

Cognitive map plasticity and various imitation strategies to extend the performance of a MAS

P. Laroque, E. Fournier, P. H. Phong and P. Gaussier

Neurocybernetic team, ETIS, CNRS (UMR 8051) / ENSEA / UCP

Université de Cergy-Pontoise, 2 rue Adolphe Chauvin, 95302 Cergy-Pontoise cedex

25th September 2005

Abstract

This paper describes the second step of a collaborative work aimed at showing how a system composed of a collection of cognitive agents can solve several non-trivial problems. This second step deals with the problem of facing dynamically changing environments, and with how to identify individual agents (to optimize imitation performance).

keywords: Embodied Intelligence, Biomimetic Autonomous Systems, Cognitive map, Learning, Neural networks

1 Introduction

In previous works, we described how cognitive maps could be used by an agent to solve problems involving contradictory goals [11], and how a simple imitation (actually, an agent-following) strategy could lead a population of such agents to dramatically increase its performance when faced to the problem of surviving in a previously unknown environment [10].

This paper reports on an evolution of our model and system in two directions. On the one hand, we study how an agent can take advantage of its ability to enforce preferred goal-reaching strategies (and forget sub-optimal ones) to adapt to a dynamically changing environment. On the other hand, we describe a mechanism of dynamic individual signature used to distinguish agents from one another; this signature can then be used to (i) determine more easily if some kind of group dynamics is emerging from the agent population, and (ii) propose another imitation strategy .

Experiments have been made that lead to think that our model and system are now mature enough to deal with complex problems such as optimisation, and to retrieve some unexpected results established for instance in spatial economics, such as *unemployment traps*.

2 Material and Method

A complete description of the system can be found in [10], so we only mention the most salient features and properties of the model.

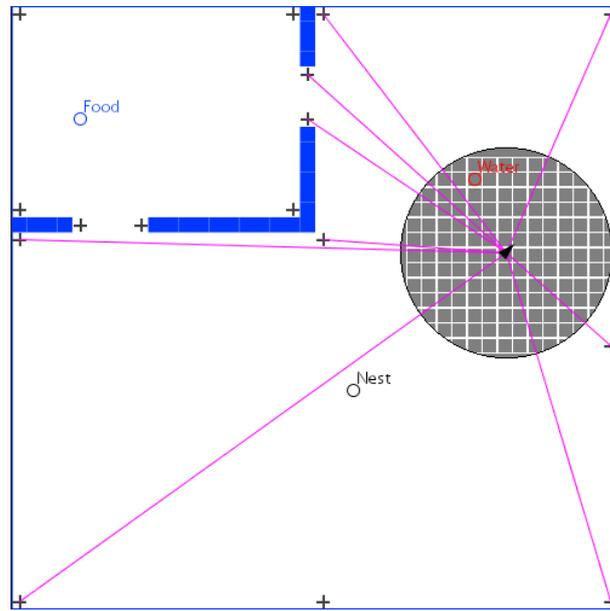


Figure 1: An animat in its environment

Animats [12, 8] live in an initially unknown environment, made of several points of interest (1):

- resources (circles labelled “water”, “food” and “nest”);
- obstacles (solid squares);
- landmarks (small crosses), visible from anywhere except when occulted by an obstacle;
- other agents if any.

The animat can only see objects that are within its visibility range (the disk that surrounds it), except for landmarks. To survive, it needs to discover, and periodically go back to, the three types of resources. Associated with each resource type is a numerical level (a percentage of satisfaction) that decreases exponentially over time. When one of the resource level falls beyond a given threshold, the animat tries to reach a previously discovered, corresponding source.

The animat possible strategies thus fall into 4 categories, that can be summarized in a subsumption-like architecture [4, 5] as in Fig. 2. During its random exploration of the unknown environment, the animat stores acquired topological knowledge in a place cells [6] map. Those place cells are then linked together, at a different level, to form the cognitive map [1, 2, 7] of the animat. When another another animat is in sight, the agent can choose (using a probabilistic decision) to follow it, in the hope it will be led to an unknown resource. When the value of one of its essential variables (hunger, thirst, or stress) reaches its bottom level, the animat needs to reach back the corresponding resource, to have the variable reset to its maximum value. Finally, the obstacle avoidance is based on a Braintenberg [3] algorithm.

