body-schema approach would not extend to understanding (and correctly predicting) how an externally controlled stick would cause an object to move.

10) Perceptual Aspects in Retrieval Tasks: Bates *et al.* [117] looked at perceptual aspects in retrieval tasks for 10-month-olds using support, string, stick, and also a hoop and crook (stick with semicircular hooked end). It was found that if the tool and desired object both had the same color and texture, then it was particularly difficult for the infant to succeed. It was surmised that the perceptual difference may help the infant to discriminate the two objects and to keep them both in mind as separate entities. A difference in both color and texture was no more helpful than a difference in one or the other.

The effect of the spatial configuration of the objects as presented to the infant was also investigated. Four types of spatial configuration were presented: 1) unbreakable contact (support and string); 2) breakable contact (hoop or crook, presented in contact so that the tool only needs to be pulled); 3) behind (hoop, or crook, surrounding object, but behind and not in contact so that the tool only needs to be pulled); and 4) beside (hoop or crook or stick, presented beside each other, so that the tool needs to be brought into contact before it is pulled). It was found that difficulty increased as follows: unbreakable contact, breakable contact, behind, and beside. The four tasks in the "breakable

contact" and "behind" groups all required the same motor action (pull the tool), yet there was a significant difference in success on them with the hoop in contact being significantly easier than the crook behind. This suggests that the infant understands the causal relation when two objects are connected and the physical contact may help the infant to remember this. It is not likely that the infant conceives of the connected objects as a single entity, because perceptual similarity of the objects is a hindrance.

11) Fitting Shapes Into Slots (Peg-in-Hole Task): The task of inserting a cylindrical peg in a cylindrical slot can be done by almost 50% of infants at 12 months, but they do not pre-orient the cylinder for insertion [128]. Instead, they press one end to the hole, and then move the other end until they find the right orientation. By 16 months, infants do pre-orient the cylinder, but not other shapes, until later (see details on context specific skills, Sections IV-C and IV-D). The 12-month-olds seem to use the reduction in degrees of freedom (DOF) strategy described in Section II-A. They first hold the peg in a fixed orientation in the hand, and move the far end of it into contact with the hole. This is a three DOF problem (which is very similar to controlling the hand in a reach towards an object). The second step is to orient the peg to be parallel to the slot, while pressing it into the hole so that the end in contact with the hole maintains contact. This is a two DOF problem. The cylinder itself has five degrees of freedom, but the sequential approach reduces the problem space. In each of the two steps, younger infants try a large amount of variation (M2) before getting the objects in the correct relationship. Repeated practice helps the infant learn the correct orientation, and so the infant tends to approach the hole with a gradually better pre-orientation in successive trials. Some robotics approaches do insert a peg in a hole in a similar two-step approach, using only a force feedback sensor [129]. Reinforcement learning has also been applied to learn peg-in-hole insertion using reactive control from force and position feedback; this work showed a gradual improvement in insertion skill somewhat similar to infants [130]. This approach was also able to relearn its strategy when given a square peg. This is similar to the robustness seen in infants; if difficulties arise (e.g., due to a strange-shaped object), infants fall back on earlier groping behaviors.

The cylinder is relatively easy as it can be inserted with its cross section in any orientation. Shapes of noncircular cross-section (e.g., triangular or rectangular) must additionally be oriented so their cross-section matches the opening. Children are remarkably bad at this task until about 26 months [128]. This seems to reveal something about the object representations they are using (see Section VI). In addition, insertion of disks in slots shows that 18-month-olds fail to preorient, even though they can well pre-orient their hand for insertion in a slot [131], showing context specificity of representations.

12) Objects With Handles: McCarty *et al.* [132] studied how infants deal with an object with a handle, and in particular, what way they grasp it. The infant experiments were done with a spoon preloaded with food and toys with a handle (bell, rattle, cow, and pig). Each object has a handle and a *goal-end* (e.g., the goal-end is the bowl of the spoon or the toy). Three different grasps were categorized as shown in Fig. 4.
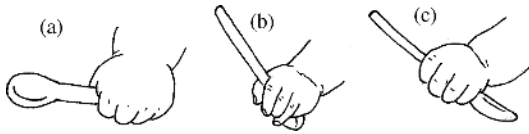
Fig. 4. (a) Grasp the handle with the thumb towards the goal-end of the object (called *radial* grip). (b) Grasp the goal-end itself (called *goal-end* grip). (c) Grasp the handle with thumb towards the non-goal-end (called *ulnar* grip).

This study looked at how infants develop the ability to use the radial grip [see Fig. 4(a)] at the start of the task. Infants develop this ability only gradually. We are not aware of any robotic work that models this development; most robotic work tends to give the robot the knowledge to grasp correctly from the outset [133].

In McCarty *et al.*'s study the preferred hand of the infants was identified in a pre-test. In the results 67% (for toys) to 70% (for spoons) of grasps used the preferred hand. The objects were presented to 9-14- and 19-month-old infants, on a stand, with the handle alternately oriented to the left and right; trials could then be categorized as *easy* if the object was presented in an orientation that would allow an overhand grasp with the preferred hand to achieve a radial grip (otherwise it was *difficult*; i.e., an overhand grasp with the preferred hand would achieve an ulnar grip). The most interesting results then concern how the infant dealt with the difficult case: 9-month-olds tended to use any of the three grasps indiscriminately; when they used a nonradial grip, they often (more than half of trials) put the handle in their mouth, and typically corrected afterwards. Fourteen-month-olds were less likely to grasp the goal-end and more likely to use an ulnar grip; however, when they used a nonradial grip they never put the handle in the mouth; instead they corrected either by rotating the wrist awkwardly, or changing to the other hand. Nineteen-month-olds used the radial grip on 86% of difficult trials, which meant that they had to suppress the tendency to use their preferred hand and use the other instead (see also [134] for the training experience that accelerates this).

Following from these results, the authors formulated a model of the development of planning in this task: 1) *Feedback-based strategy*–after an indiscriminate grasp the end of the spoon near the thumb is brought to the mouth first, and if this turns out to be the wrong end, then a correction is made and the other end put in the mouth; 2) *Partially planned strategy*–as soon as the spoon has been grasped its orientation is noted, and if it is incorrect an adjustment is made before bringing it to the mouth (this entails inhibiting the preference to bring it straight to the mouth); and 3) *Fully planned strategy*–the orientation of the spoon is noted before grasping and a grasp, which is appropriate to the goal of feeding, is selected. This model predicts that infants' actions should be slower when planning is taking place; some evidence for this was found in that the action of bringing the spoon to the mouth, when the ulnar grip had been used, was slower in 14-month-olds than in 9-month-olds [132]. Overall, we see a striking lack of planning at the earlier ages, which is inline with the idea that behavior is more affordance-based before the second year (see Section II-B); i.e., an affordance-to-grasp suggests itself and is immediately acted on without regard for later steps.

Further work has shown that the radial grasp generalizes to other tools with self-directed goals (e.g., hairbrush on self), but not to other-directed goals (spoon to feed a toy lion, hairbrush to brush toy, hammer to object) [135]. This reinforces the ideas about the context specificity of knowledge (see Sections IV-C and IV-D). From a computational point of view, it suggests that the infant is not applying the same generic planning framework for both tasks. Instead, it is likely that composition (M5) has been applied to produce a composite schema for grasping and taking the tool to the head (in the self-directed tasks). No such schema exists for other-directed tasks. We suspect that developments on the abstract track (see Fig. 1) are necessary before the common deep structure in such tasks is obvious to the child (see Section VI).

