

An Integrated Neuromimetic Model of the Saccadic Eye Movements for the Psikharpax Robot

Steve N’Guyen^{1,2}, Patrick Pirim², Jean-Arcady Meyer¹, and Benoît Girard¹

¹ Institut des Systèmes Intelligents et de Robotique
Université Pierre et Marie Curie-Paris 6, CNRS UMR 7222
4, place Jussieu

75252 Paris Cedex 05 - France

² Brain Vision Systems, 75013 Paris - France
`steve.nguyen@isir.upmc.fr`

Abstract. We propose an integrated model of the saccadic circuitry involved in target selection and motor command. It includes the Superior Colliculus and the Basal Ganglia in both cortical and subcortical loops. This model has spatial and feature-based learning capabilities which are demonstrated on various saccade tasks on a robotic platform. Results show that it is possible to learn to select saccades based on spatial information, feature-based information and combinations of both, without the necessity to explicitly pre-define eye-movement strategies.

1 Introduction

For living organisms, the ability to filter out the complex noisy sensory environment in order to focus attention on relevant events only is crucial. As this work contributes to the Psikharpax project [1] – which aims at designing a bio-inspired rat-like robot – we designed a robotic neuromimetic system capable of triggering gaze orientation movement toward salient stimuli. The combination of both Superior Colliculus (SC) and Basal Ganglia (BG) is known to be fundamental for this capability [2]. Dominey & Arbib [3–5] designed a now classical system level model of the saccadic circuitry, which provided a global explanation of the role and interactions of the implied brain regions. Nevertheless, some of their design choices are now outdated given the accumulated neurobiological data concerning this circuit. The present work thus proposes a new model of the saccadic circuitry including the integration of more up-to-date SC and BG models. This is the first neuromimetic model of saccadic circuitry with target selection learning capabilities implemented on a robotic platform. This model will be tested on various tasks, demonstrating its capability to behave correctly and to learn to associate spatial or feature-based cues to a reward.

Neurobiology of the circuit. The brain saccadic circuitry involves a number of cortical and subcortical areas, organized in two main pathways (Fig. 1, left). In




Fig. 1. (Left) Saccadic circuitry in the macaque monkey, subcortical pathway in dotted, cortical pathway in hatched. BG: basal ganglia; CBLM: cerebellum; FEF: frontal eye fields; LIP: lateral intraparietal cortex; SBG: saccade burst generator; SC: superior colliculus; TH: thalamus. (Right) Structure of the model, Sci : intermediate layers of the SC; SCs: superficial layer of the SC; lower dashed box: [6] SC model, upper-right dashed box: [7] CBTC model

the subcortical one, retinal input projects directly to the superficial layers of the superior colliculus (SC), whose deep layers then projects to the saccade burst generator (SBG), which drives the extraocular muscles. Two derivations are added to this basic circuit: a superior colliculus-thalamus-basal ganglia-superior colliculus (STBS) loop [8], which is probably involved in target selection, and the superior colliculus-cerebellum-saccade burst generator circuit, probably involved in the calibration of the system. The cortical pathway goes from the retina through the cortical visual areas, to the lateral intraparietal cortex (LIP) and the frontal eye fields (FEF); LIP and FEF then project to the subcortical pathway through the SC (minor FEF-SBG projections also exist). Finally, a cortico-baso-thalamo-cortical loop (CBTC) affects selection processes in the cortical pathway.

The visual cortex, the LIP, the FEF and the SC are organized in layers of retinotopic maps representing the visual field. In the rat, the SC maps topology seem to be linear [9], i.e. position of targets are encoded in a retinotopic space with a linear mapping (this mapping is logarithmic in primate). As the SBG is divided in four circuits – respectively responsible for the upward, downward, leftward and rightward movements – the spatial encoding of the maps has to be decomposed in burst signals, the so-called spatio-temporal transformation.

In the present work we propose a model of the saccadic circuitry including the SC and the BG interacting through different loops. We then test this model on a robotic platform in various saccade learning tasks involving spatial and feature-based information. Finally, the proposed model is discussed with regard to previous system level models of the saccadic circuitry.




