

Emergence of eye–hand coordination as a creative process in an artificial developmental agent

Adaptive Behavior
2017, Vol. 25(6) 289–314
© The Author(s) 2017
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1059712317732116
journals.sagepub.com/home/adb



Wendy Aguilar¹ and Rafael Pérez y Pérez²

Abstract

Based on Piaget's theory of cognitive development and an extension of the Engagement-Reflection model of creativity, the *Developmental Engagement-Reflection* (Dev E-R) model characterizes the early cognitive development of an agent as a creative activity. In its first version, Dev E-R uses an agent that can see and move its head; through interactions with its environment, it is able to develop elaborated behaviors consistent with Piaget's ideas. This work describes an advancement of our model. We give our agent a hand (tactile sensor) so it can detect the presence and features of an object in its environment; we also study the necessary mechanisms to coordinate its vision with the sense of touch. We report the behavior of the agent when it is granted the capacity of touching without seeing (i.e. the agent was “blind”) and when both skills, touch and sight, come together. For such purpose, we place an agent in a virtual environment and let it perform in different contexts. We analyze how new knowledge structures results from prior experiences and interactions with the environment. The outcomes from the experiments reveal that it learns new skills associated with eye–hand coordination. We observe that the arising developmental behavior resembles some of the features reported by Jean Piaget.

Keywords

Developmental robotics, creative behavior, Piaget, engagement-reflection, Dev E-R, eye–hand coordination

Handling Editor: Tom Froese, National Autonomous University of Mexico, Mexico

1. Introduction

Cognitive development is a fascinating subject that is taking the attention of many researchers in areas like developmental approaches to artificial intelligence (AI), cognitive science, philosophy, and so on (e.g. Asada et al., 2009; Guerin, 2011a; Lungarella, Mettay, Pfeiferz, & Sandiniy, 2003; Weng et al., 2001). One of its main goals is to study agents' capability of adaptation to its (sometimes changing) environment. Adaptation is seen by many academics as a necessary condition for creative behavior (e.g. Cohen, 1989; Gorney, 2007; Runco, 2007). Piaget (1936/1952) characterizes adaptation as a mechanism comprised by two processes called assimilation and accommodation: Assimilation allows children to face new situations by using their knowledge from past experiences; accommodation allows dealing with new situations by progressively modifying their expertise in order to incorporate the results of their new experiences. Clearly, these three concepts—cognitive development, adaptation, and creative behavior—are closely related

to each other. We use them as a base for our research. The *Developmental Engagement-Reflection* (Dev E-R) model represents the early cognitive development of an agent as a creative activity (Aguilar & Pérez y Pérez, 2015). Dev E-R characterizes some of the aspects that take place during the sensorimotor period described by Piaget (1936/1952). The model uses an extended version of the computational model of creativity known as Engagement-Reflection (Pérez y Pérez, 2007; Pérez y

¹Departamento de Ciencias de la Computación, Instituto de Investigaciones en Matemáticas Aplicadas y en Sistemas, Universidad Nacional Autónoma de México (UNAM), Ciudad de México, México.

²Departamento de Tecnologías de la Información, División de Ciencias de la Comunicación y Diseño, UAM Cuajimalpa, Ciudad de México, México

Corresponding author:

Wendy Aguilar, Departamento de Ciencias de la Computación, Instituto de Investigaciones en Matemáticas Aplicadas y en Sistemas, Universidad Nacional Autónoma de México (UNAM), Circuito Escolar, Ciudad Universitaria, Col. Copilco, Del. Coyoacán, Ciudad de México 01000, México.

Email: weam@turing.iimas.unam.mx

Pérez & Sharples, 2001, 2004) to represent Piaget's assimilation-accommodation adaptation process. To our knowledge, this is the first computational system that addresses cognitive development as a creative process. It proposes three characteristics for the construction of developmental agents: (1) a hedonic intrinsic motivation system based on the recovery and preservation of pleasant stimuli; (2) a knowledge representation based on affective responses, emotional reactions, and motivations; and (3) a learning mechanism driven by surprise and cognitive curiosity, based on generalization and differentiation of schemas, and on using its knowledge of past experiences to deal with new similar, but not identical, situations.

In Aguilar and Pérez y Pérez (2015), we reported the behaviors learned by the agent when it was granted with the capability of seeing its world (a living room with furniture, plants, and some toys). This article complements that work. Here, we present an advancement of Dev E-R where we incorporate a hand that allows the agent to touch its environment. Our main interest is to study which new behaviors arise as a result of adding this new sensory ability. Because, as far as we know, there are no other developmental agents that employ the Engagement-Reflection Model, we consider necessary to study further the limitations and scopes of our model. In this article, we have the hypothesis that if Dev E-R provides the bases for a framework flexible enough for implementing a developmental agent, the incorporation of new senses (such as the sense of touch) should be possible without modifying the core of the current knowledge structures and internal processes. That is, the same model must provide the necessary infrastructure to develop new abilities. These are the questions that drive this research: What modifications does Dev E-R require in order to add the sense of touch? How can the sense of touch and sense of sight be coordinated within our model? Given these new features, what new abilities will our agent develop? Our results suggest that the new skills developed by the agent resemble what Piaget describes as the first behaviors associated with hand and the first sight-touch coordination seen in newborns from ages 0 to 4 months. In this way, the inclusion of the sight-touch coordination seems to produce in our model a more elaborated cognitive development. Furthermore, this result suggests that it is possible to include new senses to our agent without substantial modifications of the model. Nevertheless, it is necessary to assess deeply the flexibility of our approach.

This article is structured as follows: section 2 presents a summary of related work; section 3 contains a description of the developmental agent and the Dev E-R model; section 4 explains the changes needed to give the sense of touch to the agent; section 5 describes the experiments we performed to test our model as well as the results obtained; in section 6, we discuss the behaviors learned by the agent in the context of Piaget's

theory, its developmental path, an evaluation of its development as a creative process, and a comparison with other methods; finally, in section 7, we present the conclusions.

2. Related work

Probably, Alan Turing (1950) is the first person to conceive the idea of developing a program that simulates an artificial baby which could be later educated like a child until it gets an adult level. Around the same time (1920s–1970s), Jean Piaget, an epistemologist, psychologist, and biologist, publishes important research papers related to the development of intelligence from infants to adults, which are still a reference to anyone interested in cognitive development (see, for example, Piaget, 1924, 1936/1952, 1954, 1985). Two decades later, during the 1960s and 1970s, Papert (1963) and Boden (1978) suggest that AI and Piaget's theory of cognitive development can be highly useful from each other.

It is until the beginning of the 1990s, that under the influence of the idea of embodiment (which states that intelligent behavior can only come from the interaction between the mind, the body, and the environment), a new area of research named as *developmental robotics* (sometimes also called *epigenetic robotics*) arose. Its focus of interest is in the intersection between robotics and developmental sciences. A review of this emerging area can be found in Asada et al. (2009); Lungarella et al. (2003); Elliott and Shadbolt (2003); Metta, Sandini, Natale, and Panerai (2001); Asada, MacDorman, Ishiguro, and Kuniyoshi (2001); Weng et al. (2001); and Sandini, Metta, and Konczak (1997).

From that decade to the present, different computational models that simulate some aspects of early cognitive development, from Piaget's perspective, have appeared. Guerin (2011a) and Stojanov (2009) present a review of these kinds of works. However, none of these systems consider the relation between cognitive development and creativity. In 2013, such relationship was discussed during the Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium called "Creativity and Cognitive Development." The discussion was around the hypothesis that the same mechanisms involved in the generation of creative artifacts are observed during cognitive development, in particular during the constant reconceptualization of one's understanding of the environment. In a paper presented in that symposium, Sandra Bruno (2013) suggests that creativity is what enables adaptation, and that in the very early stages of development, one of the first creative attitudes involves the transformation of reflexes into schema. In Stojanov and Indurkha (2013), the authors reflect about how research in developmental AI and robotics remains

completely disconnected from computational creativity, and propose to see the process of analogy as one of the common mechanisms present in both fields, as well as the consequences of this point of view. In the same line of thought, Bickhard (2013) presents a discussion about the creativity of development and the development of creativity. On one hand, he states that “development is inherently a matter of an evolutionary epistemology, and, thus, inherently involves creative construction processes in the face of new situations and problems,” and on the other hand, he discusses the different aspects of an internal evolutionary epistemology that can contribute to creativity. Other related works that were presented in this symposium include Miller (2013) and Indurkha (2013). Nevertheless, except for our work (Aguilar & Pérez y Pérez, 2013), no other computer system that models early cognitive development as a creative activity was presented. We believe that this is an important relationship that needs to be further studied, and we refer to it as development of early-creative-behavior.

3. The developmental agent and the Dev E-R model

This section provides a summary of our model introduced in Aguilar and Pérez y Pérez (2015).

The developmental agent is named Jacques, after Jacqueline, one of the daughters of Jean Piaget, who was an object of study and inspiration for him. Jacques interacts with a three-dimensional (3D) virtual world containing simple 3D models of typical objects that may be found in real life. Such objects have the following characteristics: luminous or non-luminous, static or

moving, color, and size. The agent uses a virtual camera (set in its right eye) with a field of vision of 60° to visually sense the world. The images taken are internally represented as a 180 × 120 × 3 matrix. Its field of view is divided into nine different areas. It has the capability of moving its head 10° to the right, left, up, down, up-right, down-right, up-left, and down-left. These are the physical actions, or external actions, that the agent can perform.

The general perception-action cycle that Jacques implements is summarized in Figure 1.

3.1. Motivational system

The agent simulates affective responses, emotional reactions, and motivations that push it to act. These are inspired by Piaget’s ideas, associated with the relation between affectivity and development of intelligence (Piaget, 1981). The first responses consist of intensity and valence, represented by the agent through variables that span along a scale of -1, +1, +2, wherein -1 represents disliking and +1/+2 represent 2 degrees of liking. The rest are represented internally as Boolean variables having a true value when the agent presents such emotional reactions or motivations. The situations under which these are triggered are listed in Table 1.

3.2. Knowledge representation

The agent has a memory wherein it stores all its knowledge. Particularly, the agent saves in this memory its current perception of the world (represented through the *Current-Context* structure) and actions to interact with its environment (represented through schemas).

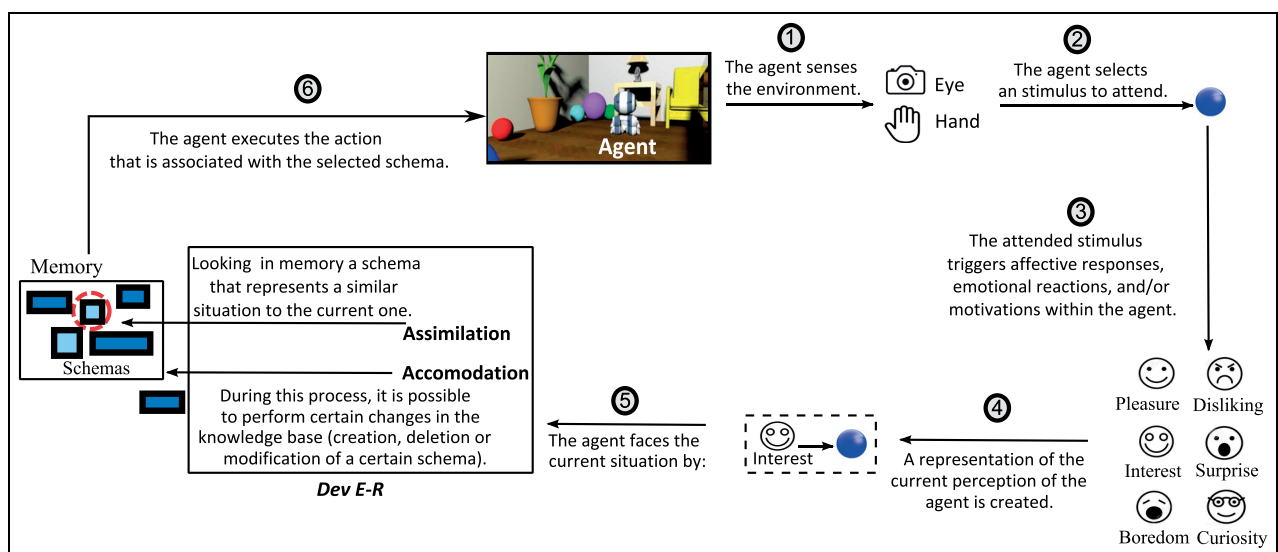


Figure 1. The perception-action cycle used in Dev E-R.

Table 1. Summary of the situations under which the different affective responses, emotional reactions, and motivations are triggered.

Name	Description
Affective responses	
<i>Pleasure</i>	It is triggered when the agent pays attention to a bright spot. It can assume three values: + 2 when the agent pays attention to a bright blob that it sees in the center of its visual field, + 1 when the element is seen in its peripheral area, and -1 when the agent loses a stimulus it likes, that is when the attended item disappears from its visual field. In other words, the -1 value is equivalent to disliking.
Emotional reactions	
<i>Interest</i>	Triggered as a result of executing an internal action named <i>show_interest_in A</i> , where <i>A</i> can be any stimuli (visual or tactile).
<i>Surprise</i>	Triggered when the agent “accidentally” recovers an object it likes, and when the current situation is faced by using agent’s experience from a similar, but not identical, situation, and even then the expectations are met.
<i>Boredom</i>	Triggered when the agent attends, in a consecutive sequence, the same object a number of times.
Motivations	
<i>Cognitive curiosity</i>	Triggered when the expectations of the agent are not met.

3.2.1. Current-Context. The *Current-Context* is a structure composed by two parts: (1) the *Current-Visual-Context*, and (2) the agent’s current expectations, which are defined as an *Expected Current-Visual-Context* (explained later in this section). The *Current-Visual-Context* is composed by two parts: (1) the features of the object that is in the center of attention of the agent (its color, size, movement, and position within the visual field); and (2) the affective responses, emotional reactions, and motivations triggered in the agent by such object. This way, the agent is able to represent its present perception of the world in terms of what the object(s) at the center of its attention (described through their features) are provoking in it, that is, whether they are causing a feeling of liking, disliking, interest, boredom, surprise, cognitive curiosity, and/or any expectation. Figure 2(a) shows an example of a *Current-Context*. With the purpose of providing a simpler, more compact notation, henceforth, current contexts composed solely by affective responses are herein represented in the form of Figure 2(b).