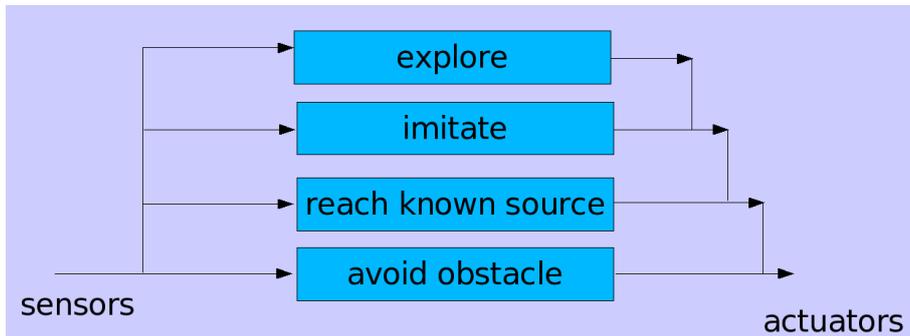


Figure 2: Possible animat strategies

2.1 Cognitive map plasticity

We showed in [11] how adding a learning rule on the building and evolution of the cognitive map could help the agent acquire significantly smarter behaviours, for instance solve contradictory goals. The equations ruling this hebbian [9] learning algorithm are as follows:

$$\begin{aligned}
 w(t+1) &= w(t) - \lambda w(t) + \alpha act_{win} \cdot act_{prevWin} \\
 w(t+1) &= w(t) - \lambda w(t)
 \end{aligned}$$

where $w(t)$ is the weight of a transition between two successively reached cells in the cognitive map and act_c the activity of cell c . The first equation rules the last used link, the second is for all other links in the map.

This development step of our model and system aims at making agents evolve in a dynamically changing environment, that is one in which some sources can disappear when visited for a long time, and others can randomly appear somewhere in the environment.

When a planning agent tries to reach a previously known source and realizes that this source has expired, two things happen: *(i)* the agent dissociates the current place cell from the ex-corresponding resource, and *(ii)* it removes the resource from its set of known resources. Since the place cell will not fire any more when the agent feels the need for this resource, there are chances that the transitions leading to this place be progressively forgotten. Similarly, when a new, matching resource is discovered, the paths leading to the resource will rapidly be reinforced, making the cognitive map evolve synchronously with the environment. This evolution is illustrated in Fig. 3, where left snapshot is taken when the (only) dynamic resource has not yet expired ($t = 15000$ time steps) whereas right picture represents the map after the agent has discovered a new matching source elsewhere in the environment ($t = 35000$ time steps). We can see that some of the paths leading to the old resource location have been almost completely forgotten, and that new paths have emerged.

This cognitive map plasticity and the ability to deal with such dynamically evolving environments makes our model able to describe situations in which the agents, instead of only adapt themselves to their environment, can try to adapt the environment to themselves.

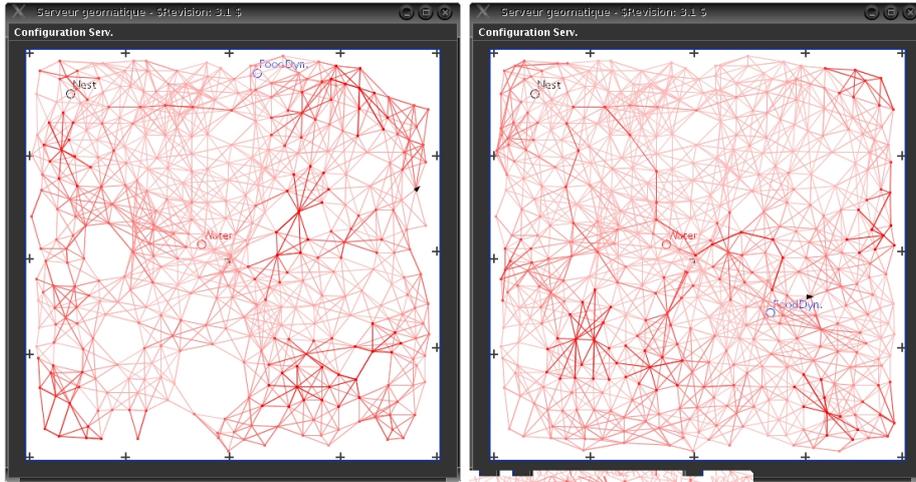


Figure 3: Cognitive map evolution induced by a changing environment

2.2 Identifying individual agents

The fact that agents are undistinguishable from one another was a problem in a variety of aspects: no agent could be sure to follow the same agent when encountering a whole group; statistical results could not separate individuals and detect if one agent spends most of its time in the vicinity of another; etc.

