Integration of navigation and action selection functionalities in a computational model of cortico-basal ganglia-thalamo-cortical loops

Benoît Girard, David Filliat, Jean-Arcady Meyer, Alain Berthoz, Agnès Guillot

Integration of navigation and action selection functionalities in a computational model of cortico-basal ganglia-thalamo-cortical loops

Benoît Girard[1,2,*], David Filliat[3], Jean-Arcady Meyer[1],

Alain Berthoz[2] and Agnès Guillot[1]

[1] *AnimatLab/LIP6, CNRS - University Paris 6*

[2] *LPPA, CNRS - Collège de France*

[3] *DGA/Centre Technique d'Arcueil*

This article describes a biomimetic control architecture affording an animat both action selection and navigation functionalities. It satisfies the survival constraint of an artificial metabolism and supports several complementary navigation strategies. It builds upon an action selection model based on the basal ganglia of the vertebrate brain, using two

[*] *correspondence to:* B. Girard, CNRS LPPA, Collège de France, 11 place Marcellin Berthelot, 75231 Paris Cedex 05, France, , *Tel.:* +33-144271391, *Fax:* +33-144271382 *E-mail:* benoit.girard@college-de-france.fr

interconnected cortico-basal ganglia-thalamo-cortical loops: a ventral one concerned with appetitive actions and a dorsal one dedicated to consummatory actions.

The performances of the resulting model are evaluated in simulation. The experiments assess the prolonged survival permitted by the use of high level navigation strategies and the complementarity of navigation strategies in dynamic environments. The correctness of the behavioral choices in situations of antagonistic or synergetic internal states are also tested. Finally, the modelling choices are discussed with regard to their biomimetic plausibility, while the experimental results are estimated in terms of animat adaptivity.

Keywords: action selection, navigation, basal ganglia, computational neuroscience

Short title: Basal ganglia model of action selection and navigation.

1 Introduction

The work described in this paper contributes to the *Psikharpax* project, which aims at building the control architecture of a robot reproducing as accurately as possible the current knowledge of the rat's nervous system (Filliat *et al.*, 2004), it thus concerns biomimetic modelling derived from data gathered with rats. The main purpose of the *Psikharpax* project is to refocus on the seminal objective advocated by the animat approach: building "a whole iguana" (Dennett, 1978), instead of designing isolated and disembodied functions.

Indeed, in the animat literature, a great deal of work is devoted to the design of isolated control architectures that provide either action selection or navigation abilities –two fun-

damental functions for an autonomous system. The main objective of robotic navigation architectures is to afford an animat with various orientation strategies, like dead-reckoning, taxon navigation, place-recognition or planning (Filliat and Meyer, 2003, Meyer and Filliat, 2003 for reviews). The main objective of action selection architectures is to maintain the animat into its "viability zone", defined by the state space of its "essential variables" (Ashby, 1952), through efficient switches between various actions (Prescott et al., 1999 for a review). Even if there is evidence that an effective animat requires the use of these two functionalities, few models attempt to integrate them, taking into account the specific characteristics of each.

On the one hand, most of the navigation models insert arbitration mechanisms typical of action selection to solve spatial issues (e.g., Rosenblatt and Payton, 1989), but they do not take into account motivational constraints.

On the other hand, action selection models always integrate navigation capacities ensuring an animat the ability to reach resources in the environment, but they typically implement only rudimentary navigation strategies –random walk and taxon navigation– (e.g., Maes, 1991, Seth, 1998).

The few models that process both navigation and action selection issues are inspired by biological considerations, indicating that the hippocampal formation, in association with the prefrontal cortex, processes spatial information (O'Keefe and Nadel, 1978), whereas the basal ganglia are hypothesized to be a possible neural substrate for action selection in the vertebrate brain (Redgrave et al., 1999).

For example, Arleo and Gerstner (2000) propose a model of the hippocampus that

elaborates an internal map with the creation of several "place cells", used by an animat to reach two different kinds of resources providing rewards. The outputs of the model are assumed to be four action cells, coding for displacements in cardinal directions, and assumed to belong to the nucleus accumbens. This nucleus, located in the ventral part of the basal ganglia, is hypothesized to integrate sensorimotor, motivational and spatial information (Kelley, 1999). In this model, it selects the actual displacement by averaging the ensemble activity of the action cells. However, the animat does not select other navigation strategies and does not have a virtual metabolism that puts constraints on the timing and efficiency of the selection of its behaviors.

Guazzelli et al. (1998) endow their simulated animat with two navigation strategies (place-recognition-triggered and taxon navigation, processed by hippocampus and prefrontal cortex) and homeostatic motivational systems (hunger and thirst, processed by hypothalamus). Here, the role of the basal ganglia is limited to computing of reinforcement signals associated with motivational states, while action selection properly occurs in the premotor cortex. Yet, in this work, there are no virtual metabolism constraints on action selection and because of the choice of a systems-interaction level of modelling, the internal operation of the modules is not specifically biomimetic.

Gaussier et al. (2000) endow a motivated robot (Koala^{TM}, K-Team) with a virtual metabolism –generating signals of hunger, thirst and fatigue– and a topological navigation capacity. A topological map is built in the hippocampus and used to build a graph of transitions between places in the prefrontal cortex, used for path planning. The motor output is assumed to be effected by action neurons in the nucleus accumbens, coding

for three egocentric motions (turn right, left, go straight). Motivational needs affect path planning by spreading activation into the prefrontal graph from the desired resources to the current location of the animat. They are transmitted to the action neurons, allowing the animat to reach one goal by several alternative paths, and to make compromises between different needs. Here, one navigation strategy only is used, while various complementary strategies coexist in animals.

These models do not entirely satisfiy the objectives of the fundamental functions, that is, dealing with survival constraints together with taking advantage of various complementary navigational strategies. Moreover, they do not exploit recent neurobiological findings concerning neural circuits devoted to the integration of these functions, involving two parallel and interconnected "cortico-basal ganglia-thalamo-cortical" loops (CBGTC, Alexander *et al.*, 1986), stacked on a dorsal to ventral axis, receiving sensorimotor (*dorsal* loop) and spatial (*ventral* loop) information.

We previously tested a computational model of action selection, inspired by the *dorsal* loop and designed by Gurney *et al.* (2001a,b, referred to here as 'GPR' after the authors'names), by replicating the Montes-Gonzalez et al. (2000) implementation in a survival task (Girard et al., 2003). To improve the survival of an artificial system in a complex environment, our objective is to add to this architecture a second circuit –simulating the *ventral* loop– which selects locomotor actions according to various navigation strategies: a *taxon* strategy, directing the animat towards the closest resource perceived, a *topological* navigation, building a map of the different places in the environment and using it for path planning, together with random *exploration*, mandatory to map unknown areas and allow-

ing the discovery of resources by chance. The interconnection of the *dorsal* and *ventral* loops is designed by means of bioinspired hypotheses. The whole model will be validated in several environments where the animat performs a simple survival task.

After describing the navigation and action selection systems and how they are interconnected, we will introduce the specific experimental setup (survival task and animat configuration). The results will concern tests on the animat's specific adaptive mechanisms and behaviors, involving topological and taxon navigation, opportunistic ability and conflict management in case of changes in the environment or internal state.

2 The control architecture

This model has been introduced in a brief preliminary form in Girard *et al.* (2004).

2.1 Navigation

The choice of the navigation model was based on functional and efficiency criteria: it had to provide the animat with the capabilities of building a cognitive map, localizing itself with respect to it, storing the location of resources and computing directions to reach these resources; these operations had to be performed in real time and had to be robust enough to cope with the physical limitations of a real robot. The navigation system proposed by Filliat (2001) was chosen as it provides the required features and has been validated on a real robot (Pioneer^{TM}, ActivMedia).

