

When Artificial Intelligence and Computational Neuroscience meet

Frédéric Alexandre, Peter Ford Dominey, Philippe Gaussier, Benoît Girard,
Mehdi Khamassi, Nicolas P. Rougier

► **To cite this version:**

Frédéric Alexandre, Peter Ford Dominey, Philippe Gaussier, Benoît Girard, Mehdi Khamassi, et al.. When Artificial Intelligence and Computational Neuroscience meet. Springer. A guided tour of artificial intelligence research, Interfaces and applications of artificial intelligence, 3, 2020, Interfaces and Applications of Artificial Intelligence, 978-3-030-06170-8. hal-01735123

HAL Id: hal-01735123

<https://hal.archives-ouvertes.fr/hal-01735123>

Submitted on 13 May 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

When Artificial Intelligence and Computational Neuroscience meet

Frédéric Alexandre, Peter F. Dominey, Philippe Gaussier, Benoît Girard, Mehdi Khamassi, and Nicolas P. Rougier

Abstract Computational Intelligence and Artificial Intelligence are both aiming at building machines and softwares capable of intelligent behavior. They are consequently prone to interactions, even if the latter is not necessarily interested in understanding how cognition emerges from the brain substrate. In this chapter, we enumerate, describe and discuss the most important fields of interactions. Some are methodological and are concerned with information representation, processing and learning. At the functional level, the focus is set on major cognitive functions like perception, navigation, decision making and language. Among the salient characteristics of the critical contributions of Computational Neuroscience to the development of intelligent systems, its systemic view of the cerebral functioning is particularly precious to model highly multimodal cognitive functions like decision making

Frédéric Alexandre

Inria Bordeaux Sud-Ouest; LaBRI UMR 5800; IMN UMR 5293; CNRS; Université de Bordeaux; Bordeaux INP; France. e-mail: Frederic.Alexandre@inria.fr

Peter F. Dominey

Inserm 1208 Stem Cell and Brain Research Institute, Human and Robot Cognitive Systems, 18, avenue du doyen Jean Lépine 69675 Bron cedex, e-mail: peter.dominey@inserm.fr

Philippe Gaussier

ETIS (ENSEA, University of Cergy Pontoise CNRS UMR 8051), 2, avenue Adolphe-Chauvin B.P. 222, 95302 Cergy-Pontoise Cedex, France, e-mail: gaussier@ensea.fr

Benoît Girard

Sorbonne Universités, UPMC Univ Paris 06, CNRS, Institute of Intelligent Systems and Robotics (ISIR), F-75005 Paris, France, e-mail: benoit.girard@isir.upmc.fr

Mehdi Khamassi

Sorbonne Universités, UPMC Univ Paris 06, CNRS, Institute of Intelligent Systems and Robotics (ISIR), F-75005 Paris, France, e-mail: mehdi.khamassi@upmc.fr

Nicolas P. Rougier

Inria Bordeaux Sud-Ouest; LaBRI UMR 5800; IMN UMR 5293; CNRS; Université de Bordeaux; Bordeaux INP; France. e-mail: Nicolas.Rougier@inria.fr

and language and to design cognitive architectures for the autonomous behavior of robots.

1 Introduction

The goal of Computational Neuroscience (CN) is to study relations between brain structures and functions by the means of information processing techniques [Marr, 1982; Schwartz, 1990; Churchland and Sejnowski, 1992; Dayan and Abbott, 2001]. This scientific domain has to deal more specifically with three major topics, neuronal processing, learning and brain functions and aims at establishing its achievements by the design of hardware and software systems to be compared with brain performances. It has consequently large overlaps with connectionism, machine learning and cognitive science but it is also different because it is specifically interested in understanding how these topics are implemented in real brains, with possible effects in neuroscience and in psychology.

Artificial Intelligence (AI) is another scientific domain overlapping CN. Though AI is not directly interested in understanding the brain, it is also a computational science aiming at building machines and softwares capable of intelligent behavior. There are consequently many fields where these two domains could cross-fertilize but others are more confrontational because each of the two domains relies on specific bases not to say dogmas.

In this Chapter, we propose to visit several fields of interest to better understand the rich relations between AI and CN. Some fields are directly related to some cognitive functions that have been deeply studied in AI and where CN has investigated and modeled the cerebral structures generally reported to be mainly involved in this cognitive function. This is the case with decision making and related processes of action selection and reinforcement learning involving basal ganglia (section 5) and also with language and the central role of the prefrontal cortex (section 6). Other fields are more methodological since CN can be also useful to propose to AI original, efficient and robust mechanisms for information representation. We begin to evoke here the problem of the level of description in computational models, from neurons to symbols (section 2), and turn to the representation of sensory information in the cortex (section 3) and to their multimodal integration in the hippocampus (section 4). Having mentioned these points of influence in computational mechanisms and in cognitive functions will be a good basis for discussion proposed in section 7.

2 From Neurons to Symbols

Initial modeling approaches of neuronal functioning [McCulloch and Pitts, 1943] and learning [Hebb, 1949] in the middle of the XX^{th} century are often considered

to be at the root of the development of both CN and connectionism (artificial neural networks without biological inspiration, presented in Chapter 12 of Volume 2), even if oldest efforts exist (as discussed by [Brunel and van Rossum, 2007]), particularly related to neuron excitability, in the context of CN.

At more microscopic levels of description, CN is also interested in understanding the behavior of single neurons through the dynamic evolution of their membrane potential, as it is the case for example in another seminal model by Hodgkin & Huxley [Nelson and Rinzel, 1995] or in other models dedicated to lower levels of description like dendrites, axons or even ion channels. Building such biophysical models is important to understand how neurons process information - how they compute - and some bottom up approaches propose to build on them up to the brain level and high level cognitive functions. This is specifically the case with huge simulation projects like the Blue Brain Project [Markram et al., 2015] and more recently the Human Brain Project, but many researchers wonder if such ascending projects are constrained enough to drive directly from sub-neuronal levels to cognition [Frégnac and Laurent, 2014; Chi, 2016]. In some way, the reciprocal criticism is sometimes given to cognitive psychology, stating that a descending approach, purely driven with functional consideration, is too vague to anchor in biological reality.

This is reminiscent of the duality in AI between the difficulty of making symbols emerge from numerical computation and the reciprocal Symbol Grounding Problem [Harnard, 1990]. This duality has been nicely addressed by D. Marr considering the visual brain as an information processing system, with the Tri-Level Hypothesis [Marr, 1982]. He proposes to define the computational level, describing the (visual) functions of the brain, the implementational level describing the underlying neuronal circuitry and, in-between, argues for an algorithmic level including the representations and processes employed by the implementational level to create the computational level.

Another property of this intermediate level is that it can be partly disconnected from the other two levels. At some moment, it can be interesting to wonder how cognitive functions are implemented by lower level mechanisms or to wonder how some mechanisms can result from a precise neuronal circuitry without having all the three levels of analysis in mind. Beyond the domain of vision, this fruitful analysis is often used by researchers exploiting a certain formalism of computation, like bayesian statisticians describing the brain [Friston, 2012] or theoreticians defining neuronal operations at the level of the population of neurons [Coombes, 2005] in reference to biological data about arrangements of neurons in repetitive circuits [Hubel et al., 1978]. Apart from the huge bottom-up projects mentioned above, most of the research in CN aiming at studying cognitive functions exploit such intermediate algorithmic levels with intermediate processing units (representing circuits or populations of neurons) and intermediate mechanisms (representing connectivity or learning rules).

Connectionism (artificial neural networks without biological inspiration) has also been used to encode, combine and manipulate symbols (Connectionist Symbol Processing [Sun and Alexandre, 1997]) and has been confronted to similar problems as AI and particularly to the Frame Problem, related to the difficulty of adequate

knowledge representation. This is also illustrated in the Searle's Chinese Room [Searle, 1980], where an agent can appear as intelligent by manipulating syntactic rules but turns out to have no knowledge about the meaning of its responses. This has been studied in Embodied AI [Pfeifer et al., 2007] by creating loops between the agent and its environment, through sensors and actuators, to create circular low-level sensorimotor relations, instead of elaborating complex high level formal rules [Brooks, 1990]. This approach is also frequently used in CN, often concerned by perception and action, with the brain seen as a way to associate them, as we are going to describe in the next sections.

3 Sensory Perception, Cortex and Unsupervised Learning

In the brain, the representation of information begins at the lower level by the representation of perceptual information received by sensors and has been primarily studied in the somatosensory system. The primary somatosensory cortex (S1 or SI) is a part of the cerebral cortex that is situated in the lateral postcentral gyrus, posterior to the central sulcus. It is, together with the primary motor cortex (anterior to the central sulcus), one of the first cortex to develop and certainly one of the most important. SI is innervated by sensory receptors (thermoreceptors, mechanoreceptors, chemoreceptors and nociceptors) that originate from both the surface of the body (e.g. skin) and from inside the body (e.g. bones and joints).

Even though the somatotopic organization of the cortex has been hypothesized more than a century ago by John Hughlings Jackson (1886) while studying epileptic patients, its existence was really demonstrated in the forties [Penfield and Boldrey, 1937; Marshall et al., 1937; Adrian, 1941]. This has been since then illustrated with the so-called sensory homunculus representation based on the work of Wilder Penfield using electrical stimulations on epileptic patients (even though this homunculus does not make justice to the amazing work of Wilder Penfield). These studies highlighted the somatotopic organization of the cortex and demonstrated a point-to-point correspondence between the surface of the body and the somatosensory cortex. However, such orderly representations are not restricted to the somatosensory cortex and, using similar approaches in the auditory, visual and motor domain [Diamond and Neff, 1957; Hubel and Wiesel, 1969; Merzenich and Kaas, 1980; Gould et al., 1986; Kaas, 1994], a large number of topographic maps have been described all over the cortex for different modalities (tonotopic maps, retinotopic maps, motor maps).

But they have different properties. Early observations of [Leyton and Sherrington, 1917] (as reported in [Lemon, 2008]) on the adult anthropoid apes demonstrated the ability of the motor cortex to recover from extensive cortical lesions. The authors hypothesized consequently the existence of a neural substrate and/or a mechanism for such extensive recovery. However, about forty years later, Hubel and Wiesel published a very influential paper [Hubel and Wiesel, 1959] that promoted the idea of fixed cortical representations following the post-natal developmental period (the so-called critical period). This hypothesis has prevailed for a long time until the

studies of [Merzenich and Kaas, 1982; Kaas et al., 1983; Kaas, 1991] provided experimental evidence for the somatosensory cortex reorganization after a peripheral nerve injury or amputation in the adult monkey.

The initial formation of these maps depends on a number of mechanisms occurring at the different stages in the development of the brain and the body and relies essentially on a complex molecular axon guidance [Tessier-Lavigne and Goodman, 1996] and a local selection of synapses and connections, based on temporal correlations [Katz and Shatz, 1996]. However, the exact mechanisms behind such organization has puzzled researchers for quite a long time. How can you preserve neighborhood relationships during the course of development such that ultimately, two neighbor skin patches tend to activate two neighbor cortical patches? Either you have to consider a genetic encoding where each cell “*knows*” where to connect or you have to consider an autonomous process that results in such orderly organization.

This latter hypothesis has been proposed in the seventies by Willshaw and von der Malsburg [Willshaw and von der Malsburg, 1976] with the idea of a self-organization. They proposed a model of the retina considering a set of cells that have short-range excitation (cooperation) and long range inhibition (competition). The distance between cells is provided by their actual position onto the cortical sheet, hence providing an explicit topology. They used this model to explain the formation of representation in V1 using a model of the retina that already possessed a topography. This model had a great influence and provided a very elegant explanation to the aforementioned question. Some years later, Kohonen [Kohonen, 1982] proposed an alternative model where he got rid of the lateral connectivity in favor of a winner takes all algorithm as well as an explicit lateral connectivity function. This model has become popular far beyond the computational neuroscience domain since it also provided an elegant solution to any vector quantization problem (see Chapter 12 of Volume 2).

