Cognitive map plasticity and 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

Abstract. This paper describes the second step of a collaborative work aimed at showing how a system composed of a collection of situated agents can solve several non-trivial planning and optimisation problems. This work deals with the problem of facing dynamically changing environments, and with how to identify individual agents (to optimize imitation performance), which should enable us, in a future third step, to make agents able to act on their environment.

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 [9], 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 an initially unknown environment [8]. One important characteristic of such a system that we wish to study is the possibility of the emergence of social, stable subgroups of agents, knowing that each stable subgroup can be seen as a possible solution to this survival problem. Nevertheless, the complexity of the model and system was limited in some important ways: only static geographical environments could be handled by agents, and they could only use one very basic imitation strategy, since there was no way to distinguish agents from one another. Those limitations did not make the model completely able to give information about the emergence of subgroups.

This paper reports on an evolution of our model and system in those two directions: on the one hand, we study how an agent can take advantage of its ability to enforce prefered goal-reaching strategies (and forget sub-optimal ones) to adapt to a dynamically changing environment. On the other hand, to show how stable social groups can emerge we designed a mechanism of dynamic individual signature used to distinguish agents from one another; this signature can then also be used to 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

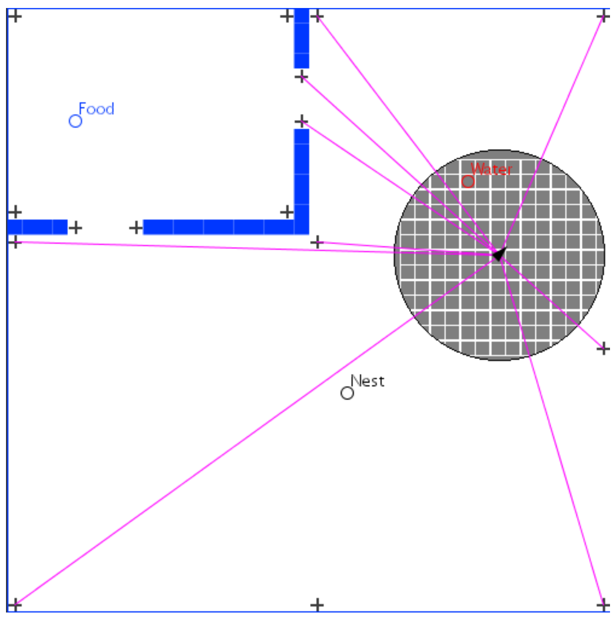

Figure 1. An animat (the small triangle) in its environment, composed of landmarks (crosses), obstacles (rectangles), resources (labelled circles) and possibly other animats. Visible landmarks are linked to selected animat, and its visibility radius is put in evidence with the grey circle around it.

unexpected results established for instance in spatial economics, such as *unemployment traps* [5].

2. Material and Method

A complete description of the system can be found in [8], so we only mention the most salient features and properties of the model.

Animats [10,4] live in an initially unknown environment, made of several points of interest (Fig. 1):

- animats (triangles);
- 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 (water, food and the animat's nest). Associated with each resource type is a numerical level (thirst, hunger and stress, a percentage of satisfaction, called in the rest of this paper *essential variables*) that decreases exponentially over time. When one of the essential variables falls beyond a given threshold, the animat tries to reach a previously discovered, corresponding source.

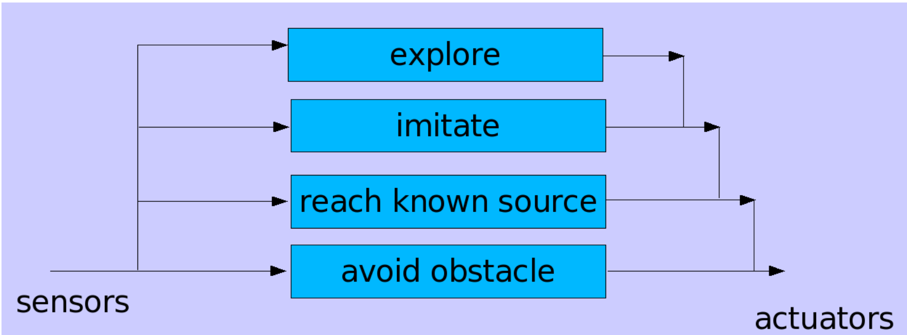

Figure 2. Possible animat strategies

The animat possible strategies thus fall into 4 categories, that can be summerized in a subsumption-like architecture [2] as in Fig. 2. During its random exploration of the unknown environment, the animat stores acquired topological knowledge in a map of transition cells [6]. Those cells are then linked together, at a different level, to form the cognitive map [3] of the animat. When 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 [1] algorithm.

