

Generalization of Integrator Models to Foraging: A Robot Study Using the DAC9 Model

Encarni Marcos, Armin Duff, Martí Sánchez-Fibla, and Paul F.M.J. Verschure

SPECS, Department of Technology, Universitat Pompeu Fabra,
Roc Boronat 138, 08018, Barcelona, Spain
{encarnacion.marcos, armin.duff, marti.sanchez, paul.verschure}@upf.edu
<http://specs.upf.edu>

Abstract. Experimental research on decision making has been mainly focused on binary perceptual tasks. The generally accepted models describing the decision process in these tasks are the integrator models. These models suggest that perceptual evidence is accumulated over time until a decision is made. Therefore, the final decision is based solely on recent perceptual information. In behaviorally more relevant tasks such as foraging, it is however probable, that the current choice also depends on previous experience. To understand the implications of considering previous experience in an integrator model we investigate it using a cognitive architecture (DAC9) with a robot performing foraging tasks. Compared to an instantaneous decision making model we show that an integrator model improves performance and robustness to task complexity. Further we show that it compresses the information stored in memory. This result suggests a change in the way actions are retrieved from memory leading to self-generated actions.

Keywords: sequence learning, decision making, integrator models.

1 Introduction

Binary perceptual tasks have been widely used to study the neural mechanisms underlying decision making [1,2]. This kind of task involve a simple decision about a feature of a stimulus that is expressed as a choice between two alternative options. Many models have been proposed to explain this decision making process predicting the relationship between reaction time and accuracy [3,4]. Most of them explain decision making as an accumulation process that takes place over time until a decision bound is reached. These models are known as integrator models and have been generally accepted as an explanation for decision making in perceptual tasks where learning is not required to successfully perform the task. Here, we investigated the interaction between an integrator model and memory in foraging tasks using a well establish cognitive architecture as a framework [5].

One largely used perceptual experimental paradigm is defined by a random-dot motion (RDM) task where humans or monkeys have to select between two

possible stimulus categories, such as leftward or rightward motion [6]. Integrator models, such as race models [3] and drift-diffusion models [4], provide a straightforward account of the speed-accuracy trade-off. These models suggest that evidence is accumulated over time until this accumulation reaches a bound, i. e. criterion level, and a decision is made. As the RDM task, many of the experimental paradigms used to study the decision making process are simple perceptual tasks where the correct performance of a trial depends exclusively on the current perceptual information, e. g. color. The proposed integrator models assume that the alternative options are known a priori and therefore learning during the task is not required. However, this would not be the case in more realistic foraging tasks where the information about different targets have to be acquired from the environment and many alternative choices might be available at each decision point. Therefore, a two-fold problem has to be solved during foraging: the appropriate learning of the environment and actions and the appropriate retrieval of information to achieve goal states (targets), i. e. sequences of perception and action need to be learned and retrieved to reach goal positions in an environment.

To study the interaction between decision making and memory, we worked in the framework of the Distributed Adaptive Control (DAC) architecture [7,5]. The decision making in DAC follows the Bayesian principle [8,9]. We extended the architecture with an integrator decision making model (DAC9; see [7,10,11] for details about previous versions of DAC), based on the race model, to investigate its implications during foraging tasks and we further compared it with the decision making in DAC (instantaneous model). We show that the integrator model resulted in a new mechanism of storing and recalling information from memory suggesting that the actions are not stored in memory but self-generated during retrieval of information. In a previous study [12], we assessed the impact of these two decision making models in the learning of event order and interval in a sequence in two foraging tasks. In the current study, we go one step further quantifying (1) the scalability of the two models with task complexity in five different foraging tasks and (2) the implications on the information stored in memory and proposing (3) a new working memory mechanism that accounts for a continuous action space.

