

**Étude des stratégies d'exploration d'un environnement inconnu
pour l'apprentissage de cartes cognitives et pour la navigation
spatiale : des données expérimentales à l'implémentation
robotique**

**Exploration Strategies in unknown environments, learning
cognitive maps for spatial navigation.**

Master Thesis
September 2011

Félix Grèzes

Host laboratory: Neurobiology of Adaptive Processes (UPMC - CNRS UMR-7102)
Institut des Systèmes Intelligents et de Robotique (UPMC - CNRS UMR-7222)

Supervisors: **Angelo ARLEO**
Stéphane DONCIEUX
Denis SHEYNIKHOVICH

Contents

- 1 - State of the Art
 - .1 Motivational Theories
 - .2 Exploratory Patterns
 - .3 Natural Tools for Spatial Learning
 - .4 Evolutionary Algorithms
- 2 - Work & Methods
 - .1 Architecture of the Simplified Model
 - .2 Neural Network Controller
 - .3 Evolutionary Algorithm
- 3 - Results
 - .1 Simple Controller
 - .2 Full Controller
 - .3 Behavior Analysis
- 4 - Conclusion and Discussion

Bibliography

1 - State of the Art

.1 Motivational Theories

Plenty of research has been done to understand the motivations behind an animal's decision during exploration. We distinguish two types of exploration studies: goal oriented exploration, where behaviors have explicit motivations, such as hunger, reproduction, fear of predation; and our focus, free exploration, where no clear motivation is present.

The consensus among the studies of free exploration is that fear and curiosity are the main driving forces behind the animal's behavior. For instance, Montgomery and Monkman (1955) argued that fear was a limiting factor in the exploratory behavior of rats, but that animal with more pressing needs, such as rats who have been sleep or food deprived, are more ready to explore the unknown environment (Montgomery, 1953). This tends to show that exploratory behavior is only observed when the curiosity drive of the animal overcomes its fear.

Moreover, a study by Halliday in 1966, showed that rats are always keen to explore a novel environment, regardless of whether or not they have already spent a lot of time exploring a different environment before hand. Halliday then argues that the animal's curiosity is just a form of fear, in this case fear of the unknown. These studies point out the causes of a rat's exploratory strategies, either choosing immediate safety by staying in one place, or long term safety by exploring and discovering the unknown environment, but do not focus on the actual strategies used by the animals.

.2 Exploratory Patterns

While for a long time it was assumed that the exploration followed a random process, more recent studies have observed recurrent patterns in the exploratory patterns of animals. One of these patterns, observable across species, is the establishment of a home base, a place where the animal often returns, even in featureless environment. (Conradt et al. 2000 and Armstrong et al.1999).

We now focus solely on the behavior of the rat, as opposed to other animals. The questions we ask ourselves are: in true free exploration, what are the rat's exploratory patterns, how do they emerge, and how can they be explained?

The first of the patterns we will focus on is the creation of a home base. When dropped in a novel environment, the rats will spend the first minutes establishing a home base, a sort of safe haven (Eilam et Golani, 1989). During this period, the animal will spend most of its time in this place, and will return there often during the exploration phase. (see Figure 1)



Figure 1: Example of home-base evidence.
The home base in this case is near the starting location.
(Image from Arleo & Rondi-Reig, 2007)

While this home base is often at the spot where the rat has been initially placed, its position can be biased by notable features of the environment, such as salient landmarks or even simple corners in the arena.

The second pattern of interest is that of looping back and forth the home base. This behavior consists of excursion leaving and returning to the home base, each excursion marked by a number of stops. (see Figure 2)




Figure 2: Example of looping around the home base.
The return path indicates path integration by the animal,
and not just tracing back the steps.
(Image from Arleo & Rondi-Reig, 2007)

Golani et al. (1993) show that these trips become longer over time, and that while different rats may have different trip duration, the rate of progression for this duration is similar across the species. They also point out two different behavior steps during these loops, first the rat moves out slowly in a zig-zag, then runs back home in a straight line to the home base.

.3 Natural tools for spatial learning

We now focus on the natural equipment that the rats possess allowing them to form an internal representation of space. It has been shown (O'Keefe and Nadel, 1978) that the hippocampal formation of the rat's brain acts as a center for the internal cognitive map.

More precisely, there are cells in the hippocampal formation whose activity is strongly correlated to the place or head direction of the animal (See Arleo and Rondi-Reig 2007). Even if the hippocampus is certainly involved in non-spatial aspects of memory, like Eichenbaum et al. argue (1999), it seems clear that the hippocampal formation has a role in extracting and using the environmental information.