2.1. Cognitive map plasticity

We showed in [9] 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 [7] learning algorithm are as follows:

$$\begin{aligned}\delta w_{ij}(t) &= -\lambda w_{ij}(t) + \alpha act_i(t).act_j(t)\\ \delta w_{kl}(t) &= -\lambda w_{kl}(t)\end{aligned} \tag{1}$$

where $w_{kl}(t)$ is the weight of a link between two successively reached cells *k* and *l* in the cognitive map, act_c the activity of cell c and we assume link *ij* is fired at time *t*.

The current state 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 formerly-corresponding resource, and *(ii)* it removes the resource from its set of known resources. Since the place cell does not fire any more when the agent feels the need for this resource, there are chances that the use of transitions leading to this place be progressively forgotten. Similarly, when a new, matching resource is discovered,

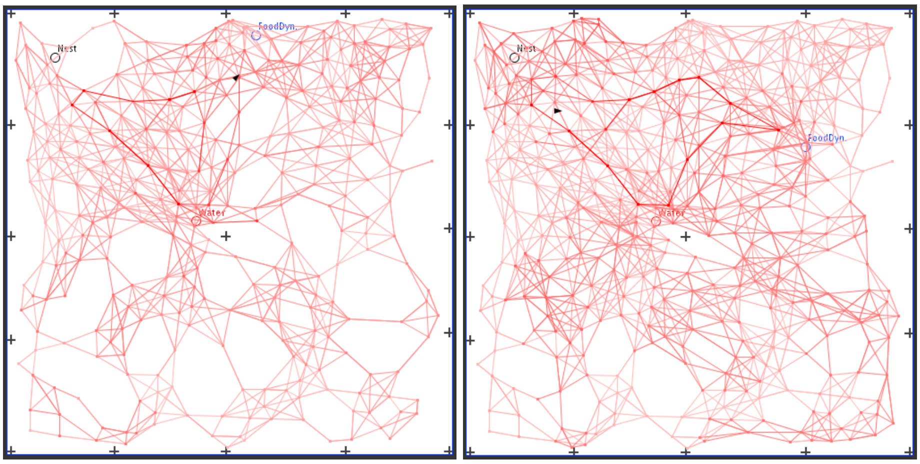

Figure 3. Cognitive map evolution induced by a changing environment. The roads leading to the food source in the left snapshot (top of figure) of the cognitive map are partially forgotten when the source has changed, whereas new paths have appeared near the new source location (right side of the right snapshot of the cognitive map).

the paths leading to the resource is rapidly 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 = 10000 time steps) whereas right picture represents the map after the agent has discovered a new matching source elsewhere in the environment (t = 25000 time steps). With respect to eq. 1, $\alpha = 5.10^{-2}$ and $\lambda = 10^{-6}$. We can see that some of the paths leading to the old resource location have been partially forgotten, and that new paths have emerged.

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 [11], which studied the evolution of a shared lexicon in a population of embodied software agents. The agents developed their vocabulary by observing a scene through digital cameras and communicating about what they had seen together.

We chose to design the signature as a two eight-bit coordinate vector, just to be able to map it on a space isomorph to the geographical space (see Fig. 4, where agents appear as triangles and signatures as diamonds), which allows to better follow the evolution of signatures, in the map and in time. 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 system-

