A Model and simulation of Early-Stage Vision as a Developmental Sensorimotor Process

Olivier L. Georgeon, Mark Cohen, Amélie Cordier

A Model and Simulation of Early-stage Vision as a Developmental Sensorimotor Process

Olivier L. Georgeon[1], Mark A. Cohen[2], Amélie V. Cordier[1]

[1] Université de Lyon, CNRS
Université Lyon 1, LIRIS, UMR5205, F-69622, France
[2] Lock Haven University, PA, USA
olivier.georgeon@liris.cnrs.fr, mcohen@lhup.edu, amelie.cordier@liris.cnrs.fr

Abstract. Theories of embodied cognition and active vision suggest that perception is constructed through interaction and becomes meaningful because it is grounded in the agent's activity. We developed a model to illustrate and implement these views. Following its intrinsic motivation, the agent autonomously learns to coordinate its motor actions with the information received from its sensory system. Besides illustrating theories of active vision, this model suggests new ways to implement vision and intrinsic motivation in artificial systems. Specifically, we coupled an autonomous sequence learning mechanism with a visual system. To connect vision with sequences, we made the visual system react to movements in the visual field rather than merely transmitting static patterns.

Keywords: Cognitive development; Intrinsic Motivation; Artificial Intelligence; Cognitive Science; Intelligent agents; Machine learning; Computer simulation.

1 Introduction

We address the question of how autonomous agents can learn *sensorimotor contingencies*—contingencies between the agent's motor actions and the signal received from the sensors. We propose a model that learns such contingencies in rudimentary settings. The agent has primitive possibilities of interaction in a two-dimensional grid, and distal sensors that reflect some remote properties of the grid. This novel learning process is driven by intrinsic motivations hard coded in the agent, and results in the agent gradually improving its capacity to exploit distal sensory information to orient itself toward targets in the grid.

The idea that visual perception is actively constructed through interaction was proposed by theories of active vision [e.g., 1]. Specifically, O'Regan and Noë [2] proposed the *sensorimotor hypothesis* to vision. They used the metaphor of a submarine controlled from the surface by engineers, but with connections that have been mixed up by some villainous marine monster. They argue that the engineers would have to learn the *contingencies* between the commands they send and the signals they receive. O'Regan and Noë's sensorimotor hypothesis to perception postulates that making sense of perception precisely consists of knowing these contingencies. Following these views, we seek to demonstrate that our agent can autonomously learn such contingencies. From an observer's viewpoint, the agent would seem to learn to make sense of its perceptions with regard to its intrinsic motivations.

To address this problem we rely on our previous work on intrinsically-motivated hierarchical sequence learning [3]. In this previous work, we implemented an original algorithm that learned regularities of interaction. We demonstrated that an agent could use this algorithm to learn sequential contingencies between *touch* interactions and *move* interactions to avoid bumping into obstacles. In this paper, we report an extension of this algorithm that allows the agent to learn contingencies when perception is not a direct feedback from motion. For example, unlike *touch* in our previous agent, vision does not directly result from motion. Yet, because our previous algorithm succeeded in learning "touch/motor" contingencies, we expect it to prove useful for learning "visio/motor" less direct contingencies. Specifically, we envision coupling the algorithm with a complementary sensory mechanism as suggested by theories of *dual process* [e.g., 4].

More broadly, this work seeks to model and simulate *ab-nihilo* autonomous learning, sometimes referred to as *bootstrapping cognition* [5]. We relate this developmental approach to Piaget's [6] notion of an *early stage* in human's ontological development (pre-symbolic). For this work, though, this early-stage notion can also fit the framework of phylogenetic evolution of animal cognition, as discussed for example by Sun [7].

Because this work focuses on autonomous cognitive development, we do not follow traditional methods of task modeling that use either a symbolic cognitive modeling approach [8] or a sub-symbolic approach (e.g., [9]). Instead, we situate ourselves in the area of studies on *intrinsic motivation* [10, 11] that seek to implement agents whose behavior organizes autonomously with no pre-assumed goal.