3.2.2. Schemas. Schemas as used herein are knowledge structures simulating the sensory-motor schemas described by Piaget. There are two types: basic and developed.

Basic schemas represent innate behaviors and tendencies observed by Piaget in babies, which are present in the agent from its initialization. These are represented as contexts associated to actions. The contexts used in the schemas are structures similar to the *Current-Context* structure. As opposed to the *Current-Context* structure, the contexts of the schemas may define the features of the object in terms of non-instantiated variables. As an example of the foregoing, Figure 3(a) shows an illustration of a basic schema, which associates the situation of feeling disliking,

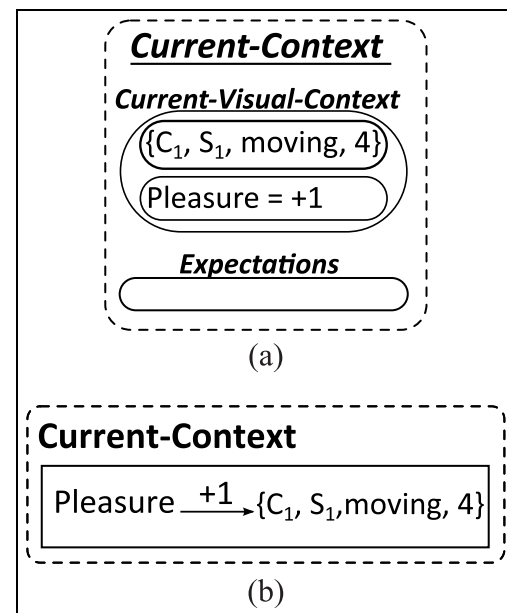


Figure 2. (a) An example of a *Current-Context* structure, which represents that the agent is seeing a pleasant object, which consists of a ball of C_1 color, and size S_1 , which is moving within position 4 of the visual field; and (b) the same *Current-Context* using the alternative notation.

triggered by an object of any color = a , of any size = b , with any movement status = c , and in any position within the visual field = d , with the action of performing a random external action (e.g. a movement of the head to the right). Henceforth, small letters shall represent non-instantiated variables. When the n features describing the object in a context are non-instantiated variables, it is said that the structure to which it is associated is an abstract schema of type T_n (as the schema shown in Figure 3(a), which is type T_4). On the contrary, if all features are instantiated variables, we have

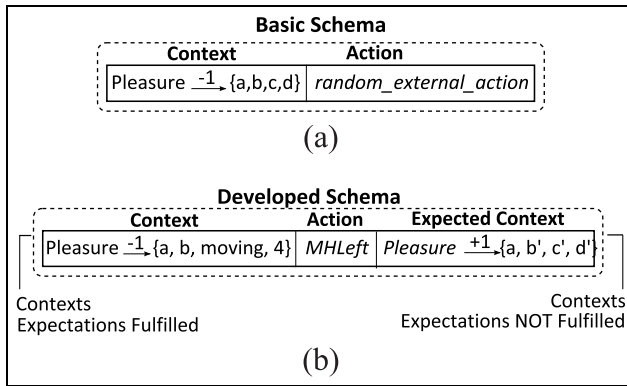


Figure 3. An example of a (a) basic schema and (b) developed schema.

a concrete schema of type T_0 . Similarly, there are intermediate schemas, types T_1, T_2, \dots, T_{n-1} , with several features in terms of non-instance variables. The concept of type is important for the implementation of the accommodation process, which is based on generalizations (modifying a schema so that it becomes more abstract, for example, changing it from T_2 to T_4) and in differentiations (modifying a schema so that it becomes more particular, for example, changing it from T_3 to T_1). These processes are described in section 3.3.2.

Developed schemas are constructed based on the interactions of the agent with its environment (explained later in this section) and represent new behaviors. These are composed by a context, an action to be executed, an expected context, and a set of contexts with which the expectations have been fulfilled (named “Contexts Expectations Fulfilled”) and others that were not fulfilled (termed “Contexts Expectations NOT Fulfilled”). Figure 3(b) shows an example of a developed schema, which associates the disliking situation triggered by a moving object of any color, any size, in position 4 within the visual field of the agent, with the action of moving the head left, and the expectation of recovering the affective response of pleasure toward that same object (two objects are considered the same if both have the same color). This is an example of a T_2 schema, as two of the features of the context (without considering the ones used in the expected context) are in terms of non-instantiated variables.

The term *visual schema* is used to define the structures whose context refers only to visual objects.

3.3. Adaptation and learning mechanisms

The adaptation and learning mechanisms are in charge of using and constructing the knowledge of the agent (represented as sensory-motor schemas).

3.3.1. General functioning. The Dev E-R model, in *Engagement* mode, searches its memory to find all

schemas whose contexts represent a similar situation to the one described in the *Current-Context*. If during this process the agent is found to know more than one way to act given the current situation, then one of those ways is preferred. The selection is performed in such a way that the developed schemas are assigned a higher probability of being chosen over the basic ones, and from the developed schemas, the one resulting in the highest number of expectations fulfilled and expected to result in the most pleasure is the one that will most likely be picked out. The final decision is made on a random basis, which considers the above information.

For instance, let us suppose that the current context of the agent is the one shown in Figure 2, and that during Engagement, the agent found three possible schemas to match: (1) a basic schema with an action where the agent continues showing interest in the moving ball; (2) a developed schema with an action where the agent moves its head to the left; with that movement, it expects to increase its pleasure with a probability p_2 (such a probability depends on its experience); (3) a developed schema, similar to the second one, with the associated action of moving the head to the left and up; it expects to increase the pleasure with a probability p_3 , where $p_2 > p_3$. The system selects at random which schema to match; however, because $p_2 > p_3$, the second structure has a higher probability to be chosen than the third one; because the agent prefers to match developed schemas rather than basic schemas, the first option in this example has the lower possibility to be used. In this way, given the same situation and the same memory state, most of the time the agent will move its head toward the left; in some other occasions, it will move its head toward the left and up; and occasionally, it will show interest toward the ball. Therefore, under equal conditions, the agent will not always behave in the same way. The simulation is non-deterministic; it depends on the knowledge that the agent has accumulated in memory.

When a schema is selected, its associated action is executed; in case the designated structure is a developed schema, then the expectations are registered in the *Current-Context*. The agent senses once more its world, updating the structure of the *Current-Context*, and the cycle continues. When no schema may be matched in the memory, that is, when the agent faces an unknown situation, then an impasse is declared. In this event, an adaptation process is required, whether by assimilation or by accommodation. These processes may be performed automatically or analytically, for instance, through an analogic reasoning. However, when the agent initiates the execution, it lacks reflexive skills to help it deal with such type of situations (based on the early sub-stages that are being modeled). Consequently, adaptation in this implementation is simulated as an automatic activity that is being performed in *Engagement* mode.

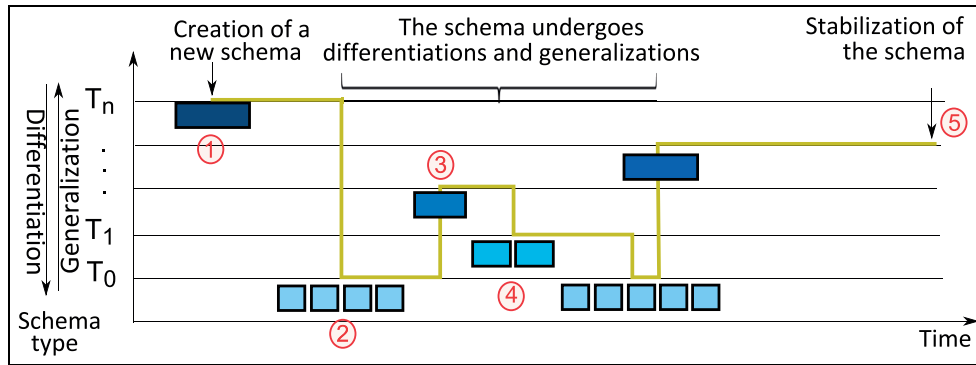


Figure 4. Summary of the accommodation processes of the agent's schemas: (1) A new T_n -type schema is formed when the agent presents an emotional reaction of surprise, that is, when it accidentally recovers an object of its interest. (2) The schema is differentiated in particular structures, type T_0 , when its associated expectations are not fulfilled in a determined percentage of times. (3) It is generalized into a schema type $T_{n1} \dots T_1$ when it is detected that one same action may lead to recovering several objects with different features. (4) If the schema continues to lead the agent to enter a conflictive cognitive state, then it is considered to be still very general and is thus differentiated once more into particular schemas. (5) This process continues until the schema is deleted or until accommodations stop because the expectations are fulfilled most of the times, that is, until it becomes stabilized.

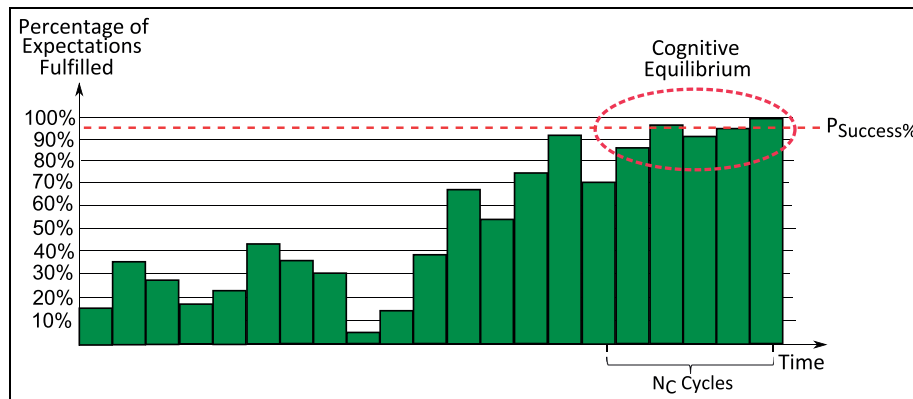


Figure 5. A depiction of the moment when the agent enters a cognitive equilibrium state. When the agent initializes and begins to create its own schemas, its expectations are fulfilled in a very low percentage of times, resulting in the necessity of accommodating its knowledge. In time, the agent is capable of interacting with its world without the need to further modify its knowledge, as the expectations were fulfilled most of the times (at least in $P_{Success}\%$ of the times, during the last N_C cycles). At this time, the agent is said to have entered a state of cognitive equilibrium.

3.3.2. Simulation of the accommodation process. Accommodation, as understood under Piaget's theory, refers to the process through which the child modifies an existing schema or creates an entirely new one to deal with an unknown object or event (Ormrod, 2012). Inspired by this concept, in the Dev E-R model, the knowledge structures are gradually created and modified by means of the generalization and differentiation processes summarized in Figure 4.

An important aspect to highlight is that at some point in time, agent's memory contains schemas of various types that may involve different senses. For a detailed explanation of the accommodation process, see Aguilar and Pérez y Pérez (2015).

3.3.3. Simulation of the cognitive equilibration process. The accommodation process continues until the moment arrives when the agent can interact with its world

during the last N_C cycles without detecting the need to modify its knowledge, as the expectations were fulfilled most of the times (at least in $P_{Success}\%$ of the times), as depicted in Figure 5. At that moment, we can say that the agent has entered a state of cognitive equilibrium, and the schemas that have been built by the agent up to that point are considered stabilized.

If any schema is built, deleted, or modified again, then the agent is considered to have entered again in a state of cognitive disequilibrium. Consequently, it will have to meet again the condition stating that the need to modify its knowledge was not detected in the last N_C cycles, as the expectations were fulfilled most of the times, and so a state of cognitive equilibrium is achieved again.

Each time the agent enters a state of cognitive equilibrium, its capability of finding partial matches between its *Current-Context* and its basic schemas and those developed schemas that have been stabilized is

activated, as described in further detail in the following section. At the beginning, the agent was not able to perform such skill.

3.3.4. Simulation of the assimilation process. Assimilation refers to the process of responding to new facts and situations according to what is already known and recoverable from the memory (Guerin, 2011b). This process is modeled in Dev E-R by seeking schemas representing situations which are similar to the one described in the *Current-Context*.

At the beginning of the execution, a 100% match must be found between both structures. However, once stabilized schemas begin to appear on the agent's knowledge base, it begins to allow partial matches, which may be presented in two ways. First, if the *Current-Context* consists of a single affective response, then any of the elements thereof is allowed to differ from any of the elements of the context associated with a schema (the type, valence, or intensity of the affective response, or in the properties of the object—color, size, movement status, or position within the visual field). Second, if the *Current-Context* consists of more than a single affective response, then each of the latter is allowed to match a different schema, or even one of said responses may be allowed to not have a match. It is important to stress that partial matches are performed only in basic schemas, as well as in developed schemas which are already considered stable.

Sometimes, assimilation leads to accommodation. For example, in cases where the agent (1) faces a current situation by performing a partial match between such situation and any of the developed schemas and (2) after applying the associated action, the expectations of the agent are fulfilled. Under such circumstances, an emotional reaction of surprise is triggered, leading to the construction of a new schema representing the way in which the agent successfully faced the new situation. With the creation of new schemas, the knowledge base of the agent suffers further accommodations, leading to a new state of cognitive disequilibrium. The agent remains at this stage until the new schemas are stabilized. At that moment, the agent enters again into a state of cognitive equilibrium. The new stabilized schemas begin to be used in finding partial matches, thus leading to the creation of new schemas and causing the agent to enter once again into a state of cognitive disequilibrium, and so on. This is how the agent simulates the process of cognitive development, going from a state of disequilibrium to a state of equilibrium, and again to disequilibrium, and so on.