This information can be classified into two categories: Allothetic and Idiothetic. Allothetic information is defined as the information that the animal obtains from external cues through its senses, e.g. lights as visual landmark, olfactory marks. Idiothetic information on the other hand is the information available to the animal based on his own internal knowledge of his movement. It is how you can approximately judge your own position while walking eyes closed. This process is called path integration and some of the the internal organs proven to play a role in it are the vestibular system and the use of optical flow.

Both types of information have weaknesses in supporting self-localization. The error in position given by allothetic information alone has the benefit of not being time dependent, it can be however place dependent, symmetries can induce ambiguities, and the error can be quite large for animals with poor eye-sight (like rats). On the other hand, the idiothetic information error is independent from the environment but accumulates with time. Our hypothesis states that both types of information are necessary for a coherent internal representation of space, and that the home base loop behavior makes optimal use of both.

.4 Evolutionary Algorithms

In evolutionary algorithms, the agents are modified through the repetition of three steps: Evaluation, Selection and Reproduction, designed to mimic the process of natural selection. Simply put, during the evaluation step, each agent is tested and given a score for each of the objectives we are testing. This score is called fitness. During the selection step, only the best individuals, the ones with the highest fitness score are kept for reproduction. During this last step, a new population is created by random mutation and crossovers from the selected individuals. The process is then repeated until termination.

In our work we used the SferesV2 framework (Mouret & Doncieux, 2010). This framework, written in C++, implements multiple evolutionary algorithms. We used a multiple objective evolutionary algorithm called NSGA2 (Deb 2002). Sferes also allows for multi-thread evaluation of individuals, a considerable boost in speed. Using powerful object oriented abstraction techniques, Sferes also allows the user to define their own genotype, phenotype, evolutionary algorithm, or any other possible aspect of the algorithm. Sferes also features a library called fastsim, designed for robotic simulations, but also useful for visualizing behaviors.

Problem and Hypothesis

Our goals, based on the previously mentioned studies, is the find optimal exploration strategies, i.e. ones that result in a coherent place cell representation of space, using

evolutionary algorithms. We will then compare these emergent strategies with the ones observed in animals, looking at the similarities. Our hypothesis is that the home-basing, round-trip loops is such an optimal behavior.

2 - Work & Methods

To test our hypothesis, we recreated an experiment adapted from the one described in the works of Eilam & Golani: in an empty, homogeneous environment, rats are released and their exploratory patterns are observed. (see Figure 3)

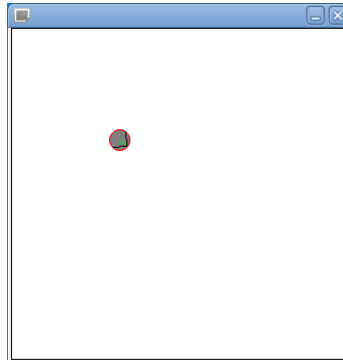


Figure 3: The virtual rat/agent in its empty arena.

All the parameters: speed, arena size etc. are tuned to reproduce the experiment setup of Eilam & Golani.

To do so, we proceeded in three steps, corresponding to part 1.3, 1.2, and 1.4: first we implemented a simplified model of the animal's internal space representation, based on the works on hippocampal place cells of Wilson & McNaughton (1993) and Sheynikhovich et al. (2009). The second step was implementing a neural network controller that used the information from the simplified model to explore the environment. Lastly we optimized the weights of this controller through the means of a evolutionary algorithm, with the goal of finding a variety of behaviors, each optimizing either security, exploration, or a combination of both.

.1 Architecture of the Simplified Model




Figure 4: Architecture Overview

Our model works by having two sets of X,Y,Θ coordinates: one representing the animal's real position and angle, and the other representing its path integration position and

angle. The real X, Y, Θ are an abstraction of all the allothetic information available through cues, that could be visual, acoustic or other. These coordinates are used to calculate the allothetic cell activities, while the path integration X, Y, Θ determine the idiothetic cell activities. Both the allothetic and idiothetic cells have Gaussian activities, with the theoretical maximum being 1.

The idiothetic cells are implemented as artificial neurons with Gaussian activity profiles whose centers cover the arena uniformly. The population activity of the idiothetic cells is meant to imitate the population activity of grid cells. These cells are pre-allocated at the beginning of the simulation.

The allothetic cells are allocated at the same time as new place cells (see below), as such their centers activities are the animal's perceived position at the time of the allocation of the cell. To emulate the rat's poor vision, the maximum value of an allothetic cells is lowered by a random amount, so that some place might need multiple cells for an accurate estimation of the position given by cues. As we have pointed out, the error in position given by the path integration coordinates increases constantly over time, and must be corrected by the position given by the external cues. This process is possible when the allothetic cell's activity give a precise enough estimation of the real position.

