Early-Stage Vision of Composite Scenes for Spatial Learning and Navigation

Olivier L. Georgeon, James Marshall, Pierre-Yves Ronot

► To cite this version:

Olivier L. Georgeon, James Marshall, Pierre-Yves Ronot. Early-Stage Vision of Composite Scenes for Spatial Learning and Navigation. First Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics, Aug 2011, Frankfurt am Main, Germany. pp.224-229, 10.1109/DEVLRN.2011.6037348 . hal-01354453

HAL Id: hal-01354453

<https://hal.science/hal-01354453>

Submitted on 20 Oct 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Early-Stage Vision of Composite Scenes for Spatial Learning and Navigation

Olivier L. Georgeon
Université de Lyon
CNRS, LIRIS, UMR5205
F-69622, France
Email:
olivier.georgeon@liris.cnrs.fr

James B. Marshall
Computer Science Department
Sarah Lawrence College
One Mead Way
Bronxville, NY 10708
Email: jmarshall@slc.edu

Pierre-Yves R. Ronot
Université Lyon 1
CNRS, LIRIS, UMR5205
F-69622, France
Email:
ronot.py@gmail.com

Abstract—Developmental theories suggest that cognitive agents develop through an initial sensorimotor stage during which they learn sequential and spatial regularities. We implemented these views in a computer simulation. Following its intrinsic motivations, the agent autonomously learns sensorimotor contingencies and discovers permanent landmarks by which to navigate in the environment. Besides illustrating developmental theories, this model suggests new ways to implement vision and navigation in artificial systems. Specifically, we coupled a sequence learning mechanism with a visual system capable of interpreting composite visual scenes by inhibiting items that are irrelevant to the agent’s current motivational state.

Index Terms—Cognitive development; Visio-spatial learning; Intrinsic motivation; Navigation.

I. INTRODUCTION

This work addresses the question of how autonomous agents can discover and exploit spatial regularities in their environment during their initial developmental stage. Cognitive scientists have often argued that cognition is underpinned by an initial developmental stage during which the agent acquires sensorimotor contingencies. These arguments can be traced back to Piaget’s [1] notion of a *sensorimotor developmental stage*. Sun [2] proposed a review of these arguments where he credited Heidegger for the founding inspiration that interactive behavior (the notion of experience in phenomenology) is prior to knowledge.

Following these arguments, we first implemented an algorithm that made artificial agents autonomously learn hierarchical sequences of interactions with their environment. We called this algorithm the *intrinsically-motivated schema mechanism* [3, 4]. This algorithm gave interesting results when the agent was put in an environment that offered sequential regularities to exploit. The agent, however, suffered when put in an environment that had a spatial structure (a two-dimensional grid). It appeared that the agent needed additional skills to deal with space. Our goal now is to investigate these needed additional skills. We begin with studying the acquisition of *permanence* of locations in space, or *topological permanence*. We take topological permanence as a preliminary step before acquiring object permanence.

We propose a model that represents a bee gathering pollen. Based on principles of intrinsic motivation (e.g., [5, 6]), the bee first learns to explore its environment, then learns to navigate between the hive and flower fields. Such navigation illustrates how the agent learns to exploit topological permanence by using visual landmarks.

More broadly, this work seeks to model and simulate *ab-nihilo* autonomous learning, sometimes also referred to as *bootstrapping cognition* [7]. We relate this developmental approach to Piaget’s [1] notion of an *early stage* in human 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 [2]. In our previous studies, we demonstrated that the intrinsically-motivated schema mechanism could be used to make an artificial agent learn *sensorimotor contingencies*—contingencies between the agent’s motor actions and the signals received from the sensors. The agent learned to use tactile perception to avoid bumping into walls [3]. The agent also used rudimentary visual perception to home to targets in the environment. The agent, however, was unable to navigate, for example, toward targets hidden behind walls. This is the kind of limitation that we wish to overcome with an agent capable of exploiting topological permanence.