4. Adding the sense of touch to the agent

To give the sense of touch to the agent, five main changes are needed:

1. Create a simulated tactile sensor that can detect (1) the presence of an object in contact with the palm of the hand (the agent can only touch one element at a time) and (2) its texture. Regarding the latter, it is possible to achieve an implementation which enables the agent to increase its ability to differentiate various tactile properties prevalent on the surface of the elements in the surroundings (see, for example, Jamali, Byrnes-Preston, Salleh, & Sammut, 2009). These may include, for instance, several degrees of roughness. This may be carried out analogous to vision (see Aguilar & Pérez y Pérez, 2015). Thus, at the beginning, the agent would only be able to differentiate between very smooth and very rough. In time, through the interaction with the virtual world, the agent would be able to differentiate several degrees of roughness. However, for simplicity purposes, in this article, we are assuming that the agent has learned to recognize a number of textures, which have been labeled as “ t_1 ,” “ t_2 ,” “ t_3 ,” and so on.
2. Add a tactile texture to each object in the environment (in this implementation, this is done by assigning them a label such as “ t_1 ”).
3. Update the motivation system in such a way that (1) an affective response of pleasure with intensity +1 is triggered when the agent has selected a tactile stimuli as its center of attention; and (2) an affective response of pleasure with intensity -1 (i.e. displeasure) is triggered when the agent loses a tactile stimulus it likes, for example, when it drops the attended object.
4. Change the *Current-Context* structure to include a new component named *Current-Tactile-Context* which is analogous to the *Current-Visual-Context*. This structure corresponds to an internal representation of what the agent is touching. The schemas with only tactile information are named *tactile schemas*.
5. Modify the attention module in such a way that when the agent is detected to have been in cognitive equilibrium for N_C cycles, and its knowledge base contains stabilized visual and tactile schemas, then the agent becomes capable of representing both visual and tactile data in its current contexts. In other words, the agent becomes capable of paying attention to what it sees and what it touches simultaneously. The term *tactile-visual schema* is used to describe those schemas referring to visual and tactile stimuli.

5. Experiments and results

We present in Aguilar and Pérez y Pérez (2015) the results obtained when the agent is granted with the ability to see but not touch the world around it. Said results are summarized herein below because these are used for the subsequent experiments. For this article, we are

interested in finding what new behaviors arise when the agent is granted the ability to touch. For this purpose, two new sets of experiments were designed. The first one consists in configuring the agent so that it could only touch but not see its world. The second one involves granting the agent both abilities: seeing and touching its environment. The description of said experiments, along with the results obtained, is reported herein below.

5.1. First set of experiments: the agent can only see its world

The first set of experiments involved letting the agent interact in three separate executions within the living room shown in Figure 6. This environment contained 3D models of pieces of furniture, plants, toys, and so on. All the objects were static, except for five balls of different colors, which began to move at different times and in different predefined directions (sometimes they rolled from left to right, and back; other times, they bounced). In these experiments, the agent was initialized with the first two schemas shown in Figure 7, which represent “innate” tendencies described by Piaget (1936/1952). The first innate conduct allowed the agent to keep its attention focused on the objects of its interest; the second one permitted the agent to perform an attempt at recovering such objects whenever they disappear. Jacques finished its execution (in one of the runs, after having remained in a state of cognitive equilibrium during the last 1000 cycles) with 29 new schemas in total (13 to recover objects it lost in the different positions within the visual field, 8 to keep those that were moving, and 8 to preserve those that were static).

When the agent used jointly the schemas it had developed, we noticed that it learned the following behaviors (listed herein in the same order in which they were constructed): (1) recovering back in its visual field the elements of agent’s interest which came out of said field, (2) following visually pleasant objects, (3) centering within the visual field the elements which were of agent’s interest and which were moving, and (4) centering in the visual field the elements that were interesting for the agent and which were static.

5.2. Second set of experiments: the agent can only touch its world

The second set of experiments consisted in configuring the agent so that it could only touch but not see its world. Its development was considered completed when it remained in a state of cognitive equilibrium during the last 250 cycles, that is, until the agent showed to have acquired new skills that enabled it to interact with the environment (recovering and preserving the tactile objects that were pleasant for the agent), using the




Figure 6. Image of the environment with which the agent interacted during the experiments.

Piaget	Agent
- Tendency to preserve pleasant stimuli.	Basic Schema₁ Context: Pleasure $+1 \rightarrow A$ Action: <i>show_interest_in A</i>
- Tendency to make an attempt to recover a pleasant object when it disappears.	Basic Schema₂ Context: Pleasure $-1 \rightarrow A$ Action: <i>random_physical_action</i>
- Reflex behavior of closing hand when it comes in contact with an object.	Basic Schema₃ Context: Pleasure $+1 \{t, \text{open_hand}\}$ Action: <i>close_hand</i>

Figure 7. Basic schemas with which the agent was initialized. These schemas represent the predefined behaviors that were known initially by the agent to interact with its world.

knowledge that it had built. The reason why we selected 250 cycles is because empirically we realized that the tactile skills that it acquires in this setup are faster to learn than the ones when it could only see. Thus, if we had used 1000 cycles as in the previous experiment, the only consequence would have been that it would take longer to finish its development. However, if we had used a smaller value than 250, there is a risk of a premature stabilization of the schemas causing the agent to be unable of predicting the consequences of its actions with an adequate precision.

Dev E-R is a research tool; therefore, the user can modify a set of parameters that influence the behavior of the agent (see Table 2). We performed several tests to study the consequences of altering their values. The following lines provide a summary of what we found out. When the first four parameters in Table 2 were set to high values, the agent took longer to develop all its schemas; so, we prefer to use low values. The parameter $P_{FailureT4}$ shapes the time that the agent spends trying to recover lost objects by using the same action that previously worked well. We observed that choosing a low percentage such as 10% is a good compromise between giving the schema time to gain knowledge about the consequences of its associated action and allowing the agent to try different actions to recover pleasant stimuli. Parameters $P_{DeleteT0}$, $P_{DeleteT4}$, and $P_{DeleteT3T2}$ affect the decision of which schemas should be deleted because

Table 2. Values and descriptions of the main Dev E-R model parameters used in the second set of experiments.

Parameter name	Description	Value
N_{UsedT4}	Minimum number of times a schema type T_4 had to be used before it could be accommodated.	4
N_{Trials}	Minimum number of times an object must be assimilated into a schema so that it is considered in the differentiation process.	2
N_{T0}	Necessary number of schemas type T_0 in order to be able to perform a generalization.	2
N_{UsedT0}	Minimum number of objects assimilated into a schema type T_0 required to delete it.	6
$P_{FailureT4}$	Minimum percentage of failure by a schema type T_4 in order to consider its accommodation.	10%
$P_{DeleteT0}$	Minimum percentage of failures needed to delete a schema type T_0 .	50%
$P_{DeleteT4}$	Minimum percentage of failure to delete a schema type T_4 .	90%
$P_{DeleteT3T2}$	Minimum percentage of failure by the schema types T_3 and T_2 so they can be deleted.	85%
$P_{SuccessObj}$	Minimum percentage of success to recover the same object so that it can be considered into the differentiation process.	60%

Dev E-R: Developmental Engagement-Reflection.

they are not good enough predicting the consequences of their associated actions. For instance, $P_{DeleteT0}$ is linked to schema type T_0 . Because these types of schemas connect actions to very specific stimuli, a failure rate greater than 50% indicates a bad accuracy of the prediction. On the contrary, $P_{DeleteT4}$ and $P_{DeleteT3T2}$ are linked to schemas that continuously are generalized and differentiated. As a result, their set of fulfilled and not fulfilled expectations is constantly changing. Thus, such schemas should only be deleted when the experience registered in them mainly represents failure cases; otherwise, valuable experience may be lost. To achieve this behavior, it is necessary to set the parameters with high values. Finally, the last parameter is related to the flexibility of the agent to deal with indeterminacy of the consequences of the agent's actions. The lower it is set, the more flexible it is. For example, if every time the agent moves its head, the attended object changes to a contiguous area of its field of vision, then a value of 100% could be used. Based on the previous observations, we initialized our agent with the parameters shown in Table 2 to perform our experiments.

The agent was also initialized with the following initial knowledge.

5.2.1. Initial knowledge. The agent was initialized with the three basic schemas shown in Figure 7. These represent the only initial behaviors that were known by the agent to interact with its world. The third one modeled the reflex behavior of closing the hand when a pleasant element comes into contact with it. Variable A represented any tangible object defined as $A = \{\text{texture}, \text{hand_status}\}$, wherein said hand status may have either of these values: *closed_hand* or *open_hand*.

Finally, the only physical (external) actions that the agent may perform were those associated with the movement of its hand (up, down, right, left, front, back, closed, and open).

5.2.2. Virtual world setup. In this set of experiments, the agent interacted again in three separate executions

within the living room of the house shown in Figure 6. In these executions, all the objects were static, except for the five balls that were moving as follows:

- When the agent had its hand open and not in contact with any object, sometimes the system randomly picked any of the five balls and placed it in its hand (so that its touch sensor could detect it during the next cycle).
- When the touch sensor was in contact with any object and the agent executed the action *close_hand*, then it was considered that the object had been grasped.
- The grasped objects moved accordingly to hand movements.
- By default, after n cycles (for this experiment, $n = 1$), the agent would automatically open its hand (unless when, during the current cycle, the action *close_hand* had been selected).
- Upon the hand being opened, the object that had been grasped could (1) remain in the same position and continue in contact with the touch sensor or (2) go back to its initial position. The selection of (1) or (2) above was made by the system on a random basis.

Regarding the tactile features that the agent could recognize on the objects it touched, this experiment was initialized with the capacity of recognizing five different textures (labeled " t_1 ," " t_2 ," " t_3 ," " t_4 ," and " t_5 "), one for each ball.

The three executions began with the agent sitting in the middle of the environment with its hand open in front of it, and five balls in positions that were out of the agent's reach. The results obtained are reported herein below.

5.2.3. Results. We should point out that, just like in the first set of experiments, the new skills acquired by the agent were based on learning how to preserve affective responses of pleasure triggered by the objects of the

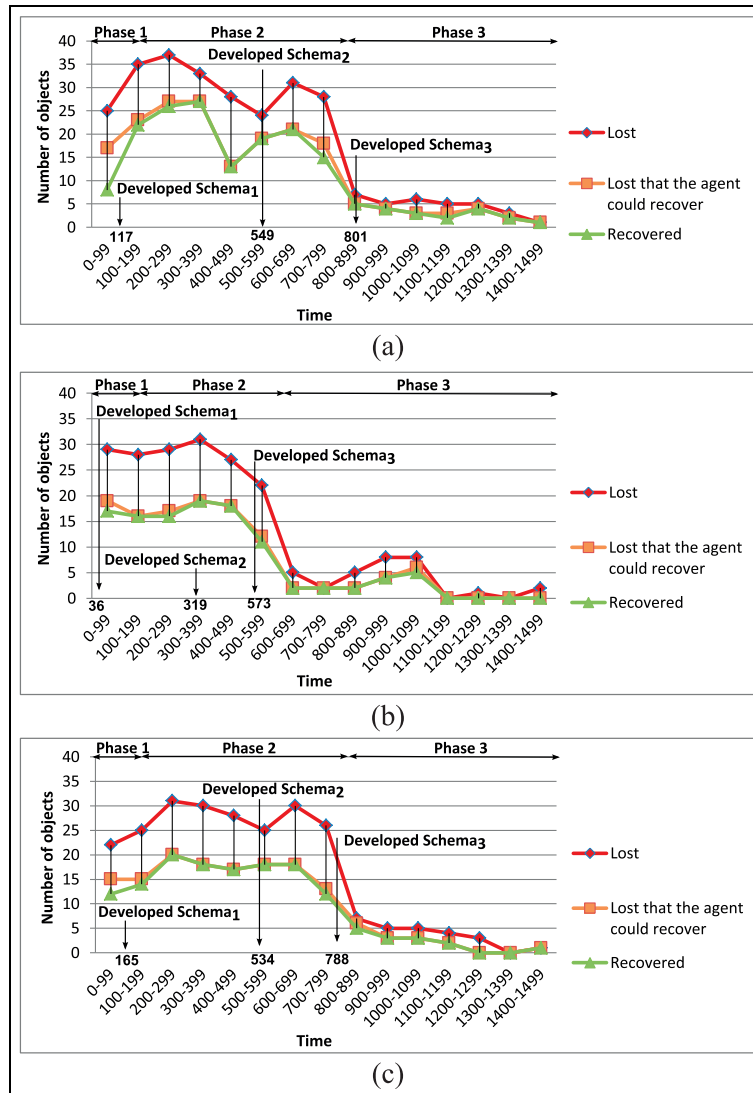


Figure 8. These graphs show the evolution of the number of pleasant objects lost by the agent, in contrast to the number of objects that the agent could recover, in each of the three executions: (a) execution 1, (b) execution 2, and (c) execution 3.

agent's interest and learning to recover them when they disappear. With this in mind, in the graphs of Figure 8, in groups of 100 cycles, the following is presented: (1) the number of pleasant objects lost by the agent, that is, which it had grasped but were lost when opening the hand (shown in red diamonds); (2) the number of objects lost that the agent could recover, as upon release by the hand, the system chose to leave them in contact with its hand (shown in orange squares); and (3) the number of objects that the agent was able to recover successfully (in green triangles). Three phases may be identified in these graphs:

1. *Phase 1* includes approximately from cycle 0 to cycle 200 in the three executions. This corresponds to a period within which the agent was able to recover the lesser number of lost objects.
2. *Phase 2* includes approximately from cycle 200 to cycle 700 or 900, depending on the execution. This

corresponds to a period within which the agent began to recover nearly all the objects that it lost and which were susceptible of being recovered.

3. *Phase 3* includes approximately from cycle 700 or 900 to cycle 1500. This third and last phase corresponds to a period within which a considerable decrease in the total number of objects lost by the agent was observed, falling to almost zero near the end of the runs.

Each stage is further discussed below.

Phase 1.

During phase 1, the agent began interacting with its world by using only its three basic schemas (shown in Figure 7). From there, it started creating its first developed schemas by generalization, differentiation, and deletion of those, where the associated expectations




Figure 9. Schema created during phase 1.

were not fulfilled most of the times. Accordingly, by the end of this first phase, the agent had built one single developed schema (the same happened throughout the three executions), which is shown in Figure 9. This first schema associates the situation of having opposed affective responses caused by the same object (unpleasantness due to the loss of an element that had been grasped, and pleasure for detecting the same object on an ongoing basis, but now with the hand open), with the action of closing the hand and the expectation of recovering the affective response of pleasure resulting from grasping again the object of interest. In other words, this schema contains in itself the knowledge that (1) the agent is able to recover the pleasant objects that were lost but which are still sensed with the open hand, and (2) it may recover said objects by closing its hand. With the creation of this new schema, the agent has learned to recover tangible objects which it is interested in. This situation leads it to a second phase.

Phase 2.

The creation of the schema constructed in the previous phase resulted in, on one hand, the agent to enter a period within which it began to recover nearly all the objects lost and which were susceptible of being recovered (see Figure 8). On the other hand, it also caused a change in the behavior of the agent's expectations, which, as shown in Figure 10 (wherein the number of expectations is contrasted to the number of expectations fulfilled), began to be fulfilled 100% of the times from the creation of said schema. These two situations (learning how to successfully recover the objects of interest and having the expectations 100% of the times on an ongoing basis) led the agent to become able to interact with its environment during N_C cycles (for these experiments, $N_C = 250$), without the need to modify its knowledge, resulting in the achievement of a status of cognitive equilibrium for the first time (indicated in the bottom portion of the graphs shown in Figure 10). As a result of having reached said state, the agent gradually became able to find partial matches among the current contexts and its developed schema (shown in Figure 9). One of said partial matches (when the agent sensed an object with its hand opened and it used the schema of Figure 9 to deal with this situation) resulted in the creation of a second schema, shown in

Figure 11. This associates the situation of having an affective response of pleasure triggered by touching any object with the open hand, with the action of closing the hand, and with the expectation of having again an affective response of pleasure caused by touching the same object, but now with the closed hand. The construction of this second schema caused the agent to begin closing its hand when it came into contact with any of the objects of interest, but not as a result of a reflex behavior (through the use of a basic schema) but rather because of a developed behavior having an expectation associated therewith.

Again, with the creation of this second schema, the agent managed to interact with its environment for N_C additional cycles without having the need to modify its knowledge, leading, for a second time, to a cognitive equilibrium state, near the end of phase 2 (see Figure 10). As a result of entering such state, the agent became able to identify partial matches between the current contexts and the two recently created developed schemas. One of such partial matches (when the agent sensed an object with its hand closed and it used the schema of Figure 11 to deal with this situation) resulted in the creation of a third schema, shown in Figure 12. The new schema associates the situation of having an affective response of pleasure triggered by touching any object with the closed hand, with the action of closing the hand, and with the expectation of maintaining the affective response of pleasure caused by holding the same object.

The construction of this third schema resulted in that, from that point onward, the agent began to maintain its hand closed when it was holding an object of its interest, which was then released when an emotional reaction of boredom was triggered (please note that this emotional reaction is triggered when the agent maintains a hold on a same object over several cycles, 10 cycles for these experiments). In other words, the agent learned to hold on to the objects of its interest, which led to the third and final stage.

Phase 3.

This is characterized in that it was a period within which a considerable decrease in the total number of objects lost by the agent was observed and resulted from the creation of a third schema through which it was able to keep a hold of the objects that were interesting for it. To show the rising of this new behavior, we have created the graphs shown in Figure 13, wherein the average number of cycles over which the agent kept a hold of an interesting object along the execution is shown.

In said graphs, we can see that, from the creation of a third schema, near the start of phase 3, the average number of cycles over which the agent kept a hold of the interesting objects was increased.

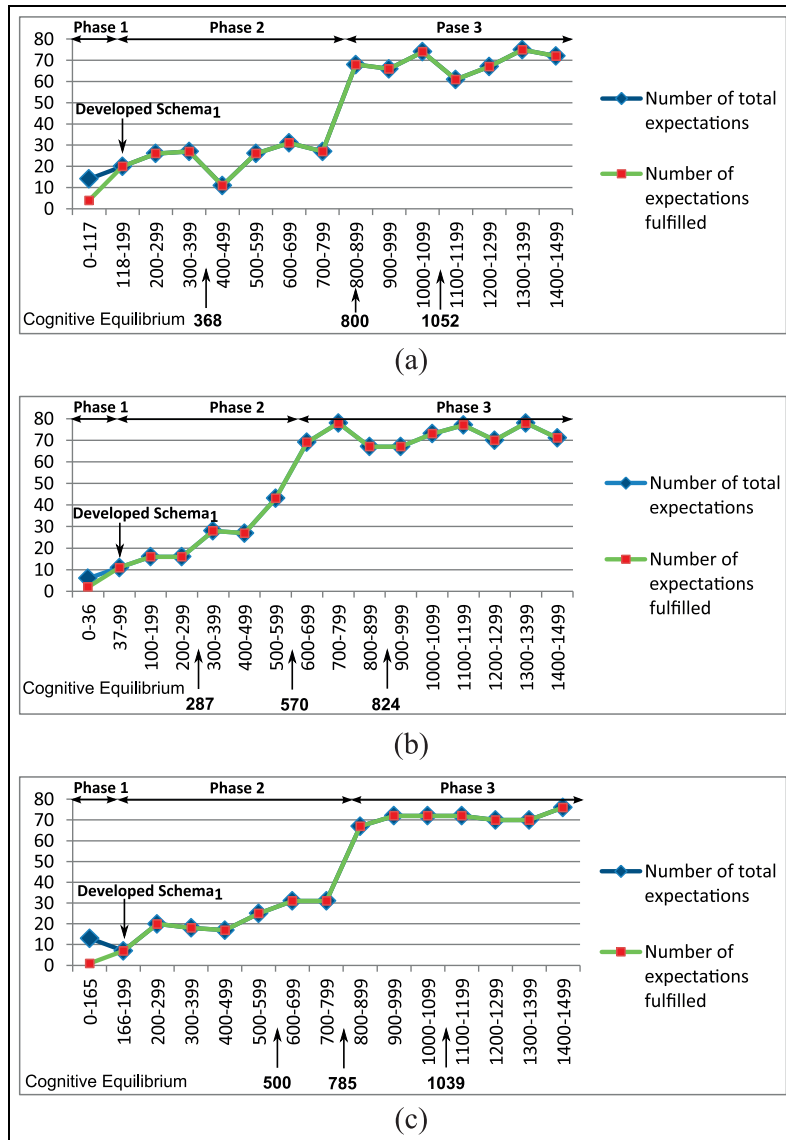


Figure 10. These graphs show the number of total expectations created by the agent in contrast to the number of expectations fulfilled in each of the three executions: (a) execution 1, (b) execution 2, and (c) execution 3.




Figure 11. Schema created during phase 2.




Figure 12. Schema created by the end of phase 2.

After interacting with its environment over 250 additional cycles, the agent once again entered a state of cognitive equilibrium, remaining in said state during 500 cycles more, thus concluding its execution.

5.2.4. Summary. In this second set of experiments in which the agent could only touch but not see its world, it learned how to

- Recover the pleasant objects that it was holding (it learned that this can only be done with the objects that are let go but which it continues to sense with the open hand),
- Hold the objects of its interest which are in contact with the palm of its hand, and

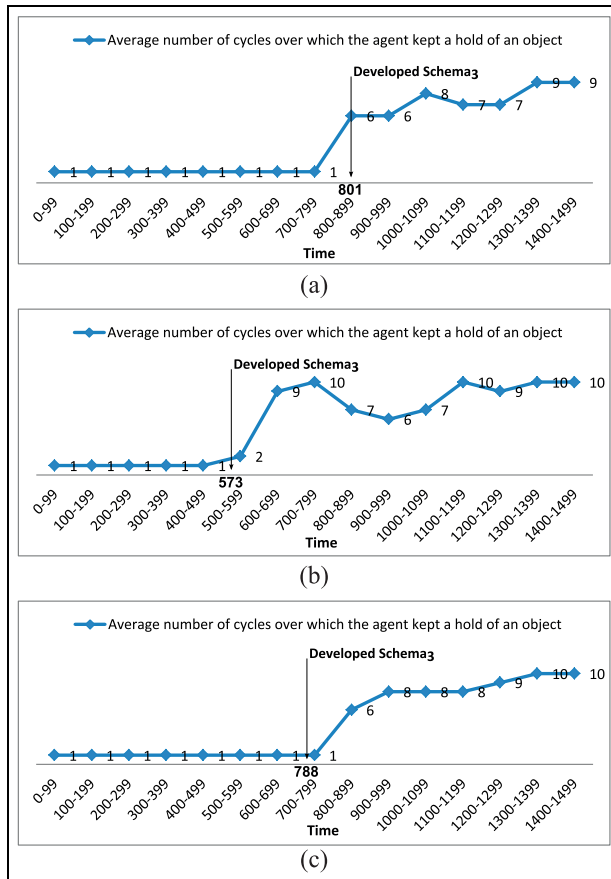


Figure 13. These graphs show the evolution in the average number of cycles over which the agent kept a hold of the interesting object along the three executions: (a) execution 1, (b) execution 2, and (c) execution 3.

- To hold on to the objects of its interest.

5.3. Third set of experiments: the agent can see and touch its world

The third set of experiments consisted in configuring the agent so that it could both see and touch its world. Development was considered completed when it remained in a state of cognitive equilibrium during the last 2000 cycles. In order to carry out the experiments, the agent was configured with the parameters of Dev E-R shown in Table 2 (except for parameter N_C , which was given the value of 2000), as well as the following initial knowledge.

5.3.1. Initial knowledge. The agent was initialized with the three basic schemas shown in Figure 7, wherein in this case, variable A represents any visual element defined as $A = \{\text{color, size, movement, position}\}$, or any tangible element defined as $A = \{\text{texture, hand_status}\}$. Additionally, its knowledge base was initialized with the schemas developed when it could

only see and when it could only touch the world around it (generated in the first and second set of experiments). In other words, initialization was started with the knowledge of how to recover and maintain the visual and tangible objects of agent's interest (including the skills of visually following the objects of interest, centering them in the visual field, as well as to hold on to the objects that come in contact with the hand and keeping them held until another object attracted its attention or until it became bored). The reason for the above was that, in this experiment, we were interested in observing the set of behaviors that arise when the agent has constructed and stabilized both its visual and tactile schemas.

Finally, physical or external actions that could be performed by the agent included all the possible head and hand movements.

5.3.2. Virtual world setup. This time, the agent was allowed to interact again in three separate executions within the living room of the house shown in Figure 6. The agent was in a sitting position in the middle of the environment, with the head looking to the front and with its hand outside the visual field. All the objects were static, except for agent's hand, which was moving in random directions at some points in time. Later in this run, the balls began to move in the same way they behaved during the second set of experiments (refer to section 5.2.2).

5.3.3. Results. Upon starting, the first observation we found was that when agent's hand came into the visual field, that part of its body caught its attention and it began to follow that luminous element with head movements (using the developed schemas that were available at initialization). This behavior is exemplified in Figure 14. It is important to note that, up to this moment, the agent sees its hand moving, not exactly because it is a part of itself (as it has not developed any knowledge structure which allows it to distinguish its own body from the rest of the objects in the environment) but because the hand catches its attention due to the color, size, and movement of the hand itself, as if it were any other element present in the virtual world.

After 500 cycles, the balls began to move to enter into contact with the agent's hand. From that moment, the behavior of the agent consisted in holding the objects that were in contact with its sensor and then releasing the items when it lost interest in them, and in following the visual elements that caught its attention. In other words, the agent continued to interact with the environment by using all the knowledge it was initialized with. After 1500 additional cycles, the agent reached a state of cognitive equilibrium, causing the schemas stored in memory to be considered stable, thus being the agent capable of (1) representing in the same




Figure 14. This image shows how the agent sees its hand moving: (a) The agent sees its hand at position 9 of its visual field. Then use one of its developed schemas to keep it. (b) The action taken, a movement of its head to the right and down, makes the hand to be placed in the center of its visual field. (c) The agent moves its hand to the right, now placing it in position 6 of his visual field. Then use one of their developed schemas to preserve it. (d) The executed action, a movement of its head to the right, causes the hand to be placed back into the center of its visual field, where it pays attention to it again.

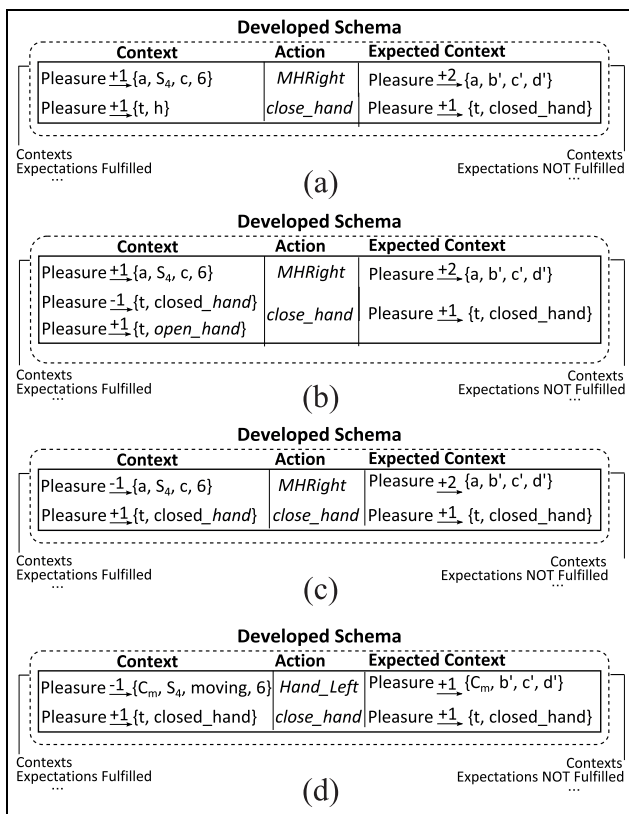


Figure 15. Schemas created when the agent was able to both see and touch its world, which helped it maintain and recover objects seen and probably touched in position 6 of its visual field: (a) preserve the visual object and the tactile object, (b) preserving the visual object while recovering the tactile object, (c) recovering the visual object (with a movement of the head) while maintaining the tactile object, and (d) recovering the visual object (with a movement of the hand) while maintaining the tactile object.

Current-Context both visual and tactile data, and (2) finding partial matches with the stabilized structures. These new capacities opened room for the creation of 32 new schemas, 4 per each position of the peripheral areas within the visual field. Figure 15 shows the four

schemas developed corresponding to position 6. The remaining schemas (the other 28) only differ in terms of position and associated actions needed to maintain and/or recover the object(s) that triggered interest (as the actions to be performed depend on the position of the visual field in which the object is seen, as illustrated in Figure 16).

The first set of schemas, shown in Figure 16(a), were created as follows. The moment when the agent saw an object that caught its attention while sensing that it was touching something with its hand (whether open or closed), a *Current-Context* was created, which included two affective responses of pleasure: one triggered by the visual element and the other triggered by the tactile object (e.g. as exemplified in Figure 17). This context represented a new situation that the agent had never faced, for which it had no schema to determine how to act. This led the agent to try to adapt its current knowledge and perception of the world to face with this new situation by finding a partial match with two schemas: one to preserve the visual object and another to preserve the tactile item (as exemplified in Figure 17). In Piaget's terminology, this means that the agent felt inclined to simultaneously maintain what it was seeing and what it was touching. When the expectations associated with both schemas were fulfilled, an emotional reaction of surprise was generated, resulting in a new structure (see Figure 17), which represents the knowledge about how to maintain simultaneously the pleasant objects that the agent is seeing and touching. For instance, when the agent saw an object in position 6 within its visual field while touching something in the palm of its hand, the schema shown in Figure 15(a) would indicate that the agent could maintain both objects by simultaneously moving the head to the right while closing its hand. In the event that the element being touched was the same as the object being seen, then the use of these schemas led to have the agent holding the object in its hand while centering it within the agent's visual field (by moving the head).

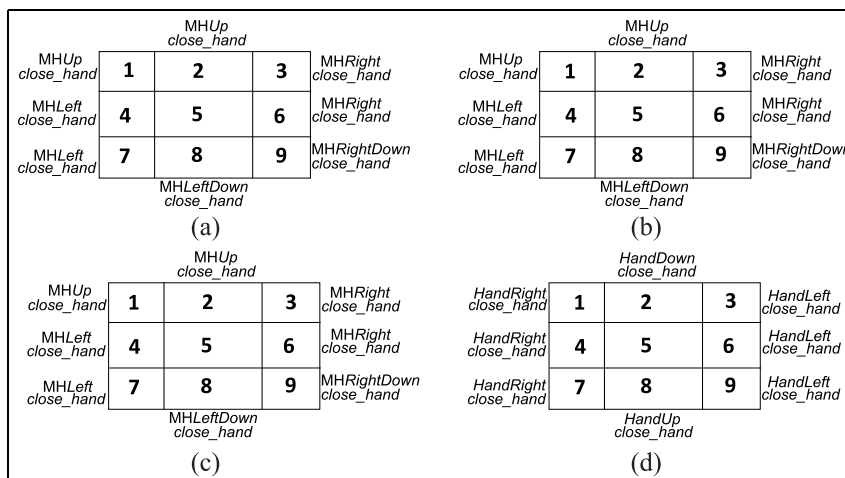


Figure 16. Set of schemas created when the agent was able to both see and touch its world, for all the positions within the visual field: (a) preserving both the visual and the tactile object, (b) preserving the visual object while recovering the tactile object, (c) recovering the visual object (with a movement of the head) while preserving the tactile object, and (d) recovering the tactile object (with a movement of the hand) while preserving the visual object.

The second and third set of schemas shown in Figures 15(b), 15(c), 16(b), and 16(c) were created similarly to the case described above. That is to say, they were constructed as a result of the agent facing an unknown situation, thus responding by making a partial match with two of its structures available in memory, as illustrated in Figures 18 and 19. However, each of them represented different behaviors. On one hand, the second set of schemas contains the knowledge about how to maintain the visual object liked while recovering the tactile object. For example, if the agent saw an object in position 6 within its visual field while letting go the object in the palm of its hand, the schema shown in Figure 15(b) indicated that the agent could again feel pleasure with both objects by simultaneously moving the head to the right while closing its hand (provided that the agent could continue to sense the desired element with the hand open). In the event, the object that was being touched by the agent was the same that was being seen by it, and then the use of these schemas caused the agent to be seen as if it was releasing and taking again the object of its interest. On the other hand, the third set of schemas represents the knowledge about how to recover the pleasant visual object while maintaining the tactile object. For example, if the agent saw its hand grabbing an object in position 6 and during that cycle, the agent chose to randomly move the hand out of its visual field unwantedly, then the use of these schemas made us see the agent pulling its hand out of and then back to the its visual field (by moving its head) while holding an object.

The creation of the fourth set of schemas is a very interesting case, since, unlike the other ones, these were constructed as a result of a partial match between only one schema in memory, allowing the agent to

find an alternative way to recover visual objects having a feature in particular: They were of a shade of blue called C_m . The way in which these arouse is illustrated in Figure 20, where we can see the moment when the agent lost a pleasant object of C_m color in position 6 while feeling pleasure for an object held in its hand, and the agent chose to use only one of the schemas whose expectation was maintaining the tactile element. During that same cycle, the agent randomly moved the hand (to the left) in such a manner that it “accidentally” caused the C_m color to come back into the agent’s visual field. At that moment, a new general schema was created, indicating the agent that any object lost from position 6 could be recovered, while maintaining the object held in hand, if the agent moved its hand to the left and closed it. In time, this schema was differentiated into one indicating that this was only possible when the object lost was of C_m color, the same color of agent’s hand. Similarly, the agent learned that there was a blur in particular (of C_m color) which could be controlled (recovered and preserved) by moving its hand. This resulted in a functional differentiation between the spot representing the hand and any other visual element.

5.3.4. Summary. Briefly, the behaviors learned when the agent was able to both see and touch its world were as follows:

- Visually following its hand by moving its head.
- Centering within its visual field (by moving its head), the object being held.
- Seeing within its visual field how its hand grabs and releases the object of interest.

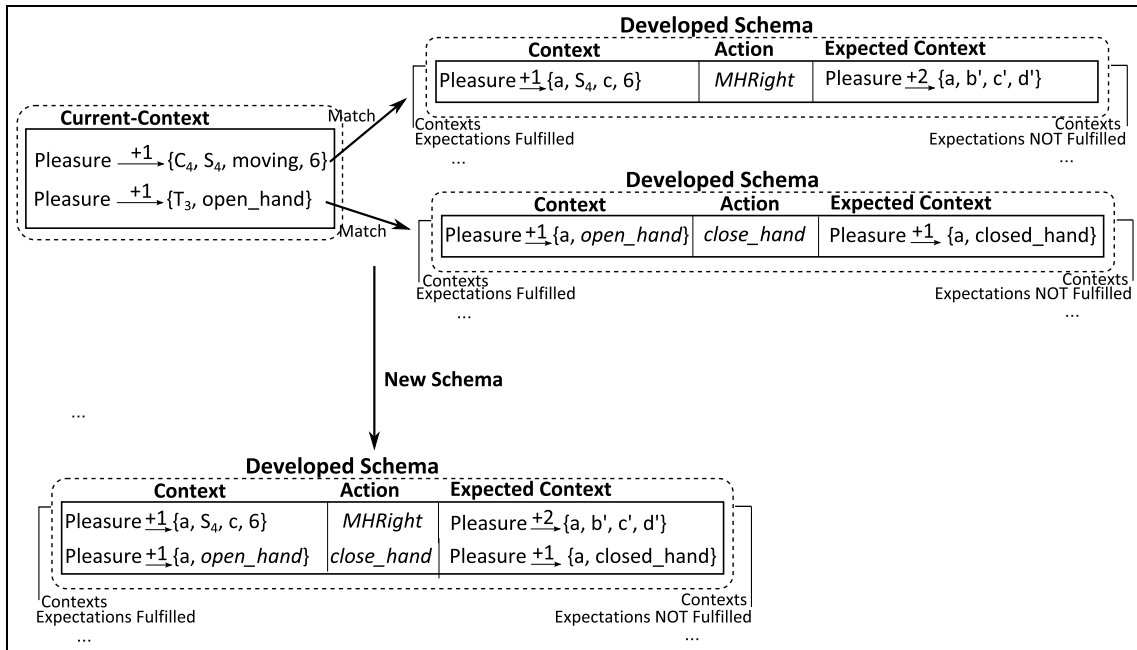


Figure 17. Exemplifies the creation of the first set of schemas developed involving both pleasure for what it sees as for what it touches.

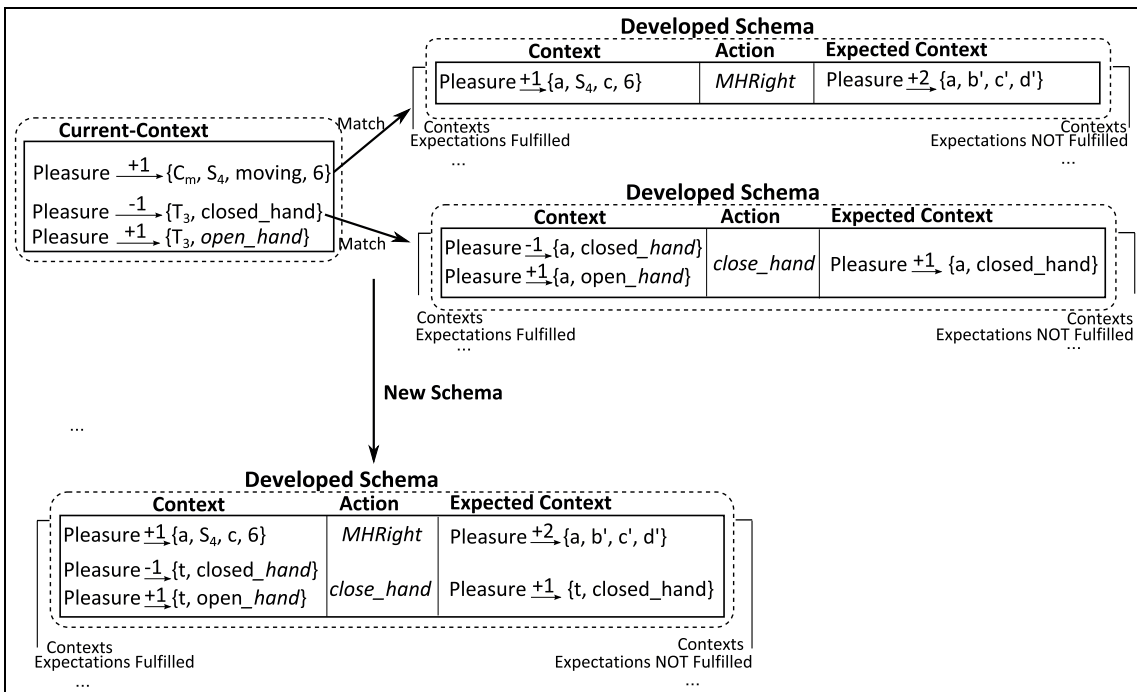


Figure 18. Exemplifies the creation of the second set of schemas developed involving both pleasure for what it sees as for what it touches.

- Seeing how its hand grabbing an object goes out of its visual field, and then recovering that image by moving its head.
- Seeing how its hand grabbing an object goes out of its visual field, and then recovering that image by moving its hand.




Figure 19. Exemplifies the creation of the third set of schemas developed involving both pleasure for what it sees as for what it touches.

5.3.5. Development from basic schemas. In this section, we present some preliminary, but interesting, results about how visual and tactile processes affect each other when the agent (1) can see and touch the world, (2) is initialized only with the three basic schemas of Figure 7, and (3) can attend one visual object and one tactile object at the same time, since the beginning of the execution.

In this experiment, visual and tactile schemas developed independently and in parallel, until the agent entered a state of cognitive equilibrium for the first time. As a result, the agent developed the following new behaviors:

- Recovering pleasant objects that went out by any area of the periphery of its visual field, by moving the head toward the right direction.
- Recovering the visual image of its hand, by moving the hand itself toward the correct direction according to the position where it was last seen.
- Closing the hand when it was touching an object, while it was seeing something (not necessarily the same object) in any area of its field of vision.

After the first stabilization, visual and tactile processes continue to develop independently to acquire skills about how to preserve pleasant visual and tactile stimuli (through partial matches performed with the previously stabilized schemas), but also they started to cooperate with each other.

Skills acquired independently of each modality are as follows:

- Visually following and centering in the visual field the objects of interest.
- Centering its hand in its visual field by moving the hand.

Once previous schemas were created, the agent started to use them in a cooperative way, in sequence one after the other, showing the following behavior:

1. Centering an object of interest in the visual field by moving the head.
2. Centering also its hand by moving the hand. This resulted in a situation where both objects were in the center of the visual field (sometimes the hand occluded the ball), and most of the times these were in contact with each other.
3. Closing the hand to hold on the object that it was seeing and touching in the center of its visual field.
4. This line of development took the agent approximately twice as long as in the previous experiments. We are working in creating strategies to reduce the time.