Here the formula detailing the inner workings of our model's allothetic cells:

$$\frac{\sigma_{xy} \otimes \sigma_{\theta}}{\sigma_{allo}} * \exp\left(\frac{(X_{cues} - X_{allo})^2 + (Y_{cues} - Y_{allo})^2}{2 * \sigma_{xy}^2}\right) * \exp\left(\frac{|\theta_{cues} - \theta_{allo}|^2}{2 * \sigma_{\theta}^2}\right)$$

The following are internal variables specific to each allothetic cell:

x_{allo} , y_{allo} and θ_{allo} are the center of activity for the cell. These are defined as the true position and angle when the cell was created.

x_{act} , y_{act} and θ_{act} are the current true position and angle. The cell activity will be maximum when the rat comes back to the cell creation spot with the same angle.

σ_{allo} represents the cell's quality. A high value emulates a cell with low maximum value, this can mean for example that vision was bad when the cell was created.

σ_{xy} and σ_{θ} are global parameters for the model. They represent the size of the allothetic cell's receptive field, respectively in terms of position and in terms of angle distance. In our experiment, all the values were set to 10.

Another important detail is how the path integration position is corrected by the allothetic information (see Figure 4). When an allothetic cell is created, we store the cues (i.e. the real coordinates) related to where the rat was at that time, but also the estimated position for that cell. It is this estimated position that is used for correcting the path integration position. An example can help understand: if we are in Berlin, walking in what we believe is East and at one point we see the Eiffel Tower, we will believe that the Eiffel Tower is East of Berlin, so next time we come across the tower and wish to go back to Berlin, we will walk West. In this case, we were so lost when we created the allothetic cell near the Eiffel Tower that it

will incorrectly recalibrate the path integration position. This indicates that if the virtual rat is to build a coherent map of the environment, it should try to never get lost, and often revisit places so that good allothetic cells outnumber the bad ones.

The place cells are the rat's internal representation of space and are allocated when no place cell is active, i.e. when the rat is in place it does not recognize, either because it has never been there, or because it is lost. The activity of a place cell is computed by a weighted sum of the activities of both allothetic and idiothetic cells. The weights are established at the creation of the place cell, and is the activity of the allothetic and idiothetic cells at that time. As such, a place cell's activity is tied to both the position given by the cues and the position given by path integration at the time of its creation. This type of one-shot learning, where the weights do not change overtime, was implemented for simplicity reasons, and can be replaced by a more biologically plausible version of Hebbian learning (Hebb, 1949), where the weights change slowly over time.

This architecture aims to reproduce experimental biological results where a rat's position can be correlated to the activity of place cells. In our implementation, if the rat is lost during the exploration phase, i.e. the cues X, Y, Θ and the path integration X, Y, Θ differ, then the place cells created will not have a strongly active point in their receptive field (see Figure 6), with the opposite happening if the rat is able to keep track of its real position (see Figure 5).




Figure 5: The receptive field of an accurate place cell after environment exploration. The maximum activity is close to 1 and near the real position at creation time. Note that the field is not a perfect Gaussian shape, which would be the case if the rat made no error in movement or perception




Figure 6: The receptive field of an inaccurate place cell after exploration
 The maximum activity is never above 0.5 and there is an ambiguity
 as to where the rat was when the cell was created.

.2 Exploration Controller

At each time step of the simulation, the virtual animal's exploratory movement is given by the output of a simple neural network. The network's neurons are classical perceptrons invented by Rosenblatt in 1957, performing a simple linear sum of the weighted inputs with the only difference being that we use a sigmoid transfer function resulting in a neural activity with values between -1 and 1. The synaptic weights are within $[-1;1]$ as well. The outputs defines the rat's movement speed and angle. The inputs of the controller are normalized to values between -1 and 1. The inputs were chosen with computability and overfitting in mind. For example, if the activity of all the place cells were included as inputs for the controller, the neural network's output calculation would take too much time, and the evolutionary algorithm optimization might have trouble finding general behaviors in a reasonable amount of time, given the hundreds of individuals that have to be evolved over many generations. With these constraints in mind, and after testing multiple possibilities, we built two versions of the controller. The first version was meant to be a proof of concept, with the second version being the real experiment.