This model emulates hippocampal and prefrontal cortex functions. It builds a dense

topological map in which nodes store the allothetic sensory input that the animat can perceive at respective places in the environment. These inputs are mean gray levels perceived by a panoramic camera in each of 36 surrounding directions, and sonar readings providing distances to obstacles in eight surrounding directions. A link between two nodes memorizes at which distance and in which direction the corresponding places are positioned relative to each other, as measured by the idiothetic sensors of odometry. The position of the animat is represented by a probability distribution over the nodes.

The model also provides an estimation of disorientation (D), which varies from 0 when the estimate of location is good, to 1 when it is poor. D increases when the robot is creating new nodes (it is in an unmapped area) and only decreases when it spends time in well known areas. The model also provides two 36-component vectors indicating which directions to follow in order to either explore unmapped areas (**Expl**) or go back to known areas in order to decrease disorientation (**BKA**). If the animat does not regularly go back to known areas when it is very disoriented, the resulting cognitive map will not be reliable. Consequently, the addition of topological navigation to an action selection mechanism will put a new constraint on the latter, the one of keeping *Disorientation* as low as possible.

We provided the model with the ability to learn the localization of resources important to survival (e.g. loading station, dangerous area) in the topological map. It is learned by associating active nodes of the graph with the type of resources encountered using Hebbian learning. By specifying the type of resource currently needed to a path planning algorithm applied on the graph, a vector **P** of 36 values is produced, representing the proximity of that resource in 36 directions spaced by 10°. Such a vector can be produced for each type of

resource res, weighted by the motivation associated to that resource $m(res)$, and combined with the other ones to produce a generic path planning vector **Plan**. The combination is processed as follows:

$$\mathbf{Plan} = 1 - \prod_{res}(1 - m(res) \times \mathbf{P}(res)) \tag{1}$$

2.2 Action Selection System

–Figure 1 around here–

The action selection model presented here is an extension of the one used in Girard *et al.* (2003), the GPR model (Gurney *et al.*, 2001a). It is a neural network model built with leaky-integrator neurons, in which each nucleus in the BG is subdivided into distinct channels each modelled by one neuron (Figure 1), and each channel associated to an elementary action. Each channel of a given nucleus projects to a specific channel in the target nucleus, thereby preserving the channel structure from the input to the output of the BG circuit. The subthalamic nucleus (STN) is an exception as its excitation seems to be diffuse. Inputs to the BG channels are *Salience* values, assumed to be computed in specific areas in the cortex, and representing the commitment to perform the associated action. They take into account internal and external perceptions, together with a positive feedback signal coming from the thalamo-cortical circuit, which introduces some persistence in the action performance. Two parallel *selection* and *control* circuits within the basal ganglia serve to modulate interactions between channels. Finally, the selection operates

via disinhibition (Chevalier and Deniau, 1990): at rest, the BG output nuclei are tonically active and keep their thalamic and motor system targets under constant inhibition. The output channel that is the less inhibited is selected, and the corresponding action executed.

A principal original feature of our model is that two parallel CBGTC loops are modelled, one selecting consummatory actions and the other appetitive actions.

2.2.1 Dorsal loop

In the BG, the *dorsal loop* implicated in the selection of motor responses in reaction to sensorimotor inputs and corresponds to the one modelled in the previous robotic studies of the GPR (Montes-Gonzalez et al., 2000; Girard et al., 2003). Here we hypothesize that it will direct the selection of non-locomotor actions, which in the present case are limited to consummatory actions (robotic equivalents of eating, resting, etc.) (Figure 2). In this loop:

- input *Saliences* are computed with internal and external sensory data;
- at the output, a "winner-takes-all" selection occurs for the most disinhibited channel, as simultaneous partial execution of both reloading behaviors doesn't make sense.

–Figure 2 around here–

2.2.2 Ventral loop

The *ventral loop* can be subdivided into two distinct subloops (Thierry et al., 2000), originating from the core and shell regions of its input nucleus (nucleus accumbens or NAcc)

(Zham and Brog, 1992). In the present work, we will only retain the core subloop (that will be henceforth also called *ventral loop*), which has been proposed to play a role in navigation towards rewarding places (Mulder *et al.*, 2004; Martin and Ono, 2000). The interactions between the hippocampus, the prefrontal cortex and the NAcc core (Thierry et al., 2000) could be the substrate of a *topological* navigation strategy. *Taxon navigation* needs sensory information only and could therefore be implemented in the *dorsal* loop. However, it was reported that the lesion of the NAcc also impairs object approach (Seamans and Phillips, 1994). This is why, in our model, this strategy will also be managed by the *ventral loop*.

To summarize, we hypothesize that this loop will direct appetitive actions (robotics equivalent for looking for food, homing, etc.), suggesting displacements towards motivated goals (Figure 2).

The *ventral loop* is very similar –anatomically and physiologically– to the circuits of the *dorsal* loop: the dorsolateral ventral pallidum plays a role similar to the GP (Maurice et al., 1997), the medial STN is dedicated to the ventral circuits (Parent and Hazrati, 1995) as well as the dorsomedial part of the SNr (Maurice et al., 1999). Thus, despite probable differences concerning the influence of dopamine on ventral and dorsal input nuclei, it is also designed by a GPR model. However, a few differences are to be noted:

- *Saliences* are computed with internal and external sensory data: the *taxon* navigation needs distal sensory inputs to select a direction and all navigation strategies are modulated by the motivations. Additional data coming from the navigation system proposes motions on the basis of a *topological* navigation strategy and map updates

of current positions;

- each nucleus is composed of 36 channels, representing allocentric displacement directions separated by 10°;

- the lateral inhibitions which occur in the nucleus accumbens core are no longer uniform as in the *dorsal* loop, but increase with the angular distance between two channels (see eqn. 7), so that close directions compete less than opposite ones;

- at the output, the selection makes a compromise among all channels disinhibited above a fixed threshold. The direction chosen by the animat is computed by a vector sum of these channels, weighted by their magnitudes of disinhibition.

2.2.3 Interconnection of Basal Ganglia loops

Interconnections between the parallel CBGTC loops is needed to coordinate their respective selection processes. This is especially true here, when selections concerning navigation taken in the *ventral* loop –like following a planned path leading to a resource– might be conflicting with behavioral choices made by the *dorsal* loop –like resting. Four main hypotheses concerning interconnections between loops have been proposed in the rat's brain. Two of them (*Hierarchical pathway* (Joel and Weiner, 1994) and *Dopaminergic hierarchical pathway* (Joel and Weiner, 2000)) were discarded because they only allow unidirectional communication from ventral to dorsal loops, whereas bidirectional or dorsal-to-ventral communication was necessary to solve our conflicts. The two remaining possibilities are (1) the *Cortico-cortical pathway*: cortical interconnections between areas implied in different loops

could allow bidirectional flows of information between loops; and (2) the *Trans-subthalamic pathway* (Kolomiets et al., 2001, 2003): the segregation of loops is not perfectly preserved at the level of the STN, some neurons belonging to one loop are excited by cortical areas belonging to other loops, thus, parts of the SNr belonging to one loop can be excited by another loop (Figure 2).

We implemented the *trans-subthalamic* hypothesis, by distributing *dorsal* STN activation to the *ventral* outputs (see eqn. 10 and Figure 2). Selection of an action in the *dorsal* loop increases activity in the *dorsal* STN, which in turn increases activation of the *ventral* outputs, preventing any movement from occuring.

The precise mathematical description of the resulting model is given in appendix A.1.

3 Experimental setup

3.1 Environment and survival task

The experiments are performed in simulated 2D environments involving, as in Girard *et al.* (2003), the presence of "ingesting" and "digesting" zones, but with the addition of "dangerous" places. The animat has to reach "ingesting" zones in order to acquire *Potential Energy* (E_P), which it should convert into *Energy* (E) in "digesting" zones, in order to use it for behavior. Note that a full load of *Energy* allows the animat to survive only $33min$. Paths to reach these zones may contain dangerous areas to avoid.

The software used is a simulator programmed in C++, developed in our laboratory.