2 Background

This work implements an *intrinsically-motivated schema mechanism* [3, 12] coupled with a distal sensory system (Figure 1). In essence, the intrinsically-motivated schema mechanism implements Piaget's [6] views that perception and action should be kept embedded in sensorimotor schemas rather than separated in the traditional perception-cognition-action loop. We use the terms schemas, interaction patterns, and sensorimotor behaviors interchangeably, and we represent them as small white rectangles (primitive schemas) and gray rectangles (learned composite schemas) in Figure 1.

The mechanism records past sequences of interaction as hierarchically organized interaction patterns (episodic memory in Figure 1). The mechanism also maintains a representation of the agent's current situation (in what we have labeled working memory). The representation of the current situation takes the form of interaction patterns, in compliance with Gibson's [13] ecological theory, according to which

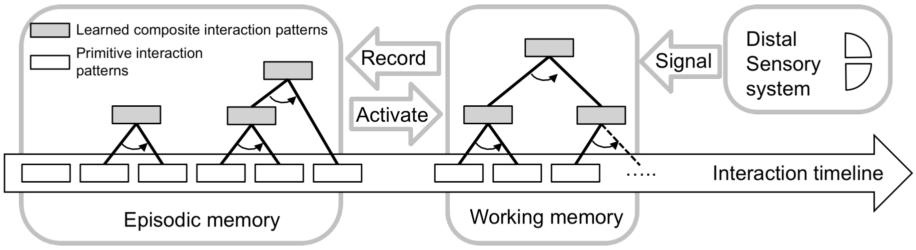

Figure 1. The intrinsically-motivated schema mechanism coupled with the distal sensory system.

agents represent the world in the form of *affordances*, i.e., possibilities of interaction. The consistency of knowledge representation across the two memory systems makes it possible to match past episodes against the current situational representation to select forthcoming behavior.

Through interaction, the mechanism learns composite behaviors, in a bottom-up way, as the agent experiences their corresponding sequences through its activity. Composite behaviors consist of hierarchically-organized subsequences of lower-level behaviors, all the way down to pre-defined primitive behaviors (the sequential relation in the schema hierarchy is represented by arc arrows between the branches of the hierarchy in Figure 1).

This schema mechanism implements intrinsic motivation by associating *satisfaction values* with schemas. The modeler initializes the agent with inborn satisfaction values associated with primitive schemas. The agent then selects newly-learned schemas based on their expected satisfaction in specific contexts, balanced with the expectation of success in such contexts. That is, satisfaction values operate in a proactive way during the selection of forthcoming behavior. In turn, the agent makes a selection of composite schemas based on the satisfaction that effectively results from enacting these schemas.

In this paper, we report an experiment where the agent has six primitive behaviors. Each primitive behavior consists of the association of primitive actuators with binary feedback. These six primitive interaction patterns are listed in the first six lines of Table 1. Similar to our previous work [3], the binary feedback corresponds to a

Table 1. Primitive actuators and sensors.

Symbols	Actuators	Sensors	Description	Intrinsic satisfaction
^ (^)	Turn left	True	Turn 90° left toward adjacent empty square	0 (indifferent)
[^]		False	Turn 90° left toward adjacent wall	-5 (dislike)
> (>)	Forward	True	Move forward	0 (indifferent)
[>]		False	Bump wall	-8 (dislike)
v (v)	Turn right	True	Turn 90° right toward adjacent empty square	0 (indifferent)
[v]		False	Turn 90° right toward adjacent wall	-5 (dislike)
*		Appear	Target appears in distal sensor field	15 (love)
+		Closer	Target approaches in distal sensor field	10 (enjoy)
x		Reached	Target reached according to distal sensor	15 (love)
o		Disappear	Target disappears from distal sensor field	-15 (hate)

proximal sense that can be thought of as *touch*. In our experiment, if the agent tries to move forward, he can either succeed and touch no wall, which leaves him indifferent; or bump into a wall, which he dislikes. When the agent turns, he receives *tactile* information about the adjacent square that he turned towards (touched a wall: dislike, or not touched: indifferent).