To solve the problem, we decided to add an individual signature to each agent. This signature is to evolve over time, when agents meet each other, in a way inspired from the *talking heads* experiments [13, 14]

[XXXX check reference XXXXX]

We chose to design the signature as a two-coordinate vector, just to be able to map it on a space isomorph to geographical space (see Fig. 4). In that respect, when a new agent appears, its initial location is chosen randomly and its initial signature is the vector of this location. The evolution of agents' signatures is ruled by their meetings with other agents: each time an animat decides to imitate another animat, its signature moves slightly closer to the imitated agent. To avoid a global convergence to a unique signature, a noise is systematically added to each agent signature at each timestep. The equation that describes the variation of signatures is as follows:

$$S_i(t+1) = S_i(t) + \delta_i d(S_i(t), S_j(t))$$

where $S_i(t)$ is the signature of agent i at time t (when the decision to imitate agent j is taken, d is the difference between the two signatures and δ_i is a decreasing function of the age of agent i (the older the animat, the less probable it will be to imitate).

3 Experiments and Results

We can compare two strategies of imitation on the emergence of subgroups: (*i*) the strategy described in [10], in which an agent decides, when encountering a

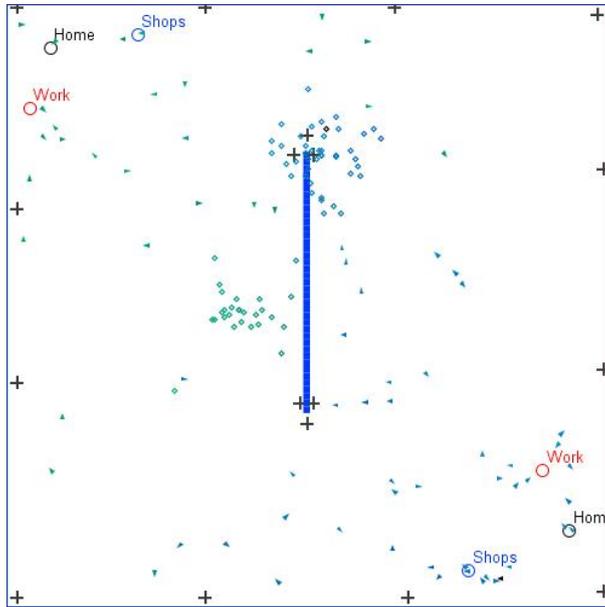


Figure 4: Agents and their signatures

bunch of other agents, to choose the one that is closer to its own direction (which we call here *imitation from azimuth*), and *(ii)* the one that relies on signatures, namely the agent will choose to imitate the one whose signature is closer to its own (*imitation from signatures*).

We compared these strategies from two points of view: *(i)* the emergence of subgroups in a “multi-village” environment, and *(ii)* the survival rate of a population.

3.1 Influence of the imitation strategy on the emergence of subgroups

We used an environment like the one shown in Fig. 4, containing two separate “villages”. For each experiment, we launched 50 agents randomly in the environment and waited for 20000 time steps (which is approximately the time needed for the agents cognitive map to pave the whole environment). Then we study the set of signatures to determine the number of subgroups formed. We repeated the experiment 42 times, 21 times with each of the two imitation strategies. The results are represented in Table 1. As expected, the possibility to distinguish individual agents leads to a more stable way to choose who to imitate, and thus to a bigger number of subgroups. What is more interesting is that, whereas the number of subgroups never exceeds the number of villages when agents imitate from azimuth, we found several cases where imitation from signatures leads to such a situation. This is a clear indication of the greater stability of the groups in this case: the “cost” to imitate someone who is not part of its own group is higher, due to the distance between the two signatures.