For the test version the only three inputs to the controller were :

- . a constant, necessary as the bias for all neurons
- . a uniform random noise within $[-1;1]$
- . the time since the last position update by the allothetic cues, in simulation time

steps

The outputs of the controller represented high-level instructions were "do nothing", "move randomly" and "return home", and the instruction associated with the highest output determines the virtual rat's action. In this test version, the agent does not need to figure out the angle home from the inputs

For the more biologically plausible exploration controller, in addition to the same inputs as the previous version, the following inputs were added:

- . the angle to the starting location, given by the rat's perceived position
- . the angle of the previous movement
- . the maximum activity of the place cells

In this version, the outputs of the controller are the explicit speed and angle of the rat movement. This version of the controller offers the possibility of showing more diverse behaviors, but to return to the home base, the angle must be extracted from the inputs.

.3 Evolutionary Algorithm

To explore the space of possible behaviors, we optimized the weights of the exploration controller. We define a population of individuals, each representing an agent (the virtual rat) in our experiment. Each individual possesses a genome that defines the weights of the rat's neural network controller. (see Figure 7)

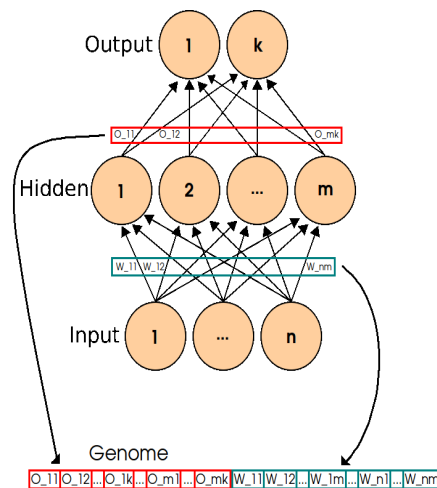


Figure 7: The neural network and the associated genome

In our experiment, the fitness of individuals is two-dimensional: we are interested in agents capable of exploring as much of the arena as possible, and agents who have an accurate internal map of what they have explored. We define the exploration fitness as the surface area covered by the virtual rat during exploration, and the map accuracy fitness as the sum of each created place cell's accuracy. This accuracy is simply defined by how far away the cell's creation place is from the center of gravity of this cell's receptive field.

To sort the individuals with the two dimensional fitness, we used the Pareto selection method. Using the entirety of the Pareto front is the only way to consider every non-dominated strategy, but extracting it is more time consuming than using other bi-objective comparison methods, such as linear combination or lexicographical order.

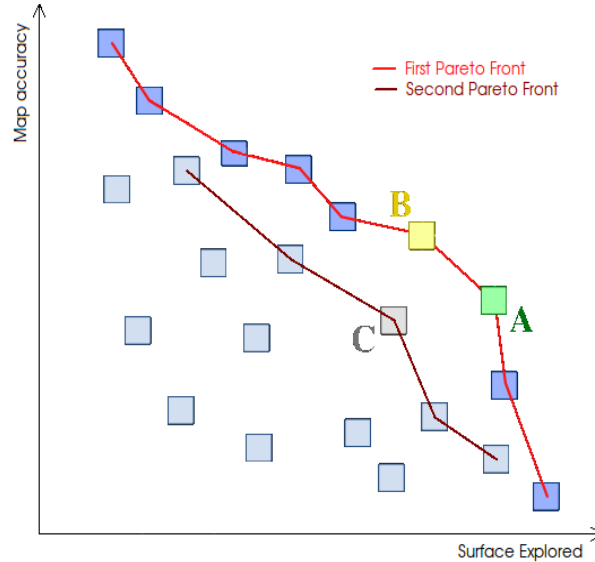


Figure 8: Pareto front
 A scores higher in map accuracy than B, but lower in surface explored, and so neither individual dominates each other

Figure 8 presents the concept of Pareto front, which is composed of individuals whose fitness is not dominated by any other individual. The NSGA2 algorithm will select for reproduction individuals on the first Pareto front. If the population of this front is too small for the selection then individuals from the second front will be chosen. If the population is too large, the algorithm will attempt to select individuals covering the largest area of the front. One important detail is that both fitness functions are directly competing against each other, so that improving an individual's accuracy fitness will necessarily decrease its surface explored fitness. Such antagonistic goals are useful at promoting a diversity of behaviors, answering the common complaint that evolutionary algorithms tend to stop on a local maximum of the fitness function.

3 - Results

Before running either version of the controller, genomes describing the sought-after home looping behavior were hand-written, proving that it was at least theoretically possible for the evolutionary algorithm to find it.

.1 Simple Controller

We first ran the simplified version of the experiment on our lab computer.