Vector quantization (VQ) refers to the modelling of a probability density function into a discrete set (codebook) of prototype vectors (a.k.a. centroids) such that any point drawn from the associated distribution can be associated to a prototype vector. Most VQ algorithms try to match the density through the density of their codebook: high density regions of the distribution tend to have more associated prototypes than low density region. This generally allows to minimize the distortion as measured by the mean quadratic error. For a more complete picture, it is to be noted that there also exist some cases where only a partition of the space occupied by the data (regardless of their density) is necessary. In this case, one wants to achieve a regular quantification *a priori* of the probability density function. For example, in some classification problems, one wants to achieve a discrimination of data in terms of classes and thus needs only to draw frontiers between data regardless of their respective density. Such vector quantization can be achieved using several methods such as variations of the *k*-means method [Macqueen, 1967], Linde–Buzo–Gray (LBG) algorithm [Linde et al., 1980] or neural network models such as the self-organizing map (SOM, a priori topology) [Kohonen, 1982], neural gas (NG, no topology) [Martinetz et al., 1993] and growing neural gas (GNG, a posteriori topol-

ogy) [Fritzke, 1995]. Nonetheless, the SOM algorithm remains the most popular in the field of computational neurosciences since it gives a plausible account on the organization of receptive fields in sensory areas where adjacent neurons share similar representations as explained previously.

However, the stability and the quality of this self-organization depends heavily on a decreasing learning rate as well as a decreasing neighbourhood function. This is a major drawback of most neural map algorithms because it is thus necessary to have a finite set of observations to perform adaptive learning starting from a set of initial parameters (learning rate, neighbourhood or temperature) at time t_i down to a set of final parameters at time t_f . In the framework of signal processing or data analysis, this may be acceptable as long as we can generate a finite set of samples in order to learn it off-line. However, from a more general point of view, it is not always possible to have access to a finite set and we must face on-line learning as for example during a robotic task. The question is thus how to achieve both stability and plasticity.

To answer this question, variants of the original SOM learning algorithm have been proposed where the time dependency has been removed (see [Rougier and Boniface, 2011] for example). Based on several experiments in both two-dimensional, high-dimensional and dynamic cases, these variants allow for on-line and continuous learning ensuring a tight coupling with the environment. Following up on these ideas, [Detorakis and Rougier, 2012, 2014] investigated the formation and maintenance of ordered topographic maps in the primary somatosensory cortex and the reorganization of representations after sensory deprivation or cortical lesion. Their model is based on neural field theory using plastic feed-forward thalamocortical connections while cortico-cortical connections drive the competition mechanism.

Beyond these limitations, the original self-organizing map by Kohonen has become ubiquitous in the artificial intelligence landscape for learning and representing simple sensory information. SOM has been and is still used in a huge number of works in both image and signal processing, pattern recognition, speech processing, artificial intelligence, etc. Hundreds of variants of the original algorithm exist today [Kaski S. and Kohonen, 1998; Oja et al., 2003] such that it is literally impossible to review all of them here. Being both simple to implement and fast to compute, it is used in a wide variety of tasks ranging from navigation, compression, encoding, feature selection, and many others.

However, for the representation of more complex multimodal information (of “objects”), some other cerebral structures are necessary, namely, the hippocampus. This structure plays a central and critical role in the integration of various sources of information as well as in the encoding of multimodal regions such as the associative cortex. We now evoke the main ingredients of these complex processes.

4 The Hippocampus for Multimodal Binding

4.1 *Functional Organization of the Hippocampus: Implication for Learning*

The hippocampal system (HS) seems to participate in a lot of cognitive functions. In humans and animals, the HS plays an important role in spatial cognition but also in more general memorization processes. One important fact is that the hippocampus (Hip), a part of the HS with a seahorse shape, is connected through the entorhinal cortex (EC) to all the cortical associative areas. This makes the HS a very special convergence point to integrate multimodal information [McClelland et al., 1995]. The bilateral resection of the hippocampi in human causes a severe anterograde amnesia while the ability to learn new skills remains intact [Scoville and Milner, 1957]. Moreover, Alzheimer disease targets first the hippocampus and induces the same kind of memory issues. All these researches suggest the HS plays a major role in the formation of new memories and more specifically in episodic and autobiographic memories. Most of these memories would next be recoded at the cortical level by mechanisms which are still not well understood.

The HS can be seen as a generic tool to index or to build hash codes of the cortical activity [Teyler and DiScenna, 1986] in order to detect and learn in a fast way new events that cannot be easily detected by cortical neurons [Cohen and Eichenbaum, 1993; Eichenbaum et al., 1994; Bunsey and Eichenbaum, 1996; Buzsáki, 2013; Buzsáki and Moser, 2013] because of the practical limitation of the neurons connectivity. As a matter of fact, "typical" neurons are connected to around 10 000 other neurons limiting the capability of one neuron to detect complex events related to the co-occurrence of signals present in different sensory cortical areas for instance¹. Using the "small world" connectivity found in the brain allows any neuron in the cortex to contact any other neuron with a quite limited number of intermediate neurons. Yet, building such a network cannot be done in one shot. The neocortex is characterized by a slow learning rate and overlapping distributed representations allowing the extraction of the general statistical structure of the environment, while the hippocampus learns rapidly, using separated representations to encode the details of specific events while suffering minimal interference [O'Reilly and Rudy, 2000] providing both structures a quite complementary role in learning. This dual system for learning and memorization could be used in AI to limit the learning to the statistic of "important" events.

This fast learning capability of the Hip is supported by the properties of the Long Term Potentiation (LTP) found in the granular cells of the dentate gyrus (DG) and in the pyramidal cell of the Cornu Ammonia areas (CA). The specific organization of the recurrent connections in the CA3 region has been exploited in numerous models of auto associative memories [Marr et al., 1971; Hopfield, 1982; McNaughton and

¹ We can imagine this limitation is due to wiring issues: if some neurons were connected to all the neurons in other cortical areas, the size and weight of the brain would increase dramatically due to the space need for the dendritic trees of these neurons.

Morris, 1987] in order to explain pattern completion or pattern retrieval and even for the learning of sequences of events. In these models, the dentate gyrus (DG), one of the major inputs to CA3 is supposed to perform pattern separation or pattern orthogonalization of the activities coming from and through EC [Marr et al., 1971; McNaughton and Morris, 1987; Treves and Rolls, 1994]. More recent theories propose that the hippocampus is a predictive auto-encoder [Gluck and Myers, 1993] playing a major role in detecting new complex events and allowing their learning thanks to its reciprocal connections with the septal nuclei which controls the acetylcholine (ACh) neuromodulation [Hasselmo et al., 1995, 1996]. The ACh provided by the medial septum seems to mediate the learning and memory capabilities in the rest of the brain. For instance, it has been shown that lesions of the HS disrupt the acquisition of long-latency conditioned responses [Berger and Thompson, 1978] such as the eyeblink conditioning learned in the cerebellum [Thompson, 1986; Kim et al., 1995].

Hence, the neurobiological results teach us the brain uses at least two memory systems: the cortex learns on the long term but slowly the statistical structures of our interactions with the environment while the HS learns quickly (but for a maximum of few weeks) some compressed codes allowing to detect novelty and to control cortical learning.

4.2 The Hippocampus in Navigation Tasks: an Example of Multimodal Integration

The discovery of place cells in the hippocampus [O'Keefe and Dostrovsky, 1971; Morris et al., 1982] and later the discovery of grid cells [Fyhn et al., 2004] in the dorso median entorhinal cortex (dMEC) has emphasized the role of the hippocampus in spatial cognition. In this framework, the HS is supposed to play a specific role in spatial cognition and even to constitute a "cognitive map" [O'Keefe and Nadel, 1978]. A lot of works have focused on explaining how neurons in Hip can become specific to a given place. Some of the models start from visual information and show that a code built from the concatenation of several visual inputs is sufficient to recognize one place [Arleo and Gerstner, 2000; Milford et al., 2004; Krichmar et al., 2005]. The use of more specific codes to recognize one place as a constellation of landmark x azimuth couples (using conjunctive cells) is also used to improve generalization capabilities [Zipser, 1985; Bachelder and Waxman, 1994; Gaussier and Zrehen, 1995] when the animal is moving. Other models are using the distance instead of the azimuth [Burgess et al., 1997] or can use both since the elevation can be seen as a distance measure [Giovannangeli et al., 2006]. Yet, most of these models suppose that some place-action associations can be learned presumably thanks to the output of the Hip in the direction of the basal ganglia to control the direction of the action to be performed in a given place (see next section for more details on the basal ganglia).

The nature of the inputs and their coding can have very important impact on how place recognition can be used. For instance, if the azimuth information is coded thanks to a 1 dimension neural field (a bubble of activity centered on the neuron associated to a given azimuth) then the landmark-azimuth conjunctions are sensitive to the azimuth variation between the learned place and the actual place. The resulting place cells exhibit very large and reliable place fields [Gaussier et al., 2000; Giovannangeli et al., 2006]. In a room of 10x10m, the place field can easily be 2x2 m and can scale with the size of the environment. It is really an interesting property since after the learning of few place-action associations in the vicinity of a "goal" location, the competition between actions allows reaching the goal place from never visited places and from locations far away from the learned places². In this case, the capability of one neuron to recognize one place is no more important. It is only the rank of the place cells activity level (their firing rate) that matters: the action is chosen according to the winning place (competitive process). This is an important change from classical place cells models since the issue is no more to recognize one place but to be able to reach that place. In real world conditions, being sure of recognizing one place can be quite difficult when landmarks are moved or occluded (the effect of these changes on the cells activities can be higher than when moving away from the learned place). Yet, the rank of the place cells will remain the same: all the neurons are usually impacted in the same way if landmarks are randomly hidden or displaced. For instance, using 40 visual landmarks while learning places allows maintaining a good homing behavior even if more than half of the landmarks are hidden (basically 2 to 3 well recognized landmarks are sufficient). This generalization capability can also be very interesting to find a shortcut or to perform a detour according to the experimental situation. Moreover, if the landmark x azimuths conjunctions are selected in a 180 degrees image instead of a 270 degrees panorama then the cell activities are no more place cells but view cells and could explain the recording of "view cells" in the monkey hippocampus [Rolls and O'Mara, 1995] as opposed to the place cells in the rat. Yet, the place fields found in the CA region of the Hip look like really small (size about 20 cm) as compared to the large place fields described here. One hypothesis could be that those place cells would be located before the Hip in the ventromedial part of EC and that they would not be considered as place cells because of their broad receptive field [Quirk et al., 1992]. The small place fields in the Hip would be explained by the result of a competition layer allowing to obtain place field with a size equal to the mean distance between the learned places. In a new environment these cells will continue to react and provide some activities allowing to recognize the new places as a compound code. However, these visual place cells cannot explain how place cells can be maintained in the dark.

This approach is sometimes opposed to models supposing place cells are primarily built from path integration information [McNaughton et al., 1996; Touretzky and Redish, 1996; Redish and Touretzky, 1997] and that the HS is dedicated to spatial cognition. These models suppose path integration is computed in a discrete way

² to the extent that these places still belong to the same visual environment i.e. that they are in the area surrounded by the visual landmarks used for learning.

using a pre-existing two dimensional grid of neurons to store successive positions [Wan et al., 1994; Touretzky and Redish, 1996; Samsonovich and McNaughton, 1997; Redish and Touretzky, 1997; Clark and Taube, 2012]. Starting from a given place, the integration of velocity signal and direction of movement (supported by head direction cells) allows predicting a new position and then an update from place to place.

Hence, even if these different models differs in the way they suppose the place cells are built, all of them emphasize the role of the hippocampus in merging multimodal information: i.e exteroceptive information (vision, tactile, odor...) and proprioceptive information (speed, orientation,...). Robotics experiments provide an interesting way to test the limitations and the emergent properties of these different models.