Fig. 2. Feature to location transformation. Assuming that each feature channel (e.g. each color) is modulating the whole activity of a feature map (e.g. color map), each perceived object is represented by a gaussian activity which location is retinotopically preserved. The activity amplitude represents the feature amplitude (e.g. quantity of a given color in the object).

2 Material & Methods

2.1 Model

Our model (Fig. 1, right) selects the target of the upcoming saccade based on its location in the visual field and on its features; it can learn which locations and features to favor, using temporal-difference learning (TD). It is then able to generate the saccadic orders driving the motors.

The previously evoked problem of transformation from the topological encoding of the selected target position in the colliculus output (Motor map on Fig. 1, right) into a two dimensional output command is solved using the superior colliculus model proposed in [6], simplified to use linear maps of the visual field. This model is fed with a map merging both location and feature information (Fusion map) – on which the most salient target is selected using a model of the subcortical basal ganglia circuit – and is based on the basal ganglia model proposed in [7], in which the cortical components have been removed and the position of the thalamus changed. According to neuroanatomy of the STBS loop, the target location information is provided by the visual apparatus of the robot (see Sect. 2.2) and corresponds to the direct projections of the retina onto the SC superficial layers. The visual feature detection is fed into the cortical part of our model, where the selection of the relevant features is operated by a CBTC loop using the [7] model³.

Concerning visual features, the model’s inputs are a family of maps, each of them encodes in retinotopic coordinates the activity of neurons responding to a specific feature only (Fig. 2, left). The selection of the visual features to favor in the CBTC is based on feature channels receiving the sum of the activity of the feature maps (Fig. 2, middle). On the output, the interaction with both FC

³ Available in Python on ModelDB <http://senselab.med.yale.edu/modeldb/>

and BG creates new modulated channels (FC tends to amplify channels and BG tends to disinhibit only the maximum channel). The resulting channel amplitude is the FC activity minus GPi activity, a strong channel will be amplified by FC and fully disinhibited by BG, where a weak channel will be slightly amplified by FC but inhibited by BG. Here, the learning capability allows the biasing of BG disinhibition toward the rewarded features. Finally, normalized feature maps are multiplied by these channel values in order to obtain modulated feature maps. The global Fusion map which represents saliency (Fig. 2, bottom right) is obtained by summing all the modulated feature maps and feeds the previously described STBS loop.

The parameters of both BG models were adapted as to operate correctly with the number of channels required by our implementation (720 in STBS and 16 in CBTC, see below).

The strength of the projections of the inputs of the basal ganglia circuits (in gray on Fig. 1, right) is learnt using the basic implementation of the neuro-mimetic actor-critic algorithm proposed in [10]. This allows the system to learn that a location (resp. a feature) is more rewarded than another one, and thus to bias the selection process when multiple targets are perceived.

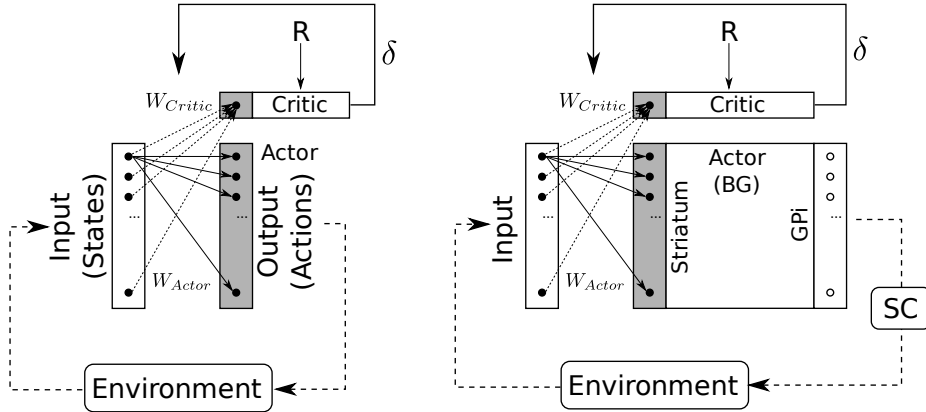


Fig. 3. (Left) Classical Actor-Critic. (Right) Implementation of the Actor-Critic learning in the Basal Ganglia model. Eligibility traces, acting as memory, are not represented