Figure 9: Simple experiment
Evolution of the Pareto front

On figure 9, every Pareto optimal individual of all 500 generations is plotted. Our results suggest that the evolutionary algorithm does not seem to be capable of systematically improving the behaviors. This can be seen from the fact that the Pareto front is not moving toward the ideal point of maximum exploration and accuracy. The three visibly separate group of individuals correspond to identical individuals, in the sense that they always execute the same action for the duration of exploration (except for the small, center group). Figure 10 show what action each groups corresponds to.

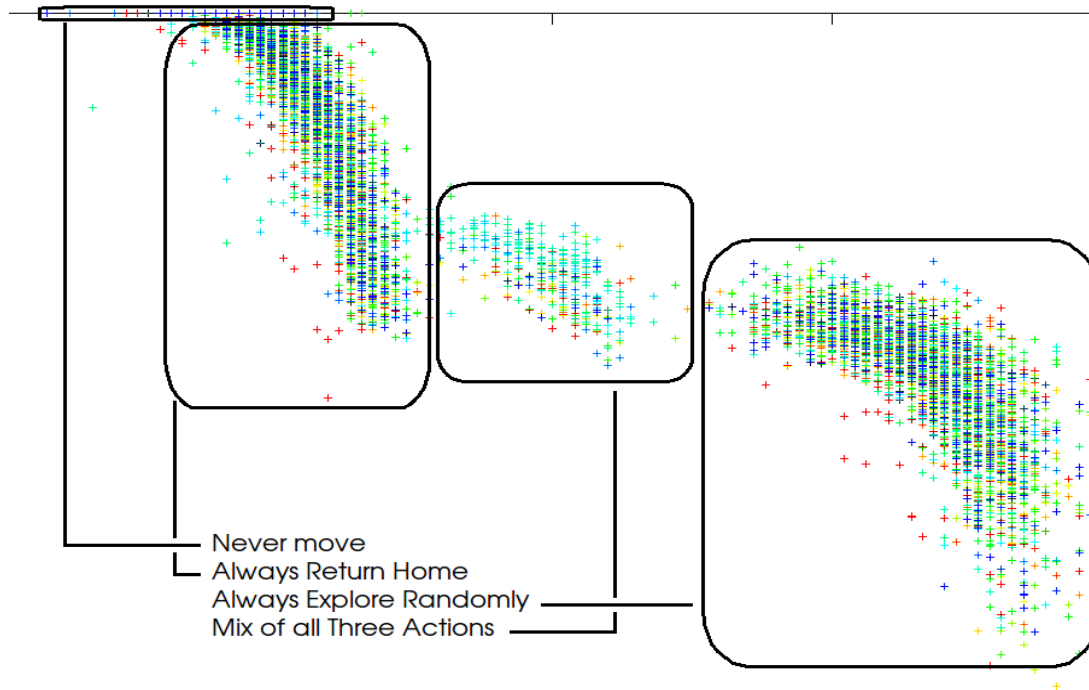


Figure 10: Groups of similarly acting individuals

Our analysis of individuals from different groups suggest that, among individuals of the same group, the differences in fitness is solely explained by random chance. Indeed, the individuals of each group visible above, except for the ones in the groups mixing actions, where identical to all other individuals of the same group. The group whose behavior is a mix of different actions could potentially hide the home base looping behavior. The strategies among these individuals varied greatly: moving for a certain period of time then stopping completely, randomly returning home. Unfortunately these individuals were rare and not present on each generation, which means their behavior was not transmitted to the offspring, and so no individuals showed the complete home base looping behavior for the entire exploration time. That such basic strategies could be optimal might indicate that something is wrong with the implemented simple version of the exploration controller, or that the parameters chosen were inadequate. In the full experiment, the individuals were evaluated 4x more, so as to reduce the variance in fitness that was apparent here.

.2 Full Controller

We ran the real experiment, with angle and speed outputs, on the ISIR cluster computer. With a population size of 100 individuals, evolution took place at the approximate rate of five generations per hour. Here is the results of the experiment after 50 generations.




Figure 10: Full experiment
Evolution of the Pareto front

Unlike with simple controller experiment, Figure 10 shows us that with the full controller, no visible groups of individuals is formed. This might be because the actions now allow for a more continuous variation. And now, the strategies clearly improve over time, as can be seen by the latest generations covering a larger surface and improving closer to the ideal point. Certain optimal behaviors appear very rapidly, but as can be expected, these are simple and extreme, in the sense they only optimize one of the two fitness.

Behavior 1: Stay in Place

The simplest strategy for the rat is to never move. By staying in the same spot, it will never get lost, thus maximizing the accuracy of its internal map. This strategy is easy to find by evolution or even random sampling (which is what the first generation is), since all that is needed is a zero or close to zero speed during the entire duration of exploration. The end result is a virtual rat turning around itself in a very tight angle. This behavior can be compared to the first step of the home base behavior observed in rats, where the animal spends a long time at the starting location before exploring.