6. Discussion

According to Piaget, the hand (together with the mouth, the eye, and the ear) is one of the most crucial instruments used by intelligence. Its core task is grasping, which is developed following five stages—which do not correspond to defined ages, but whose sequence is fundamental, except for stage number 3 (Piaget, 1936/

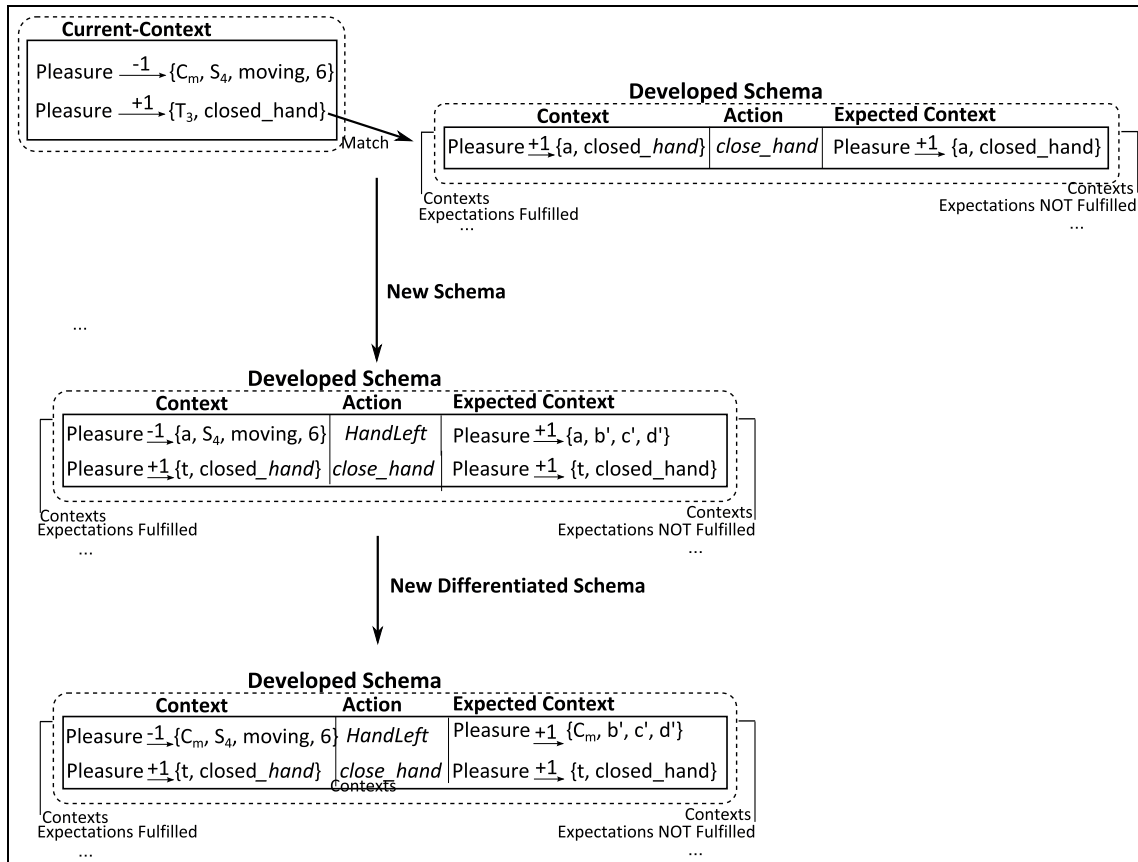


Figure 20. Exemplifies the creation of the fourth set of schemas developed involving both pleasure for what it sees as for what it touches.

1952). In this section, we shall discuss the results obtained within the context of said theory.

6.1. Emergence of the first eye–hand coordination

6.1.1. Stage 1: impulsive movements and pure reflex. During this first stage, the newborn closes the hand upon feeling pressure on the palm, as a result of the grasping reflex that we are all born with. The newborn also moves by impulse both arms, hands, and fingers. In our agent, we observed that, at the beginning of the executions, it kept on randomly moving its hand and closing it whenever its tactile sensor detected the presence of an object. The action of closing its hand was caused by the use of the basic schema representing the grasping reflex. This behavior was noticed for a short time (between 36 and 165 cycles) until the agent faced for the first time a situation wherein, while holding an object, opened the hand automatically and the object, now released, remained in contact with the hand. At that moment, a *Current-Context* was created, describing a new situation: displeasure due to having lost the object that was being held but also pleasure as the open hand was still in contact with said item. That is, the agent was facing opposing emotions (see Figure

9). In this situation, during *Engagement*, the agent had the option to form a partial match between the *Current-Context* with any of the three basic schemas available in its memory, as shown in Figure 21. The result of (1) selecting the second basic schema and (2) choosing the random action of closing the hand is the formation of the new schema illustrated in Figure 9. The same occurs when it is selected the third basic schema. With the creation of this new structure, the agent has learned to recover tangible objects which it is interested in. This situation leads the agent to move on to the second stage of grasp development.

6.1.2. Stage 2: grasping just for grasping and vision is adapted to hand movements. In this second stage, the baby manages to grasp and maintain the objects in its hand without seeing them and without attempting to take them to its mouth. During this period, coordination between vision and general hand movements starts to appear. That is, baby manages to visually follow their own hands, but can only keep them within the visual field by moving the eyes around, not the hands. In other words, vision adapts to hand movements, but a reciprocal action by the hands is not yet certain.

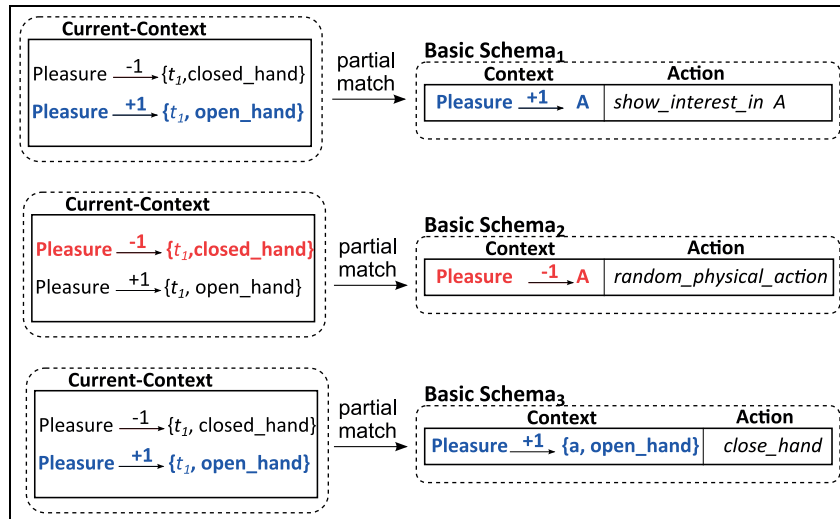


Figure 21. Illustration of the potential partial matches between a *Current-Context* representing opposing emotions and basic schemas. The partial matches are considered so because only one of the two affective responses in the *Current-Context* corresponds to the single affective response available in *Basic Schema₁*, *Basic Schema₂*, and *Basic Schema₃*. Matches are shown in red and blue.

In our agent, we observed that, once it learned how to recover the interesting objects that it released, and once this knowledge became stabilized after repeatedly releasing and recovering different elements, it learned how to keep a hold on them. Additionally, in the third set of experiments, as soon as we activated agent's capacity of sight and touch, it applied all the knowledge acquired from the previous experiments and began to visually follow its hand by moving the head. However, at that point, the hand represented to the agent just a spot, as the rest of the objects in its surroundings.

6.1.3. Stage 3: coordination between pressure and suction, and limitation of hand movements into the visual field. During this third stage, the baby manages to grasp objects and take them to the mouth, as well as taking the elements that its mouth is sucking. Regarding vision, during this stage, babies already exert influence on hand movements. For example, looking at the hand seems to increase its activity, or, on the contrary, it may limit its movements within the visual field. With this step forward in baby's development, it can be noticed that, when infant's hand randomly appears within its visual field, the hand tends to remain on sight. This is the first sign of a reciprocal adaptation, with the hand tending to preserve and repeat movements seen by the eye, while the eye tends to look at everything the hand does. In other words, hand tends to assimilate into its own schemas the visual domain, just as the eye assimilates into its own schemas the hand domain.

To this regard, it is quite interesting to notice how our agent assimilated visual data into its tactile schemas, and vice versa. We can see this in the structures shown

in Figure 15, where we notice that the contexts are now formed by emotional reactions of pleasure and disliking, both toward visual and tactile elements, with the associated actions including head and hand movements, and the same occurs with the expected context. The formation of these new structures led the agent to learn how to center within its visual field (by moving its head), the object being held; watch, in the center of the visual field, how its own hand releases and then recovers the object of its interest; watch how its own hand goes out of the visual field grasping an object and then coming back into the visual field by moving the head; and watch how this last behavior can be performed by moving the hand instead of the head. The use of this set of new behaviors by our agent allowed us to watch it interact with the objects grasped, most of the times within the center of the visual field. Moreover, by creating the last schema, a functional differentiation between the spot represented by the hand and any other one began to manifest. This means that the agent learned that there was a blur, having very particular visual characteristics, which could be controlled by moving its hand.

Agent's development came that far. During the fourth stage, babies get the ability to take an object when they simultaneously see their hands and the desired object; while during the fifth and last stage, they acquire the ability to take any seen object without limitations related to the position of the hand (it can even be outside of their visual field). We believe that our agent may achieve the last two stages of the development of eye–hand coordination, if we grant it capacities of differentiation between means and ends and goal-oriented behaviors. This shall constitute part of our subsequent efforts.

6.2. Developmental path

The developmental path of new behaviors is essential in Piaget’s theory. For this reason, in this section, we shall discuss the path followed by all the behaviors learned by the agent.

In the first and second set of experiments, wherein the agent could only see or touch the world around it, we could see that the earliest skills it acquired were the result (1) of having lost the objects of interest, (2) of their “accidental” recovery (by using one of the basic knowledge structures available to the agent), and (3) of the generalization/differentiation of such experiences. Accordingly, the agent first learned how to recover into its visual field the items that had gone out thereof, and to recover the elements that were held and then released. The developmental path of these two new behaviors is illustrated in the top portion of Figure 22.

Later, our agent entered for the first time into a state of cognitive equilibrium, thus triggering its capacity of finding partial matches with the stabilized schemas. In other words, the agent became capable of using its built knowledge about how to recover interesting objects to face unknown situations. The foregoing resulted in the acquisition of two new behaviors: (1) visually following the interesting objects and (2) holding an object that

was touched by the palm of the hand (see the middle portion of Figure 22).

Thereafter, the agent entered for the second time into a state of cognitive equilibrium, during which it learned two new abilities (as a result of making partial matches between its current situation and past experiences): (1) centering into the visual field static objects of its interest and (2) maintaining a hold on an object (see the bottom portion of Figure 22).

Finally, during the third set of experiments, the agent entered once more into a state of equilibrium and remained in that state for several cycles. As a result, its visual and tactile schemas were considered as stable, thus becoming capable of representing, in its *Current-Context*, objects that caught its attention, both visual and tactile. This led the agent to learn four new behaviors (which appeared again as a result of partial matches): (1) centering within its visual field (by moving its head), the object being held; (2) watching, in the center of the visual field, how its own hand releases and then recovers the object of its interest; (3) watching its own hand grasping an object going out of the visual field and then coming back by moving the head; and (4) watching its own hand grasping an object going out of its visual field and then coming back in sight by

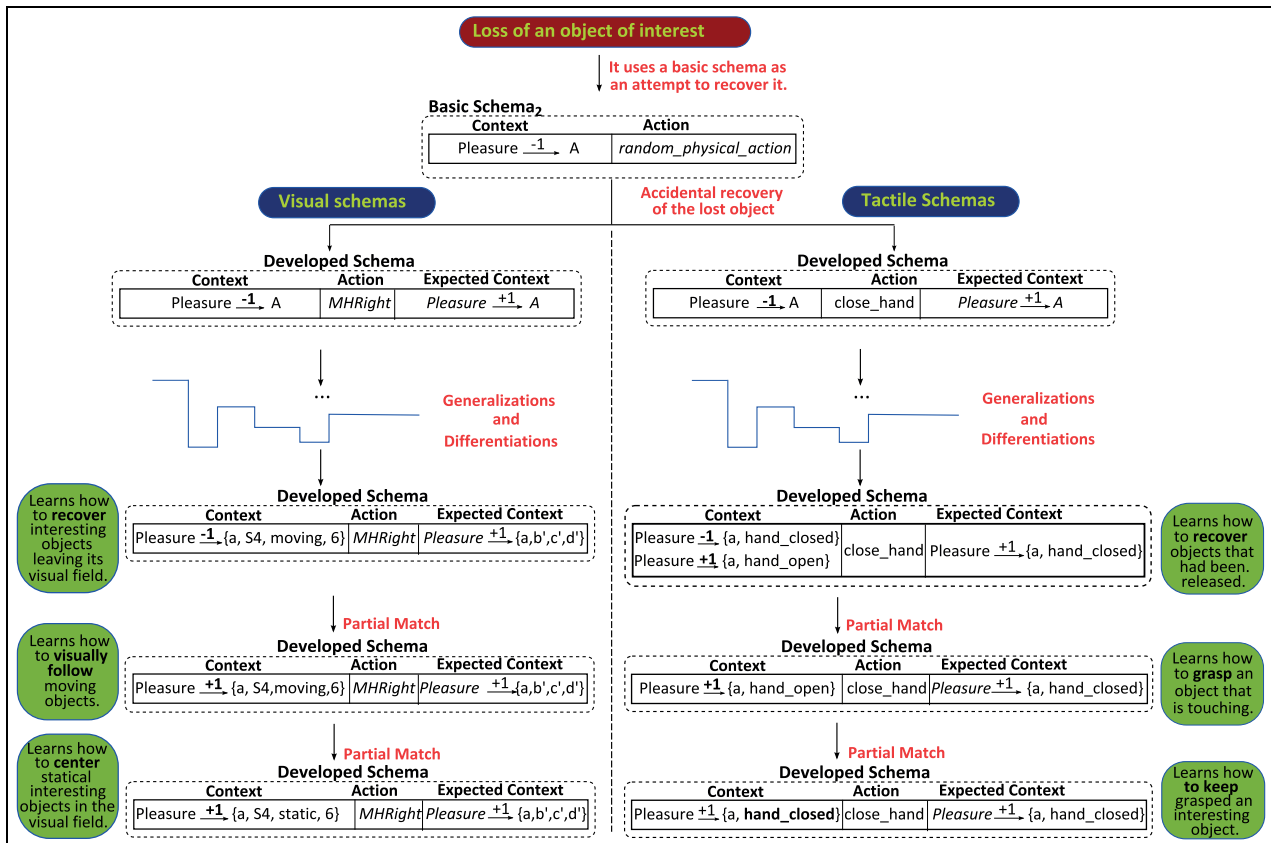


Figure 22. Exemplifies the developmental path of the abilities acquired by the agent.

moving the hand (Figures 17–20 show the origin of these skills).