D. Tool Use Example: Transport Using a Spoon and Bowl

Many of the behaviors described so far have been building up the necessary knowledge for tool use by understanding the properties of individual objects (see Section V-B), and the effects of various actions on various objects in relationships (see Section V-C). Self-feeding from a bowl using a spoon is "proper" tool use and is common in human cultures. Connolly and Dalgleish [136] studied two groups of infants longitudinally, at monthly intervals; one from 12–16 months, and the other from 18–23 months. They outlined four stages in the development of this behavior: stage 1) repeating one part of the feeding sequence, such as putting the spoon into and out of the bowl, or into and out of the mouth; stage 2) performing the outline of the correct action sequence spoon-to-dish-to-mouth, but not effectively loading food on the spoon or unloading in the mouth; stage 3) effective performance of loading and unloading within the sequence; stage 4) incorporation of correction routines (e.g., check if food has been successfully loaded, if not, return to the bowl or pick up food that has dropped during the transfer to mouth).

The behavior of stage 1 can be called *play*, where the goal of feeding was not pursued, and the means is done for its own sake (mechanism of repetition, M1); in addition, the infant would sometimes pass the spoon from hand to hand, bang it in the dish, or on the table, or drop it to the floor, or rub it against his/her own head (mechanism of variation and selection, M2). Though these activities were not directly in the service of feeding, they did serve to increase the infant's knowledge of these actions, and their effects in the feeding context (this is the role of play, as described in Section II-B). Sometimes goal-directed behavior was observed, but there was a lack of understanding of the purpose of the spoon: the younger children were sometimes observed putting their spoon into and out of the dish repeatedly, while taking food from the dish with their other hand. In the behavior of stage 2, the younger infants did not seem to understand the need to load the spoon. The behavior is learned by imitation (see Section II-D), so they have some knowledge of the sequence of the operations before understanding their individual purposes. To effectively learn a component part, stage 3 is an example of means–end behavior, which tends to follow the pattern of accidental success leading to acquisition of the appropriate schema (where M1 leads to M3), followed by intentional repetition with variation and refinement (M1 and M2), and later understanding

(see Section II-B3). The correction of errors in the sequence occurs first for those elements at the end of the sequence, and latest for those at the beginning of the sequence; e.g., by 18 months the remaining errors are only in the earlier stages [137]; this is probably due to the younger infants' limits in general planning abilities as noted in Section II-C.

In terms of the component skills, the authors outlined four principal problems for the infant: controlling the spoon in the hand, loading food on the spoon, taking it to the mouth without losing the food, and unloading in the mouth. A number of "behavior categories" were devised to code the observations of the infants in various activities such as grasp employed, trajectory to mouth, loading method, etc. Loading of the spoon, for example, included 1) dipping-in motion (spoon lifted and lowered, sometimes repeatedly); 2) side-to-side motion across the dish; 3) scooping motion towards the infant; and 4) dropping the spoon in the bowl and sometimes picking it up. Overall, the results showed that younger infants had more varieties of hand grasps, and less stable movement strategies; with age came increasing consistency in the actions used (mechanism of variation and selection, M2). Also the behavior categories used changed; e.g., for loading the spoon, younger infants preferred dipping-in, while older ones preferred scooping with a wrist rotation, or the side-to-side motion (which was often effective in trapping food against the side of the bowl) (see Siegler's changing strategies discussed in Section IV-C). The overall pattern of movements became smoother and more direct, and the time needed to perform individual components of the action decreased (see modularization, M6). In terms of hand grasps, older infants used fewer inappropriate grasps (e.g., ulnar grasps are inappropriate because the arm gets in the way when trying to bring the food to the mouth), and, furthermore, older infants used more flexible grasps; *flexible* here means that the spoon can be manipulated with finger movements, as opposed to a rigid grip, which only permits wrist movement (these general features are also seen in the progression from novice to expert in adults [115, Ch. 4, 5]).

We also note a strong similarity with the task of learning to drink from a cup [15, pp. 72]. Initially the child grabs the cup and pulls it to the mouth in a single step. With practice, the child slows this down and puts in a number of stopping points to rebalance the cup so the liquid does not spill, and also adjusts the head position, bringing head to cup, and monitoring how the cup is moving towards mouth (see decomposition, M4). With more practice, it becomes a smooth motion where monitoring of the level is done continuously during the motion (see modularization, M6).

In this behavior, we see how the mechanisms of schema development need to work together over a relatively long time to eventually produce efficient spoon-feeding skill. We should also point out that this task is relatively simple because it does not require mental representation of unseen parts, which poses more severe difficulty for infants (see, e.g., [138]). There does not yet exist a computational model of this type of tool use development. There is currently a large gap between what robots are capable of by developmental and nondevelopmental approaches. For example, a nondevelopmental approach has been used to produce a pancake-making robot, using a spatula and pan [133]

(a task which is well beyond young children), whereas one of the more sophisticated developmental toolusers [139] merely learns which of a set of given tools is useful for moving a hockey puck towards the robot on a smooth table, and this when the tool is already placed on the table. Of course, the developmental work has advantages; e.g., it can autonomously rediscover how to use a tool if it becomes partly broken, whereas the pancake-maker lacks a mechanism for this. We believe that developmental robotics is broadly heading in the right direction, by focusing on mechanisms of development in simple tasks. We should not be surprised that there do not yet exist developmental robots that can acquire spoon use (a difficult task).

Having completed our description of the infant behavioral studies, we can now reflect on how developmental robots compare with infants (and what is lacking). Most of the examples of infant object–object behaviors described here have not been attempted in robots in a developmental way, there is more developmental robotics work on earlier behaviors such as stereotypical movement (see Section V-A), or grasping (see Section V-B1); it is important that future work links these up to later more advanced behaviors. At present, most computational models of the development of advanced behaviors implement these as isolated episodes of learning, where the starting point is largely handcrafted (see also models beyond infancy [11]). Developmental robotics lacks examples of longitudinal developments, autonomously building on each other, as infants show (also called "ongoing emergence" [140]). A notable exception is the modeling of motor developments in early infancy by shaping [141]; what remains to be tackled is the application of similar shaping techniques to the development of object–object behaviors. To implement this in robotics would require that each learning episode builds object–object knowledge, which forms a basis from which the next learning episode can begin (without additional input from a programmer at that point). This is closely linked to the issue of representational developments (discussed next) and has not figured largely in developmental robotics thus far.

VI. INTERNAL REPRESENTATIONS (ABSTRACT TRACK)

This section briefly looks at changes in internal representations (upper track, Fig. 1), using the observable ability to transfer as a way to deduce what representations may be in use. Transfer of specific skills to similar related scenarios or objects is very important for robust tool use. The evidence from Section V suggests that improvements in this ability during development can be explained by increases in the world knowledge within the system, rather than some generic developing "transfer ability." We have seen from Brown's study on retrieval [126] (see Section V-C9) that children transfer very well when they understand the causal relationships in the particular task; Brown has also shown that they do not transfer on more abstract tasks where the relationships are not understood according to any of their prior knowledge and therefore seem arbitrary to them. This message is reinforced by a further study of 3- to 5-year-olds [142], which points out that the ability to transfer is not directly dependent on age, instead it is dependent on the level of representation achieved; young children can achieve a deep representation of causal relationships in tasks

involving simple physical manipulations, and therefore can easily transfer in these. Older children can achieve a deeper representation in a wider variety of domains, and hence can show transfer in more domains. Therefore, the observable ability to transfer could serve as a proxy for deducing something about the unobservable internal representations. Using this, we could say that representations seem to develop in (at least) the following three ways.