In both STBS and CBTC loops, reinforcement learning is computed as follows (cf. Fig. 3 (right)):

We first compute the TD-error δ

$$\delta = R_t + (\gamma \times V_t) - V_{t-1} \quad \text{with} \quad V_t = W_{Critic} \cdot Input_t$$

V_t being the estimated value function at time t , W_{Critic} the Critic input weights, $Input_t$ the input matrix and γ the discount factor, $\gamma = 0.7$ for all. We then




Fig. 4. (Left) Picture of the Psikharpax robotic platform.(Right) Schematic of the experimental protocol. Example for a fixation cue followed by a black screen and then 2 targets.

update the Critic weights using eligibility traces:

$$W_{Critic} \leftarrow W_{Critic} + \eta \times \delta \times E_{Critic} \quad \text{with} \quad E_{Critic} \leftarrow \alpha \times E_{Critic} + Input_{t-1}$$

η is the learning rate and is fixed to 0.0005 and α is the eligibility trace decay factor fixed to 0.9 (these parameter values are identical in both BG circuits). The Critic’s weights vector size is N , the size of the *Input* vector, so connexions are “all-to-one”. We then compute the Action vector, which is the weighted input:

$$A_t = W_{Actor} \cdot Input_t$$

Actor’s weights are updated following:

$$W_{Actor} \leftarrow W_{Actor} + \eta \times \delta \times E_{Actor} \quad \text{with} \quad E_{Actor} \leftarrow \alpha \times E_{Actor} + Input_{t-1} \otimes A'_{t-1}$$

$$\text{and} \quad A'_{t-1} = 1 - GPI_{t-1}$$

Actor’s weights size is $N \times N$, connexions are “all-to-all”. So here, compared to “traditional” reinforcement learning (cf. Fig. 3 (left)), we can see that the “States” are inputs to be selected, “Actions” are weighted inputs according to the reward bias. The Basal Ganglia performs the selection on this weighted input and then the GPI project back to the SC, eventually triggering an action.

2.2 Experimental settings

Experiments were conducted using our robotic platform Psikharpax (cf. Fig. 4, left). This platform is equipped with 2 miniatures VGA cameras (Field of view of approximately 60° horizontally and 45° vertically) each mounted on a pan-tilt mechanism driven by servomotors. Visual processing was real time computed using 2 BVS BIPS⁴ processors. This system provides objects’ position and various features such as movement (speed and direction), color, saturation, luminance and edges (with edges orientations and corners) but for this experiment we only used position and color.

⁴ <http://www.bvs-tech.com>

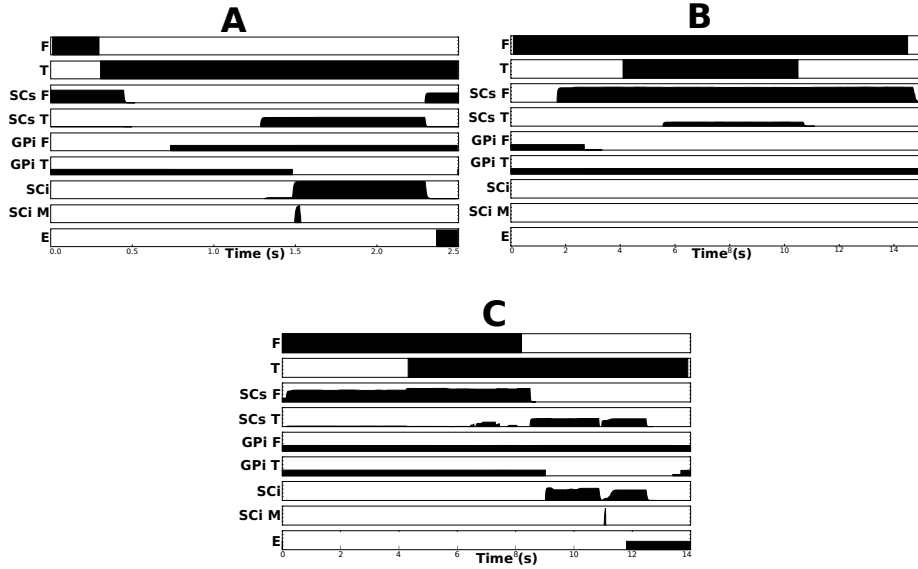


Fig. 5. (A) Saccade task. (B) Fixation task. (C) Overlap task. System activity in F: fixation cue. T: target cue. SCs F: SCs unit at the center of the field of view. SCs T: SCs unit at the target position. GPi F: GPi unit corresponding to SCs F. GPi T: GPi unit corresponding to SCs T. SCi: SCi unit corresponding to SCs T. SCi M: SCi unit on the motor layer corresponding to SCi. E: eye position.