Behavior 2: Constant Exploration

Another simple behavior is the polar opposite of the previous one. Now the rat is always moving in a new direction, thus getting lost very fast but maximizing the surface explored. While the simplest form of this strategy, moving randomly, is present in early generations, it is optimized through the generations as to only change direction when hitting an arena wall. This lowers the chances of revisiting a place the rat has already explored.




Figure 11: Evolution of a strategy

The most exploratory behavior improves over time.

Not visible on the figure, the gen 50 strategy spends less time stuck on walls.

Figure 11 shows us improvement in strategy. The red line is a sample path from the most exploratory strategy from the first generation. It explores recklessly but goes over its own path, while the strategy from the 50th generation, represented by the sample blue path, make wider angles to avoid staying in the same area. This strategy never establishes an accurate map of the environment, except at the starting location, and is never observed on real rats. Since no internal map is accurately made of the environment, it can not even be called an exploratory strategy.

Behavior 3: Spring-like Spiraling

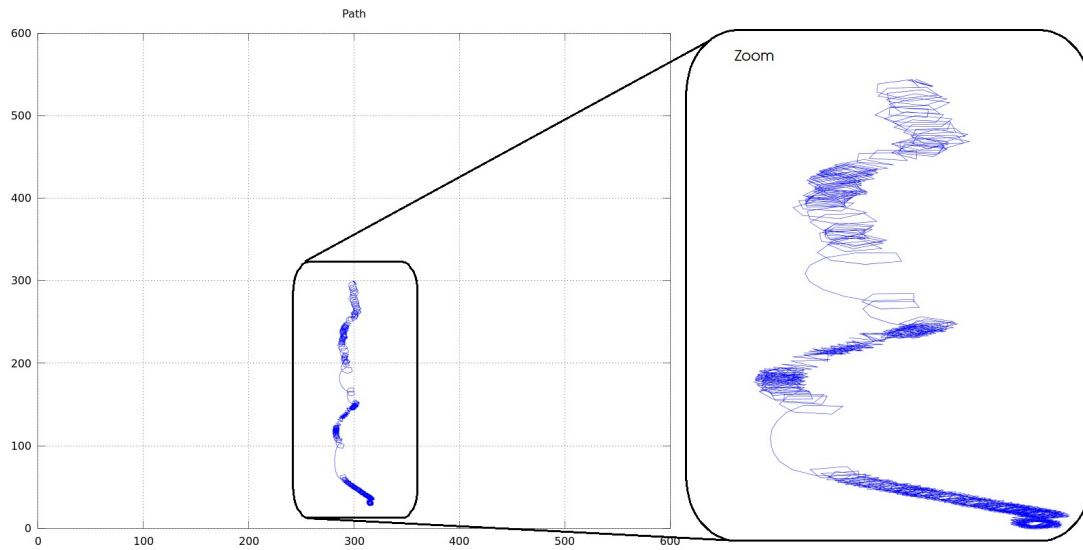


Figure 12: Spiraling Behavior

Appearing in the later generations, figure 12 shows a sample path of this behavior consisting of loops that push the rat in given direction. By passing over already explored spots, the rats ensures that it never gets lost, but still explores portions of the environment. This is interesting because the characteristics are somewhat similar to the desired home base loops behavior we are searching for, but much easier to find, since all that is needed is a more or less constant change in angle movement. However it does not show explicit attempts by the rat to go straight home.

The average error in estimating the true position for this behavior is shown in Figure 13.