We derived inspiration from the *dual process hypothesis* for vision (e.g., [8]) that suggests that vision is processed through two parallel pathways: a sensorimotor pathway that is sequential in nature, and an iconic pathway. We also derive inspiration from studies of real bees. Figure 1 shows the agent’s architecture. The part of this architecture that comes from our previous work is reviewed in the next section. The new elements are presented in the subsequent sections.

II. THE SEQUENTIAL SYSTEM

In essence, the *intrinsically-motivated schema mechanism* implements Piaget’s [1] views that perception and action should be considered embedded in sensorimotor schemas rather than separated in the traditional perception-cognition-action loop. Schemas represent interaction patterns between the agent and the environment, and therefore encompass both

Figure 1. The agent's architecture.

the agent's actions and the environment's regularities. In Figure 1, these interaction patterns are represented along the interaction timeline. For example, the two trapezoids at the top of the timeline represent the agent turning to the right. Each trapezoid represents what is seen by each eye while this primitive schema is enacted, as further explained in section V.

The algorithm records past sequences of interactions as hierarchically organized schemas (in the hierarchical sequential sensorimotor memory, part of the sequential system in Figure 1). The algorithm also maintains a representation of the agent's current situation (in the sensorimotor sequential integration system). This representation of the current situation also takes the form of schemas, which complies with Gibson's [9] ecological theory because schemas represent the agent's situation in the form of affordances, i.e., possibilities of interaction. The consistency of knowledge representation across these two systems makes it possible to match past episodes against the current situational representation in order to select appropriate behavioral response (this matching is represented by the double arrow between the sensorimotor sequential integration system and the premotor follow-up system).

Through interaction, the mechanism learns composite schemas (represented in the sequential system as a series of sub-schemas) in a bottom-up way, as the agent experiences the schemas' corresponding sequences. Composite schemas consist of hierarchically-organized subsequences of lower-

level schemas, all the way down to pre-defined primitive schemas (In Figure 1, arc arrows represent the sequential relations in the schema hierarchy).

Lastly, the algorithm implements a novel mechanism of intrinsic motivation: schemas have satisfaction values that operate in a proactive way during the selection of the next schema to try to enact. The modeler initializes the agent with inborn satisfaction values associated with primitive schemas. These settings represent the agent's behavioral tendencies or intrinsic motivations. Through the agent's activity, the schema mechanism "recursively" organizes the agent's behaviors as increasingly elaborated schemas. These schemas capture and exploit the environment's regularities to increase the agent's satisfaction. For example, in our previous experiments, we gave negative satisfaction values to primitive schemas that represented bumping into walls, and positive satisfaction values to primitive schemas that represented the target enlarging in the visual field, which resulted in the agent learning elaborated composite schemas for homing to targets.

III. THE MODEL

Now that we have presented the sequential mechanism, we can introduce our current model. Our bee has three kinds of effectors (represented in the sensorimotor system in Figure 1): a) motion effectors make the bee move, either one square forward, turn 45° to the left, or turn 45° to the right. b) Manipulation effectors: collect pollen from a flower or

deposit pollen into the nest. c) Eye saccade effectors: control the eye saccades, as we will further explain in section IV.

The bee also has three kinds of receptors: a) touch receptors provide feedback from the environment, and indicate whether she bumped into a wall when trying to move forward, or rubbed against a wall when turning. b) Chemical receptors (for “taste”) indicate if she has found pollen or the nest. Chemical receptors have a hard-coded (inborn reflex) connection to manipulation effectors, which causes pollen to be collected from flowers, or to be deposited into the nest. c) Visual receptors generate signals that reflect features in the visual scene. In our previous work, visual receptors only sent signals to the sequence learning mechanism—signals that reflected only dynamic features (changes in the visual field). We have now added *static signals* sent to the static system. Static signals reflect features in the visual field that the bee assumes to denote permanent properties of the world, as we will explain in section IV.