The visual input related to target positions is a retinotopic map composed of 120×23 IPDS neurons [7] with a linear geometry representing a sensory field of 240° by 45° (notice that the total sensory field is wider than the field of view for future usage). All the SC maps in the model have the same size, geometry and neurons model. Each unit of the Striatum part of the BG (in the STBS loop) receives input from 4 SC units. And reciprocally the GPi projects back to 4 SC units, trying to mimic the “funnel” property of SC-BG connectivity [11]. Hence the BG in the STBS loop is composed of 60×12 (720) channels. In the non-spatial loop (CBTC), colors are decomposed in 16 values, thus the corresponding BG entails 16 channels.

The behavior of the system was tested in a number of tasks by putting the robot’s head in front of a 100cm diameter wide LCD screen at a distance of 65cm (Fig. 4, right) on which we simply displayed colored fixation and targets cues over a black background according to experimental protocols classically used with animals [12].

3 Results

3.1 Basic behavior

We first tested the operation of the system by reproducing basic behaviors described in [12], namely the “saccade task”, the “fixation task” and the “overlap task”. These tasks were learned using a conditioning paradigm according to which the system received a positive reward when it succeeded and a negative one when it failed. One can notice that, as our model doesn’t contain any working memory, we cannot reproduce the “delayed task”. Results are summarized in figure 5.

- Saccade task: The basic saccade behavior is well reproduced as, when the target appears, the corresponding GPi unit disinhibits the SCi Motor area and then triggers a saccade.
- Fixation task: The system succeeds in learning to continuously fixate the fixation cue while ignoring the distractor. The activity of the GPi at the location of the distractor shows no disinhibition, preventing any activity on the motor layer of the SCi and thus no saccade.
- Overlap task: Here again the system successfully learned to perform this task. The GPi unit corresponding to the target starts to disinhibit only when the fixation cue disappears. The resulting saccade is not triggered when the target appears (target+fixation) but when the fixation disappears (target only).

3.2 Elaborated behavior

As the system behaves correctly on simple minimal tasks, we now proceed to some more complex experiments.

We previously described our model as composed of 2 separated loops. The STBS loop dealing with spatial knowledge and the CBTC loop dealing with non-spatial knowledge. As both loops are provided with independent learning capabilities, it should be possible to learn to associate a specific place or a specific color to a reward.

Spatial reward: The corresponding procedure calls upon a black screen step (2.5s) followed by a fixation step (4s) during which a central cue (red square) appears until the robot fixates it. Then the fixation cue disappears and 2 targets appear (two colored disks, one actual target and one distractor) for 6s maximum until a saccade is done. Then the procedure restarts with the black screen and so on. A positive reward is given if a saccade to the desired location is done (always left or always right) and a negative one if a saccade to the wrong location is done. Averaging saccades⁵ are not rewarded. Color of targets are randomly alternated

⁵ both target are selected simultaneously and provoke a saccade to the center of gravity. This behavior is observed in animals.