4.3 Grid Cells in the Entorhinal Cortex : Information Compression and Coding

The seminal finding of grid cells³ in EC [Hafting et al., 2005] has reinforced the idea of the HS as a cartesian map. The models based on a direct computation of place recognition from the association of the dynamical memory of the departure place and a speed vector have been transformed to use the grid cells instead of the place cells [McNaughton et al., 2006; Fuhs and Touretzky, 2006; Burak and Fiete, 2009]. In these attractor models, grid activity is explained by the folding of a 2-dimensional Cartesian map representing the physical environment (a torus which has the advantage that it avoids the side effects related to the borders of the map). As pointed out by [Burak and Fiete, 2006], the first continuous attractor models [Fuhs and Touretzky, 2006; McNaughton et al., 2006; Samsonovich and McNaughton, 1997] work correctly only if the activity bubble moves exactly in register with the rat position. In other words, the network must precisely integrate rat velocity (which, in turn, must fit the environment discretization used in the simulation). If speed or movement direction does not correspond exactly to the parameters used for the discretization of the grid (in terms of angle and distance⁴), there is an accumulation of errors inducing a rapid blurring of the grid activity [Burak and Fiete, 2006]. [Burak and Fiete, 2009] proposed a solution for this issue. Yet, this model has still a lot of constraints on the network connectivity to work correctly. But above all, these models suppose that the entorhinal cortex and/or the hippocampus are devoted to navigation (see "the hippocampus as a cognitive map" [O'Keefe and Nadel, 1978]) and still need visual information for the (re)calibration of the path integration system.

However, if the hippocampus is not dedicated to spatial computation how to explain grid cell activities in the EC [Hafting et al., 2005]? [Gaussier et al., 2007;

³ grid cells: cells with a spatial firing frequency related to the wandered distance and direction of the animal movements.

⁴ distance = speed * time.constant of the computation time in the hippocampal loop

Jauffret et al., 2015] propose that grid cell activity results from a special case of the compression of the cortical activity projected onto the EC in order to build a hash code usable by the HS to recover information and detect complex novel states. Hence, grid cell activity would result from the simple projection and merging of a long-distance path integration onto the dorso medial entorhinal cortex (dMEC). Using a kind of modulo projection such that the same cell is activated from neurons associated to different distances allows the building of a very compact code with grid activity able to differentiate correctly different locations if the modulo factors are prime.

Recent studies on the lateral entorhinal cortex (LEC) and the way it can merge visual information for instance could reconcile both approaches. However, it is a matter of debate whether the hippocampus performs the path integration by itself or is just a generic structure to build a compact code of some path integration performed outside the hippocampus [Gaussier et al., 2007]. Anyhow, these works have led to efficient bio-inspired architectures [Gaussier and Zrehen, 1995; Gaussier et al., 1998; Krichmar et al., 2005; Milford and Wyeth, 2008] merging path integration and visual information in order to build robust place cells and to recalibrate path integration [Arleo and Gerstner, 2000; Gaussier et al., 2007; Jauffret et al., 2015].

4.4 Implication of the Hippocampus in Planning and Transition Recognition

Using direct place-action associations either with a strict recognition of place (need of learning a large number of places to pave regularly the environment) or with a competition mechanism (allowing to build a Voronoi tessellation of the environment and playing with the generalization properties of place cells) allows obtaining good performances in repetitive tasks (learning an habitual behavior).

When the goals can change, the place-action associations need to be relearned (long procedure) or need to be duplicated for each potential context or goal (with as many motor mappings as the number of goals). Reinforcement learning lacks some plasticity to explain specific learning capabilities such as latent learning [Tolman, 1948]. In this case, it is useful to learn independently the graph or the cognitive map connecting the known places so as to use them for planning the route in the direction of any known place without the need for more learning [Mataric, 1991; Schmajuk and Thieme, 1992; Guazzelli et al., 1998]. If we suppose a fronto-parietal cognitive map is ultimately built from the known places using Hebbian learning (neighbor places are connected to each other), several update rules allow to see the shortest path to reach a given goal location. For instance, if the connection weights are constant (or inversely proportional to the distance or to the difficulty to reach a place) and if the neuron activity for one node on the cognitive map is the maximum of the incoming activities then after the diffusion of the goal activity onto the map,

following the maximum gradient of activity allows taking the shortest path to reach that goal location.

Yet, there is a need for introducing an algorithm to read this gradient information since at a given place, the rat has not access to the activity in future places (until it reaches one of them). This problem is usually solved by using an ad hoc algorithm having a global view of the cognitive map to select the next place to be reached⁵ but as far as we suppose the neurons are only performing local computations there is a clear homonculus issue. One solution is to use a vicarious trial and error approach [Schmajuk and Thieme, 1992], where the animal is checking the different alternative directions before choosing the most active one. This solution is correct if important changes can be perceived from the decision point (i.e. the junction where the choice has to be made).

Unfortunately, in a real size environment, it is more likely that corridors or paths will look like quite identical for over long distances making this approach inefficient. One solution is to suppose the hippocampus role might be to predict transitions of multimodal events [Schmajuk, 1991; Grossberg and Merrill, 1996; Banquet et al., 1997]. Working with "transition cells" instead of classical steady state place cells solves the issue of the cognitive map readout by building a graph of transition cells and allowing a given transition cell to be connected to a single action [Revel et al., 1998; Gaussier et al., 2002; Banquet et al., 2005; Hirel et al., 2013]. As a matter of fact, building a cognitive map is no longer learning to connect for instance place "A" to place "B" then next to place "C" but is learning instead that transition "AB" is connected to transition "BC". If the current recognized place is "A" and "A" is connected to "B" and "D" for instance then the transitions "AB" and "AD" will have the same activity when being in A (both transition are predicted at the same level). If a diffusion activity from the goal to the different transitions on the map adds a little more activity to "AB" than "AD" then the action associated to "AB" will win and the agent can move in the correct direction since one transition is always associated to one unique direction of movement.

Coming back to the hippocampus, learning temporal transitions implies to have access both to the previous and the current places and in parallel to be able to integrate the movement from A to B (i.e. the heading direction of the animal [Sharp, 1999] when going from "A" to "B"). In our case, we consider two kinds of short term memories inside the HS. First, the recurrent connections in the dentate gyrus (DG) between the granular cells and the mossy cells could allow building a temporal trace that the CA3 pyramidal cells can use to learn when a new place will be reached (let's say when "B" will be reached from the memory trace of the previous place "A"). Next, CA1 neurons could use a more rustic short term memory relying on EC3 pyramidal cells (and built from EC2 activities) to build transition cells in CA1. Then, CA1 activity could be propagated to the fronto-parietal network to build long term cognitive maps and also to the nucleus accumbens (ACC) to learn and propose the different possible transitions. Finally, the neurons in the ACC receiving the activities from the fronto-parietal network and the activity from EC1 could decide

⁵ yet selecting the correct action can be tricky when the place fields are not regular enough

about the transition to be selected. Interestingly, recent results on "time cells" in the hippocampus [Naya and Suzuki, 2011; Kraus et al., 2013; Pfeiffer and Foster, 2013; Eichenbaum, 2014] are coherent with this view of the hippocampus as a system to predict temporal transitions [Hirel et al., 2013].

In conclusion, the HS is mainly known and studied in rodents for its implication in navigation tasks while in primates and humans it is known to participate also in higher cognitive functions such as the building of autobiographical memories. The specific architecture and connectivity of the hippocampal system makes it important for the building of declarative (or explicit) memories as opposed to the procedural (or implicit) memories directly stored in the cortical and subcortical structures. The HS capability to build a spatio-temporal code and to perform fast or even one shot learning is a crucial element to build place codes for navigation tasks, to predict transitions between multi modal states and even to implement some timing properties for sequences learning and recognition. The complementary role of the HS and cortex in the building of different kind of memories is certainly a key element for the understanding of the human intelligence and our capability to filter the huge amount of data our brain faces during real life interactions. It is clear that mimicking the way the hippocampus is interacting with the cortical areas and the basas ganglia could be helpful in AI when facing the management of big data in real time conditions. Moreover, neurobiological data and robotics experiments show that at least for navigation tasks, recognizing perfectly a place is not necessary to reach that place in a robust way. Hence recognizing a place or an object is much more being able to come back to that place or to grasp and act on that object. As opposed to a purely passive recognition scheme, a sensory-motor approach of cognition has strong impact on the way information has to be coded and emphasizes the importance of taking into account the action selection in all the stages of the cognitive processes (i.e. even in the design of a pattern recognition system).

5 Action selection, reinforcement learning and the basal ganglia

5.1 The Basal Ganglia as a Central Action Selection Device in the Brain

The basal ganglia are a group of inter-connected subcortical nuclei, which receive massive convergent input from most regions of cortex, hippocampus and amygdala and output to targets in the thalamus and brainstem. It is thus considered as a privileged region in the brain having access to diverse information (from sensorimotor, emotional, motivational and reward information to associative memory and episodic memory) and directly influencing motor regions [Alexander et al., 1990; Voorn et al., 2004]. It has even been hypothesized to implement some sort of dimensionality reduction – formalized as a similar process to a Principal Component

Analysis – in order to sort out the most important and relevant features among the large amount of information to which an individual is confronted in order to decide which motor response should be performed at a given moment [Bar-Gad et al., 2000].

One of the main theories of the role of the basal ganglia is that it constitutes a neural substrate of a central action selection device within the brain [Mink, 1996; Redgrave et al., 1999; Gurney et al., 2001a,b]. The theory explicitly makes interesting links with the problem of module selection within distributed, modular control architectures in the field of Engineering. It argues that a central selection device minimizes the number of connections as well as human prior knowledge. While a distributed selection architecture requires inhibitory interconnections between all modules, a central selection device only requires connections between each module and itself. While a subsumption architecture [Brooks, 1986] requires a preprogrammed, fixed priority scheme, a central selection device can compare modules with a common currency (thus enabling learning mechanisms as described below).

These features have attracted the attention of several researchers in Artificial Intelligence and Robotics, who wanted to study the potential benefits of neuro-inspired basal ganglia action selection models for artefacts compared to Engineering methods [Prescott et al., 1999; Girard et al., 2003; Khamassi et al., 2005; Girard et al., 2005; Khamassi et al., 2006; Girard et al., 2008]. In particular, Girard and colleagues have shown that such basal ganglia models have some persistence properties which enable a robot to save more energy than a classical winner-takes-all mechanism [Girard et al., 2003]. The anatomical loops that the basal ganglia form with the cortex and the thalamus provide a feedback mechanism to the action selection mechanism, enabling a selected channel (or action) to have a slight bonus over competing channels during the competition at the next time step. Hence the persistence in action selection. Such a feedback is particularly relevant to avoid "hesitations" in the system when two channels have the same saliency and are thus oscillatorily selected one after the other, resulting in the absence of any displacement by the robot. In other words, such a persistence mechanism can help solving the Buridan donkey paradox, which in extension to a discussion raised by Aristotle suggests that a donkey having to choose between food and water, and being as hungry as thirsty, would remain unable to decide, immobile and would die from starvation.

Another interesting property of action selection mechanisms in the basal ganglia is that they operate through disinhibition rather than excitation of the selected action [Chevalier and Deniau, 1990]. Basal ganglia output nuclei are indeed tonically inhibiting their motor targets, so that the selection of an action results in the suppression of the inhibition of the corresponding channel (hence a disinhibition). This results in a faster action initiation compared to alternative mechanisms where action selection would result in the initiation of a motor command. Moreover, it has been shown in the oculomotor domain that such inhibition mechanisms enable to prevent the blocking of voluntary sight orientation movements by compensatory eye movements due to the vestibulo-ocular reflex [Berthoz, 2002].

The basal ganglia is not only fed with cortical input information (i.e. neurotransmission) but is also strongly modulated by neuromodulators such as dopamine, no-

radrenaline, serotonin and acetylcholine. These neuromodulators have been hypothesized to perform a meta-control or meta-learning process on top of basal ganglia action selection mechanisms [Doya, 2002]. They have been shown to both affect neural plasticity and instantaneous gain modulations of information transmission [Servan-Schreiber et al., 1990; Reynolds et al., 2001]. One of the advantages of such neuromodulation mechanisms is that they enable to reduce the combinatorial explosion in the number of required channels to represent different variations of the same action. For instance, rather than representing a different channel for the same action performed with different response vigors, with different speeds of movements, or in relation to different contexts and goals, the neuromodulatory system can learn to perform different levels of modulations on the same action channel [Niv et al., 2007; Humphries et al., 2012].