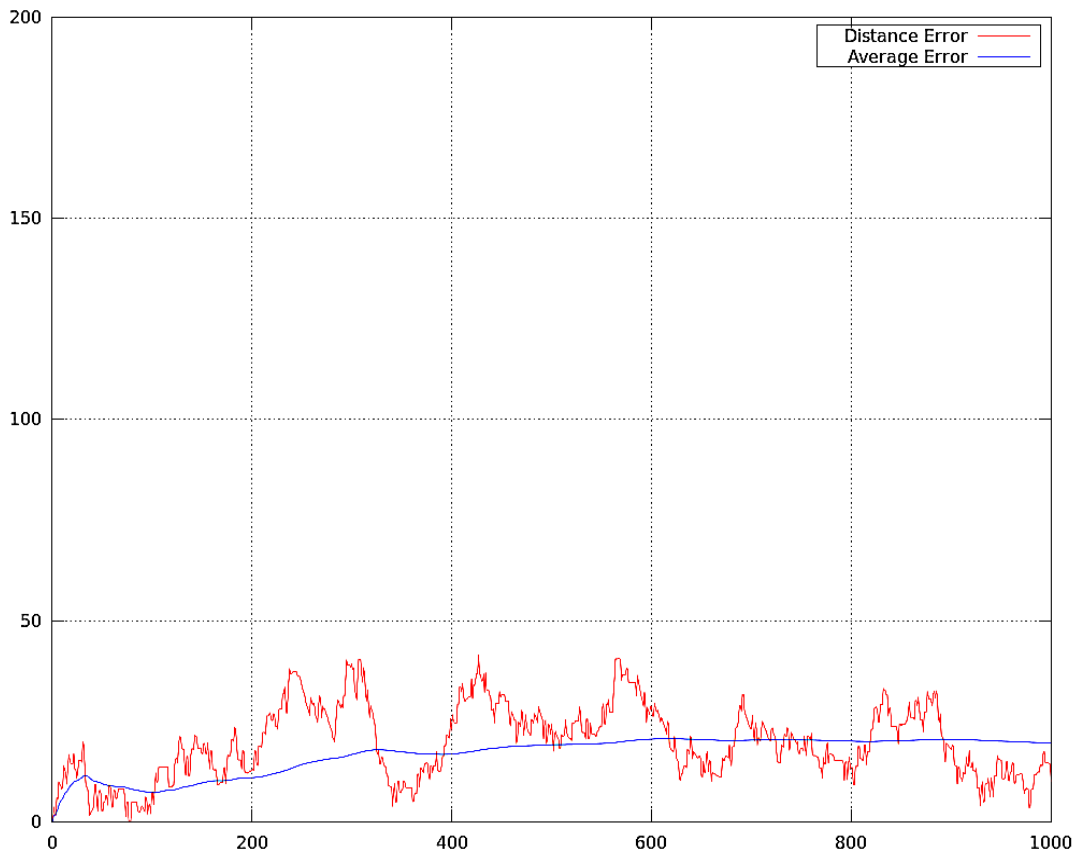


Figure 13: True Position Estimation Error for Spiraling Behavior
 X axis: exploration time in simulation iterations.
 Y axis: Distance error in cm.

The average error for this strategy is very low, almost as low as with the strategy that always stays home. This is in addition to the fact that this individual was always moving across the arena, unlike those expressing pure home base behavior. Such a strategy could very well be optimal in our implementation, the reason being that it can be simply tweaked by changing the size of its loops emphasize precision or exploration.

Others

Other types of strategies were found by evolution, but did not offer interesting behaviors. For example some individuals spends a certain amount of time at the home base, before leaving it without returning. This type of behavior is easy to modulate to be more or less exploratory / secure, by staying for a longer or shorter time at the starting location, and a such are present wherever more complex behaviors have not emerged.

4 - Conclusions and Discussion

Our first goal was find the optimal exploratory behavior for our model, controller and fitness functions. It seems that the spiraling behavior might be close to this optimal, as it is the only one we found capable of both exploring the environment while keeping an accurate internal map of the places the rat visits. The only other behavior whose average error between true and estimated position did not continuously grow was the home base behavior, where the rat stayed in the starting spot for the duration of the exploration.

Our second was comparing the behaviors we found with those experimentally observed on live rats. Experimental data shows that some animals spend almost all their time at the home base, which one of our behaviors seems to mimic. However we did not find a behavior in which the rat returns to the home base, instead the spiral looping behavior does not seem to need such drastic actions to keep a coherent internal map. In order to continue our search, changes need to be made with the simplified place cell model, or with the fitness functions, or probably both.

Our model has trouble incorporating angles in the activity of cells. More precisely, it is very unlikely for the animal to return to a position X,Y already visited, while having the same angle as the first time. This means that the activity of place cells is often lower than what needed. For the experiments, the influence of angle on cell activity was artificially lowered, but this meant that the angle movement error had to be kept to almost zero, because the true angle could not be strongly inferred from allothetic cell activity. This had a profound influence on all the other parameters of the experiment because angle error is a lot more severe than distance error. In the end this probably made the model too far away from reality. To fix this issue, we should probably modify the model so that two type of cells are present, on type reacting only to X,Y position, and the other to angle direction. This would be based on works on head direction cells (see Taub et al. 1990).