Walls and obstacles are made of segments colored on a 256 level grayscale. The effects of lighting conditions are not simulated: the visual sensors have a direct access to the color. The three type of resources are represented by $50cm \times 50cm$ squares of specific colors: the "ingesting" (E_p), "digesting" (E) and "dangerous" (DA) areas are respectively gray (127), white (255) and dark gray (31). They can be used by the animat when the distance between their centre and the centre of the animat is less than $70cm$ (i.e. when they occupy more than 60° of the visual field). The other gray objects have no impact on survival but help the navigation system discriminating places.

3.2 The animat

The animat is circular ($30cm$ diameter), and translation and rotation speeds are $40cm.s^{-1}$ and $10°.s^{-1}$ respectively. Its simulated sensors are:

- an omnidirectional linear camera providing the color of the nearest segment for every 10° surrounding sector,
- eight sonars with a $5m$ range, a directional incertitude of ±5° and a $\pm 10cm$ distance accuracy,
- encoders measuring self-displacements with an error of ±5% of the measured distance,
- a compass with a ±10° range of error of estimated direction.

The sonars are used by a low level obstacle avoidance reflex which overrides any decision taken by the BG model when the animat comes too close to obstacles. The navigation

model uses the camera, encoders and compass inputs. The BG model uses the camera input to compute nine external variables:

- Three 36-component vectors, **Prox**(DA), **Prox**(E_P) and **Prox**(E) providing the proximity of each type of resource in each direction. This measure is related to the angular size of the resource in the visual field with a 10° resolution, as it is obtained by counting the number of contiguous pixels of the resource color in a 7 pixels window centered on the direction considered. These vectors are the basis of the taxon navigation strategy.

- Three variables, $mProx(DA)$, $mProx(E_P)$ and $mProx(E)$ which are the max values of the components of $Prox$ vectors.

- Three Boolean variables, $A(DA)$, $A(E_P)$ and $A(E)$, which are true if the corresponding $mProx$ value is one (i.e. if the resource is less than $70cm$ away and thus usable).

These purely sensory inputs are completed by the vectors produced by the topological navigation system: the path planning vector **Plan**, the exploration vector **Expl** and the "go back to known areas" vector **BKA**.

The animat has four internal variables: *Energy* and *Potential Energy*, which concern the survival task (see 3.1), *Fear*, which is a constant, fixing the strength of the repellent effect of "dangerous areas" and *Disorientation*, which is provided by the topological navigation system (see 2.1). From these variables are derived four motivations used in *saliences* computations and in the weighting of the **Plan** vector (eqn. 1). The motivations to go back to known areas and to flee dangerous areas are respectively equal to the *Disorientation*

and *Fear* variables, while the motivation to reach *Energy* and *Potential Energy* resources are more complex:

$$\begin{aligned} m(DA) &= F \\ m(BKA) &= D \\ m(E) &= (1-E)\sqrt{1-(1-E_P)^2} \\ m(E_P) &= 1-E_P \end{aligned} \tag{2}$$

The variables used to compute saliences in each loop are summarized in Figure 2, and the details of these computations are given in appendix A.2.

4 Experiments

Three different experiments are carried out in simple environments in order to test the adaptive mechanisms the animat is provided with.

Experiment 1 tests the *efficiency of the navigation/action selection models interface*. An animat capable of topological navigation has to survive in an environment containing one resource of *Energy* and one resource of *Potential Energy* which cannot be seen simultaneously. It is compared to an animat using the taxon strategy only, the use of the topological navigation is expected to improve the survival time.

Experiment 2 tests adaptive action selection in a *changing environment*: on the one hand, the animat has to use a taxon strategy in order to reach newly appeared resources; on the other hand, it has to forget the location of exhausted resources to head towards abundant ones.

Experiment 3 tests adaptive action selection in case of *antagonistic or synergetic internal states*: on the one hand, in a situation where two paths lead to a resource and the shortest one includes a dangerous area; on the other hand, in a situation where a short path leads to one resource only, while a longer one leads to two resources satisfying two different needs.

In experiments 2 and 3, the animat is provided with a previously built map of the environment in order to allow statistical comparison of runs with identical initial conditions.

4.1 Experiment 1: Efficiency of the navigation/action selection interface

In this experiment, an animat traverses the environment ($7m \times 9m$) depicted in Figure 3: it contains one resource of E and one resource of E_P, but it is impossible to see one resource from the vicinity of the other. In the first model configuration (*condition A*), the animat uses both object approach and topological navigation strategies, whereas in the other one (*condition B*), the animat uses object approach only. The "reactive" animat (*condition B*), following taxon strategy only, has to rely on random exploration to find hidden resources. In contrast, after a first phase of random exploration and map building, the animat in *condition A* should be able to reach desired resources using its topological map.

–Figure 3 around here–

Ten tests, with a four-hour duration limit, are run for both animats. *Energy* and *Potential Energy* are initially set to 1. The comparison of the median of survival durations

for both sets shows that in *condition A*, the animat is able to survive significantly longer ($p < 0.01$, U-test, see Table 1) than the animat in *condition B*.

–Table 1 around here–

In (Girard et al., 2003), action selection was only constrained by the virtual metabolism. Here, the addition of the *topological navigation* system generates a new constraint of limiting *Disorientation*. Yet it does not affect the efficiency of action selection, as the life span of animats is enhanced.

4.2 Experiment 2: Changing environment

–Figure 4 around here–

This experiment takes place in the $6m \times 6m$ environment depicted in Figure 4, where the second *Potential Energy* resource is not always present.

4.2.1 New resources: Coordination of the navigation strategies

In this case, the second *Potential Energy* resource is not present during the mapping phase, so that when the animat reaches the first intersection, it perceives a new resource that is unknown by the topological navigation system. The topological and the taxon strategies are thus competing, the first one suggesting to move to the distant resource (E_P1) and the second to the newly appeared and closer resource (E_P2). For all tests, the animat is initially placed on the same location shown in Figure 4 and lacks *Potential Energy* ($E = 1$ and $E_p = 0.5$). The tests are stopped when the animat activates the $ReloadE_P$ action.

The control experiment consisting of ten tests in which resource E_P2 is not added,

results in a repeatable behavior of the animat: it goes directly to E_P1 and activates the $ReloadE_P$ action when close enough to E_P1. Three series of fifteen tests, with different weightings of the salience computations (variations of eqn. 17 in appendix using the weights of Table 2), are compared by counting how many times the animat chose one resource versus the other. The results are summarized in Table 2.

–Table 2 around here–

The first weighting corresponds to the configuration used in the previous experiment (eqn. 17). The path planning weight is larger than the taxon strategy one. As a result, the animat often ignores the new resource and chooses the memorized one. When the relative importance of the two strategies is modulated by progressively lowering the path planning weight, the behavior of the animat is modified and an opportunistic behavior, where it prefers the new and closest resource, can be obtained.

Consequently, if our control architecture does not intrinsically exhibit an opportunistic or a pure planning behavior, it can easily be tuned to generate the desired balance between these two extremes.

4.2.2 Exhausted resources: Forgetting mechanism

In this situation, resource E_P2 is present during mapping but is removed during the tests. The animat then has to "forget" its existence in the map in order to go to the other resource.

Fifteen tests are carried out, with the animat initially placed on the same start location (see Figure 4) lacking *Potential Energy* ($E = 1$ and $E_p = 0.5$). The tests are stopped when

the animat activates the $ReloadE_P$ action.

The animat first goes to the closest E_P resource coded by the topological navigation system: the near but absent E_P2 resource. The forgetting mechanism (implemented by the Hebbian rule used to link resources with locations on the map) allows the animat to finally leave this area and to reach resource E_P1. The time necessary to forget E_P2 is estimated by subtracting the duration of the most direct path leading from the start position to E_P1 *via* E_P2 ($46s$) to the duration of each test. The mean duration is $178s$ ($\sigma = 78$), i.e. 2 minutes and 58 seconds (max value 5 minutes). It is a bit long (almost 10% of the 33 minutes survival duration with a full charge of *Energy*), but it can be reduced by simply modifying the gain of the Hebbian rule.