Finally, the basal ganglia are anatomically organized into different territories which form different parallel loops with different cortical and thalamic territories [Alexander et al., 1990; Voorn et al., 2004]. Such organization is well conserved through evolution [Redgrave et al., 1999; Prescott et al., 1999; Stephenson-Jones et al., 2011]. Since each basal ganglia territory and group of nuclei is seen as roughly organized in the same manner as other territories [Gerfen and Wilson, 1996], such a parallelism permits a reuse of the same selection mechanisms applied to different action domains (locomotor, oculomotor, etc.), different levels of selection (movement selection, action selection, action plan selection, strategy selection, goal selection). These parallel loops are thus considered as engaged in different cognitive functions such as motor, associative, limbic and oculomotor [Alexander et al., 1990; Haber et al., 2000; Uylings et al., 2003]. An architecture with two simulated basal ganglia loops has for example been successfully used in a robotic task to select among both appetitive actions (directions of movement) and consumatory ones (stops at different reloading stations) and to coordinate them [Girard et al., 2005]. Furthermore, the limbic loop having a privileged access to reward as well as other emotional information from the amygdala, hypothalamus and brainstem, and being in a position of influence over other loops through neuromodulatory projections, this suggests that the limbic loop of the basal ganglia may play a central role in learning [Graybiel, 1998; Bornstein and Daw, 2011; Ito and Doya, 2011; Khamassi and Humphries, 2012; van der Meer et al., 2012].

5.2 The Basal Ganglia as a Center for Reinforcement Learning

Animals' ability to learn from their own experience and errors, in particular in the context of sparse reward and punishment signals, is considered to rely on reinforcement learning processes [Doya, 2000; Foster et al., 2000; Balleine and O'Doherty, 2010; Khamassi and Humphries, 2012; van der Meer et al., 2012; Palminteri et al., 2015]. The most central theory, developed in the field of Artificial Intelligence, currently considers that such learning relies on: 1) the competition between actions, resulting in action selection as a function of the actions' relative probabilities; 2) the

anticipation of the value of rewards and punishments that could follow the execution of the action; 3) the computation of a *reward prediction error* comparing what was expected with what is actually obtained; 4) the use of such a reward prediction error as a feedback (*i.e.* positive, negative or null reinforcement signal) to update either the probability of the performed action or the predictive value associated to the action and to the stimuli present in this context [Sutton and Barto, 1998].

This formalism can be seen as an extension of the Rescorla-Wagner model [Rescorla and Wagner, 1972] developed in Psychology, in which learning requires prediction errors to explain various properties of associative learning during animals classical conditioning. Prediction errors can indeed explain the *blocking* phenomenon – when a stimulus B cannot be associated with a reward if it is presented together with a stimulus A which is already fully predictive of the reward –, and cases of *overexpectation* – when the concomitant presentation of two reward predictive stimuli influences behavior as if they were adding up, to form a stronger prediction.

A particular subgroup of RL algorithms implementing what is called Temporal-Difference (TD) learning extends the Rescorla-Wagner model in that prediction error signals contain three terms rather than two. The Rescorla-Wagner indeed compares past expectation with present outcome (*e.g.* reward). The TD learning rule adds to this comparison a term representing future expectations of reward. As a consequence, a reinforcement signal can be computed even before the reward is attained by comparing temporally consecutive expectations of reward – hence the term *Temporal-Difference*: *e.g.* when an action leads to a situation or state where reward expectations are higher than previous ones, this action should be reinforced.

Since nearly twenty years, this theory has provided Neuroscientists with formal tools which contributed to important breakthroughs in the understanding of neural correlates of learning. Reinforcement Learning models turned out to be able to explain a wide range of adaptive behaviors experimentally observed both in humans (*e.g.* [Frank et al., 2009; Balleine and O’Doherty, 2010]) and in non-human animals (*e.g.* [Yin and Knowlton, 2006; Khamassi and Humphries, 2012]). This formalism also enabled to explain a variety of neural correlates of learning [Schultz et al., 1997; Khamassi et al., 2008]. The most striking example and probably the most central in the field is the observation that phasic responses of dopaminergic neurons (which send massive neuromodulatory projections to the basal ganglia [Haber et al., 2000]) follow the profile of reward prediction errors as they are formalized by the RL theory: an increase in activity when the outcome of action is better than expected; a decrease in activity when it is worse than expected; an absence of response when it meets the expectations [Schultz et al., 1997]. In addition, the third term of the learning rule mentioned above enables TDRL models to account for reward anticipation signals in the rat ventral striatum (the main nucleus in the limbic part of the basal ganglia) [Khamassi et al., 2008] as well as in some dopaminergic neurons [Bellot et al., 2012].

The accumulation of neurophysiological results corroborated by this computational theory has also enabled to establish that the learning of reward values and action values depends on plasticity in projections from the cortex to the basal gan-

glia (in particular to the striatum, the main input structure of the basal ganglia), and that these adjustments depend on dopaminergic signals sent from the substantia nigra pars compacta and the ventral tegmental area [Barto, 1995; Houk et al., 1995; Schultz et al., 1997; Reynolds et al., 2001]. Numerous computational models of the basal ganglia were derived from these experimental results [Houk et al., 1995; Schultz et al., 1997; Doya, 2000; Joel et al., 2002; Khamassi et al., 2005; Frank, 2005; Guthrie et al., 2013; N’Guyen et al., 2014], and were built on the central assumption mentioned above that the basal ganglia play a critical role in action selection [Redgrave et al., 1999; Gurney et al., 2001a,b].

Strikingly, this field of investigation has entertained an important dialog between AI and Neuroscience. One of the main recent examples is the interest that neuroscientists have gained for the distinction between different types of TDRL algorithms in the field of AI, namely: Actor-Critic, Q-learning, SARSA. These three algorithms use different information in their learning rule, which is respectively based on state value (independent from the action), the maximal action value in the current state, the value of the action chosen to be performed in the current state at the next timestep. These three variations of TDRL thus lead to different profiles of reward prediction error signals. As a consequence, several neuroscience groups have designed experiments to specifically investigate whether reinforcement signals in the brain are consistent with either Actor-Critic, Q-learning or SARSA. Contradictory results have been obtained by different groups so far [Morris et al., 2006; Roesch et al., 2007; Bellot et al., 2012]. There could exist differences between species (monkeys and rats) or between experimental configurations (presentation of single versus pairs of reward predicting stimuli). Future dialogs between Neuroscience and AI thus promise fertile exchanges along this line of research.

Finally, in the last decade computational neuroscientists have discovered that the classical distinction between learning strategies considered in AI and Machine Learning, namely model-based and model-free RL, also applies to Neuroscience by capturing different experimentally observed behavioral strategies in mammals, and related brain activities in different networks involving different basal ganglia territories [Daw et al., 2005; Khamassi and Humphries, 2012; Dollé et al., 2018]. More precisely, it turns out that mammals often start learning a task by trying to build an internal model of the task states, actions and transitions between them. This enables initial flexible behavior – which in particular enables to quickly adapt to task changes – in parallel to the slower acquisition of model-free local action values in the background. Once the latter learning process has converged – if the stability and familiarity of the task permit this long convergence – it starts expressing its learned behavioral sequences. This permits to free the parts of the brain which are responsible for model-based decisions (including the prefrontal cortex), thus enabling quicker but at the same time more rigid action selection. If the task changes after a long period of stability, it is more difficult for mammals to break their habits and adapt to the new task contingencies.

Nevertheless, this line of Neuroscience research promises interesting future inspiration for AI by investigating how the brain efficiently coordinates these different types of learning. For instance, Daw and colleagues have suggested that the relative

uncertainty within each learning system could help the brain decide which system should control behavior at any given moment [Daw et al., 2005]. In addition, more recent models can explain a variety of animal behavior in different tasks by employing a meta-controller which meta-learns which learning system was the most efficient in each state of each task [Dollé et al., 2018]. This suggests principles for the coordination of multiple learning systems which could inspire Artificial Intelligence in return. Recent applications of these principles to Robotics suggest that this can work on real robots in a variety of tasks [Caluwaerts et al., 2012; Renaudo et al., 2014]. While classical AI-based robots usually try to solve a given problem with a single algorithm (either planning or reinforcement learning), trying to make this algorithm the best possible on the considered solution, it turns out that different problems require different algorithms [Kober et al., 2013]. Here the neuro-inspired solution suggests that a system or cognitive architecture coordinating multiple learning processes (as it is the case in the basal ganglia) may benefit from the advantages of each process/algorithm, and may learn by itself which one is the most appropriate in each given situation [Caluwaerts et al., 2012; Renaudo et al., 2015; Khamassi et al., 2016]. While this may lead to good but suboptimal performance in a given problem (as is the case for mammals), this may enable the agent to adapt to many different situations, which could be of great potential interest for Artificial Intelligence research. In return, more recent developments in AI and Robot learning permitting to efficiently and dynamically tune the exploration-exploitation trade-off in reinforcement learning [Wang et al., 2016; Khamassi et al., 2018] as well as to bootstrap learning with prioritized experience replay [Schaul et al., 2015] could greatly inspire future improvements of reinforcement learning models in Neuroscience [Caze et al., 2018].

6 Language and the prefrontal cortex

As the most recent arrival in the long evolution of the primate cortex, the prefrontal cortex (PFC) is considered to one of the pillars of higher cognitive function [Fuster, 1991; Goldman-Rakic, 1987; Miller and Cohen, 2001; Wang et al., 2015]. Generally speaking, there is a transition in the posterior to anterior extent of the cortex from sensory-motor functions posteriorly to progressively more integrated and abstract functions as we move more anterior in cortex, culminating in prefrontal cortex [Fuster, 1991]. Thus, PFC has been a candidate for numerous computational models [Bastos et al., 2012; Dominey et al., 1995; Duncan, 2001; O'Reilly and Frank, 2006]. Complimentary to its place as a highly associative area, one of the principal neurophysiological characteristics of prefrontal cortex is the density of local recurrent connections [Goldman-Rakic, 1987]. A second neurophysiological characteristic of the prefrontal cortex is its privileged relation with the basal ganglia and thalamus. We will see how these characteristics can lead to impressive computational capabilities.

The computational power of recurrent networks has traditionally been demonstrated in a number of domains [Douglas et al., 1995; Hermans and Schrauwen, 2012; Pearlmutter, 1995]. One of the great challenges in modeling recurrent networks concerns how to adapt the recurrent connection weights, as it is difficult to assign credit to recurrent connections [Pearlmutter, 1995]. When an input excites the network, activation can circulate through the recurrent connections numerous times before the output is generated. If the output is incorrect, and one wants to modify a recurrent connection in order to reduce the error, one must keep in memory the different roles of this connection throughout the numerous cycles of activation. A number of technical solutions that can involve cutting of the recurrent history to simplify the credit assignment, or unrolling the recurrent cycles, have been employed [Elman, 1990; Jordan, 1986; Pearlmutter, 1995]. The resulting recurrent neural network (RNN) models have a long and rich history in cognitive science [Cleeremans and McClelland, 1991; Elman, 1991].

An alternative approach is to retain the rich temporal dynamics within recurrent network, with no cutoff, by maintaining the recurrent connections fixed, and modifying connections between the recurrent units and the output units. This approach was first invented by Dominey and colleagues [Dominey, 1995; Dominey et al., 1995]. Later it was again independently developed by Maass as the liquid state machine [Maass et al., 2002], and by Jaeger as the echo state network [Jaeger, 2001; Jaeger and Haas, 2004]. These three approaches have been integrated under the title of reservoir computing [Lukosevicius and Jaeger, 2009].

The initial motivation for these recurrent networks was to understand how the prefrontal cortex encodes sequential structure. Barone and Joseph [Barone and Joseph, 1989] studied neural activity in the prefrontal cortex of monkeys that had been trained to perform a sequence learning task that involved watching the presentation of a visual sequence on a response button board, and then after a short delay, reproducing the sequence by touching the buttons on the board in the same order that they were presented. They observed that neurons in the dorsolateral prefrontal cortex (DLPFC) displayed two characteristic responses to stimuli in the sequence task. First, as had previously been observed, the neurons were spatially selective, with preferences for stimuli in particular locations in the retinal image. The second characteristic was new, and revolutionary: many of these neurons also displayed a “sequence rank” effect, that is, they had preferences for stimuli that had appeared first, second or last in the input sequence. Thus, the spatial selectivity in many neurons was modulated by the rank or order of the element in the sequence. This indicated that DLPFC embodies a mechanism for discriminating the order of items in a perceptual sequence.