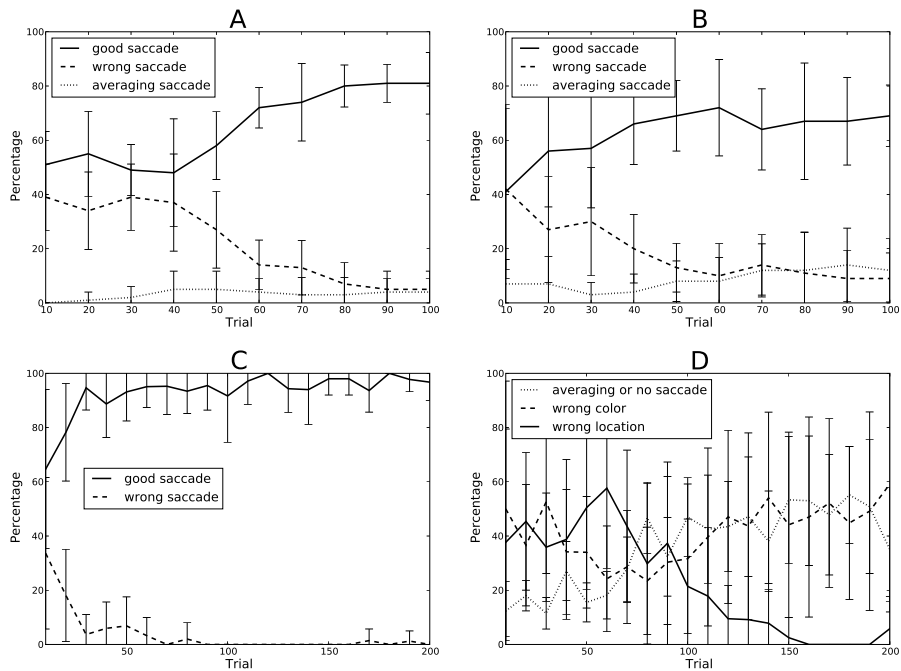


Fig. 6. A: Results of the “spatial reward” task. B: results of the “color reward” task. C: Performance of the “spatial/color” task. D: details of types of errors in task “spatial/color” for negative cases

(blue or green) between trials. Fig. 6A shows the mean results obtained from 10 runs of 100 trials each. We can see that the learning starts to correctly bias the behavior from around the 50th trial, and stabilizes with a mean performance of 80%. We can also notice that the learned behavior doesn’t fully block other behaviors, as a low level of both wrong saccades and averaging saccades still remains which is largely due to perception variability.

Non-spatial reward: The procedure for this task is the same than the preceding one, but we now reward a color based saccade (always green or always blue with randomly alternating position between trials). Here again, the system successfully learned the desired behavior quite similarly to the “spatial reward” task (cf. Fig. 6B). However, the mean performance here is slightly lower and variability higher. This can be explained by the competition between spatial and non-spatial learning. Observed behavior shows that spatial learning seems to have a stronger impact on behavior, so that having the correct colored target consecutively located at the same place will trigger spatial learning which will degrade further performance. Indeed, as the fusion map influenced by the CTBC

loop projects to the SBTS loop, the spatial learning has got the “final word” and thus is more directly involved in behavior.

Combination of spatial and non-spatial: Here we combine the 2 preceding tasks by rewarding only a specific color at a specific location. This experiment seems more difficult as, in this case, spatial and non spatial learning are competing, but performance of good saccade in positive case – i.e. when a good saccade is possible – rises quickly to near 95% and, after 100 trials, wrong saccades (saccade to both wrong location and wrong color) almost never appear again (cf. Fig. 6C). So, here, the difficulty is more to learn not to saccade when a negative case appears – i.e. when no good saccade is possible – as the percentage of correct behavior in this case does not exceed 50%. Fig. 6D shows the details of the type of error occurring in negative cases. After 100 trials, saccades to the wrong location fall to near zero. Thus, at the end of the learning, most of the errors are saccades to the right place but when the wrong color appears. This result is in accordance with the “color task” where we observed a dominance of the spatial loop (SBTS).

4 Discussion

We proposed an integrated model of Superior Colliculus and Basal Ganglia based on recent models [6, 7]. To the best of our knowledge, this is the first system level neuromimetic model of the saccadic circuitry implemented in a robot, that includes subcortical selection of spatial information (in a STBS loop) and cortical selection of features (in a CBTC loop) with learning capabilities.

Moreover we demonstrated basic functioning capabilities on classical target selection tasks. Results have shown to be efficient although the proposed system is purely reactive, i.e. doesn’t involve any working memory.

We also observed that in our model the spatial learning has a slightly stronger impact on behavior than the non-spatial one, because it is more directly involved in motor commands. Indeed, we can see in Figure 1 (right) that the spatial learning occurs at a lower level, i.e. nearer the output motor map, and thus has the “final word” on selection. Even if this property should be adjusted by appropriate weighting, the question whether animals can learn more easily a spatial reward than a non-spatial one has to be addressed, as this is what the current version of our model predicts.

The basic learning algorithm used here proved to be sufficient when the system had to learn simple sensorimotor associations with easily discriminable stimuli and without long behavioral sequences leading to the reward. Naturally, to deal with more complex situations, the learning algorithm must be improved [13].

Moreover, following computational considerations, we decided to separate the STBS and the CBTC learning capabilities. This solution has proved to be efficient, even if our model clearly lacks a cortical spatial selection loop, involving

the FEF, which would allow cognitive control on top of our reactive model and would probably affect the spatial predominance.

Contrary to previous work, our model does not explicitly contain any exploration behavior module. Indeed at the beginning of the learning task no bias already exists between targets and thus a systematic averaging saccade should be done, but inherent noise in “real world” sometime provokes varying delay in target detection or variations in color perception value. Explicit exploration, which can be view as noise in selection is here replaced by implicit exploration, i.e. natural noise in environment. So here, the robotic implementation naturally solves the exploration problem allowing for a more parsimonious model. This solution is not said to be the best one, but has proven to be sufficient for the considered tasks.

4.1 Related work

The seminal model of Dominey & Arbib [3–5] has memory and sequence learning capabilities that we have not replicated yet, and is thus more complete with regards to these points. Nevertheless, it is outdated on at least two aspects. First, while they integrate BG models in cortical loops only, we took into account the now clearly identified STBS loop, which can operate faster than the cortical circuitry. Second, their basal ganglia model is oversimplified: it is based on the direct/indirect interpretation of the BG connectivity [14], from which it keeps the direct pathway only. Consequently, simultaneously presented targets do not interact in the BG circuitry, preventing the resolution of conflicts. Their SC motor layer thus requires an ad hoc winner-takes-all mechanism. Our BG model, using most of the known BG connectivity, solves such conflicts.

More recently, Brown et al. [15] proposed a model including a CBTC loop dedicated to saccade strategy selection, and a STBS loop dedicated to target selection. Here again they include working memory mechanisms that we have not yet included. Their strategies specify whether the saccade will be based on the fixation point, the target position or target features. We showed here that such an explicit strategy selection is not necessary for the basic tasks they tested and that we implemented, our cortical feature selection circuit modulating the subcortical location selection circuit is sufficient. Moreover, their STBS does not stand on its own as it is strictly driven by cortical inputs, while our decorticated system is still able to learn and generate location- or fixation-based saccades. The details of their BG circuitry also suffer from limitations, discussed in details in [16].

Finally, Chambers et al. [11] proposed a model integrating both the subcortical and cortical pathways, where a single BG model dedicated to location-based selection integrates FEF and SC inputs. Using the various positive feedback loops of this circuitry, they show that manipulating the level of dopamine in their BG model generates reaction time and saccade size modifications reminiscent of Parkinson’s disease patient behavior. Unfortunately, they rely on a single thalamus module, receiving both SC and FEF inputs, while it is clearly

established that CBTC loops run through the ventral anterior and ventral lateral nuclei and the STBS loops through the pulvinar and intralaminar nuclei.

4.2 Future work

It has been explained that, unlike previous works, the proposed model doesn't contain any working memory and thus cannot reproduce memory related tasks. Working memory related activity has been demonstrated in quasi-visual cells in the FEF, the LIP, the SC and the BG. Many models of these cells have been proposed (e.g. [17, 18]) Therefore, we will add such a capability in future work.

Moreover, the SC is known to integrate vision, audition and tactile information [19]. In this work we only used visual information in order to test our model, but as our robotic platform is also equipped with auditory [20] and tactile [21] capabilities, we intend to extend it to these modalities.

Acknowledgements

This work has been funded by the EC Integrated Project ICEA (Integrating Cognition, Emotion and Autonomy), IST-027819-IP.

References

1. Meyer, J.A., Guillot, A., Girard, B., Khamassi, M., Pirim, P., Berthoz, A.: The Psikharpax project: Towards building an artificial rat. *Robotics and autonomous systems* **50**(4) (2005) 211–223
2. Moschovakis, A., Scudder, C., Highstein, S.: The microscopic anatomy and physiology of the mammalian saccadic system. *Prog Neurobiol* **50** (1996) 133–254
3. Dominey, P.F., Arbib, M.A.: A cortico-subcortical model for generation of spatially accurate sequential saccades. *Cereb Cortex* **2** (1992) 153–175
4. Dominey, P., Arbib, M., Joseph, J.P.: A model of corticostriatal plasticity for learning oculomotor associations and sequences. *J Cogn Neurosci* **7** (1995) 311–336
5. Dominey, P.F.: Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning. *Biol Cybern* **73** (1995) 265–274
6. Tabareau, N., Bennequin, D., Berthoz, A., Slotine, J.J., Girard, B.: Geometry of the superior colliculus mapping and efficient oculomotor computation. *Biol Cybern* **97**(4) (2007) 279–292
7. Girard, B., Tabareau, N., Pham, Q., Berthoz, A., Slotine, J.J.: Where neuroscience and dynamic system theory meet autonomous robotics: a contracting basal ganglia model for action selection. *Neural Netw* **21**(4) (2008) 628–641
8. McHaffie, J., Stanford, T., Stein, B., Coizet, V., Redgrave, P.: Subcortical loops through the basal ganglia. *Trends Neurosci* **28**(8) (Aug 2005) 401–407
9. Siminoff, R., Schwassmann, H., Kruger, L.: An electrophysiological study of the visual projection to the superior colliculus of the rat. *J Comp Neurol* **127** (1966) 435–444
10. Montague, P.R., Dayan, P., Sejnowski, T.J.: A framework for mesencephalic dopamine systems based on predictive hebbian learning. *J Neurosci* **16**(5) (1996) 1936–1947

11. Chambers, J., Gurney, K., Humphries, M., Prescott, A.: Mechanisms of choice in the primate brain: a quick look at positive feedback. In Bryson, J., Prescott, T., Seth, A., eds.: *Modelling Natural Action Selection: Proceedings of an International Workshop*, Brighton, UK, AISB Press (2005) 45–52
12. Hikosaka, O., Sakamoto, M., Usui, S.: Functional properties of monkey caudate neurons. I. Activities related to saccadic eye movements. *J Neurophysiol* **61**(4) (1989) 780–798
13. Khamassi, M., Lachèze, L., Girard, B., Berthoz, A., Guillot, A.: Actor-critic models of reinforcement learning in the basal ganglia: From natural to artificial rats. *Adaptive Behavior* **13**(2) (2005) 131–148
14. Albin, R.L., Young, A.B., Penney, J.B.: The functional anatomy of disorders of the basal ganglia. *Trends Neurosci* **18**(2) (1995) 63–64
15. Brown, J., Bullock, D., Grossberg, S.: How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Netw* **17**(4) (2004) 471–510
16. Girard, B., Berthoz, A.: From brainstem to cortex: Computational models of saccade generation circuitry. *Prog Neurobiol* **77** (2005) 215–251
17. Droulez, J., Berthoz, A.: A neural network model of sensoritopic maps with predictive short-term memory properties. *Proc Natl Acad Sci* **88** (1991) 9653–9657
18. Mitchell, J., Zipser, D.: Sequential memory-guided saccades and target selection: a neural model of the frontal eye fields. *Vision Res* **43** (2003) 2669–2695
19. Meredith, M.A., Stein, B.E.: Visual, auditory, and somatosensory convergence on cells in superior colliculus results in multisensory integration. *J Neurophysiol* **56**(3) (Sep 1986) 640–662
20. Bernard, M., N’Guyen, S., Pirim, P., Gas, B., Meyer, J.A.: Phonotaxis behavior in the artificial rat psikharpax. In: *International Symposium on Robotics and Intelligent Sensors, IRIS2010*, Nagoya, Japan (2010)
21. N’Guyen, S., Pirim, P., Meyer, J.A.: Tactile texture discrimination in the robot-rat psikharpax. In: *BIO SIGNALS 2010, Third International Conference on Bio-Inspired Systems and Signal Processing*, Valencia, Spain (2010)