This shows that the ability to forget, which is necessary to survive in environments where resources are exhaustible, operates correctly.

4.3 Experiment 3: Antagonistic or synergetic internal states

4.3.1 Antagonistic internal states: Fear vs reloading need

–Figure 5 around here–

A first experiment is run in an environment ($10m \times 6m$) containing two E_P resources and a dangerous area blocking direct access to the closest one (Figure 5). The *Dangerous Areas* affect the planning algorithm of the topological navigation system in an inhibitory manner. A path planning vector leading to dangerous areas is computed, multiplied by the level of *Fear* and subtracted to the other planning vectors: the term $-m(DA) \times \mathbf{P}(DA)$

is added to the computation of **Plan** described in eqn. 1.

The animat initially lacks *Potential Energy* and its level of *Fear* is fixed ($E = 1$, $E_P < 1$, $F = 0.2$). When the *Dangerous Area* is absent, the animat systematically chooses the closest resource (E_P1). However, when it is present, this inhibits the drive to go towards the E_P1 resource and the final choice of the E_P resource should thus depend on the importance of the lack of energy.

–Table 3 around here–

Two series of 20 tests are carried out in order to induce conflicts between internal states depending on *Fear* and E_P, respectively with a moderate ($E_P = 0.5$) and a strong ($E_P = 0.1$) lack of E_P. As illustrated in Table 3, the inhibition generated by the *Dangerous Area* in the first case is strong enough and the animat, despite the longer route, selects E_P2. In the second one, the need for *Potential Energy* is stronger and the animat, despite the danger, selects E_P1. These two opposite tendencies are significantly different (Fischer's exact probability test, $p < 0.01$).

This experiment shows that the animat may take risks in emergency situations and avoid them otherwise. But, more generally, it shows that it can exhibit, in an identical environmental configuration, different behavioral choices adapted to its conflicting internal needs, an essential property for a motivated animat.

4.3.2 Synergetically interacting motivations

–Figure 6 around here–

This task is inspired by a T-maze experiment proposed in Quoy *et al.* (2002) in order

to study the behavior generated by the coupling of two motivations. The left branch of the T contains one E_P resource while the right one contains both an E and an E_P resource (Figure 6). The length of the right branch is varied so that the ratio of the right branch length to the left branch length is 1, 1.5 or 2. The animat is initially placed in the lower branch of the T, with a motivation for both E and E_P ($E = 0.5$ and $E_P = 0.5$). The test stops when the animat activates the $ReloadE_P$ action. In such a situation, the animat is expected to systematically prefer the right branch, even if it is longer, because choosing the left only satisfies the E_P need, while choosing the right can satisfy both E and E_P needs.

–Table 4 around here–

Three series of fifteen tests are carried out with branch length ratio values of 1, 1.5 and 2, with an animat that needs both E and E_P. As long as the ratio is not too high, the cumulated activation generated by the two resources on the right is higher than the drive generated by the single E_P resource on the left (Table 4, ratio 1 and 1.5). However, when the two resources on the right are too far away, the drive they generate is attenuated by distance and the animat becomes more and more attracted by the resource on the left (Table 4, ratio 2).

The Gaussier *et al.* (2000) model of navigation integrates the notion of "preferred path" by reducing the apparent distance between two nodes of the map when they are often used. This allows the right branch to become preferred and thus systematically chosen over time. Future development of our model should include such a habit learning capability.

5 Discussion

The proposed biomimetic model integrates both navigation and action selection, in taking into account the specificities of both survival constraint and variety of navigation strategies. Simulations in benchmark environments validate 1) the survival advantage of using path planning strategies, 2) the benefits of simultaneously using taxon and planning strategies along with the necessity of being able to forget when operating in changing environments, and 3) the capability of the model to behave adaptively in case of conflicting and synergetic motivations.

5.1 From Rattus rattus...

How the brain coordinates the interface between spatial maps, motivation, action selection and motor control systems is of timely interest. The rat brain is widely investigated in this purpose, but many issues remain to be clarified. By synthesizing observed mechanisms in a behaving artificial system, our work helps to formulate several questions.

For example, our model points out limitations about the current neurobiological knowledge concerning the actual role of NAcc core channels: do they represent, as in our model and in e.g., Strösslin (2004), competing directions of movements? In Experiment 2.1, the level of opportunism is fixed and does not adapt to changing conditions (whereas taxon navigation is less reliable in poor lighting conditions), as the *ventral* loops selects one direction taking into account all the navigation strategies. This could be changed by having it selecting among the strategies the most adapted one before a *dorsal* loop selects the

direction of motion based on the chosen strategy suggestion only. Such coding has recently received support by the work of Mulder *et al.* (2004), on the basis of electrophysiological recordings in hippocampal output structures associated with the NAcc and a nucleus of the dorsal stream (ventromedial caudate nucleus). Another and more complex role may also be considered: NAcc core could interface goals, their location, their amount and the corresponding motivations with information coming from several neural structures like other limbic structures or CBGTC loops (Dayan and Balleine, 2000).

Likewise, our model questions the putative substrates of interactions between CBGTC loops and their mode of operation, a subject of active current research. We may have implemented the trans-subthalamic hypothesis in an exaggerated manner. In fact the overlap of STN projections from various loops is rather limited (Kolomiets et al., 2003), while in our model they extensively reach the whole output of the *ventral loop*. This choice was indeed convenient for the role attributed here to the dorsal and ventral channels, respectively coding for immobile and mobile actions. Recent results relative to interactions at the level of BG output projections to dopaminergic nuclei in rats (Mailly et al., 2003) shed a new light on the *dopamine hierarchical pathway* and could be the basis of an alternative model. In the GPR, varying the dopamine level affects directly the ability to select, therefore, the possibility that one loop may modulate the dopamine level of another one could be the basis of an alternative mechanism for a loop to shunt another loop. One cannot finally exclude the possibility that the resolution of selection conflicts in the CBGTC loops is not only managed in the BG but also in downstream brainstem structures, for example in the reticular formation (Humphries *et al.*, this issue).

5.2 ...to Psikharpax

In Experiment 1, the planning animat (*condition A*) sometimes dies because of a imperfect hand-tuning of the salience computations, which causes it to stop to reload too far away from resources. The basal ganglia, in interaction with the dopaminergic system, is supposed to be the neural substrate for reinforcement learning. In order to avoid such problems in the future, we are now adding such a mechanism of automatic optimization of salience computations to our model (Khamassi et al., 2004).

As mentioned in introduction, this work contributes to the *Psikharpax* project, which aims at building an artificial rat (Filliat *et al.*, 2004). As it evolves, this artificial rat will be endowed with more than the few motivations taken into account here, in the aim to improve the actual autonomy of current robots, often devoted to a single task. The development of polyvalent artifacts working in natural environments is indeed promising for many applications in the home or in the office, as well as future space programs with unmanned missions. Our work also helps assessing the operational value of the biomimetic models used for this purpose.

A Appendix: Mathematical model description

A.1 GPR structure

Activation (a) of every neuron of the model:

$$\tau \frac{da}{dt} = I - a \tag{3}$$

where: I: input of the neuron, τ: time constant ($\tau = 25ms$). Corresponding output (y):

$$y = \begin{cases} 0 & \text{if } a < \epsilon \\ m \times (a - \epsilon) & \text{if } \epsilon \leq a < \epsilon + 1/m \\ 1 & \text{if } \epsilon + 1/m \leq a \end{cases} \tag{4}$$

Values of ϵ and m for each nucleus in Table 5.