Additionally, we want our agent’s intrinsic motivation to vary with the agent’s internal state, to give the agent an incentive to navigate between different places in the environment. To achieve this, we first considered a mechanism that would change the primitive schemas’ inborn satisfaction values depending on the agent’s internal state. When trying to implement this approach, however, several problems arose. First, it was not clear how to readjust the satisfaction values of past experience in sequential memory. Secondly, the agent had trouble dealing with the complexity of the visual scene. The agent’s motivational state should, indeed, guide the way the agent explores the visual scene, but it was not clear how this could be implemented. To support variation in motivations, and, accordingly, elaborated vision of composite scenes, we looked for inspiration from studies of real bees and other hymenopterans. Specifically, we focus on the following three principles:

- 1) Hymenopterans use visual landmarks to navigate between nesting and foraging areas [10].
- 2) Hymenopterans estimate the distance they travel between landmarks (based on time and on visual flow), and use this estimation for navigation [11].
- 3) Hymenopterans use visual fixation and saccades (through body and head movements) to observe and recognize elements of the visual scene [12].

Notably, hymenopterans use additional navigation methods (e.g., solar compass based on sky light polarization, and path integration) that we chose not to implement to keep our model general. Some studies have also proposed anatomical architectures of the bee’s brain that inspired our agent’s architecture [13]. These studies drove us to implement a second pathway that would process static features in the visual field and associate such static features with the agent’s motivational state. We call this second pathway the static system (Figure 1).

The static system relies on the assumption that landmarks can be individually distinguished and that landmarks are static in the animal's world. This assumption is consistent with experiments on the hymenopterans' iconic memory [10].

Accordingly, in our experiment, each landmark has a unique color that the agent can distinguish. We also implemented a decay function in the static feature memory. Our bee uses this decay to estimate the distance she has traveled since passing a landmark until she reaches the pollen or the nest. In the static system in Figure 1, icons composed of three elements represent the association of static visual features (a square element), motivation (a hive or flower element), and distance traveled (a clock element).

We found that the static system could impact the agent’s behavior by sending an inhibitory signal to the sensorimotor sequential integration system (open arrow in Figure 1), as we explain in the next section.

IV. THE VISUAL SYSTEM

The agent's visual system consists of two eyes that detect colored squares (grid cells) in the environment. Each eye is controlled by a *semiautomatic saccade mechanism* that makes the eye “fix” *singularities*. A singularity is any square salient from the uniform background (i.e., any colored square that contrast with the uniform dark green walls). The saccade mechanism is semiautomatic in that it results from the conjunction of a bottom-up and a top-down process. The bottom-up process makes the eye automatically scan the visual field in search of singularities. The top-down process allows the motivational system to *inhibit* some known squares. That is, the motivational system prevents the saccade mechanism from fixing squares that the motivational system knows and assesses as irrelevant to the current context and motivation (based on the square's static visual features). Figure 2 illustrates these principles.

The scanning starts from the agent's longitudinal axis and goes sideward up to a maximum 45° angle (to the left for the left eye, and to the right for the right eye). If it finds a non-inhibited square, the eye fixes this square and returns a corresponding signal. This signal reflects both a static feature

Figure 2. The agent's visual system.

and a dynamic feature. The static feature is an iconic representation of the square: in our case a single-pixel colored “icon”. The dynamic feature indicates either if this square *appeared* in the visual field, *got closer* (enlarged), or *disappeared* from the visual field during the last interaction cycle.