A. Coarse to Fine

In some situations, infants generalize very well and immediately (for example, supports or sticks [2, Obs. 152, 160]). The fact that these generalizations can happen immediately after the skill is first learned suggests that the objects were already represented in the same way (e.g., a coarse representation of a long object); once the skill is learned for one, it can generalize to all. Sometimes infants over-generalize, e.g., *scale errors* [143], [144], where behavior is generalized to objects of incompatible sizes (such as a too large tool in a container), or the attempt to insert incompatible shapes in holes; this again suggests a coarse representation, which might ignore some details of shape and scale, but which is strongly linked to functional use. The development of representations seems to follow a path from coarse to fine, with initial representations capturing rough shapes, and the detail on objects only being gradually elaborated later.

B. Context Specific to General

In some situations, infants do not generalize well at all, for example, in the way a spoon is grasped for self-feeding, or for directing to another object (see Section V-C12), or placing the hand in a slot versus posting a disk in a slot (see Section V-C11). Much of an infant's learning is quite task-specific. Examples of lack of generalization suggest that high-level representation is not that well-developed (i.e., the high-level similarity between tasks is not apparent to the infant) and suggest that it is important to spend an extended period focusing on task-specific learning. This then needs to be followed by some process of representational redescription (M7), which can find a higher level abstraction common to a number of concrete behaviors. This higher level may, for example, capture causal understanding of the behavior, and when it is achieved generalization in other domains becomes possible, and understanding of demonstrated actions becomes possible as well.

C. Integration of Fragmentary Representations

Kellman and Arterberry explain that "perception leads to multiple representations that may be recruited for different tasks" [19, pp. 262]. Part of the work of development is to connect these up to produce more generic and reliable world models. We have seen examples of this already in the connection between haptic object representation and visual object representation (see Section V-B5). Additionally, Kaufman *et al.* [145] describe how the two separate visual processing streams in the infant brain (dorsal or ventral) are responsible for different tasks. The dorsal route seems to be primarily used for knowledge related to grasping (a practical representation), while the ventral is for representation and recognition of the whole object; yet these must be integrated to allow grasp

knowledge to be associated with an object representation. It may be at quite a late age (maybe 9 months [145]) that infants can integrate the information from the two streams. Both before and after this there is further evidence of integrating fragments. Surprisingly advanced perceptual competences are shown by 4-month-olds in perceiving the 3D form of rotating wireframes [19, pp. 168], yet this seems to constitute only a fragmentary understanding of objects because they do not "complete" solid 3D objects until 6 months [146]. Even at 18 months, fragmentary representations based on view dependent images and parts of objects seem to be still in use, and then there is a period of rapid change where 3D whole-object geometric representations are built by 24 months [147]. The picture emerging from the literature suggests that object representations may undergo a long and complex developmental trajectory. At the same time, we can see advantages of fragmentary task specific representations. They provide a simple space that is appropriate to a particular task, and when another seems more appropriate it is possible to switch representation (see also [148] on the need for multiple representations).

VII. REFLECTION AND RECOMMENDATION

In this section, we first reflect on the psychological results to summarize the salient points about how the overall development works (see Section VII-A); we then formulate some succinct guidelines for developmental roboticists who wish to model similar developmental trajectories (see Section VII-B).

A. Reflection on Infant Development

In reflecting on the examples above we can see the two tracks of sensorimotor skill and representation developing (Fig. 1). From this we extract the following main ideas:

1) Innate Knowledge Is Fragmentary and Incomplete: Innate knowledge of the physical world seems to be given in a fragmentary form; it is not given from the outset in the useful form that an adult has, but rather the evolutionary endowment seems to provide constraints and boosts for the development of world knowledge at various times. It is given in a form that presumes a prolonged development process in concert with the environment. This can work in complex ways where the innate fragments may be creating opportunities for the necessary environmental interactions (see Section V-B1) or providing fragmentary representations to bootstrap the development of knowledge of objects (see Section VI). The fact that physical knowledge is not given in a "final" form might be important to ensure that the knowledge eventually developed is linked to sensorimotor experiences of the infant, and hence more practically useful.

2) Infants Learn Slowly, But Thoroughly: Infants spend months practicing individual actions in varying circumstances, and gathering good knowledge about how to apply an action, and its expected effects.[20] We see this in the way that the period dominated by affordance-based play must precede goal directed planning (see Section II-B and also poor planning in

---

[20]This is compatible with the principle of "developmental gradualness" [104], which describes "particular skills and abilities appearing initially in rudimentary forms and in highly specific contexts, and then gradually becoming more complex and wide-ranging over time."

Section V-C12), and furthermore, play may sometimes need to resurface when perception-action knowledge is inadequate (see Section V-C6). Expertise and flexibility on a task come from extensive practice with the elementary actions comprising the task. During this time (in addition to the environmental circumstances varying), small variations are tried out, and the effects of those variations are learned. Behaviors learned in this slow manner are well grounded. This highlights the importance of achieving robustness and variety for controlling elementary skills, as these will come into play later when these skills may be constituents of more complex behavior.

This slowness explains why we often see some (fortuitous) success in a particular behavior before a fuller understanding is achieved some time later (e.g., the support, see Section V-C3); many complex skills are learned in a crude outline before the constituent parts are properly refined (e.g., the spoon, see Section V-D). The early generation of experiences through this approach provides the training data that improve the behavior.

There is a link between this general slowness and the acquisition of physical world knowledge above (see Section VII-A1). Piaget said "to understand is to invent" and so it makes sense for the genetic "preprogramming" to only provide a fragmentary outline, which guides the development of the knowledge; to achieve a thorough understanding of the physical world knowledge of it is necessary for the individual to gather significant experience with the component fragments from which they can then themselves build the necessary concepts (e.g., the building of knowledge of objects, see Section VI). When the general representations of objects and space are built in this way they are more useful because they are so closely connected to the actions that can manipulate them.

This process appears to be facilitated by a "schedule" for development that forces more time to be spent on earlier tasks; e.g., the fact that the pincer grasp arrives relatively late (see Section V-B3) means that significant time prior to this is spent on coarser grasps, where a coarser representation of objects is adequate; (see Section V-B5 for the way in which perception and action may help to bootstrap each other's development); furthermore, in language, acquisition of vocabulary proceeds very rapidly once it starts, but it does not begin until significant interaction with objects is complete.

3) Generalization Depends on Representation: The ability to generalize and transfer to new situations is dependent on the underlying representations in use, and sometimes infants are surprisingly poor at this and seem to have knowledge that is locked in context. We have seen in Section VI that in some cases some of the early representations facilitate certain types of transfer (e.g., the stick), but in other cases the ability to transfer appears relatively late because it takes a long time for new appropriate representations to develop (e.g., handled objects, see Section V-C12). The processes underlying this development are hinted at in Section VI, such as representational redescription (M7), see also Section VII-B2), but we know very little about how these processes work. They seem to be slow processes that come into play after extensive experience with more primitive context specific representations (so there is a link between this and the previous points).

Nevertheless, a lot of tool using behavior can happen without the need for advanced representations of objects, which are independent of specific tasks. Task specific learning seems to account for most observations quite well. Popular perceptions of the intellectual abilities underlying tool use sometimes overemphasize the notion of "sudden insight" and anecdotes of dramatic inventions may often turn out to have simpler explanations on closer inspection; i.e., they may be minor generalizations from very similar behavior that was practiced extensively [115, pp. 308], [28].