–Table 5 around here–

In each module (D1 and D2 striatum subparts, STN, EP/SNr, GP, VL, TRN and cortical feedback), the input of each channel i is defined by the equations 5 to 14, where N: number of channels, S_i: salience of channel i, λ: dopamine level (0.2).

$$I^i_{D1} = (1 + \lambda)S_i - \sum_{\substack{j=0 \\ j \neq i}}^{N} y^i_{D1} \tag{5}$$

$$I^i_{D2} = (1 - \lambda)S_i - \sum_{\substack{j=0 \\ j \neq i}}^{N} y^i_{D2} \tag{6}$$

In our model of the *ventral* loop, lateral inhibitions (sum terms in eqn. 5 and 6) increase with the angular difference between two channels. They are replaced in the *ventral* loop by the following LI term:

$$LI^i = \sum_{\substack{j=0 \\ j \neq i}}^{N} \frac{|i-j| mod(N/2)}{N/2} \times y^i_{(D1 \text{ or } D2)} \tag{7}$$

$$I^i_{STN} = S_i - y^i_{GP} \tag{8}$$

$$I^i_{EP} = -y^i_{D1} - 0.4\, y^i_{GP} + 0.8 \sum_{j=0}^{N} y^j_{STN} \tag{9}$$

The *trans-subthalamic* pathway is modelled by a modified input for the *ventral* EP/SNr (v and d stand for *ventral* and *dorsal*):

$$\begin{aligned} I^i_{EPv} = &- y^i_{D1v} - 0.4\, y^i_{GPv} \\ &+ 0.8 \sum_{j=0}^{N} y^j_{STNv} + 0.4 \sum_{j=0}^{N} y^j_{STNd} \end{aligned} \tag{10}$$

$$I^i_{GP} = -y^i_{D2} + 0.8 \sum_{j=0}^{N} y^j_{STN} \tag{11}$$

$$I^i_{VL} = y^i_P - y^i_{EP} - 0.13 \sum_{\substack{j=0 \\ j \neq i}}^{N} y^j_{TRN} \tag{12}$$

$$I^i_{TRN} = y^i_{VL} + y^i_P \tag{13}$$

$$I_P^i = y_{VL}^i \tag{14}$$

A.2 Salience computations

The modification to the GPR model proposed in Girard *et al.* (2003) consisted in allowing, for the computation of saliences, the use of sigma-pi neurons and non-linear transfert function applied to the inputs. This was kept in the present model and is the origin of the square roots and multiplications in the following equations.

A.2.1 Experiments 1 and 2

Dorsal loop saliences (E and E_P reloading actions):

$$\begin{aligned} S_E = 0.4 \times P_E + 1.2 \times A(E) \times m(E) \\ + 0.6 \times mProx(E) \times m(E) \end{aligned} \tag{15}$$

$$\begin{aligned} S_{E_P} = 0.4 \times P_{E_P} + A(E_P) \times m(E_P) \\ + 0.2 \times mProx(E_P) \times m(E_P) \end{aligned} \tag{16}$$

Ventral loop salience for each direction i:

$$
\begin{aligned}
S_i = {} & 0.2 \times P_i + W_{plan}\sqrt{\mathbf{Plan}_i} \\
& + 0.55\sqrt{\mathbf{Prox}(\mathbf{E})_i} \times m(E) \\
& + W_{taxon}^{E_p}\sqrt{\mathbf{Prox}(\mathbf{E_P})_i} \times m(E_P) \\
& + 0.4 \times \mathbf{BKA}_i \times m(BKA) \\
& + \mathbf{Exp}_i \times (0.25 \\
& \quad + 0.05 \times (1 - mProx(E_P)) \times m(E_P) \\
& \quad + 0.05 \times (1 - mProx(E)) \times m(E))
\end{aligned}
\tag{17}
$$

Where W_{plan} and $W_{taxon}^{E_p}$ are respectively set to 0.65 and 0.55, except in experiment 4.2.1, where they take the values recorded in Table 2.

A.2.2 Experiment 3.1

Saliences of the *dorsal* loop computed as in experiments 1 and 2. *Ventral* saliences modified to include the avoidance of dangerous areas:

$$\begin{aligned} S_i = {} & 0.2 \times P_i + 0.45\sqrt{\mathbf{Plan}_i} \\ & + 0.35\sqrt{\mathbf{Prox}(\mathbf{E})_i} \times m(E) \\ & + 0.35\sqrt{\mathbf{Prox}(\mathbf{E_P})_i} \times m(E_P) \\ & + 0.19 \times (1 - \mathbf{Prox}(\mathbf{DA})_i) \times m(DA) \\ & + 0.4 \times \mathbf{BKA}_i \times m(BKA) \\ & + \mathbf{Exp}_i \times (0.05 \\ & \quad + 0.05 \times (1 - mProx(E_P)) \times m(E_P) \\ & \quad + 0.05 \times (1 - mProx(E)) \times m(E)) \end{aligned} \tag{18}$$

A.2.3 Experiment 3.2

Experiment 3.2 showed that the weight of the *dorsal* computations had to be lowered:

$$\begin{aligned} S_E = {} & 0.4 \times P_E + 0.9 \times A(E) \times m(E) \\ & + 0.1 \times mProx(E) \times m(E) \end{aligned} \tag{19}$$

$$\begin{aligned} S_{E_P} = {} & 0.4 \times P_{E_P} + 0.9 \times A(E_P) \times m(E_P) \\ & + 0.1 \times mProx(E_P) \times m(E_P) \end{aligned} \tag{20}$$

The *ventral* salience computations from experiments 1 and 2 risked stopping the animat too far from resources. As this problem arose systematically in experiment 3.2, the term $(0.65\sqrt{\mathbf{Plan}_i})$ term was changed for $(0.55\sqrt{\mathbf{Plan}_i} \times (1 - mProx(E)) \times (1 - mProx(E_P))$.

Acknowledgments

We thank Sidney Wiener for valuable discussions and proofreading of the manuscript. This research has been funded by the LIP6 and the Project *Robotics and Artificial Entities* (ROBEA) of the French Centre National de la Recherche Scientifique.

References

Alexander, G. E., DeLong, M. R., and Strick., P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, 9:357–381.

Arleo, A. and Gerstner, W. (2000). Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity. *Biological Cybernetics*, 83:287–299.

Ashby, W. R. (1952). *Design for a brain*. Chapman and Hall.

Chevalier, G. and Deniau, M. (1990). Disinhibition as a basic process of striatal functions. *Trends in Neurosciences*, 13:277–280.

Dayan, P. and Balleine, B. (2000). Reward, motivation and reinforcement learning. *Neuron*, 36:285–298.

Dennett, D. (1978). Why not a whole iguana? *Behavioral and Brain Sciences*, 1:103–104.

Filliat, D. (2001). *Cartographie et estimation globale de la position pour un robot mobile autonome*. PhD thesis, LIP6/AnimatLab, Université Paris 6, France.

Filliat, D., Girard, B., Guillot, A., M.Khamassi, Lachèze, L., and Meyer, J.-A. (2004). State of the artificial rat psikharpax. In Schaal, S., Ijspeert, A., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J.-A., editors, *From Animals to Animats 8*, pages 2–12. MIT Press, Cambridge, MA.

Filliat, D. and Meyer, J.-A. (2003). Map-based navigation in mobile robots. i. a review of localization strategies. *Journal of Cognitive Systems Research*, 4(4):243–282.

Gaussier, P., Leprêtre, S., Quoy, M., Revel, A., Joulain, C., and Banquet, J.-P. (2000). Experiments and models about cognitive map learning for motivated navigation. In Demiris, J. and Birk, A., editors, *Interdisciplinary Approaches to Robot Learning*, volume 24 of *Robotics and Intelligent Systems*, pages 53–94. World Scientific.

Girard, B., Cuzin, V., Guillot, A., Gurney, K. N., and Prescott, T. J. (2003). A basal ganglia inspired model of action selection evaluated in a robotic survival task. *Journal of Integrative Neuroscience*, 2(2):179–200.

