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

Olivier L. Georgeon¹, Mark A. Cohen², and Amélie V. Cordier¹

¹ Université de Lyon, CNRS

Université Lyon 1, LIRIS, UMR5205, F-69622, France

² Lock Haven University, PA, USA

{olivier.georgeon,amelie.cordier}@liris.cnrs.fr, mcohen@lhup.edu

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 intrinsically motivated schema 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 through 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. The 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 towards targets within the environment.

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* of 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 of perception posits that making sense of perception precisely consists of knowing these contingencies.

To implement these views, we rely on our previous work regarding intrinsically-motivated hierarchical sequence learning [3]. In this previous work, we implemented an original algorithm that learned regularities of interaction. This algorithm was inspired both by constructivist schema mechanisms (e.g. [4, 5]) and by principles of intrinsic motivation [6, 7]. The algorithm implements an innovative way to associate these two notions by incorporating intrinsic motivation into the schema mechanism, with schemas representing hierarchical sequences of interaction. Intrinsic motivation makes it possible to address the scalability issues of traditional schema mechanisms by driving the selection of schemas.

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 to 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., 8].

More broadly, this work seeks to model and simulate *ab-nihilo* autonomous learning, sometimes referred to as *bootstrapping cognition* [9]. We relate this developmental approach to Piaget’s [10] notion of an *early stage* in human’s ontological development (pre-symbolic). For this work, though, this early-stage notion can also fit within the framework of phylogenetic evolution of animal cognition, as discussed for example by Sun [11].

2 The Model

In this model, 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 proximal sense that can be thought of as *touch*. If the agent tries to move forward, he can either succeed and touch no wall,

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)

or fail and bump into a wall. Succeeding in moving forward is associated with the intrinsic satisfaction of 0 (indifference) while bumping is associated with -5 (dislike). When the agent turns, he receives *tactile* information about the adjacent square that he turned towards; touched a wall (-8, dislike), or not touched (0, indifferent).

In addition, we have now implemented the rudimentary distal sensory system depicted in Figure 1a. This system consists of two *eyes* that detect the blue color in the environment (*target* in Figure 1a). 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 monochromatic 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 many random combinations 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 [12, 13]. Specifically, we retained the two following principles:

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)*" [13, 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.

Fig. 1. **a)** The agent's vision of its environment (left). **b)** The diagonal strategy (top-right). **c)** The tangential strategy (bottom-right).

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 Figure 1b. 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 Figure 1c. 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.

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.

3 The Experiment

We use the *vacuum environment* implemented by Cohen [14] as an experimental test bed for our agent. Figure 1.a) 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 quarter-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, resulting in the blue square disappearing. 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 agent¹. 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 has learned the diagonal strategy² and where it has learned the tangential strategy³. The website also provides a detailed analysis of activity traces to discuss the learning process⁴. An interactive demonstration is also available⁵ where the visitor can interact with the agent by adding food on the grid (this online demonstration is based on a slightly different configuration than the experiment reported here; in particular, the agent is in a continuous world rather than a grid, but the underlying algorithm remains the same).

Different 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 learning performance varies with the initial settings and the environmental conditions during training. Our goal was not to optimize the learning performance but to qualitatively demonstrate the nature of this developmental learning mechanism. In particular, 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 [15]. Our model is also 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) [12].

4 Conclusion

This work demonstrates a technique for implementing vision as an embodied process in an intrinsically motivated 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 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

¹ <http://liris.cnrs.fr/ideal/>

² <http://e-ernest.blogspot.com/2011/01/ernest-82-can-find-his-food.html>

³ <http://e-ernest.blogspot.com/2011/01/tengential-strategy.html>

⁴ <http://e-ernest.blogspot.com/2011/01/tangential-strategy-details.html>

⁵ <http://liris.cnrs.fr/ideal/demo/ernest83/Seca.html>

natural organisms and suggest new techniques to implement vision in intrinsically motivated robots.

Acknowledgments. This work was supported by the Agence Nationale de la Recherche (ANR) contract RPDOC-2010-IDEAL. We gratefully thank Pr. Alain Mille and Jonathan H. Morgan for their useful comments, and Olivier Voisin for his implementation of the online demonstrations.

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, 939–1031 (2001)
3. Georgeon, O.L., Morgan, J.H., Ritter, F.E.: An Algorithm for Self-Motivated Hierarchical Sequence Learning. In: *International Conference on Cognitive Modeling*, Philadelphia, PA, pp. 73–78 (2010)
4. Drescher, G.L.: *Made-up minds, a constructivist approach to artificial intelligence*. MIT Press, Cambridge (1991)
5. Arkin, R.: Motor schema-based mobile robot navigation. *The International Journal of Robotics Research* 8, 92–112 (1987)
6. 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)
7. Oudeyer, P.-Y., Kaplan, F.: Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation* 11, 265–286 (2007)
8. Norman, J.: Two visual systems and two theories of perception: An attempt to reconcile the constructivist and ecological approaches. *Behavioral and Brain Sciences* 25, 73–144 (2002)
9. Dennett, D.: *Brainchildren. Essays on designing minds*. Penguin Books (1998)
10. Piaget, J.: *The construction of reality in the child*. Basic Books, New York (1937)
11. Sun, R.: Desiderata for cognitive architectures. *Philosophical Psychology* 17, 341–373 (2004)
12. Shuster, C., Barlow, R.B., Brockmann, J.: *The american horseshoe crab*. Harvard University Press, Harvard (2004)
13. Barlow, R.B., Hitt, J.M., Dodge, F.A.: Limulus vision in the marine environment. *Biological Bulletin* 200, 169–176 (2001)
14. Cohen, M.A.: Teaching agent programming using custom environments and Jess. *AISB Quarterly* 120, 4 (2005)
15. Newell, A.: *Unified Theories of Cognition*. Harvard University Press, Cambridge (1990)