In Figure 2, the gray areas represent the surface covered by each eye’s saccade. In this scenario, the agent moved forward from cell a4 to b4. The left eye found no singularity in its 45° angular span, assuming that square f5 (blue) was inhibited by the static system. During the previous interaction cycle, however, square c6 (turquoise) disappeared from the visual field. The left eye, therefore, sent no static feature, but did send a *disappear* dynamic feature (represented by the agent’s white left eye with a “o” symbol in Figure 2). Simultaneously, the right eye fixed square e3 (green) which appeared larger than it did during the previous interaction cycle. The right eye, therefore, signaled a *green icon* static feature and a *closer* dynamic feature (represented by the agent’s green right eye and the “+” symbol in Figure 2). Because the saccade stopped before reaching square e1 (yellow), the square remained unseen, even though it was within the 45° visual span (dotted line).

This visual mechanism has two advantages: first, it facilitates the agent’s processing of composite visual scenes by inhibiting irrelevant items; second, it provides a place to implement the effects of the agent’s motivational states on the agent’s behavior. By simply inhibiting the landmarks that are known to be farther away from the target of current interest than the agent’s current estimated position, we get our agent to navigate from landmark to landmark toward the target, guided by the agent’s homing tendency. The agent associates each landmark with the time traveled between the landmark and a specific target (in number of interaction cycles). This time is stored and subsequently adjusted in static memory. The effects of this mechanism are further explained in the next section with an example experiment.

V. THE EXPERIMENT

We use the grid environment represented in Figure 3, adapted from Cohen’s [14] work. The bee is initially in d1. Dark green squares are walls that the bee will bump into if she tries to fly through them (f1-f6, c6-f6, and i6-l6). Square f6 (turquoise) is the wall corner that the bee can distinguish as a singularity. Walls also surround the grid so the bee cannot exit the grid. She does not know *a priori* what possibility of interaction each element of the grid offers to her. Her sequential system only gives her intrinsic motivation to visit and *taste* each singularity.

Square e5 (violet) is the nest where she can deposit pollen. Blue squares are flowers from which she can collect pollen (h3, j4, k7, l8, l9). All these flowers look and feel the same to her. Flowers are removed from the grid after she collects pollen from them. She cannot collect pollen from another flower until she has deposited her previous load back into the nest. Other colored squares are singularities that she can fly over and use as landmarks (b2, a9, h8, g1, i2, l5). The

Figure 3. The experimental grid.

interaction cycle number is displayed in square l6.

A representative run is provided as a video online¹ and the corresponding trace is displayed in Figure 5 (formatted with a method developed from a previous study [15]). The trace shows the first 600 interaction cycles during which the bee explored the environment and harvested the five blue flowers, gathering the pollen into the nest between successive visits to each flower. Time goes upwards and the trace is split into four columns. Figure 4 shows the legend for a sample portion. The *left* and *right* squares and trapezoids indicate the bee’s perception of her environment: either the color currently seen by the corresponding eye, or white if no eye fixation. The trapezoid shapes reflect the shift in the visual field when the bee is turning. Red rectangles indicate wall bumping, and dark green trapezoids indicate wall rubbing while turning (both perceived through touch). On the left side of the trace, vertical lines indicate the bee’s estimated distance to the target. The bee estimates this distance when she recognizes nearby landmarks whose distance to the target was learned earlier. The line is violet when the bee is looking for the nest and blue when she is looking for flowers. Violet hive icons represent pollen-depositing manipulations, and blue flower icons represent pollen-collecting manipulations.

At the beginning (steps 0-60), the trace shows babbling behavior as the bee learns sensorimotor contingencies (as discussed in our previous study [3]). She rubs and bumps into walls on steps 4 through 14. She discovers the nest on step 17. The trace shows the yellow landmark (b2) passing back and forth through the visual field during steps 21 to 38 until she reaches it on step 39. The blue line remains on the left because she has not yet estimated any distance to a flower.

After step 60, she exhibits a more confident and systematic exploration behavior and finds the first flower (k7) on step 97. This causes the other flowers to be inhibited from her visual field. She passes by h8 (yellow-green) on step 104. Because this landmark has not been visited yet, it has no estimated distance to the nest attached (no violet line).