3 The Model

Now that we have presented the intrinsically motivated schema mechanism, we can examine the issue of implementing an agent that learns to coordinate actions with sensory inputs, when both are not *a priori* connected together within sensorimotor schemas.

As a start, we have implemented the rudimentary distal sensory system depicted in Figure 2. This system consists of two *eyes* that detect the blue color in the environment (*target* in Figure 2). Each eye has a *visual field* that covers 90° of the agent's surrounding environment. The two visual fields overlap in the line straight in front of the agent, including the agent's location. Each eye generates a single integer value that indicates the amount of blue color detected, also reflecting the distance to a blue square if there is only one. This distal sensory system can thus be seen as a rudimentary monochromic visual system with a resolution of two pixels.

We want the visual system to forward visual information to the agent's situational representation in working memory to inform the selection of behavior. Inspired by the dual process argument, we first considered an iconic approach that would incorporate the two-pixel icon provided by the visual system as an element of context in our existing schemas. This approach proved inefficient because it added too much random combination in the schema mechanism and the agent was unable to learn the contingency between the perceived icons and the actions. Moreover, the combinatorial growth would be prohibitive with wider icons.

To better support contingency learning, we looked for inspiration from studies of biological primitive organisms. We found useful insights from the limulus (horseshoe crab), an archaic arthropod whose visual system has been extensively studied [14]. Specifically, we retained the two following principles:

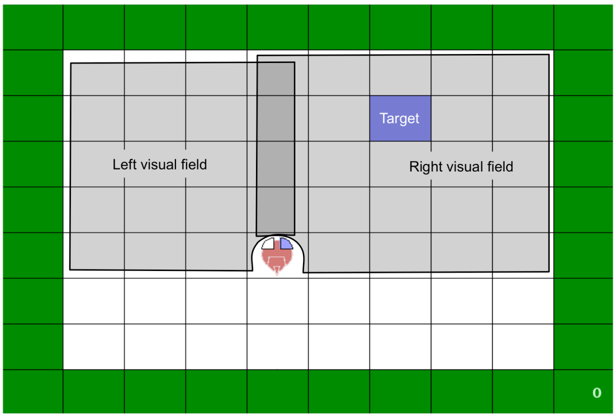

Figure 2. The agent's vision of its environment.

1) Sensibility to movement: the signal sent to the brain does not reflect static shape recognition but rather reflects changes in the visual field. A horseshoe crab's "*eye is highly sensitive to images of crab-size objects moving within the animal's visual range at about the speed of a horseshoe crab (15 cm/s)*" [15, p. 172].

2) Visio-spatial behavioral proclivity: male horseshoe crabs move toward females when they see them with their compound eyes, whereas females move away from other females.

Similar to our agent, horseshoe crabs' eyes are fixed to their body and have poor resolution (roughly 40*25 pixels).

From these insights, we modeled the visual system so that it updated the schema mechanism only when a change occurred in the visual field. We identified four different signals that each eye would generate: *appear*, *closer*, *reached*, and *disappear*. Additionally, to generate a visio-spatial behavioral proclivity, the schema mechanism receives an additional intrinsic satisfaction associated with each of these signals. These four signals are listed with their satisfaction values in the last four lines of Table 1. For example, an eye sends the signal *closer* when the amount of blue color has increased in this eye's visual field over the last interaction cycle, meaning the square has gotten closer (with this regard, our agent's visual acuity is more than two pixels because the agent can detect the enlargement of the target's span in the visual field). We associate the *closer* signal with a positive inborn satisfaction value (10) to generate the proclivity to move toward blue squares.

With these settings (as reported in Table 1), we expect our agent to learn to coordinate its actions with its perception and orient itself toward the blue square. We must note that nothing tells the agent *a priori* that moving would, in some contexts, get it closer, or that turning would shift the blue color in the visual field. These are the kind of contingencies that the agent will have to learn through experience.

After an initial learning phase, we expect to see different behaviors emerge. One possible behavior is the *diagonal strategy* depicted in the left side of Figure 3. This behavior consists of alternatively moving forward and turning toward the blue square until the agent becomes aligned with the blue square. At this point, the agent will continue to move forward.

Another possible behavior is the *tangential strategy* depicted in the right side of Figure 3. The tangential strategy consists of approaching the blue square in a straight line. The trick with the tangential strategy is that the agent cannot accurately predict when he should turn toward the blue square before he passed it. The tangential strategy thus consists of moving on a straight line until the blue square disappears from the visual field, then returning one step backward, and then turning toward the blue square.

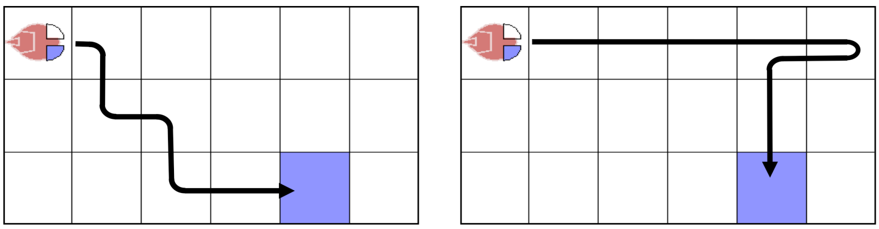

Figure 3. Two possible strategies: *diagonal strategy* (left) and *tangential strategy* (right).

Of course, such specific strategies have probably little to do with real horseshoe crabs. These strategies would only arise due to the coupling of our environment's topological structure with our agent's sensorimotor system, intrinsic motivations, and learning skills.

4 The Experiment

We use the *vacuum environment* implemented by Cohen [16] as an experimental test bed for our agent. This environment is based on Russell and Norvig's [17] general vacuum cleaner environment. Figure 3 shows the agent in this environment. Filled squares around the grid are walls that the agent will bump into if he tries to walk through them. The agent's eyes are represented by quarters of circles that turn a blue color when they detect a blue square; the closer the blue square, the more vivid the eye's color. When both eyes send a *reached* signal, this signal triggers an additional systematic *eating* behavior (with no additional satisfaction value) that makes the agent "eat" the blue square and makes the blue square disappear. The observer can thus see the agent's behavior as a quest for food. The observer can click on the grid to insert a new blue square when the agent has eaten the previous one.

We provide online videos of different runs of the agent . At the beginning, these videos show the agent acting frantically because it has not yet learned the contingency between its actions and its perceptions. The agent picks random behaviors when it has no knowledge of what to do in a specific context. It learns to categorize contexts in terms of possibilities of behavior, in parallel with learning interesting composite behaviors. After the initial pseudo random activity, the videos show the agent more often orienting itself toward the blue square. After eating one or two blue squares, the agent starts to stick to a specific strategy. Our website shows example videos where the agent learned the diagonal strategy[1] and where it learned the tangential strategy[2]. Traces of these runs are reported in Table 2 and Table 3, and discussed below to provide a deeper view of the learning process.

In the traces, the numbers indicate the interaction-cycle counter also displayed in the bottom-right corner of the video. The symbols used to represent the agent's primitive behaviors are ^ (turn left), > (try to move forward), v (turn right). The symbols that represent the eyes' signals are * (appear), + (closer), x (arrived), o (disappear). The eye signals are represented on each side of a | character, the left eye signal being on the left and the right eye signal on the right. When none of the eyes sends a signal, the | character is omitted. The entire sensorimotor pattern is displayed within parentheses when a *true* feedback is received and within angle brackets when a *false* feedback is received. For example, the agent bumped a wall on step 8 and 9, and turned toward an adjacent wall on steps 11 and 24. Simultaneously, a blue square appeared in the left eye's visual field on step 10, and disappeared on step 11.