Thus, it was then possible to trace a developmental path, which allows us to observe how the construction of new behaviors depend on and derive from known experiences.

6.3. Development as a creative process

We are interested in the study of the creative process employing computers. Surprisingly, it is hard to find computational research projects that focus on studying the genesis of the creative process. As far as we know, this is the first computational model that attempts to contribute in that area that we refer to as early-creative behavior. As an interesting characteristic, the ER-Model model that we employ in this work has been originally used to develop an automatic storyteller. So, one can picture a continuous that represents computer models of creativity at different states during the development of skills, where Dev E-R is located in one extreme while our storyteller is located in the opposite extreme. In this way, we can compare both systems. The main purpose of our storyteller is to develop a coherent narrative where conflicts start to rise, until they reach a climax, and then they are sorted out; so, conflicts are essential to progress a plot. It is interesting to notice that in Dev E-R, conflicts are the force that pushes the agent to produce new schemas. Thus, in a sense, in ER-Model, conflicts also arise, reach a climax, and then they decay. In our storyteller, knowledge structures are represented in terms of emotional links and tensions between characters, while in Dev E-R such structures are represented as emotional links between the agent and its environment. This seems to suggest that emotions and conflicts might play an important role in the representation of the agents' world. The capacity of matching partial knowledge structures in memory provides to both, Dev E-R and our storyteller, the opportunity to deal with novel situations. In the same way, there are important differences between both systems: Our storyteller has a sophisticated process of reflection, while in Dev E-R the advancement of reflection is part of the agent's development, and the storyteller represents much more abstract information than Dev E-R (e.g. the general structure of the story).

In Aguilar and Pérez y Pérez (2014) it is presented a set of useful criteria to assess whether the behaviors generated by an agent may be considered as creative. We will discuss each of them to evaluate the results obtained in this article.

6.3.1. Novelty. A behavior is considered novel if it did not exist explicitly in the initial database of knowledge of the agent (Pérez y Pérez, 2014). In the experiments,

the agent was initialized with three basic schemas. By the end of the executions, it had constructed at least 58 new structures that did not exist in the initial database of knowledge of the agent. Hence, under this criterion, all behaviors the agent learned are considered novel.

6.3.2. Utility. A behavior is considered useful if it serves as basis for the construction of new knowledge that gradually leads the agent to acquire new skills that are typical of the following stage of development (cf. Pérez y Pérez, 2014). Thus, the knowledge structures developed by the agent are considered useful, as they allowed it to go from predefined or innate behaviors (typical of the first substage of the sensory-motor period) to body-based behaviors (typical of the second substage of the sensory-motor period) and to behaviors involving external objects (typical of the third substage of the sensory-motor period).

6.3.3. Emergence. Following Steels (1990), a behavior emerges when its origin may not be traced back directly to the components of the system, but rather, it is the result of the way in which such components interact with each other. In the Dev E-R model, the learning of different behaviors depends on a number of factors, notably, (1) environmental properties, (2) physical characteristics of the agent, and (3) current knowledge. In this way, because the new behaviors are not pre-programmed and they are all context-dependent, we may conclude that the behaviors learned by the agent emerged as a result of the way in which the different system components interacted with each other.

6.3.4. Motivations. Amabile and Collins (1999) distinguished two types of creativity: (1) intrinsically motivated and (2) extrinsically motivated. Following them, in this work a behavior developed by an agent is considered creative if it appears as a result of an intrinsic or extrinsic motivation. Our model represents this feature because the emotional reaction of surprise and the intrinsic motivation of cognitive curiosity trigger in the agent the need to modify or construct new schemas.

6.3.5. Adaptation to a new environment. The ability to adapt ourselves to our environment has been traditionally deemed as a condition needed for truly creative behavior (Runco, 2007); similarly, Leonora Cohen describes adaptation as the closest synonym of creativity (Runco, 2007). In Piaget's theory, adaptation is defined in terms of the processes of assimilation and accommodation. We suggest that, in order to consider a behavior developed by an agent as creative, it must acquire such behavior as a result of a process of adaptation to the environment.

The schemas developed by the agent were created as a consequence of it facing unknown situations and

reacting to them: (1) by assimilating the new circumstance to the previously acquired knowledge (through a process of searching the memory for a schema, which represents a situation similar to the one being faced in the *Current-Context*) or (2) by accommodating the knowledge in such a way that it may adjust to the new experience (thus creating a new schema, or differentiating, generalizing, or deleting an existing one). Therefore, we suggest that they are originated as a result of the agent's adaptation to its world.

6.4. Comparison with other methods

In this section, we compare four characteristics of developmental agents: (1) initial state and knowledge representation, (2) motivational components, (3) learning processes, and (4) stage transition mechanisms.

6.4.1. Initial state and knowledge representation. Developmental agents can be classified into two groups: those that start interacting with the environment from *tabula rasa* (i.e. they are initialized with an empty "mind," and therefore all its knowledge is the result of learning through their experiences and sensorial perceptions) and those that start with "innate" knowledge. Examples of the first group include Perotto and Alvares (2006), Modayil and Kuipers (2007), and Mugan and Kuipers (2008). Typically, these agents work with raw data, in contrast to the ones of the second group which usually abstract the world into discrete states and actions (Guerin & McKenzie, 2008). Our agent belongs to the latter.

Among the most representative works that belong to the second category is Drescher's (1991) work and some subsequent systems based on it such as Holmes and Isbell (2005), Guerin and McKenzie (2008), and Lee et al. (2012), which although they report that their system does not start with any explicit innate knowledge, it has the pre-programmed behaviors of palmar reflex and holding the objects until someone removes them or they are dropped by accident. Commonly, these agents define their knowledge structures as schemas that are comprised by three main parts: pre-conditions (named by some authors as context), an action, and post-conditions (named by some authors as result). Contexts and results represent some condition of the world (e.g. visual, tactile, and proprioceptive information). In this way, the schema represents that, when the pre-conditions are satisfied, the agent might perform the action; as a result, the world is modified according to the post-conditions. A typical example of these kinds of structures is shown in Figure 23. Drescher's (1991) schema mechanism and Holmes and Isbell's (2005) improvement of it are initialized with one schema for each action the agent can perform, and with empty contexts and results. This initial "empty" knowledge structures are points of departure for building contentful schemas. The same holds for Guerin and McKenzie's

(2008) agent, which is initialized in a similar way, although four extra schemas that model reflexes (e.g. sucking, grabbing, gazing) are also included. Our agent represents its knowledge as a context, an action, and a result, and it is initialized with one schema that represents the palmar reflex, but, unlike the others, it (1) defines the contexts in terms of more general attributes such as affective responses, emotional reactions, and current motivations, and (2) includes two schemas to model the tendency to preserve and to recover pleasant stimulus. That is, our agent perceives its world in terms of representations of emotions and motivations triggered by the stimuli present in the environment, and not in terms of the particular features of the objects it sees and touches. We believe that these characteristics make our model more flexible. The following lines elaborate this idea. Figure 23 illustrates a schema that shows that when a green object is located in a given position (x,y) , then the agent can perform the action "reach"; the features describing the context are very specific. We wonder why does a green object located on (x,y) trigger such an action? We claim that the decision of what to do depends on more elaborated circumstances that can be (at least partially) represented in terms of affective reactions. If a green (or blue or white, or big or small) object produces curiosity, attraction, need, and so on, then the agent probably attempts to reach it. However, the same object, under different circumstances, might generate fear, anxiety, nervousness, and so on; in such cases, the agent might try to avoid it. Schemas with very specific information, like the one presented in Figure 23, can hardly represent the complexity of the circumstances surrounding an object; more abstract structures, like the ones used by Dev E-R, show more flexibility and adaptability to diverse situations. For instance, our agent can generalize that any item, not just those in green, can be reached by the same action (c.f. Figure 23); in the same way, the capacity of performing partial matches allows dealing with new situations. Some models attempt to sort out this problem by employing a predefined affective reaction. For instance, in the work of Law, Shaw, Earland, Sheldon, and Lee (2014), each time the agent sees a new object, the system triggers novelty. However, this generalization hardly reflects what one observes in the world. Following Pérez y Pérez and Sharples (2001) and Küger et al. (2011), we suggest that models of developmental agents should represent the same context at different levels of abstraction. Specific contexts allow registering the type of information necessary for implementing, for example, a robot's arm, while abstract representations characterize more elaborated contexts. In this way, if we are able to link them, we can exploit the best of both approaches.

6.4.2. Motivational components. Since one of the main characteristics of developmental agents is that they are

Pre-conditions	Action	Post-conditions
Green object at (17.5, 72.4)	Reach to (17.5, 72.4)	Hand at (17.5, 72.4) Green object at (17.5, 72.4) Touch sensation

Figure 23. The typical structure of the schemas defined in Drescher's work and all subsequent models based on it.

Source: Law et al. (2014).

not given any explicit goal or task, then how and why should they perform any action and learn new skills? This is where intrinsic motivations play an important role in these kinds of systems. Some researchers have explored the idea of using novelty as a driver (see, for example, Law et al., 2014), others have used curiosity (see, for example, Oudeyer, Kaplan, & Hatner, 2007), others have explored the idea of having an intrinsic motivator that rewards the discovery of environment affordances (see, for example, Hart & Grupen, 2011), and some others have proposed to use emotions such as happiness, fear, and sadness to shape behavior (see, for example, Ahn & Picard, 2006; Gao & Edelman, 2016). In the case of our agent, conflict is what moves it to act. In particular, this happens (1) when Jacques is in a state of pleasure and then something happens that prevents it to continue in that satisfactory situation, and (2) when its expectations differ from "reality." So, a very important characteristic of our model is that if there is no conflict, there is no learning. As a consequence of this, our agent only builds new schemas when it faces these kinds of situations while interacting with the environment, in contrast to other agents which work, finding all possible contexts and results caused by the execution of all actions (see, for example, Drescher, 1991) or by maximizing a reward such as novelty. Currently, Dev E-R works in a constrained environment where search heuristics and optimization procedures might get some similar results; however, in a more open environment, where an agent can find countless different situations, we believe that our approach will be more effective because of the use of the contextual information to generate and differentiate schemas. In future works, we will run some experiments to test our assumption.

6.4.3. Learning processes. Early sensorimotor learning can be considered as a process that discovers the consequences of actions, as well as the conditions that these consequences depend on. Several learning methods have been proposed to solve this problem. For instance, Drescher (1991) uses marginal attribution; McClelland (1995), Shultz et al. (1995), Parisi and Schlesinger (2002), and Chaput (2004), among others, use neural networks, and some others like Guerin and McKenzie (2008) and Hart and Grupen (2011) use reinforcement learning. In this research, we propose a different

approach which consists in implementing learning (an essential mechanism for adaptation and cognitive development) as a creative process. Using this new approach, our agent learns new skills by (1) using a process of generalization and differentiation of schemas, and (2) using its knowledge of past experiences to deal with new similar, but not identical, situations (i.e. using its knowledge of past experiences to deal with new similar, but not identical, situations).

The first one differs from other approaches in that in our model, new schemas start as a generalization of a sole experience, which then may split into particular schemas, and then they may continue in this process of generalization and differentiation until they stabilize (as summarized in Figure 4). In contrast, works such as Sheldon and Lee (2011) and Guerin and McKenzie (2008) can only modify schemas in one direction, from particular to general. An exception is the Constructivist Anticipatory Learning Mechanism (CALM) system (Perotto et al., 2007), which implements three methods for learning: differentiation, adjustment, and integration. The main differences are that (1) in the CALM system, the processes of differentiation and generalization occur every time an expectation is not met, while in Dev E-R they only occur when the used schema exceeds a certain percentage of failure, giving the agent more flexibility to respond to indeterminacies in the environment; (2) generalization (called adjustment in CALM) replaces, by the undefined symbol #, those properties in the expectation that are different from the perceived ones, while in Dev E-R the replacement is performed based on the properties registered in the two extra structures of successful and failure contexts attached to the schemas, which the CALM system does not have; (3) the generalization process only occurs when a differentiation cannot be performed because all the elements in the context are specified, while in Dev E-R it can happen at any time of the execution; and (4) when a differentiation is performed, the more general schema is always preserved, causing the CALM system to require more memory than Dev E-R. Furthermore, since the CALM system has not been tested in environments and with agents which develop the very first abilities related with vision and touch, it is difficult to make a direct comparison with Dev E-R.

The second learning strategy differs from other works in that most of them only perform a total match between the current situation and the schema's context; and from the ones that do partial matches, due to its very specific knowledge representation, as previously discussed, are limited to discover irrelevant sensor values (e.g. Guerin & McKenzie, 2008). On the contrary, our agent benefits from its more general knowledge representation making it possible, for example, to learn to preserve pleasant objects by using its experience to recover them.

6.4.4. Stage transition mechanisms. Most computational works that try to mimic children development are focused on learning mechanisms within a single developmental stage. Among the most recent and complete projects is that of Law et al. (2014). In their model, they use physical constraints to deal with the complexity of identifying which motor movements cause which effects. These prevent the agent to learn everything at the same time and also help it to bootstrap between stages. For instance, at the start of the experiment, their robot can only move its eyes. Once it achieves control over such a skill, a new competence is activated: first, its capacity to move its head; then, its ability to move its neck; then the shoulder, torso, and so on.

The Dev E-R model uses constraints to let the agent learn new skills of increasing complexity but, unlike the previous example, restrictions are applied to cognitive capacities rather than physical ones. Jacques does this by detecting when it has reached into a state of cognitive equilibrium, which activates the possibility of performing partial matches with its new stabilized schemas. That is, its cognitive capability of using its stable previous experience with similar, not identical, situations is released. So, in Dev E-R, cognitive development emerges as a consequence of the agent going from a state of equilibrium to disequilibrium and back to equilibrium, as Piaget suggests.

We believe that both kinds of constraints, physical and cognitive, can complement each other in such a way that a more complete agent can be created. This is a strategy we think may help to scale our model when more senses are included or when the number of actions increases.

7. Conclusion

Dev E-R is a computational model of early cognitive development, implemented as a creative process. It was inspired by the theories of Jean Piaget (1936/1952) and Leonora Cohen (1989). In a prior paper (Aguilar & Pérez y Pérez, 2015), it was explored its functionality in an artificial agent which could only see (but not touch) the world around it. In this article, our main interest was to study its potential by using it in an agent which could only touch (but not see) the world around it, and in a separate agent which could do both, seeing and touching. The results from the experiments described herein have allowed us to observe the generality of said model, in the sense that it, based on the sensorial capabilities of the agent, was able to learn, on one hand, new behaviors associated with vision and, on the other hand, new behaviors associated with the sense of touch, and finally, the agents showed new behaviors based on both touching and seeing. These latter behaviors represent the first eye–hand coordination skills identified by Piaget. Moreover, the developmental path followed by

the agent matches the one described by the renowned researcher. Hence, Dev E-R is a model which allows us to study relevant aspects of the development of an agent from a novel perspective. This model includes the possibility of observing how the environment and the sensorial capabilities of the agents affect its development, while enabling us to follow, on a step-by-step basis, the construction of its knowledge. The results obtained in this article are quite exciting, although there is still much work to be done.

Acknowledgment

We want to thank Leonardo Sánchez Bojorquez, who greatly helped in the programming of some of the experiments.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work has been supported by the National Council of Science and Technology in México (CONACYT), project number 181561.

References

- Aguilar, W., & Pérez y Pérez, R. (2013, March 25–27). *A computer model of a developmental agent to support creative-like behavior*. Paper presented at the AAAI Spring Symposium: Creativity and (Early) Cognitive Development: A Perspective from Artificial Creativity, Developmental AI and Robotics, Stanford, CA.
- Aguilar, W., & Pérez y Pérez, R. (2014, June 10–13). *Criteria for evaluating early creative behavior in computational agents*. Proceedings of the Fifth International Conference on Computational Creativity, Ljubljana, Slovenia.
- Aguilar, W., & Pérez y Pérez, R. (2015). *Dev E-R: A computational model of early cognitive development as a creative process*. *Cognitive Systems Research*, 33, 17–41.
- Ahn, H., & Picard, R. (2006, April 18–21). *Affective cognitive learning and decision making: The role of emotions*. Proceedings of the 18th European Meeting on Cybernetics and Systems Research, Vienna, Austria.
- Amabile, T. M., & Collins, M. A. (1999). Motivation and creativity. In R. J. Sternberg (Ed.), *Handbook of creativity* (pp. 197–312). Cambridge, UK: Cambridge University Press.
- Asada, M., Hosoda, K., Kuniyoshi, Y., Ishiguro, H., Inui, T., Yoshikawa, Y., . . . Yoshida, C. (2009). Cognitive developmental robotics: A survey. *IEEE Transaction on Autonomous Mental Development*, 1, 12–34.
- Asada, M., MacDorman, K., Ishiguro, H., & Kuniyoshi, Y. (2001). Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robotics and Autonomous Systems*, 37, 185–193.

- Bickhard, M. H. (2013, March 25–27). *The creativity of development and the development of creativity*. Paper presented at the AAAI Spring Symposium: Creativity and (Early) Cognitive Development: A Perspective from Artificial Creativity, Developmental AI and Robotics, Stanford, CA.
- Boden, M. A. (1978). *Artificial intelligence and Piagetian theory*. AAAI Fall Symposium on Cognitive Embodied Systems (Technical Report FS-05-05). Palo Alto, CA: AAAI Press.
- Bruno, S. (2013, March 25–27). *Creativity as a necessity for human development*. Paper presented at the AAAI Spring Symposium: Creativity and (Early) Cognitive Development: A Perspective from Artificial Creativity, Developmental AI and Robotics, Stanford, CA.
- Chaput, H. H. (2004). *The constructivist learning architecture: A model of cognitive development for robust autonomous robots* (PhD thesis). AI Laboratory, The University of Texas at Austin, Austin, TX.
- Cohen, L. (1989). A continuum of adaptive creative behaviors. *Creativity Research Journal*, 3, 169–183.
- Drescher, G. L. (1991). *Made-up minds: A constructivist approach to artificial intelligence*. Cambridge, MA: MIT Press.
- Elliott, T., & Shadbolt, N. (2003). Developmental robotics: Manifesto and application. *Philosophical Transaction: Mathematical, Physical and Engineering Sciences*, 361, 2187–2206.
- Gao, Y., & Edelman, S. (2016). Happiness as an intrinsic motivator in reinforcement learning. *Adaptive Behavior*, 24, 292–305.
- Gorney, E. (2007). *Dictionary of creativity: Terms, concepts, theories and findings in creativity research*. Retrieved from <http://creativity.netslova.ru/Adaptation.html>
- Guerin, F. (2011a). Learning like a baby: A survey of artificial intelligence approaches. *The Knowledge Engineering Review*, 26, 209–236.
- Guerin, F. (2011b). Sensorimotor schema(s). In N. Seel (Ed.), *Encyclopedia of the sciences of learning*. Berlin, Germany: Springer Science + Business Media. pp. 3039–3041.
- Guerin, F., & McKenzie, D. (2008, July 30–31). *A Piagetian model of early sensorimotor development*. Proceedings of the Eighth International Conference on Epigenetic Robotics Modeling Cognitive Development in Robotic Systems, Brighton, UK.
- Hart, S., & Grupen, R. (2011). Learning generalizable control programs. *IEEE Transactions on Autonomous Mental Development*, 3, 216–231.
- Holmes, M., & Isbell, C. (2005). Schema learning: Experience-based construction of predictive action models. *Advances in Neural Information Processing Systems*, 17, 585–592.
- Indurkha, B. (2013, March 25–27). *Thinking like a child: The role of surface similarities in stimulating creativity*. Paper presented at the AAAI Spring Symposium: Creativity and (Early) Cognitive Development: A Perspective from Artificial Creativity, Developmental AI and Robotics, Stanford, CA.
- Jamali, N., Byrnes-Preston, P., Salleh, R., & Sammut, C. (2009, December 2–4). *Texture recognition by tactile sensing*. Australian Conference on Robotics and Automation, Sydney, New South Wales, Australia.
- Küger, N., Geib, C., Piater, J., Petrick, R., Steedman, M., Wörgötter, F., . . . Dillmann, R. (2011). Object-action complexes: Grounded abstractions of sensory-motor processes. *Robotics and Autonomous Systems*, 59, 740–757.
- Law, J., Shaw, P., Earland, K., Sheldon, M., & Lee, M. (2014). A psychology based approach for longitudinal development in cognitive robotics. *Frontiers in Neurobotics*, 8, 1–19.
- Lee, M., Law, J., Shaw, P., & Sheldon, M. (2012, November 7–9). *An infant inspired model of reaching for a humanoid robot*. IEEE International Conference on Development and Learning and Epigenetic Robotics, San Diego, CA.
- Lungarella, M., Metta, G., Pfeifer, R., & Sandini, G. (2003). Developmental robotics: A survey. *Connection Science*, 15, 151–190.
- McClelland, J. L. (1995). A connectionist perspective on knowledge and development. In T. Simon, & G. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling*. Hillsdale, NJ: Lawrence Erlbaum Associates. pp. 157–204.
- Metta, G., Sandini, G., Natale, L., & Panerai, F. (2001, November 22–24). *Development and robotics*. Proceedings of IEEE-RAS International Conference on Humanoid Robots, Tokyo, Japan.
- Miller, M. S. P. (2013, March 25–27). *The neural proposition: Structures for cognitive systems*. Paper presented at the AAAI Spring Symposium: Creativity and (Early) Cognitive Development: A Perspective from Artificial Creativity, Developmental AI and Robotics, Stanford, CA.
- Modayil, J., & Kuipers, B. (2007, July 22–26). *Autonomous development of a grounded object ontology by a learning robot*. Twenty-Second AAAI Conference on Artificial Intelligence, Vancouver, British Columbia, Canada.
- Mugan, J., & Kuipers, B. (2008, June 24–26). Continuous-domain reinforcement learning using a learned qualitative state representation. 22nd International Workshop on Qualitative Reasoning, Boulder, CO.
- Ormrod, J. E. (2012). *Essentials of educational psychology: Big ideas to guide effective teaching*. Boston, MA: Pearson Education, Inc.
- Oudeyer, P. Y., Kaplan, F., & Hatner, V. (2007). Intrinsic motivation system for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11, 265–286.
- Papert, S. (1963). Etude compare de l'intelligence chez l'enfant et chez le robot [Comparative study of intelligence in children and in robots]. In L. Apostel, J. Grize, S. Papert, & J. Piaget (Eds.), *La filiation des structures [The filiation of structures]* (pp. 131–194). Paris, France: Presses Universitaires de France.
- Parisi, D., & Schlesinger, M. (2002). Artificial life and Piaget. *Cognitive development*, 17, 1301–1321.
- Pérez y Pérez, R. (2007). Employing emotions to drive plot generation in a computer-based storyteller. *Cognitive Systems Research*, 8, 89–109.
- Pérez y Pérez, R. (2014, June 10–13). *The three layers evaluation model for computer-generated plots*. Proceedings of the Fifth International Conference on Computational Creativity, Ljubljana, Slovenia.
- Pérez y Pérez, R., & Sharples, M. (2001). MEXICA: A computer model of a cognitive account of creative writing.

- Journal of Experimental and Theoretical Artificial Intelligence*, 13, 119–139.
- Pérez, y, Pérez, R., & Sharples, M. (2004). Three computer-based models of storytelling: BRUTUS, MINSTREL and MEXICA. *Knowledge Based Systems Journal*, 17, 15–29.
- Perotto, F. S., Buisson, J., & Alvares, L. (2007, November 4–7). *Constructivist Anticipatory Learning Mechanism (CALM): Dealing with partially deterministic and partially observable environments*. Proceedings of the Seventh International Conference on Epigenetic Robotics, Piscataway, NJ.
- Perotto, F. S., & Alvares, O. L. (2006, May 8–12). *Learning regularities with a constructivist agent*. Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, Hakodate, Japan. New York, NY: ACM.
- Piaget, J. (1924). *Judgment and reasoning in the child*. London, England: Routledge & Kegan Paul.
- Piaget, J. (1952). *The origins of intelligence in children* (M. Cook, Trans.). London, England: Routledge & Kegan Paul. (French version published in 1936)
- Piaget, J. (1954). *The construction of reality in the child* (M. Cook, Trans.). New York, NY: Basic Books.
- Piaget, J. (1981). *Intelligence and affectivity: Their relationship during child development* (T. A. Brown, & C. E. Kaegi, Trans. & Ed.). Oxford, UK: Annual Reviews.
- Piaget, J. (1985). *The equilibration of cognitive structures: The central problem of intellectual development*. Chicago, IL: The University of Chicago Press.
- Runco, M. A. (2007). *Creativity theories and themes: Research, development, and practice*. Cambridge, MA: Academic Press.
- Sandini, G., Metta, G., & Konczak, J. (1997, November). *Human sensorimotor development and artificial systems*. Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Intellectual Human Activity Support for Applications, Wakoshi, Japan.
- Sheldon, M., & Lee, M. (2011, August 24–27). *PSchema: A developmental schema learning framework for embodied agents*. IEEE International Conference on Development and Learning, Frankfurt am Main, Germany.
- Shultz, T. R., Schmidt, W. C., Buckingham, D., & Mareschal, D. (1995). Modeling cognitive development with a generative connectionist algorithm. In T. Simon, & G. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling*. Hillsdale, NJ: Lawrence Erlbaum Associates. pp. 205–261.
- Steels, L. (1990). *Towards a theory of emergent functionality*. From Animals to Animats: First International Conference on Simulation of Adaptive Behaviour, Paris, France, 24–28 September, pp. 451–461. Cambridge, MA: Bradford Books.
- Stojanov, G. (2009, January 1). *History of usage of Piaget's theory of cognitive development in AI and robotics: A look backwards for a step forwards*. Proceedings of the Ninth International Conference on Epigenetic Robotics, Venice, Italy.
- Stojanov, G., & Indurkha, B. (2013, March 25–27). Creativity and cognitive development: The role of perceptual similarity and analogy. Paper presented at the AAAI Spring Symposium: Creativity and (Early) Cognitive Development: A Perspective from Artificial Creativity, Developmental AI and Robotics, Stanford, CA.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460.
- Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., & Thelen, E. (2001). Autonomous mental development by robots and animals. *Science*, 291, 599–600.

About the Authors



Wendy Aguilar studied a PhD, MSc, and BSc in Computer Science at the National Autonomous University of Mexico (UNAM). Her research interests include computational creativity, developmental robotics, and computer vision. She has collaborated with the Robot Vision Group at the University of Alicante, Spain. Currently, she works as a Research Assistant at the Department of Computer Science at UNAM. She also collaborates as a tutor and lecturer in the Master Program in Computer Science at UNAM.



Rafael Pérez y Pérez studied BSc in Electronics and Computers at Universidad Iberoamericana in México City, MSc in Knowledge-Based Systems, and DPhil in Artificial Intelligence at Sussex University. His research has focused in computer models of creativity. In 2006, he founded the Interdisciplinary Group on Computational Creativity. He is a full-time Professor at the Autonomous Metropolitan University at Cuajimalpa (Universidad Autónoma Metropolitana, Cuajimalpa) in México City. Currently, he is the Chair of the Association for Computational Creativity.