We conclude this reflection by asking: What are infants good at and what are they bad at? They seem to be good at building on what they know. Once they have acquired a skill, even crudely, they will try it out in varied situations, and refine it and improve it and specialize it for new situations that did not produce quite what they expected (leading to robustness and generalization). They are good at assimilating new results and relating them to what they already know (provided there is some relation). They seem to be bad at making big leaps to new tasks that do not build on what they already know; there are tasks that are beyond them at certain ages, and it can take several months for them to acquire the necessary precursors before they can attempt them. They are, however, good at innovation; when presented with a task that is beyond them they will try a large range of strategies, and even if they do not succeed, they may discover something new through play.

B. Direction for Roboticists

This section offers advice for those who want to make tool-using robots that have the kind of robustness and generalization that children have (i.e., able to cope with changes to tools and materials, and to find appropriate ways to do a job without explicit detailed instructions).

Despite our incomplete knowledge of how biological systems achieve tool use, we can outline how artificial systems might be constructed to tackle the problem in a similar way. Starting with a small set of innate sensorimotor schemas (see Section V-A), a bootstrapping process can be initiated by which new sensorimotor schemas develop through the interaction of innate schemas with objects in the world (see Sections V-B and V-C) by means of the six mechanisms M1–M6. In that process, the preconditions and effects of the schemas are refined and become more and more predictive. Eventually they can be utilized by a planning machinery (which is to a large degree innate) for the purpose tool use.

We believe that, when designing developing artificial cognitive systems, for some aspects, it is acceptable to take artificial shortcuts, but for others one should be more careful to closely follow the biological approach. For both planning and social aspects it would seem acceptable to take advantage of the possibilities artificial systems offer; i.e., a planning system can be made available, and social demonstrations can be made directly available (through human-provided motions, for example) without the need to observe or interact with a social partner. However, both the schemas and the representations of the world, must develop slowly and autonomously, and *this should not be shortcut by direct coding*. To emulate this development, it is valuable

for roboticists to attempt to emulate the tasks that infants really can do; this avoids making robots do overly sophisticated things (which might lead the roboticist to use mechanisms that are inflexible and not generalizable). For example, if starting with behaviors achieved only at two years of age, one might need to code advanced representations, and thereby miss out on coding the processes that build those representations (missing out on one of the core mechanisms of development).

To emulate the two tracks of development, as seen in Fig. 1, the following is suggested.

1) Start With Few Schemas, To Get A Lot: We have seen that a small number of sensorimotor schemas, when applied in the world, can lead naturally, by means of the mechanisms M1–M6, to a wide variety of schemas. The smallness of the initial set may be important to simplify the state-space exploration in early development, and the gradual process of additions may be important to allow them to be well grounded. By "well grounded" we mean that they must be refined through extensive practice in varied situations. For roboticists, this requires us to build systems that can generate large amounts of varying and meaningful experience and the patience to let the robot "play" for a long time.

2) Representations Must Develop Gradually: The cognitive architecture must allow representations to develop (see Section VI), by processes such as representational redescription (M7), in order to facilitate generalization and transfer. The system may use unsophisticated representations in the early stages of development (e.g., simple internal reproduction of perceptions with little abstraction). There must then be an ongoing process of upgrading the representations in use so as to capture more generic and abstract world knowledge. This is likely to require some scaffolding in the form of certain innate representational fragments that help the system generate more sophisticated representations, as in the human case (see Section VI). This must be a gradual process; if overly advanced representations are designed at an early stage, then there is a danger that they will be inflexible and nonextensible. For this reason, we should not expect the early system to perform advanced tasks; it must spend a long time on simple tasks.

3) Interaction Between the Concrete and Abstract Tracks: A particular challenge is to establish mechanisms such as representational redescription (M7) that allow development on the abstract track, while also synchronizing with the concrete track. This requires an ongoing modification and refinement of internal representations through the experience provided by the sensorimotor schemas and the adaptation of these schema to the restructured internal knowledge representation. This is a very complex task since it is very difficult to observe the change of internal representations. Establishing such processes in developing robot systems can actually help to understand this (for two examples and a more detailed discussion, see [13]). To model development on (and interaction between) the two tracks, it is clear that both symbolic and sub-symbolic representations are necessary; a key open question is what techniques should be used to bridge the gap between these two; we see various proposals in recent work [37], [58], [149].

4) Guiding Examples and Benchmarks for Development: We provided a general outline of the development of sensorimotor schemas of infants (see Fig. 2) as well as a number of concrete stages of development in solving certain tasks. In particular, we have devoted considerable attention to object-object behaviors (see Section V-C), which comprise a major portion of infant behaviors, and are clear precursors to tool use; we have identified these object-object behaviors as being insufficiently explored in developmental robotics so far. The general outline given here might serve as a guide for the overall developmental process to be realized, and the concrete examples can serve as benchmarks for truly cognitive behavior in artificial agents.

Reflecting on the development of tool use in infants as outlined in this paper we have noted the crucial importance of developments in perception and action capabilities, and the seamless progression between this and the beginnings of tool use; this forces us to be keenly aware of the conceptual and technical hurdles still to be addressed in achieving the same in artificial systems. Nevertheless, we believe that it will eventually be possible to design artificial systems that develop advanced and stable tool use capabilities by equipping them with: 1) a small initial set of sensorimotor schemas; 2) a suitable architecture in which the mechanisms M1–M7 operate; and 3) large amounts of experiences generated by applying the sensorimotor schemas to objects in the world.

REFERENCES

[1] F. Guerin, "Learning like baby: A survey of AI approaches," *Knowl. Eng. Rev.*, vol. 26, no. 2, pp. 209–236, 2011.

[2] J. Piaget, *The Origins of Intelligence in Children* (in French:1936) Transl.:1952. London, U.K.: Routledge & Kegan Paul.

[3] L. B. Smith, "Dynamic systems, sensori-motor processes and the origins of stability and flexibility," in *Toward a Unified Theory of Development: Connectionism and Dynamic Systems Theories Re-Considered*, J. Spencer, M. Thomas, and J. McClell, Eds. Oxford, UK: Oxford. Univ. Press, 2009.

[4] K. W. Fischer, "A theory of cognitive development: The control and construction of hierarchies of skills," *Psychol. Rev.*, vol. 87, no. 6, pp. 477–531, 1980.

[5] J. J. Lockman, "A perception-action perspective on tool use development," *Child Develop.*, vol. 71, no. 1, pp. 137–144, 2000.

[6] "Preface," in *Action and Thought: From Sensorimotor Schemes to Symbolic Operations*, G. Forman, Ed. New York, NY, USA: Academic, 1982.

[7] R. A. Brooks, "Intelligence without representation," *Artif. Intell.*, vol. 47, pp. 139–159, 1991.

[8] M. Johnson, *The BODY in the MIND: The Bodily Basis of Meaning, Imagination, and Reason*. Chicago, IL, USA: Univ. Chicago Press, 1987.

[9] A. Stoytchev, "Some basic principles of developmental robotics," *IEEE Trans. Auton. Mental Develop.*, vol. 1, no. 2, pp. 122–130, 2009.

[10] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379–423, 1948.

[11] D. Mareschal, "Computational perspectives on cognitive development," *Wiley Interdisc. Rev. Cog. Sci.*, vol. 1, no. 5, pp. 696–708, 2010.

[12] K. W. Fischer and R. W. Hencke, "Infants' construction of actions in context: Piaget's contribution to research on early development," *Psychol. Sci.*, vol. 7, no. 4, pp. 204–210, 1996.

[13] N. Krüger, M. Popovic, L. Bodenhagen, D. Kraft, and F. Guerin, "Grasp learning by means of developing sensorimotor schemas and generic world knowledge," in *Proc. AISB Convention, Comput. Models Cognitive Develop.*, 2011, pp. 23–31.

[14] A. Clark and A. Karmiloff-Smith, "The cognizer's innards: A psychological and philosophical perspective on the development of thought," *Mind & Language*, vol. 8, no. 4, pp. 487–519, 1993.

[15] J. Bruner, "The growth and structure of skill," in *Mechanisms of Motor Skill Development*, K. J. Connolly, Ed. New York, NY, USA: Academic, 1970, pp. 63–92.

[16] R. F. A. Cox and A. Smitsman, "The planning of tool-to-object relations in young children," *Develop. Psychobiol.*, vol. 48, no. 2, pp. 178–186, 2006.

[17] B. Resende, E. Ottoni, and D. Fragaszy, "Ontogeny of manipulative behavior and nut-cracking in young capuchin monkeys (cebus apella): A perception-action perspective," *Develop. Sci.*, vol. 11, no. 6, pp. 828–840, 2008.

[18] E. Thelen, "Rhythmical behavior in infancy: An ethological perspective," *Develop. Psychol.*, vol. 17, no. 3, pp. 237–257, 1981.

[19] P. Kellman and M. Arterberry, *The Cradle of Knowledge*. Cambridge, MA: MIT Press, 1998.

[20] J. Piaget, *The Construction of Reality in the Child* (in French, 1955). London, U.K.: Routledge & Kegan Paul, 1937.

[21] A. Diamond and J. Gilbert, "Development as progressive inhibitory control of action: Retrieval of a contiguous object," *Cognitive Develop.*, vol. 4, no. 3, pp. 223–249, 1989.

[22] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini, "Developmental robotics: A survey," *Connect. Sci.*, vol. 15, no. 4, pp. 151–190, 2003.

[23] M. H. Lee, Q. Meng, and F. Chao, "Staged competence learning in developmental robotics," *Adapt. Behavior*, vol. 15, no. 3, pp. 241–255, 2007.

[24] L. Berthouze and M. Lungarella, "Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom," *Adapt. Behavior*, vol. 12, no. 1, pp. 47–64, 2004.

[25] M. D. Schmill, T. Oates, and P. R. Cohen, "Learning planning operators in real-world, partially observable environments," in *Proc. 5th Int. Conf. AI Plan. Schedul.*, 2000, pp. 246–253.

[26] J. Bruner, "On voluntary action and its hierarchical structure," *Int. J. Psychol.*, vol. 3, no. 4, pp. 239–255, 1968.

[27] M. Lopes and P.-Y. Oudeyer, "Guest editorial active learning and intrinsically motivated exploration in robots: Advances and challenges," *IEEE Trans. Auton. Mental Develop.*, vol. 2, no. 2, pp. 65–69, 2010.

[28] P. Willatts, "Development of problem-solving strategies in infancy," in *Children's Strategies: Contemporary Views of Cognitive Development*, D. Bjorklund, Ed. Mahwah, NJ: Lawrence Erlbaum, 1990, pp. 23–66.

[29] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd ed. Englewood Cliffs, NJ, USA: Prentice-Hall, 2003.

[30] P. Mounoud and T. G. R. Bower, "Conservation of weight in infants," *Cognition*, vol. 3, no. 1, pp. 29–40, 1974.

[31] H. H. Chaput, "The constructivist learning architecture: A model of cognitive development for robust autonomous robots," Ph.D. dissertation, AI Laboratory, Univ. Texas, Austin, TX, USA, 2004.

[32] E. Sahin, M. Çakmak, M. R. Dogar, E. Ugur, and G. Ücoluk, "To afford or not to afford: A new formalization of affordances toward affordance-based robot control," *Adapt. Behavior*, vol. 15, no. 4, pp. 447–472, 2007.

[33] E. Ugur, E. Oztop, and E. Sahin, "Goal emulation and planning in perceptual space using learned affordances," *Robot. Auton. Sys.*, vol. 59, no. 7–8, pp. 580–595, 2011.

[34] J. Modayil and B. Kuipers, "Autonomous development of a grounded object ontology by a learning robot," in *Proc. 22nd AAAI Nat. Conf. Artif. Intell.*, 2007, vol. 2, pp. 1095–1101.

[35] P. Fitzpatrick and G. Metta, "Grounding vision through experimental manipulation," *Philos. Trans. Royal Soc. A: Math., Physic. Eng. Sci.*, vol. 361, pp. 2165–2185, 2003.

[36] A. Stoytchev, "Robot tool behavior: A developmental approach to autonomous tool use," Ph.D. dissertation, College of Computing, Georgia Instit. Technol., Atlanta, GA, USA, 2007.

[37] S. Hart and R. Grupen, "Learning generalizable control programs," *IEEE Trans. Auton. Mental Develop.*, vol. 3, no. 3, pp. 216–231, Sep. 2011.

[38] E. Thelen, G. Schöner, C. Scheier, and L. B. Smith, "The dynamics of embodiment: A field theory of infant perseverative reaching," *Behavioral Brain Sci.*, vol. 24, no. 1, pp. 1–86, 2001.

[39] M. Carpenter, K. Nagell, and M. Tomasello, "Social cognition, joint attention, and communicative competence from 9 to 15 months of age," *Monographs Soc. Res. Child Develop.*, vol. 63, no. 4, pp. i–174, 1998.

[40] M. Tomasello, *The Cultural Origins of Human Cognition*. Cambridge, MA, USA: Harvard Univ. Press, 1999.

[41] A. Meltzoff and K. Moore, "Persons and representation: Why infant imitation is important for theories of human development," in *Imitation in infancy*, J. Nadel and G. Butterworth, Eds. Cambridge, UK: Cambridge Univ. Press, 1999, pp. 9–35.

[42] G. L. Drescher, *Made-Up Minds, A Constructivist Approach to Artificial Intelligence*. Cambridge, MA, USA: MIT Press, 1991.

[43] G. Stojanov, C. Balkenius, J. Zlatev, H. Kozima, K. Dautenhahn, and C. Breazeal, Eds., "Petitagé: A case study in developmental robotics," in *Proc. Epigenetic Robot. 1*, 2001, pp. 37–44.

[44] F. Perotto, J. Buisson, and L. Alvares, "Constructivist anticipatory learning mechanism (CALM): Dealing with partially deterministic and partially observable environments," in *Proc. 7th Int. Conf. Epigenetic Robot.*, Piscataway, NJ, USA, 2007, pp. 117–127.

[45] R. E. Fikes and N. J. Nilsson, "STRIPS: A new approach to the application of theorem proving to problem solving," *Artif. Intell.*, vol. 2, no. 3-4, pp. 189–208, 1971.

[46] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artif. Intell.*, vol. 101, no. 1-2, pp. 99–134, May 1998.

[47] R. S. Sutton, Verification, the Key to AI, 2006 [Online]. Available: http://www.cs.ualberta.ca/~sutton/IncIdeas/KeytoAI.html, unpublished

[48] T. Zimmerman and S. Kambhampati, "Learning-assisted automated planning: Looking back, taking stock, going forward," *AI Mag.*, vol. 24, no. 2, pp. 73–96, 2003.

[49] J. Mugan and B. Kuipers, "Autonomous learning of high-level states and actions in continuous environments," *IEEE Trans. Auton. Mental Develop.*, vol. 4, no. 1, pp. 70–86, 2012.

[50] J. J. Gibson, *The Ecological Approach to Visual Perception*. Mahwah, NJ, USA: Lawrence Erlbaum Associates, 1986.

[51] J. Modayil and B. Kuipers, "Autonomous development of a grounded object ontology by a learning robot," in *Proc. 22nd AAAI Conf. Artif. Intell.*, Vancouver, BC, Canada, 2007, pp. 1095–1101.

[52] D. Kraft, R. Detry, N. Pugeault, E. Başeski, F. Guerin, J. Piater, and N. Krüger, "Development of object and grasping knowledge by robot exploration," *IEEE Trans. Auton. Mental Develop.*, vol. 2, no. 4, pp. 368–383, 2010.

[53] A. Murata, L. Fadiga, L. Fogassi, V. Gallese, V. Raos, and G. Rizzolatti, "Object representation in the ventral premotor cortex (area F5) of the monkey," *J. Neurophys.*, vol. 78, no. 4, pp. 2226–2230, 1997.

[54] D. Caligiore, A. Borghi, D. Parisi, and G. Baldassarre, "Tropicals: A computational embodied neuroscience model of compatibility effects," *Psychol. Rev.*, vol. 117, no. 4, pp. 1188–228, 2010.

[55] R. Held and A. Hein, "Movement-produced stimulation in the development of visually guided behavior," *J. Compar. Physiol. Psychol.*, vol. 56, no. 5, pp. 872–876, 1963.

[56] J. G. Bremner, *Infancy*. Cambridge, MA., USA: Blackwell, 1994.

[57] D. Mareschal, M. Johnson, S. Sirois, M. Spratling, M. Thomas, and G. Westermann, *Neuroconstructivism, Vol. I: How the Brain Constructs Cognition*. Oxford, U.K.: Oxford Univ. Press, 2007.

[58] N. Krüger, C. Geib, J. Piater, R. Petrick, M. Steedman, F. Wörgötter, A. Ude, T. Asfour, D. Kraft, D. Omrčen, A. Agostini, and R. Dillmann, "Object-action complexes: Grounded abstractions of sensorimotor processes," *Robot. Auton. Syst.*, vol. 59, pp. 740–757, 2011.

[59] I. Bratko, "Autonomous discovery of abstract concepts by a robot," in *Adaptive and Natural Computing Algorithms*, A. Dobnikar, U. Lotric, and B. Ster, Eds. Berlin, Germany: Springer, 2011, vol. 6593, Lecture Notes in Computer Science, pp. 1–11.

[60] E. Thelen and L. B. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA, USA: MIT Press, 1994.

[61] T. Wiesel and D. Hubel, "Ordered arrangement of orientation columns in monkeys lacking visual experience," *J. Comp. Neurol.*, vol. 158, pp. 307–318, 1974.

[62] B. Chapman, M. Stryker, and T. Bonhoeffer, "Development of orientation preference maps in ferret primary visual cortex," *J. Neurosci.*, vol. 15, pp. 6443–6453, 1996.

[63] R. S. Siegler and S. Ellis, "Piaget on childhood," *Psychol. Sci.*, vol. 7, no. 4, pp. 211–215, 1996.

[64] Z. Chen, R. S. Siegler, and M. W. Daehler, "Across the great divide: Bridging the gap between understanding of toddlers' and older children's thinking," *Monographs Soc. Res. Child Develop.*, vol. 65, no. 2, pp. i–105, 2000.

[65] R. S. Siegler, *Emerging Minds : The Process of Change in Children's Thinking*. Oxford, U.K.: Oxford Univ. Press, 1996.

[66] S. Schaal, "Dynamic movement primitives – A framework for motor control in humans and humanoid robots," in *Proc. 2nd Int. Symp. Adapt. Motion Animals Machines (AMAM)*, 2003.

[67] S. K. U. Zibner, C. Faubel, I. Iossifidis, and G. Schöner, "Dynamic neural fields as building blocks of a cortex-inspired architecture for robotic scene representation," *IEEE Trans. Auton. Mental Develop.*, vol. 3, no. 1, pp. 74–91, 2011.

[68] O. Sporns and G. M. Edelman, "Solving bernstein's problem: A proposal for the development of coordinated movement by selection," *Child Develop.*, vol. 64, no. 4, pp. 960–981, Aug. 1993.

[69] I. C. Uzgiris and J. M. Hunt, *Assessment in Infancy : Ordinal Scales of Psychological Development*.   Champaign, IL, USA: Univ. Illinois Press, 1975.

[70] C. von Hofsten, "An action perspective on motor development," *Trends Cognitive Sci.*, vol. 8, no. 6, pp. 266–272, 2004.

[71] A. L. Ciancio, L. Zollo, E. Guglielmelli, D. Caligiore, and G. Baldassarre, "Hierarchical reinforcement learning and central pattern generators for modeling the development of rhythmic manipulation skills," in *Proc. Conf. Develop. Learning Epigenetic Robot.*, 2011.

[72] M. Rolf, J. J. Steil, and M. Gienger, "Online goal babbling for rapid bootstrapping of inverse models in high dimensions," in *Proc. Conf. Develop. Learning Epigenetic Robot.*, 2011.

[73] E. Thelen, "Rhythmical stereotypies in normal human infants," *Animal Behaviour*, vol. 27, no. 3, pp. 699–715, Dec. 1979.

[74] C. K. Rovee and D. T. Rovee, "Conjugate reinforcement of infant exploratory behavior," *J. Experiment. Child Psychol.*, vol. 8, no. 1, pp. 33–39, 1969.

[75] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, pp. 1330–1334, Nov. 2000.

[76] R. Horaud and F. Dornaika, "Hand-eye calibration," *Int. J. Robot. Res.*, vol. 14, no. 3, pp. 195–210, 1995.

[77] R. Saegusa, G. Metta, and G. Sandini, "Own body perception based on visuomotor correlation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, Oct. 2010, pp. 1044–1051.

[78] P. Rochat, "Self-perception and action in infancy," *Experiment. Brain Res.*, vol. 123, no. 1/2, pp. 102–109, 1998.

[79] T. G. R. Bower, *Development in Infancy*.   San Francisco, CA, USA: W.H. Freeman, 1982.

[80] A. Streri and J. Féron, "The development of haptic abilities in very young infants: From perception to cognition," *Infant Behavior Develop.*, vol. 28, no. 3, pp. 290–304, Sep. 2005.

[81] F. Wörgötter, A. Agostini, N. Krüger, N. Shyloa, and B. Porr, "Cognitive agents – A procedural perspective relying on the predictability of object-action-complexes (OACs)," *Robot. Auton. Syst.*, vol. 57, no. 4, pp. 420–432, 2009.

[82] C. von Hofsten, "Eye-hand coordination in the newborn," *Develop. Psychol.*, vol. 18, no. 3, pp. 450–461, 1982.

[83] J. Bruner, "Organization of early skilled action," *Child Develop.*, vol. 44, no. 1, pp. 1–11, 1973.

[84] E. Thelen, D. Corbetta, K. Kamm, J. P. Spencer, K. Schneider, and R. F. Zernicke, "The transition to reaching: Mapping intention and intrinsic dynamics," *Child Develop.*, vol. 64, no. 4, pp. 1058–1098, 1993.

[85] D. C. Witherington, "The development of prospective grasping control between 5 and 7 months: A longitudinal study," *Infancy*, vol. 7, no. 2, pp. 143–161, Apr. 2005.

[86] C. von Hofsten, "Developmental changes in the organization of pre-reaching movements," *Develop. Psychol.*, vol. 20, no. 3, pp. 378–388, 1984.

[87] A. Gordon, "Development of the reach to grasp movement," in *Insights into the Reach to Grasp Movement, Advances in Psychology*, K. Bennett and U. Castiello, Eds.   Amsterdam, The Netherlands: Elsevier, 1994, vol. 105.

[88] P. McDonnell, "The development of visually guided reaching," *Attention Perception Psychophysics*, vol. 18, pp. 181–185, 1975.

[89] R. K. Clifton, D. W. Muir, D. H. Ashmead, and M. G. Clarkson, "Is visually guided reaching in early infancy a myth?," *Child Develop.*, vol. 64, no. 4, pp. 1099–1110, 1993.

[90] M. Schlesinger, D. Parisi, and J. Langer, "Learning to reach by constraining the movement search space," *Develop. Sci.*, vol. 3, pp. 67–80, 2000.

[91] M. Schlesinger and D. Parisi, "Multimodal control of reaching-simulating the role of tactile feedback," *IEEE Trans. Evol. Comput.*, vol. 5, no. 2, pp. 122–128, 2001.

[92] F. Chaumette and S. Hutchinson, "Visual servo control, Part i: Basic approaches," *IEEE Robot. Autom. Mag.*, vol. 13, no. 4, pp. 82–90, Dec. 2006.

[93] F. Chaumette and S. Hutchinson, "Visual servo control, Part ii: Advanced approaches," *IEEE Robot. Autom. Mag.*, vol. 14, no. 1, pp. 109–118, Mar. 2007.

[94] L. Natale, F. Nori, G. Sandini, and G. Metta, "Learning precize 3d reaching in a humanoid robot," in *Proc. IEEE 6th Int. Conf. Develop. Learn. (ICDL)*, 2007, pp. 324–329.

[95] A. Mishra, Y. Aloimonos, and L. F. Cheong, "Active segmentation with fixation," in *Proc. IEEE 12th Int. Conf. Comput. Vision*, 2009, pp. 468–475.

[96] M. Begum and F. Karray, "Visual attention for robotic cognition: A survey," *IEEE Trans. Auton. Mental Develop.*, vol. 3, no. 1, pp. 92–105, 2011.

[97] J. Fagard, "Linked proximal and distal changes in the reaching behavior of 5- to 12-month-old human infants grasping objects of different sizes," *Infant Behavior Develop.*, vol. 23, no. 3–4, pp. 317–329, Mar. 2000.

[98] C. V. Hofsten and S. Fazel-Zandy, "Development of visually guided hand orientation in reaching," *J. Experiment. Child Psychol.*, vol. 38, no. 2, pp. 208–219, 1984.

[99] M. E. McCarty, R. K. Clifton, D. H. Ashmead, P. Lee, and N. Goubet, "How infants use vision for grasping objects," *Child Develop.*, vol. 72, no. 4, pp. 973–987, 2001.

[100] B. C. L. Touwen, "A study on the development of some motor phenomena in infancy," *Develop. Med. Child Neurol.*, vol. 13, no. 1, pp. 435–446, 1971.

[101] C. P. Johnson and P. A. Blasco, "Infant growth and development," *Pediatr. Rev.*, vol. 18, pp. 224–242, 1997.

[102] J. Steffen, R. Haschke, and H. Ritter, "Experience-based and tactile-driven dynamic grasp control," in *Proc. IEEE/RSJ Conf. Intell. Robot. Syst.*, 2007, pp. 2938–2943.

[103] M. Madry, D. Song, and D. Kragic, "From object categories to grasp transfer using probabilistic reasoning," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2012, pp. 1716–1723.

[104] E. W. Bushnell and J. P. Boudreau, "Motor development and the mind: The potential role of motor abilities as a determinant of aspects of perceptual development," *Child Develop.*, vol. 64, no. 4, pp. 1005–1021, 1993.

[105] H. Ruff, "Role of manipulation in infants: Responses to invariant properties of objects," *Develop. Psychol.*, vol. 18, no. 5, pp. 682–691, 1982.

[106] A. Bierbaum, M. Rambow, T. Asfour, and R. Dillmann, "Grasp affordances from multi-fingered tactile exploration using dynamic potential fields," in *Proc. IEEE/RAS Int. Conf. Humanoid Robots (Humanoids)*, 2009, pp. 168–174.

[107] M. Meier, M. Schopfer, R. Haschke, and H. Ritter, "A probabilistic approach to tactile shape reconstruction," *IEEE Trans. Robot.*, vol. 27, no. 3, pp. 630–635, 2011.

[108] E. W. Bushnell and J. P. Boudreau, "Exploring and exploiting objects with the hands during infancy," in *The Psychobiology of the Hand*, K. Connolly, Ed.   Cambridge, U.K.: Cambridge Univ. Press, 1998, pp. 144–161.

[109] T. Striano and E. W. Bushnell, "Haptic perception of material properties by 3-month-old infants," *Infant Behavior Develop.*, vol. 28, no. 3, pp. 266–289, Sep. 2005.

[110] K. S. Bourgeois, A. W. Khawar, S. A. Neal, and J. J. Lockman, "Infant manual exploration of objects, surfaces, and their interrelations," *Infancy*, vol. 8, no. 3, pp. 233–252, 2005.

[111] R. S. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile sensing — from humans to humanoids," *IEEE Trans. Robot.*, vol. 26, pp. 1–20, 2010.

[112] J. Sinapov, V. Sukhoy, R. Sahai, and A. Stoytchev, "Vibrotactile recognition and categorization of surfaces by a humanoid robot," *IEEE Trans. Robot.*, vol. 27, no. 3, pp. 488–497, 2011.

[113] J. Windau and W.-M. Shen, "An inertia-based surface identification system," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2010, pp. 2330–2335.

[114] N. Jamali and C. Sammut, "Material classification by tactile sensing using surface textures," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2010, pp. 2336–2341.

[115] *Stone Knapping: The Necessary Conditions for a Uniquely Hominin Behaviour*, V. Roux and B. Bril, Eds.   London, U.K.: McDonald Institute Monograph Series, 2005.

[116] P. Willatts, "Pulling a support to retrieve a distant object," *Develop. Psychol.*, vol. 35, no. 3, pp. 651–667, 1999.

[117] E. Bates, V. Carlson-Luden, and I. Bretherton, "Perceptual aspects of tool-using in infancy," *Infant Behavior Develop.*, vol. 3, no. 2, pp. 181–190, 1980.

[118] B. Koslowski and J. S. Bruner, "Learning to use a lever," *Child Develop.*, vol. 43, no. 3, pp. 790–799, 1972.

[119] T. Erez and W. Smart, "What does shaping mean for computational reinforcement learning?," in *Proc. 7th IEEE Int. Conf. Develop. Learning*, Aug. 2008, pp. 215–219.

[120] M. Kimmerle, L. A. Mick, and G. F. Michel, "Bimanual role-differentiated toy play during infancy," *Infant Behavior Develop.*, vol. 18, no. 3, pp. 299–307, 1995.

[121] S. J. Lederman and R. L. Klatzky, "Hand movements: A window into haptic object recognition," *Cognitive Psychol.*, vol. 19, no. 3, pp. 342–368, Sep. 1987.

[122] V. Mohan, P. Morasso, G. Metta, and G. Sandini, "A biomimetic, force-field based computational model for motion planning and bimanual coordination in humanoid robots," *Auton. Robots*, vol. 27, no. 3, pp. 291–307, 2009.

[123] J. Steffen, C. Elbrechter, R. Haschke, and H. Ritter, "Bio-inspired motion strategies for a bimanual manipulation task," in *Proc. 10th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)*, , Dec. 2010, pp. 625–630.

[124] P. R. Zelazo and R. B. Kearsley, "The emergence of functional play in infants: Evidence for a major cognitive transition," *J. Applied Develop. Psychol.*, vol. 1, no. 2, pp. 95–117, 1980.

[125] L. Fenson, J. Kagan, R. B. Kearsley, and P. R. Zelazo, "The developmental progression of manipulative play in the first two years," *Child Develop.*, vol. 47, no. 1, pp. 232–236, Mar. 1976.

[126] A. L. Brown, "Domain-specific principles affect learning and transfer in children," *Cognitive Sci.*, vol. 14, no. 1, pp. 107–133, 1990.

[127] M. Hoffmann, H. Marques, A. Arieta, H. Sumioka, M. Lungarella, and R. Pfeifer, "Body schema in robotics: A review," *IEEE Trans. Auton. Mental Develop.*, vol. 2, no. 4, pp. 304–324, Dec. 2010.

[128] H. Örnkloo and C. von Hofsten, "Fitting objects into holes: On the development of spatial cognition skills," *Develop. Psychol.*, vol. 43, no. 2, pp. 404–416, 2007.

[129] T. Yamashita, I. Godler, Y. Takahashi, K. Wada, and R. Katoh, "Peg-and-hole task by robot with force sensor: Simulation and experiment," in *Proc. Int. Conf. Ind. Electron., Control, Instrument.*, 1991, vol. 2, pp. 980–985.

[130] V. Gullapalli, "Skillful control under uncertainty via direct reinforcement learning," *Robot. Auton. Syst.*, vol. 15, no. 4, pp. 237–246, 1995.

[131] S. Street, K. James, S. Jones, and L. Smith, "Vision for action in toddlers: The posting task," *Child Develop.*, vol. 82, no. 6, pp. 2083–2094, 2011.

[132] M. E. McCarty, R. K. Clifton, and R. R. Collard, "Problem solving in infancy: The emergence of an action plan," *Develop. Psychol.*, vol. 35, no. 4, pp. 1091–1101, 1999.

[133] M. Beetz, U. Klank, I. Kresse, A. Maldonado, L. Mösenlechner, D. Pangercic, T. Rühr, and M. Tenorth, "Robotic roommates making pancakes," in *Proc. 11th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)*, Bled, Slovenia, Oct. 26–28, 2011, pp. 529–536.

[134] M. E. McCarty and R. Keen, "Facilitating problem solving performance among 9- and 12-month-old infants," *J. Cognition Develop.*, vol. 6, no. 2, pp. 209–228, 2005.

[135] M. E. McCarty, R. K. Clifton, and R. R. Collard, "The beginnings of tool use by infants and toddlers," *Infancy*, vol. 2, no. 2, pp. 233–256, 2001.

[136] K. Connolly and J. Dalgleish, "The emergence of tool using skills in infancy," *Develop. Psychol.*, vol. 25, no. 6, pp. 894–912, 1989.

[137] K. Connolly and J. Dalgleish, "Individual patterns of tool use by infants," in *Motor Development in Early and Later Childhood: Longitudinal Approaches*, A. Kaverboer, B. Hopkins, and R. Geuze, Eds. Cambridge, U.K.: Eur. Sci. Foundation-Cambridge Univ. Press, 1993, pp. 174–204.

[138] B. Achard and C. von Hofsten, "Development of the infant's ability to retrieve food through a slit," *Infant Child Develop.*, vol. 11, no. 1, pp. 43–56, 2002.

[139] A. Stoytchev, "Learning the affordances of tools using a behavior-grounded approach," in *Towards Affordance-Based Robot Control*, E. Rome, J. Hertzberg, and G. Dorffner, Eds. Berlin, Germany: Springer, 2008, vol. 4760, Lecture Notes in Computer Science, pp. 140–158.

[140] C. Prince, N. Helder, and G. Hollich, L. Berthouze, F. Kaplan, H. Kozima, H. Yano, J. Konczak, G. Metta, J. Nadel, G. Sini, G. Stojanov, and C. Balkenius, Eds., "Ongoing emergence: A core concept in epigenetic robotics," in *Proc. Int. Conf. Epigenetic Robot.*, 2005, Lund Univ. Cognitive Studies, pp. 63–70.

[141] J. Law, M. Lee, M. Hülse, and A. Tomassetti, "The infant development timeline and its application to robot shaping," *Adapt. Behavior*, vol. 19, no. 5, pp. 335–358, 2011.

[142] A. L. Brown, M. J. Kane, and C. H. Echols, "Young children's mental models determine analogical transfer across problems with a common goal structure," *Cognitive Develop.*, vol. 1, pp. 103–121, 1990.

[143] K. S. Rosengren, I. T. Gutiérrez, K. N. Anderson, and S. S. Schein, "Parental reports of children's scale errors in everyday life," *Child Develop.*, vol. 80, no. 6, pp. 1586–1591, 2009.

[144] K. Casler, A. Eshleman, K. Greene, and T. Terziyan, "Children's scale errors with tools," *Develop. Psychol.*, vol. 47, no. 3, pp. 857–866, 2011.

[145] J. Kaufman, D. Mareschal, and M. H. Johnson, "Graspability and object processing in infants," *Infant Behavior Develop.*, vol. 26, no. 4, pp. 516–528, Dec. 2003.

[146] K. C. Soska and S. P. Johnson, "Development of three-dimensional object completion in infancy," *Child Develop.*, vol. 79, no. 5, pp. 1230–1236, Sep. 2008.

[147] L. B. Smith, "From fragments to geometric shape: Changes in visual object recognition between 18- and 24- months," *Current Directions Psychol. Sci.*, vol. 18, no. 5, pp. 290–294, Oct. 2009.

[148] M. Minsky, "Logical versus analogical or symbolic versus connectionist or neat versus scruffy," *AI Mag.*, vol. 12, no. 2, pp. 34–51, Apr. 1991.

[149] H. Liu, "A fuzzy qualitative framework for connecting robot qualitative and quantitative representations," *IEEE Trans. Fuzzy Syst.*, vol. 16, no. 6, pp. 1522–1530, Dec. 2008.

Frank Guerin received the Ph.D. degree from Imperial College, London, in 2002.

Since August 2003, he has been a Lecturer in computing science at the University of Aberdeen, Scotland. He is interested in infant sensorimotor development, especially means–end behavior and precursors to tool use.

Dr. Guerin is a member of The Society for the Study of Artificial Intelligence and Simulation of Behaviour, for which he has served as a committee member and co-chair of the annual convention.

Norbert Krüger received the M.Sc. degree from the Ruhr-Universität Bochum, Germany, and the Ph.D. degree from the University of Bielefeld, Germany.

He is a Professor at the Mærsk Mc-Kinney Møller Institute, University of Southern Denmark, Odense M. He leads the Cognitive Vision Lab, which focuses on computer vision and cognitive systems, and in particular the learning of object representations in the context of grasping.

Dirk Kraft received the Diploma degree in computer science from the University of Karlsruhe, Germany, in 2006, and the Ph.D. degree from the University of Southern Denmark, Odense M, in 2009.

He is currently an Assistant Professor at the Mærsk McKinney Møller Institute, University of Southern Denmark. His research interests include cognitive systems, robotics, and computer vision.