---

[1] http://e-ernest.blogspot.com/2011/01/ernest-82-can-find-his-food.html
[2] http://e-ernest.blogspot.com/2011/01/tengential-strategy.html

Table 2. Trace of diagonal strategy learning.

```
11 2(> |+) 3(v*|o) 4(vo| ) 5(v) 6(^) 7(>) 8[>] 9[>] 10(^*| ) 11[vo| ]
12(^*| ) 13(^o|*) 14(^ |o) 15(>) 16(v |*) 17(v*|o) 18(>+| ) 19(^o|*)
20(v*|o) 21(>+| ) 22(>+| ) 23(>o| ) 24[v] 25[>] 26(^) 27(>) 28[v] 29(v
|*) 30(v*|o) 31(>+| ) 32(^o|*) 33(v*|o) 34(>+| ) 35(^o|*) 36(v*|o)
37(>+| ) 38(^o|*) 39(> |+) 40(^ |o) 41(v |*) 42(> |+) 43(v*| ) 44(^o| )
45(> |o) 46(^) 47(>) 48(v) 49(^) 50(v) 51(v |*) 52(> |+) 53(v*|o) 54(>+|
) 55(^ |*) 56(>+|+) 57(^o| ) 58(v*| ) 59(>x|x) 60(vo|o) 61(v) 62(v)
63(v) 64(v) 65(v) 66(v) 67(v) 68(v) 69(v |*) 70(> |+) 71(v*|o) 72(>+| )
73(^o|*) 74(> |+) 75(v*|o) 76(>+| ) 77(^o|*) 78(> |+) 79(v*|o) 80(>+| )
81(^ |*) 82(>x|x) 83(>o|o) ...
```

The beginning of the trace in Table 2 shows pseudo random behavior. On steps 21 to 23, the agent learned the sequence consisting of getting closer twice, then loosing the detection. On step 36 to 39, it learns a step of the *stairstep progression* necessary for the diagonal strategy: turn-right, move-forward, turn-left, move-forward, with the appropriate visual signals associated. It eats its first blue square on step 59. By clicking on the grid, we inserted the second blue square on step 69. From then on, the agent enacted a perfect diagonal strategy and reached the second blue square on step 82. The video shows that the agent continues enacting the diagonal strategy to reach subsequent blue squares.

Table 3 reports the trace of a run where the agent learned the tangential strategy. In this run, the agent learned the behavior consisting of moving forward and getting closer to the blue square (on the right eye) from the beginning (from step 2 to 6). Due to the positive satisfaction experienced, it sticks to this behavior until it loses the detection on step 7. A partial component of the returning sequence is then discovered on steps 7 to 11, but the agent then gets away from the blue square. This partial returning sequence is then enacted again in steps 47 through 51, and followed by a *move forward* and *getting closer* in both eyes on step 52, but then, the agent moved away again. The agent finally learns to stick to moving forward when in alignment with the blue square during steps 60 to 62. It finally eats its first blue square on step 62. By clicking on the grid, we added the second blue square on step 75. From then on, the agent was able to put together all the subsequences of the tangential strategy that it has learned. It enacts a perfect tangential strategy to reach the second blue square on step 88. The video shows that the agent then keeps enacting the tangential strategy to reach subsequent blue squares.

Table 3. Trace of tangential strategy learning.

```
1 2(> |+) 3(> |+) 4(> |+) 5(> |+) 6(> |+) 7(> |o) 8(v |*) 9(v*|o) 10(>+|
) 11(^ |*) 12(^o| ) 13(^ |o) 14(^*| ) 15(>o| ) 16(>) 17(>) 18(^*| )
19(vo| ) 20(v) 21(^) 22(>) 23(^*| ) 24(^o|*) 25(^ |o) 26(^) 27(v) 28(v
|*) 29(^ |o) 30(^) 31(>) 32(^*| ) 33(^o|*) 34(v*|o) 35(>+| ) 36(^o|*)
37(v*|o) 38(>+| ) 39(^ |*) 40(v |o) 41(>o| ) 42(>) 43(^*| ) 44(^o|*)
45(> |+) 46(> |+) 47(> |o) 48(v |*) 49(v*|o) 50(>+| ) 51(^ |*) 52(>+|+)
53(v |o) 54(vo| ) 55(v |*) 56(v*| ) 57(v |o) 58(^ |*) 59(>+|+) 60(>+|+)
61(>+|+) 62(>x|x) 63(>o|o) 64(v) 65(v) 66(v) 67[v] 68(^) 69[v] 70(v)
71(v) 72(v) 73[v] 74[>] 75(^*| ) 76(^o|*) 77(> |+) 78(> |+) 79(> |+)
80(> |+) 81(> |+) 82(> |o) 83(v |*) 84(v*|o) 85(>+| ) 86(^ |*) 87(>+|+)
88(>x|x) 89(^o|o) ...
```