¹ <http://e-ernest.blogspot.com/2011/03/ernest-91-gathers-pollen.html>

Figure 5. The activity trace.

Figure 4. The legend of the trace.

The bee gets an estimated distance to the nest on step 115 when arriving in a9's (green) vicinity. Although already known as farther to the nest than a9, f6 (turquoise) is not inhibited at this point because f6 is not sufficiently refreshed in static feature memory. This causes her to move away again from the nest toward f6 from step 120, before returning to a9 on step 132, and finally finding her way back to the nest via b2 (yellow).

After depositing the pollen on step 163, she uses the estimated distance to pollen now attached to b2 (yellow), a9 (green), and h8 (yellow-green) to navigate again toward the northeast field where she finds the second flower (l8) on step 206. She fumbles again on her way back to the nest. On the third round, the distance values and inhibition mechanism are then settled and she navigates a direct way out to the northeast field; she finds l9 (flower); then she navigates a direct way back to the nest (steps 320-363), using the appropriate landmarks (h8, a9, b2).

In the fourth round, she navigates to h8 again (step 396). From here, she turns in search of flowers, which she sees in the southeast field (h4) and reaches on step 405. Then she navigates back to f6 (turquoise). When arriving in f6's vicinity from the southeast field, though, a9 is hidden behind the walls, and h8 is inhibited because it is estimated to be farther away from the nest (as learned when going to the northeast field). Because she has no uninhibited landmarks at this point, she starts spinning in place (steps 426-447) until h8 decays enough to become uninhibited. Then she finds her way back to the nest via h8.

In the last round, she finds the last flower (i2) via landmark g1. She finds her way back with less waiting because the landmarks' distance and inhibition values start to re-adjust to the southeast field, and she reaches the nest on step 597.

These results demonstrate that the agent can learn both sequential regularities in its sensorimotor interaction with the environment, and topological permanence of landmarks and areas of interest. The agent was able to reach targets of interest spread around in a first area, then explore a second area and readapt to the topology of the second area.

VI. RELATED WORK AND DISCUSSION

While this work is inspired by methods of landmark navigation in mobile robots (e.g., [16]), it goes beyond mere navigation in that it simultaneously supports exploration and route discovery based on the agent's intrinsic motivations.

Some studies also associate landmark navigation with the construction of cognitive maps [17], which we also envisage in future developments.

Our work also differs from traditional cognitive modeling [18] in that we do not yet seek to model how cognitive agents reflect upon symbols. Our agent's architecture, as depicted in Figure 1, is only cognitive insofar as animals such as bees are considered cognitive. This architecture focuses on a pre-symbolic learning stage that we consider important because it keeps the agent's knowledge grounded in the agent's activity.

We also drew lessons from studies that examined the hymenopterans' brain from a functional perspective (e.g., [13]). These studies highlight three main brain regions: the *mushroom body*, which we relate to our sensorimotor integration system, the *premotor region*, which we relate to our sequential system, and the *protocerebrum*, which we relate to our static system. They also identify specific neurons (*ventral unpaired medial neurons*) involved in reward-based learning that we relate to our motivational associative system. They report two pathways from the mushroom body: one toward the premotor region that we relate to our *update* pathway to the sequential system, and one toward the protocerebrum that we relate to our static feature pathway. They report the absence of a direct significant connection between the premotor region and the protocerebrum, which supports our similar modeling choice.

The question of whether hymenopterans use cognitive maps for navigation is still considered open, although many specialists doubt the possibility [11]. Our work illustrates how hymenopterans may navigate without cognitive maps.

VII. CONCLUSION

This work demonstrates that it is possible to implement an agent that autonomously learns and exploits topological permanence in rudimentary settings. We did so in an architecture that associates an intrinsically-motivated sequence learning mechanism with a static visual mechanism. The static visual mechanism processes composite visual scenes, and identifies and memorizes landmarks in association with internal motivational states. The visual system impacts the agent's behavior by inhibiting irrelevant items in the visual field. Through exploration, the agent finds targets of interest and learns the route toward such targets via intermediary landmarks. This work opens the way to more complex models in which an agent will learn topological maps and have more control over its ocular saccades. Such developments inform our understanding of navigation in natural organisms, and suggest new techniques to implement vision and navigation in autonomous robots.