# groups	azimut	signature
1	15	0
2	6	14
3	0	7

Table 1: Influence of the imitation strategy on the formation of subgroups (2 villages)

# agents	azimut	stdev	signature	stdev
40	15	1.89	18.42	2.23
50	16.28	3.58	17.57	4.35
70	27	5.08	31.14	7.99
average	19.43		22.38	

Table 2: Influence of the imitation strategy on the survival rate of the populations

3.2 Influence of the imitation strategy on the survival rate of a population

In this series of experiments, we used the same environment and successively launched 40, 50, then 70 agents, for both of the imitation strategies, and counted the number of agents that died for not having found all of the three types of resource. Each of the experiments has been conducted 7 times, and the results are summarized in Table 2. Average number of lost agents is almost always greater with signature-based imitation strategy, but what can seem even more curious is that the standard deviation is significantly higher. Observing what happens during those simulations, we saw that this was due to some subgroups creation around one or two agents that *did not find the three resource types* at the time the group emerged. Consequently, since the chance for an agent to follow its peers is high, the result is of type “all or nothing”: either the resources are discovered by one of the group members, and the whole group survives, or they are not and the whole group disappears.

This phenomenon is to be put in parallel with another, in the spatial economics field, known as “unemployment traps”: they are well-defined portions of urban territories in which the unemployment rate is significantly higher than anywhere around. Although it might be possible for people living there to find a job a few kilometers ahead, things are like if people didn’t try to move outside this small region.

4 Perspectives

[XXXXX déplacement des ressources excédentaires sur les “grandes routes”; ajout des points de rencontre comme centre d’intérêt? Conclusion sur l’autopoïèse? XXXXX]

References

- [1] V. Babeau, P. Gaussier, C. Joulain, A. Revel, and J.P. Banquet. Merging visual place recognition and path integration for "cognitive" map learning. In *The Sixth International Conference on the Simulation of Adaptive Behaviour, SAB 2000*, Paris, september, 2000.
- [2] J.P. Banquet, P. Gaussier, J.C. Dreher, C. Joulain, and A. Revel. *Cognitive Science Perspectives on Personality and Emotion*, chapter Space-Time, Order and Hierarchy in Fronto-Hippocamal System: A Neural Basis of Personality. Elsevier Science BV Amsterdam, 1997.
- [3] V. Braitenberg. *Vehicles: experiments in sythetic psychology*. Cambridge MA: MIT Press, 1984.
- [4] R. A. Brooks. A robust layered control system for a mobile robot. *IEEE journal of Robotics and Automation*, 40:201–211, 1981.
- [5] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, R.A. 2(1):14–23, March 1986.
- [6] N. Burgess, M. Recce, and J. O'Keefe. A model of hippocampal function. *Neural Networks*, 7(6/7):1065–1081, 1994.
- [7] N. Cuperlier, M. Quoy, Ph. Laroque, and Ph. Gaussier. Transition cells and neural fields for navigation and planning. In J.R. Alvarez J. Mira, editor, *IWINAC05*, Lecture Notes in Computer Science, pages 346–355. Springer, june 2005.
- [8] A. Drogoul and J-A. Meyer, editors. *Intelligence Artificielle Située*. Hermes, 1999.
- [9] D. Hebb. *The organization of behavior*. Wiley, New York, 1949.
- [10] P. Laroque, N. Cuperlier, and P. Gaussier. Impact of imitation on the dynamics of animat populations in a spatial cognition task. In *IAS-8*, Amsterdam, 2004.
- [11] P. Laroque, M. Quoy, and P. Gaussier. Learning and motivational couplings promote smarter behaviors of an animat in an unknown world. In *European Workshop on Learning Robots, EWLRL*, pages 25–31, Prague, september 2002.
- [12] J-A. Meyer and S.W. Wilson. From animals to animats. In MIT Press, editor, *First International Conference on Simulation of Adaptive Behavior*. Bardford Books, 1991.
- [13] L. Steels. A case study in the behavior-oriented design of autonomous agents. In *SAB'94*, pages 445–451, 1994.
- [14] L. Steels. A selectionist mechanism for autonomous behavior acquisition. *this issue*, 1996.