These observations motivated us to develop a recurrent network that received retinotopic (spatially organized) inputs, and combined these inputs with mixed excitatory and inhibitory connections. We reasoned that this would lead to a form of mixed selectivity combining spatial location and sequence rank, as observed in the primate. This is indeed what we observed, thus demonstrating that such recurrent networks have significant sequence learning capability, and that their coding corresponds to that seen in the prefrontal cortex [Dominey et al., 1995].

More recently we have used this same type of reservoir model to solve a rather complex task where the system should search for the one rewarded target amongst four possibilities, then repeat that response for several rewards, before starting anew with the search for the new rewarded target. Model neurons showed a remarkable similarity to neurons recorded in the primate anterior cingulate cortex, including the non-linear mixture of task relevant parameters [Enel et al., 2016]. It is now considered that these recurrent networks with feedback have universal computing properties [Maass et al., 2007, 2002]. Intuitively, the recurrent connections project the inputs into an infinitely high dimensional space that incorporates past history. Modifiable readout neurons can then be trained to select the appropriate representation for the task at hand.

Such universal computing should be appropriate for language learning. Language learning can be characterized as learning from paired sentence-meaning examples how to generate the corresponding meaning for a new sentence. In order to have a simplified version of such a task, we looked to neurolinguistic tasks used to measure human language comprehension. Caplan and colleagues developed a task that exploited nine different sentence types, and measured patients ability to perform thematic role assignment, or to determine who did what to whom [Caplan et al., 1985]. Typical sentences tested included

1. Dative passive: The elephant was given to the monkey by the rabbit.
2. Subject-Object relative: The elephant that the monkey hit hugged the rabbit.

Patients were asked to read such sentences, and then by pointing to pictures, indicate in order the agent, the object and the recipient of the described actions. Patients with lesions in the left hemisphere, in the region surrounding the sylvian fissure, displayed deficits in using these grammatical cues to determine who did what to whom. Instead, they relied on the so called canonical order and indicating that the first mentioned noun was the agent, second object, third recipient [Caplan et al., 1985].

We reasoned that this thematic role assignment problem could be considered as a sequence processing problem, where the system should take the input sentence and reorder the nouns (if necessary) into the agent, object, recipient order, based on the grammatical function words like “was”, “to” and “by”. In this context, meaning can be represented in a predicate-argument form such as “predicate(agent, object, recipient)”, and sentences are represented as sequences of words. As illustrated in figure 1, these grammatical words are processing the recurrent network (that we posit to be in BA47 PFC region), and the open class words (nouns and verbs) are held in a working memory (that we posit to be in prefrontal cortex BA44). Through learning, the system associates different states of activity in the recurrent network with selection of different elements in the working memory, thus linking different sentence forms with different re-ordering of the open class elements into the agent, object, recipient order, as required for Caplan’s task [Dominey et al., 2003].

For example, “It was the cat¹ that the dog² chased³” becomes “The dog² chased³ the cat¹”. From an abstract perspective, this corresponds to an abstract structure ABC-BAC which is a sequence of six elements where the second triplet is a system-

atically transformed version of the first [Dominey et al., 1998]. Interestingly this model made a strong prediction about the equivalent processing of linguistic sequences (i.e. sentences) and non-linguistic abstract sequences in the brain: both sentences, and non-linguistic abstract sequences that required a systematic re-ordering of certain elements should recruit a common brain network for structure processing, while language should require additional processing to integrate semantic contents.

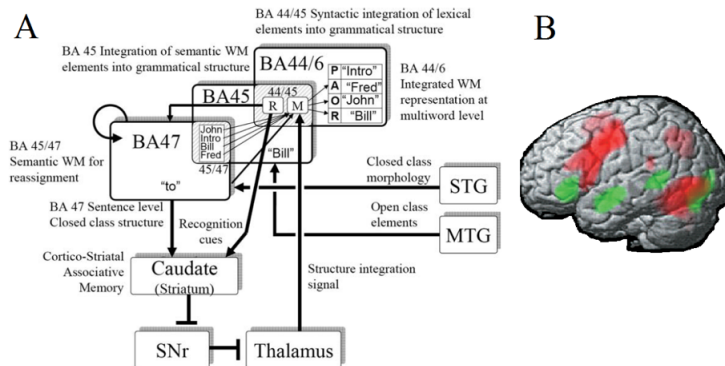


Fig. 1 A. Neural network model of fronto-striatal system for sentence and non-linguistic abstract sequence processing (from [Dominey et al., 2009]). B. Neural activity for sentence and abstract sequence processing (from [Hoen et al., 2006]). Common activity for sentences and sequences in red (corresponding to recurrent network and structure mapping), and language specific activity in green (corresponding to integration of semantic content into grammatical structure via ventral pathway and BA45).

One of the remarkable features of language processing, and thus a key property for any model of language processing is the ability to learn grammatical structure from a limited number of examples, and to generalize this learning to new grammatical sentences. By using more optimized learning techniques to learn the associations between activity in the recurrent network and the corresponding responses, we have demonstrated such generalization. Figure 2 illustrates the updated reservoir model for language comprehension, and generalization performance on untrained sentence types in a parameter exploration where network parameters are systematically varied.

An interesting property of the behavior of the network can be seen in Figure 3. The two panels illustrate activity in the readout neurons that code for the semantic role of the first noun in two different grammatical types. In the second sentence, this noun is the agent or subject of the main (second) verb and the object of the relative (first) verb. The two sentences are identical up to the arrival of the fourth word. In the subject-subject sentence depicted in A, the model accurately predicts what happens and there is little change in the output indicated at the arrow. In the less frequent subject-object sentence in B, the model's prediction must be updated at the arrival of the fourth word causing a large visible change in activity. This is

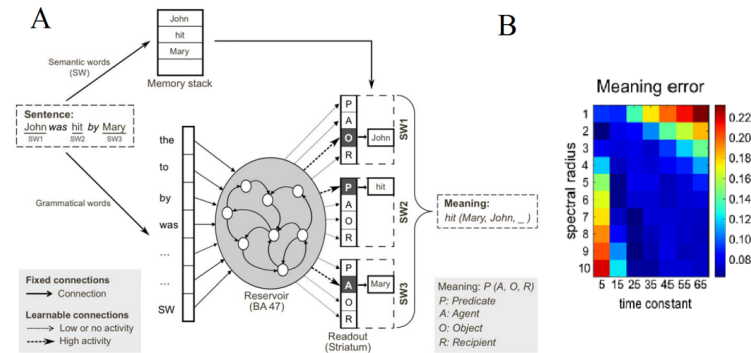


Fig. 2 Simplified reservoir model for sentence comprehension. A. Grammatical words are input to the reservoir. Through learning, connections from reservoir to the readout can associate patterns of activity in the reservoir with activation of neurons that represent the semantic role (predicate, agent, object or recipient) for each semantic word in the sentence. B. Performance error in comprehension in a test of generalization to new untrained constructions. Note an extended parameter space with good generalization.

precisely the kind of activity change that is seen in the human brain in response to lower frequency grammatical structures, referred to as the syntactic positive shift or P600 (a positivity that comes 600ms after the offending word) [Frisch et al., 2002; Hagoort and Brown, 2000].

We have seen that recurrent network models of PFC can perform complex tasks like language comprehension. Recent work at the forefront between computational neuroscience and machine learning has contributed to the concept that the prefrontal cortex is a recurrent network that has universal coding properties. Rigotti and colleagues have shown that recurrent reservoir networks inherently display mixed selectivity in their single units, and that this mixed selectivity is precisely the high dimensional coding required for solving complex cognitive tasks [Fusi et al., 2016; Rigotti et al., 2013]. Interestingly this is the same kind of mixed selectivity that we modeled (mixing spatial location and sequence rank) in our first instantiation of the reservoir concept [Dominey et al., 1995], where a network of neurons connected by fixed recurrent connections provides a universal high dimensional encoding, and modifiable connections allow the system to learn to associate these representations with the desired output. Future research should attempt to further understand these neurocomputational systems, and address questions concerning how the encoding of task related context can render these systems even more powerful.

7 Conclusion

AI and CN are two scientific domains developing their own formalisms but they are both interested in understanding how cognitive functions can emerge from com-

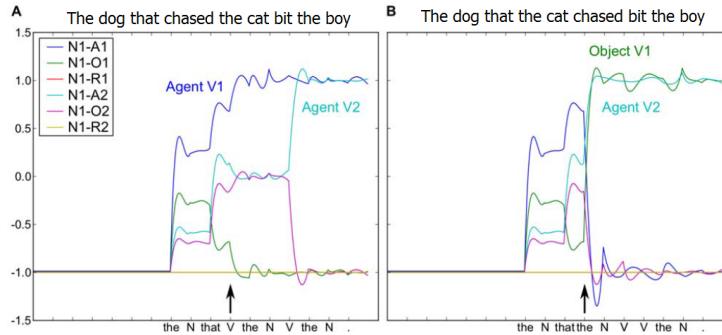


Fig. 3 Processing of Subject-Relative and Object-Relative sentences in corpus where subject-relatives are more frequent than object-relatives. A. Subject-relative. For the word “V” following “that”, there is relatively small change in the readout neurons, indicating that the predictions of the model were essentially confirmed. B. Object-relative. For the “the” following “that”, there is a significant shift in activity, corresponding to a reassignment of the most probable coded meaning. From [Hinaut and Dominey, 2013].

putational mechanisms. Specifically, CN takes a strong inspiration from the brain circuitry, which is undoubtedly a choice source of information to study the emergence of cognitive functions. Accordingly, CN can provide AI with original and first-hand mechanisms related to cognition. In addition, Neuroscience is a very dynamical domain extracting more and more information from the mysterious brain, and CN is perpetually renewed by following these progresses and sometimes contributing to them. In this Chapter, we have proposed some illustrations on important questions on which the domains can interact.

The problem of representation of information is a central issue in AI, from the Physical Symbol System hypothesis to more recent questions evoked in most Chapters of Volume 1, about the representation of complex objects or events [Russell and Norvig, 2003]. Whereas the existence in the cortex of topographic maps of neurons responding each to preferential stimuli (cf. section 3) can give arguments for localized representations at the symbolic and subsymbolic levels, the dynamics of recurrent networks in the PFC (cf. section 6) and the hippocampus (cf. section 4) is rather reminiscent of distributed encoding. It is also important to mention that the corresponding neuronal models are associated to well explored learning rules that allow to build both kinds of representation from sampling in the environment. Concerning the elaboration of information representation, a related topic is about the opposition between top down and bottom up approaches in AI (cf. section 2). Pieces of evidence have been proposed here that Marr’s algorithmic level could be fed with sensory information in an ascending way and controlled in the opposite way by structures like the PFC (cf. section 6).

Whereas learning is a central mechanism in cognition, it has a particular status in AI. Sometimes it is not directly addressed and it is believed that information can be efficiently injected in a cognitive system as formalized knowledge. Sometimes

it is the central topic of an algorithm in Machine Learning that is intended to solve the considered cognitive task by processing iteratively data received in experimental cases. From this duality between adaptation by integration of knowledge or data [Sun and Alexandre, 1997], CN argues that knowledge injected into a cognitive system can correspond to the very slow learning, at the scale of species evolution, of the structures and characteristics of the neuronal circuitry and proposes, for adaptation at shorter time constants, a large variety of biologically inspired learning algorithms that we have evoked throughout this Chapter, particularly mentioning their strong links to classical algorithms in Machine Learning presented in Chapter 12 of Volume 1 and Chapter 12 of Volume 2, and their cross-fertilization. It is thus notable that the recently developed Deep Learning approach evoked in Chapter 12 of Volume 2, which offers among the most efficient artificial learning systems, makes strong references to neuroscience and in particular to the visual system [Yamins and DiCarlo, 2016], although we must also relativize this type of analogy. Specifically, the effectiveness of these systems is based largely on their training from very large corpus of examples, whereas natural learning generally has fewer trials to adapt.

Furthermore, as it can be observed in behaving animals (or robots !), an intelligent behavior often corresponds to the capacity to adapt to a variety of tasks and not to be a specialist for only one task. This is clearly one domain where AI is only starting out (but consider the domains of Lifelong Machine Learning and of Artificial General Intelligence) and where CN is more mature, allowing to coordinate multiple ways of learning in modular networks, as it has been evoked several times above.