ACKNOWLEDGMENT

This work was supported by the *Agence Nationale de la Recherche* (ANR) contract RPD0C-2010-IDEAL.

REFERENCES

- [1] J. Piaget, *The construction of reality in the child*. New York: Basic Books, 1937.
- [2] R. Sun, "Desiderata for cognitive architectures," *Philosophical Psychology*, vol. 17, pp. 341-373, 2004.
- [3] O. L. Georgeon, J. H. Morgan, and F. E. Ritter, "An Algorithm for Self-Motivated Hierarchical Sequence Learning," in *International Conference on Cognitive Modeling*, Philadelphia, PA, 2010, pp. 73-78.
- [4] O. L. Georgeon and F. E. Ritter, "An Intrinsically-Motivated Schema Mechanism to Model and Simulate Emergent Cognition," State College, PA, The Pennsylvania State University, <http://e-ernest.blogspot.com/2011/01/intrinsically-motivated-schema.html>
- [5] P.-Y. Oudeyer and F. Kaplan, "Intrinsic motivation systems for autonomous mental development," *IEEE Transactions on Evolutionary Computation*, vol. 11, pp. 265-286, 2007.
- [6] D. S. Blank, D. Kumar, L. Meeden, and J. Marshall, "Bringing up robot: Fundamental mechanisms for creating a self-motivated, self-organizing architecture," *Cybernetics and Systems*, vol. 32, 2005.
- [7] D. Dennett, "Brainchildren. Essays on designing minds," *Penguin Books*, 1998.
- [8] J. Norman, "Two visual systems and two theories of perception: An attempt to reconcile the constructivist and ecological approaches," *Behavioral and Brain Sciences*, vol. 25, pp. 73-144, 2002.
- [9] J. J. Gibson, *The ecological approach to visual perception*. Boston: Houghton-Mifflin, 1979.
- [10] T. S. Collett, P. Graham, R. A. Harris, and N. Hempel-de-Ibarra, "Navigational memories in ants and bees: Memory retrieval when selecting and following routes," *Advances in the Study of Behavior*, vol. 36, pp. 123-172, 2006.
- [11] R. Wehner, B. Michiel, and P. Antonsen, "Visual navigation of insects: Coupling of egocentric and geocentric information," *Journal of Experimental Biology*, vol. 199, pp. 129-140, 1996.
- [12] M. F. Land, "Motion and vision: Why animals move their eyes," *Journal of comparative physiology A Sensory neural and behavioral physiology*, vol. 185, pp. 341-352, 1999.
- [13] R. Cotterill, "Cooperation of basal ganglia, cerebellum, sensory cerebrum and hippocampus: Possible implications for cognition, consciousness, intelligence and creativity," *Progress in Neurobiology*, vol. 64, pp. 1-33, 2001.
- [14] M. A. Cohen, "Teaching agent programming using custom environments and Jess," *AISB Quarterly*, vol. 120, p. 4, 2005.
- [15] O. L. Georgeon, A. Mille, T. Bellet, B. Mathern, and F. E. Ritter, "Supporting activity modelling from activity traces," *Expert Systems*, in press.
- [16] L. Smith, A. Philippides, P. Graham, B. Baddeley, and P. Husbands, "Linked local navigation for visual route guidance," *Adaptive Behavior*, vol. 15, pp. 257-271, 2007.
- [17] D. Pierce and B. Kuipers, "Map learning with uninterpreted sensors and effectors.," *Artificial Intelligence*, vol. 92, pp. 169-227, 1997.
- [18] A. Newell, *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press, 1990.