2 Materials and Methods

2.1 Cognitive Architecture

The DAC architecture has already proven its suitability to study the problems encountered in biology helping to investigate perception, cognition and behavior in foraging situations in which the access to real neuronal and behavioural data is difficult [5]. DAC is based on the assumption that learning consists of the interaction of three layers of control: reactive, adaptive and contextual, as illustrated in Fig. 1. The reactive layer provides pre-wired responses that allows for a simple interaction with the environment and accomplish simple automatic behaviours. The adaptive layer provides mechanism for the classification of the

sensory events (internal representations) and the shaping of responses in simple tasks as in classical conditioning [13]. The internal representations (prototypes; see [14] for details) generated by the adaptive layer are stored in the contextual layer as couplets of sensory-motor states and used to plan future behaviour, as in operant conditioning [15].




Fig. 1. Schematic representation of the DAC architecture. It is based on the assumption that behavior results from three tightly coupled layers of control: reactive, adaptive and contextual. Abbreviations mean: a_r , reactive layer action; a_a , adaptive layer action, a_c , contextual layer action.

In this study, we mainly focused on the contextual layer of DAC that provides mechanisms for memorizing and recalling information. It consists of two memory structures: the short-term memory (STM) and the long-term memory (LTM), for permanent storage of information. During learning, pairs of prototype-action are stored in the STM as the robot interacts with the environment. When a goal state is reached, i. e. reward or punishment, the content of the STM is copied into the LTM and the STM is reset. The LTM has sequences of pairs of prototype-action that lead the robot to goal states. The prototype-action pairs that form a sequence are called segments. During the recall process, the prototypes stored in LTM are matched against the generated prototypes from ongoing sensory events. The degree of matching of segment l in sequence q determines the input to its, so called, *collector* unit, c_{lq} :

$$c_{lq} = (1 - d(e, e_{lq}))t_{lq} \quad (1)$$

where $d(e, e_{lq})$ is calculated as the Euclidean distance between stored *prototype* e_{lq} and current *prototype* e and t_{lq} is called *trigger*. The trigger value biases the sensory matching process of the segments and allows chaining through a sequence, i. e. its default value is 1 and it is set to a higher value if the previous segment $l - 1$ is activated.

The activity of the collectors contribute to the action proposed by the contextual layer. We only consider the collectors' activity that satisfy both conditions:

(1) its activity is above a certain threshold (θ^C), (2) its activity is inside a predefined percentage range from the maximum collector's activity, i.e. the collectors compete in an E%-Max Winner Take All (WTA) mechanism [16]. The actual action proposed by the contextual layer (a_c) is calculated as:

$$a_c = \sum_{l,q \in LTM} \pm \frac{c_{lq} H(c_{lq} - \theta^C)}{\delta_{lq}} a_{lq} \quad (2)$$

where $H(\cdot)$ is a step function that is 0 for values lower than θ^C and is 1 for values higher than θ^C , δ_{lq} is the distance measured in segments between the selected segment l and the last segment in the sequence, i.e. the distance to the goal state and a_{lq} is the action stored in segment l of sequence q . By doing this division the segments closer to the goal state have more impact on the contextual action. The sign is positive if the segment belongs to an appetitive sequence and negative if it belongs to an aversive sequence.

The actions triggered by each of the three different layers are filtered by priority, giving more priority to reactive actions (a_r), then to contextual actions (a_c) and finally to adaptive actions (a_a). The one that takes the control of the motor is stored in STM and afterwards in LTM.

2.2 Integrator Models

Many integrator models have been proposed, but mainly, in all of them, the change in the accumulation of evidence in favour of one alternative ($x_i(t)$) can be described as:

$$\frac{dx_i}{dt} = \mu E_i(t) + \xi \quad (3)$$

where μ is the growth rate of the accumulation, $E_i(t)$ is the internal estimate of evidence at time t and ξ is a Gaussian noise with mean of zero and variance of σ^2 . The proposed models consider the variables $x_i(t)$, $E_i(t)$, μ and ξ in a different manner. We implement a rise-to-threshold model based on the race model. The race model [3] suggests that there are separate variables $x_i(t)$ for each option that accumulate evidence independently until one of them reaches a decision bound and a decision is made.