Girard, B., Filliat, D., Meyer, J.-A., Berthoz, A., and Guillot, A. (2004). An integration of two control architectures of action selection and navigation inspired by neural circuits in

the vertebrate: the basal ganglia. In Bowman, H. and Labiouse, C., editors, *Connectionist Models of Cognition and Perception II*, volume 15 of *Progress in Neural Processing*, pages 72–81. World Scientific, Singapore.

Guazzelli, A., Corbacho, F. J., Bota, M., and Arbib, M. A. (1998). Affordances, motivations and the worls graph theory. *Adaptive Behavior: Special Issue on biologically inspired models of spatial navigation*, 6(3/4):435–471.

Gurney, K., Prescott, T. J., and Redgrave, P. (2001a). A computational model of action selection in the basal ganglia. i. a new functional anatomy. *Biological Cybernetics*, 84:401–410.

Gurney, K., Prescott, T. J., and Redgrave, P. (2001b). A computational model of action selection in the basal ganglia. ii. analysis and simulation of behaviour. *Biological Cybernetics*, 84:411–423.

Joel, D. and Weiner, I. (1994). The organization of the basal ganglia-thalamocortical circuits: open interconnected rather than closed segregated. *Neuroscience*, 63:363–379.

Joel, D. and Weiner, I. (2000). The connections of the dopaminergic system with the striatum in rats and primates: An analysis with respect to the functional and compartmental organization of the striatum. *Neuroscience*, 96(3):452–474.

Kelley, A. E. (1999). Neural integrative activities of nucleus accumbens subregions in relation to learning and motivation. *Psychobiology*, 27:198–213.

Khamassi, M., Girard, B., Berthoz, A., and Guillot, A. (2004). Comparing three critic models of reinforcement learning in the basal ganglia connected to a detailed actor part in a s-r task. In Groen, F., Amato, N., Bonarini, A., Yoshida, E., and Kröse, B., editors, *Proceedings of the Eighth International Conference on Intelligent Autonomous Systems (IAS8)*, pages 430–437. IOS Press, Amsterdam, The Netherlands.

Kolomiets, B. P., Deniau, J. M., Glowinski, J., and Thierry, A.-M. (2003). Basal ganglia and processing of cortical information: functional interactions between trans-striatal and trans-subthalamic circuits in the substantia nigra pars reticulata. *Neuroscience*, 117(4):931–938.

Kolomiets, B. P., Deniau, J.-M., Mailly, P., Glowinski, A. M. J., and Thierry, A.-M. (2001). Segregation and convergence of information flow through the cortico-subthalamic pathways. *Journal of Neuroscience*, 21(15):5764–5772.

Maes, P. (1991). A bottom-up architecture for behavior selection in an artificial creature. In Meyer, J.-A. and Wilson, S., editors, *From Animals to Animats 1*, pages 478–485. MIT Press, Cambridge, MA.

Mailly, P., Charpier, S., Menetrey, A., and Deniau, J.-M. (2003). Three-dimensional organization of the recurrent axon collateral network of the substantia nigra pars reticulata neurons in the rat. *Journal of Neuroscience*, 23(12):5247–5257.

Martin, P. D. and Ono, T. (2000). Effects of reward anticipation, reward presentation, and

spatial parameters on the firing of single neurons recorded in the subiculum and nucleus accumbens of freely moving rats. *Behavioural Brain Research*, 116(1):23–28.

Maurice, N., Deniau, J.-M., Glowinski, J., and Thierry, A.-M. (1999). Relationships between the prefrontal cortex and the basal ganglia in the rat: physiology of the cortico-nigral circuits. *Journal of Neuroscience*, 19(11):4674–4681.

Maurice, N., Deniau, J.-M., Menetrey, A., Glowinski, J., and Thierry, A.-M. (1997). Position of the ventral pallidum in the rat prefrontal cortex-basal ganglia circuit. *Neuroscience*, 80(2):523–534.

Meyer, J.-A. and Filliat, D. (2003). Map-based navigation in mobile robots. ii. a review of map-learning and path-planning strategies. *Journal of Cognitive Systems Research*, 4(4):283–317.

Montes-Gonzalez, F., Prescott, T. J., Gurney, K. N., Humphries, M., and Redgrave, P. (2000). An embodied model of action selection mechanisms in the vertebrate brain. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S. W., editors, *From Animals to Animats 6*, volume 1, pages 157–166. The MIT Press, Cambridge, MA.

Mulder, A. B., Tabuchi, E., and Wiener, S. I. (2004). Neurons in hippocampal afferent zones of rat striatum parse routes into multi-pace segments during maze navigation. *European Journal of Neuroscience*, 19(7):1923–1932.

O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford.

Parent, A. and Hazrati, L.-N. (1995). Functional anatomy of the basal ganglia. ii. the place of subthalamic nucleus and external pallidum in basal ganglia circuitry. *Brain Research Reviews*, 20:128–154.

Prescott, T. J., Redgrave, P., and Gurney, K. N. (1999). Layered control architectures in robot and vertebrates. *Adaptive Behavior*, 7(1):99–127.

Quoy, M., Laroque, P., and Gaussier, P. (2002). Learning and motivational couplings promote smarter behaviors of an animat in an unknown world. *Robotics and Autonomous Systems*, 38:149–156.

Redgrave, P., Prescott, T. J., and Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, 89(4):1009–1023.

Rosenblatt, J. K. and Payton, D. (1989). A fine-grained alternative to the subsumption architecture for mobile robot control. In *Proceedings of the IEEE/INNS International Joint Conference on Neural Networks*, volume 2, pages 317–324. Washington, DC.

Seamans, J. K. and Phillips, A. G. (1994). Selective memory impairments produced by transient lidocaine-induced lesions of the nucleus accumbens in rats. *Behavioral Neuroscience*, 108:456–468.

Seth, A. K. (1998). Evolving action selection and selective attention without actions, attention, or selection. In Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S. W., editors, *From Animals to Animats 5*, pages 139–147. The MIT Press, Cambridge, MA.

Strösslin, T. (2004). *A Connectionist Model of Spatial Learning in the Rat*. PhD thesis, EPFL – Swiss Federal Institute of Technology, Swiss.

Thierry, A.-M., Gioanni, Y., Dégénetais, E., and Glowinski, J. (2000). Hippocampo-prefrontal cortex pathway: anatomical and electrophysiological characteristics. *Hippocampus*, 10:411–419.

Zahm, D. S. and Brog, J. S. (1992). Commentary: on the significance of the core-shell boundary in the rat nucleus accumbens. *Neuroscience*, 50:751–767.

Table 1: Comparison (U-test) of experiments testing median survival duration of animats in conditions A (taxon navigation only) and B (taxon and topological navigation).

Durations (s)	**Median**	**Range**
A	14431.5	2531 : 17274
B	4908.0	2518 : 8831
U test	U = 15	$p < 0.01$

Table 2: Resource choice depending on the relative weighting of the two navigation strategies in the salience computation. W_{plan} and $W_{taxon}^{E_p}$: weights related to planning and taxon navigation strategies respectively (see eqn. 17).

Weights		**Choices**	
$\mathbf{W_{plan}}$	$\mathbf{W_{taxon}^{E_P}}$	$\mathbf{E_P1}$	$\mathbf{E_P2}$
0.65	0.55	13	2
0.55	0.55	7	8
0.45	0.55	2	13

Table 3: Resource choice depending on the initial E_P level.

Internal state		**Incidence of choices**	
F	**E_P**	**E_P1**	**E_P2**
0.2	0.1	13	7
0.2	0.5	2	18
Fisher's test		**p¡**	0.01

Table 4: Branch choices depending on the length ratio.

	Incidence of first choice	
Ratio	**Left**	**Right**
1	3	12
1.5	4	11
2	8	7

Table 5: Parameters of the transfer functions of the GPR model.