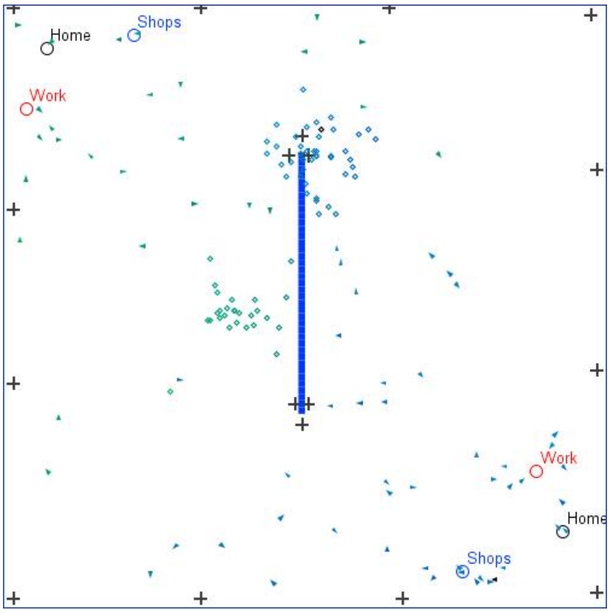

Figure 4. Agents and their signatures

atically added to each agent signature at each timestep. The equation that describes the variation of signatures when an imitation decision is taken is as follows:

$$S_i(t+1) = S_i(t) + \delta_i(S_j(t) - S_i(t))$$

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

3. Experiments and Results

We can compare two strategies of imitation on the emergence of subgroups: *(i)* the strategy described in [8], in which an agent decides, when encountering a bunch of other agents, to choose the one that is closer to its own direction (which we call here *imitation from azimut*), 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*).

One of the interesting characteristics of our system is that it is possible to deal with problems having several (families of) solutions. For instance, if we group together the three types of resources in a small area of the map, and we do the same in another area of the same map, then the trajectories of the agents can stay closely related to one of the two possible “villages”.

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

# subgroups	azimut	signature
1	15	0
2	6	14
3	0	7
total experiments	21	21

Table 1. Influence of the imitation strategy on the formation of subgroups (2 villages). We report here the number of experiments that led to the emergence of 1, 2 or 3 subgroups, when imitation is based on the azimut (left column) or on the signature (right column).

# agents	azimut	signature
40	15.0 +/- 1.5	18.4 +/- 2.4
50	16.3 +/- 3.1	17.6 +/- 4.9
70	27 +/- 4.8	31.1 +/- 8.3
average	19.4	22.4

Table 2. Influence of the imitation strategy on the survival rate of the populations. Average number (and standard deviation) of surviving agents. The much higher standard deviation for signature-based imitation is due to some global group disappearance.

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 tile 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, and reported, for each experiment, the number of subgroups that have emerged. The results are represented in Table 1. As expected, the possibility to distinguish individual agents leads to a more stable way to choose which agent 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 azimut, 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.

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 summerized 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 could be put in related to another process, observed 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 [5]. 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. In our simulation, altough the needed resource is in the reach of the group, it is only seldom discovered because of the strength of the intra-group link.

4. Conclusion

The current step of our model and system allows to solve non trivial planning and optimisation problems, with only very little assumptions on the initial knowledge of the agents. They can adapt themselves to a changing environment, share a partial kowledge of the problem with each other, handle multiple, contradictory goals and find different solutions to the problem when multiple solutions exist.

It should now be possible to see how the optimisation can be pushed one step further, by making agents able to act on their environment instead of just adapting themselves to it: for instance, carrying some resource from "natural" sources to locations near important paths could dramatically enhance the performance of the global system, as far as the average "satisfaction" level of the agents (the average of an animat's essential variables values) is concerned.

References

[1] V. Braitenberg. *Vehicles: experiments in sythetic psychology*. Cambridge MA: MIT Press, 1984.

[2] 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.

[3] 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.

[4] A. Drogoul and J-A. Meyer, editors. *Intelligence Artificielle Située*. Hermes, 1999.

[5] F. Gaschet and N.Gaussier. *Concentration économique et ségrégation spatiale*, chapter 12, "Les échelles du mauvais appartement spatial au sein de l'agglomération bordelaise", pages 221–241. M.A. Buisson and D. Mignot, de boeck university edition, 2005.

[6] P. Gaussier, A. Revel, J.P. Banquet, and V. Babeau. From view cells and place cells to cognitive map learning: processing stages of the hippocampal system. *Biological Cybernetics*, 86:15–28, 2002.

[7] D. Hebb. *The organization of behavior*. Wiley, New York, 1949.

[8] 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.

[9] 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, EWLR*, pages 25–31, Prague, september 2002.

[10] 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.

[11] L. Steels. *The Talking Heads Experiment. Volume 1. Words and Meanings*. Laboratorium, Antwerpen, 1999.