Possibly because of the previous issue, the simple experiment showed that identical behaviors can have wildly different fitnesses. This make the emergence by evolution of different strategies more difficult, and while increasing the number of evaluations in the full experiment seems to have helped, it did not completely clear the problem. This increase has a steep cost in evaluation time, so a better solution needs to be found. One idea would be to restrict the fitness to some intervals. For instance, the behavior of always staying home, while optimal in a strict sense, does not interest us, and it seems improbable that it would evolve into something else of interest. Setting a minimum exploration fitness could be useful. Likewise with a minimum accuracy fitness, although our second experiment showed that among the worst accuracy fitness strategies, there was room for improvement through evolution. Another possible change involves the accuracy fitness. If the rat is constantly staying in the same place, it stops allocating place cells to that place because it has maximized the knowledge it has of that position. It is neither exploring, no improving the accuracy of its internal map, only preventing it from decreasing. In other words, perhaps we should use a fitness were the individual that explore a tiny bit is strictly superior to the one that stays home for the whole duration of the evaluation. Also, the individuals maximizing the surface covered fitness displayed behaviors never observed on real rats. This might lead to the conclusion that that fitness is not really biologically plausible, and could be improved.

Another area of improvement would be more thorough research on the inputs of the controllers. While some testing was done, the current group of inputs is more a result of using what worked first, and not a real in depth search. More inputs would, at least theoretically, allow for new diverse behaviors.

The final step of this research would be a robotic simulation, where the outputs of the controller would be mechanical decisions, such as wheel movement. The SFERES module Fastsim is made for this very purpose, unfortunately technical bugs and lack of time to fix them postponed this experiment.

Acknowledgments

I would like to thank Angelo Arleo and the entire ANC team for welcoming me amongst them, helping me during the entire internship, and introducing me to other lab members and researchers. Thank you Denis Sheynikhovich, Luca L. Bologna, Louis-Emmanuel Martinet, Jean-Baptiste Passot, Jérémie Pinoteau & Eléonore Duvelle.

I would also like to thank Stéphane Doncieux and Jean-Baptiste Mouret for their help with using Sferes and the ISIR cluster computer.

Bibliography

- Arleo J, Rondi-Reig L (2007) Multimodal sensory integration and concurrent navigation strategies for spatial cognition in real and artificial organisms, *J Integr Neurosci*, 6(3):327-66.
- Armstrong JD, Huntingford FA, Herbert NA (1999) Individual space use strategies of wild juvenile Atlantic salmon, *J Fish Biol* 55:1201-12.
- Conradt L, Bodsworth EJ, et al. (2000) Non-random dispersal in the butterfly *Maniola jurtina*: implications for metapopulation models. *Proc Biol Sci* 267(1452):1505-10.
- Eichenbaum H, Dudchenko P, et al. (1999) The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron* 23(2):209-26.
- Eilam D and Golani I (1989) Home base behavior of rats (*Rattus norvegicus*) exploring a novel environment. *Behav Brain Res* 34(3):199-211.
- Golani I, Benjamini Y, et al. (1993) Stopping behavior: constraints on exploration in rats (*Rattus norvegicus*). *Behav Brain Res*, 53(1-2):21-33.
- Halliday MS (1966). The effect of previous exploratory activity on the exploration of a simple maze. *Nature* 209:432-3.
- Hebb, D. O. (1949). *The organization of behavior*. Wiley, New York.
- Deb K. et al. (2002) *IEEE transactions on evolutionary computation*, vol. 6, NO. 2,
- Montgomery KC (1953). The effect of activity deprivation upon exploratory behavior. *J Comp Physiol Psychol* 46(6):438-41.
- Montgomery, KC (1953). The effect of the hunger and thirst drives upon exploratory behavior. *J Comp Physiol Psychol* 46(5):315-9.
- Montgomery KC and Monkman JA (1955). The relation between fear and exploratory behavior. *J Comp Physiol Psychol*, 48(2):132-6.
- Mouret, J.-B. and Doncieux, S. (2010). Sferes_v2: Evolvin' in the Multi-Core World. *WCCI 2010 IEEE (CEC)*. 4079-4086.
- O'Keefe J (1990). A computational theory of the hippocampal cognitive map. *Prog Brain Res* 83:301-12.
- Rosenblatt, Frank (1957), *The Perceptron--a perceiving and recognizing automaton*. Report 85-460-1, Cornell Aeronautical Laboratory.
- Sheynikhovich et al. (2009) Is There a Geometric Module for Spatial Orientation? Insights From a Rodent Navigation Model. *Psychological Review*, Vol. 116, No. 3, 540–566
- Taube, JS, Muller RU, Ranck JB Jr. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. *J. Neurosci.* 10(2): 420–435.
- Wilson M. A, McNaughton B. L. (1993) Dynamics of the Hippocampal Ensemble Code for Space. *Science* Vol. 261, 1055-1058