A related characteristic of central interest in intelligent behavior is autonomy, also too poorly dealt with in AI. Here also, CN is some steps beyond because its systemic view is more adapted to implement a versatile agent able to adapt by itself to changing and unknown conditions and also because recent models of the loops between the basal ganglia and the PFC [Koechlin et al., 2003; O'Reilly et al., 2010] are beginning to address the functions of self-evaluation of performances and cognitive control, not to mention consciousness, which are also fundamental ingredients to autonomous behavior. Nevertheless, these approaches are still far from proposing an architecture of control making the agent fully autonomous and able to exploit previously elaborated knowledge in new circumstances, and to identify these circumstances, particularly in the case of unstationary environment and a lot of work remains to be done in that direction.

Our capacity to adapt to completely unexpected situations is certainly a central explanation to the fact that our brain today spends most of its time making symbolic manipulations (with natural language, with mathematics, but also with digital devices) it was not necessarily designed for at the origin. Understanding how the systemic arrangement of cerebral structures presented here produces this general purpose information processing system specially interested in symbolic analysis is certainly another rich domain of interaction between AI and CN. Particularly, inspiration from natural sciences can bring to AI an original understanding of this problem, focusing for example on developmental issues (how skills can be installed one after the other to produce an increasingly mature system) and also on the importance

of social interactions (imitation, teaching) and their impact on the development of the cerebral system.

Whereas general intelligence is clearly an important goal of CN-AI interactions but remains on a long-term perspective, we have also shown in this Chapter that very concrete interactions are already existing for cognitive functions corresponding to more classical goals of AI, like perception, navigation, decision making (including action selection, reasoning, planning) and language. We have explained here that CN can provide AI with useful mechanisms and principles of information representation, by deciphering the brain circuitry involved in these functions. In addition, CN can offer another decisive contribution for helping computational models to master more widely these functions. Their development is certainly linked to the better integration of multimodal sensorimotor flows and to their interconnections, one with the others. The systemic approach of CN is clearly an asset in that direction.

In conclusion, CN has already demonstrated interesting contributions to AI at the methodological level, for information representation, processing and learning and also at the functional level for the implementation of a variety of cognitive functions [Hassabis et al., 2017]. Other types of contribution are particularly precious because they exploit the unique capability of CN to develop an integrated approach of brain modeling, in a systemic view. They should be encouraged for the development of AI in domains like autonomous robotics, for multimodal cognitive functions like decision making and language and also for general intelligence.

References

- Adrian, E. D. (1941). Afferent discharges to the cerebral cortex from peripheral sense organs. *Journal of Physiology London*, 100.
- Alexander, G., Crutcher, M., and DeLong, M. (1990). Basal ganglia-thalamocortical circuits: parallel substrates for motor, oculomotor, "prefrontal" and "limbic" functions. *Prog Brain Res*, 85:119–146.
- Arleo, A. and Gerstner, W. (2000). Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity. *Biological Cybernetics*, 83(3):287–299.
- Bachelder, I. A. and Waxman, A. M. (1994). Mobile robot visual mapping and localization: A view-based neurocomputational architecture that emulates hippocampal place learning. *Neural networks*, 7(6):1083–1099.
- Balleine, B. W. and O'Doherty, J. P. (2010). Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology*, 35(1):48–69.
- Banquet, J., Gaussier, P., Dreher, J. C., Joulain, C., Revel, A., and Günther, W. (1997). Space-time, order, and hierarchy in fronto-hippocampal system: A neural basis of personality. In Matthews, G., editor, *Cognitive Science Perspectives on Personality and Emotion*, volume 124, pages 123–189, Amsterdam. North Holland.

- Banquet, J.-P., Gaussier, P., Quoy, M., Revel, A., and Burnod, Y. (2005). A hierarchy of associations in hippocampo-cortical systems: Cognitive maps and navigation strategies. *Neural Computation*, 17(6):1339–1384.
- Barone, P. and Joseph, J. P. (1989). Prefrontal cortex and spatial sequencing in macaque monkey. *Experimental Brain Research*, 78(3):447–464.
- Barto, A. (1995). Adaptive critics and the basal ganglia. In Houk, J., Davis, J., and Beiser, D., editors, *Models of Information Processing in the Basal Ganglia*, chapter 11, pages 215–232. MIT Press, Cambridge, MA.
- Bar-Gad, I., Havazelet-Heimer, G., Goldberg, J. A., Ruppin, E., and Bergman, H. (2000). Reinforcement-driven dimensionality reduction – a model for information processing in the basal ganglia. *J Basic Clin Physiol Pharmacol*, 11:30520.
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4):695 – 711.
- Bellot, J., Sigaud, O., and Khamassi, M. (2012). Which temporal difference learning algorithm best reproduces dopamine activity in a multi-choice task? In Ziemke, T., Balkenius, C., and Hallam, J., editors, *From Animals to Animats 12*, volume 7426 of *Lecture Notes in Computer Science*, pages 289–298. Springer Berlin / Heidelberg.
- Berger, T. W. and Thompson, R. F. (1978). Neuronal plasticity in the limbic system during classical conditioning of the rabbit nictitating membrane response. i. the hippocampus. *Brain research*, 145(2):323–346.
- Berthoz, A. (2002). *Le Mouvement*. Odile Jacob, Paris, France.
- Bornstein, A. M. and Daw, N. D. (2011). Multiplicity of control in the basal ganglia: computational roles of striatal subregions. *Current opinion in neurobiology*, 21(3):374–380.
- Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, R.A-2:14–23.
- Brooks, R. A. (1990). Elephants don’t play chess. *Robotics and Autonomous Systems*, 6(1):3 – 15.
- Brunel, N. and van Rossum, M. C. W. (2007). Lapicque’s 1907 paper: from frogs to integrate-and-fire. *Biological Cybernetics*, 97(5):337–339.
- Bunsey, M. and Eichenbaum, H. (1996). Conservation of hippocampal memory function in rats and humans. *Nature*, 379:255–257.
- Burak, Y. and Fiete, I. (2006). Do we understand the emergent dynamics of grid cell activity? *J Neurosci*, 26(37):9352–4; discussion 9354.
- Burak, Y. and Fiete, I. R. (2009). Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput Biol*, 5(2):e1000291.
- Burgess, N., Donnett, J. G., Jeffery, K. J., John, O., et al. (1997). Robotic and neuronal simulation of the hippocampus and rat navigation. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 352(1360):1535–1543.
- Buzsáki, G. (2013). Cognitive neuroscience: time, space and memory. *Nature*, 497(7451):568–569.

- Buzsáki, G. and Moser, E. I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature neuroscience*, 16(2):130–138.
- Caluwaerts, K., Staffa, M., S., N., Grand, C., Dollé, L., Favre-Félix, A., Girard, B., and Khamassi, M. (2012). A biologically inspired meta-control navigation system for the psikharpax rat robot. *Bioinspiration and Biomimetics*, 7(2):025009.
- Caplan, D., Baker, C., and Dehaut, F. (1985). Syntactic determinants of sentence comprehension in aphasia. *Cognition*, 21(2):117 – 175.
- Caze, R., Khamassi, M., Aubin, L., and Girard, B. (2018). Hippocampal replays under the scrutiny of reinforcement learning models. *Journal of Neurophysiology*, page in revision.
- Chevalier, G. and Deniau, J. (1990). Disinhibition as a basic process in the expression of striatal functions. *Trends Neurosci*, 13(7):277–280.
- Chi, K. R. (2016). Neural modelling: abstractions of the mind. *Nature*, 531.
- Churchland, P. and Sejnowski, T. J. (1992). *The Computational Brain*. Cambridge, MA: MIT Press.
- Clark, B. J. and Taube, J. S. (2012). Vestibular and attractor network basis of the head direction cell signal in subcortical circuits. *Front. Neural Circuits*, 6(7):10–3389.
- Cleeremans, A. and McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of experimental psychology. General*, 120(3):235–253.
- Cohen, N. and Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system*. MA:MIT.
- Coombes, S. (2005). Waves, bumps and patterns in neural field theories. *Biol. Cybern.*, 93:91–108.
- Daw, N., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci*, 8(12):1704–1711.
- Dayan, P. and Abbott, L. F. (2001). *Theoretical neuroscience: computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.
- Detorakis, G. and Rougier, Nicolas, P. (2012). A Neural Field Model of the Somatosensory Cortex: Formation, Maintenance and Reorganization of Ordered Topographic Maps. *PLoS ONE*, 7(7):e40257.
- Detorakis, G. I. and Rougier, N. P. (2014). Structure of receptive fields in a computational model of area 3b of primary sensory cortex. *Frontiers in Computational Neuroscience*, 8:26.
- Diamond, I. T. and Neff, W. D. (1957). Ablation of temporal cortex and discrimination of auditory patterns. *Journal fo Neurophysiology*, 20.
- Dollé, L., Chavarriaga, R., Guillot, A., and Khamassi, M. (2018). Interactions of spatial strategies producing generalization gradient and blocking: A computational approach. *PLoS computational biology*, 14(4):e1006092.
- Dominey, P. F. (1995). Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning. *Biological Cybernetics*, 73(3):265–74.

- Dominey, P. F., Arbib, M. A., and Joseph, J.-P. (1995). A Model of Cortico-Striatal Plasticity for Learning Oculomotor Associations and Sequences. *Journal of Cognitive Neuroscience*, 7(3):311–336.
- Dominey, P. F., Hoen, M., Blanc, J.-M., and Lelekov-Boissard, T. (2003). Neurological basis of language and sequential cognition: Evidence from simulation, aphasia, and {ERP} studies. *Brain and Language*, 86(2):207 – 225.
- Dominey, P. F., Inui, T., and Hoen, M. (2009). Neural network processing of natural language: Ii. towards a unified model of corticostriatal function in learning sentence comprehension and non-linguistic sequencing. *Brain and Language*, 109(2-3):80–92.
- Dominey, P. F., Lelekov, T., Ventre-Dominey, J., and Jeannerod, M. (1998). Dissociable processes for learning the surface structure and abstract structure of sensorimotor sequences. *Journal of Cognitive Neuroscience*, 10(6):734–51.
- Douglas, R., Koch, C., Mahowald, M., Martin, K., and Suarez, H. (1995). Recurrent excitation in neocortical circuits. *Science*, 269(5226):981–985.
- Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr Opin Neurobiol*, 10(6):732–739.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Netw*, 15(4-6):495–506.
- Duncan, J. (2001). An adaptive coding model of neural function in prefrontal cortex. *Nature Reviews Neuroscience* 2, 2(11):820–829.
- Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews Neuroscience*, 15(11):732–744.
- Eichenbaum, H., Otto, T., and Cohen, N. (1994). Two functional components of the hippocampal memory system. *Behav. Brain Sci.*, 17.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2):179–211.
- Elman, J. L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7:195–225.
- Enel, P., Procyk, E., Quilodran, R., and Dominey, P. F. (2016). Reservoir computing properties of neural dynamics in prefrontal cortex. *PLoS Comput Biol*, 12(6):1–35.
- Foster, D., Morris, R., and Dayan, P. (2000). A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus*, 10(1):1–16.
- Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: a neurocomputational account of cognitive deficits in medicated and nonmedicated parkinsonism. *Journal of Cognitive Neuroscience*, 17(1):51–72.
- Frank, M. J., Doll, B. B., Oas-Terpstra, J., and Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature neuroscience*, 12(8):1062–1068.
- Frégnac, Y. and Laurent, G. (2014). Where is the brain in the human brain project? *Nature*, 513.
- Frisch, S., Schlesewsky, M., Saddy, D., and Alpermann, A. (2002). The {P600} as an indicator of syntactic ambiguity. *Cognition*, 85(3):B83 – B92.

- Friston, K. (2012). A Free Energy Principle for Biological Systems. *Entropy (Basel, Switzerland)*, 14(11):2100–2121.
- Fritzke, B. (1995). A growing neural gas network learns topologies. In Tesauro, G., Touretzky, D., and Leen, T., editors, *Advances in Neural Information Processing Systems 7*, pages 625–632. MIT Press, Cambridge MA.
- Fuhs, M. C. and Touretzky, D. S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience*.
- Fusi, S., Miller, E. K., and Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. *Current Opinion in Neurobiology*, 37:66 – 74. Neurobiology of cognitive behavior.
- Fuster, J. M. (1991). Chapter 10 the prefrontal cortex and its relation to behavior. In Holstege, G., editor, *Role of The Forebrain in Sensation and Behavior*, volume 87 of *Progress in Brain Research*, pages 201 – 211. Elsevier.
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., and Moser, M.-B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688):1258–1264.
- Gaussier, P., Banquet, J., Sargolini, F., Giovannangeli, C., Save, E., and Poucet, B. (2007). A model of grid cells involving extra hippocampal path integration, and the hippocampal loop. *Journal of integrative neuroscience*, 6(03):447–476.
- Gaussier, P., Joulain, C., Banquet, J.-P., Leprêtre, S., and Revel, A. (2000). The visual homing problem: an example of robotics/biology cross fertilization. *Robotics and autonomous systems*, 30(1):155–180.
- Gaussier, P., Moga, S., Quoy, M., and Banquet, J.-P. (1998). From perception-action loops to imitation processes: A bottom-up approach of learning by imitation. *Applied Artificial Intelligence*, 12(7-8):701–727.
- Gaussier, P., Revel, A., Banquet, J.-P., and Babeau, V. (2002). From view cells and place cells to cognitive map learning: processing stages of the hippocampal system. *Biological Cybernetics*, 86:15–28.
- Gaussier, P. and Zrehen, S. (1995). Perac: A neural architecture to control artificial animals. *Robotics and Autonomous Systems*, 16(2):291–320.
- Gerfen, C. and Wilson, C. (1996). The basal ganglia. In Swanson, L., Björklund, A., and Hökfelt, T., editors, *Handbook of chemical neuroanatomy*, volume Vol 12: Integrated Systems of the CNS, Part III, chapter Chapter II, pages 371–468. Elsevier Science B.V.
- Giovannangeli, C., Gaussier, P., and Banquet, J. (2006). Robustness of visual place cells in dynamic indoor and outdoor environment. *International Journal of Advanced Robotic Systems*, 3(2):115–124.
- Girard, B., Cuzin, V., Guillot, A., Gurney, K., and Prescott, T. (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. (2005). Integration of navigation and action selection functionalities in a computational model of cortico-basal ganglia-thalamo-cortical loops. *Adaptive Behavior*, 13(2):115–130.
- Girard, B., Tabareau, N., Pham, Q., Berthoz, A., and Slotine, J.-J. (2008). Where neuroscience and dynamic system theory meet autonomous robotics: a contracting basal ganglia model for action selection. *Neural Networks*, 21(4):628–641.

- Gluck, M. A. and Myers, C. E. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus*, 3(4):491–516.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. *Comprehensive Physiology*.
- Gould, J. H. I., Cusick, C. G., Pons, T. P., and Kaas, J. H. (1986). The relationship of corpus callosum connections to electrical stimulation maps of motor, supplementary motor, and frontal eye fields in owl monkeys. *Journal of Comparative Neurology*, 247.
- Graybiel, A. (1998). The basal ganglia and chunking of action repertoires. *Neurobiol Learn Mem*, 70(1-2):119–136.
- Grossberg, S. and Merrill, J. (1996). The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. *Journal of Cognitive Neuroscience*, 8:257–277.
- Guazzelli, A., Bota, M., Corbacho, F. J., and Arbib, M. A. (1998). Affordances, motivations, and the world graph theory. *Adaptive Behavior*, 6(3-4):435–471.
- Gurney, K., Prescott, T., and Redgrave, P. (2001a). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological Cybernetic*, 84(6):401–410.
- Gurney, K., Prescott, T., and Redgrave, P. (2001b). A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour. *Biol Cybern*, 84(6):411–423.
- Guthrie, M., Leblois, A., Garenne, A., and Boraud, T. (2013). Interaction between cognitive and motor cortico-basal ganglia loops during decision making: a computational study. *Journal of neurophysiology*, 109(12):3025–3040.
- Haber, S., Fudge, J., and McFarland, N. (2000). Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *J Neurosci*, 20(6):2369–2382.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436:801–806.
- Hagoort, P. and Brown, C. M. (2000). {ERP} effects of listening to speech compared to reading: the p600/sps to syntactic violations in spoken sentences and rapid serial visual presentation. *Neuropsychologia*, 38(11):1531 – 1549.
- Harnard, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42:335–346.
- Hassabis, D., Kumaran, D., Summerfield, C., and Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245–258.
- Hasselmo, M. E., Schnell, E., and Barkai, E. (1995). Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region ca3. *The Journal of neuroscience*, 15(7):5249–5262.
- Hasselmo, M. E., Wyble, B. P., and Wallenstein, G. V. (1996). Encoding and retrieval of episodic memories: role of cholinergic and gabaergic modulation in the hippocampus. *Hippocampus*, 6(6):693–708.
- Hebb, D. (1949). *The Organization of Behavior : A Neuropsychological Theory*. Wiley, New York.

- Hermans, M. and Schrauwen, B. (2012). Recurrent kernel machines: Computing with infinite echo state networks. *Neural Comput.*, 24(1):104–133.
- Hinault, X. and Dominey, P. F. (2013). Real-time parallel processing of grammatical structure in the fronto-striatal system: a recurrent network simulation study using reservoir computing. *PLoS ONE*, 8(2):e52946.
- Hirel, J., Gaussier, P., Quoy, M., Banquet, J.-P., Save, E., and Poucet, B. (2013). The hippocampo-cortical loop: spatio-temporal learning and goal-oriented planning in navigation. *Neural Networks*, 43:8–21.
- Hoen, M., Pachot-Clouard, M., Segebarth, C., and Dominey, P. F. (2006). When Broca experiences the Janus syndrome: an ER-fMRI study comparing sentence comprehension and cognitive sequence processing. *Cortex*, 42(4):605–623.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558.
- Houk, J., Adams, J., and Barto, A. (1995). A model of how the basal ganglia generate and use neural signals that predict reinforcement. In Houk, J., Davis, J., and Beiser, D., editors, *Models of Information Processing in the Basal Ganglia*, chapter 13, pages 249–270. MIT Press, Cambridge, MA.
- Hubel, D. and Wiesel, T. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of physiology*, 148(3):574–591.
- Hubel, D. and Wiesel, T. (1969). Anatomical demonstration of columns in the monkey striate cortex. *Nature*, 221.
- Hubel, D. H., Wiesel, T. N., and Stryker, M. P. (1978). Anatomical demonstration of orientation columns in macaque monkey. *J. Comp. Neur.*, 177:361–380.
- Humphries, M., Khamassi, M., and Gurney, K. (2012). Dopaminergic control of the exploration-exploitation trade-off via the basal ganglia. *Frontiers in Neuroscience*, 6:9.
- Ito, M. and Doya, K. (2011). Multiple representations and algorithms for reinforcement learning in the cortico-basal ganglia circuit. *Current opinion in neurobiology*, 21(3):368–373.
- Jaeger, H. (2001). The "echo state" approach to analysing and training recurrent neural networks-with an erratum note. *Bonn, Germany: German National Research Center for Information Technology GMD Technical Report*, 148:34.
- Jaeger, H. and Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *science*, 304(5667):78–80.
- Jauffret, A., Cuperlier, N., and Gaussier, P. (2015). From grid cells and visual place cells to multimodal place cell: a new robotic architecture. *Frontiers in Neurobotics*, 9(1).
- Joel, D., Niv, Y., and Ruppin, E. (2002). Actor-critic models of the basal ganglia: new anatomical and computational perspectives. *Neural Netw*, 15(4-6):535–547.
- Jordan, M. I. (1986). Serial order: A parallel, distributed processing approach. Technical Report 8604, Institute for Cognitive Science, University of California, San Diego.
- Kaas, J. (1991). Plasticity of sensory and motor maps in adult mammals. *Annual review of neuroscience*, 14(1).

- Kaas, J., Merzenich, M., and Killackey, H. (1983). The reorganization of somatosensory cortex following peripheral nerve damage in adult and developing mammals. *Annual Review of Neuroscience*, 6:325–356.
- Kaas, J. H. (1994). *Aotus: The owl monkey*, chapter The organization of sensory and motor cortex in owl monkeys. Baer, J. F. and Weller, R. E. and Kakoma, I.
- Kaski S., Jangas, J. and Kohonen, T. (1998). Bibliography of self-organizing map papers: 1981-1997. *Neural Computing Surveys*.
- Katz, L. C. and Shatz, C. J. (1996). Synaptic activity and the construction of cortical circuits. *Science*, 274.
- Khamassi, M., Girard, B., Clodic, A., Devin, S., Renaudo, E., Pacherie, E., Alami, R., and Chatila, R. (2016). Integrator of action, joint action and learning in robot cognitive architectures. *Intellectica*, 2016/1:169–203.
- Khamassi, M. and Humphries, M. (2012). Integrating cortico-limbic-basal ganglia architectures for learning model-based and model-free navigation strategies. *Frontiers in Behavioral Neuroscience*, 6:79.
- Khamassi, M., Lacheze, L., Girard, B., Berthoz, A., and Guillot, A. (2005). Actor-critic models of reinforcement learning in the basal ganglia: from natural to artificial rats. *Adaptive Behavior*, 13:131–148.
- Khamassi, M., Martinet, L.-E., and Guillot, A. (2006). Combining self-organizing maps with mixtures of experts: application to an actor-critic model of reinforcement learning in the basal ganglia. In *From Animals to Animats 9, LNAI 4095*, pages 394–405. Berlin, Heidelberg: Springer-Verlag.
- Khamassi, M., Mulder, A., Tabuchi, E., Douchamps, V., and Wiener, S. (2008). Anticipatory reward signals in ventral striatal neurons of behaving rats. *European Journal of Neuroscience*, 28:1849–1866.
- Khamassi, M., Velentzas, G., Tsitsimis, T., and Tzafestas, C. (2018). Robot fast adaptation to changes in human engagement during simulated dynamic social interaction with active exploration in parameterized reinforcement learning. *IEEE Transactions on Cognitive and Developmental Systems*, page in revision.
- Kim, J. J., Clark, R. E., and Thompson, R. F. (1995). Hippocampectomy impairs the memory of recently, but not remotely, acquired trace eyeblink conditioned responses. *Behavioral neuroscience*, 109(2):195.
- Kober, J., Bagnell, J. A., and Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, pages 1238–1274.
- Koechlin, E., Ody, C., and Kouneiher, F. (2003). The Architecture of Cognitive Control in the Human Prefrontal Cortex. *Science*, 302(5648):1181–1185.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43.
- Kraus, B. J., Robinson, R. J., White, J. A., Eichenbaum, H., and Hasselmo, M. E. (2013). Hippocampal “time cells”: time versus path integration. *Neuron*, 78(6):1090–1101.
- Krichmar, J. L., Seth, A. K., Nitz, D. A., Fleischer, J. G., and Edelman, G. M. (2005). Spatial navigation and causal analysis in a brain-based device modeling cortical-hippocampal interactions. *Neuroinformatics*, 3(3):197–221.

- Lemon, R. (2008). An enduring map of the motor cortex. *Experimental Physiology*, 93:798–802.
- Leyton, S. and Sherrington, C. (1917). Observations on the excitable cortex of the chimpanzee, orang-utan and gorilla. *Quarterly Journal of Experimental Physiology*, 11:135–222.
- Linde, Y., Buzo, A., and Gray, R. (1980). An algorithm for vector quantization design. *IEEE Trans. on Communications*.
- Lukosevicius, M. and Jaeger, H. (2009). Survey: Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.*, 3(3):127–149.
- Maass, W., Joshi, P., and Sontag, E. D. (2007). Computational aspects of feedback in neural circuits. *PLoS Comput Biol*, 3(1):1–20.
- Maass, W., Natschläger, T., and Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Comput.*, 14(11):2531–2560.
- Macqueen, J. B. (1967). Some methods of classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297.
- Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A., Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., Kahou, G. A., Berger, T. K., Bilgili, A., Buncic, N., Chalimourda, A., Chindemi, G., Courcol, J.-D., Delalondre, F., Delattre, V., Druckmann, S., Dumusc, R., Dynes, J., Eilemann, S., Gal, E., Gevaert, M. E., Ghobril, J.-P., Gidon, A., Graham, J. W., Gupta, A., Haenel, V., Hay, E., Heinis, T., Hernando, J. B., Hines, M., Kanari, L., Keller, D., Kenyon, J., Khazen, G., Kim, Y., King, J. G., Kisvarday, Z., Kumbhar, P., Lasserre, S., Le Bé, J.-V., Magalhães, B. R. C., Merchán-Pérez, A., Meystre, J., Morrice, B. R., Muller, J., Muñoz Céspedes, A., Muralidhar, S., Muthurasa, K., Nachbaur, D., Newton, T. H., Nolte, M., Ovcharenko, A., Palacios, J., Pastor, L., Perin, R., Ranjan, R., Riachi, I., Rodríguez, J.-R., Riquelme, J. L., Rössert, C., Sfyarakis, K., Shi, Y., Shillcock, J. C., Silberberg, G., Silva, R., Tauheed, F., Telfont, M., Toledo-Rodriguez, M., Tränkler, T., Van Geit, W., Díaz, J. V., Walker, R., Wang, Y., Zaninetta, S. M., DeFelipe, J., Hill, S. L., Segev, I., and Schürmann, F. (2015). Reconstruction and Simulation of Neocortical Microcircuitry. *Cell*, 163(2):456–492.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Henry Holt and Co., Inc., New York, NY, USA.
- Marr, D., Willshaw, D., and McNaughton, B. (1971). Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 262(841):23–81.
- Marshall, W. H., Woolsey, C. N., and Bard, R. (1937). Cortical representation of tactile sensibility as indicated by cortical potentials. *Science*, 85.
- Martinetz, T. M., Berkovich, S. G., and Schulten, K. J. (1993). Neural-gas network for vector quantization and its application to time-series prediction. *IEEE Trans. on Neural Networks*, 4(4):558–569.

- Mataric, M. J. (1991). Navigating with a rat brain: a neurobiologically inspired model. In *From Animals to Animats; Proceedings of the First International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge, Mass.
- McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419.
- McCulloch, W. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5:115–133.
- McNaughton, B., Barnes, C., Gerrard, J., Gothard, K., Jung, M., Knierim, J., Kudrimoti, H., Qin, Y., Skaggs, W., Suster, M., et al. (1996). Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *The Journal of Experimental Biology*, 199(1):173–185.
- McNaughton, B., Battaglia, F. P., Jensen, O., Moser, E. I., and Moser, M.-B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nat Rev Neurosci*, 7:663–678.
- McNaughton, B. L. and Morris, R. G. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in neurosciences*, 10(10):408–415.
- Merzenich, M. and Kaas, J. (1982). Reorganization of mammalian somatosensory cortex following peripheral nerve injury. *Trends in Neurosciences*, 5:434–436.
- Merzenich, M. M. and Kaas, J. H. (1980). *Progress in psychobiology and physiological psychology*, chapter Principles of organization of sensory-perceptual systems in mammals. Sprague, J. M. and Epstein, A. N.
- Milford, M. J. and Wyeth, G. F. (2008). Mapping a suburb with a single camera using a biologically inspired slam system. *IEEE Transactions on Robotics*, 24(5):1038–1053.
- Milford, M. J., Wyeth, G. F., and Rasser, D. (2004). Ratslam: a hippocampal model for simultaneous localization and mapping. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 1, pages 403–408. IEEE.
- Miller, E. K. and Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.*, 24:167–202.
- Mink, J. W. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Progress in neurobiology*, 50(4):381–425.
- Morris, G., Nevet, A., Arkadir, D., Vaadia, E., and Bergman, H. (2006). Midbrain dopamine neurons encode decisions for future action. *Nat Neurosci*, 9(8):1057–1063.
- Morris, R., Garrud, P., Rawlins, J., and O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, 297:24.
- Naya, Y. and Suzuki, W. A. (2011). Integrating what and when across the primate medial temporal lobe. *Science*, 333(6043):773–776.
- Nelson, M. and Rinzal, J. (1995). *The Hodgkin-Huxley model.*, chapter 4, pages 27–51. Bower and Beeman. The book of Genesis. Springer, New York.

- N'Guyen, S., Thurat, C., and Girard, B. (2014). Saccade learning with concurrent cortical and subcortical basal ganglia loops. *Frontiers in Computational Neuroscience*, 8:48.
- Niv, Y., Daw, N., Joel, D., and Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology (Berl)*, 191(3):507–520.
- Oja, M., Kaski, S., and Kohonen, T. (2003). Bibliography of self-organizing map papers: 1998-2001 addendum. *Neural Computing Surveys*.
- O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 34(1):171–175.
- O'Keefe, J. and Nadel, N. (1978). *The hippocampus as a cognitive map*. Clarendon Press, Oxford.
- O'Reilly, R. C. and Frank, M. J. (2006). Making Working Memory Work: A Computational Model of Learning in the Prefrontal Cortex and Basal Ganglia. *Neural Comp.*, 18(2):283–328.
- O'Reilly, R. C., Herd, S. A., and Pauli, W. M. (2010). Computational models of cognitive control. *Current Opinion in Neurobiology*, 20(2).
- O'Reilly, R. C. and Rudy, J. W. (2000). Computational principles of learning in the neocortex and hippocampus. *Hippocampus*, 10(4):389–397.
- Palminteri, S., Khamassi, M., Joffily, M., and Coricelli, G. (2015). Medial prefrontal cortex and the adaptive regulation of reinforcement learning parameters. *Nature Communications*, 6:8096.
- Pearlmutter, B. A. (1995). Gradient calculations for dynamic recurrent neural networks: a survey. *IEEE Trans Neural Netw*, 6(5):1212–1228.
- Penfield, W. and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex as studied by electrical stimulation. *Brain*, 60.
- Pfeifer, R., Bongard, J., and Grand, S. (2007). *How the body shapes the way we think: a new view of intelligence*. Bradford Books. MIT Press.
- Pfeiffer, B. E. and Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447):74–79.
- Prescott, T., Redgrave, P., and Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, 7:99–127.
- Quirk, G. J., Muller, R. U., Kubie, J. L., and Ranck, J. (1992). The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. *The Journal of Neuroscience*, 12(5):1945–1963.
- Redgrave, P., Prescott, T., and Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, 89(4):1009–1023.
- Redish, A. D. and Touretzky, D. S. (1997). Cognitive maps beyond the hippocampus. *Hippocampus*, 7(1):15–35.
- Renaudo, E., Girard, B., Chatila, R., and Khamassi, M. (2014). Design of a control architecture for habit learning in robots. In Duff, A., Lepora, N., Mura, A., Prescott, T., and Verschure, P., editors, *Biomimetic and Biohybrid Systems, Third International Conference, Living Machines 2014*, pages 249–260.
- Renaudo, E., Girard, B., Chatila, R., and Khamassi, M. (2015). Respective advantages and disadvantages of model-based and model-free reinforcement learning

- in a robotics neuro-inspired cognitive architecture. *Procedia Computer Science*, 71:178–184.
- Rescorla, R. and Wagner, A. (1972). A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In Black, A. and Prokasy, W., editors, *Classical Conditioning II: Current Research and Theory*, pages 64–99. Appleton-Century-Crofts, New-York.
- Revel, A., Gaussier, P., Lepretre, S., and Banquet, J. (1998). Planification versus sensory-motor conditioning: what are the issues ? In *SAB'98: From animals to animats 5*, pages 129–138.
- Reynolds, J., Hyland, B., and Wickens, J. (2001). A cellular mechanism of reward-related learning. *Nature*, 413(6851):67–70.
- Rigotti, M., Barak, O., Warden, M. R., Wang, X. J., Daw, N. D., Miller, E. K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451):585–590.
- Roesch, M., Calu, D., and Schoenbaum, G. (2007). Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nat Neurosci*, 10(12):1615–1624.
- Rolls, E. T. and O'Mara, S. M. (1995). View-responsive neurons in the primate hippocampal complex. *Hippocampus*, 5(5):409–424.
- Rougier, N. P. and Boniface, Y. (2011). Dynamic Self-Organising Map. *Neurocomputing*, 74(11):1840–1847.
- Russell, S. and Norvig, P. (2003). *Artificial Intelligence - A Modern Approach*. Prentice-Hall, Upper Saddle River, New Jersey.
- Samsonovich, A. and McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. (2015). Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schmajuk, N. (1991). A neural network approach to hippocampal function in classical conditioning. *Behavioral Neuroscience*, 105(1):82–110.
- Schmajuk, N. and Thieme, A. (1992). Purposive behavior and cognitive mapping: a neural network model. *Biological Cybernetic*, 67:165–174.
- Schultz, W., Dayan, P., and Montague, P. (1997). A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599.
- Schwartz, E. L. (1990). *Computational Neuroscience*. MIT Press, Cambridge, MA, USA.
- Scoville, W. B. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, 20(1):11.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3:417–457.
- Servan-Schreiber, D., Printz, H., and Cohen, J. (1990). A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science*, 249(4971):892–895.
- Sharp, P. E. (1999). Comparison of the timing of hippocampal and subicular spatial signals: implications for path integration. *Hippocampus*, 9(2):158–172.

- Stephenson-Jones, M., Samuelsson, E., Ericsson, J., Robertson, B., and Grillner, S. (2011). Evolutionary conservation of the basal ganglia as a common vertebrate mechanism for action selection. *Current Biology*, 21(13):1081–1091.
- Sun, R. and Alexandre, F. (1997). *Connectionist-Symbolic Integration : From Unified to Hybrid Approaches*. Lawrence Erlbaum Associates.
- Sutton, R. and Barto, A. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Tessier-Lavigne, M. and Goodman, C. S. (1996). The molecular biology of axon guidance. *Science*, 274.
- Teyler, T. J. and DiScenna, P. (1986). The hippocampal memory indexing theory. *Behavioral neuroscience*, 100(2):147.
- Thompson, R. F. (1986). The neurobiology of learning and memory. *Science*, 233(4767):941–947.
- Tolman, E. (1948). Cognitive maps in rats and men. *The Psychological Review*, 55(4):189–208.
- Touretzky, D. S. and Redish, A. D. (1996). Theory of rodent navigation based on interacting representations of space. *Hippocampus*, 6(3):247–270.
- Treves, A. and Rolls, E. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, 4(3):374–391.
- Uylings, H., Groenewegen, H., and Kolb, B. (2003). Do rats have a prefrontal cortex? *Behav Brain Res*, 146(1-2):3–17.
- van der Meer, M., Kurth-Nelson, Z., and Redish, A. D. (2012). Information processing in decision-making systems. *The Neuroscientist*, 18(4):342–359.
- Voorn, P., Vanderschuren, L., Groenewegen, H., Robbins, T., and Pennartz, C. (2004). Putting a spin on the dorsal-ventral divide of the striatum. *Trends Neurosci*, 27(8):468–474.
- Wan, H., Touretzky, D., and Redish, A. (1994). Towards a computational theory of rat navigation. In Mozer, M., Smolensky, P., Touretzky, D., Elman, J., and Weigend, A., editors, *Proc. of the 1993 Connectionist Models Summer School*, pages 11–19. Lawrence Erlbaum Associates.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., and Botvinick, M. (2016). Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*.
- Wang, L., Li, X., Hsiao, S. S., Lenz, F. A., Bodner, M., Zhou, Y. D., and Fuster, J. M. (2015). Differential roles of delay-period neural activity in the monkey dorsolateral prefrontal cortex in visual-haptic crossmodal working memory. *Proc. Natl. Acad. Sci. U.S.A.*, 112(2):E214–219.
- Willshaw, D. J. and von der Malsburg, C. (1976). How patterned neural connections can be set up by self-organization. *Proceedings of the Royal Society London B*, 194.
- Yamins, D. L. K. and DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3):356–365.
- Yin, H. and Knowlton, B. (2006). The role of the basal ganglia in habit formation. *Nat Rev Neurosci*, 7(6):464–476.

Zipser, D. (1985). A computational model of hippocampal place fields. *Behavioral neuroscience*, 99(5):1006–1018.