GPR Module	ϵ	**m**
D1 Striatum	0.2	1
D2 Striatum	0.2	1
STN	-0.25	1
GP	-0.2	1
EP/SNr	-0.2	1
Ctx	0	1
TRN	0	0.5
VL	-0.8	0.62

Figure Captions

Figure 1: The GPR model. Nuclei are represented by boxes, each circle in these nuclei represents an artificial leaky-integrator neuron. On this diagram, three channels are competing for selection, represented by the three neurons in each nucleus. The second channel is represented by gray shading. For clarity, the projections from the second channel neurons only are represented, they are similar for the other channels. White arrowheads represent excitations and black arrowheads, inhibitions. D1 and D2: neurons of the striatum with two respective types of dopamine receptors; STN: subthalamic nucleus; GP: globus pallidus; EP/SNr: entopedoncular nucleus and substantia nigra pars reticulata; VL: ventrolateral thalamus; TRN: thalamic reticular nucleus. Dashed boxes represent the three subdivisions of the model proposed by its authors (Selection, Control of selection and thalamo-cortical feedback or TCF), note that these subdivisions appear on the simplified sketch of Figure 2.

Figure 2: Final model structure. Input variables are exhaustively listed, 36-component vectors are in bold type. The excitatory projections from the STN of the dorsal loop to the EP/SNr of the ventral loop, which are the substrate for loops coordination, are highlighted.

Figure 3: Experiment 1 environment. Initial position and orientation are represented by the schematic animat. E: *Energy* resource; E_P: *Potential Energy* resource.

Figure 4: Experiment 2 environment. Initial position and orientation are represented by the schematic animat. E_P: *Potential Energy* resource; E_P2 is absent in some experiments, see text.

Figure 5: Experiment 3 environment. Initial position and orientation are represented

by the schematic animat. E_P1,2: *Potential Energy* resources; DA: dangerous area.

Figure 6: The three environments of experiment 4. The ratio of the right branch length to the left branch length varied between 1 and 2. Initial position and orientation is represented by the schematic animat. E_P1,2: *Potential Energy* resources; E: *Energy* resource.

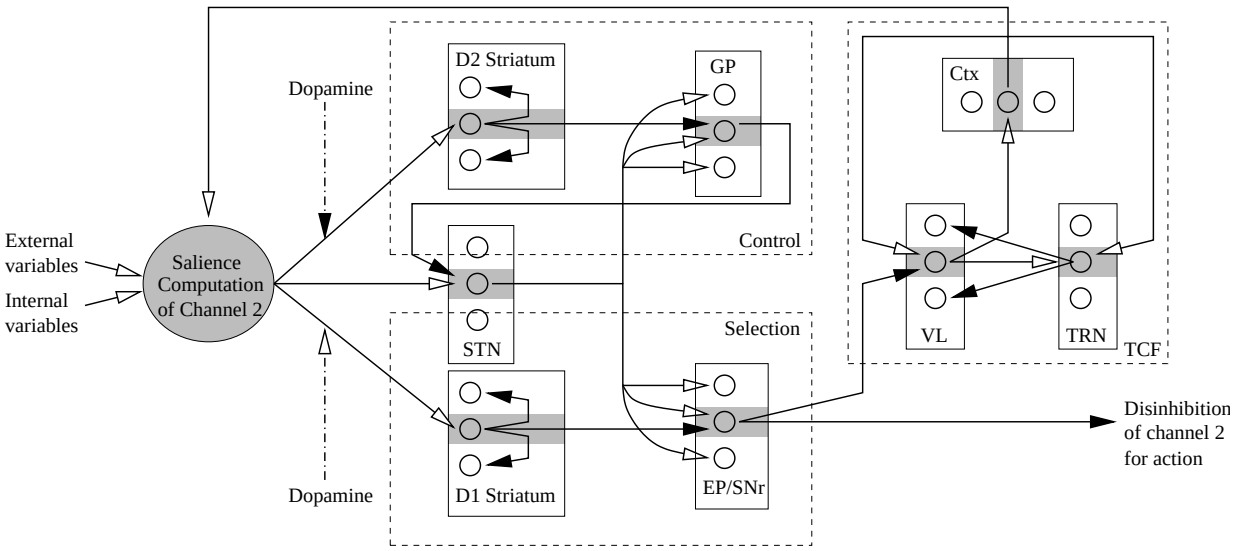
Dopamine
D2 Striatum
GP
Ctx
External
variables
Internal
variables
Salience
Computation
of Channel 2
Control
STN
Selection
VL
TRN
TCF
Disinhibition
of channel 2
for action
Dopamine
D1 Striatum
EP/SNr

Figure 1:

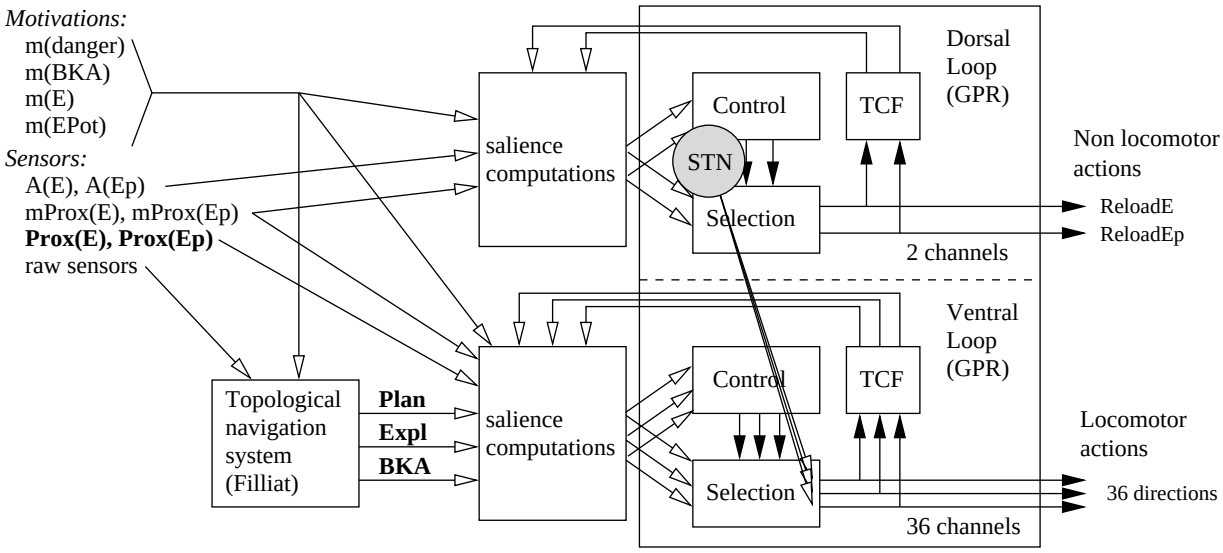
Motivations:
m(danger)
m(BKA)
m(E)
m(EPot)
Sensors:
A(E), A(Ep)
mProx(E), mProx(Ep)
Prox(E), Prox(Ep)
raw sensors
Topological navigation system (Filliat)
Plan
Expl
BKA
salience computations
Control
STN
Selection
TCF
Dorsal Loop (GPR)
2 channels
Non locomotor actions
ReloadE
ReloadEp
salience computations
Control
Selection
TCF
Ventral Loop (GPR)
36 channels
Locomotor actions
36 directions

Figure 2:

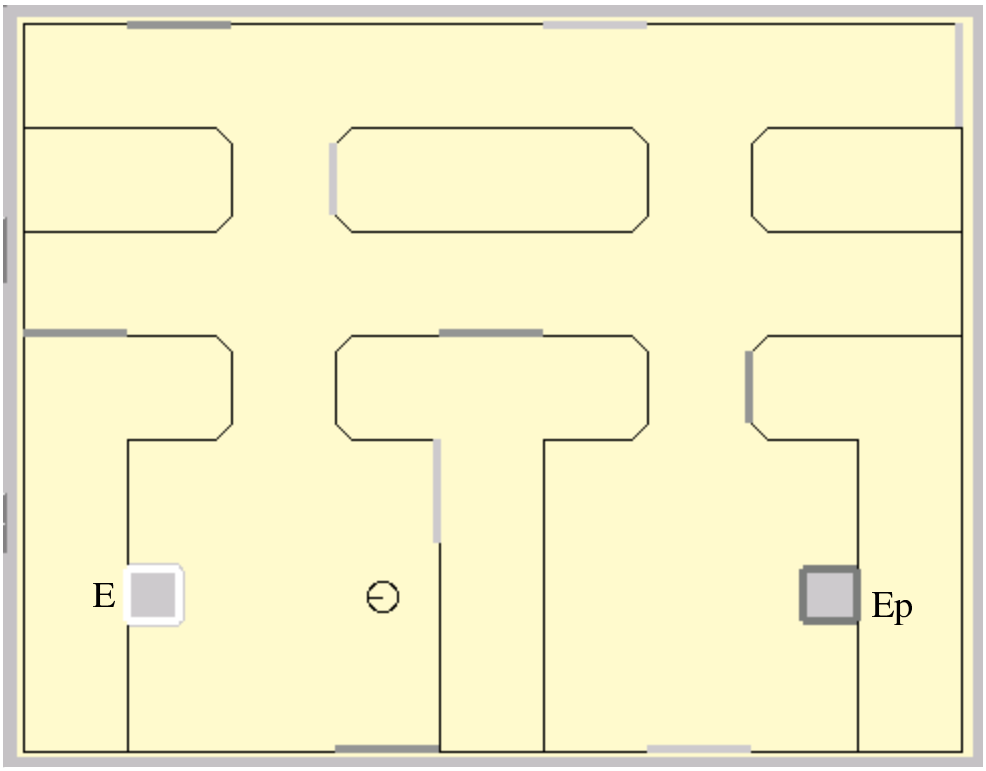

Figure 3:

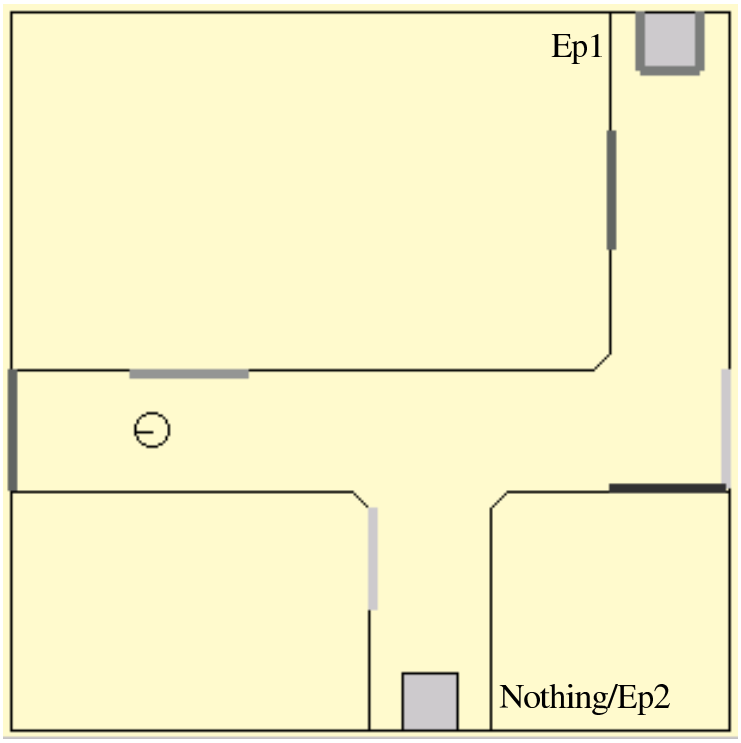

Figure 4:

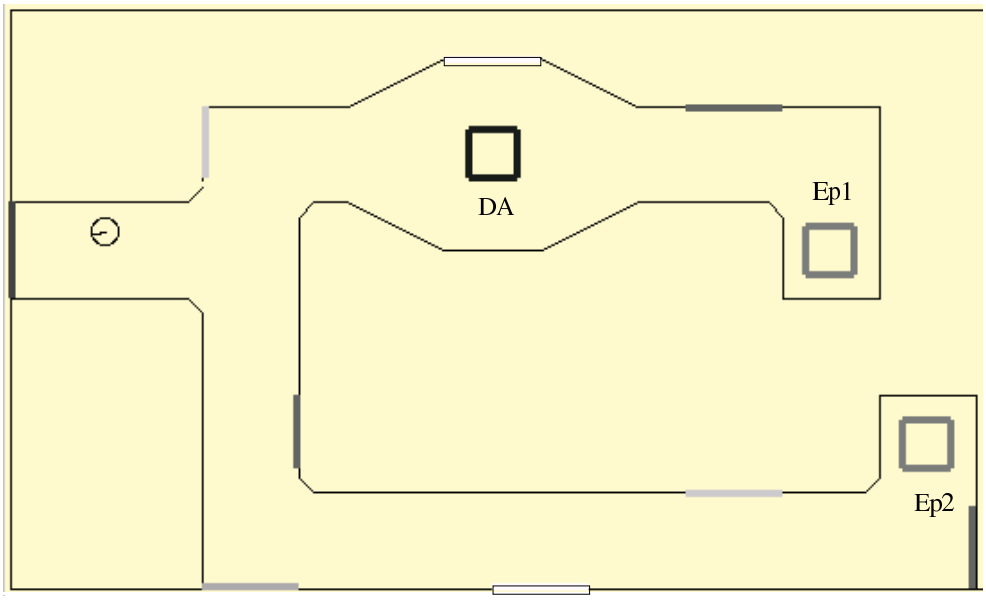

Figure 5:

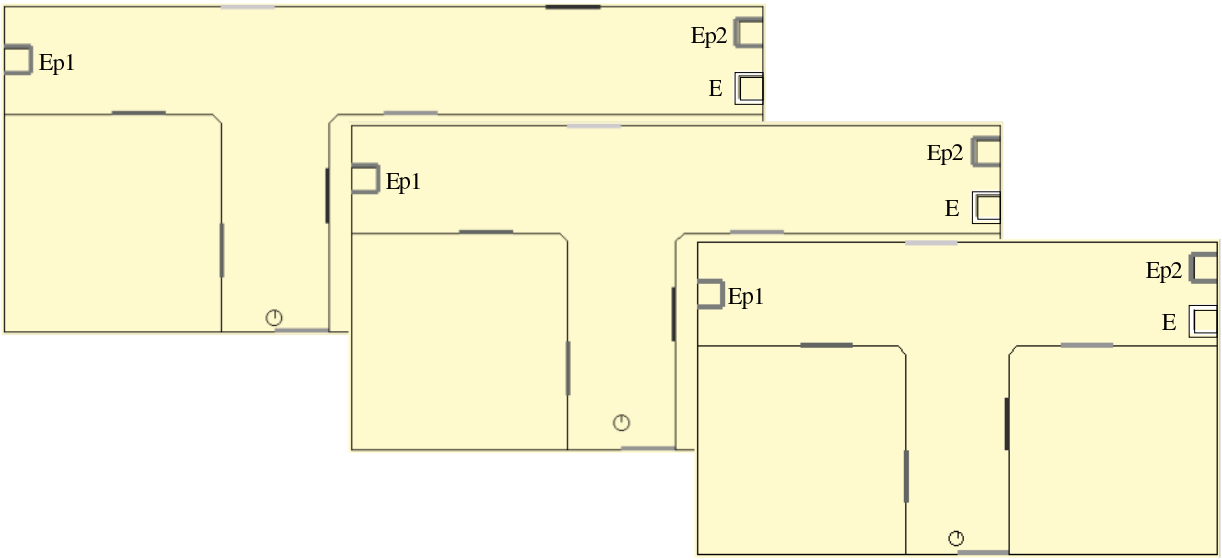

Figure 6: