

Some Basic Principles of Developmental Robotics

Alexander Stoytchev, *Member, IEEE*

Abstract—This paper formulates five basic principles of developmental robotics. These principles are formulated based on some of the recurring themes in the developmental learning literature and in the author’s own research. The five principles follow logically from the verification principle (postulated by Richard Sutton) which is assumed to be self-evident. This paper also gives an example of how these principles can be applied to the problem of autonomous tool use in robots.

Index Terms— Artificial intelligence, developmental robotics, intelligent robots, learning systems, principles, psychology, robots, robot programming.

I. INTRODUCTION

DEVELOPMENTAL robotics is one of the newest branches of robotics [1]–[3]. The basic research assumption of this field is that “true intelligence in natural and (possibly) artificial systems presupposes three crucial properties: *embodiment* of the system, *situatedness* in a physical or social environment, and a prolonged *epigenetic developmental process* through which increasingly more complex cognitive structures emerge in the system as a result of interactions with the physical or social environment” [2].

Many fields of science are organized around a small set of fundamental laws (e.g., Newton’s laws in Physics or the fundamental laws of Thermodynamics). Progress in a field without any fundamental laws tends to be slow and incoherent. Once the fundamental laws are formulated, however, the field can thrive by building upon them. This progress continues until the laws are found to be too insufficient to explain the latest experimental evidence. At that point the old laws must be rejected and new laws must be formulated so the scientific progress can continue.

In some fields of science, however, it is not possible to formulate fundamental laws because it would be impossible to prove them, empirically or otherwise. Nevertheless, it is still possible to get around this obstacle by formulating a set of basic principles that are stated in the form of postulates or axioms, i.e., statements that are presented without proof because they are considered to be self-evident. The most famous example of this approach, of course, is Euclid’s formulation of the fundamental axioms of Geometry.

Developmental robotics is still in its infancy and it would be premature to try to come up with the fundamental laws or ax-

ioms of the field. There are some recurring themes in the developmental learning literature and in the author’s own research, however, that can be used to formulate some basic principles. These principles are neither laws (as they cannot be proved at this point) nor axioms (as it would be hard to argue at this point that they are self-evident and/or form a consistent set). Nevertheless, they can be used to guide future research until they are found to be inadequate and it is time to modify or reject them. Five basic principles are described below.

II. THE VERIFICATION PRINCIPLE

Developmental robotics emerged as a field partly as a reaction to the inability of traditional robot architectures to scale up to tasks that require close to human levels of intelligence. One of the primary reasons for scalability problems is that the amount of programming and knowledge engineering that the robot designers have to perform grows very rapidly with the complexity of the robot’s tasks. There is mounting evidence that pre-programming cannot be the solution to the scalability problem. The environments in which the robots are expected to operate are simply too complex and unpredictable. It is naive to think that this complexity can be captured in code before the robot is allowed to experience the world through its own sensors and effectors.

For example, consider the task of programming a household robot with the ability to handle all possible objects that it can encounter inside a home. It is simply not possible for any robot designer to predict the number of objects that the robot may encounter and the contexts in which they can be used over the robot’s projected service time.

There is yet another fundamental problem that pre-programming not only cannot address, but actually makes worse. The problem is that programmers introduce too many hidden assumptions in the robot’s code. If the assumptions fail, and they almost always do, the robot begins to act strangely and the programmers are sent back to the drawing board to try and fix what is wrong. The robot has no way of testing and verifying these hidden assumptions because they are not made explicit. Therefore, the robot is not capable of autonomously adapting to situations that violate these assumptions. The only way to overcome this problem is to put the robot in charge of testing and verifying everything that it learns.

To make this point more clear consider the following example. Ever since the first autonomous mobile robots were developed [4], [5] robots have been avoiding obstacles. After almost 30 years of mobile robotics research, however, there is still not a single robot today that really “understands” what an obstacle is and why it should be avoided. The predominant approach in robotics is still to ignore this question entirely and to make the hidden assumption that an obstacle is equivalent to a

Manuscript received February 14, 2009; revised July 27, 2009. First published date August 11, 2009; current version published October 23, 2009. A preliminary version of this paper appeared in the NIPS 2006 Workshop on Grounding Perception, Knowledge, and Cognition in Sensori-Motor Experience.

The author is with the Developmental Robotics Laboratory, Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA (e-mail: alexs@iastate.edu).

Digital Object Identifier 10.1109/TAMD.2009.2029989

short reading on one of the proximity sensors (e.g., laser range finder, sonar, etc.). The problem with this assumption is that it is true only in some specific situations and not true in general. As a result, outdoor robots cannot distinguish grass from steel rods using a laser range finder. A better approach would be for the robot to try to drive over the detected object at low speed and see if it is displaceable and, thus, traversable. Some recent studies have used this approach successfully for robot navigation [6].

After this introduction, the first basic principle can be stated. It is the so-called *verification principle* that was first postulated by Richard Sutton in a series of on-line essays in 2001 [7], [8]. The principle is stated as follows:

The Verification Principle: *An AI system can create and maintain knowledge only to the extent that it can verify that knowledge itself* [8].

According to Sutton, “the key to a successful AI is that it can tell for itself whether or not it is working correctly” [8]. The only reasonable way to achieve this goal is to put the AI system in charge of its own learning using the verification principle. In other words, all AI systems and AI learning algorithms should follow the motto: *No verification, no learning*. If verification is not possible for some concept then the AI system should not be hand-coded with that concept.

Sutton also points out that the verification principle eventually will be adopted by many AI practitioners because it offers fundamental practical advantages over alternative methods when it comes to scalability. Another way of saying the same thing is: “Never program anything bigger than your head” [8]. Thus, the verification principle stands for autonomous testing and verification performed by the robot and for the robot. As explained above, it would be unrealistic to expect the robot programmers to fix their robots every time the robots encounter a problem due to a hidden assumption.

Sutton was the first researcher in AI to state the verification principle explicitly. However, the origins of the verification principle go back to the ideas of the *logical positivists* philosophers of the 1930s. The two most prominent among them were Rudolf Carnap and Alfred Ayer. They both argued that statements that cannot be either proven or disproven by experience (i.e., metaphysical statements) are meaningless. Ayer defined two types of verifiability, “strong” and “weak,” which he formulated as follows:

“A proposition is said to be verifiable, in the strong sense of the term, if and only if, its truth could be conclusively established in experience. But it is verifiable, in the weak sense, if it is possible for experience to render it probable” [9, p. 37].

Thus, in order to verify something in the “strong” sense, one would have to physically perform the verification sequence. On the other hand, to verify something in the “weak” sense, one does not have to perform the verification sequence directly, but one must have the prerequisite sensors, effectors, and abilities to perform the verification sequence if necessary.

For example, a blind person may be able to verify in the “strong” sense the statement “this object is soft” by physically touching the object and testing its softness. He can also verify

this statement in the “weak” sense as he is physically capable of performing the verification procedure if necessary. However, the same blind person will not be able to verify, neither in the “strong” nor in the “weak” sense, the statement “this object is red” as he does not have the ability to see and thus to perceive colors. In Ayer’s own words:

“But there remain a number of significant propositions, concerning matters of fact, which we could not verify even if we chose; simply because we lack the practical means of placing ourselves in the situation where the relevant observations could be made” [9, p. 36].

The verification principle is easy to state. However, once a commitment is made to follow this principle the implications are far-reaching. In fact, the principle is so different from the practices of traditional autonomous robotics that it changes almost everything. In particular, it forces the programmer to re-think the ways in which learnable quantities are encoded in the robot architecture as anything that is potentially learnable must also be autonomously verifiable.

The verification principle is so profound that the remaining four principles can be considered as its corollaries. As the connection may not be intuitively obvious, however, they will be stated as principles.

III. THE PRINCIPLE OF EMBODIMENT

An important implication of the verification principle is that the robot must have the ability to verify everything that it learns. Because verification cannot be performed in the absence of actions the robot must have some means of affecting the world, i.e., it must have a body.

The principle of embodiment has been defended many times in the literature, e.g., [10]–[15]. It seems that at least in robotics there is a consensus that this principle must be followed. After all, there are not any robots without bodies.

Most of the arguments in favor of the embodiment principle that have been put forward by roboticists, however, are about justifying this principle to its opponents (e.g., [11] and [12]). The reasons for this are historical. The early AI systems, or as Brooks calls them Good Old Fashioned AI (GOF AI), were disembodied and their learning algorithms manipulated data in the computer’s memory without the need to interact with the external world. The creators of these early AI systems believed that a body is not strictly required as the AI system could consist of just pure code, which can still learn and perform intelligently. This reasoning is flawed, however, for the following reason: code cannot be executed in a vacuum. The CPU, the memory bus, and the hard disk play the role of the body. While the code does not make this assumption it is implicitly made for it by the compiler which must know how to translate the code into the target machine language.

As a result of this historic debate most of the arguments in favor of embodiment miss the main point. The debate should not be about whether or not to embrace the principle of embodiment. Instead, the debate should be about the different ways that can be used to program truly embodied robots. Gibbs makes a similar observation about the current state of the art in AI and robotics:

“Despite embracing both embodiment and situatedness in designing enactive robots, most systems fail to capture the way bodily mechanisms are truly embedded in their environments” [15, p. 73].

Some of the arguments used to justify the embodiment principle can easily be explained from the point of view of the verification principle. Nevertheless, the connection between the two has not been made explicit so far. Instead of rehashing the debate in favor of embodiment, which has been argued very eloquently by others, e.g., [10], [15], I am only going to focus on a slightly different interpretation of embodiment in light of the verification principle.

In my opinion, most arguments in favor of the embodiment principle make a distinction between the body and the world and treat the body as something special. In other words, they make the body/world boundary explicit. This distinction, however, is artificial. The only reason why the body may seem special is because the body is the most consistent, the most predictable, and the most verifiable part of the environment. Other than that, there should be no difference between the body and the external world. To the brain, the body may seem special, but that is just because “the brain is the body’s captive audience” [16, p. 160]. In other words, the body is always there and we can’t run away from it.

According to the new interpretation of the embodiment principle described here, the body is still required for the sake of verification. However, the verification principle must also be applicable to the properties of the body, i.e., the properties of the body must be autonomously verifiable as well. Therefore, the learning and exploration principles that the robot uses to explore the external world must be the same as the ones that it uses to explore the properties of its own body.

This interpretation reduces the special status of the body. Instead of treating the body as something special, the new interpretation treats the body as simply the most consistent, the most predictable, and the most verifiable part of the environment. Because of that the body can be easily distinguished from the environment. Furthermore, in any developmental trajectory the body must be explored first.

Distinguishing the body from the external world should be relatively easy because there are certain events that only the owner of the body can experience and no one else. Rochat [17] calls these events *self-specifying* and lists three such events: 1) efferent-afferent loops (e.g., moving one’s hand and seeing it move); 2) double touch (e.g., touching one’s two index fingers together); 3) vocalization behaviors followed by hearing their results (e.g., crying and hearing oneself cry). These events are characterized by the fact that they are multimodal, i.e., they involve more than one sensory or motor modality. Also, these events are autonomously verifiable because you can always repeat the action and observe the same result.

Because the body is constructed from actual verifiable experience, in theory, it should be possible to change one’s body representation. In fact, it turns out that this is surprisingly easy to do. Some experiments have shown that the body/world boundary is very pliable and can be altered in a matter of seconds [18], [19]. For example, it comes as a total surprise for many people to realize that what they normally think of as their own body is

just a phantom created by their brains. There is a very simple experiment which can be performed without any special equipment that exposes the phantom body [18]. The experiment goes like this: a subject places his arm under a table. The person conducting the experiment sits right next to the subject and uses both of his hands to deliver simultaneous taps and strokes to both the subject’s arm (which is under the table) and the surface of the table. If the taps and strokes are delivered synchronously then, after about two minutes, the subject will have the bizarre sensation that the table is part of his body and that part of his skin is stretched out to lie on the surface of the table. Similar extensions and re-mappings of the body have been reported by others [19]–[21].

The conclusions from these studies may seem strange because typically one would assume that embodiment implies that there is a solid representation of the body somewhere in the brain. One possible reason for the phantom body is that the body itself is not constant, but changes over time. Our bodies change with age. They change as we gain or lose weight. They change when we suffer the results of injuries or accidents. In short, our bodies are constantly changing. Thus, it seems impossible that the brain should keep a fixed representation for the body. If this representation is not flexible then sooner or later it will become obsolete and useless.

Another possible reason for the phantom body is that it may be impossible for the brain to predict all complicated events that occur within the body. Therefore, the composition of the body must be constructed continuously from the latest available information. This is eloquently stated by Damasio:

“Moreover, the brain is not likely to predict how all the commands—neural and chemical, but especially the latter—will play out in the body, because the play-out and the resulting states depend on local biochemical contexts and on numerous variables within the body itself which are not fully represented neurally. What is played out in the body is constructed anew, moment by moment, and is not an exact replica of anything that happened before. I suspect that the body states are not algorithmically predictable by the brain, but rather that the brain waits for the body to report what actually has transpired” [16, p. 158].

The author’s previous work [22] describes a computational representation for a Robot body schema (RBS). This representation is learned by the robot from self-observation data. The RBS representation meets the requirements of both the verification principle and the embodiment principle as the robot builds a model for its own body from self-observation data that is repeatedly observable.

The benefits of self-observation learning have also been pointed out by others. For example, Chaminade *et al.* [23] tested the hypothesis that sensorimotor associations for hands and fingers learned from self-observation during motor babbling could be used to bootstrap imitative abilities.

IV. THE PRINCIPLE OF SUBJECTIVITY

The principle of subjectivity also follows quite naturally from the verification principle. If a robot is allowed to learn and maintain only knowledge that it can autonomously verify for itself,

then it follows that what the robot learns must be a function of what the robot has experienced through its own sensors and effectors. In other words, learning must be a function of experience. As a consequence, two robots with the same control architectures, but with different interaction histories with a given object could have two unique representations for the same object, i.e., the two representations will be subjective. The subjectivity, or uniqueness, is due to the capabilities of the robots (sensorimotor limitations) and their specific interaction histories (experiential limitations).

Ayer was probably the first one to recognize that the verification principle implies subjectivity. He observed that if all knowledge must be verifiable through experience, then it follows that all knowledge is subjective as it has to be formed through individual experiences [9, p. 125–126]. Thus, what is learned depends entirely on the capabilities of the learner and the history of interactions between the learner and the environment, or between the learner and its own body. Furthermore, if the learner does not have the capacity to perform a specific verification procedure, then the learner would never be able to learn something that depends on that procedure (as in the blind person example given above). Thus, subjectivity may be for developmental learning what relativity is for physics—a fundamental limitation that cannot be avoided or circumvented.

The subjectivity principle captures very well the subjective nature of object affordances. A similar notion was suggested by Gibson who stated that a child learns “his scale of sizes as commensurate with his body, not with a measuring stick” [24, p. 235]. Thus, an object affords different things to people with different body sizes; an object might be graspable for an adult, but may not be graspable for a child. Noë has recently given a modern interpretation of Gibson’s ideas and has stressed that affordances are also skill relative:

“Affordances are animal-relative, depending, for example, on the size and shape of the animal. It is worth noting that they are also skill-relative. To give an example, a good hitter in baseball is someone for whom a thrown pitch affords certain possibilities for movement. The excellence of a hitter does not consist primarily in having excellent vision. But it may very well consist in the mastery of sensorimotor skills, the possession of which enables a situation to afford an opportunity for action not otherwise available” [25, p. 106].

In robotics, Brooks [26], [27] was the first one to argue that the robot’s *Merkwelt*—the perceptual world in which the robot lives—is different from the programmer’s *Merkwelt*. Brooks strongly objects to the common practice in robotics of using perceptual abstraction which “reduces the input data so that the program experiences the same perceptual world (*Merkwelt*) as humans.” [26]. Thus, it is hopeless for a human to try to impose his own *Merkwelt* on the robot because “each animal species, and clearly each robot species with their own distinctly non-human sensor suites, will have their own different *Merkwelt*. [...]he *Merkwelt* we humans provide our programs is based on our own introspection. It is by no means clear that such a *Merkwelt* is anything like what we actually use internally” [26]. Instead, the robot should be allowed to use its own sensorimotor apparatus to

build its own perceptual world, which would be autonomously verifiable by the robot.

From what has been said so far one can infer that the essence of the principle of subjectivity is that it imposes limitations on what is potentially learnable by a specific agent. In particular, there are two types of limitations: sensorimotor and experiential. Each of them is discussed below along with the adaptation mechanisms that have been adopted by animals and humans to reduce the impact of these limitations.

A. Sensorimotor Limitations

The first limitation imposed on the robot by the subjectivity principle is that what is potentially learnable is determined by the sensorimotor capabilities of the robot’s body. In other words, the subjectivity principle implies that all learning is pre-conditioned on what the body is capable of doing. For example, a blind robot cannot learn what the meaning of the color red is because it does not have the ability to perceive colors.

While it may be impossible to learn something that is beyond the sensorimotor limitations of the body, it is certainly possible to push these limits farther by building tools and instruments. It seems that a common theme in the history of human technological progress is the constant augmentation and extension of the existing capabilities of our bodies. For example, Campbell outlines several technological milestones which have essentially pushed one body limit after another [28]. The technological progression described by Campbell starts with tools that augment our physical abilities (e.g., sticks, stone axes, and spears), then moves to tools and instruments that augment our perceptual abilities (e.g., telescopes and microscopes), and it is currently at the stage of tools that augment our cognitive abilities (e.g., computers and PDAs).

Regardless of how complicated these tools and instruments are, however, their capabilities will always be learned, conceptualized, and understood *relative to our own sensorimotor capabilities*. In other words, the tools and instruments are nothing more than prosthetic devices that can only be used if they are somehow tied to the pre-existing capabilities of our bodies. Furthermore, this tool-body connection can only be established through the verification principle. The only way in which we can understand how a new tool works is by expressing its functionality in terms of our own sensorimotor repertoire. This is true even for tools and instruments that substitute one sensing modality for another. For example, humans have no natural means of reading magnetic fields, but we have invented the compass which allows us to do that. The compass, however, does not convert the direction of the magnetic field into a modality that we cannot interpret, e.g., infrared light. Instead, it converts it to human readable form with the help of a needle.

The exploration process involved in learning the functional properties or affordances of a new tool is not always straight forward. Typically, this process involves active trial and error. Probably the most interesting aspect of this exploration is that the functional properties of the new tool are learned in relation to the existing behavioral repertoire of the learner.

The related work on animal object exploration indicates that animals use stereotyped exploratory behaviors when faced with

a new object [29], [30]. This set of behaviors is species specific and may be genetically pre-determined. For some species of animals, these tests include almost their entire behavioral repertoire: “A young corvid bird, confronted with an object it has never seen, runs through practically all of its behavioral patterns, except social and sexual ones” [30, p. 44].

Unlike crows, adult humans rarely explore a new object by subjecting it to all possible behaviors in their behavioral repertoire. Human object exploration tends to be more focused, although that is not always the case with human infants [29]. Nevertheless, an extensive exploration process similar to the one displayed by crows can sometimes be observed in adult humans as well. This process is easily observed in the members of technologically “primitive” societies when they are exposed to an object for the first time from a technologically advanced society [31, p. 246].

In a previous paper the author described a method for autonomous learning of object affordances by a robot [32]. The robot learns the affordances of different tools in terms of the expected outcomes of specific exploratory behaviors. The affordance representation is inherently subjective as it is expressed in terms of the behavioral repertoire of the robot (i.e., it is skill relative). The affordance representation is also subjective because the affordances are expressed relative to the capabilities of the robot’s body. For example, if an object is too thick to be grasped by the robot, the robot learns that the object is not graspable even though it might be graspable for a different robot with a larger gripper [33].

B. Experiential Limitations

In addition to sensorimotor limitations, the subjectivity principle also imposes experiential limitations on the robot. Experiential limitations restrict what is potentially learnable simply because learning depends on the history of interactions between the robot and the environment, i.e., it depends on experience. Because experience is a function of time, this limitation is essentially due to the finite amount of time that is available for any type of learning. One interesting corollary of this is that: the more intelligent the life form, the longer it has to spend in the developmental stage.

Time is a key factor in developmental learning. By default developmental learning requires interaction with the external world. There is a limit on how fast this interaction can occur, which ultimately restricts the speed of learning. While the limitation of time cannot be avoided it is possible to speed up learning by relying on the experience of others. The reason why this does not violate the subjectivity principle is because verification can be performed in the “weak sense,” and not only in the “strong sense.” Humans, for example, often exploit this shortcut. Ever since writing was invented, we have been able to experience places and events through the words and pictures of others. These vicarious experiences are essential for us.

Vicarious experiences, however, require some sort of basic overlap between our understanding of the world and that of others. Thus, the following question arises: if everything that is learned is subjective, then how can two different people have a common understanding about anything? Obviously this is not a big issue for humans because, otherwise, our civilization will

not be able to function normally. Nevertheless, many philosophers have grappled with this fundamental question.

To answer this question without violating the basic principles that have been stated so far, we must allow for the fact that the representations that two agents have may be functionally different, but nevertheless they can be qualitatively the same. Furthermore, the verification principle can be used to establish the qualitative equivalence between the representations of two different agents. This was well understood by Ayer who stated the following:

“For we define the qualitative identity and difference of two people’s sense-experiences in terms of the similarity and dissimilarity of their reactions to empirical tests. To determine, for instance, whether two people have the same colour sense we observe whether they classify all the colour expanses with which they are confronted in the same way; and, when we say that a man is colour-blind, what we are asserting is that he classifies certain colour expanses in a different way from that in which they would be classified by the majority of people” [9, p. 132].

Another reason why two humans can understand each other even though they have totally different life experiences is because they have very similar physical bodies. While no two human bodies are exactly the same, they still have very similar structure. Furthermore, our bodies have limits which determine how we can explore the world through them (e.g., we can only move our hands so fast). On the other hand, the world is also structured and imposes restrictions on how we can explore it through our actions (e.g., an object that is too wide may not be graspable). Because we have similar bodies and because we live in the same physical world, there is a significant overlap which allows us to have a shared understanding. Similar ideas have been proposed in psychology and have been gaining popularity in recent years [15], [25], [34], [35].

Consequently, experience must constantly shape or change all internal representations of the agent over time. Whatever representations are used they must be flexible enough to be able to change and adapt when new experience becomes available. A good amount of experimental evidence suggests that such adaptation takes place in biological systems. For example, the representation of the fingers in the somatosensory cortex of a monkey depends on the pattern of their use [36]. If two of the fingers are used more often than other fingers then the number of neurons in the somatosensory cortex that are used to encode these two fingers will increase [36].

The affordance representation described in [32] is influenced by the actual history of interactions between the robot and the tools. The affordance representation is pliable and can accommodate the latest empirical evidence about the properties of the tool. For example, the representation can accommodate tools that can break—a drastic change that significantly alters their affordances.

V. THE PRINCIPLE OF GROUNDING

While the verification principle states that all things that the robot learns must be verifiable, the grounding principle describes what constitutes a valid verification. Grounding is very

important because if the verification principle is left unchecked it can easily go into an infinite recursion. At some point there needs to be an indivisible entity which is not brought under further scrutiny, i.e., an entity which does not require additional verification. Thus, figuratively speaking, grounding puts the brakes on verification.

Grounding is a familiar problem in AI. In fact, one of the oldest open problems in AI is the so-called *symbol grounding problem* [37]. Grounding, however, is also a very loaded term. Unfortunately, it is difficult to come up with another term to replace it with. Therefore, for the purposes of this document, the term grounding is used only to refer to the process or the outcome of the process which determines what constitutes a successful verification.

Despite the challenges in defining what constitutes grounding, if we follow the principles outlined so far, we can arrive at the basic components of grounding. The motivation for stating the embodiment principle was that verification is impossible without the ability to affect the world. This implies that the first component that is necessary for successful verification (i.e., grounding) is an action or a behavior.

The action by itself, however, is not very useful for the purposes of successful verification (i.e., grounding) because it does not provide any sort of feedback. In order to verify anything, the robot needs to be able to observe the outcomes of its own actions. Thus, the second component of any verification procedure must be the outcome or outcomes that are associated with the action that was performed.

This leads us to one of the main insights of this section, namely, that grounding consists of Act-Outcome (or Behavior-Observation) pairs. In other words, grounding is achieved through the coupling of actions and their observable outcomes. Piaget expressed this idea when he said that “children are real explorers” and that “they perform experiments in order to see.” Similar ideas have been proposed and defended by others, e.g., [10], [15], [24], [25], [29], [35], [38], and [39].

Grounding of information based on a single act-outcome pair is not sufficient, however, as the outcome may be due to a lucky coincidence. Thus, before grounding can occur the outcome must be replicated at least several times in the same context. If the act-outcome pair can be replicated, then the robot can build up probabilistic confidence that what was observed was not just due to pure coincidence, but that there is a real relationship that can be reliably reproduced in the future.

Thus, grounding requires that action-outcome pairs be coupled with some sort of probabilistic estimates of repeatability. Confidence can be built up over time if multiple executions of the same action lead to the same outcome under similar conditions. In many situations, the robot should be able to repeat the action (or sequence of actions) that were executed just prior to the detection of the outcome. If the outcome can be replicated, then the act-outcome pair is worth remembering as it is autonomously verifiable. Another way to achieve this is to remember only long sequences of (possibly different) act-outcome pairs which are unlikely to occur in any other context due to the length of the sequence. This latter method is closer to Gibson’s ideas for representing affordances.

Stating that grounding is performed in terms of act-outcome pairs coupled with a probabilistic estimate is a good start, but leaves the formulation of grounding somewhat vague. Each action or behavior is itself a very complicated process that involves multiple levels of detail. The same is true for the outcomes or observations. Thus, what remains to be addressed is how to identify the persistent features of a verification sequence that are constant across different contexts. In other words, one needs to identify the sensorimotor invariants. Because the invariants remain unchanged they are worth remembering and thus, can be used for grounding.

While there could be potentially infinite number of ways to ground some information, this section will focus on only one of them. It is arguably the easiest one to pick out from the sensorimotor flux and probably the first one to be discovered developmentally. This mechanism for grounding is based on detection of temporal contingency.

Temporal contingency is a very appealing method for grounding because it abstracts away the nature and complexity of the stimuli involved and reduces them to the relative time of their co-occurrence. The signals could come from different parts of the body and can have their origins in different sensors and actuators.

Temporal contingency is easy to calculate. The only requirement is to have a mechanism for reliable detection of the interval between two events. The events can be represented as binary and the detection can be performed only at the times in which these signals change from 0 to 1 or from 1 to 0. Furthermore, once the delay between two signals is estimated it can be used to predict future events. Based on this, the robot can easily detect that something does not feel right even if the cause for that is not immediately identifiable. A more sophisticated model of contingency detection is used by the Neo system [40], which maintains contingency tables for all pairs of its streams.

Timing contingency detection is used in [41] and [42] to detect which perceptual features belong to the body of the robot. In order to do that, the robot learns the characteristic delay between its motor actions (efferent stimuli) and the movements of perceptual features in the environment (afferent stimuli). This delay can then be used to classify the perceptual stimuli that the robot can detect into “self” and “other.”

Detection of temporal contingency is very important for the normal development of social skills as well. In fact, it has often been suggested that contingency alone is a powerful social signal that plays an important role in language acquisition [43], learning to imitate [44], social perception [45], and more generally in social cognition [46]. Temporal contingency has also been used to explain why you cannot tickle yourself [47]. Watson [48] proposed that the “*contingency relation* between a behavior and a subsequent stimulus may serve as a social signal beyond (possibly even independent of) the signal value of the stimulus itself.” This might be a fruitful area of future robotics research.

VI. THE PRINCIPLE OF INCREMENTAL DEVELOPMENT

The principle of incremental development recognizes the fact that it is impossible to learn everything at the same time. Before

we learn to walk, we must learn how to crawl. Before we learn to read we must learn to recognize individual letters. There is no way around that. Development is incremental and cumulative in nature. The “continuous development and integration of new skills” is a process that has often been called *Ongoing Emergence* [49]. The gradual accumulation of knowledge and skills during development is similar to what Michael Tomasello has called “the ratchet effect” [50]. Once a skill has been learned it can readily be adopted by the individual or by other members of the society to discover new skills, thus, the ratchet has gone up one notch.

Every major developmental theory either assumes or explicitly states that development proceeds in stages [51]–[53]. These theories, however, often disagree about what causes the stages and what triggers the transitions between them. Variations in the timing of these stages have also been observed between the members of the same species. Therefore, the age limits set by some of these theories about what developmental milestone should happen when must be treated as rough guidelines and not as fixed rules.

E. J. Gibson (who was J. J. Gibson’s wife) has even expressed some doubts about the usefulness of formulating stages in developmental learning: “*I want to look for trends in development, but I am very dubious about stages. [...] To repeat, trends do not imply stages in each of which a radically new process emerges, nor do they imply maturation in which a new direction exclusive of learning is created.*” [38, p. 450] Others have embraced the idea of stages, but only as far as they are useful in forming certain developmental invariants: “Although the stages correspond roughly to age levels (at least in the children studied [by Piaget]), their significance is not that they offer behavior norms for specific ages but that the sequence of the stages and the transition from one stage to the next is invariant.” [54, p. 37]. Still others, have suggested that an information-processing perspective might be the only common organizing principle in the literature on infant perception and cognition [55].

Regardless of what causes the stages (or the appearance of stages), one of the most important lessons that roboticists can draw from developmental studies is that the final outcome depends not just on the stages, but on their relative order and duration. Some really good lessons can be learned from an inspirational area of research that compares and contrasts the developmental sequences of different organisms. Comparative studies between primates and humans are useful precisely because they expose the major developmental differences between different species that follow Piaget’s sequence in their development [29], [56].

For example, the time during which autonomous locomotion emerges after birth in primates varies significantly between different species [29], [56]. In chimpanzees this is achieved fairly rapidly and then they begin to move about the environment on their own. In humans, on the other hand, independent locomotion does not emerge until about a year after birth. An important consequence of this is that human infants have a much longer developmental period during which they can manually explore and manipulate objects. They tend to play with objects, rotate them, chew them, throw them, relate them to one another, and bring them to their eyes to take a closer look. In contrast, chim-

panzees are not as interested in sitting down and manually exploring objects because they learn to walk at a much younger age. To the extent that object exploration occurs in chimpanzees it usually is performed when the objects are on the ground [29], [56]. Chimpanzees rarely pick up an object in order to bring it to the eyes and explore it [29]. The full implications of these developmental differences on the overall intelligence of the two species are still being debated. Perhaps developmental robotics can help clarify this in the near future.

Another interesting result from comparative studies is that object exploration (and exploration in general) seems to be self-guided and does not require external reinforcement. What is not yet clear, however, is what process initiates exploration and what process terminates it.

The principle of *incremental development* states that exploration is self-regulated and always proceeds from the most verifiable to the least verifiable parts of the environment. In other words, the exploration is guided by an attention mechanism that is continually attracted to parts of the environment that exhibit medium levels of verifiability. When some parts of the environment are explored fully they begin to exhibit perfect levels of verifiability and thus are no longer interesting. Therefore, the exploration process can chart a developmental trajectory without external reinforcement because what is worth exploring next depends on what is being explored now. Note that the “environment” may refer to the robot’s body as well. For example, once a robot has mastered its body movements, object manipulation movements begin to have medium levels of contingency and should draw the robot’s attention automatically.

The previous section described how temporal contingency can be used for successful verifiability (i.e., grounding). This section builds upon that example, but also takes into account the level of contingency that is detected. At any point in time the parts of the environment that are the most interesting, and thus worth exploring, exhibit medium levels of contingency. To see why this might be the case consider the following example.

In his experiments with infants, Watson (1985) observed that the level of contingency that is detected by the infants is very important. For example, he observed that 16-week-old infants only paid attention to imperfect contingencies. In his experiment, the infants watched a TV monitor which showed a woman’s face. The TV image was manipulated such that the woman’s face would become animated for 2-s intervals after the infant kicked with his legs. The level of this contingency was varied by adjusting the timing delay between the infants’ kicking movements and the animation. Somewhat surprisingly, the infants in this study paid more attention to faces that did not show the perfect contingency (i.e., faces that did not move immediately after the infants’ kicking movements). This result led Watson to conclude that the infant’s attentional mechanisms may be modulated by an inverted U-shaped function based on the contingency of the stimulus [48].

An attention function that has these properties seems ideal for an autonomous robot. If a stimulus exhibits perfect contingency then it is not very interesting as the robot can already predict everything about that stimulus. On the other hand, if the stimulus exhibits very low levels of contingency then the robot cannot learn a predictive model of that stimulus which makes

that stimulus uninteresting as well. Therefore, the really interesting stimuli are those that exhibit medium levels of contingency.

E. J. Gibson reached conclusions similar to those of Watson. She argued that perceptual systems are self-organized in such a way that they always try to reduce uncertainty. Furthermore, this search is self-regulated and does not require external reinforcement, i.e., it is *intrinsically motivating* [57], [58]:

“The search is directed by the task and by intrinsic cognitive motives. The need to get information from the environment is as strong as to get food from it, and obviously useful for survival. The search is terminated not by externally provided rewards and punishments, but by internal reduction of uncertainty. The products of the search have the property of reducing the information to be processed. Perception is thus active, adaptive, and self-regulated” [38, p. 144].

Thus, the main message of this section is that roboticists should try to identify attention functions for autonomous robots that have properties similar to the ones described above. This seems to be a promising area of future research.

VII. AN EXAMPLE: DEVELOPMENTAL SEQUENCE FOR AUTONOMOUS TOOL USE

This section provides an example that uses the five principles described above to formulate a developmental sequence. This sequence can be used by autonomous robots to acquire tool using abilities. By following this sequence, a robot can explore progressively larger chunks of the initially unknown environment that surrounds it. Incremental development is achieved by detecting regularities that can be explained and replicated with the sensorimotor repertoire of the robot. This exploration proceeds from the most predictable to the least predictable parts of the environment.

The developmental sequence begins with learning a model of the robot’s body since the body is the most consistent and predictable part of the environment. Internal models that reliably identify the sensorimotor contingencies associated with the robot’s body are learned from self-observation data. For example, the robot can learn the characteristic delay between its motor actions (efferent stimuli) and the movements of perceptual features in the environment (afferent stimuli). By selecting the most consistently observed delay the robot can learn its own efferent-afferent delay. Furthermore, this delay can be used to classify the perceptual stimuli that the robot can detect into “self” and “other” [41].

Once the perceptual features associated with the robot’s body are identified, the robot can begin to learn certain patterns exhibited by the body itself. For example, the features that belong to the body can be clustered into groups based on their movement contingencies. These groups can then be used to form frames of reference (or body frames) which in turn can be used to both control the movements of the robot as well as to predict the locations of certain stimuli [59].

During the next stage, the robot uses its body as a well defined reference frame from which the movements and positions of environmental objects can be observed. In particular, the robot can learn that certain behaviors (e.g., grasping) can reliably cause an environmental object to move in the same way as some part of the robot’s body (e.g., its wrist) during subsequent robot behaviors. Thus, the robot can learn that the grasping behavior is necessary in order to control the position of the object reliably. This knowledge is used for subsequent tool-using behaviors. One method for learning these first-order (or binding) affordances is described in [33].

Next, the robot can use the previously explored properties of objects and relate them to other objects. In this way, the robot can learn that certain actions with objects can affect other objects, i.e., they can be used as tools. Using the principles of verification and grounding the robot can learn the affordances of tools. The robot can autonomously verify and correct these affordances if the tool changes or breaks [32].

VIII. SUMMARY

This paper proposed five basic principles of developmental robotics. These principles were formulated based on some of the recurring themes in the developmental learning literature and in the author’s own research. The five principles follow logically from the verification principle (postulated by Richard Sutton) which is assumed to be self-evident.

The paper also described an example of how these principles can be applied to autonomous tool use in robots. The author’s previous work describes the individual components of this sequence in more details [59], [60].

ACKNOWLEDGMENT

The author would like to thank A. Silvescu who was one of the graduate students in my Developmental Robotics class in fall 2005. He was the first person to tell me about Ayer’s book and its relevance to the current debates in developmental robotics.

REFERENCES

- [1] J. Weng, J. McClelland, A. Pentland, O. Sporns, I. Stockman, M. Sur, and E. Thelen, “Autonomous mental development by robots and animals,” *Science*, vol. 291, no. 5504, pp. 599–600, Jan. 2001.
- [2] J. Zlatev and C. Balkenius, “Introduction: Why epigenetic robotics?,” in *Proc. 1st Int. Workshop Epigenetic Robot.*, Sep. 2001.
- [3] M. Asada *et al.*, “Cognitive developmental robotics: A survey,” *IEEE Trans. Auton. Mental Develop.*, vol. 1, no. 1, pp. 12–34, May 2009.
- [4] H. P. Moravec, “Obstacle avoidance and navigation in the real world by a seeing robot rover,” Ph.D. dissertation, Stanford University, Stanford, CA, 1980.
- [5] N. J. Nilsson, “Shakey the Robot, AI Center” Menlo Park, CA, 1984, Tech. Rep. 323.
- [6] E. Ugur, M. R. Dogar, M. Cakmak, and E. Sahin, “The learning and use of traversability affordance using range images on a mobile robot,” in *IEEE Int. Conf. Robot. Autom.*, Rome, Italy, Apr. 10–14, 2007.
- [7] R. Sutton, Verification 2001 [Online]. Available: <http://www.cs.ualberta.ca/~sutton/Incldeas/Verification.html>
- [8] R. Sutton, Verification, the Key to AI 2001 [Online]. Available: <http://www.cs.ualberta.ca/~sutton/Incldeas/KeytoAI.html>
- [9] A. J. Ayer, *Language Truth and Logic*. New York: Dover, 1952.

- [10] F. Varela, E. Thompson, and E. Rosch, *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press, 1991.
- [11] R. Brooks and L. Stein, "Building brains for bodies," *Auton. Robots*, vol. 1, no. 1, pp. 7–25, Nov. 1994.
- [12] R. Brooks, C. Breazeal, M. Marjanovic, B. Scassellati, and M. Williamson, "The Cog project: Building a humanoid robot," in *Comput. Metaphors, Analogy, Agents*, C. Nehaniv, Ed. New York: Springer, 1999, pp. 52–87.
- [13] A. Clark, *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press, 1997.
- [14] R. Pfeifer and C. Scheier, *Understanding Intelligence*. Cambridge, MA: MIT Press, 1999.
- [15] R. W. Gibbs, Jr, *Embodiment and Cognitive Science*. Cambridge, MA: Cambridge University Press, 2006.
- [16] A. R. Damasio, *Descartes' Error: Emotion, Reason and the Human Brain*. New York: Gosset/Putnam Press, 1994.
- [17] P. Rochat, "Five levels of self-awareness as they unfold early in life," *Consciousness Cogn.*, vol. 12, pp. 717–731, 2003.
- [18] V. Ramachandran and S. Blakeslee, *Phantoms in the Brain: Probing the Mysteries of the Human Mind*. New York: William Morrow, 1998.
- [19] A. Iriki *et al.*, "Coding of modified body schema during tool use by macaque postcentral neurons," *Neuroreport*, vol. 7, pp. 2325–30, 1996.
- [20] J. Botvinick and J. Cohen, "Rubber hands 'feel' touch that eyes can see," *Nature*, vol. 391, pp. 756–756, 1998.
- [21] H. Ishibashi, S. Hihara, and A. Iriki, "Acquisition and development of monkey tool-use: Behavioral and kinematic analyses," *Canadian J. Physiol. Pharmacol.*, vol. 78, pp. 958–966, 2000.
- [22] A. Stoytchev, "Computational model for an extendable robot body schema" Georgia Inst. Tech., Atlanta, GA, College of Computing Technical Report, GIT-CC-03-44, Oct. 21, 2003.
- [23] T. Chaminade, E. Oztog, G. Cheng, and M. Kawato, "From self-observation to imitation: Visuomotor association on a robotic hand," *Brain Res. Bulletin*, vol. 75, no. 6, pp. 775–84, Apr. 2008.
- [24] J. J. Gibson, *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin, 1979.
- [25] A. Noë, *Action in Perception*. Cambridge, MA: MIT Press, 2004.
- [26] R. A. Brooks, "Achieving artificial intelligence through building robots" Massachusetts Inst. Tech., Cambridge, MA, Tech. Rep. AI Memo 899, 1986.
- [27] R. A. Brooks, "Intelligence without representation," *Artificial Intelligence*, vol. 47, no. 1–3, pp. 139–159, Jan. 1991.
- [28] B. Campbell, *Human Evolution: An Introduction to Man's Adaptations*, 3rd ed. New York: Aldine, 1985.
- [29] T. G. Power, *Play and Exploration in Children and Animals*. Mahwah, NJ: Lawrence Erlbaum, 2000.
- [30] K. Lorenz, "Innate bases of learning," in *Learning as Self-Organization*, K. H. Pribram and J. Kinga, Eds. Mahwah, NJ: Lawrence Erlbaum, 1996.
- [31] J. M. Diamond, *Guns, Germs, and Steel: The Fates of Human Societies*. New York: Norton, 1999.
- [32] A. Stoytchev, "Behavior-grounded representation of tool affordances," in *Proc. IEEE ICRA*, Barcelona, Spain, 2005, pp. 3071–3076.
- [33] A. Stoytchev, "Toward learning the binding affordances of objects: A behavior-grounded approach," in *Proc. AAAI Symp. Develop. Robot.*, Mar. 21–23, 2005, pp. 17–22.
- [34] A. Glenberg, "What memory is for," *Behav. Brain Sci.*, vol. 20, no. 1, pp. 1–55, 1997.
- [35] J. K. O'Regan and A. Noë, "A sensorimotor account of vision and visual consciousness," *Behavioral Brain Sci.*, vol. 24, no. 5, 2001.
- [36] X. Wang, M. Merzenich, K. Sameshima, and W. Jenkins, "Remodeling of hand representation in adult cortex determined by timing of tactile stimulation," *Nature*, vol. 378, pp. 71–75, Nov. 1995.
- [37] S. Harnad, "The symbol grounding problem," *Physica D*, vol. 42, pp. 335–346, 1990.
- [38] E. J. Gibson, *Principles of Perceptual Learning and Development*. New York: Appleton-Century-Crofts, 1969.
- [39] D. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control," *Neural Netw.*, vol. 11, no. 7–8, pp. 1317–1329, Oct./Nov. 1998.
- [40] P. R. Cohen, M. S. Atkin, T. Oates, and C. R. Beal, "Neo: Learning conceptual knowledge by sensorimotor interaction with an environment," in *Proc. 1st Int. Conf. Auton. Agents*, Marina del Ray, CA, 1997, pp. 170–177.
- [41] A. Stoytchev, "Toward video-guided robot behaviors," in *Proc. 7th Int. Conf. Epigenetic Robot.*, NJ, Nov. 5–7, 2007, pp. 165–172.
- [42] P. Michel, K. Gold, and B. Scassellati, "Motion-based robotic self-recognition," in *Proc. IEEE RSJ IROS*, Sendai, Japan, 2004.
- [43] M. Goldstein, A. King, and M. West, "Social interaction shapes babbling: Testing parallels between birdsong and speech," *Proc. Nat. Acad. Sci.*, vol. 100, no. 13, Jun. 2003.
- [44] S. S. Jones, "Infants learn to imitate by being imitated," in *Proc. 5th Int. Conf. Develop. Learning*, Bloomington, IN, Jun. 3, 2006.
- [45] D. Wolpert, K. Doya, and M. Kawato, "A unifying computational framework for motor control and social interaction," *Philos. Trans. Roy. Soc. B: Biol. Sci.*, vol. 358, no. 1431, pp. 593–602, Mar. 2003.
- [46] C. D. Frith, *Making up the Mind: How the Brain Creates Our Mental World*. Malden, MA: Wiley-Blackwell, 2007.
- [47] S.-J. Blakemore, D. M. Wolpert, and C. D. Frith, "Why can't you tickle yourself?," *NeuroReport*, vol. 11, no. 11, pp. R11–6, Aug. 2000.
- [48] J. S. Watson, "Contingency perception in early social development," in *Social Perception in Infants*, T. M. Field and N. A. Fox, Eds. Norwood, N. J: Ablex Pub. Corp, 1985, pp. 157–176.
- [49] C. G. Prince, N. A. Helder, and G. J. Hollich, "Ongoing emergence: A core concept in epigenetic robotics," in *Proc. 5th Int. Workshop Epigenetic Robotics*, Nara, Japan, Jul. 22–24, 2005, pp. 63–70.
- [50] M. Tomasello, *The Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press, 2001.
- [51] J. Piaget, *The Origins of Intelligence in Children*. New York: International Universities Press, 1952.
- [52] S. Freud, *New Introductory Lectures on Psychoanalysis*. New York: Norton, 1965.
- [53] J. Bowlby, *Attachment*. New York: Basic Books, 1969, vol. 1.
- [54] P. H. Wolff, "The developmental psychologies of Jean Piaget and psychoanalysis," *Psychol. Issues*, vol. 2, no. 1, pp. 3–181, 1960.
- [55] L. B. Cohen and C. H. Cashon, "Infant perception and cognition," in *Comprehensive Handbook of Psychol., Ser. Develop. Psychol. II. Infancy*, R. Lerner, A. Easterbrooks, and J. Mistry, Eds. New York: Wiley, 2003, vol. 6, pp. 65–89.
- [56] M. Tomasello and J. Call, *Primate Cognition*. New York: Oxford University Press, 1997.
- [57] A. Barto and O. Şimşec, "Intrinsic motivation for reinforcement learning systems," in *Proc. 13th Yale Workshop Adaptive Learning Syst.*, New Haven, CT, 2005, pp. 113–118.
- [58] P.-Y. Oudeyer, F. Kaplan, and V. Hafner, "Intrinsic motivation systems for autonomous mental development," *IEEE Trans. Evol. Comput.*, vol. 11, no. 2, pp. 265–286, Apr. 2007.
- [59] A. Stoytchev, "Robot tool behavior: A developmental approach to autonomous tool use," Ph.D. dissertation, College of Computing, Georgia Institute of Technology, Atlanta, GA, 2007.
- [60] A. Stoytchev, "Five basic principles of developmental robotics," presented at the NIPS 2006 Workshop on Grounding Perception, Knowledge, and Cognition in Sensori-Motor Experience, Whistler, BC, Canada, Dec. 8, 2006.



Alexander Stoytchev (S'00–M'07) received the M.S. and Ph.D. degrees in computer science from the Georgia Institute of Technology, Atlanta in 2001 and 2007, respectively.

He is currently an Assistant Professor of Electrical and Computer Engineering and Director of the Developmental Robotics Laboratory at Iowa State University, Ames. His research interests are in the areas of developmental robotics, autonomous robotics, computational perception, and machine learning.