Other runs show that the agent always learns a strategy within the first hundred steps, and that the most frequently found strategy is the diagonal strategy, with the settings defined in Table 1. The experiment therefore demonstrates that the agent always succeeds in learning sensorimotor contingencies.

The pre-defined primitive satisfaction values impact the frequency of the resulting strategies; for example, a negative satisfaction value for turning would favor the tangential strategy because it incites the agent to minimize turning. In addition, the location where the experimenter introduces new blue squares influences the resulting strategy. For example, placing blue squares next to walls disfavor the tangential strategy because the agent would often bump the wall, which it does not like. The experimenter thus plays a role in training the agent.

5 Results and discussion

This work demonstrates a technique for implementing vision as an embodied process in an artificial agent. In this technique, the visual system does not send static images to the central system, but rather sends signals denoting change in the visual field. Notably, this technique allows the agent to see static objects, because changes in the visual field can result from the agent's own movements. This technique is consistent with studies of horseshoe crabs that show male horseshoe crabs can orient themselves toward static females because their visual system reacts to female-like objects that appear to be moving (relatively to their own speed) in a uniform background (a sandy shallow ocean bottom or beach) [14].

By succeeding in learning the contingencies between actions and sensors, our model illustrates a sensorimotor theory of vision [2]. O'Regan and Noë argue that this approach would also be consistent with complex visual systems, including human vision. This argument rests upon experiments that show that human vision requires eye movements [e.g., 18]. In future studies, we will further investigate this theory by endowing our agent with the capacity to move its eyes independently from its body, and with better eye resolution, so it can learn contingencies between eye saccades and the received visual signals.

This work also demonstrates that our intrinsically motivated schema mechanism can be used as a general framework to model perception and action. The principles of this study can be followed to model other senses, based on the detection of features other than basic changes in the visual field. Moreover, this framework helps model aspects of ontogenetic development of cognition in a novel way because:

a) The agent's goals emerge from low-level drives. The agent does not encode strategies or task procedures defined by the programmer, but rather autonomously constructs a strategy, as opposed to traditional cognitive models [8]. This accounts for theories of cognition that see natural cognitive systems not as being born with an explicit final goal to achieve, but rather pushed by primitive drives [e.g., 6].

b) The agent's instances are capable of "individuating" themselves through their experience, i.e., acquiring their own cognitive individuality that was not encoded in their "genes". This accounts for the role of individual experience in cognitive development.

6 Related Work

Few studies have investigated the sensorimotor hypothesis to vision. Moller & Wolfram [19] proposed a model that learned to distinguish between dead-ends and corridors in visual scenes. Their agent internally simulated its own way through the scene, based on its own experience previously acquired through interaction, to discover a dead-end or a corridor. Moller and Wolfram's objectives, however, differ from our own in that they focused on visual processing of complex images rather than on initial sensorimotor contingency learning. Their work, nevertheless, motivates us to endow our agent with the capacity to simulate courses of action internally to more efficiently distinguish scenes and choose forthcoming behaviors.

Other studies refer more loosely to the sensorimotor hypothesis of vision to categorize images of objects through representing the objects' functions [e.g., 20, 21].

In the domain of modeling the visual system of natural organisms, Barlow and coauthors [15] proposed a computerized model of the horseshoe crab's visual neural system. This model, however, did not simulate the horseshoe crab's behavior, nor did it functionally represent the connection between the animal's perceptions and its behaviors.

7 Conclusion

This work demonstrates that it is possible to implement a sensorimotor theory of vision in rudimentary settings. To do so, we proposed a novel intrinsically motivated schema mechanism coupled with a visual system that detects dynamic features in the visual field. The agent has primitive drives that push it to learn to coordinate motor actions with sensory information to enact behaviors that, in turn, increase the agent's fulfillment of its primitive drives. This work opens the way to more complex models where the eye's resolution will be increased and where the agent will have the capacity to move its eyes independently from its body. Such developments inform our understanding of visual systems in natural organisms and suggest new techniques to implement vision in autonomous robots.

Acknowledgments. This work was supported by the Agence Nationale de la Recherche (ANR) contract RPDOC-2010-IDEAL. We gratefully thank Pr. Alain Mille for his useful comments.

References

1. Findlay, J., Gilchrist, I.: Active Vision: The Psychology of Looking and Seeing. Oxford University Press, USA (2003)
2. O'Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences **24** (2001) 939–1031

3. Georgeon, O.L., Morgan, J.H., Ritter, F.E.: An Algorithm for Self-Motivated Hierarchical Sequence Learning. International Conference on Cognitive Modeling, Philadelphia, PA (2010) 73-78
4. Norman, J.: Two visual systems and two theories of perception: An attempt to reconcile the constructivist and ecological approaches. Behavioral and Brain Sciences **25** (2002) 73-144
5. Dennett, D.: Brainchildren. Essays on designing minds. Penguin Books (1998)
6. Piaget, J.: The construction of reality in the child. Basic Books, New York (1937)
7. Sun, R.: Desiderata for cognitive architectures. Philosophical Psychology **17** (2004) 341-373
8. Newell, A.: Unified Theories of Cognition. Harvard University Press, Cambridge, MA (1990)
9. Brooks, R.: Intelligence without representation. Artificial Intelligence Journal **47** (1991) 139–159
10. Blank, D.S., Kumar, D., Meeden, L., Marshall, J.: Bringing up robot: Fundamental mechanisms for creating a self-motivated, self-organizing architecture. Cybernetics and Systems **32** (2005)
11. Oudeyer, P.-Y., Kaplan, F.: Intrinsic motivation systems for autonomous mental development. IEEE Transactions on Evolutionary Computation **11** (2007) 265-286
12. Georgeon, O.L., Ritter, F.E.: An Intrinsically-Motivated Schema Mechanism to Model and Simulate Emergent Cognition. The Pennsylvania State University, State College, PA (2011)
13. Gibson, J.J.: The ecological approach to visual perception. Houghton-Mifflin, Boston (1979)
14. Shuster, C., Barlow, R.B., Brockmann, J.: The american horseshoe crab. Harvard University Press, Harvard, MA (2004)
15. Barlow, R.B., Hitt, J.M., Dodge, F.A.: Limulus vision in the marine environment. Biological Bulletin **200** (2001) 169-176
16. Cohen, M.A.: Teaching agent programming using custom environments and Jess. AISB Quarterly **120** (2005) 4
17. Russell, S., Norvig, P.: AI: A modern approach. Prentice-Hall, Englewoods Cliffs, NJ (1995)
18. Schlingensiepen, K.-H., Campbell, F.W., Legge, G.E., Walker, T.D.: The importance of eye movements in the analysis of simple patterns. Vision Research **26** (1986) 1111-1117
19. Moller, R., Wolfram, S.: Bootstrapping Cognition from Behavior--A Computerized Thought Experiment. Cognitive Science **32** (2008) 504-542
20. Bogoni, L.: More than just shape: a representation for functionality. Artificial Intelligence in Engineering **12** (1998) 337-354
21. Stark, L., Bowyer, K.: Achieving generalized object recognition through reasoning about association of function to structure. IEEE Transactions on Pattern Analysis and Machine Intelligence **13** (1991) 1097-1104