Our implementation of the race model consisted of a number of independent variables that compete to take the control of the robot. Each variable accumulated evidence in favour of one action, such as right or left. When the value of a variable grew above a criterion level, i. e. decision bound, the action associated with it was performed by the robot. The change in the activity of the variables within a time step dt was defined as:

$$da_i(t) = \begin{cases} dt(\mu_r a_{r_i} + \mu_a a_{a_i} + \mu_c a_{c_i} + \xi) & , \text{ if } t - t_{la} > T_{ref} \\ 0 & , \text{ if } t - t_{la} \leq T_{ref} \end{cases} \quad (4)$$

where a_r, a_a, a_c are the actions triggered by the reactive, adaptive and contextual layer respectively, $i \in \mathbb{N}^N$ and it is the subindex of the N different possible actions, μ_r, μ_a, μ_c are the mean growth rates of the variables units, ξ is a Gaussian noise term with a mean of zero and a variance of σ^2 , t_{la} is the time at which the last action was executed and T_{ref} is the refractory period. In our experiments $dt = 1ms$ and $\xi = 0$. When the value of a_i reaches a predefined threshold the associated action is executed. In biology, the refractory period is the amount of time a excitable membrane needs to be ready for a second stimulus once it returns to the resting state. Consistent with this, the T_{ref} term referred to the amount of time necessary to start again the competition between actions after one of them was executed.

2.3 Foraging Tasks

The mobile agent was simulated in C++ and wSim [17] using the 3D Open Graphics Library approximating a Kephra robot ¹. Different previous studies have proven the validity of this simulated robot with respect to a real one [17,18]. The robot has a radio of 5.5 cm and 8 proximity sensors and 8 light sensors. The values captured by both light and proximity sensors decay exponentially. The proximity sensors measure the distance to obstacles while the light sensors measure the intensity of light sources. The robot is equipped with a color camera with a visual angle of 45 deg. of amplitude. The image from the camera is color separated such that there are three channels: red, green and blue, each of them with a resolution of 36x36 pixels. Except otherwise specified the camera is always pointing to the floor with a tilt angle of -60 deg. with respect to the horizontal axis. The robot translates with a speed of $0,1 \times robotradius$ and it rotates with a speed of 10 deg.

To study the interaction between an integrator model and memory we defined a number of foraging tasks where not only perceptual but also memory information was essential to achieve a performance about chance. The tasks had different rated complexity to assess how the decision making models scaled to it (Figure 2). In all the environments the goal of the task was to go to the light source, i. e. reward. Every trial started from one of the positions shown in Figure 2, randomly selected. The trial ended when the robot hit the light or collided with the wall. A successful trial ended when the light was hit. The environments contained colored patches that served as cues. The light was detected by the light sensors of the robot. However, the light was not strong enough to trigger a reactive action from the side patches. The adaptive layer used reactive layer sub-threshold activity to generate the prototypes and to learn the associations between prototype-action. Once the prototypes were stable the contextual layer started storing sequences of prototype-action that led to a goal state. The goal state occurs when the robot reaches the light or collides with the wall. When a collision occurred it was stored as an undesirable state in memory and had a negative influence on the action proposed by the contextual layer.

¹ K-Team, Lausanne, Switzerland.

The complexity of the tasks was rated taking into account the number of patches and how ambiguous they were as follows:

$$TC = \frac{n_p}{n_c} n_a \quad (5)$$

where n_p is the number of patches, n_c is the number of different colors and n_a is the number of different turning angle amplitudes needed to be learned. This measure was then useful to compare the robot performance in each of the tasks for the two proposed models. The complexity of the task 1 is 3, the complexity of the task 2 is 4.5 and 5, 7.5 and 11.7 for the tasks 3, 4 and 5, respectively.




Fig. 2. Foraging tasks ordered by task complexity. (a) Task 1. Unambiguous restricted open arena, $TC = 3$. (b) Task 2. Ambiguous restricted open arena, $TC = 4.5$. (c) Task 3. Ambiguous restricted open arena, $TC = 5$. (d) Task 4. Ambiguous restricted open arena, $TC = 7.5$. (e) Task 5. Ambiguous restricted open arena, $TC = 11.7$.

The first task was an unambiguous restricted open arena foraging task, i.e. no context information was needed because the location of the target was uniquely predicted by the color patches (Fig. 2a). Therefore, this task could be correctly solved by the adaptive layer, but still we tested the performance at the contextual layer level. The rest of the four tasks consisted of ambiguous restricted open arenas, because, in all cases, context information was needed to reach a performance above chance. The actions associated with the patches closest to the light were not unique but depended on the previous context (Fig. 2b, 2c, 2d and 2e) and therefore the problem could only be solved at the contextual layer level. These experiments allowed us to study the detailed performance of each model and its dynamics as well as to evaluate the results in tasks where different kind of actions were required. In all the foraging tasks, when the contextual layer was enabled, the actions from the reactive and the adaptive layers were deactivated to avoid any influence they could have on the results.

To test the system for robustness we added 5% of noise to the motors, following a Gaussian random distribution and we varied the initial position of each trial according to a two dimensional normal distribution with mean 0 and variance $0,1 \times$ robot radius. Moreover, to assess the impact of the camera noise in the information stored in memory, in Task 2 (Fig. 2b and 2e), we added noise to the hue sensed by the camera from 0% to 10% in steps of 1%, following a Gaussian random distribution. For every condition, we ran 10 experiments with 1000 trials each.

To investigate what was the impact on memory of the interaction between the decision making models and memory itself we calculated the degree of compression of information in memory through the entropy of the stored information as follows:

$$E_M = - \sum_{s \in S} p(s) \log_2(p(s)) \quad (6)$$

where s is one segment in memory and $p(s)$ is the probability that the segment is selected in a current experiment. This measurement allows us to assess the amount of information needed to encode a visual stimulus in memory.

3 Results

In this study, we investigated the generalization of an integrator model in foraging tasks. We designed a number of foraging tasks with increased complexity to assess the generalization of the integrator model in more realistic tasks. The results in these tasks suggested a new mechanism to store and recall information from memory. We further tested the implications of this new mechanism in a foraging task and we show that it resulted in a more optimal way of learning and exploiting the environment.

3.1 Foraging Tasks

In all the tasks, we recorded the performance of the robot after LTM acquisition. As shown in Fig. 3, as task complexity increased the performance of the robot decreased dramatically in the case of the instantaneous model where it dropped to a mean value of 0.55 for the most complex task. It kept stable in the integrator model, maintaining a mean value of performance above 0.9.

In order to evaluate the impact of the camera noise in the information stored in memory due to the influence of each decision making models, we used Task 2 ($TC = 4.5$) because it was the simplest one that requires the use of the contextual layer. For clarity we also report here the performance of the robot with varying camera noise [12]. As previously reported in [12], the performance of the robot decreased as the camera noise increased in both models (Fig. 4a). The difference between the performance of the two models was significantly different along the different values of camera noise (ks-test, $p < 0.01$). From 0% of camera noise to 6% the instantaneous model was incrementally more affected by the noise than the integrator model. However, from 6% to 10% the noise had an important

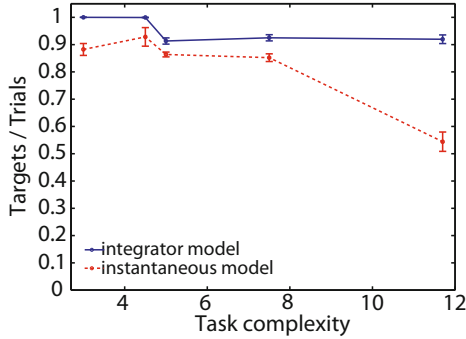


Fig. 3. Robot performance along different task complexity. Ratio targets/trials with instantaneous and integrator models as task complexity increases. Bars represent means \pm sem.

impact on the integrator model, resulting in a smaller difference in performance with respect to the instantaneous model. Once the noise of the camera started to critically affect the sequentiality of the actions the performance decayed in both models with a similar slope (Fig. 4a).

To assess the impact of both models at the memory level we calculated the entropy of the stored information, E_M (see Eq. 6). As shown in Fig. 4b, E_M with the integrator model was higher along the different camera noise compared to the instantaneous model (ks-test: $p < 0.001$). Moreover, the dynamics in both cases were opposite: E_M decreased as camera noise increased in the integrator model whereas it increased as camera noise increased in the case of the instantaneous model. Low values of E_M means that segments of memory respond to a small fraction of the stimuli resulting in a higher number of segments in memory. The opposite occurs for high values of E_M . Therefore, the integrator model compressed the memory and less number of segments were necessary to encode same stimulus. As a drawback, explicit representation of time in memory, i. e. the number of steps needed to cross a patch, is lost.

3.2 Self-generated Actions

The compression of information in memory due to the use of the race model changes the way information is stored in memory suggesting a new mechanism to recall it. Instead, of a recall of actions from memory it suggests the recall of goals. Consequently, we hypothesize that the actions are self-generated rather than stored in memory. During the recall period, visual information is retrieved from memory and actions are performed depending on the position of perceptual target with respect to the robot. When the information is selected from memory, we distinguish between two different recall methods: (1) the next prototype in the sequence is retrieved, i.e. sub-goals are progressively achieved; (2) the prototype associated with the final goal is retrieved. In both cases, the retrieved information is stored in working memory and the robot searches for it. To do




Fig. 4. Performance and entropy along different camera noise. (a) Ratio Targets / Trials distribution for instantaneous and integrator models. (b) Entropy of the memory along different camera noise. In both figures bars represent means \pm sem.

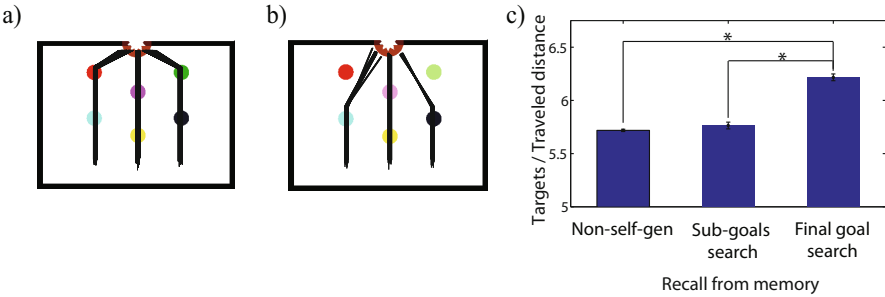


Fig. 5. Robot trajectories and performance. (a) Trajectories generated during sub-goals search. (b) Trajectories generated during final goal search. (c) Ratio Targets / Travelled distance for the integrator model with non-self-generated actions recall, sub-goals search recall and final goal search recall. Bars represent means \pm sem.

so, it moves the tilt angle of the camera from -60 deg. to -20 deg. and rotates over its own axis. In this way, the robot can see the visual cues that are far away from its current position. Once the robot sees the sub-goal or final goal it moves again the tilt angle from -20 deg. to -60 deg. and goes straight to the goal, i. e. self-generating the actions. We test this new way of retrieving information from memory in the Task 1 and compare the results to the non-self-generated actions investigated in the section. As shown in Fig. 5a and 5b, the robot follows a different path depending on the mechanism it uses to recall information from memory. We observed that the ratio targets divided into travelled distance was significantly higher in the case of the final goal search mechanism of self-generated actions (Fig. 5c; ks-test: $p < 0.001$), i. e. the robot follows a shorter path to hit the target. On the contrary, there is no significant difference between the sub-goal search and the non-self-generated actions (ks-test: $p > 0.05$). This result shows an optimal way of using the compressed information of the memory when the

actions are not stored but self-generated during the retrieval of information from memory. It results in more flexibility in the actions to be taken and allows to account for a continuous action space.

4 Conclusions

We tested the implications of an integrator decision making model in sequence learning tasks with multiple alternatives using a cognitive architecture that we called DAC9, evolving from previous implementations [7]. We compared the results with a Bayesian decision making model which is thought to be optimal for action selection. As a framework we used a robot based architecture which allowed us to understand the behavioural and architectural implications of these alternative models during foraging tasks. We showed that the race model has a more robust task-related performance when perceptual noise is added [12] and when task complexity increases compared to the Bayesian model. Moreover, the race model also implied a compression of information in memory suggesting an alternative way of storing information, i. e. only perceptual information is acquired and the actions are self-generated during recall. The self-generation of actions during the retrieval from memory shows a mechanism able to account for a continuous action space.

In a previous study [12], we reported the differences in the storage of information in memory due to both models and the impact they have in performance. Here, we quantified the difference in the information stored in memory by calculating the memory entropy. We showed that the entropy is higher in the integrator model than in the instantaneous model. In the instantaneous model the actions are continuously recalled and performed. Therefore, in this case, the robot executes a number of actions, generally greater than one, each time it crosses a visual cue. In the case of the integrator model, we proposed a new mechanism to optimally use the information from memory. We implemented a goal oriented mechanism that retrieves visual cues from memory instead of actions. Once the visual cue, i. e. goal, is selected from memory the robot searches for it in the environment. Whenever the robot sees the goal it goes towards it. This new mechanism can be seen as the storage of an abstract object in memory, i. e. a door. If a person wants to leave a room she/he has to first localize the door and then go towards it.

So far the integrator models have been used to explain simple decision making tasks, such as RDM task [6] or the countermanding task [19]. The implementation of the integrator model was based on the race model. Generally, the race model has been mainly used to explain behavior in a countermanding task [19], predicting probability of failure and reaction time. Here, we showed the implication of this decision making model in a more general framework. We observed that it has an important impact on how the memory is constructed and therefore on how the information is used later on.

The main assumptions we made in our proposal of self-generated actions during the recall from memory is that visual cues can always be seen from the

current position of the robot. However, in wide open field environments, when this is not the case, our assumption would fail. In those situations, we would rely on head direction accumulator [20] cells. The heading direction information would be stored in memory together with the visual information. During the recall from memory the actions would be also self-generated. Similar to the search of visual prototypes tested in this study, the robot would rotate around its own axis until its current head direction is equal or close enough to the retrieved head direction.

Physiological studies have shown that granular and pyramidal cells in the hippocampus encode information with high sparsity (low entropy), i. e. neurons respond to a small fraction of stimuli [21]. In contrast, cells in the PFC have shown to be selective to particular cues with less sparsity (higher entropy) than the hippocampus and also with distinct temporal profile [22]. We observed that these two mechanisms of encoding memory have some similarity with the implications shown in this study due to the two decision making models, i. e. higher entropy in the integrator model compared to the instantaneous model. One could speculate that there is a distributed control system for sequence learning involving the hippocampus and the PFC connected to an external area which accumulates evidence, as found in the superior colliculus [23], the lateral intraparietal area [24], the frontal eye fields [25] and the PFC itself [26].

Acknowledgements. This work was supported by the EU projects Synthetic Forager-SF (EC-FP7/217148), eSMC (EC-FP7/270212) and Goal-Leaders (EC-FP7/270108).

References

1. Gold, J.I., Shadlen, M.N.: The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574 (2007)
2. Smith, P.L., Ratcliff, R.: Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168 (2004)
3. Logan, G.D., Cowan, W.B.: On the ability to inhibit thought and action: A theory of an act of control. *Psychol. Rev.* 91, 295–327 (1984)
4. Ratcliff, R., Rouder, J.N.: Modelling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356 (1998)
5. Verschure, P.F.M.J., Voegtlin, T., Douglas, R.J.: Environmentally mediated synergy between perception and behavior robots. *Nature* 425, 620–624 (2003)
6. Shadlen, M.N., Newsome, W.T.: Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1936 (2001)
7. Verschure, P.F.M.J., Althaus, P.: A real-world rational agent: unifying old and new ai. *Cognitive Science* 27, 561–590 (2003)
8. Bayes, T.: An essay towards solving a problem in the doctrine of chances. *Transactions of the Royal Society* 53, 370–418 (1763)
9. Verschure, P.F.M.J.: Distributed adaptive control: A theory of the mind, brain, body nexus. In: *BICA* (in press, 2012)

10. Duff, A., Fibla, M.S., Verschure, P.F.M.J.: A biologically based model for the integration of sensory-motor contingencies in rules and plans: A prefrontal cortex based extension of the distributed adaptive control architecture. *Brain Res. Bull.* (2010)
11. Mathews, Z., i Badia, S.B., Verschure, P.F.: Pasar: An integrated model of prediction, anticipation, sensation, attention and response for artificial sensorimotor systems. *Information Sciences* 186(1), 1–19 (2012)
12. Marcos, E., Duff, A., Sánchez-Fibla, M., Verschure, P.F.M.J.: The neuronal substrate underlying order and interval representations in sequential tasks: a biologically based robot study. In: *IEEE World Congress on Computational Intelligence* (2010)
13. Pavlov, I.P.: *Conditioned reflexes: an investigation of the physiological activity of the cerebral cortex*. Oxford University Press (1927)
14. Duff, A., Verschure, P.F.M.J.: Unifying perceptual and behavioral learning with a correlative subspace learning rule. *Neurocomputing* 73(10-12), 1818–1830 (2010)
15. Thorndike, E.: *Animal intelligence*. Macmillan, New York (1911)
16. de Almeida, L., Idiart, M., Lisman, J.E.: A Second Function of Gamma Frequency Oscillations: An E%-Max Winner-Take-All Mechanism Selects Which Cells Fire. *J. Neurosci.* 29(23), 7497–7503 (2009)
17. Wyss, R.: *Sensory and motor coding in the organization of behavior*. PhD thesis, ETHZ (2003)
18. Wyss, R., König, P., Verschure, P.: A model of the ventral visual system based on temporal stability and local memory. *PLoS Biol.* 4 (2006)
19. Hanes, D.P., Schall, J.D.: Neural control of voluntary movement initiation. *Science* 274, 427–430 (1996)
20. Mathews, Z., Lechón, M., Calvo, J.M.B., Dhir, A., Duff, A., Badia, S.B.: Verschure, P.F.M.J.: Insect-like mapless navigation based on head direction cells and contextual learning using chemo-visual sensors. In: *Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009*, pp. 2243–2250. IEEE Press (2009)
21. Jung, M.W., McNaughton, B.L.: Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus* 3, 165–182 (1993)
22. Asaad, W.F., Rainer, G., Miller, E.K.: Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21, 1399–1407 (1998)
23. Ratcliff, R., Hasegawa, Y.T., Hasegawa, R.P., Smith, P.L., Segraves, M.A.: Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J. Neurophysiol.* 97, 1756–1774 (2007)
24. Roitman, J.D., Shadlen, M.N.: Response of neurons in the lateral intraparietal area during combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489 (2002)
25. Gold, J.I., Shadlen, M.N.: The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J. Neurosci.* 23, 632–651 (2003)
26. Kim, J.N., Shadlen, M.N.: Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* 2, 176–185 (1999)