SPATIAL LEARNING AND NAVIGATION IN NEURO-MIMETIC SYSTEMS

MODELING THE RAT HIPPOCAMPUS

THÈSE N° 2312 (2000)

PRÉSENTÉE AU DÉPARTEMENT D'INFORMATIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

ANGELO ARLEO

Laurea in Scienze dell'Informazione, Università degli Studi di Milano,
Facoltà di Scienze Matematiche, Fisiche e Naturali, Italia
de nationalité italienne

soumise à l'approbation du jury:

Prof. Roger Hersch, président
Prof. Wulfram Gerstner, directeur de thèse
Dr. Neil Burgess, corapporteur
Prof. Ariane Etienne, corapporteur
Prof. Jean-Arcady Meyer, corapporteur
Prof. Roland Siegwart, corapporteur
Dr. Paul Vershure, corapporteur

Lausanne, EPFL
2000

SPATIAL LEARNING AND NAVIGATION IN NEURO-MIMETIC SYSTEMS

MODELING THE RAT HIPPOCAMPUS

To my family
To Valérie

Abstract

Intelligent behavior in complex environments involves many basic capabilities. In particular, to achieve cognitive navigation, both biological organisms and artificial agents need to be endowed with spatial learning capacities, that is, they must be able to locate themselves within the environment, to perform complex target-directed behavior (e.g., to reach feeder locations), and to avoid penalties (e.g., collisions with obstacles).

Neurophysiological and behavioral experiments on rodents suggest that the *Hippocampus* is the structure of the brain that provides a neural basis for space coding. Indeed, hippocampal *place cells* exhibit an explicit representation of the animal's location within the environment. In addition, neurons encoding the orientation of the rat's head in the azimuthal plane have been recorded from the hippocampal formation, i.e., *head-direction cells*. Place coding and directional sense are crucial capabilities for solving spatial learning tasks. This endorses the hypothesis that the hippocampus plays a functional role in rodent navigation, and that it supports spatial cognition and spatial behavior.

The overall objective of this work is to study the neurophysiological mechanisms underlying rodent spatial behavior, and to emulate these processes to design an autonomous navigating agent. In particular, we focus on cognitive navigation and we develop a neuro-mimetic system modeling biological spatial learning. We address the two following questions:

(i) How do rodents establish a space code of their environment based on available sensory inputs and on their continuous interaction with the world? We put forward a modular neural architecture based on the functional properties as well as the anatomical interconnections of the brain regions involved in space representation. We model hippocampal place cells and head-direction cells as two neural populations. These two systems are strongly coupled and interact with each other to form a unitary space learning system. Allothetic (e.g., visual and auditory signals) and idiothetic inputs (e.g., vestibular and proprioceptive signals) are combined to establish stable place and direction codes. Unsupervised Hebbian learning is employed to allow the agent to build its worldview based on its own experience.

(ii) How can cognitive navigation be accomplished based on the above spatial knowledge? The hippocampal space representation model is used as a basis for achieving target-oriented behavior. Hippocampal place cells drive an extra-hippocampal population of action neurons. Synaptic efficacy between place cells and action cells is modified as a function of dopaminergic target-related reward signals. This results in an ensemble action cell activity that provides goal-directed navigation.

Biologically inspired solutions have been shown to be a useful methodology for developing flexible and self-contained artificial navigation systems. The experimental evaluation of the model has been done by implementing it on a mobile robot. We emphasize the importance of continuous interaction between the robot and the environment. This results in an incremental and dynamic development of the navigation system, and enables the robot to adapt its lifelong behavior according to situations that it has never experienced before.

The present study has produced some insights in the neurophysiological processes involved in animal spatial cognition (i.e., functional and anatomical predictions). From a robotic viewpoint, we have endowed an artificial agent with animal-like exploration and self-localization capabilities, which allows the robot to accomplish effective target-oriented navigation exploiting its interaction with the world.

Sommario

La navigazione di un agente autonomo in ambienti complessi richiede capacità cognitive. Sia gli organismi biologici che gli agenti artificiali devono essere dotati di capacità di apprendimento per mostrare un comportamento cognitivo. L'interazione agente-ambiente è una componente fondamentale dell'apprendimento: Essa permette lo sviluppo della conoscenza necessaria per adattare le proprie azioni a situazioni che non sono prevedibili.

La navigazione cognitiva si basa sulla capacità di auto-localizzazione da parte di un agente: quest'ultimo deve utilizzare le proprie risorse sensoriali per determinare la propria posizione nello spazio. L'informazione sensoriale può provenire da fonti esterne (e.g., stimuli visivi) oppure da fonti interne (e.g., stimuli vestibolari). Gli organismi biologici integrano questi due tipi di informazione per costruire una rappresentazione spaziale adeguata. I processi neuro-fisiologici che determinano tale capacità cognitiva non sono ancora stati completamente identificati.

Gli esperimenti neuro-fisiologici e comportamentali condotti sui topi da laboratorio rivelano che l'*ippocampo* è la struttura cerebrale sulla quale si basano le capacità di apprendimento spaziale degli animali. Infatti, le *cellule ippocampali di posizione* (i.e., *hippocampal place cells*) codificano esplicitamente la posizione di un animale nell'ambiente, mentre le *cellule ippocampali di direzione* (i.e., *head-direction cells*) ne rappresentano l'orientazione nel piano orizzontale. La capacità di codificare posizione e orientazione nello spazio bidimensionale è determinante per il processo di auto-localizzazione. Ciò conferma l'ipotesi che l'ippocampo ha una funzione importante nei processi di navigazione cognitiva degli animali.

Gli obiettivi principali di questa tesi sono: *(i)* Lo studio dei meccanismi neuro-fisiologici che governano l'attività delle cellule ippocampali; *(ii)* L'emulazione di tali processi per lo sviluppo di un agente artificiale autonomo. In particolare, lo scopo è di realizzare un sistema neuronale basato sulle proprietà funzionali e anatomiche delle strutture cerebrali coinvolte nel processo di navigazione cognitiva.

È stato sviluppato un modello ippocampale in cui due popolazioni neuronali (i.e., cellule di posizione e di direzione) interagiscono al fine di permettere l'apprendimento di rappresentazioni spaziali. Tale sistema integra le informazioni sensoriali esterne ed interne per ottenere un'adeguata stabilità a livello della rappresentazione. L'apprendimento non-supervisionato (i.e., *unsupervised Hebbian learning*) è utilizzato per permettere al sistema di acquisire conoscenza grazie alla propria esperienza.

La rappresentazione appresa è utilizzata per permettere la navigazione da un punto qualsiasi dello spazio verso una posizione obiettivo. Una popolazione extra-ippocampale di *cellule d'azione* (i.e., *action cells*) è utilizzata per determinare il comportamento del sistema. L'attività di tali cellule è determinata dalle cellule di posizione attraverso delle connessioni sinaptiche. L'entità di tali connessioni è modificata tramite apprendimento "per rinforzo" (i.e., *reinforcement learning*). Un segnale di tipo dopaminergico è utilizzato per modificare tali connessioni in funzione dell'obiettivo. Tale metodo permette di ottenere un'attività neuronale a livello delle cellule d'azione che codifica una mappa di navigazione.

Il modello è stato implementato su un robot mobile per effettuarne la validazione sperimentale. L'enfasi data all'interazione robot-ambiente permette lo sviluppo incrementale del sistema di navigazione. Inoltre, il robot acquisisce la capacità di generalizzazione necessaria per adattare il proprio comportamento a situazioni mai affrontate in precedenza.

Per concludere, tale lavoro ha permesso di chiarire alcuni meccanismi che sono alla base della navigazione cognitiva degli animali e di formulare delle predizioni di carattere funzionale e anatomico concernenti le strutture cerebrali coinvolte in tali processi.

Acknowledgments

I wish to express my deep gratitude to the many people who, in one way or another, have contributed to the achievement of this thesis[1]. First of all, I would like to thank my doctoral advisor Prof. Wulfram Gerstner: His effective and constant oversight helped me a lot throughout these three years of work. I would also like to acknowledge the members of the thesis committee: Prof. Roger Hersch, Dr. Neil Burgess, Prof. Ariane Etienne, Prof. Jean-Arcady Meyer, Prof. Roland Siegwart, and Dr. Paul Vershure.

I am particularly grateful to Prof. Dario Floreano for his support and his useful advices during the first year of my thesis. I would also like to thank the members of the Microprocessors and Interfaces Laboratory (LAMI), especially the Prof. Jean-Daniel Nicoud and the secretary Marie-Jo Pellaud, for having provided me with all the necessary facilities to work comfortably.

I am also in debt to the system administrators Yuri López de Meneses, Olivier Carmona, Stéphane Ecuyer, and Benjamin Barras, who managed to maintain all the computers working properly 24 hours a day. Similar to computers, physical robots may be a source of enormous time wasting if they are not maintained properly. For this reason, I wish to thank K-team, and in particular Francesco Mondada, for having provided me with a reliable robotic platform such as the Khepera miniature robot. I would also like to thank André Guignard and Alcherio Martinoli for having taken care of all the Khepera robots "living" at LAMI.

I have been really pleased to work as a member of the Centre for Neuro-mimetic Systems (MANTRA) directed by Prof. Wulfram Gerstner. A big thanks to all my colleagues for their friendship and support: Alix Herrmann (*mahalo nui loa!*), Thomas Strösslin, Mona Spiridon, Werner Kistler, Silvio Borer, Pierre-Edouard Sottas, and Edgar Oropesa Farras. A special thanks to Monique Dubois for her kindness and efficiency in taking care of all the administrative issues. I would also like to thank Fabrizio Smeraldi and Stéphane Hug for their crucial contribution to one part of my thesis.

During this work I had the opportunity to study something completely new to me: The neurobiology of the spatial behavior of animals. I am deeply indebted to Catherine Brandner for all the interesting discussions we had, which helped me a lot to understand the language as well as the scientific approach adopted by people from this domain.

I wish to address a BIG thanks to those persons who shared with me my "non-thesis time", making me really appreciate my stay in Lausanne: My friends Yann, Carine, Ivo, Kerstin, Thomas, Anne, Luca, Nathalie, Abram, Natalia, Guy, Claudio, and Laurent.

I would like to express my sincere gratitude to a person who has been determinant for my professional as well as my personal growth during the last six years: *Muchas gracias José!*

Dulcis in fundo, I wish to thank the most important people in my life, those persons who really make me happy every day. My family: *Grazie Giuseppe, Filomena, Luisa, Mauro e Aurora*, and Valérie: *Merci pour tout!*

[1]This Ph.D. project has been supported by the Swiss National Science Foundation, nr. 21-49174.96.

We realize that what we are accomplishing is a drop in the ocean.
But if that drop were not in the ocean, it would be missed.

Mother Teresa.

Contents

Part I

Introduction

Chapter 1

Introduction

1.1 Motivation

For successful spatial behavior, both animals and autonomous artifacts need to continuously interact with their environment. An agent, be it a robot or an animal, senses its world via different perceptual capabilities (e.g., vision, olfaction, touch), and has to employ such multimodal information to adapt its behavior. Reaching an interesting location (e.g., a food source), returning home, finding short-cuts, adopting efficient exploration strategies: In order to perform these tasks autonomously and effectively, a navigating agent needs to be endowed with spatial learning capabilities. In particular, it has to able to locate itself in the environment as well as to perform target-oriented behavior.

Cognitive neuroscience defines navigation as the capability of planning and performing a path from a current position $\vec{p}(t)$ towards a desired location $\vec{p}_{goal}$ [114, 97]. Implicit in this definition is the ability of a navigating agent to adapt its goal-directed behavior to the complexity of the task to be solved.

Approaching a target location $\vec{p}_{goal}$ is easy if the latter is either directly visible or identified by a visible cue. In this case, a simple landmark-guidance behavior can be adopted to solve the task, i.e., *taxon navigation*. Taxon strategies can be understood by means of simple stimulus-response behavior, that is, the agent must associate a single response to a single stimulus [309]. This yields a simple reactive strategy: Given a stimulus (e.g., a spotlight), the agent orients towards it and approaches it.

A second type of navigation based on stimulus-response learning is *praxic navigation*. The agent navigates towards a goal by executing a specific motor sequence acquired by training [272]. This strategy is appropriate, for instance, when the target is identified by a sequence of specific cues. Then, instead of single orienting responses as in taxon navigation, the agent can learn more elaborate sequences of stimulus-action associations.

Target locations are often either not directly identified by any cue (or sequences of cues), or simply "hidden" with respect to the agent's sensory capabilities. In this case, goal-oriented behavior requires more complex spatial learning, i.e., *locale navigation* [256]: A representation of the spatio-temporal properties of the environment is built to enable the agent to locate itself and navigate based on locally-available sensory information. This type of representation is also referred to as *cognitive map* [361], in the sense that it requires high-level information processing to encode the knowledge about a familiar environment, and to use such knowledge for flexible behavior (e.g., planning short-cuts).

Do animals have access to internal models of their environments? Do animals use these models

for cognitive-like navigation? Neurophysiological as well as behavioral findings support the hypothesis that they actually do. Indeed, as the complexity and the perceptual capabilities of biological organisms increase, an explicit spatial representation appears to be employed to support navigation in complex task-environment contexts [361, 255, 256].

The most studied neurophysiological structure of mammals exhibiting such a spatial representation property is the *hippocampus* [255, 256, 209, 243, 324, 278, 394, 323, 166]. This area of the brain has been thought to mediate spatial coding ever since experimental evidence for location-sensitive neurons (*place cells*) in the hippocampus of freely-moving rats was found [255]. In addition, neurons whose activity is tuned to the orientation of the rat's head in the azimuthal plane have also been observed in the hippocampal formation. These neurons (*head-direction cells*) endow rodents with an allocentric compass system. Place coding and directional sense are crucial for solving complex spatial learning tasks. This supports the hypothesis that the hippocampus plays a functional role in rodent navigation, and provides a neural basis for spatial cognition and spatial behavior.

1.2 Towards Neuro-Mimetic Spatial Learning

Within the multidisciplinary research field concerning the design of artificial intelligent agents, the task of developing autonomous navigating systems is still an open problem [74, 87]. The complexity of this task lies inherently in the concept of *autonomy*: An autonomous agent should have a completely self-contained control system to adapt its lifelong behavior to all possible situations it might face.

The traditional Artificial Intelligence approach, based on predefining internal accurate models of the world to endow robots with human-like symbolic capabilities, has been recognized unsuitable for such a task [53]. Firstly, real complex environments are often unpredictable, making it impossible to design a built-in knowledge to associate the most appropriate action for every situation. Additionally, the agent's sensory-motor system and the interaction agent-environment are typically corrupted by noise whose distribution is often unknown. Finally, predefined models are intrinsically biased since they tend to reflect the worldview determined by the human sensory system and to inherit the structure of linguistic descriptions used to formulate them [88].

Recent research on autonomous robots has moved towards a novel non-symbolic approach termed *behavior-based robotics* [54, 190, 89]. Capabilities such as reactivity, flexibility and real-time acting have received considerably more attention than optimality and completeness. The idea is to let the agent build up its own worldview and behavior by means of its own experience (i.e., *learning*). Based on the idea that "the world is its own best model" [53], the behavior-based approach stresses the importance of continuous interaction between the robot and its environment for an incremental and dynamic development of the control system [190]. The principles for designing behavior-based robots often take inspiration from basic behavior mechanisms observed in biological systems and from well-established neurophysiological adaptive strategies such as neural plasticity [27, 73, 204, 268, 52, 269, 338]. This research direction has been shown to be a promising way to realize autonomous navigating systems. Several behavior-based learning frameworks, such as *reinforcement learning* and *evolutionary techniques*, have successfully addressed the problem of designing adaptive systems to autonomously navigate in unpredictable real environments [218, 217, 72, 115, 176, 191, 104, 86, 248]. Most of the proposed solutions, are based on *reactive behavior*: Agents learn to map incoming stimuli to actions in order to perform the best control policy to accomplish their task, without building any internal spatial model of the environment. As a consequence, the behavior-based paradigm can be employed to capture the functions undertaken by the taxon system in biological agents, although it does not allow us to scale up to more complex tasks such as cognitive (i.e., locale) navigation [186].

The issue of building an internal model of the world to be used for autonomous navigation has been

largely investigated in robotics. Map-learning research has produced two principal approaches, namely the *metric paradigm* [232, 231, 94, 95] and the *topological* paradigm [178, 197, 407]. In the former, the geometrical features of the world are modeled accurately. One of the most popular of such methods consists of representing space by means of a two-dimensional evenly-spaced grid called *occupancy grid* [231, 232]. Each grid cell estimates the occupancy probability of the corresponding area of the world. Topological maps are more compact representations in which spatial relations between relevant locations in the environment are modeled by means of a topological graph [178, 197, 308].

Since occupancy grids reproduce the geometrical structure of the environment explicitly, they are easy to learn and maintain. On the other hand, this approach is limited by its vulnerability to errors that affect the metric information (i.e., robot's position and distance to obstacles). In addition, building occupancy grids is very expensive in terms of memory and time: to accurately model a complex environment, the resolution of the occupancy grid must be high, which results in heavy computation. Topological maps are more qualitative representations of the world, so that they are not necessarily vulnerable to errors in the metric information. Also, the complexity of the learned graph is directly related to the world complexity optimizing the use of time and space resources. However, since pure topological maps rely on a sensory pattern recognition process, distinct places producing equivalent sensory patterns might disrupt the self-localization system.

Since these two paradigms exhibit complementary strengths and weaknesses [359], several models have been put forth to integrate both representations into a hybrid system [68, 95, 356, 359, 12]. However, systems engineered so far are not as robust, flexible, and adaptable as biological space-learning systems. Therefore, similar to the behavior-based idea of modeling biological solutions, moving towards a biologically-founded approach to achieve cognitive navigation offers the attractive prospect of developing autonomous systems that directly emulate mammalian navigation abilities.

1.3 Aims of this Thesis

The overall objective of this work is to study the neurophysiological mechanisms underlying rodent spatial behavior, and to emulate these mechanisms to design an autonomous navigating agent. In particular, we will focus on cognitive (i.e., locale) navigation and we will develop a neuro-mimetic system suitable for modeling the components of biological spatial learning. We address the two following questions:

(i) How do rodents establish a space code of their environment based on available sensory inputs and on their continuous interaction with the world? We will model the properties of hippocampal place cells and head-direction cells, and we will realize a neural architecture based upon the anatomical interrelations between the brain areas involved in space learning.

(ii) How can cognitive navigation be accomplished based on the above spatial knowledge? We will develop a navigation mechanism to achieve goal-directed behavior by employing the above place and direction representations for reward-based action learning.

The experimental evaluation of the developed model will be done by implementing the system on a mobile robotic platform (Sec. 1.4). We stress the importance of continuous interaction between the robot and its environment (situated artificial intelligence paradigm [371, 270]) to produce an incremental experience-dependent development of the spatial learning system.

An understanding of the functional role of spatial representations in autonomous agents is of central importance to the field of both neurobiology and robotics. Research in this field may contribute to new developments and cross-disciplinary insights:

- Exploration of bio-inspired models of spatial learning may lead to a better understanding of the mechanisms supporting intelligent navigation capabilities of animals. Neuro-mimetic robotics offers a useful tool to validate new hypotheses concerning functionalities in neurophysiological processes. The fact that artifacts are simpler and more "experimentally transparent" than biological systems (e.g., a rat), makes them appealing for understanding the nature of the underlying mechanisms of animal behavior.

- Conversely, modeling biological solutions may provide efficient and robust navigation strategies and then lead to an immediate applicational pay-off in designing more powerful and adaptive artificial autonomous agents. Modeling the fact that biological systems can acquire their internal models incrementally and on-line according to the requirements of the given task-environment context, may provide interesting solutions to current robotics.

1.4 Methods

1.4.1 Neural Networks

The system we use to endow the agent with space learning and navigation capabilities relies on the artificial neural network paradigm[1] [146, 144, 36]. A neural network is a massively parallel distributed system suitable for storing and processing complex information, and making it available for use [144].

The elementary constituents of a biological neural network are simple computing units referred to as neurons. Each neuron i receives a large number of input connections termed *dendrites*, and transmits its response through the *axon* output connection. The computational power of a neural network derives from the massive interconnections between its neurons. *Synapses* are the elementary components that mediate the interaction between neurons. The most common type of synapse converts a presynaptic electrical signal into a chemical signal and then back into a postsynaptic electrical signal [325]. In particular, a presynaptic process liberates a neural *transmitter* that diffuses across the synaptic junction and acts on the postsynaptic dendrite.

An artificial neural network is a machinery designed to model the adaptive way in which the brain processes information to perform a given task. In particular, a neural network resembles the brain in two aspects [144]: *(i)* It acquires the necessary knowledge to solve a problem through a *non-symbolic learning process*; *(ii)* Storing this knowledge occurs through the modification of the weights of the interconnections between neurons (i.e., *synaptic plasticity*).

Therefore, artificial neural networks offer a suitable tool for designing experience-dependent navigation systems, and allow us to model the functional properties as well as the anatomical interconnections of the brain regions involved in animals' spatial learning.

1.4.1.1 Rate Coding Model

We employ a highly simplified neuronal model in which the firing activity r_i of each neuron i is measured by temporally averaging the number of spikes emitted by i during a time window Δt [2, 146]. The concept of rate coding has been largely employed in artificial neural systems. A signal I forms the input of a neuron i and yields an output $r_i = f(I)$ representing the mean firing rate of i. The function f is termed *transfer function*. An example of commonly used transfer function is the sigmoid, such that the activity r_i tends asymptotically to 1 for large input signals [146].

[1] Artificial neural networks are also referred to as connectionist networks, neurocomputers, and parallel distributed processors.

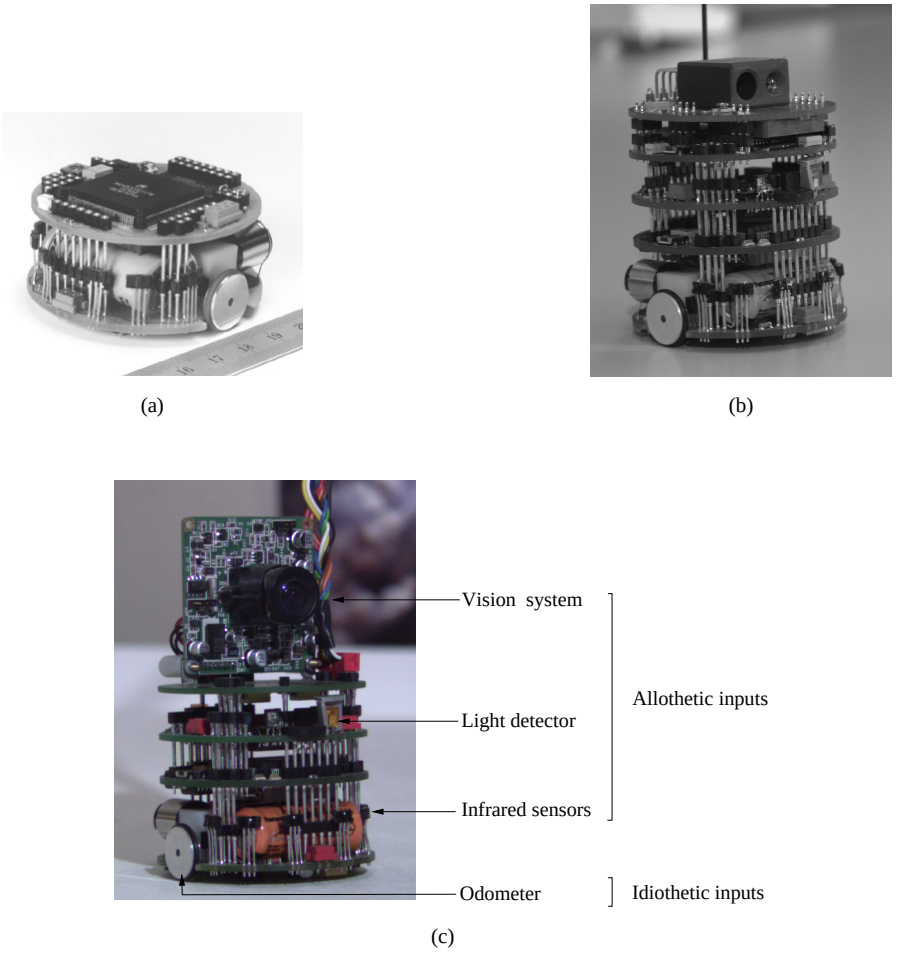

Figure 1.1: (a) The mobile Khepera miniature robot (basic configuration). (b) One of the configurations used in this thesis: the robot is equipped with a linear camera on the top. Images consist of 64-pixel one-dimensional arrays. The azimuthal visual range is of approximately 36°. (c) The second configuration employed: the linear vision turret is replaced by a black and white CCD video camera. The two-dimensional view field covers about 90° in the horizontal plane and 60° in the vertical plane. The image resolution is 768×576 pixels. In this configuration, the Khepera's sensory system involves *(i)* allothetic information, i.e., infrared readings, light detection, vision, and *(ii)* idiothetic information, i.e., dead-reckoning.

The main limitation of this neuronal model is that it cannot capture the fast processing speed exhibited by real neurons (e.g., cells in the human visual cortex [358]). That is, reaction times in neural cells are often too short to tolerate temporal averages. Nevertheless, the rate code model permits simple analytical calculations which results in acceptable computation times even in the case of large neural populations. This is suitable for the type of neural processing adopted in this work.

1.4.2 Experimental Setup

1.4.2.1 Agent: The Khepera Robot

In order to validate the spatial learning system developed in this thesis, we employ a physical implementation on a mobile Khepera miniature robot[2].

[2]The Khepera has been designed and developed at the Microprocessors and Interfaces Laboratory of the Swiss Federal Institute of Technology Lausanne (EPFL) in 1993 [230]. Nowadays, it is a largely diffused commercial platform produced and distributed by K-Team S.A. (www.k-team.com).

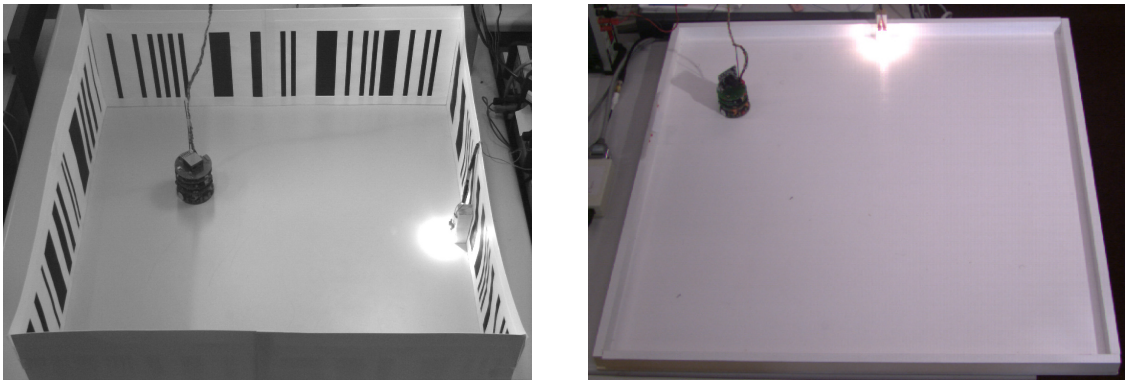

(a) (b)

Figure 1.2: (a) The $60 \times 60\,cm$ arena with the Khepera robot inside. Black and white stripes provides the input for the linear visual system. (b) The $80 \times 80\,cm$ arena used for the robot equipped with the CCD camera. Visual stimuli are provided from standard laboratory background. In both cases a camera above the environment is used to monitor the robot's behavior.

The Khepera robot has a cylindrical shape with a diameter of $55\,mm$, and, in the basic configuration, it is $36\,mm$ tall (Fig. 1.1 (a)). Two DC motors drive two wheels independently providing the robot with non-holonomic motion capabilities. The Khepera CPU board consists of a Motorola MC68331 processor with $128\,Kb$ of EEPROM and $256\,Kb$ of RAM. An A/D converter supports the acquisition of analog sensory signals (e.g., infrared readings). A standard RS232 serial-port provides the possibility of connecting the robot to an external computer in order to elaborate and process sensory signals[3].

The Khepera has an open modular architecture: additional turrets can be superimposed to extend its basic configuration. Figs. 1.1 (b, c) show the two types of configurations used in this work. The sensory capabilities of the agent are provided by: *(i)* Eight infrared sensors to detect obstacles[4] (i.e., tactile-like perception) and measure ambient light; six of the infrared sensors cover the frontal 180° of the robot, while the remaining two sensors cover approximately 100° on the back side; *(ii)* A light detector placed in the front of the robot (additional turret); *(iii)* A vision system (either a linear camera, Fig. 1.1 (b), or a two-dimensional CCD camera, Fig. 1.1 (c)); *(iv)* An odometer to map wheel turns into linear displacement (i.e., dead-reckoning). As shown in Fig. 1.1 (c), signals provided by infrared sensors, light detector, and visual system, form the allothetic sensory information perceived by the robot through its interaction with the external world. On the other hand, the internal movement-related signal provided by the odometer constitutes the idiothetic sensory input.

Due to its physical and functional characteristics, Khepera offers a suitable platform for designing experiments similar to those performed by ethologists and neurophysiologists on rodents. Indeed, its small size allows us to work with experimental setups similar to those used with rats. Furthermore, its sensory system (i.e., vision, light detection, tactile-like perception, and odometry) is appropriate to emulate the perceptual capabilities of rats.

In this thesis we will develop a "high-level" controller to enable the robot to behave based on its own experience (i.e., learning). However, we also endow the Khepera with a low-level hard-coded reactive module providing obstacle avoidance. This behavior is triggered whenever the proximity sensors detect an obstacle, and takes the control (over the high-level system) to always keep the robot in a safe state.

[3]The robot can also be used in a totally stand-alone configuration by using its own processor to treat data. However, since we have to interpret two-dimensional video streams provided by a CCD camera module, we need an external frame grabber.

[4]In average, Khepera's infrared sensors can detect an object at a maximum distance of about $4 - 5\,cm$.

1.4.3 Experimental Apparatus

Figs. 1.2 (a, b) show the two experimental arenas used for experiments performed in this thesis. Both arenas are open-field environments in which the agent can freely move and sense its world. The $60 \times 60\,cm$ arena of Fig. 1.2 (a) is employed for experiments with the linear-vision system. Walls are covered by *randomly distributed* black and white stripes of variable width providing the visual patterns to the system. The $80 \times 80\,cm$ arena of Fig. 1.2 (b) is used for the two-dimensional vision system. The arena is placed within a standard laboratory background. Small walls prevent the robot from running outside the environment. In both cases, a light source is located on one of the arena walls. In addition, the robot's behavior is monitored by means of a video camera above the experimental environment.

1.5 Roadmap of this Dissertation

This thesis is organized in three main parts: *(i)* Spatial orientation: Head-direction cells; *(ii)* Space representation: Hippocampal place cells; *(iii)* Spatial behavior: Goal-oriented navigation. Each part provides a self-contained description of the theoretical and experimental concepts related to its main topic. In particular, all parts have the same following structure:

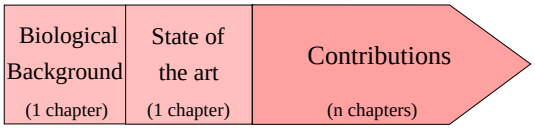

The first chapter of each part gives the background for understanding the kind of neurophysiological mechanisms to be modeled. The second chapter reviews previous research work. The remaining n chapters contain the modeling and experimental work done in this thesis.

Spatial Orientation: Head-direction Cells (Part II)

Chapter 2 describes the neurophysiological properties of direction-sensitive neurons (i.e., *head-direction cells*) as well as the anatomical interconnections between the areas of the brain in which they have been observed.

Chapter 3 describes several models that have been put forth to capture the direction-selectivity property of head-direction neurons.

Chapter 4 is the first in which we present our computational approach for modeling the rodent directional sense. We also give experimental results obtained by validating the model with our robotic implementation.

Chapter 5 addresses the interrelation between external (i.e., allothetic) and internal (i.e., idiothetic) signals and their relative influence upon head-direction cells. In particular, we perform some experiments similar to those done with rodents to investigate in which conditions one type of information (either external or internal) becomes predominant.

Space Representation: Hippocampal Place Cells (Part III)

Chapter 6 describes the hippocampal involvement in spatial learning. The properties of location-sensitive neurons (i.e., *place cells*) are described and related experimental findings are reviewed. Several brain regions involved in spatial learning are briefly introduced, and a general view of the principal neural pathways connecting the hippocampal formation to the rest of the brain is given. Also,

anatomical lesion data are reviewed which give some hints about the functional role of the involved brain areas.

Chapter 7 provides a rather extensive review of previous models of animals' space representation capabilities. We do not only present approaches focusing on hippocampal place cells, but also models focusing on the animals' capability of navigating based on self-motion information only (i.e., path integration). Furthermore, we give some results obtained by implementing a specific approach to space coding, namely the model by Burgess, Recce, and O'Keefe [62].

Chapter 8 presents our contribution in interpreting visual data in order to achieve self-localization. We describe the allothetic pathway of our hippocampal model. Starting from real images taken by the robot during exploration, we establish a population of visually-driven place cells. Experimental results obtained with both the linear and the two-dimensional vision setup are presented.

Chapter 9 describes our model of path integration (i.e., the idiothetic pathway of the system). Results are given to emphasize the intrinsic limitation of path integration, i.e., its cumulative self-localization error over time.

Chapter 10 combines the allothetic and idiothetic representations to establish a population of place cells in the hippocampus. Reported results demonstrate that combining external and internal signals yields stability in the system. As for head-direction cells, we study the interaction between extrinsic and intrinsic cues to control hippocampal place cells.

Spatial Behavior: Goal-oriented Navigation (Part IV)

Chapter 11 discusses the goal-directed navigation capabilities of animals and reviews experimental data on rodents. Also, it describes the role of dopaminergic neurons in reward-based learning.

Chapter 12 describes previous models of hippocampal-based navigation. It also presents results obtained by implementing two of the reviewed approaches: the model by Abbott, Blum, and Gerstner [1, 46, 122], and the navigation part of the model by Burgess, Recce, and O'Keefe (the space-code part has been implemented in Chapter 7).

Chapter 13 proposes our action learning scheme to allow the robot to acquire navigational maps based on the hippocampal representation established in Part III. Goal-oriented navigation is achieved by means of reward-based learning.

This dissertation terminates with Chapter 14 in which the achievements and the limitations of this thesis are discussed, and where our future directions of research are described.

Part II

Spatial Orientation: Head-direction Cells

Foreword

Spatial orientation is crucial for both biological and artificial systems involved in navigation. We define *self-orienting agents* those navigating systems capable to establish their own orientation in the azimuthal plane autonomously and on-line.

Experimental findings suggest that animals and humans are self-orienting agents: During navigation they can autonomously monitor their directional heading with respect to an allocentric spatial reference. Neurophysiological data show the existence of an *internal directional sense* as a basis for spatial behavior [210, 97, 353]. In particular, experiments show that ants [382, 77, 238], bees [35], spiders [131, 316], rodents [224, 223, 273, 100, 314, 3], dogs [21, 30, 313], and humans [30, 304, 111, 102] rely on an internally maintained directional reference when homing by *path integration*. The latter is a route-based mechanism in which self-motion signals are integrated over time in order to infer the location and the heading of a navigating agent relative to a departure point [223, 201, 98]. This implies estimating the rotational as well as the translational component of motion continuously [314].

In order to assess their rotational displacements over time (i.e., *directional tracking*), animals like desert ants and bees estimate their current azimuthal heading by using the sun as an absolute directional reference [380, 383, 78]. By contrast, mammals mostly rely on idiothetic inertial cues in order to measure rotations [98, 102]: They combine and integrate vestibular stimuli, proprioceptive information, and efference copy of motor commands over time [42, 210, 102]. However, measuring angular displacements based on inertial self-motion information only, induces a directional sense affected by cumulative drift [198, 210, 102]. For this reason, when available, allothetic information (e.g., visual cues) has to be incorporated in order to recalibrate the heading representation [354, 210, 166].

In this part, we study the neurophysiological structures involved in self-orientation, and we address the issue of modeling their basic functionalities:

- In Chapter 2 we provide the biological background: We describe the properties and the functional role of direction-sensitive neurons termed *head-direction cells*.
- In Chapter 3 we review some existing approaches to model the biological directional system.

Contributions (the work presented in this part of the thesis has been published in [15, 17]):

- In Chapter 4 we put forth a new computational approach for the rodent directional sense, and we give experimental results obtained by validating it on a real mobile Khepera robot.
- In Chapter 5 we study the influence of external signals (e.g., visual cues) upon our directional system. We perform experiments (with the mobile robot) similar to those employed by ethologists and neurophysiologists (with rodents) to study the control of extrinsic cues upon head-direction cells.

Chapter 2

Biological Head Direction Cells

Experimental recordings from freely-moving rodents show the presence of *head-direction cells*, i.e., neurons which fire as a function of the animal's heading in the azimuthal plane [279, 353, 347, 388, 229, 71, 41]. Each head-direction cell fires maximally only when the animal's head is oriented in a specific direction, regardless of the orientation of the head relative to the body, of the animal's behavior, and of the animal's spatial location. Thus, the ensemble activity of head-direction cells works as an allocentric compass providing a neural basis for the rodent's directional sense [210].

2.1 Anatomical Background

Direction-sensitive neurons have been recorded from the following brain regions: *(i)* The deep layers of the postsubicular cortex[1] (poSC) [279, 353], *(ii)* the anterodorsal nucleus (ADN) of the anterior thalamus (AT) [347, 348, 41, 166], *(iii)* the lateral mammillary nuclei (LMN) [181], *(iv)* the laterodorsal nucleus (LDN) of the thalamus [229, 228], *(v)* the posterior parietal (PP) and the posterior cingulate[2] (PC) cortices [69, 71, 70, 213], *(vi)* and the dorsal striatum (caudate nucleus) [397, 226].

Although not completely understood yet, some of the anatomical mutual connections between the above brain areas (Fig. 2.1 (a)) have been characterized [353, 282]. The postsubiculum (poSC), the anterodorsal nucleus (ADN), and the lateral mammillary nuclei (LMN) are all interconnected [138, 29]. The laterodorsal nucleus (LDN) and the postsubiculum (poSC) are also interconnected [138]. The posterior cingulate cortex (PC) is interconnected with both ADN and poSC [138, 29]. Finally, the posterior parietal cortex (PP) is also interconnected with poSC, projects onto posterior cingulate cortex, and receives afferent inputs from LDN, and from the visual, sensory-motor, and somato-sensory cortices [138, 174, 60].

[1] The postsubiculum (poSC), also termed dorsal presubiculum, is a six-layered cortical area (like other cortical regions in the retrohippocampal formation). Almost all observed head-direction cells were located in the deep layers IV-VI [353].

[2] The posterior cingulate cortex is also known as retrosplenial cortex.

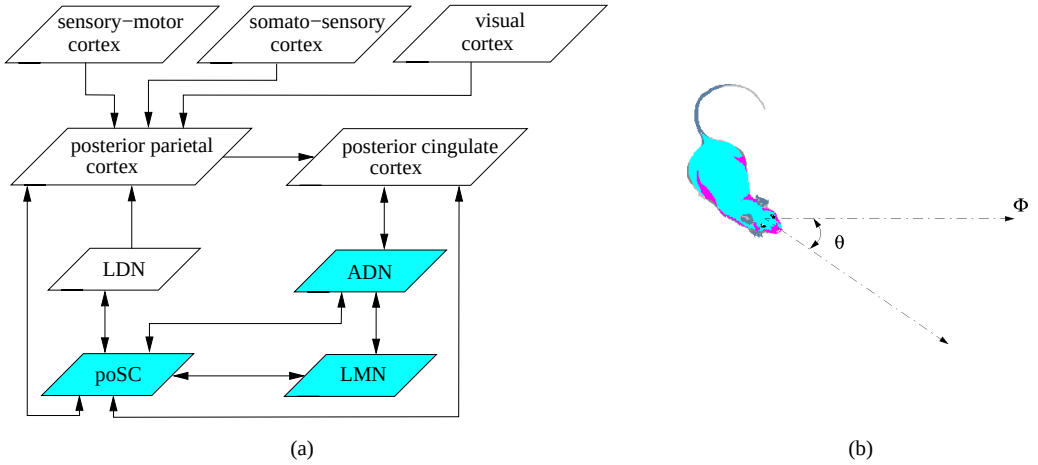

Figure 2.1: (a) A simplified overview of the brain areas containing head-direction cells and their anatomical interconnections. Glossary: poSC: postsubiculum, ADN: anterodorsal nucleus, LMN: lateral mammillary nuclei, LDN: laterodorsal nucleus. (b) The ensemble head-direction activity codes for the allocentric animal's heading θ with respect to a reference direction Φ.

2.2 Neurophysiological Properties of Directional Cells

Let i be a head-direction cell. Its spiking frequency $r_i(\theta)$ varies as a function of the angle θ between the midline of the animal's head and an arbitrary reference direction Φ (Fig. 2.1 (b)) [353]. Cell i has its own *unique* preferred direction θ_i, that is $r_i(\theta)$ is maximal only when the animal's heading θ is equal to θ_i. The mean spike-frequency distribution of a head-direction cell i is approximately a Gaussian curve[3] centered in the preferred direction θ_i and with an average standard deviation σ ranging from 20° to 60° [353, 210] (Fig. 2.2 (a)). The amplitude of the peak is characteristic for each cell [353]. Importantly, the preferred direction of each cell is invariant with respect to the animal's location, that is, for a given cell i, θ_i remains constant over the whole environment. As a consequence, since the set of all preferred directions $\Theta := \{\, \theta_i \mid \forall i \,\}$ covers uniformly 360°, the ensemble cell activity acts as an internal neural compass [166].

Data show that directional cells are not topographically organized [353, 348, 318, 405]. That is, two distinct head-direction cells i and j with similar preferred directions θ_i and θ_j are not necessarily neighboring units of the network.

A key property of directional cells is the stability and consistence of the representation [354, 130, 350, 282]: The difference between preferred directions of distinct directional cells remains constant, despite cue manipulation, disorientation of the animal, or exposure to different environments [166, 327, 405]. Thus, if cell i and j have preferred headings θ_i and θ_j, respectively, then any event inducing a shift in the firing direction θ_i, $\Delta\theta_i$, will produce an equivalent shift in the preferred direction θ_j, $\Delta\theta_j$, such that the difference $\theta_i - \theta_j$ will remain constant in magnitude and sign [210]. This rotational coupling endows the system with high stability.

Although head-direction cells are primarily correlated with the animal's spatial orientation, the head angular velocity ω plays an important role in modulating their firing activity [213]. Experiments show that most of the cells in the lateral mammillary nuclei (LMN) are strongly correlated with both head direction and angular velocity [181]. Thus, the mean firing activity of a LMN cell i is a function of θ as well as ω, that is $r_i(\theta, \omega)$. Similar cells have also been observed in the posterior parietal cortex (PP), the somato-sensory cortex, and the visual cortex [210, 376]. In the postsubicular cortex (poSC), cells whose activity depends on the *sign* of the angular head velocity have been found [353, 318], but only in a

[3] Or alternatively, a triangular curve [353].

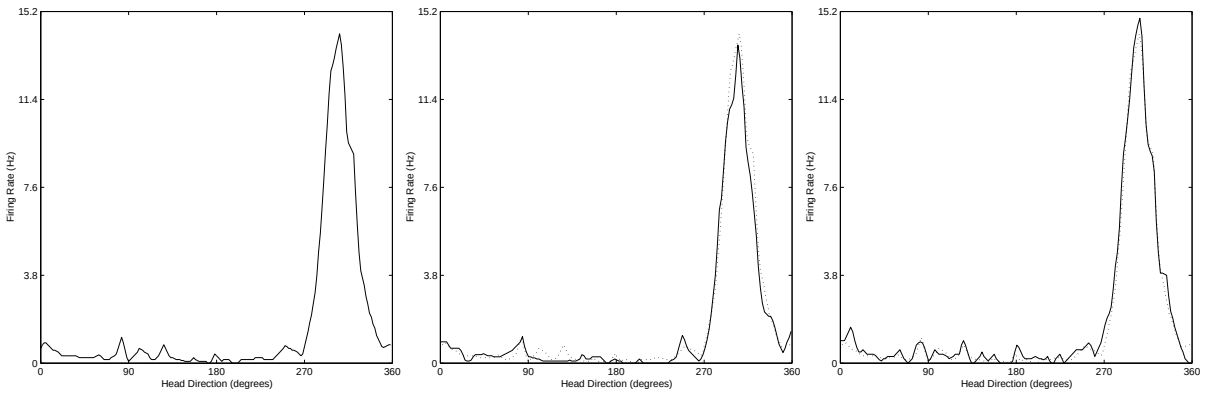

(a) (b) (c)

Figure 2.2: (a) Tuning curve of a postsubicular cell recorded by Blair and Sharp (1995) [41]. (b) The poSC cell exhibits similar preferred directions during both head still (dotted line) and counterclockwise head rotations (solid line). (c) The same holds for clockwise turnings (solid line) *vs* no rotations. Data courtesy T. Blair and P. Sharp.

very small percentage [282]. Finally, ADN head-direction cells tend to increase their average firing rate proportionally to the magnitude of the head angular velocity, whereas poSC cells, typically, do not [41].

Interestingly, cells in the anterodorsal nucleus (ADN) tend to anticipate the animal's "future" direction during head turning [41, 355][4]. In particular, ADN cells shift their preferred direction as a function of ω such that their activity may temporally anticipate head directions by a time delay τ. Let i be an ADN cell, and let θ denote the present animal's heading. The fact that cell i shifts its preferred direction θ_i proportionally to ω can be expressed by the following equation [41]

$$\theta_i(\omega) = \theta - \omega\tau_i \tag{2.1}$$

where $\tau_i \geq 0$ is the time delay of cell i. Velocity ω is assumed to be positive for clockwise turning and negative for counterclockwise turning. According to Eq. 2.1, cell i encodes the present orientation θ whenever the head is not turning (Fig. 2.3 (a)). During counterclockwise turns, $\omega < 0$, the relation $\theta_i \geq \theta$ holds, and cell i predicts θ by shifting its preferred direction to the right (Fig. 2.3 (b)). By contrast, during clockwise turns, $\omega > 0$, cell i tends to fire maximally at $\theta_i \leq \theta$, that is it tends to shift its preferred direction to the left (Fig. 2.3 (c)).

The physical constant τ_i defines the anticipatory property of cell i. Earlier experimental data suggested for τ_i a value of approximately $30\,ms$ constant over the entire ADN population [41]. However, recent findings show that ADN cells have "individual" specific time delays ranging between $0\,ms$ and $100\,ms$ [39]. In addition, recent results show that ADN cells have double-peak firing curves [40], where the deviation between the two peaks is proportional to the constant τ_i of the cell. These double-peaked ADN curves deform during rotation such that the "anticipating" peak (e.g., the left peak if $\omega > 0$) rises while the other decreases [282].

To summarize, postsubicular cortex (poSC), anterodorsal nucleus (ADN), and lateral mammillary nuclei (LMN) may be considered as primary structures with respect to the directional selectivity property. By contrast, cells observed in the posterior parietal (PP) cortex, in the posterior cingulate (PC) cortex, and in the laterodorsal nucleus (LDN) exhibit less clear first-order correlation to head direction [282]. When recorded in a radial maze [71], PP and PC cells show correlation to both directional and behavioral information [213]. For instance, their activity depends on whether the animal is moving inward or

[4]Recent experimental findings [38, 336] show that also cells in the lateral mammillary nuclei (LMN) exhibit an anticipatory directional coding.

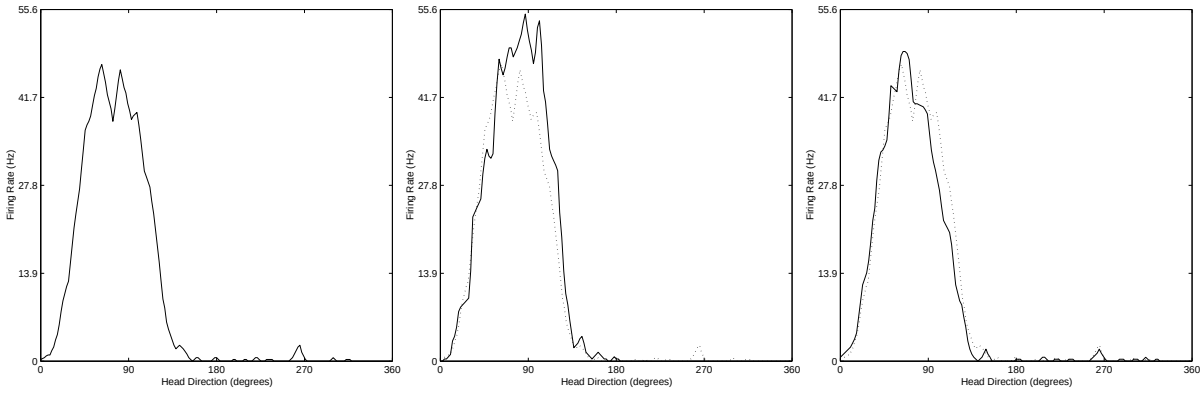

(a) (b) (c)

Figure 2.3: (a) A sample of mean firing rate of an ADN cell recorded by Blair and Sharp (1995) [41] during no head rotation. (b) The tuning curve of the same ADN cell tends to shift to the right (solid line) during counterclockwise turns. (c) The cell shifts its preferred firing direction to the left (solid line) during clockwise turnings. Data courtesy T. Blair and P. Sharp.

outward the maze, and whether the animal is turning or moving straight [71]. Furthermore, PP and PC head-direction cells may show multiple peaks during cue-manipulation experiments [69, 70, 282]. LDN cells show a strong dependence on allothetic cues in order to become tuned to directional cues in a radial maze [228]. Indeed, directionality in LDN cell activity is established only after an exposure period (about one minute) of the animal to the environment in light conditions. If light is subsequently switched off, orientation selectivity persists for approximately two minutes only [228, 282].

2.3 Anatomical Lesion Data

A large body of experimental work has been done to study how anatomical lesions impair the directional selectivity of head-direction cells. The following is a brief review of results:

Lesions to the postsubicular cortex. Data show that after poSC has been damaged, ADN thalamic cells are still directional selective [129, 351]. However, ADN cells become insensitive to allothetic cues and depend on vestibular input only [282].

Lesions to the anterodorsal nucleus. The directional selectivity of poSC cells is seriously impaired when damaging ADN [128].

Vestibular lesions. The activity of ADN thalamic head-direction cells is severely disrupted after vestibular lesions [335].

Lesions to the laterodorsal nucleus. Damages to LDN do not disrupt head-direction cells in poSC [126].

These results suggest that thalamic directional cells (ADN) are related primarily to vestibular inputs, and indirectly to allothetic signals via the postsubiculum (poSC). Indeed, after poSC lesions they are still tuned to the animal's head direction, but they do not respond to external stimuli. On the other hand, postsubicular cells seem indirectly related to inertial signals via the ADN cell activity, while they might be more involved in incorporating allothetic information coming from afferent projections from the parietal cortex [405]. As far as we know, there are no published results concerning lesions to the lateral mammillary nuclei, and to parietal and retrosplenial cortices.

2.4 Allothetic *vs* Idiothetic Cues

Head-direction cells continuously update their ensemble activity to track the animal's heading. How is this update mechanism carried out? What are the sensory inputs involved? Several experiments have been done to establish which cues primarily control biological directional cells [354, 210, 166, 350, 130, 102].

The preliminary hypothesis concerning the influence of the Earth's geomagnetic field upon directional firing activity has been deserted at once. Neurophysiological data clearly show that cells in poSC, ADN, and LMN can be controlled by salient visual cues *within* the experimental environment (e.g., a white card on the wall of the arena) [210, 229, 354, 166, 102]. For instance, rotating a dominant visual cue by an angle δ induces an approximately equivalent shift δ in the preferred direction of all cells i, that is $\Theta' = \{ \theta_i + \delta \mid \forall i \}$ [279, 354, 130, 348, 181].

Despite their dependence upon visual stimuli, poSC, ADN, and LMN cells exhibit clean directional selective properties even in the absence of allothetic polarizing cues [351]. In particular: *(i)* If head-direction cells have been established in light conditions, they maintain stable directional curves either after removing visual cues or in the dark [354, 210, 229]; *(ii)* If an animal enters an environment in complete darkness, head-direction cells are carried forward from the previous environment, and they continue to show normal directional firing [282].

According to this evidence, it has been hypothesized that idiothetic inertial cues (e.g., vestibular and proprioceptive signals) might play a major role in controlling head-direction cells. Thus, directional activity would be primarily given by integrating angular displacements of the head over time [210]. The existence of neurons (in the posterior parietal, somato-sensory, and visual cortices) whose activity is correlated to the head's angular velocity supports this theory [210, 376]. However, self-motion information is vulnerable to cumulative drift over time [198, 210, 102]. Thus, during exploration, long lasting associations might be learned to form a stable mapping between allothetic and idiothetic representations. Afterwards, this learned mapping might enable the animal to use stable visual cues to calibrate inertial information [210, 166].

Experiments with disoriented rats have been carried out to study how the animal's experience influences the strength of the control exerted by visual cues on head-direction cells [350]. These studies show that the more a visual cue is considered as stable by the animal, the stronger the control exerted by the cue on directional cell activity [166].

To summarize, although the actual mechanism underlying biological head-direction cells is still unclear, a general consensus is now emerging which identifies egomotion information as a primary input. Nevertheless, animals seem to rely on external reliable cues in order to maintain their head-direction system consistent over time [210, 229, 327, 166, 102, 405].

Chapter 3

Modeling Head-direction Cells: State of the Art

Since early neurophysiological results in the 80s showed the existence of directional cells in the rat brain [279, 353, 347, 388, 229, 71, 41], several approaches have been proposed to model the orientation selectivity property of these neurons. Despite their differences, all these models postulate a primary role of inertial stimuli (e.g., vestibular and proprioceptive information) in determining the dynamics of biological directional cells. Nevertheless, different models focus on distinct underlying aspects, and propose either a more theoretical, neurophysiologically plausible, or functional approach. In this chapter, we survey the most relevant approaches which have characterized the research on directional cells in the last decade.

3.1 McNaughton *et al.* (1991)

McNaughton *et al.* [210] propose the first mechanism to integrate vestibular signals in order to update the directional representation. In particular, they employ a linear associative neural network to replace the mathematical integration of angular velocity. The input of the network consists of the current heading information $H(t)$, and the current angular displacement $H'(t)$[1]. Then, the linear associative mapping defined as

$$\varphi : H \times H' \rightarrow H$$

provides the next angular heading $H(t + \Delta t)$.

McNaughton *et al.* consider a population H of directional cells encoding $H(t)$, and a population of angular velocity cells H' encoding $H'(t)$. In addition, since linear independence of the inputs is a necessary condition for linear mapping, they propose an intermediate group of neurons in which all possible joint values HH' are represented by a set of linearly independent vectors. Each cell belonging to the HH' population is tuned to both a specific heading and a specific angular velocity.

At each time step the system is in a particular state $H(t)$ and the $H'(t)$ signal is used to induce the shift towards the new state $H(t + \Delta t)$. The HH' population acts as a neural associative look-up table of

[1]In fact, $H'(t)$ is the angular velocity over an arbitrary discrete time Δt, that is $H'(t) = dH/dt$.

angular velocity integrals, and determines the state transition $H(t) \rightarrow H(t + \Delta t)$ by projecting onto H cells properly. Below is a functional view of the system

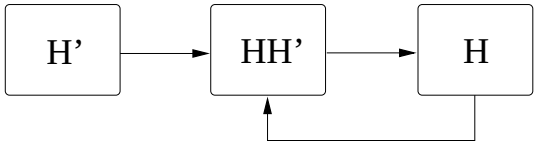

For instance, let $H(t) = 90°$ and $H'(t) = 45°/\Delta t$ be the current heading state and the current angular velocity, respectively. Then, the cell $i \in$ HH' encoding the joint feature $(H(t), H'(t))$ is maximally active. Cell i projects to the H population so that it will activate the directional neuron coding for $H(t + \Delta t) = 135°$.

With respect to the issue of calibrating directional cells by means of external cues, McNaughton *et al.* assume a population of local-view cells (i.e., direction-selective place cells) projecting to the H population and suggest to apply Hebbian learning to correlate local-view cell activity to H cell activity.

The hypothesis postulated by McNaughton *et al.* is the first plausible theory explaining both the update mechanism underlying head-direction cells and the influence of extrinsic signals. However, they address the problem on a rather abstract level, without accounting for neural connections and dynamics. Indeed, they do not report any implementation of the model into a concrete directional system.

3.2 Skaggs *et al.* (1995)

Skaggs *et al.* [327] propose a refinement of the model suggested by McNaughton and colleagues. In particular, they focus on the neural architecture of the directional system. A one-dimensional attractor network scheme [393, 9, 96, 165, 172, 173] is used to consistently model the population of directional cells. That is, head-direction cells are coupled by intrinsic connections, so that nearby cells are linked by strong excitatory synapses, and distant cells are connected by strong inhibitory projections. A simplified scheme of the network follows

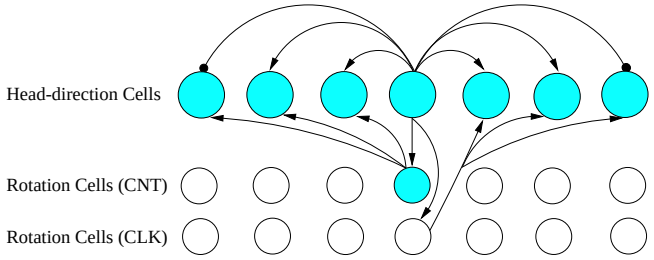

Due to the attractor nature, at each time step the directional system exhibits a stable representation consisting of a single localized cluster of active cells. The attractor scheme also determines the shift of activity over the head-direction population. Given a stable state, if an excitatory external signal is applied to cells on one side of the peak, the activity will gradually shift towards the side at which the input has been applied.

Skaggs *et al.* consider two populations of "rotation" cells whose role is similar to that of HH' cells in the model by McNaughton and colleagues. One group of rotation cells is responsible for clockwise turns, the other for counterclockwise turns. Each rotation cell is active only if the current heading is

equal to its preferred direction and the head is turning according to its preferred angular velocity sign. Clockwise rotation cells project to head-direction cells neighboring them on the right. Counterclockwise rotation cells project to head-direction cells neighboring them on the left. Then, during clockwise turns, for instance, clockwise rotation cells will excite directional cells to the right of the current peak, and the hill of activity will shift rightward.

Skaggs *et al.* also consider a set of visual feature detectors, each of which responds maximally to a specific visual cue located at a specific egocentric bearing angle. These neurons project to head-direction cells, and Hebbian learning is used to modify synaptic connections. The activity of visual feature detectors is used to correct head-direction activity. Note that, since visual detectors are head-centered (egocentric) and not world-centered (allocentric) orientation selective, Skaggs *et al.* assume that detected visual cues are infinitely distant from the animal's motion area.

To summarize, Skaggs *et al.* model the neural mechanism underlying directional cell activity and focus on how the peak of the population activity may shift to track head rotations. However, the proposed model is still an abstract description of the head-direction system and it is not supported by any simulations or implementations.

3.3 Zhang (1996)

Zhang [405] stresses the central role of the intrinsic dynamics of the head-direction ensemble. The model includes a fully-connected neural network having a one-dimensional ring topology. Zhang considers the distribution of the preferred directions over the entire population as continuous and uniform, and focuses on the average behavior of directional cells. In particular, an analytical framework is formulated accounting for: *(i)* Gaussian-shaped directional tuning curves. *(ii)* The stationary self-sustaining dynamics of the system. A pure attractor network is constructed by defining both a weight distribution function describing the average synaptic strength between directional cells, and an activation sigmoid function relating the average firing activity of cells to their average inputs. *(iii)* A dynamical shift model. Another weight distribution function is defined which extends the weight function used for the static case by introducing an antisymmetric component (anisotropy in the connection matrix). This modifies the system's dynamics by producing a shift of the hill of activity over the entire population. In addition, the weight distribution function is such that the shape of the activity profile is invariant with respect to system rotation.

In Zhang's model, the velocity of the shift is determined by a factor $\gamma(t)$ which defines the magnitude of the antisymmetric component in the weight distribution. It is shown that if $\gamma(t)$ is taken proportional to the angular head velocity $\omega(t)$, then the network works as a perfect velocity integrator. On the other hand, if $\gamma(t)$ also contains a component proportional to the angular head acceleration ω', then the directional shift over the network tends to lead the true head direction by a constant time delay τ. Zhang argues that this might explain experimental data on anterior thalamic directional cells (ADN) [41].

The model also accounts for head-direction calibration based on external signals. In particular, Zhang considers one local-view detector responding to one specific local view. Then, Hebbian learning is used to correlate the activity of this detector to the head-direction ensemble activity.

To recapitulate, Zhang puts forth a mathematical framework to explain the neurophysiological properties of head-direction cells. An explicit analytical formulation of the head-direction dynamics as a continuum is proposed, in which the descriptions of both the static and the dynamic behaviors in terms of the symmetry of the synaptic weight distribution are unified. Zhang reports the first simulation results of a directional system based on a pure attractor scheme. The mathematical description of the directional shift might explain the anticipatory behavior of ADN cells [41]. However, the model does not explain the fact that some ADN cells tend to increase their firing rate at high turning velocity [41]. In addition, Zhang does not explicitly transpose his model into an anatomically plausible implementation, that is he

does not separate the role of distinct biological areas involved in direction representation. Finally, he does not report any results concerning the tracking capability of the system on realistic data.

3.4 Blair (1996)

Blair [37] suggests that postsubiculum (poSC) and anterior thalamus (AT) might form a thalamocortical circuit to integrate angular displacements over time. In particular, he focuses on the role of anticipatory head-direction cells (termed AHD cells in the model) lying in AT. AHD cells, which predict future headings, project to poSC cells (termed PHD cells in the model) which code for the current animal's heading. PHD neurons can be seen as the equivalent of the H cells in the model by McNaughton *et al.* [210]. AHD cells inform PHD cells about when to update their representation of the current head-direction. PHD cells do nothing but receiving afferent inputs from AHD cells. In the model, AHD cells anticipate PHD cells due to an assumed transmission delay ($5\,ms$) between AT and poSC cells. Furthermore, the strength of AHD $\rightarrow$ PHD synapses is weak so that each AHD cell must fire several spikes in order to trigger a PHD cell. Below is a simplified overview of the model

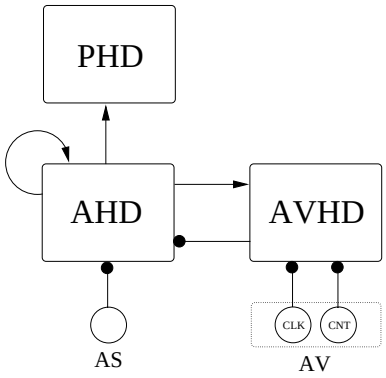

The state-shifting behavior of the system relies on the mutual interaction between AHD cells: Lateral excitatory and inhibitory projections are used to make the system rotate according to head angular velocity. In particular, each AHD projects excitatory synapses to its neighbors. Inhibition is obtained via two group of cells, namely AVHD$_{CLK}$ and AVHD$_{CNT}$, which fire as a function of the head-direction and the sign of the angular velocity (similar to HH' cells in the model by McNaughton *et al.*). Blair postulates that AVHD cells might be located in the reticular thalamic nucleus. In order to correlate the AVHD cell activity to angular velocity, Blair considers two velocity cells, named AV$_{CLK}$ and AV$_{CNT}$, coding for the sign of the angular velocity. These cells are assumed to be located either in the postsubiculum, or in the retrosplenial cortex. The angular velocity cell AV$_{CLK}$ inhibits AVHD$_{CNT}$ cells, whereas AV$_{CNT}$ inhibits AVHD$_{CLK}$ cells.

Each AHD cell i excites the corresponding AVHD$_{CLK}$ cell which in turn inhibits the AHD cell $i-1$ (i.e., the cell neighboring i on the left). On the other hand, each AHD cell i excites the AVHD$_{CNT}$ cell which inhibits the AHD cell $i+1$ (i.e., the cell neighboring i on the right). During clockwise turns, for instance, AVHD$_{CLK}$ are active whereas AVHD$_{CNT}$ are silent. The AHD cell i will directly excite its neighbors $i-1$ and $i+1$, but cell $i-1$ will be inhibited by the AVHD$_{CLK}$ cell. As a consequence, the activity will shift towards cell $i+1$ making the system rotate clockwise.

The velocity of the system is determined by an angular speed cell AS, whose activity is inversely proportional to the magnitude of angular velocity. The AS cell inhibits AHD cells, so that during fast

turning, for instance, propagation through the AHD layer is weakly inhibited and the system rotates quickly. Blair predicts that the AS cell might be located in the mammillary bodies.

Blair's hypothesis is the first in which AT and poSC cells form an angular velocity dead-reckoning circuit. The model stresses the importance of the anticipatory behavior of AT cells, in particular the fact that the AT representation leads the poSC representation.

The model fits some properties of AT and poSC biological cells. Nevertheless, the system produces directional tuning curves that are much narrower than real ones. Likewise, Blair only validates the tracking performance of the system for a simple case (one single 60° back and forth rotation), whereas he does not report any results on more realistic data. Finally, the system does not provide any calibration mechanism to reset vestibular information in order to keep the tracking error bounded over time.

3.5 Redish, Elga, and Touretzky (1996)

Redish, Elga, and Touretzky [284] further develop the hypothesis of attractor network (formulated by Skaggs *et al.* and simulated by Zhang) to maintain stability in the head-direction system. In particular, they utilize a coupled attractor network in order to create a directional system characterized by a Gaussian attractor state [96]. The attractor network consists of two groups of neurons: One pool of excitatory units E, and one pool of inhibitory cells I. Neurons in both pools have evenly distributed preferred directions. Each excitatory cell $i \in$ E with preferred direction θ_i projects strong exciting synapses to neurons in both E and I having preferred directions close to θ_i. Each inhibitory neuron $j \in$ I weakly inhibits all cells in both E and I pools, whereas cells in E and I having preferred directions close to θ_j are inhibited slightly more. The dynamics of the attractor module is such that both excitatory and inhibitory pools are characterized by a single stable state at each time. Below is a scheme of the coupled attractor network

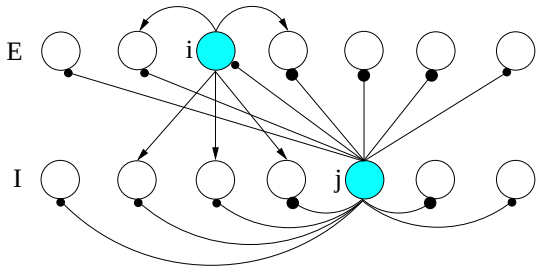

Redish, Elga, and Touretzky propose a head-direction system involving primarily the postsubiculum (poSC) and the anterior thalamic nuclei (AT). Both poSC and AT are characterized by the above coupled attractor architecture. Interconnections between poSC and AT are between their excitatory pools E_{poSC} and E_{AT}, respectively. A set of synapses, termed matching connections, interlinks neurons in E_{poSC} with neurons in E_{AT} having equivalent preferred directions. Moreover, a set of projections, namely left- and right-offset connections, is responsible for updating the head-direction representation. Each unit $i \in \mathrm{E}_{poSC}$ with preferred direction θ_i has a left-offset connection to unit $j_l \in \mathrm{E}_{AT}$ such that $\theta_j = \theta_i - \delta$, and has a right-offset connection to unit $j_r \in \mathrm{E}_{AT}$ such that $\theta_j = \theta_i + \delta$. The offset δ is the same for all units and it is equal to $\delta = 10^\circ$.

All offset synapses are modulated by the head angular velocity ω. Thus, during rightward turnings, right-offset connections have a strength proportional to the magnitude of angular velocity $|\omega|$, whereas left-offset connections have strength zero. By contrast, during leftward rotations, left-offset connections

are proportional to $|\omega|$ and right-offset connections are zero. When the head is still, both left- and right-offset connections have strength zero, and matching synapses synchronize poSC and AT representations.

Similar to Blair's model, the system by Redish, Elga, and Touretzky involves an inhibitory input to AT coming from the mammillary bodies (MB). This input is assumed to be proportional to the magnitude of the angular velocity. MB cells work as a gain control mechanism to compensate for modulated offset connections. This allows the system to maintain the shape of the hill of AT activity nearly unchanged during rotations (otherwise, the combined input from offset and matching connections would distort it). A scheme of the entire system is sketched below

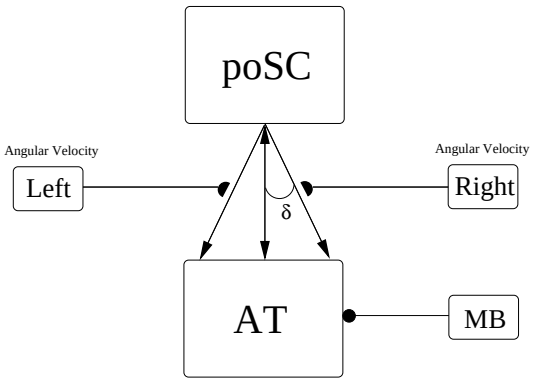

Simulation results show biologically plausible directional tuning curves as well as effective head-tracking capabilities (in particular, the authors report tracking performance over $20\,s$ of simulation). Also, in the model, AT activity anticipates poSC activity by a constant leading time τ of approximately $10\,ms$ (capturing the anticipatory property of real anterior thalamic neurons).

As discussed by Redish (1997) [282], a major problem for this theory is its discrepancy with anatomical lesion data. Indeed, after lesions to the postsubiculum (poSC), anterior thalamic cells (AT) are still directional selective [129, 351]. By contrast, the model predicts that lesions to poSC cells would disrupt the entire system. Also, the authors do not report any explicit results concerning the calibration of the cumulative tracking error by means of extrinsic cues. Finally, as pointed out by Redish (1997) [282], the model requires the existence of poSC $\rightarrow$ AT synapses modulated by head angular velocity. An alternative system would involve a population of cells firing as a function of both head-direction and angular velocity (i.e., HH' and AVHD cells in the models of McNaughton *et al.* and Blair, respectively) to drive AT cell activity.

Chapter 4

A Model of the Rodent Directional Sense

We model biological head-direction cells by means of a neural architecture in which internal (e.g., vestibular) and external (e.g., visual) cues are combined to establish a stable head-direction representation. The dynamics of the system is primarily controlled by idiothetic stimuli which determine the directional selectivity property. Allothetic information is used to occasionally modify the system's dynamics and calibrate head-direction cell activity. Fig. 4.1 shows a functional overview of the model. The architecture may be understood as consisting of three functional modules: *(i)* An angular velocity integrator, *(ii)* a neural substrate forming the output of the system, and *(iii)* a module incorporating external signals to achieve system calibration. The model is implemented and tested on a robot.

4.1 Integrating Angular Velocity

In order to integrate the angular velocity ω over time, we consider a directional circuit involving three neural populations, namely HAV, ADN, and LMN (Fig. 4.1).

Head Angular Velocity Cells. A directional system, whose dynamics is primarily controlled by inertial self-motion signals, needs to assess angular displacements continuously. In both rodents and primates, there exist cells coding for angular velocity based on vestibular and sensory-motor information [212]. Cells that are maximally active when the head is turning in one specific direction, as well as cells responding maximally only when the head is not turning, have been observed in the posterior parietal, retrosplenial, and posterior sensory-motor cortices [210, 376].

We take a population HAV$= \{ \Omega_+, \ \Omega_=, \ \Omega_- \}$ of cells whose activity is correlated to the sign and the magnitude of the angular velocity ω. Cells Ω_+ and Ω_- fire proportionally to $|\omega|$ only during clockwise and counterclockwise turning, respectively. Cell $\Omega_=$ fires only when the robot is not rotating.

LMN Directional Cells. LMN cells model neurons in the lateral mammillary nuclei which are tuned to the animal's heading θ as well as the head angular velocity ω [181]. We discretize the continuous angular space $\Theta = [0^\circ, 360^\circ]$ in $S = 180$ steps of 2° each. Then, we define a population LMN$=$ $\{ \{ l_z, l_s, l_m, l_f \} \mid 1 \leq l \leq S \}$ of $S \times 4 = 720$ cells with evenly distributed preferred directions θ_l. In particular, we take four cells l_i for each direction θ_l. That is, the magnitude of the angular velocity $|\omega|$ is discretized by LMN cells in four possible states which are ω_z (zero), ω_s (slow), ω_m (medium), and ω_f (fast). We assume that $\omega_z = 0^\circ/sec$, $\omega_s = 100^\circ/sec$, $\omega_m = 200^\circ/sec$, and $\omega_f = 400^\circ/sec$. A LMN

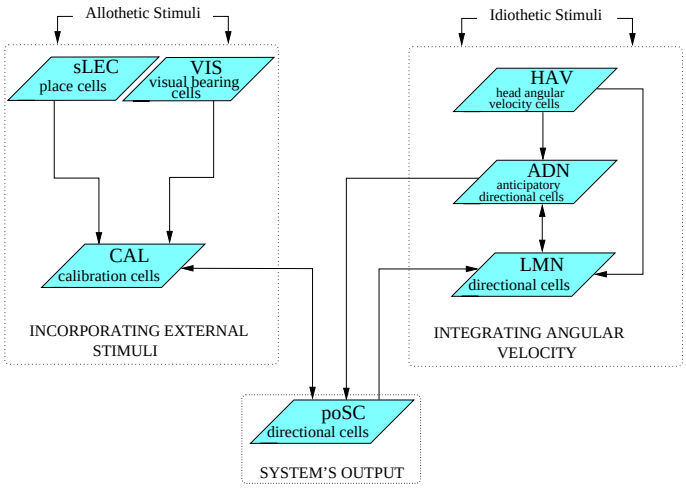

Figure 4.1: A functional overview of our directional system. Idiothetic inputs (right-hand pathway) are integrated over time by the circuit including ADN, LMN, and HAV. Allothetic signals (left-hand pathway) are incorporated to calibrate head-direction cells by the circuit made of sLEC, VIS, and CAL cells. Finally, poSC cells form the output of the system.

cell fires maximally when the agent's heading is equal to its preferred direction *and* the agent is turning at a specific angular velocity. For instance, the activity of cell $l_m \in$ LMN is maximal when $\theta = \theta_l$ *and* $|\omega| = \omega_m$.

ADN Directional Cells. Directional cells in the anterodorsal nucleus have the property of anticipating the animal's future direction during head turning, whereas they encode the current direction when the head is still. This anticipatory behavior depends on a constant time delay τ [41, 355].

According to recent experimental findings [37], we take τ_a as characteristic for each cell $a \in$ ADN. In other words, we assume that some ADN cells may anticipate head-direction more than others. We consider five distinct time delays $\tau_1 < \cdots < \tau_5$ for each preferred direction θ_a, such that $\tau_1 = 20\,ms$, $\tau_2 = 40\,ms$, $\tau_3 = 60\,ms$, $\tau_4 = 80\,ms$, and $\tau_5 = 100\,ms$. Then, we define a population ADN$=\{\{a_1,\ldots,a_5\} \mid 1 \leq a \leq S\}$ in which five ADN cells a_j (having different time delays τ_j) code for each direction θ_a. Every a_j cell anticipates θ by a different angle when $|\omega| > 0$, but all preferred directions θ_{a_j} are equal to θ when $\omega = 0$.

4.1.1 Idiothetic-based Dynamics

LMN and ADN interact with each other in order to integrate angular motion over time, that is, to update the direction representation according to HAV cell activity. In this section, we describe the synaptic connections and the neural activity of HAV, LMN, and ADN populations.

HAV Cell Activity. The activity of a cell $\Omega_i \in$ HAV is correlated to the sign and magnitude of the angular velocity ω according to

$$r_{\Omega_+} = \begin{cases} \frac{\omega}{\omega_{max}} & \omega > 0 \\ 0 & otherwise \end{cases}, \; r_{\Omega_-} = \begin{cases} \frac{|\omega|}{\omega_{max}} & \omega < 0 \\ 0 & otherwise \end{cases}, \; r_{\Omega_=} = \begin{cases} 1 & \omega = 0 \\ 0 & otherwise \end{cases} \tag{4.1}$$

Thus, $r_{\Omega_+} \in [0,1]$, $r_{\Omega_-} \in [0,1]$, and $r_{\Omega_=} \in \{0,1\}$.

ADN Cell Activity. ADN cells play a major role in determining the system's dynamics. They induce the synchronized shifting behavior enabling directional cells to track rotations. ADN cells are driven by

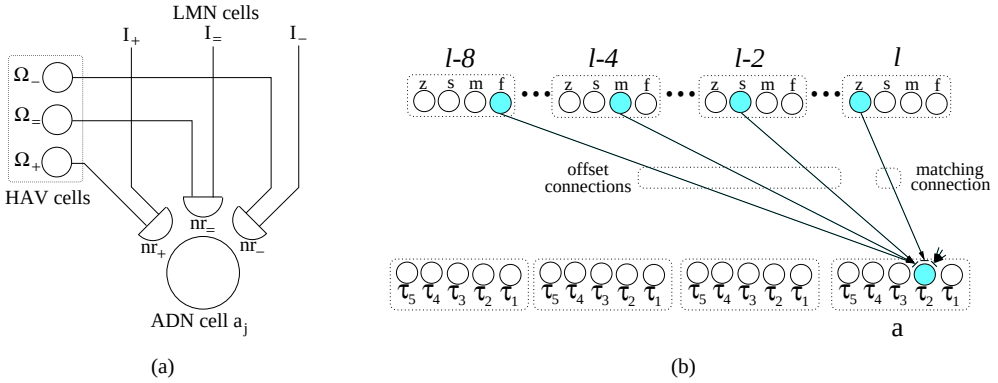

Figure 4.2: (a) Each ADN cell receives afferents through three distinct receptors nr_+, $nr_=$, nr_-. A neuroreceptor nr_i is open to inputs I_i from LMN cells only when the corresponding angular cell ω_i is active. (b) Synaptic inputs to an ADN cell a_j, with $j = 2$.

LMN and HAV cells. In particular, LMN feeds ADN with information about the current heading $\theta(t)$, whereas HAV determines the direction and velocity of the shift of ADN activity.

Let a_j be an ADN cell, where $1 \leq a \leq S$ and $1 \leq j \leq 5$. In the current implementation, all ADN afferent connections have synaptic weight equal to 1. Cell a_j has three neural receptors nr_+, $nr_=$, nr_-, which correspond to positive, negative, and zero ω, respectively. Each receptor nr_i receives inputs I_i from LMN cells and it is gated by the activity of cell $\omega_i \in$ HAV (Fig. 4.2 (a)).

In order to model the ADN anticipatory property (Eq. 2.1), we employ LMN $\rightarrow$ ADN *matching* and *offset* synaptic projections (Fig. 4.2 (b)) [327], such that the higher τ_j, the more a cell a_j is informed in advance about the current heading $\theta(t)$ during agent rotation, and the more the cell may anticipate the future direction $\theta(t')$. On the other hand, projections LMN$\rightarrow$ADN are such that ADN cells a_j may code for the present heading when the agent is not rotating.

Let l and a be LMN and ADN cells such that $\theta_l = \theta_a$. Each a_j neuron, with $1 \leq j \leq 5$, receives:

(i) One matching input from cell l_z on its receptor $nr_=$;

(ii) Three left-offset inputs on its nr_+ neuroreceptor from cells $(l-j)_s$, $(l-2j)_m$, and $(l-4j)_f$, respectively;

(iii) Three right-offset inputs on its nr_- neuroreceptor from cells $(l+j)_s$, $(l+2j)_m$, and $(l+4j)_f$, respectively.

The activity $r_{a_j}(t)$ of an ADN cell is determined by its LMN input. The sign of the angular velocity ω decides which receptor nr_i is currently open, while the magnitude of ω determines which LMN cell is currently firing through nr_i. For instance, if $\omega > 0$ and $|\omega| = \omega_m = 200°/sec$, then cell a_j will listen to its nr_+ input and will fire according to

$$r_{a_j}(t) = r_{(l-2j)_m}(t) \cdot w_{jl} \tag{4.2}$$

where w_{jl} is the connection from cell $(l-2j)_m$ to cell a_j.

Due to the above connecting scheme, the higher τ_j, the more cell a_j is informed in advance about the current heading $\theta(t)$ during agent rotations. In other words, the higher τ_j, the more cell a_j will anticipate future directions. The rotational speed of the system depends on the anticipatory behavior of ADN cells and then on the angular velocity ω as discretized by LMN activity. Moreover, since receptors nr_i are gated by the continuous HAV activity, the larger $|\omega|$, the more frequently cell a_j will listen to its LMN afferents inputs. Then, the velocity of the shift in ADN activity is proportional to the continuous $|\omega|$.

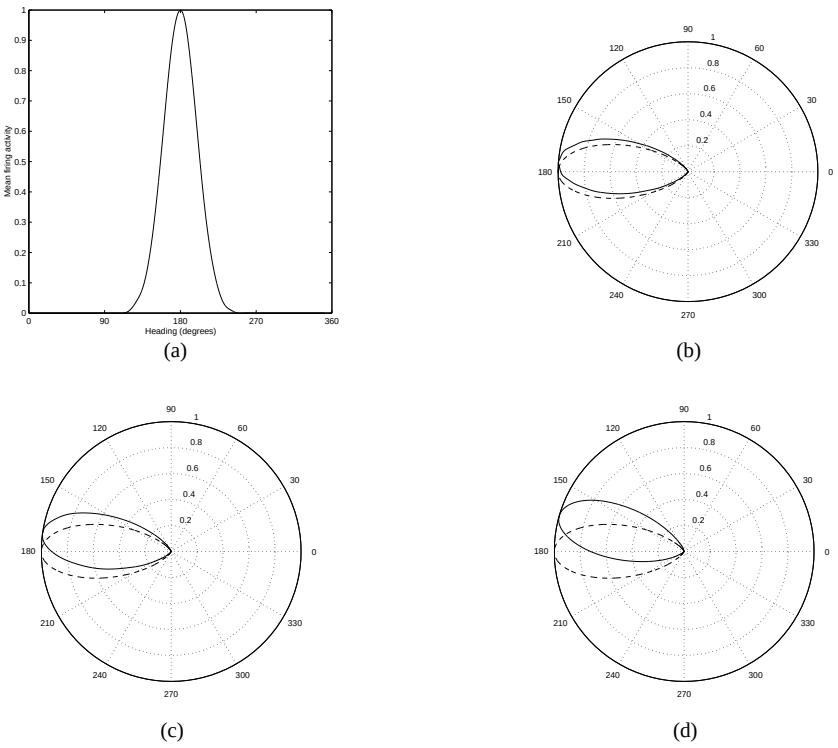

Figure 4.3: (a) Tuning curve of a cell $a_j \in$ ADN of our model with $\tau_j = 40\,ms$. When the robot is not rotating, the preferred direction is $\theta_a = 180^\circ$. (b) Polar tuning curve (solid line) of cell a_j during counterclockwise turning with $|\omega| = 100^\circ/sec$. The cell anticipates the future heading (dashed line) by about 4°. (c) During counterclockwise turning with $|\omega| = 200^\circ/sec$, cell a_j shifts its preferred direction of about 8° (solid line). (d) During counterclockwise turning with $|\omega| = 400^\circ/sec$, cell a_j shifts its preferred direction of about 16° (solid line).

Fig. 4.3 shows some results obtained by recording the firing activity of a cell a_j in the ADN layer of the model. The time delay of the recorded cell is $\tau_j = 40\,ms$. When the robot is not rotating (Fig. 4.3 (a)), cell a_j exhibits a preferred direction $\theta_a = 180^\circ$. Figs. 4.3 (b, c, d) show the polar tuning curves of cell a_j during counterclockwise turning, when $|\omega| = \omega_s = 100^\circ/sec$, $|\omega| = \omega_m = 200^\circ/sec$, and $|\omega| = \omega_f = 400^\circ/sec$, respectively. Cell a_j anticipates the robot's future heading by shifting its preferred direction (solid curve) of approximately 4°, 8°, and 16°, respectively, which is consistent with Eq. 2.1.

LMN Cell Activity. LMN cells are correlated with the heading $\theta(t)$ as well as with the angular velocity ω. Each cell $l_i \in$ LMN receives afferents from five ADN cells a_j (with, $\tau_1 < \cdots < \tau_5$) and modulates its firing activity according to HAV cell activity.

ADN cells inform LMN cells about future headings. Let l and a be LMN and ADN cells, respectively, such that $\theta_l = \theta_a$. Each cell $l_i \in \{\, l_z,\, l_s,\, l_m,\, l_f \,\}$ receives five input connections w_{ij}, one from each cell a_j, where $1 \leq j \leq 5$ (Fig. 4.4 (a)). Connections w_{ij} are fixed and are defined according to

$$w_{ij} = \begin{cases} W_{strong} & j = 1 \\ W_{weak} & 2 \leq j \leq 5 \end{cases} \tag{4.3}$$

where W_{weak} and W_{strong} are two constant terms such that $W_{strong} \gg W_{weak}$, and $\sum_j w_{ij} = 1$. The activity of a LMN cell l_i is given by

$$r_{l_i}(t) = \begin{cases} \sum_{j=1}^{5} w_{ij} \cdot r_{a_j}(t) & |\omega| = \omega_i \\ 0 & otherwise \end{cases} \tag{4.4}$$

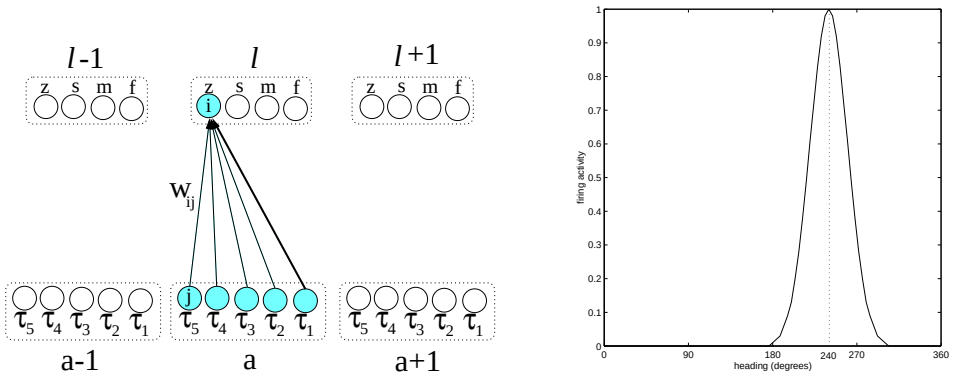

(a) (b)

Figure 4.4: (a) Synaptic projections from ADN to LMN in our model. (b) Tuning curve of a head-direction cell $l_i \in$ LMN with preferred direction $\theta_l = 240°$.

where $\omega_i \in \{\omega_z, \omega_s, \omega_m, \omega_f\}$. Due to the distinct time delays τ_j associated with each preferred direction, each LMN cell l_i is actually informed about arrival at θ_l by a sequential activation on its ADN inputs. We assume that only when *all* ADN inputs a_j have sequentially predicted arrival at θ_l, cell l_i is enabled to fire fire according to Eq. 4.4. Note that, due to Eq. 4.3, $r_{l_i}(t) \approx r_{a_1}(t)$. That is, cell l_i needs the sequential activation of all anticipatory inputs a_j to be able to fire, but its firing rate is mostly determined by the most recent input a_1. Fig. 4.4 (b) shows a result obtained by recording the firing activity of a cell $l_i \in$ LMN in our robotic implementation. The cell's preferred direction is $\theta_l = 240°$.

4.2 Interpreting the Directional Output

poSC Directional Cells. We consider a population of cells poSC $= \{\, p \mid 1 \leq p \leq S \,\}$, with evenly distributed preferred directions θ_p, forming the output of the directional system (Fig. 4.1). Thus, at each time t, the poSC ensemble activity $\mathcal{R}^{poSC}(t) = \{\, r_p(t) \mid 1 \leq p \leq S\}$ provides the estimate $\overline{\theta}(t)$ of the robot's allocentric heading $\theta(t)$.

Experimental results suggest that postsubicular cells are indirectly related to internal cues through thalamic directional cells [128, 405]. Our poSC cells are primary driven by ADN cells. In particular, ADN $\rightarrow$ poSC projections are defined according to the same scheme used for ADN $\rightarrow$ LMN connections. Also, poSC activity is based on the same mechanism used to drive LMN cells. That is, a cell p is enabled to fire only when all its ADN inputs have been sequentially activated. Fig. 4.5 (a) shows the tuning curve of a cell p recorded from the poSC layer of our model. Cell p responds maximally when the robot's heading is approximately equal to $\theta_p = 160°$.

In order to interpret the poSC ensemble activity as output of the directional system, we apply *population vector coding* [121, 394]. We compute the estimation of the robot's heading $\overline{\theta}(t)$ according to

$$\overline{\theta}(t) = \arctan\big(\sum_p \sin(\theta_p)\, r_p(t) \,/\, \sum_p \cos(\theta_p)\, r_p(t)\big) \tag{4.5}$$

where $r_p(t)$ is the firing rate of poSC cell p, and θ_p is its preferred direction. Fig. 4.5 (b) shows a population activity as recorded from our poSC layer during robot turning. According to Eq. 4.5, the ensemble activity codes for the allocentric heading $\overline{\theta}(t) \approx 180°$.

When the robot first enters the arena, the system is initialized with respect to an arbitrary absolute direction Φ. We set the population activity of poSC, LMN, and ADN cells as a Gaussian distribution centered at Φ and with $\sigma = 20^o$. That is, the activity of a cell i is initialized as $r_i(t) = \exp(-d_i(t)^2 \,/\, 2\sigma^2)$,

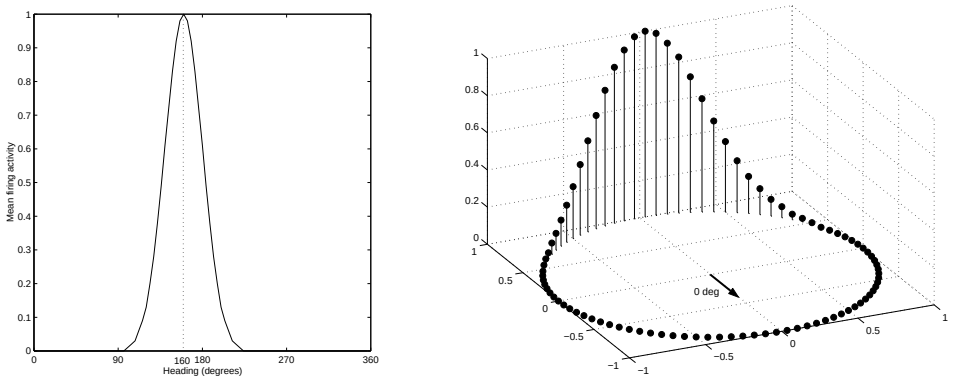

(a) (b)

Figure 4.5: (a) Tuning curve of a directional cell $p \in$ poSC of our model. The cell's preferred direction is approximately $\theta_p = 160^\circ$. (b) A sample of poSC population activity coding for a robot's heading of approximately 180°.

where the deviation $d_i(t)$ is given by

$$d_i(t) = \begin{cases} \min(|\,\theta_i - \Phi\,|,\ |\,\Phi - \theta_i + S\,|) & \theta_i > \Phi \\ \min(|\,\theta_i - \Phi\,|,\ |\,\theta_i - \Phi + S\,|) & otherwise \end{cases} \tag{4.6}$$

As the robot starts rotating, the system starts shifting its internal representation to track angular displacements.

4.3 Using Allothetic Cues for Head-direction Calibration

Neurons in the postsubicular cortex (poSC) are correlated with external directional information [129, 351, 405]. They receive afferents from the posterior parietal (PP) cortex as well as the laterodorsal nucleus (LDN), two structures which are strongly related to allothetic signals. We assume that cells in the poSC layer of our model are responsible for incorporating allothetic inputs (e.g., visual cues) in order to maintain the heading representation consistent over time.

We let the robot estimate its egocentric bearing α relative to a weak light source L which is located on one of the walls of the arena (Fig. 1.2). The idea is to use the bearing signal α for calibrating the directional system. However, since L is not an infinitely distant cue, it does not provide an absolute directional reference (as the sun does for desert ants and bees [380, 383, 78]). Thus, compared to the absolute framework defined by Φ, the egocentric bearing α is not invariant with respect to spatial location (Fig. 4.6). As a consequence, we need to combine the bearing information α with some place coding information. We consider a calibrating system (Fig. 4.1) made of:

(i) A population of visual bearing cells (VIS), whose activity encodes the robot's egocentric bearing $\alpha(t)$ at time t;

(ii) A population of place cells (sLEC) coding for the robot's current spatial location $\vec{p}(t)$ based on visual information only;

(iii) A population of calibration cells (CAL), driven by VIS and sLEC cells, firing as a function of $\alpha(t)$ and $\vec{p}(t)$. We call these neurons calibration cells (CAL) because their activity is directly used to calibrate head-direction cells.

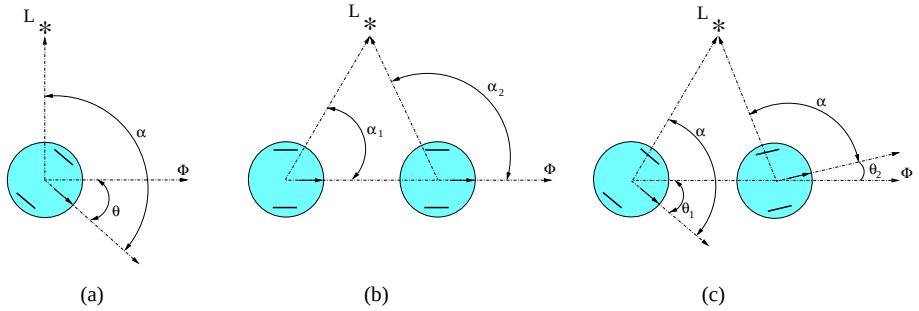

Figure 4.6: (a) A light source L is used as an external polarizing cue. The robot (dark-grey circle) estimates its egocentric bearing α to calibrate head-direction cells. With respect to Φ, the bearing α is not invariant with respect to the robot's location. (b) At two distinct positions the robot has the same allothetic heading θ, but two distinct egocentric bearings α_1, α_2. (c) At two distinct positions the robot has the same bearing α, but two different headings θ_1, θ_2.

Visual Bearing Cells. A cell $v \in$ VIS fires as a function of the current bearing angle α between the robot and the light cue L. Biological cells that respond maximally only when an external stimulus arrives from a particular egocentric direction have been observed in the inferior parietal cortex, the internal medullary thalamic lamina and the superior collicus [210, 327, 166]. For the robotic implementation, we consider a population VIS $= \{ v \mid 1 \leq v \leq S \}$ of $S = 180$ cells with preferred directions evenly distributed over $360°$. In order to interpret the VIS ensemble activity $\mathcal{R}^{VIS}(t) = \{ r_v(t) \mid 1 \leq v \leq S\}$, we apply population vector coding (Eq. 4.5).

The robot detects light by means of eight infrared sensors and one light-sensor facing forward. These nine sensory inputs $\vec{s}(t) = (s_1(t), \ldots, s_9(t))$ have to be interpreted to drive VIS activity $\mathcal{R}^{VIS}(t)$. To deal with such a noisy and high-dimensional input, we trained a feed-forward neural network $\mathcal{N}$ to approximate the mapping function $\mathcal{M} : \mathcal{S} \rightarrow \mathcal{R}^{VIS}(\alpha)$, where $\mathcal{S}$ is the input sensory space. The idea is to let the robot learn $\mathcal{M}$ only once by off-line supervised learning. Network $\mathcal{N}$ is trained by using gradient descent back-propagation [294]. The network's input consists of the sensory reading vector $\vec{s}(t) = (s_1(t), \ldots, s_9(t))$. The output is the VIS cell population activity $\mathcal{R}^{VIS}(t)$ (see Appendix B for more details).

Vision-driven Place Cells. We employ place-coding information to map the egocentric bearing α into an allocentric directional framework. In particular, we consider place cells in the superficial lateral entorhinal cortex (sLEC) of our hippocampal model (Part III), whose activity depends on visual cues only. The firing activity $r_s(\vec{p})$ of a sLEC cell s is a function of the robot's spatial position $\vec{p}$ within the arena, regardless of its allocentric heading θ. Fig. 4.7 shows two samples of sLEC place fields recorded in our model after spatial learning.

Calibration Cells. We consider a population CAL $= \{ c \mid 1 \leq c \leq S \}$ of $S = 180$ calibration cells. CAL cells play a central role in the calibration process. They form the "neural boundary" between the allothetic vision-based representation and the idiothetic movement-related representation. On the one hand, CAL cells receive afferents from both VIS and sLEC cells. On the other hand, CAL cells are interconnected with poSC cells. In particular, each CAL cell c receives an input connection from one and only one poSC cell p, such that $\theta_c = \theta_p$. By contrast, cell c projects efferents onto all poSC cells (Fig. 4.8 (a)). Thus, calibration takes first place in poSC layer. Then, poSC propagates the calibration signal to LMN, which in turn starts driving ADN cells based on recalibrated signals.

Calibration cells (CAL) are characterized by a bimodal firing behavior depending on whether learning (i.e., correlating allothetic and idiothetic information) or calibration (i.e., updating directional representation based on external cues) is taking place. In both cases, CAL cell activity relies on a *winner-take-all* competitive scheme [295]: At time t, only the most active cell c^* (i.e., $r_{c^*} \geq r_c \ \ \forall c \in$ CAL) is

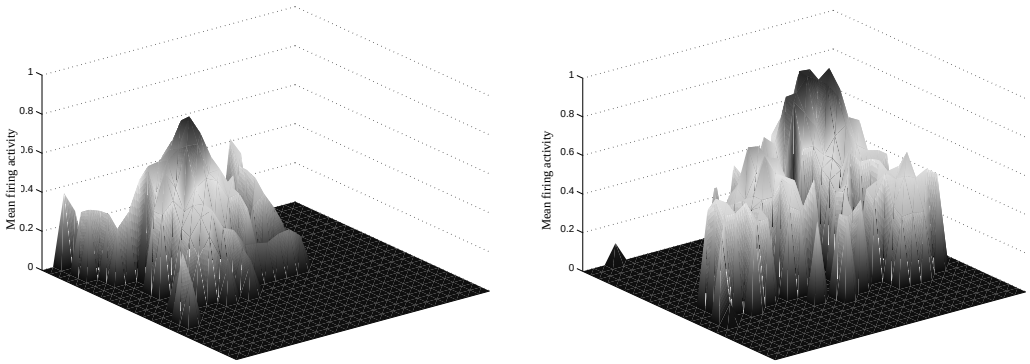

(a) (b)

Figure 4.7: Two examples of sLEC place fields recorded from our hippocampal model (Part III). Diagrams represent the firing rates $r_{s_1}(\vec{p})$ and $r_{s_2}(\vec{p})$ of cells s_1 and s_2 as a function of the robot's spatial position $\vec{p}$ within the arena. The peaks of the two tuning curves identify two rather localized regions of the space, which endows the robot with self-positioning capabilities.

enabled to fire, whereas all other cells are forced to remain silent. For sake of clearness, in the following we discuss learning and calibration separately.

4.3.1 Correlating Allothetic and Idiothetic Signals

We let the robot explore the environment and apply LTP[1] correlational learning to modify synaptic connections VIS → CAL, sLEC → CAL, and CAL → poSC. In particular, the aim is to associate the heading $\overline{\theta}(t_1)$ encoded by poSC cells at time t_1, with the bearing angle $\alpha(t_1)$ *and* the spatial location $\vec{p}(t_1)$ encoded by VIS and sLEC cells at time t_1, respectively. As a consequence, if at time t_2 the agent is at position $\vec{p}(t_2) \approx \vec{p}(t_1)$ with bearing $\alpha(t_2) \approx \alpha(t_1)$, it may calibrate its directional cells by recalling the memorized activity pattern corresponding to $\overline{\theta}(t_1)$. In other words, Hebbian learning enables the system to associate allothetic and idiothetic representations based on the robot's experience. CAL cells function as a long-term memory device: they allow the agent to store "snapshots" of poSC cell activity and use the combined signal $(\alpha, \vec{p}\,)$ to recall this activity for calibration.

During learning, CAL cell activity depends on poSC cells only, whereas all inputs coming from VIS and sLEC cells are inhibited. In particular, each CAL cell c is driven by the corresponding poSC cell p, such that $\theta_c = \theta_p$ (Fig. 4.8 (a)). The $p \rightarrow c$ connection is fixed and has a synaptic strength $w_{cp} = 1$. The activity r_c of each CAL cell c is given by

$$r_c(t) = \begin{cases} 0 & \exists c^* : r_{c^*}(t) > r_c(t) \\ r_p(t) & otherwise \end{cases} \tag{4.7}$$

Thus, during learning, CAL cells are rigidly coupled with poSC cells and function as directional neurons.

Connections CAL → poSC, VIS → CAL, and sLEC → CAL are initialized to zero, and changed on-line by means of one-shot Hebbian learning[2]

$$\Delta w_{ij} = \nu \, r_i \, r_j \, (1 - w_{ij}) \tag{4.8}$$

where i and j are the post- and presynaptic neuron, respectively, and $\nu = 1$ is the learning rate. At time t, learning is triggered only if:

[1] See Sec. 6.5 for a brief description of the long-term potentiation (LTP) learning mechanism.

[2] "One-shot" means that once a connection has been learned, it will not be modified any further.

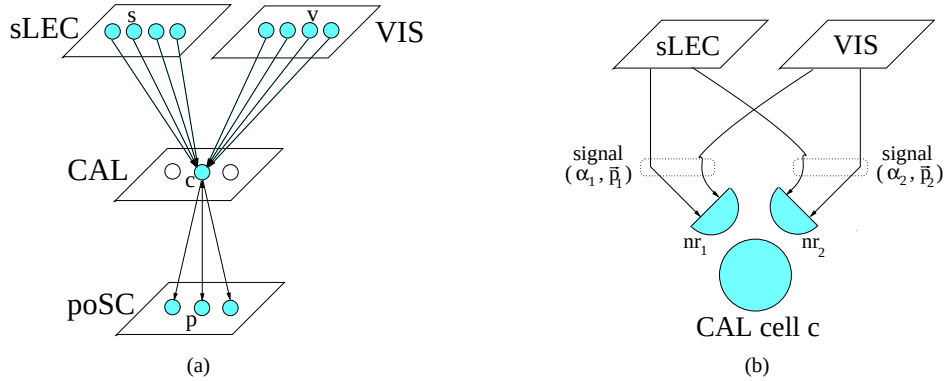

Figure 4.8: (a) Each cell $c \in$ CAL receives input connections from all VIS and sLEC cells. Moreover cell c receives one input from the corresponding cell $p \in$ poSC. Cell c projects to all poSC cells. (b) Combined signals $(\alpha_i, \vec{p}_i)$, which have been learned at distinct time steps t_i, enter cell c via distinct neuroreceptors nr_i.

(i) Relation $t - t' \leq \mathcal{T}$ holds, where t' is the last calibration time, and $\mathcal{T}$ is a fixed temporal threshold. This prevents the system from correlating allothetic and idiothetic signals when the deviation $|\theta(t) - \overline{\theta}(t)|$ might be large.

(ii) Relations $\sigma_v \leq \Sigma_v$ and $\sigma_s \leq \Sigma_s$ hold, where $\sigma_v(t)$ is the VIS activity variance around the center of mass $\mu_v(t)$, and $\sigma_s(t)$ is the sLEC activity variance around $\mu_s(t)$, respectively. Σ_v and Σ_s are two fixed thresholds. We assume that only when σ_v and σ_s variances fall below a threshold, the informations carried by VIS and sLEC are suitable for calibration.

Bearing-place associations $(\alpha_i, \vec{p}_i)$, which have been established at different learning times t_i, have to be maintained separate in order to achieve effective calibration. Thus, each CAL cell c has several neuroreceptors nr_i to keep separate signals $(\alpha_i, \vec{p}_i)$ that are temporally distinct (Fig. 4.8 (b)).

4.3.2 Calibrating Head-direction Cells

During non-learning, CAL cell activity is determined by external cues via connections VIS $\rightarrow$ CAL and sLEC $\rightarrow$ CAL previously created. In particular, the activity r_c of cell $c \in$ CAL is defined by

$$r_c(t) = \begin{cases} 0 & \exists\, c^* : r_{c^*}(t) > r_c(t) \\ \Big(1 + \exp\big(-\beta(\sum_j r_j(t) w_{cj}^{nr_i} - \frac{1}{2}\sum_j r_j)\big)\Big)^{-1} & otherwise \end{cases} \tag{4.9}$$

where j varies over VIS and sLEC cells, $w_{cj}^{nr_i}$ represents an afferent synapse to receptor nr_i, and β is taken equal to 1.

CAL cell activity is maximal only when the robot has a specific bearing α relative to L *and* it is at a specific spatial location $\vec{p}$. Thus, CAL cells enable the robot to recognize previously learned place-orientation contexts. As a consequence, whenever there exists a maximally active CAL cell (i.e., $\exists c$: $r_c(t) \geq \varepsilon$), shifting the directional cell activity towards CAL cell activity results in calibrating the robot's head-direction system.

Calibrating Technique n. 1. Calibration is accomplished by means of the learned CAL $\rightarrow$ poSC projections. Indeed, after one-shot Hebbian learning (Eq. 4.8), synapses encode the poSC ensemble activity as memorized when learning occurred. Thus, if at time t there exists a maximally active cell $c \in$ CAL, poSC activity may be simply calibrated by:

(i) Making cell c inhibit all ADN efferents to poSC;

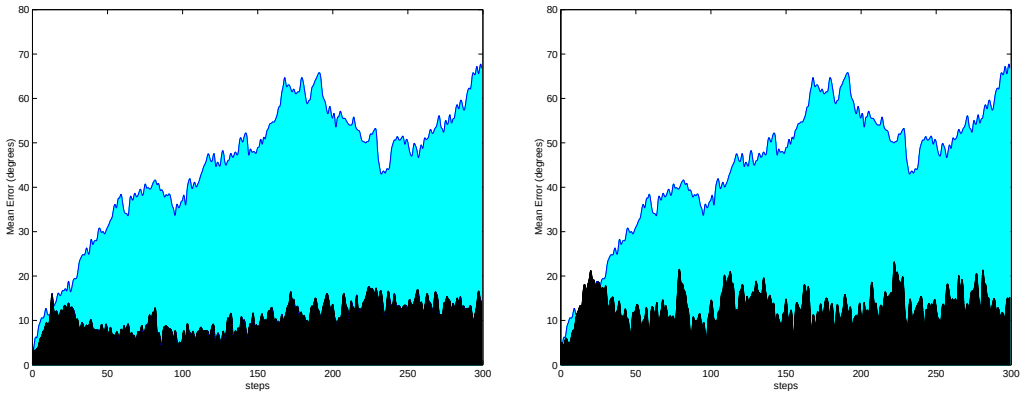

(a) (b)

Figure 4.9: Calibrating the internal directional system by using allothetic signals helps to keep bounded the tracking error over time. (a) The uncalibrated error (grey region) versus the calibrated error (black region) when applying the first calibration technique (Eq. 4.10). (b) The uncalibrated error versus the calibrated error when using the second calibrating technique (Eq. 4.11).

(ii) Updating poSC activity according to

$$r_p(t) = w_{pc} \cdot \max(r_c(t),\, 1) \qquad \forall p \in poSC \tag{4.10}$$

where w_{pc} weights the synapse connecting c to p.

Let $\overline{\theta}_c(t)$ and $\overline{\theta}_p(t)$ be the headings encoded by CAL and poSC activity at time t, respectively. Eq. 4.10 forces a shift in the poSC activity towards CAL activity such that, after calibration, the relation $\overline{\theta}_p(t) = \overline{\theta}_c(t)$ holds.

Calibrating Technique n. 2. Eq. 4.10 gives CAL cells the highest priority whenever a previously learned place-bearing situation is recognized. A slightly different approach would consist of considering a linear combination of idiothetic and allothetic signals to calibrate poSC activity. Let $\overline{\theta}_p^{new}$ denote the calibrated poSC representation. We take

$$\overline{\theta}_p^{new} = \arctan\left(\Big(\gamma \sin(\overline{\theta}_p) + (1-\gamma)\sin(\overline{\theta}_c)\Big) \,/\, \Big(\gamma \cos(\overline{\theta}_p) + (1-\gamma)\cos(\overline{\theta}_c)\Big)\right) \tag{4.11}$$

The weighting factor $\gamma(\lambda,\, \xi)$ depends on the confidence λ of the robot relative to idiothetic signals at time t, as well as on the discrepancy ξ between the internal and external representations of the robot heading at time t. In particular, we take

$$\gamma(\lambda, \xi) = \exp\Big(-\frac{1}{2}\Big(\Big(\frac{\lambda - \mu_\lambda}{\sigma_\lambda}\Big)^2 + \Big(\frac{\xi - \mu_\xi}{\sigma_\xi}\Big)^2\Big)\Big) \tag{4.12}$$

where $\mu_\lambda = \mu_\xi = 1$, and $\sigma_\lambda = \sigma_\xi = 0.5$. We do not model explicitly λ and ξ. Rather, we simply take λ as a function of the time step t_c at which the last calibration occurred, that is, $\lambda(t_c) = \exp(-t_c^2/2\sigma_t^2)$, with $\sigma_t = 40$. Thus, we simply assume that the confidence of the robot relative to idiothetic signals decreases non linearly with respect to the time from last calibration. Similarly, we take $\xi = \exp(-(d-\mu_d)^2/2\sigma_d^2)$, where d is the deviation between internal and external representations, and $\sigma_d = 70$.

4.3.3 Measuring the Tracking Error Over Time

In order to evaluate the advantages of using external cues to calibrate idiothetic information, we run a series of n experiments all starting at $t = 0$. At each step the robot updates its orientation randomly and the head-direction system is used to estimate its allocentric heading. The above learning scheme is employed to correlate external and internal cues on-line. In each trial i we monitor the deviation between the actual robot's heading $\theta_i(t)$ and the estimation $\overline{\theta}_i(t)$ encoded by the directional system (Eq. 4.5) over time. The mean tracking error at time t is defined as

$$e(t) = \frac{1}{n} \sum_{i=1}^{n} \mid \theta_i(t) - \overline{\theta}_i(t) \mid \tag{4.13}$$

The robot turns with a constant angular velocity $\omega \approx 100^{\circ}/sec$, and rotates, at each time step, by an angle $\Delta\theta$ drawn from $[-90^{\circ}, 90^{\circ}]$. For simplicity, the robot turns on the spot during the whole trial at a location where the necessary learning condition $\sigma_s \leq \Sigma_s$ holds. The experiment includes $n = 10$ trials, each of which consists of 300 steps.

We compare the error $e(t)$ of a non-calibrated system with the error $e_c(t)$ of a calibrated one. Furthermore, we evaluate the system's performance when using the two above calibration techniques. Fig. 4.9 (a) shows the uncalibrated mean error $e(t)$ and the calibrated error $e_{c_1}(t)$ when using the first calibration technique (Eq. 4.10). Fig. 4.9 (b) shows the uncalibrated error $e(t)$ versus the calibrated error $e_{c_2}(t)$ when using the second calibration technique (Eq. 4.11). Error $e(t)$ grows continuously because of its cumulative nature. On the contrary, calibrated errors $e_{c_1}(t)$ and $e_{c_2}(t)$ remain bounded over time. Although error $e_{c_1}(t)$ is in average slightly smaller than $e_{c_2}(t)$, the two calibration techniques yield similar results. Note that Eq. 4.10 can be directly implemented by means of a neural network, whereas Eq. 4.11 is merely an algorithmic technique. On the other hand, Eq. 4.11 accounts for two important factors (i.e., how much the agent trusts idiothetic information, and how much internal and external signals diverge from each other), which makes it suitable for studying the interaction between allothetic and idiothetic cues to control the head-direction dynamics (Chapter 5).

In Fig. 4.10 we present a result concerning the head-tracking capability of the system when using allothetic-based calibration (technique n. 1, Eq. 4.10). According to the above experimental setup, we let the robot randomly update its orientation during trials of 50 time steps. Solid lines represent the robot's nominal heading θ, whereas dashed lines represent the heading $\overline{\theta}$ as encoded by the poSC population activity. Due to the combination of inertial and visual stimuli, the directional system tracks the robot heading rather effectively over time.

4.4 Discussion

Head-direction cells play a major role in spatial learning: they provide an allocentric neural compass supporting cognitive navigation. Experiments suggest that lesions to directional cells in the postsubicular cortex seriously disrupts the performances of animals in spatial tasks [352].

We put forth a functional model to reproduce the neurophysiological properties of direction-sensitive neurons. Our approach is consistent with previous models (Chapter 3) in that it postulates the primary role of inertial stimuli in determining the dynamics of head-direction cells. However, we stress the importance of incorporating extrinsic signals into the system in order to improve the stability of the representation. This is necessary for continuously tracking the agent's direction in real sensory-environment contexts.

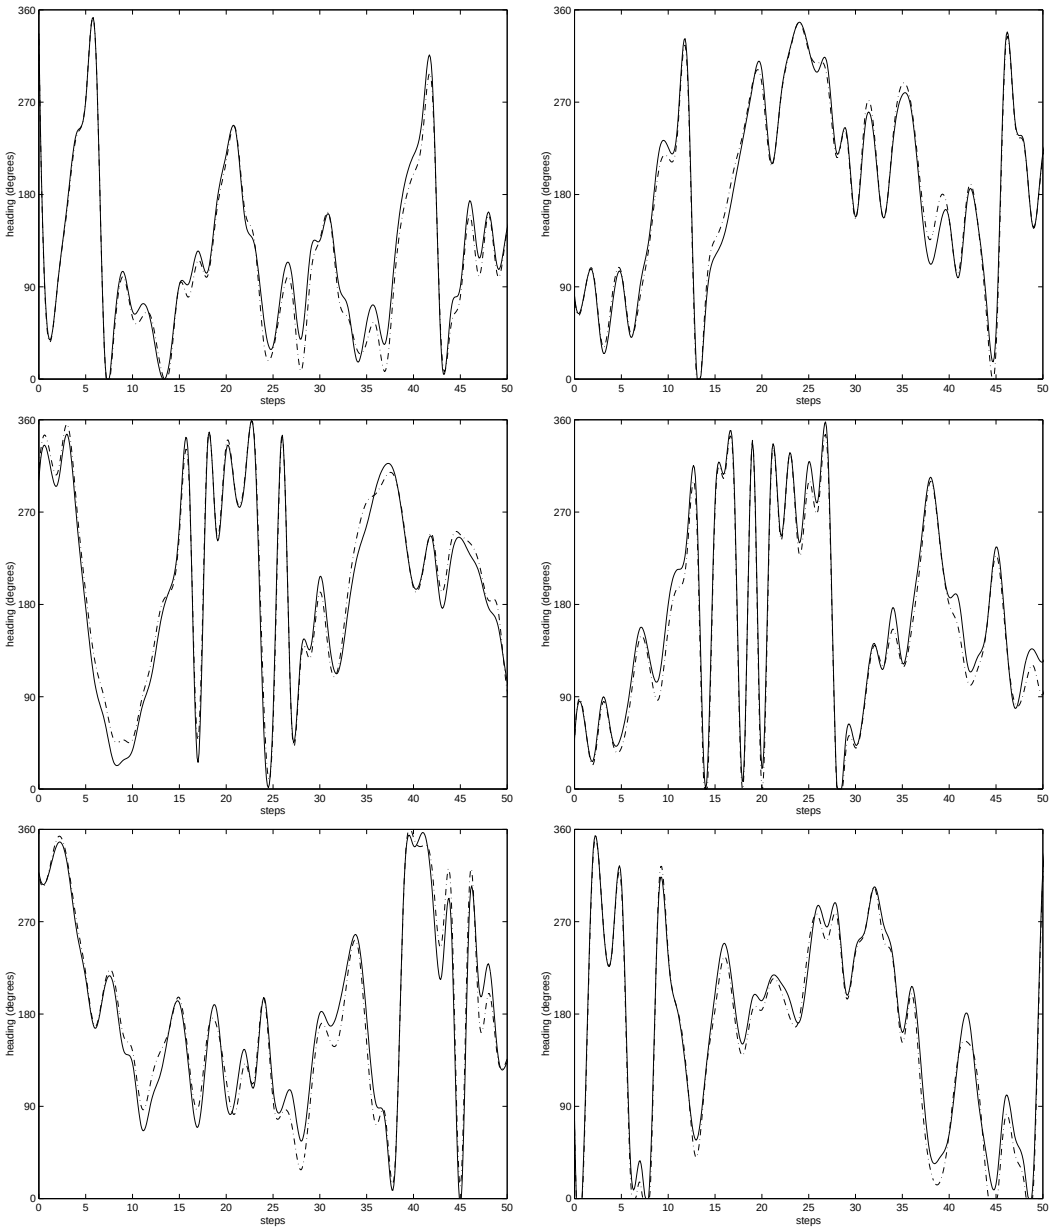

Figure 4.10: Some samples of head-direction tracking capability. Solid lines are the robot's real heading θ, whereas dashed lines are the orientations $\overline{\theta}$ as estimated by the head-direction system. The robot is let update its orientation θ at each time step by randomly turning of an angle $\Delta\theta$ uniformly drawn from $[-90^\circ, 90^\circ]$. Head-direction cell activity is used to track the allocentric heading θ over time. The angular velocity ω is constant and equal to $100\ ^\circ/sec.$ Each trial consists of 50 time steps.

Our basic directional circuit includes the anterodorsal thalamic nucleus (ADN), the lateral mammillary nuclei (LMN), and the postsubiculum (poSC), three brain regions primarily involved in the biological directional sense.

ADN and LMN form the angular velocity integrator of the model. They are continuously informed about angular displacements by a population of head angular velocity (HAV) cells, whose activity is correlated to the sign and the magnitude of angular velocity. LMN and ADN cells interact with each other to integrate HAV cell activity over time. LMN provides ADN with the current heading, whereas

ADN informs LMN cells about future headings. Similar to Blair [37] and Redish, Elga, and Touretzky [284], we emphasize the anticipatory property of ADN cells. However, in contrast to those approaches, we do not consider the anticipatory time-delay as invariant over the ADN population. Rather, according to the abstract hypothesis formulated by Blair and Sharp [41], we consider a mechanism in which each preferred direction is encoded by a cluster of ADN neurons having different anticipatory time-delays (which is consistent with recent experimental findings [39]). As a consequence, each LMN cell is informed about arrival at its preferred direction by a sequential activation of its ADN afferents. A LMN cell fires only when all its ADN inputs have been sequentially activated.

Cells in poSC form the output of our directional system. They are indirectly driven by idiothetic stimuli via ADN cells. In contrast to [37, 284], our poSC cells do not drive ADN cells directly. Rather, poSC neurons are responsible for incorporating allothetic information. They are related to external directional cues via a hypothetical population of neurons termed calibration (CAL) cells. CAL cell activity combines egocentric bearing information with allocentric place coding information to calibrate the poSC cell activity. The calibration signal is then propagated towards the ADN-LMN velocity integrator via poSC $\rightarrow$ LMN projections.

A consequence of the architecture of the model is that lesions to the poSC region do not disrupt ADN directional tunings; rather, they make them insensitive to external cues. Lesions to our ADN cells, instead, destroy poSC direction selectivity. Impairing HAV cells (which encode vestibular signals) disrupts ADN activity. These data are consistent with anatomical lesion findings reported in Sec. 2.3. Experimental results obtained by validating the system on a real robot show effective head-tracking capabilities over time. The model induces plausible Gaussian-shaped directional tuning curves.

As described in Sec. 2.2, the ensemble activity of biological head-direction cells always rotates consistently. For each pair of directional cells, the difference between their preferred directions is a constant across different environments and despite animal disorientation or cue-based system rotation. The dynamics of our directional cells is consistent with this property (note that we do not use an explicit attractor network scheme as [327, 405, 284]). If an external cue induces a shift in the population activity, the system rotates rigidly as a whole, maintaining the mutual deviations between preferred directions unchanged.

Some predictions may be formulated based on the model: *(i)* Lesions to LMN cells would disrupt the entire directional system; *(ii)* Lesions to poSC cells would not impair ADN directional firings, rather would make ADN cells insensitive to allothetic calibration; *(iii)* Lesions to either ADN cells or HAV cells would disrupt LMN direction selectivity; *(iv)* "Calibration cells", whose activity encodes both visual directional reference and place coding information, would exist; *(v)* Lesions to entorhinal place cells (in particular, to the superficial layer of the lateral entorhinal cortex) would prevent the directional system from learning relationships between local-views and headings effectively; *(vi)* As initially suggested by Blair and Sharp (1995) [41], different ADN anticipatory time-delays would provide a "safety mechanism" for integrating angular velocity.

A limitation of the model is the fact that it does not account for recent experimental results revealing double-peaked ADN firing curves [40]. Another limitation is the system's dependence on the initial state. Indeed, head-direction activity is initialized according to a Gaussian distribution with respect to an arbitrary absolute direction. That is, an initial stable state is assumed and it does not emerge as an intrinsic property of the system. A future extension of the model will consist of adding lateral center/surround connections (similar to an attractor network scheme) in order to allow the system to *(i)* settle to a stable state given any initial random activity, *(ii)* cope with intrinsic noise (in the current model, synaptic connections are assumed to be noise-free).

Chapter 5

Studying the Interaction between Allothetic and Idiothetic Cues to Control Head-direction Cells

An important issue, addressed by a large body of experimental work [103, 210, 229, 397, 350, 130, 166, 390, 102, 168, 78], concerns the way external and internal signals interact with each other to optimize spatial behavior. Experiments on insects and rodents have studied the interrelation between visual and inertial cues by considering either conflict conditions or mutual confirmation situations [98]. Even if sometimes experimental findings are inconsistent (e.g., [130, 229] *vs* [397, 70]), two principal results have received a general consensus: *(i)* The influence of a visual cue does not depend on a fixed parameter, but reflects the confidence of the animal with respect to this type of information [166]; *(ii)* For small divergences between visual cues and egomotion, animals tend to shift their internal representation according to external references. Alternatively, for large deviations, they tend to rely on their inertial frame of reference [207, 98].

5.1 Learning Landmark Stability

In the previous Chapter (Sec. 4.3), we have employed an external landmark (a weak light source) to calibrate the directional representation. Correlational LTP learning has been applied to incorporate allothetic information into the system, and to keep the tracking error bounded over time (Figs. 4.9 and 4.10). These experiments rely on the assumption that the light cue provides a stable reference and exerts a strong control over the head-direction dynamics. Here, we address the question of how the influence of the light cue may vary as a function of its stability during robot training.

5.1.1 Experiment n. 1

Given the same experimental setup as in Chapter 4 (i.e., a square arena with a weak light source L on one of its walls), we run a series of 15 training trials, of 100 steps each. At the beginning of a trial, the robot is placed at the center of the arena (the nest), and its orientation is initialized randomly. At each step, the robot rotates by a random angle from $[-90^\circ,\ 90^\circ]$, and uses the head-direction system to track

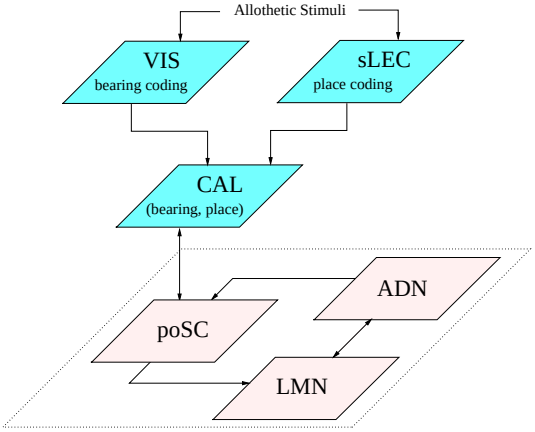

Figure 5.1: The calibration system combines two types of allothetic signals: The bearing information provided by visual cells (VIS), and the place coding information encoded by place cells in the superficial lateral entorhinal cortex (sLEC). These two informations are combined to drive calibration cell (CAL) activity which is used to calibrate poSC directional cells.

its heading based on idiothetic signals only. Internal and external cues are correlated by updating VIS $\rightarrow$ CAL connections (Fig. 5.1) according to

$$\Delta w_{cv}^{nr_n} = \nu \, r_c \, r_v \, (1 - w_{cv}^{nr_n}) \tag{5.1}$$

where nr_n is a "special" neuroreceptor for VIS $\rightarrow$ CAL connections when the robot is at the nest location[1], $w_{cv}^{nr_n}$ is the synaptic weight from cell $v \in$ VIS to cell $c \in$ CAL, and r_v, r_c are their firing activities, respectively. The learning rate ν is taken equal to 0.4.

We want to evaluate the strength of the correlation between extrinsic and intrinsic signals with respect to the stability of landmark L during training. Each training trial is characterized by the number of times the light L is relocated during that session. In particular, if m denotes the trial number, we relocate the light source $m - 1$ times. We measure the stability s_L of L by

$$s_L = 1 - \frac{m-1}{M-1} \tag{5.2}$$

where $M = 15$ is the total number of trials. As a consequence, $s_L = 1$ in the first trial, while $s_L = 0$ when $m = M = 15$.

The external cue L is relocated by modifying its orientation φ_L relative to an arbitrary frame of reference whose origin is the center of the arena. The light is always relocated on one of the walls of the arena. Let $\Delta\varphi_L$ indicate the angle by which the light cue is rotated, and let t_L be the time at which relocation occurs. Both $\Delta\varphi_L$ and t_L are randomly selected during a session.

After each 100-step training trial m, we measure the strength S_m of the correlation between external and internal representations by counting the number of CAL cells that have been strongly associated to at least one VIS cell:

$$S_m = \sum_{c \in CAL} \mathcal{H}\Big(\sum_{v \in VIS} \mathcal{H}(w_{cv} - \varrho) - 1 \Big) \tag{5.3}$$

where $\mathcal{H}$ is the Heaviside function, and $\varrho = 0.8$ is a threshold used to define a strong connection. Fig. 5.2 (a) shows $\overline{S}_m$ (obtained by averaging S_m over 10 experiments) as a function of the light cue stability s_L. As expected, the less stable L, the lower $\overline{S}_m$. That is, correlational learning (Eq. 5.1) allows the system to incorporate allothetic information only if the robot "trusts it" based on its own experience.

[1]Since, for simplicity, we let the robot rotate on the nest spot, we do not account for sLEC $\rightarrow$ CAL projections.

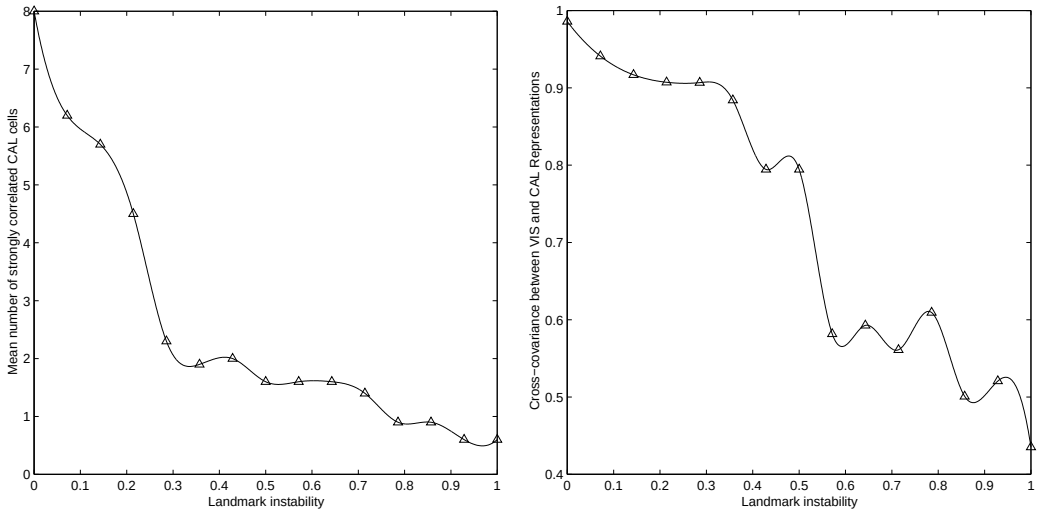

(a) (b)

Figure 5.2: (a) This diagram shows $\overline{S}_m$ (i.e., the mean number of CAL cells that have been strongly correlated to at least one VIS cell) as a function of the landmark L instability $1 - s_L$. (b) The mean cross-covariance between $\theta_c(t)$ and $\theta_v(t)$ representation as a function of $1 - s_L$. Both (a) and (b) diagrams suggest that the less stable L, the smaller its influence upon the directional system.

5.1.2 Experiment n. 2

In this experiment, we employ the same training technique used in Experiment 1 to establish VIS $\rightarrow$ CAL projections. After training, we run a series of 15 test trials. During the first trial, we employ the VIS $\rightarrow$ CAL projections learned during the first training trial; during the second trial, we employ the VIS $\rightarrow$ CAL projections established during the second training trial; and so on. At the beginning of each test trial the robot is located at the nest position and given a random orientation. Then, it makes a complete turn on the spot by steps of 4°. At each step, CAL cells are driven by VIS cells, and fire according to

$$r_c(t) = \begin{cases} 0 & \exists\, c^* : r_{c^*}(t) > r_c(t) \\ \Big(1 + \exp\big(-\beta(\sum_j r_j(t) w_{cj}^{nr_n} \; - \; \frac{1}{2}\sum_j r_j)\big)\Big)^{-1} & otherwise \end{cases} \tag{5.4}$$

where j varies over VIS cells, $w_{cj}^{nr_n}$ represents afferents to receptor nr_n, and β is taken equal to 1.

Let $\theta_c(t)$ and $\theta_v(t)$ be the angles encoded by CAL and VIS ensemble activities at step t, respectively. $\theta_c(t)$ and $\theta_v(t)$ are computed by population coding (Eq. 4.5). In order to evaluate the influence of landmark L over our head-direction system, we look at the correlation between $\theta_c(t)$ and $\theta_v(t)$. Fig. 5.2 (b) displays the mean cross-covariance between $\theta_c(t)$ and $\theta_v(t)$ as a function of the stability of landmark L. The diagram has been obtained by averaging over 10 test experiments. The curve can be interpreted as the strength of the control of external stimuli over the system. Again, the less stable L, the less correlated $\theta_c(t)$ and $\theta_v(t)$, that is, the less L influences the system.

5.1.3 Experiment n. 3

For this experiment, we take inspiration from neurophysiological tests done by Knierim *et al.* [166], and we attempt to reproduce equivalent results with our robotic implementation.

Knierim *et al.* trained two groups of rats to search for randomly-tossed chocolate pellets in a high-walled cylindrical arena. A white card, covering approximately 90° of the wall, provided the only avail-

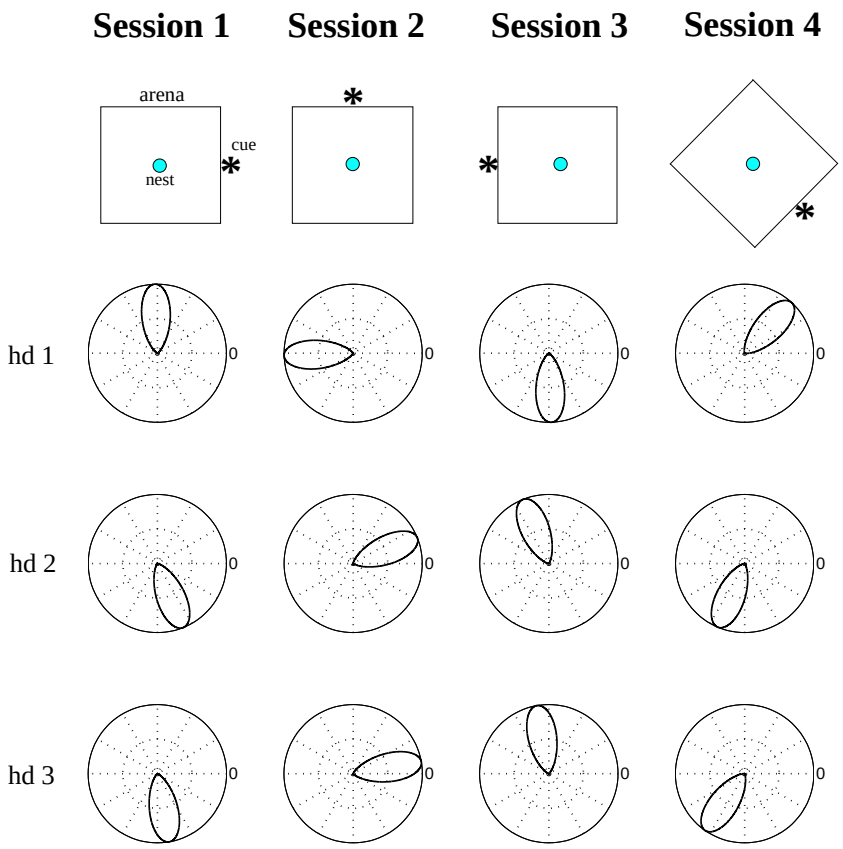

Figure 5.3: Head-direction cells recorded from the robot trained under non disorientation conditions. At the beginning of each trial, the robot is disoriented and then placed at the nest location. Tuning curves show that head-direction cells rotate their preferred direction to follow the external reference. This implies a strong influence of allothetic signals over the robot's directional system.

able external cue. Rats had 2 training sessions of 15 minutes per day for $1-4$ weeks. Some animals underwent disorientation before each training trial: the aim being to disrupt their internal directional sense before training and yield an inconsistent relationship between the cue card and the animal's inertial system. Rats from another group were not disoriented before training: the objective being to maintain a perceptual consistency between the experimental arena and the external world, such that animals could perceive the cue card always in the same location, i.e., as a stable visual reference. After training, rats underwent several recording trials of $10-15$ minutes each. All animals were disoriented before each recording trial. The cue card was either replaced or maintained at the initial location before each trial.

Given our experimental setup, we run 2 series of 10 training sessions each. Each session is about 15 minutes long. At the beginning of each trial the robot is located at the center of the arena (i.e., the nest). Before starting exploration, the robot makes a complete turn on the nest spot by steps of 4° each. Let nr_n be the neuroreceptor deserved to VIS $\rightarrow$ CAL connections when the robot is at the nest location. At each step the robot applies Hebbian learning (Eq. 5.1) to update the strength of VIS $\rightarrow$ CAL synapses arriving to nr_n. We take a learning rate ν equal to 0.2. After that a 360° turn has been completed, the robot starts exploring the arena randomly[2].

In the first training series, the robot is not "disoriented" at the beginning of each trial: its idiothetic

[2]During this phase, spatial learning takes place to create a population of place cells (see Part III).

directional system is not disrupted before entering the arena. Thus, the robot's inertial frame of reference remains consistent across different training trials. Alternatively, in the second training series the robot is "disoriented" before each trial: its idiothetic directional representation is initialized randomly before entering the arena. As a consequence, with respect to the robot's inertial system, the external light is not a stable cue across different sessions.

After training, we run 2 series of 4 recording sessions each. In the first series, we record head-direction cells from the robot that did not undergo disorientation before training. In the second, we take the robot that has been disoriented before training. At the beginning of each recording session the robot is disoriented (by random initializing its directional sense) and entered the arena at the nest location. Then, the robot tries to orient itself by using the external light as a polarizing cue. Thus, it starts rotating on the spot by steps of 4° and uses the learned VIS $\rightarrow$ CAL connections to determine the activity of every cell $c \in$ CAL according to Eq. 5.4.

If there exists a maximally active cell c (i.e., $\exists c \in \text{CAL} : r_c(t) \geq \varepsilon$), then the robot's directional system is reset according to the heading encoded by c. Otherwise, the robot cannot polarize its directional cells based on external information, and maintains its directional sense as randomly initialized at the beginning of the session.

Fig. 5.3 shows the effects of non disorienting training. The tuning curves of three poSC cells recorded over the four sessions of the first recording series are shown. The light cue is located east, north, west, and south-east, at the beginning of each trial, respectively. Results show that the light source has a strong control over the directional system. Indeed, the preferred direction of head-direction cells tends to rotate to follow the external reference consistently. Note that all cells shift their preferred direction such that the mutual deviation between their preferred headings remains constant.

Fig. 5.4 shows the tuning curves of three poSC cells recorded from the robot trained under disorientation conditions. Head-direction cells hd_1, hd_2, and hd_3 are not controlled by the light cue. Despite the fact the light maintains its east location across the first three sessions, the cells' preferred directions vary as a function of their random initialization. That is, the robot fails to polarize its directional system relative to the light cue after being disoriented before recording. Interestingly, between the third and the fourth session the firing directions of cells hd_1, hd_2, and hd_3 do not change, while the light location does. Results shown in Figs. 5.3 and 5.4 are consistent with experimental findings on rodents reported by Knierim *et al.* [166].

5.2 Studying Conflict Situations

In order to further investigate the relative weight of allothetic and idiothetic stimuli, we study their mutual interaction in conflict conditions. We take inspiration from experiments on rodents by Etienne *et al.* [103]. They examined the homing behavior of golden hamsters, as a function of self-motion and visual cues, during hoarding trips within a circular arena. Hamsters were guided as directly as possible from the nest (a box located at a fixed peripheral position) towards a feeding location in the center of the arena. After uptaking of food[3], animals had to return home autonomously. A unique visual cue was made available, i.e., a weak light source. Other allothetic cues (e.g., tactile and olfactory stimuli) were eliminated. Different experiments investigated different conflict situations induced by either shifting the light source from its standard position, or disorienting the animal during food uptake. As a general result, Etienne *et al.* found that the larger the conflict between visual and self-motion cues, the smaller the influence of allothetic information [103]. In particular, they showed that in $\pm 90^\circ$ conflict situations, hamsters tend to rely mostly on external cues, while they give higher priority to internal signals when the

[3]When hoarding food, hamsters tend to turn around the feeder location. Collecting food takes up to 30 seconds depending on subjects.

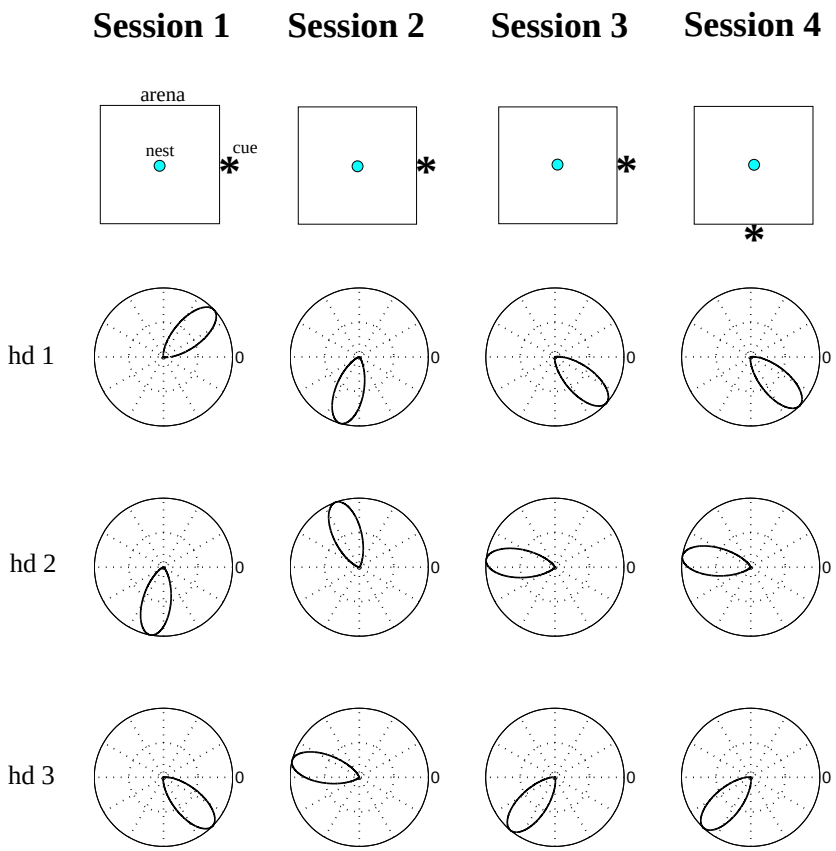

Figure 5.4: Effects of disorienting training on head-direction cells. The robot fails to calibrate its directional system with respect to the allothetic cue. The light cue is perceived such as an unstable landmark and does not influence directional cells.

conflict increases ($180°$). Similar results have been reported by later studies on rats [211, 167].

5.2.1 Experiment n. 1

Given our experimental arena, we take the peripheral east location as robot's nest, and the center of the arena as feeder position. We evaluate the robot's homing behavior when a conflict occurs between its self-generated homing vector and an *unstable* light L [103]. The task consists of: *(i)* An outward journey during which the robot is "guided" from the nest to the feeder[4]. *(ii)* A "hoarding" phase: We simulate the hamster hoarding behavior by letting the robot rotate randomly at the feeder location for t_h steps, with t_h drawn from $[10,\ 30]$ uniformly. At each step, the robot rotates by an angle from $[-90°,\ 90°]$. *(iii)* A backward journey: The robot uses its internal homing vector to return to the nest.

We run 10 training trials during each of which the robot performs one hoarding excursion. For simplicity, we let the robot correlate external and internal representations (Eq. 5.1) only during hoarding, that is when it executes its rotatory behavior at the arena center. In order to provide the robot with an unstable directional landmark, at the beginning of each training trial the location of L is shifted randomly. After training, we run a series of 10 test trials. At the beginning of each test, L is located at the location opposite the nest (i.e., west peripheral position). Then, at the end of the outward journey, L is shifted by $90°$ clockwise, which creates a conflict between internal and external representations.

Fig. 5.5 (a) reports the robot's homing behavior as recorded during test. Dots represent the peripheral

[4]Note that, the nest-to-food journey is not exactly the same from trial to trial. Indeed, the robot's trajectory around the optimal direct path changes randomly at each excursion.

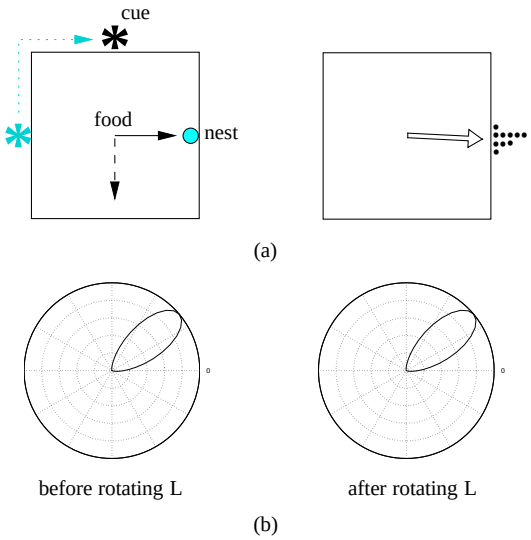

Figure 5.5: (a) Left: L (grey-black star) is an unstable cue during training. During test, L is rotated by 90° clockwise at the end of the journey towards the feeder (center of the arena). In order to return home (grey circle), the robot might follow either its idiothetic representation (solid arrow), or its allothetic representation (dashed arrow). Right: The white arrow represents the average homing behavior over 10 testing trials (dots). (b) Left: A sample of poSC cell activity recorded before cue rotation. Right: the same poSC cell recorded after cue rotation. The preferred direction has not been influenced by L.

positions reached by the robot at the end of the backward journey. The white arrow indicates the average homing vector. The mean deviation relative to the nest is about -3°, while the mean dispersion around it is about 4°. Fig. 5.5 (b) shows the preferred direction of a cell $p \in$ poSC as recorded before and after cue rotation. As a result, the robot does not rely on landmark L to infer its homing orientation. This is consistent with our previous results in that an unstable external cue tends to have a slight influence over the robot's directional system.

5.2.2 Experiment n. 2

Given the same experimental task, we create a conflict between a *stable* landmark L and self-motion information [103]. We train the robot through 10 training sessions by keeping the light L at a fixed location (west). The aim is to yield a strong correlation between the idiothetic representation and the information provided by a familiar external cue. In this case, the robot should return home in an opposite direction with respect to L.

Thereafter, we run 3 test series of 10 trials each. During test, L is turned off during the outward journey, and switched on during hoarding. In the first test series, L is presented to the robot shifted by 90° clockwise with respect to its standard position (Fig. 5.6 (a1)). In the second series, L is shifted 90° counterclockwise (Fig. 5.6 (a2)). Finally, in the third series, L is rotated by 180° (Fig. 5.6 (a3)).

Results in Figs. 5.6 (a1, a2) suggest a strong control exerted by L: the average homing vector (white arrow) deviates from the nest direction by reflecting the effect of the directional cue L. In the first test series (a1), we obtain a mean deviation of about -63° with a mean dispersion of about 16° around the mean. In the second series (a2), the mean deviation is about 60° with a mean dispersion of about 15°. Fig. 5.6 (a3) reveals an inverted order of priority between internal and external signals when L is shifted by 180°: the robot, in average, returns home by following its self-generated homing direction. The average deviation relative to the nest is about 10° with a mean dispersion of about 12°. Thus, the familiar directional cue L influences the robot's behavior when a $\pm 90^\circ$ conflict occurs. On the contrary, the influence of L vanishes when the conflict is further increased (180°).

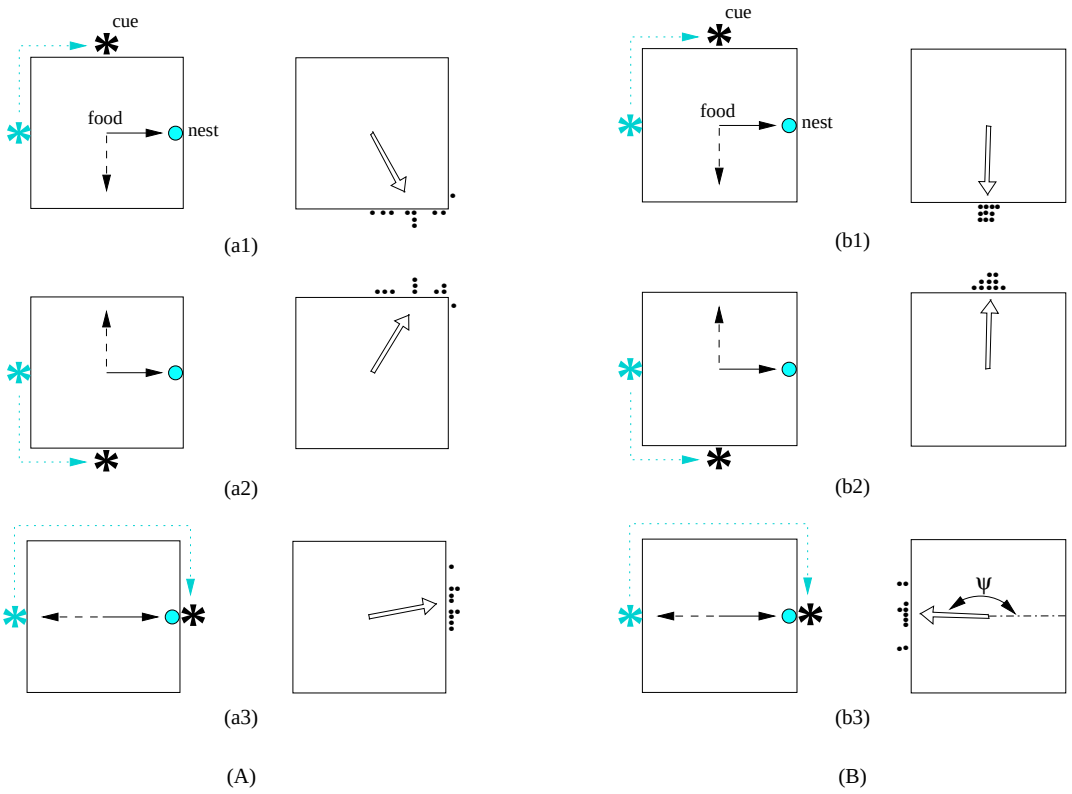

Figure 5.6: (A) Light L (grey-black star) remains stable during training. During test, L is switched off when the robot is moving towards the food location (outward journey). During hoarding, L is presented to the robot shifted by 90° clockwise (a1), 90° counterclockwise (a2), and 180° (a3). The white arrow indicates the resulting mean homing behavior of the robot. L has a strong control in (a1) and (a2), whereas it has a slight influence in (a3). (B) In this experiment, the robot undergoes disorientation as soon as it reaches the feeder location. This induces a strong predominance of the external cue L when shifted by 90° clockwise (b1), 90° counterclockwise (b2), and 180° (b3).

5.2.3 Experiment n. 3

In the previous experiment, homing strongly depends on external stimuli for $\pm 90^\circ$ rotations of L. Nevertheless, self-generated information still affects the robot's trajectory remarkably (Figs. 5.6 (a1, a2)). In this experiment, we ask the question: How does the robot behave in conflict situations when its egomotion representation has been partially disrupted? Or, to put it differently: How does the robot weight external and internal conflicting information when it does not trust egomotion signals?

As in the previous experiment, we let the robot solve the hoarding task during 10 training sessions. During learning, L is kept at its standard west location. After training, we run 2 series of 10 test trials each. The light is switched off during the outgoing journey. As soon as the robot reaches the feeder location, it undergoes "disorientation": At each of d steps, with d drawn from $[50,\ 100]$ uniformly, the robot rotates by $\Delta\theta$ degrees, with $\Delta\theta$ drawn from $[-90^\circ,\ 90^\circ]$[5]. After disorientation, the light L is turned on. In the first test series, L is presented to the robot shifted by 90° clockwise (Fig. 5.6 (b1)). In the second, L is rotated by 90° counterclockwise (Fig. 5.6 (b2)).

As expected, experimental results suggest that the familiar landmark L exerts a stronger control when idiothetic signals have been disrupted through disorientation. In the first series, we obtain a mean

[5]We call this phase "disorientation" because, due to the large number of rotations, the confidence of the robot relative to idiothetic signals (i.e., λ factor in Eq. 4.11) decreases.

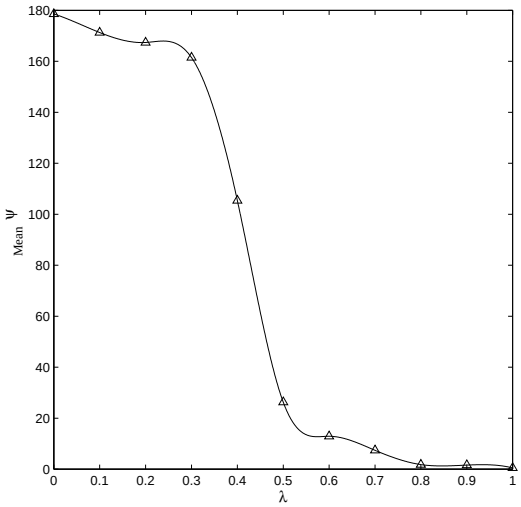

Figure 5.7: The response of the robot to stable allothetic cues varies as a function of its confidence λ with respect to idiothetic signals. This diagram shows the deviation $\overline{\Psi}$ (averaged over 10 trials) between the actual and the self-generated homing vector, for λ varying within $[0,\ 1]$, in the case of 180° conflict situation.

homing direction of about -91° with a standard deviation of 3.3°. In the second series, the average homing direction is about 90°, with standard deviation 6°.

5.2.4 Experiment n. 3a

Results obtained in Experiments 1, 2, and 3 are consistent with those reported by Etienne *et al.* when testing hamsters in similar conflict conditions [103]. In this last experiment, we investigate the 180° conflict situation in the case of a disoriented agent (not reported by [103]). The aim is to show that, even for large divergences between internal and external cues, a stable landmark L may influence the homing behavior when the robot does not trust egomotion.

The result in Fig. 5.6 (b3) confirms this hypothesis: the mean homing direction reflects a predominance of the external directional information (in contrast to Fig. 5.6 (a3)). We obtain a mean homing direction which deviates from the nest of about 178°, with a mean dispersion of about 15°. Let Ψ be the divergence between the actual homing direction and the self-generated homing vector. We argue that Ψ is a function of the robot's confidence about its idiothetic representation. Fig. 5.7 shows a result obtained by measuring the mean deviation $\overline{\Psi}$ (over 10 trials) when varying λ within $[0,\ 1]$.

5.3 Discussion

In this chapter we have studied the interrelation between extrinsic and intrinsic stimuli to control the directional allocentric representation provided by head-direction cells. In particular, we have addressed two questions: How does the control exerted by an external landmark reflect the stability of the cue [166]? How do animals assign an order of priority to allothetic and idiothetic information in conflict situations [103]?

To answer the first question we have carried out a series of experiments showing that, by adopting LTP correlational learning (Eq. 5.1), only stable external cues (i.e., landmarks that remain sufficiently steady during training) will exert enough control over our directional system. Indeed, an unstable external signal would result in weak correlations between allothetic and idiothetic representations, and

would slightly affect the head-direction dynamics. This is consistent with neurophysiological findings on rodents [350, 166].

To address the second question we have studied the homing behavior of the robot as a function of conflicting external and internal references. The two categories of signals are combined linearly in order to control the head-direction dynamics (Eq. 4.11). In particular, we adopt a weighting factor $\gamma_{\lambda,\xi}(t)$ that depends on the confidence λ of the robot relative to idiothetic signals at time t, as well as the discrepancy ξ between the internal and external representation of the robot's heading at time t. Consistently with behavioral data and results from single cell recordings [103, 211, 167], our system tends to rely mainly on allothetic cues for relatively small conflicts, whereas it gives higher priority to self-motion signals when the conflict increases. Nevertheless, we argue that an important role in weighting directional signals is played by factor λ, i.e., by how much the agent trusts its egomotion signals. As a natural result, the model predicts that even in a large conflict situation (180°) a completely disoriented agent will rely on stable external landmarks to return home.

Part III

Space Representation: Place Cells

Foreword

As the complexity of the task and the perceptual capabilities of biological organisms increase, an explicit spatial representation of the environment appears to be employed to support navigation [361]. In rodents, *hippocampal place cells* exhibit such a spatial representation property [255]. This has given rise to the hypothesis that the hippocampus plays a functional role in rodent navigation, and that it provides a neural basis for cognitive spatial behavior [256].

The presence of location-sensitive cells elicits at least two eminent questions: *(i)* How is the spatial selectivity property of these neurons established (i.e., spatial learning)? *(ii)* How do animals utilize such a place field space representation to perform navigation in complex task-environment contexts (i.e., cognitive navigation)? In this part of the thesis we address the first question, that is, how a population of hippocampal place cells can be generated based on noisy sensory information. The second question will be treated in Part IV, in which hippocampally dependent goal-oriented navigation will be studied.

In the following five chapters we study the hippocampal involvement in spatial memory[6], and we put forward a computational model of its basic functional properties.

- In Chapter 6 we provide the biological background: we describe the principal brain areas involved in space coding and we review the neurophysiological properties of hippocampal place cells.
- In Chapter 7 we review some of the previous hippocampal models. We also present some results obtained by implementing one of them: the model by Burgess and colleagues (1994) [62].

Contributions: We propose a hippocampal model for spatial learning in which *(i)* allothetic and idiothetic representations are combined to establish stable place cells, *(ii)* head-direction cells (Part II) and place cells interact with each other in order to achieve spatial learning (the work described in this part of the thesis has been published in [13, 14, 16, 15]):

- In Chapter 8 we model the visual pathway of our hippocampal model. We address the issue of self-localization based on real visual stimulation.
- In Chapter 9 we present the idiothetic pathway of the model (i.e., path integration).
- In Chapter 10 we combine the two above representations in order to form place cells in the hippocampus proper of the model.

As for the head-direction system (Part II) we present results obtained by validating the model with a real mobile Khepera robot.

[6] Experimental findings, especially in humans, suggest the hippocampus is involved in a larger class of memory named episodic memory. However, this thesis focuses on the hippocampal role in space coding. Examples of studies on the hippocampal involvement in episodic memory can be found in [196, 288, 392, 263, 290, 298, 365, 289, 373, 31, 192].

Chapter 6

Biological Place Cells

A place cell is a neuron whose firing activity is primarily correlated with the animal's spatial location[1]. A place cell shows action potentials only when the animal is in a specific region of the environment, which defines the *place field* of the cell. Since the spike frequency of a place cell i encodes a specific exocentric location $\vec{p}_i$ (the center of the place field), the animal may rely on the ensemble place cell activity to achieve self-localization. In other words, the place cell population provides an allocentric space code supporting the animal's cognitive behavior. Due to this appealing property, place cells have been the subject of a huge body of research. Nowadays, one of the most studied brain regions is the area in which place cells were first discovered by O'Keefe and Dostrovsky in 1971 [255]: the hippocampus.

6.1 Anatomical Background

Location-sensitive neurons have been recorded from the following brain areas of freely-moving rats: *(i)* The hippocampus proper [255, 251, 254, 260, 209, 271, 258, 243, 240, 357, 212, 324, 394, 257, 194, 403, 166, 195, 253, 227, 132, 133, 329, 25, 216], *(ii)* the dentate gyrus [162], *(iii)* the medial entorhinal cortex [23, 278, 228], *(iii)* the subiculum [23, 244, 323, 319], and *(iv)* the parasubiculum [349]. Although neurons in all these areas exhibit similar spatial correlates, their firing does not always rely upon the same determinants, which results in different neurophysiological properties (Sec. 6.2).

The presence of place cells in the above regions suggests that these structures might form a neural circuitry to support spatial learning. To deduce their basic functionalities, it is worth characterizing their anatomical interconnections. Fig. 6.1 is a simplified representation of the mutual projections between space code related areas[2] [8, 399, 400, 397, 60]. Note that, since head-direction cells play a fundamental role in spatial learning, the regions involved in the directional representation (see Chapter 2) have been included in the scheme.

The hippocampal formation, in rodents, occupies a large volume of the corresponding cerebral hemisphere, and encompasses the thalamic area (see Appendix A for some anatomical images). In the early 1970s, the hippocampal anatomical organization was understood as a lamellar structure, formed by several transverse slices stacked along the longitudinal hippocampal axis [10]. These slices were supposed to have stereotyped cytoarchitecture and connection schemes, and to operate as independent functional

[1] In particular, to the location of the animal's head.

[2] Connection schemes in Figs. 6.1, 6.2 and 6.3, are not intended to be exhaustive.

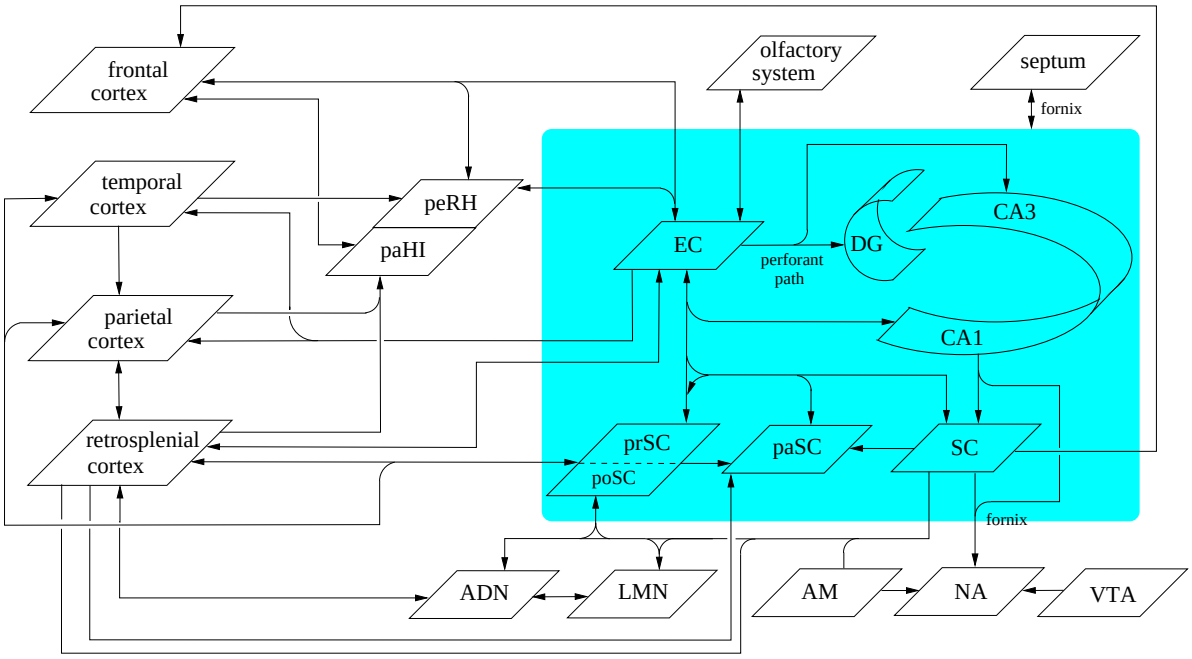

Figure 6.1: Anatomical interconnections between brain areas related to space representation (for sake of clearness, the displayed mutual projections are a subset of those reported in the text). Glossary: EC: entorhinal cortex, DG: dentate gyrus, SC: subiculum, paSC: parasubiculum, prSC: presubiculum, poSC: postsubiculum, peRH: perirhinal cortex, paHI: parahippocampal cortex (postrhinal cortex), ADN: anterodorsal thalamic nucleus, LMN: lateral mammillary nuclei, AM: amygdala, NA: nucleus accumbens, VTA: ventral tegmental area. The hippocampus proper consists of CA3-CA1 areas. The hippocampal formation (shaded area) includes the hippocampus proper, DG, EC, SC, paSC, and prSC. Adapted from Burgess *et al.* [60], Redish and Touretzky [285], Amaral and Witter [8], and Witter [397].

units. According to this theory, processing in the hippocampal formation would take place along the longitudinal axis only [10]. However, recent anatomical data have refuted the lamellar hypothesis and indicated that the hippocampal formation is a complex three-dimensional structure with functional differentiations along both longitudinal and transversal axis [8, 397].

The hippocampal formation (darkgrey area in Fig. 6.1) contains the hippocampus, the entorhinal cortex (EC), the subiculum (SC), the parasubiculum (paSC), and the presubiculum (prSC), whose dorsal part forms the postsubiculum (poSC) [8]. The hippocampus (double C-shaped structure in Fig. 6.1) includes the dentate gyrus (DG) and the hippocampus proper (or *cornu ammonis* CA). The latter consists of 4 subregions CA1-CA4, but CA1 and CA3 are the most distinguishable areas.

Hippocampal afferents. Two major inputs enter the system: *(i)* Signals from neocortical areas converge onto the entorhinal cortex which in turn projects to the hippocampus via the perforant path. These signals carry information coming from most of the unimodal and multimodal associative areas. As a consequence, the hippocampal formation is the recipient of highly processed multimodal sensory information [289]. *(ii)* Inputs from subcortical areas reach the hippocampus via the fornix fiber bundle. Signals from the thalamus, the hypothalamus, the brainstem, and the amygdala, probably concern arousal, emotional, and autonomic information [60]. The subcortical cholinergic and GABA-ergic inputs from the septal region [142] modulate the ensemble hippocampal activity, and seem to be responsible for generating the hippocampal theta rhythm [64, 219] (Sec. 6.4).

Intrinsic hippocampal circuit. Fig. 6.2 is a simplified scheme of the intrinsic hippocampal connections [8, 5, 397, 60]. In particular, we focus on the propagation of neocortical inputs through the hippocampal formation. *(i)* As a first step, the highly processed information from neocortical areas reaches

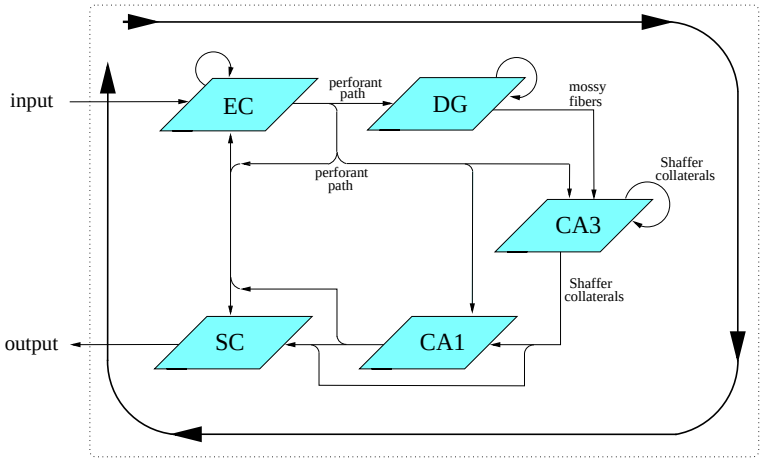

Figure 6.2: The intrinsic circuitry of the hippocampal formation (prSC and paSC are not included in this connection scheme).

the entorhinal cortex. Entorhinal cells, via the perforant path, project to DG granule cells, CA3 and CA1 cells, and subicular cells. Furthermore, EC exhibits intrinsic connections. *(ii)* The dentate gyrus sends efferents to CA3 via the mossy fibers. Synaptic projections from DG to CA3 are very selective: each granule cell projects approximately onto 14 pyramidal cells only [5]. DG has also intrinsic projections: granule cells generate collateral synapses that terminate on the polymorphic DG region [8]. *(iii)* CA3 pyramidal cells form a recurrent network through the Shaffer collaterals, but the latter fiber bundle is also used to synapse CA1 and subicular cells. *(iv)* CA1 neurons send their output to entorhinal as well as subicular cells via the angular bundle [6]. *(v)* Finally, SC projects onto the entorhinal cortex.

As indicated in Fig. 6.2 by the thick arrow, the hippocampal circuit can be coarsely approximated by a feed-forward loop [8]: information enters the loop via EC, proceeds towards DG, then to CA3 and CA1, and finally arrives to SC which closes the loop by projecting to EC.

Hippocampal efferents. The subiculum forms the major output of the hippocampal formation [8, 397] (nevertheless, CA3 and CA1 regions also send an output directly to subcortical areas, e.g., the lateral septum [344]). Until the mid-1970s it was believed that the hippocampal output was predominantly carried by the fornix. Though, recent studies have shown that an important pathway for the hippocampal outflow consists of the non-fornical projection to the deep layers of the entorhinal cortex (dEC). From dEC, information is sent to a variety of cortical areas [397, 153].

The entorhinal cortex [278, 397] is a six-layered structure divided into superficial (I,II, and III) and deep layers (IV, V, and VI)[3]. We will refer to the former as sEC, and to the latter as dEC. In rats, in which the entorhinal cortex occupies the rear end of the cerebral hemisphere, a distinction is also made between lateral and medial areas of the entorhinal cortex, LEC and MEC, respectively (Fig. 6.3).

EC cortical afferents. As previously mentioned, the entorhinal region is the main "cortical gate" of the hippocampal formation, in the sense that it receives sensory signals from the neocortex and conveys such information to the hippocampus via the perforant path [8, 278, 397]. Cortical afferents to the entorhinal cortex reach both superficial and deep areas (Fig. 6.3 (a)): *(i)* Afferents to sEC form the main input source for the information subsequently sent to the hippocampus. An eminent cortical input to sEC is via the perirhinal (peRH) and parahippocampal (paHI) cortices[4] [397]. These afferents distribute to

[3]Superficial and deep regions are separated by the lamina dissecans, a cell-poor area generally associated to layer IV [153].

[4]The parahippocampal cortex is well identified in primates, whereas a precise parahippocampal homologous has not been delineated in rodents. Nevertheless, Burwell *et al.* [63] have suggested that the primate parahippocampal cortex might be the equivalent of the postrhinal cortex in the rat.

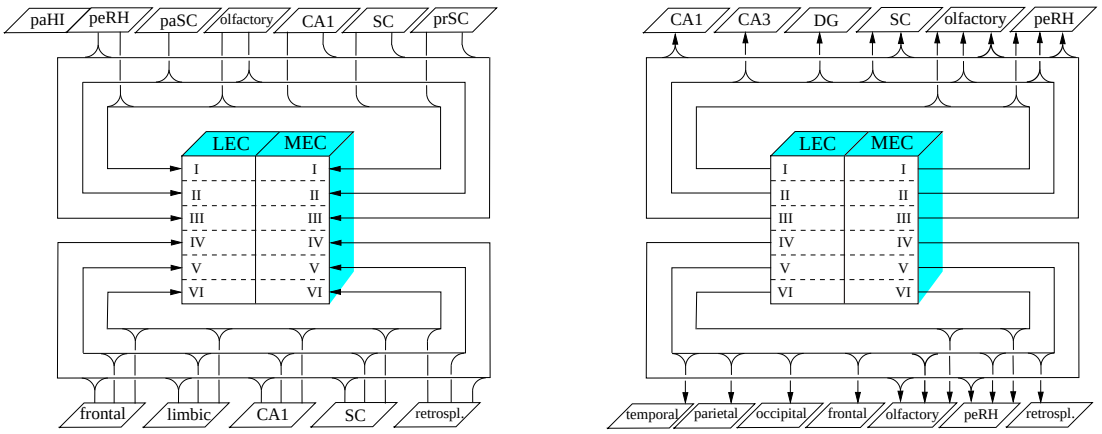

(a) (b)

Figure 6.3: A simplified connection scheme representing a subset of the afferent (a) and efferent (b) EC projections described in the text (based on results reported by [398, 8, 80, 278, 401, 397, 153]).

layers I and III of both LEC and MEC. The perirhinal and parahippocampal cortices convey information from most of the associative areas (visual, auditory, and somatosensory), as well as from the parietal, temporal, frontal, and retrosplenial lobes. Another prominent input to the rodent sEC originates from the olfactory system [175]. This input affects layers I and II of both LEC and MEC[5]. *(ii)* The deep entorhinal cortex receives cortical afferents from the limbic system [80], the retrosplenial cortex (which projects almost exclusively onto MEC) [401, 397], and from the frontal cortex [60].

EC cortical efferents [80, 153]. The entorhinal cortex projects primarily to perirhinal (peRH), infralimbic, prelimbic, orbitofrontal, and olfactory cortices. Minor EC projections reach the temporal, frontal, retrosplenial, occipital, and parietal areas. Fig. 6.3 (b) is a highly-simplified representation of the laminar segregation of EC cortical projections. As reported by Insausti *et al.* [153], neuroanatomical tracing experiments show that projections to temporal, frontal, occipital, retrosplenial, and parietal regions mainly arise from layer V of EC. Efferents to peRH originate from all layers, but predominantly from layers II and III of LEC. Olfactory structures receive inputs from all EC layers, albeit layers III and V dominate in the rostral and caudal EC parts, respectively. Finally, all EC layers project to prelimbic and infralimbic cortices, with layers V and VI being prominent for prelimbic targets, and layer III dominating the efferent source to infralimbic areas.

EC hippocampal afferents. The hippocampus proper projects to the entorhinal region via CA1 only (an early hypothesis about the existence of CA3 projections to EC [344] has been refuted by later experimental data [8]). The subiculum (SC) projects to EC as well. Concerning the laminar segregation of the hippocampal afferents to EC, it was early reported that only the deep layers of the entorhinal cortex were reached by CA1 and SC projections [344, 137]. However, subsequent studies [397] have indicated that layers I-III of EC (especially of the MEC) are innervated by an important projection from CA1 and SC (Fig. 6.3 (a)). Another important input to sEC comes from presubicular and parasubicular cortices. Presubicular afferents terminate predominantly in layers I and III of MEC, whereas parasubicular inputs mainly reach layer II of both LEC and MEC.

EC hippocampal efferents. As previously mentioned, EC projects to the hippocampus via the perforant path. In particular, the perforant projection arises from the superficial layers II and III [398, 278]. Layer II mainly innervates the dentate gyrus and CA3, but also sends collaterals to the subiculum [397]. Layer III predominantly projects to CA1 and subiculum. Electrophysiological data indicate that the di-

[5]Only the caudal region of MEC might not be reached by olfactory afferents.

rect influence of sEC upon CA3-CA1 regions is as strong as the influence relayed via DG [402, 278]. Also, EC sends a (weak) output to the presubicular and parasubicular cortices [398, 401].

EC intrinsic projections. Internal links connect the deep layers of the entorhinal cortex (dEC) to the superficial zone (sEC) [8]. In addition, the lateral entorhinal cortex (LEC) strongly projects to the medial entorhinal area (MEC) [278].

The subicular complex consists of three principal subregions, namely the subiculum (SC), the presubiculum (prSC), and the parasubiculum (paSC).

Subicular afferents. SC receives input projections predominantly from CA1 and sEC [6]. The presubiculum (prSC) receives afferents from the subiculum, the posterior parietal as well as the temporal lobes, the retrosplenial cortex, the anterodorsal thalamic nucleus, the laterodorsal thalamic nucleus, and the lateral mammillary nuclei. The parasubiculum (paSC) is reached by projections from the subiculum, the presubiculum, and the retrosplenial cortex [60].

Subicular efferents. prSC and paSC project to layers III and II of EC, respectively [8]. SC projects to both LEC and MEC, prSC, paSC, the medial prefrontal cortex, the retrosplenial cortex, the septal complex, the nucleus accumbens[6] (NA), the mammillary nuclei, the amygdala (AM), the hypothalamus, and the thalamic nuclei [8, 397].

The parietal cortex (in particular its posterior region) plays an important role in spatial cognition and it strongly interacts with the hippocampal formation. While the latter exhibits an allocentric representation of space, parietal lobes seem to be implicated in representing egocentric spatial frameworks (see [60] for a review concerning parietal-hippocampal interactions). One of the most convincing hypotheses about the parietal-hippocampal interaction suggests that information coming from different receptors (e.g., eyes) reaches the parietal lobes (via the sensory cortices) in which is represented by different abstract egocentric frameworks. Then, these multiple egocentric representations converge onto the hippocampal formation in which are translated into an allocentric spatial reference frame. According to this theory, the parietal cortex and the hippocampus would cooperate for solving spatial tasks by addressing the egocentric and allocentric components of the task, respectively [374, 60, 31, 192].

Parietal afferents. The posterior parietal cortex (PP) combines multimodal allothetic signals (e.g., visual, auditory, and somatosensory inputs) as well as multimodal idiothetic information (e.g., vestibular, proprioceptive, eye-position and eye-velocity signals) [60]. Indeed, it receives cortical inputs from sensory-motor, visual, and somatosensory areas. Furthermore, PP receives projections from the posterior cingulate cortex (retrosplenial cortex), the temporal cortex, the parasubiculum, the postsubiculum, the laterodorsal (LDN) and lateroposterior (LPN) thalamic nuclei [282]. Finally, the hippocampal formation projects onto the parietal lobe via the entorhinal cortex [397].

Parietal efferents. The posterior parietal cortex projects to the posterior cingulate cortex, the parasubiculum, the postsubiculum, and the superior colliculus [60]. Finally, the interaction between PP and the hippocampal formation occurs via the perirhinal and parahippocampal (postrhinal) cortices, which receive a direct input from PP and project to EC [63].

6.2 Neurophysiological Properties of Place Cells

Let i be a place cell whose firing activity $r_i(\vec{p})$ is a function of the animal's current position $\vec{p}$. The mean spiking frequency distribution of i forms a two-dimensional place field with a peak at the preferred location $\vec{p}_i$ and with smoothly falling off edges in all directions [255]. Fig. 6.4 shows two samples of place fields recorded from CA3-CA1 regions of a freely-moving rat.

[6] In particular, SC projects to NA via the fornix fiber bundle. Note that, also the hippocampal CA1 region contributes to this projections [282].

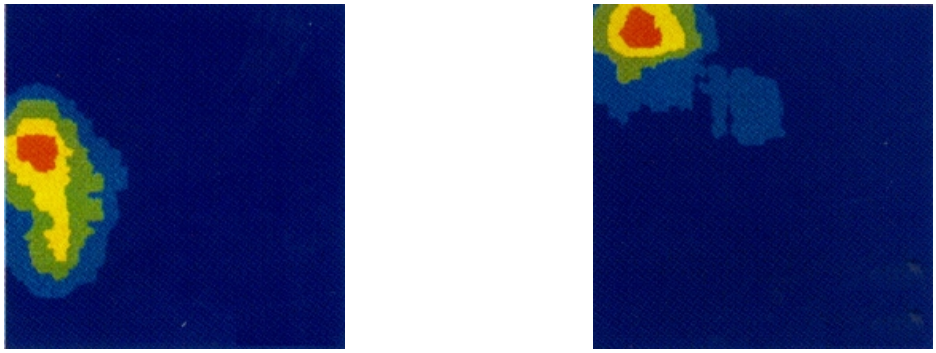

(a) (b)

Figure 6.4: Two samples of CA3 (a) and CA1 (b) place fields recorded from a freely-moving rat in a square arena. The red region indicates the area in which the cell is maximally responding. By contrast, when the animal is visiting a dark-blue region, the cell remains silent. Adapted from O'Keefe and Burgess 1996 [253].

A typical hippocampal place field encodes a region a few times the animal's size [60]. Such an imprecise single-cell coding is compensated for by the ensemble place cell representation $\mathcal{R}(t) = \{r_i(t) \,|\, i\}$. Place cells tend to cover the environment densely and uniformly, which results in highly overlapping place fields [254, 357]. As a consequence, accurate space coding may be achieved by taking into account the ensemble, rather than single-cell, firing activity. Spatial locations can be determined by appropriately averaging the population place cell activity (i.e., *population vector decoding*) [121, 243, 241, 394, 299].

A large body of behavioral, neurophysiological, and neuroanatomical studies, has been done over the last three decades in order to *(i)* understand the nature of the space coding property of hippocampal place cells (e.g., how are sensory inputs processed to produce spatial code?), *(ii)* study the influence of environmental manipulations upon the hippocampal representation, *(iii)* investigate behavioral environmental determinants, *(iv)* identify the specific functionalities of anatomical areas through lesion experiments. Fig. 6.5 shows the variety of environmental setups employed to perform behavioral and neurophysiological tests: *(i)* Water maze [233, 236, 307, 237, 385, 25], *(ii)* cylindrical and rectangular open-field arenas [224, 76, 101, 243, 240, 238, 324, 394, 314, 166, 195, 132, 133, 253, 4, 282], *(iii)* hole board circular platform [22, 24, 208], *(iv)* T, Y, and "+" mazes [251, 254, 24, 271, 258, 403, 195], *(v)* radial arm maze [220, 272, 209, 357, 229, 194, 403, 226, 195, 227], *(vi)* linear, and rectangular tracks [257, 194, 132, 329, 330, 25, 216].

All these experiments have produced a huge amount of data providing insights about spatial learning capabilities of rodents. Below, is a review of some of the most relevant properties emerged from the above studies[7].

6.2.1 CA3-CA1 Pyramidal Place Cells

Tuning shape of place fields. A typical place field can be coarsely approximated by a two-dimensional single-peak Gaussian surface [242, 60]. However, cells encoding peripheral locations show crescent-shaped fields hugging the arena walls [243]. Place cells may also exhibit multi-peak fields within a single environment[8] [254, 209, 258, 281, 394].

Rapid learning of place fields. Establishing a place field representation in a novel environment takes a relatively short time. Early recording experiments [147] showed that no experience is necessary for tuning up place cell firing: most of the observed cells exhibited place coding as robust on the first visit

[7] For an experiment-by-experiment review, see [282]. For a review of early experiments on rodents, refer to [256, 114].
[8] However, multi-peak place fields have been observed only in a very small percentage of cases.

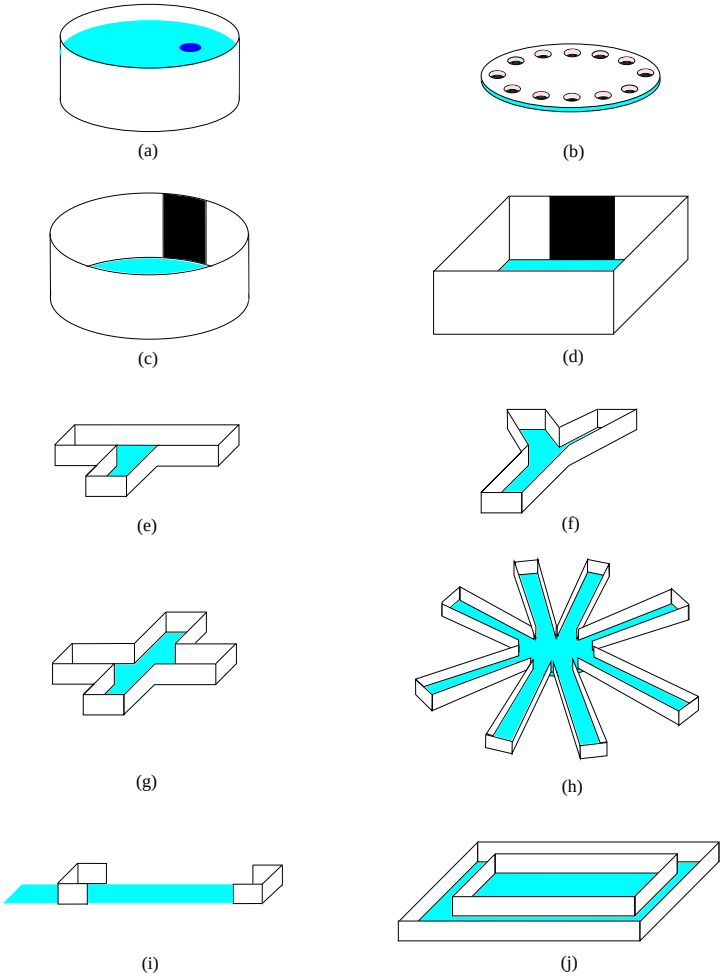

Figure 6.5: Experimental mazes: (a) Morris's water maze, (b) hole board circular platform, (c, d) cylindrical and rectangular enclosed arenas, (e, f, g) T, Y, and "+" mazes, (h) radial 8-arm maze, (i, j) linear and rectangular tracks.

to their field as on subsequent visits. Nonetheless, later experiments [394, 18] suggest that at least 10 to 30 minutes of exploration are necessary in order to generate stable place fields. Also, some cells seem to develop their firing fields more quickly (a few minutes) than others (30 minutes) [346].

Experience-dependent place field reshaping. Recent neurophysiological findings [215, 216] show that, as the animal experiences several times a route, CA1 cells tend to asymmetrically expand their (initially symmetric) field and to shift their field center backwards with respect to the rat's direction of motion. Such an asymmetric expansion property does not persist: *(i)* Across different environments. Let i be a CA1 cell whose place field has been "expanded" in an experienced route. If i will exhibit a place field in a novel route, the field will be initially symmetric and it will re-expand based on the experience in the new environment. *(ii)* Across different sessions in the same apparatus. The effects of place field expansion disappear at the begin of each daily recording session.

Place field directionality. Experimental data show that place cells have non-directional place fields (i.e., their firing activity does not depend on head direction) when the animal randomly moves over two-dimensional open environments [209, 239, 195, 253]. On the other hand, data show that place cells have directional place fields when the rat moves along fixed trajectories, such as in linear track mazes [132, 215], rectangular track mazes [216], radial narrow arm mazes [209], "+" mazes [195], and fixed paths within open field environments [195].

Place field distribution and absence of neural topology. CA3-CA1 firing fields tend to cover the whole environment uniformly, without differentiating areas with respect to their potential relevance (e.g., feeder location) [254, 394]. Experimental findings show that a given place cell i can have place fields in several but not all environments [177]. In particular, results suggest that i has a mean probability of 0.3 to exhibit at least one place field in a given arena [357]. As a consequence, a very large number of cells participate, in average, to the representation of an environment. This results in a dense population of highly overlapping place fields [367]. Importantly, CA3-CA1 place cells are not topographically organized, in the sense that there is no relationship between the physical place field topology and the anatomical place cell arrangement[9] [254, 177, 357, 240]. Thus, two cells i and j coding for neighboring locations $\vec{p}_i$ and $\vec{p}_j$, respectively, are not necessarily anatomically adjacent.

Uncorrelation between place field topologies across environments. Experiments involving single-cell recordings in different environments show that the spatial relationships between place cells are not preserved across environments [177, 243, 240, 357, 278]. Let i and j be two place cells encoding two neighboring locations $\vec{p}_i^{\,1}$ and $\vec{p}_j^{\,1}$, respectively, in the first environment. If i and j are also active in the second environment and code for positions $\vec{p}_i^{\,2}$ and $\vec{p}_j^{\,2}$, respectively, in general: *(i)* $\vec{p}_i^{\,1} \neq \vec{p}_i^{\,2}$ and $\vec{p}_j^{\,1} \neq \vec{p}_j^{\,2}$, *(ii)* $\vec{p}_i^{\,2}$ and $\vec{p}_j^{\,2}$ are not neighboring locations.

Place code replaying during sleep. Recording hippocampal cells during sleep shows that *(i)* cells that have been activated during the last session are more active during subsequent sleep episodes than others [266], *(ii)* cells with overlapping place fields during the last session (i.e., cells with temporally correlated firing) exhibit correlated reactivation during sleep [395, 329].

Differences between dorsal and ventral hippocampus. Although, place cells have been mostly recorded from the dorsal hippocampus, location sensitive neurons have also been found in the ventral (temporal) hippocampus [275]. However, spatial selectivity seems to occur more strongly and frequently in the dorsal area than in the ventral region of the hippocampus [163]. Also, typical ventral place fields are larger than receptive fields recorded dorsally [207].

6.2.2 Dentate Gyrus' Place Cells

Neurophysiological experiments with the radial arm maze show that DG's granule cells show clean location selectivity [162]. Similar to pyramidal cells, granule neurons exhibit directional place fields with the radial arm maze, and there exists an important relation between the DG cell firing and the theta rhythm [330] (Sec. 6.4).

6.2.3 Subicular Place Cells

In contrast to CA3-CA1 place cells, SC neurons have the property of maintaining a similar place field topology across distinct environments: *(i)* When recorded in two geometrically different arenas (e.g., a square box and a cylinder) providing diverse visual cues (e.g., different wall colors), a typical cell $i \in$ SC exhibits similar firing patterns (i.e., field location, firing rate, and field size) in the two recording chambers [323, 319, 321]. For instance, if i codes for a location $\vec{p}_i^{\,1}$ near the east wall of the cylinder, i will also respond maximally when the animal is at a location $\vec{p}_i^{\,2}$ near to the east wall of the square box. *(ii)* When recorded in two square environments of different size, a cell $i \in$ SC tends to show similar overall firing patterns such that its place field tends to expand or shrink to fit the size of the current environment [321].

SC cells have broader receptive fields than place cells in the hippocampus proper. Moreover, SC cells show directional tuning even in situations in which CA3-CA1 do not (e.g., open-field spatial tasks) [323]. Also, subicular cell activity seems to be strongly modulated by self-motion information [207].

[9] As reported in Sec. 2.2, head-direction cells have the same property [353, 348, 318, 405].

6.2.4 Entorhinal Place Cells

Available neurophysiological data on spatially selective cells in EC mainly concern the superficial layers of the medial entorhinal cortex (MEC) [278]. In contrast to CA3-CA1 place cells (which have a very low firing rate outside their firing fields), entorhinal cells tend to be noisier and to fire relatively rapidly over the entire environment. Nevertheless, their firing activity is maximal only at a single location. Like CA3-CA1 representation, MEC place fields are evenly distributed and tend to cover the environment uniformly. Also, physical place field topology is not preserved at the anatomical level (i.e., two MEC cells encoding adjacent locations are not necessary adjacent neurons).

Similar to SC, MEC exhibits place cells that are not affected by changes in the environmental shape (e.g., from square to circular arenas). Thus, MEC place fields are not independent across different environments. Rather, moving the animal from a square to a cylindrical chamber, yields highly related MEC firing topologies. Also, a MEC cell that is active in the first recording apparatus is likely to be active in the second as well. This is in contrast to CA3-CA1 place cell behavior: for instance, a cell $i \in$ CA3-CA1 that is active in the first environment, might either become silent or encode a completely unrelated location in the second arena [278].

6.3 Anatomical Lesion Data

Lesion studies provide some clues about the functionalities of the anatomical areas involved in spatial learning and navigation:

Lesions to the hippocampus. Data show that hippocampal rats exhibit impaired cognitive navigation (e.g., they are incapable to solve the hidden-platform water maze). However, hippocampal animals seem capable to perform simpler spatial behavior like taxon navigation (e.g., they can learn the visible-platform water maze) [236, 237]. In addition, more recent studies show that hippocampal animals may learn the hidden-platform water maze under particular conditions [386]. Hippocampal lesions impair spatial learning in the radial arm maze [156]. Finally, damaging the hippocampus does not seem to disrupt path integration in rats [4].

Lesions to the fimbria-fornix[10]. Similar to hippocampal animals, fornix-lesioned rats cannot learn the hidden-platform maze, whereas they can solve the visible (or cued) platform task [93, 341, 385]. Also, lesions to the fimbria-fornix impair spatial performances in the Y-maze task [150] as well as in the radial arm maze [156].

Lesions to the dentate gyrus (DG). Damaging DG produces deficits for learning the hidden water maze. Nevertheless, lesions done several weeks after training (e.g., 12 weeks) do not prevent the animal from solving the task [340, 208]. Animals with DG lesions are impaired to solving the hole-board circular platform task as well as the radial arm maze [208].

Lesions to the caudate nucleus. Experimental data show that damaging the caudate nucleus impairs navigation to both visible- and cued-platform in the water maze (i.e., taxon navigation) [264, 203], and also impairs navigation in the radial arm maze (in particular, impairing praxic navigation) [272].

Lesions to the subiculum (SC). Animals with SC lesions exhibit random-like spatial behavior in the water maze [237].

Lesions to the entorhinal cortex (EC). EC damages produce deficits in solving spatial learning tasks and also tend to diminish spatial selectivity of hippocampal pyramidal cells [220, 261, 307, 127, 278].

Lesions to the nucleus accumbens (NA). Damaging NA impairs spatial behavior in the hidden-platform

[10] Fimbria-fornix lesions mainly disrupt *(i)* the cholinergic input to the hippocampus, *(ii)* connections from the subiculum (SC) to the nucleus accumbens (NA), *(iii)* connections from the postsubiculum (poSC) to the lateral mammillary nuclei (LMN).

maze but not in the visible-platform maze [341]. Moreover, combined hippocampal and accumbens lesions produce a decrease of the animal's locomotion [387].

Lesions to the posterior parietal cortex (PP). Animals with posterior parietal lesions are impaired for solving visible- as well as cued-platform mazes [174]. Also, data on humans and monkeys show that PP lesions produce deficits in spatial cognition (see [60] for a data review).

Lesions to the posterior cingulate cortex. Lesions to the retrosplenial cortex prevent animals from directly approaching a hidden escape platform [340].

Lesions to the postsubiculum (poSC). Damaging poSC yields deficits in solving the radial arm maze as well as the water maze [352].

Lesions to the anterodorsal thalamic nucleus (ADN). Lesions to ADN severely disrupt navigation capabilities in the water maze [341].

6.4 The Theta Rhythm

The hippocampal EEG exhibits two characteristic patterns depending on the animal's ongoing behavior: *(i)* During locomotor activity (e.g., walking, running, swimming, jumping) the hippocampus is timed by a regular sinusoidal signal of $7-12\,Hz$ termed *theta rhythm* [135, 219, 257, 330]. Recent experiments show that the theta rhythm is also observable during passive locomotion of the animal [119]. Also, the theta rhythm occurs during sensory scanning as well as in REM sleep [65]. *(ii)* During non-motion behavior (e.g., eating, drinking, grooming) the hippocampal EEG exhibits a rather irregular activity with large amplitude and broad frequency spectrum [375, 66].

There exists a phase correlation between the theta rhythm and hippocampal place cell firing [257, 330]. Let i be a place cell. As the animal goes through the place field of i on a linear path, the theta phase φ_i at which i discharges shifts systematically: Every time the rat enters the field, i starts firing at the same phase φ_i^l late in the theta period. Then, as the animal runs through the field, i tends to fire earlier and earlier in the cycle. If φ_i^e is the firing phase of i when the rat exits the field, the relation $\varphi_i^l - \varphi_i^e \leq 360^\circ$ holds [257]. This phase shift phenomenon is termed *phase precession*. Such a temporal firing property of a place cell i provides more information than the solely firing rate r_i: measuring the phase φ_i allows us to estimate the position of the animal inside the place field of cell i [62].

Both CA3 and CA1 cells exhibit phase precession. All place cells start firing at the same theta phase, and, in average, precession occurs over 5 to 10 theta cycles [257, 330, 49]. DG granule cells (projecting to CA3 cells) undergo phase shifting as well, but they tend to precess in a small number of theta cycles than CA3-CA1 cells [330].

Neuroanatomical results suggest that the medial septum might be involved in modulating temporal processing in the hippocampus. In particular, the cholinergic and GABA-ergic septal efferents seem to be responsible for driving the theta rhythm [135, 64, 219, 49]. Lesion data confirm this hypothesis: damaging the medial septum disrupts the theta rhythm [396].

6.5 Hippocampal Synaptic Plasticity

To accomplish its role in cognitive behavior, the hippocampus has to provide rapid on-line storage of spatio-temporal correlations extracted from its highly processed inputs. How is information stored in the hippocampus? Which are the mechanisms underlying hippocampal learning? Activity-dependent synaptic plasticity in the hippocampal formation offers a suitable neurochemical system for implementing associative learning [235].

Long-term potentiation (LTP) is generally used to indicate the synaptic modification which is the basis of hippocampal learning [44]. LTP is defined as a persistent enhancement of the synaptic efficacy induced by electrical stimulation [235]. The resulting synaptic modification persists at least one hour, but it can last for several days in certain conditions [43]. The best understood mechanism underlying synaptic long-term potentiation is known as NMDA-mediated LTP. NMDA (N-methyl-D-aspartate) is a postsynaptic glutamatergic[11] receptor whose activation induces LTP (by leading a dendritical flow of Ca^{2+} ions) [56].

LTP can be understood as a Hebb-like mechanism [145]: it is local and temporally correlational, that is, it depends on the firing interrelation between the presynaptic neuron i and the postsynaptic neuron j [214, 45, 58, 206]. Let w_{ji} denote the strength of the synapse from i to j. In hippocampal slices, LTP occurs (i.e., w_{ji} is enhanced) only if the activity of i precedes the activity of j by less than approximately $200\,ms$ [183, 1]. On the other hand, if the postsynaptic neuron j precedes the presynaptic neuron i, either w_{ji} is unchanged (no LTP) or long-term depression (LTD)[12] occurs [85]. Another important property of NMDA-dependent LTP is the fact that it is asymmetric [1], in the sense that when i precedes j, only the efficacy w_{ji} is potentiated whereas the connection w_{ij} is not changed.

Experimental findings (mainly from pharmacological studies) strongly support the idea that the hippocampal NMDA-mediated LTP is relevant for spatial learning. Blockade of NMDA receptors results in impaired spatial learning capabilities [234]. Animals with malfunctioning NMDA receptor are unable to solve the water maze task [326, 372]. Mice that have been genetically engineered to have deficient LTP exhibit instable CA1 place fields between recording sessions [292] and are severely impaired to learn the circular platform escape task [19].

To conclude, recently it has been postulated that NMDA-dependent hippocampal synaptic plasticity plays a more general role in episodic memory rather than merely providing a basis for associative spatial learning [235].

6.6 Which are the Hippocampal Place Field Determinants?

As for head-direction cells (Sec. 2.4), several studies have been done to identify which information primarily determines place cell activity (e.g., distal *vs* local cues) [254, 240, 243, 258, 50, 278, 32, 403, 253, 79]. In particular, research has focused on the interrelation between allothetic and idiothetic inputs [148, 324, 277, 166, 390, 322].

Distal vs local cues. Neurophysiological data show that hippocampal place fields strongly depend on distal visual cues [254, 220, 258, 278, 403, 79]: rotating distal landmarks induces corresponding place field rotations. In general, low control is exerted by local landmarks (e.g., a cylinder array inside the arena) [79]. However, whenever a reward location is directly identified by one or several local cues, place fields tend to account for local landmark information [76, 32, 33, 133, 282]. Also, local non-visual allothetic cues have a strong influence upon place fields of blind and deaf rats [148, 306].

Geometrical determinants. The geometry of the environment seems to influence hippocampal place cell activity directly. Recordings data show that rescaling the experimental arena may result in rescaled place fields [240]. Also, animals placed in rectangular arenas with high cue-free walls develop place fields reflecting metric information [253]. The location encoded by the firing peak is determined by the distance to two orthogonal walls, or by the proportion of the distance between opposite walls. Place fields become stretched or exhibit multiple peaks when the environment is linearly stretched. O'Keefe and Burgess [253] suggest that this is due to the fact that place fields are normally determined by two or

[11] The glutamate is the main excitatory neurotransmitter in the hippocampus.

[12] Long-term depression yields a persistent decrease of the synaptic efficacy.

more independent metric subcomponents that become pulled apart when the geometry of the environment is changed.

Idiothetic determinants. Despite their dependence on external signals, place cells exhibit clean location selectivity even when allothetic cues are absent. This suggests that hippocampal cells are also influenced by internal movement-related signals, since these are the only available cues in the absence of environmental landmarks. Experiments show that removing visual landmarks does not perturb the activity of hippocampal cells, which continue to show stable place fields [254, 148, 271, 240, 258, 212, 278]. Also, in visually symmetric environments (e.g., a high-walled circular arena with two identical cue cards placed in two diametrically opposite positions), most place cells have asymmetric firing patterns, i.e., they exhibit distinct place fields in locations that are visually identical [324]. A strong support for the idiothetic basis for place fields comes from the fact that *(i)* place cells maintain unchanged place fields if light is extinguished, *(ii)* when the animal is introduced in the arena in complete darkness, place cells are established and persist in subsequent light conditions [277, 194]. Another hint in the favor of the contribution of self-motion signals (e.g., efference copy of motor commands), is the fact that hippocampal place cells become silent under conditions of movement restraint [209, 108]. Finally, it has been shown that place fields can be controlled by vestibular as well as optical-flow signals explicitly activated by the experimenter [322].

Dependence of place fields on the interrelation between entry position and salient allothetic cues. Experiments in visually symmetric environments suggest that the location at which the animal is placed at the beginning of a recording session (i.e., the entry position) plays a fundamental role in controlling place fields and prevents CA3-CA1 cells from doubling their firing patterns [324]. This supports the idea that a mnemonic process relying on internal cues might be employed by the animal to disambiguate vision. In particular, experimental data show that visual cues and the entry point are combined to influence place field locations [324].

Relationship between the stability of distal cues and place cell control by allothetic signals. Both behavioral and neurophysiological findings [32, 166, 159, 158] indicate that the more stable an allothetic cue is perceived by the animal, the higher its influence upon place cell dynamics. Consistently, head-direction cells are strongly controlled only by orienting cues that remain stable across sessions [166] (Sec. 2.4 and Chapter 5). When environmental manipulations (e.g., cue card rotations) occur within the same recording session (rather than between separate sessions), place cells tend to be controlled by slightly changing cues (e.g., small cue rotations). By contrast, major environmental changes (e.g., large cue rotations) occurring suddenly do not induce corresponding place field updates [293]. Similar results have been found when studying rodents' homing behavior in conflict situations between idiothetic and allothetic signals [103, 98] (Chapter 5).

Place field changes induced by adding a barrier into the arena. Once the animal has experienced an environment (i.e., it has established a place field representation), introducing an either opaque or transparent high barrier affects the activity of those cells whose place field intersects the obstacle [240]. In particular, the activity tends to diminish and place fields to be disrupted. However, the influence of the barrier vanishes whenever its height is such that the animal's motion is not affected. This suggests that locomotion-related information is relevant for establishing and maintaining place fields.

Dependence of the place field representation on reward locations. As previously mentioned, place fields tend to distribute uniformly over the space independently from reward-related locations [254, 394]. Nevertheless, pyramidal cells whose activity becomes correlated to a reward location have been observed [209, 391, 51, 239]. Data concerning this dual pyramidal cell behavior are still controversial: Speakman and O'Keefe [334] report that hippocampal place fields are independent from the target (i.e., they do not change when the target location changes). By contrast, the experiment by Breese and colleagues [51] shows the presence of cells whose first-order correlate is the reward location (i.e., they change their place

field as the target location varies)[13]. Finally, Kobayashi and colleagues [169] report that hippocampal cell firing is independent from the reward in most of the cases (80%), whereas the remaining cells tend to gradually shift their field location towards the goal location.

Latent spatial learning. The previous issue is closely related to the concept of latent learning, that is the fact that a rat establishes a representation of the environment even in the absence of any explicit reward [361, 256]. Whenever a neither hungry nor thirsty animal is put in a novel environment, it tends to learn it (i.e., to build a representation) motivated by exploration only. As a result of latent learning, when a feeder location is added into an environment previously experienced by the animal without any reward and the animal is made hungry, target-oriented navigation is learned very quickly (i.e., the animal learns immediately how to reach the reward location from any point in the environment)[14].

Place field remapping. The hippocampus can develop more than one place field representation for the same environment [258, 277, 50, 322, 166]. Remapping can take place between and within sessions whenever the current spatial representation becomes inconsistent with respect to the perceptual context (e.g., replacing a familiar white card with a new black card [50]). Note that a change in the perceptual context might occur without any change in the environment: if the animal perceives a world which is no longer congruous with its internal spatial code, a new place field map might be developed [166]. Also, hippocampal remapping can result from a change in the behavioral task performed by the animal in a totally familiar environment [195].

Non-spatial place cell correlates. Location may be considered as the first-order correlate of place cell firing [243]. Nevertheless, other factors, like animal's speed, orientation, and rotation, may also influence the activity of pyramidal cells [209, 391]. Furthermore, CA3-CA1 cells have been studied in non-spatial tasks such as odor as well as auditory discrimination [92, 263, 298]. These experiments involve stimulus-response-reward learning, e.g., animals have to learn to go to a specific reward location given a specific set of odors [92]. Results show that pyramidal cells are involved in this kind of non-spatial mnemonic process. Moreover, experimental data acknowledge the influence of the *context* on pyramidal activity: a cell responding to a stimulus in a specific task, might not encode the same stimulus in a different context [298]. The fact that place cells may be task-sensitive has been explicitly demonstrated by a series of experiments in which rats were trained to solve different tasks within the same environment [195]. As a result, each task was associated to a different place cell representation (e.g., a cell might have different place fields in different tasks). In addition, in a complex task, place cells may encode specific components belonging to the behavioral sequence for solving the task [92, 391, 298, 133]. For instance, in the odor discrimination experiment, a place cell might encode the sniff port location when the animal approaches that region to sample the odor, but not when the rat passes through that same region to reach the reward location after sampling [263].

These findings indicate that the hippocampus might play a more general role than merely being responsible for spatial learning. Indeed, hippocampal pyramidal cells might be involved in a more general class of memory (e.g., episodic memory) related to perceptual, behavioral, and reward-related variables [196, 91, 389]. Nevertheless, space might provide the contextual framework necessary to encode events [256, 247, 388].

[13]Note that, the fact that a cell i encodes the target location $\vec{p}_{goal}$ does not imply that i works as a goal cell driving the animal towards $\vec{p}_{goal}$ from any other position.

[14]Tolman [361] interprets latent learning as a form of learning which occurs without any obvious reinforcement and which is not explicitly displayed by the animal until when it is stimulated to do it.

Chapter 7

Modeling Spatial Cognition: State of the Art

The ability of animals to navigate in complex task-environment contexts has been the subject of a large body of research over the last decades. Due to its spatial representation properties (Chapter 6), the hippocampus has been studied and modeled intensively. In the first part of the chapter (Sec. 7.1), we review the principal theories attempting to explain the hippocampal role, and we describe some models of the mechanisms beneath place cell firing.

Another subject that has received a large interest is the ability of animals to navigate based on self-motion information only. This navigation mechanism, termed path integration[1], implies that the animal maintains a continuously updated vector pointing home. It allows, for instance, desert ants to forage 250 meters far away from the colony along winding routes, and then to return home along direct paths [381]. In the second part of this chapter (Sec. 7.2), we review several models explaining such a route-based mechanism. Some of these models identify the hippocampus as having a fundamental role in path integration.

7.1 Modeling Hippocampal Activity

The multivariate properties of the hippocampus have given rise to a large number of theories to explain its functionalities. The question these theories attempt to answer is: What does the hippocampus actually do? The hypotheses proposed during the last century are as differentiated as the hippocampal properties: *(i)* Early theories by Papez in 1937 [265], and Isaacson in 1974 [154], attribute to the hippocampus the function of mediating emotions. *(ii)* Olton and colleagues in 1979 [259] differentiate between reference memory (i.e., between-task information) and working memory (i.e., within-task information), and postulate that the hippocampus is primary involved in working memory. *(iii)* Gray in 1982 [134] puts forth a theory in which the hippocampal-subicular circuit acts as a predictor-comparator system: the hippocampus would predict future states and the subiculum would compare them with subsequent sensory information. *(iv)* Hirsh in 1974 [149] proposes a theory according to which the hippocampal activity underlies context retrieval: the hippocampus would provide the context to support stimulus-response

[1] Path integration will be the topic of Chapter 9.

learning. Nadel and Willner in 1980 [246] and Jarrard in 1993 [157] also postulate a role for the hippocampus in contextual conditioning. *(v)* Marr in 1971 [196] formalizes an alternative view in which the hippocampus is seen as an auto-associative memory. This theory holds that the hippocampus is a temporary memory device capable of storing and retrieving patterns of neocortical activity. *(vi)* O'Keefe and Nadel in 1978 [256] put forth the cognitive map theory for the hippocampus. This theory identifies the hippocampus as the rodent brain region providing a cognitive representation of the environment.

The two latter hypotheses have mainly influenced researchers studying the hippocampal functionalities over the last years. Indeed, both theories have served as a framework for associative-memory models [182, 392, 290, 365, 289, 373, 184, 328, 366, 280] as well as cognitive map models [296, 317, 363, 62, 377, 308, 61, 285, 117, 369]. In particular, Recce and Harris (1996) [280] have argued that the associative-memory theory by Marr [196] and the cognitive map theory by O'Keefe and Nadel [256] are consistent with each other. Indeed, learned neural representations encoding the agent's spatio-temporal knowledge of given task-environment contexts can be thought of as episodic-like memories, thus the auto-associative functionalities of the hippocampus can be used to store and retrieve this information. We agree with such an unified view of the hippocampal role. However, we emphasize the importance of place cell activity in spatial learning, which supports the cognitive map theory by O'Keefe and Nadel. Since we mainly focus on space representation based on hippocampal activity, we review some of the most relevant models attempting to explain the mechanisms underlying place cell firing. Note that, in all the approaches described below, place coding supports goal-oriented navigation. For sake of clearness, in this chapter we do not discuss the navigation part of these models. They will be reviewed in Part IV, in which goal-directed behavior of animals is considered.

7.1.1 Sharp (1991)

In this model place fields are generated based on local-view pattern classification by competitive learning [317]. The system consists of a three-layer neural network in which all units of a layer project to all units of the above layer through Hebb-like synapses.

Cells in the first layer act as metric sensory cells responding to specific stimuli (e.g., distance to a cue). The input layer involves two types of units: Type 1 cells fire as a function of the distance to specific external cues. Type 2 cells encode the distance of a cue as well as its angle relative to the agent's heading. Thus, the activity of type 2 cells depends on the agent's current location and orientation. Both types of sensory units become responding to a specific visual cue by a stochastic assignment done the first time the agent enters an environment.

Each cell i in the middle layer receives afferent projections from all units j in the sensory cell layer. Connection weights w_{ij} are initialized randomly such that they are all positive and their sum is normalized. The activity r_i of a cell i is given by $\sum_j w_{ij} r_j$, where $r_j \in \{0, 1\}$ is the firing rate of an input cell j. There are 60 cells in the middle layer, divided in three winner-take-all clusters [296]. Only one cell i^* per cluster will fire at any time, such that i^* receives the largest input $\sum_j w_{i^*j} r_j$. Synaptic weights w_{i^*j} to each winner cell i^* are strengthened by Hebbian learning. Sharp associates the middle layer of the model to the entorhinal cortex.

The pattern of activity of cells in the middle layer (i.e., 3 active cells at any time t) is propagated to the output level of the model: the hippocampus. In this level, there is only one cluster of 20 cells whose firing activity is determined according to the same winner-take-all scheme defined above. Also, synaptic enhancement occurs as before. Below is a representation of the model:

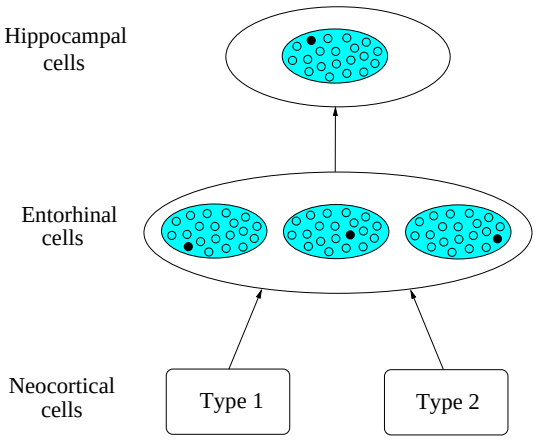

Sharp reports results obtained by simulating a circular environment with 8 evenly distributed landmarks on the edge of the arena. At each step the simulated rat computes distances and angles of all 8 cues relative to its current position and heading, respectively. This information is used to drive cells in the first layer of the model, and is propagated through the network to generate hippocampal place cell firing. Reported receptive fields are similar to real hippocampal place fields. A limitation of the model is that it does not capture data from experiments in the absence of visual landmarks (e.g., when light is extinguished) suggesting that rats are able to maintain place fields even without visual information [277, 194] (Chapter 6).

7.1.2 Burgess, Recce, and O'Keefe (1994)

Burgess, Recce, and O'Keefe put forth a hippocampal model in which allothetic sensory information activates a neural layer of entorhinal cells (EC), and then propagates through the network to form place fields in CA1-CA3 and in the subiculum (SC) [62, 61].

At the sensory level, there is a population of cells with broad tuning curve responses to the distance of external cues. A discrete set of landmarks, arranged around the edge of a square arena, is considered. For each landmark l, 15 sensory cells i respond maximally to different distances from l. Each cell i has a triangular tuning curve peaked at the preferred distance d_i^l from l:

$$r_i = \begin{cases} 0 & |d^l - d_i^l| \geq 6D \\ \lfloor (7D - |d^l - d_i^l|)/2D \rfloor & \text{otherwise} \end{cases}$$

where r_i is the activity of cell i, d^l is the current distance of landmark l, D is a unitary distance-step of $14\,cm$, and $\lfloor x \rfloor$ indicates the integer part of x. Thus, for each landmark l, the information of its distance to the agent is distributed over the ensemble activity of the sensory cell array associated to l.

Each cell $e \in$ EC receives hardwired projections from two sensory cells i and j, and multiplies their firing rates to produce an approximate radial-basis function place field. In particular, each cell e is driven by a pair of sensory cells i and j such that: *(i)* i and j respond to two distinct cues l_i and l_j, respectively, *(ii)* the tuning curves of cells i and j overlap near to the centroid location with respect to l_i and l_j. The activity of cell e is taken equal to $r_e = \lfloor r_i r_j / 2 \rfloor$.

EC cells project to the CA1-CA3 layer of the model through binary connections. The input to a cell $p \in$ CA1-CA3 is the sum of the spikes fired by its afferent EC neurons. A connection w_{pe} is switched on whenever the pre- and postsynaptic cells fire at the maximum spike frequency in the same time step (i.e., Hebbian learning). Cells in the CA1-CA3 layer are arranged in 5 clusters of 50 units each, and competitive learning [296] is applied within each cluster: At each time step, only the four cells

with the largest input fire a number of spikes proportional to their input, the others remain silent. As a consequence of competition, CA1-CA3 cells have smaller place fields than EC cells.

CA1-CA3 units project to SC cells through connections that are formed according to the same Hebbian learning scheme as before. Competitive learning also occurs within the SC layer. However, SC cells are arranged in 10 groups of 25 each, that is, each SC cell has to compete with fewer cells than each CA1-CA3 cell. This results in larger SC place fields. Below is a functional representation of the model:

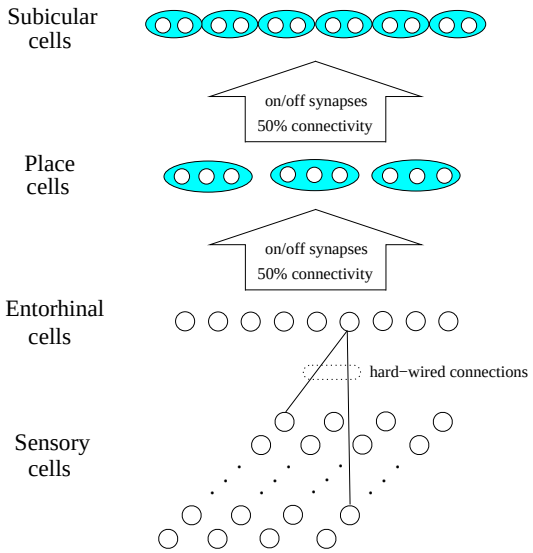

The model produces plausible EC, CA1-CA3, and SC place fields. Also, due to the metric information (i.e., distance to cues) used to determine the place cell dynamics, neurons in the model exhibit stretched (or multi-peaks) tuning curves when the environment is linearly stretched, which is consistent with experimental data [253].

The system incorporates the phase-coding relationship between place cell firing and the theta rhythm θ (Sec. 6.4). The frequency of θ is taken equal to $10\,Hz$, and each time step is equal to $1/30\,s$. Each θ cycle is divided in early, middle, and late phase. In the model, phase precession is generated at the level of EC firing: For each pair of cues l_i and l_j driving a cell $e \in$ EC, the agent estimates the angle α between its heading and the average direction of the two landmarks. If α is ahead of the agent then cell e fires at a late phase of θ, if α is behind then e fires early, otherwise e fires at a middle phase. This results in a firing activity such that a place cell active at a late θ phase, for instance, has a field centered ahead of the agent. This is consistent with the phase precession phenomenon described in Sec. 6.4.

A limitation of the model is that external landmarks are assumed to be perfectly distinguishable, and that a prewired sensory system underlies place cell formation. Also, the approach does not take into account idiothetic information (i.e., path integration) in order to enable the simulated agent to exhibit stable place fields in the absence of external cues.

In Fig. 7.1 we show some results obtained by implementing a slightly modified version of the model by Burgess, Recce, and O'Keefe. The only difference concerns the fact that we do not employ preconfigured connections between sensory cells and entorhinal cells. Rather, they are built incrementally by a simple unsupervised learning mechanism: For each location explored by the agent, simultaneously active sensory cells associated to pairs of landmarks are connected together to a newly recruited EC cell. A place is considered adequately represented whenever the number of EC cells encoding that position reaches a fixed threshold. This allows the system to bound the number of EC cells recruited during exploration.

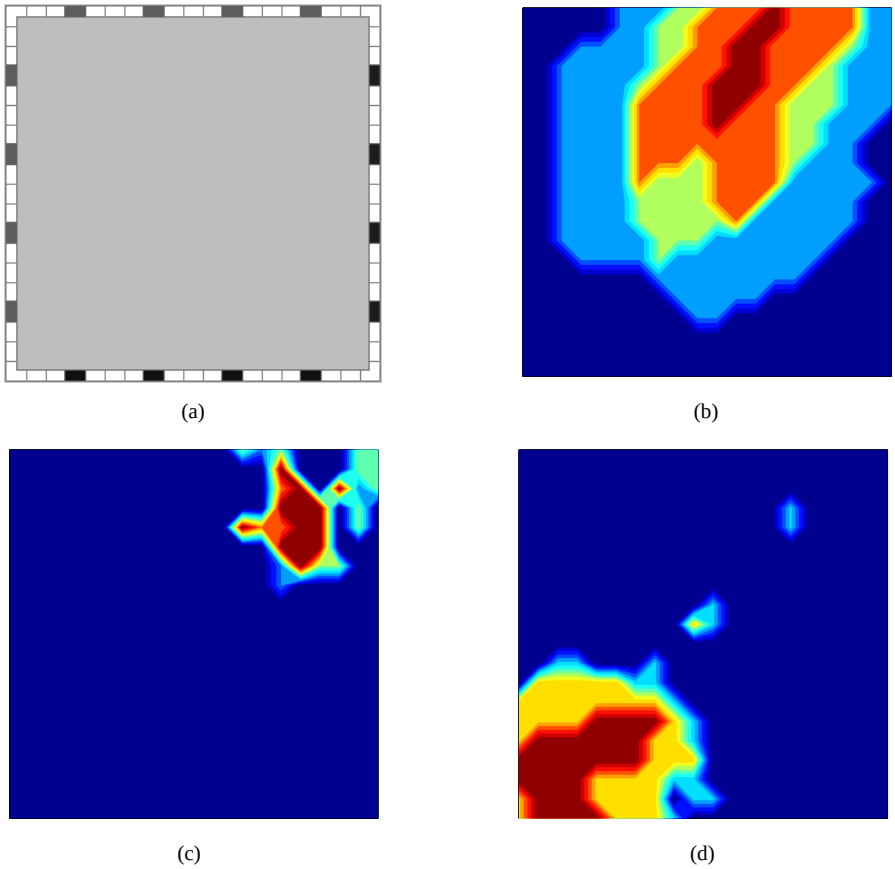

(a) (b) (c) (d)

Figure 7.1: Some results obtained by implementing a modified version of the model by Burgess, Recce, and O'Keefe [62]. (a) The simulated $60 \times 60\,cm$ arena with 16 cues evenly distributed around its edge. Results: (b) Sample of entorhinal receptive field. (c) Place field recorded from the CA1 layer of the model. (d) Place field of a subicular cell.

7.1.3 Wan, Redish, and Touretzky (1994, 1997)

Wan, Redish, and Touretzky formulate a comprehensive theory of rodent navigation based on the interaction among separate modules encoding head-direction, path integration, and place recognition [377, 286, 285, 283]. Each of these representations mediates associations of internal and external stimuli. Below is a diagram of the model:

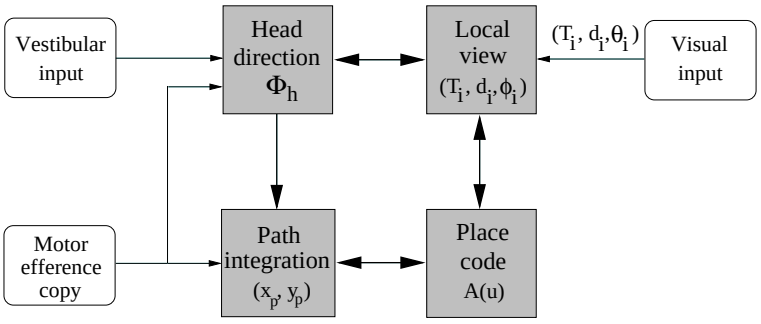

First, visual input consists of triplets (T_i, d_i, θ_i). Thus, the agent is assumed to be able to identify the type of visual landmark T_i, and to compute its egocentric polar coordinates d_i, θ_i. This latter information is combined with the head-direction signal Φ_h to compute the allocentric landmark bearing ϕ_i. The

information (T_i, d_i, ϕ_i) forms a local view relative to landmark i. Second, a representation (x_p, y_p) of the animal's position $\vec{p}$ within an internal coordinate system is updated by path integration. Third, place code is established by means of radial-basis units tuned to *(i)* the type, distance, and allocentric bearing (T_i, d_i, ϕ_i) and (T_j, d_j, ϕ_j) of pairs of distinct landmarks i and j, respectively, *(ii)* the retinal angle $\alpha_{ij} = \theta_i - \theta_j$ between landmarks i and j, *(iii)* the path integration coordinates (x_p, y_p). The activity $A(u)$ of a place unit u is computed by a "fuzzy conjunction" of these three informations in which terms drop out when the corresponding input is not available (e.g., in the dark $A(u)$ depends solely on path integration). More precisely, $A(u)$ is given by $A(u) = C(u) \cdot L(u)$, where $C(u)$ is a two-dimensional gaussian tuning to path integrator coordinates centered at (x_u, y_u), and $L(u)$ is a product of univariate gaussians of the inputs provided by the local view module (each input consists of (T_i, d_i, ϕ_i), (T_j, d_j, ϕ_j), and α_{ij}).

Training proceeds as follows: The first time the agent enters an unfamiliar environment it is told its head-direction Φ_h and its path integrator coordinates (x_p, y_p). For each newly explored position (x'_p, y'_p) (i.e., each position that is not already encoded by existent place units), a new place unit u' is recruited such that: *(i)* Its response function $C(u')$ is tuned to the path integrator coordinates (x'_p, y'_p); *(ii)* Two visible landmarks i and j are randomly selected, and the unit's response functions encoding (d_i, ϕ_i) and (d_j, ϕ_j) are tuned to these values; *(iii)* Two different landmarks k and l are chosen, and the unit's response function encoding α_{kl} is tuned to the bearing difference $\theta_k - \theta_l$. This process is iterated for each newly visited position. After training, all multiple representations converging to the place code interact with each other to maintain a consistent space knowledge.

Computer simulations reported by the authors replicate several neurophysiological as well as behavioral experiments. A limitation of the model is the fact that during training (i.e., when tuning place units) the path integrator is assumed to be coherent over time. Indeed, there is no mechanism to correct the path integration coordinates (x_p, y_p) simultaneously with tuning place cells during exploration.

7.1.4 Schölkopf and Mallot (1995)

In this model, a topological representation of the environment is built incrementally by generating a view graph [308]. The input of the model is constituted by sequences of views and movements. Vertices of the graph correspond to views, while edges connect distinct views that are experienced by the agent in immediate temporal sequence. Also, each arc of the graph from node i to node j is labeled by the egocentric movement (e.g., go left) necessary to the agent to perceive the view v_j associate to node j when seeing v_i in node i. Once the view graph has been learned, places can be easily recovered from the topological view representation.

Schölkopf and Mallot propose a neural network architecture to implement their view-graph model. Below is a sketch of the network:

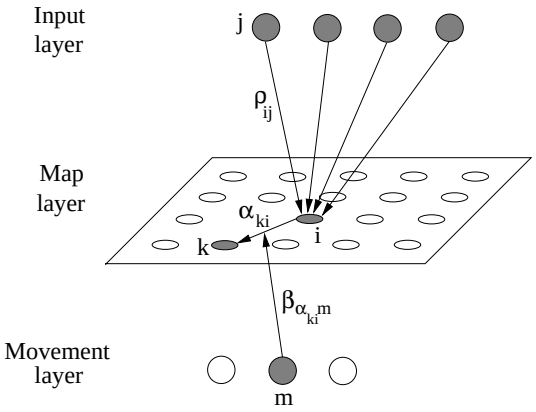

The network consists of an input layer, a map layer, and a movement layer. Views j in the input layer must be linked to units i in the map layer. This is achieved by updating weights ρ_{ij}. After learning, perception of a view will be represented by the activity of the associated map unit, i.e., the unit whose input weights are most closely tuned to the presented view. In order to learn topology, units i and k in the map layer are interconnected by a link α_{ki}, whose strength is set according to the temporal sequence of the views represented by these units. Then, movement labels are learned in terms of modulatory connections $\beta_{\alpha_{ki}m}$ from cells m in the movement layer to links α_{ki} in the map layer. Three movement units m are considered (encoding left, right, and back[2] actions), and, during learning, modulatory links $\beta_{\alpha_{ki}m}$ are updated.

The authors provide computer simulation results obtained by validating the model in a corridor maze in which places correspond to junction locations. A nice property of the model is that, since movement information to go from one place to another becomes stored in the topological graph, the space code can be maintained even without any visual input. An intrinsic limitation of the approach is that all views taken by the agent are assumed to be perfectly distinguishable.

7.1.5 Trullier and Meyer (1997)

In this model, the hippocampus is seen as an hetero-associative network that learns temporal relationships between successive configurations of stimuli [370, 368, 369]. Thus, exploration is a process by which the agent experiences sequences of places. Learning occurs to store these sequences onto a topological graph by transforming temporal relationships into spatial relationships. The authors assume that CA3 is the anatomical locus for this topological graph.

Nodes of the graph represent visited places, and edges connect sequentially experienced locations. Similar to Schölkopf and Mallot [308], head-direction information is used to label the edge between nodes i and j with the direction of movement experienced when arriving in j from i. In order to create the graph incrementally (i.e., to add new nodes as exploration proceeds), the authors adopt a mechanism by which the animat can recruit place cells according to the local complexity of the environment. They consider wall corners as landmarks, and assume that the agent is always capable to recognize them perfectly. Place cells p are tuned to the distances to a certain number n of landmarks

$$act_p = \prod_{i=1}^{n} exp\Big(- (d_i - d_i^*)^2/\sigma^2 \Big)$$

where: d_i, $0 \leq d_i \leq 1$, is given by $d_i = exp(-\overline{d}_i)$ with $\overline{d}_i$ being the distance between the current agent's location and landmark i; d_i^*, $0 \leq d_i \leq 1$, is given by $d_i^* = exp(-\overline{d}_i^*)$ where $\overline{d}_i^*$ is the distance between the location where cell p has been recruited and landmark i; σ defines the size of the place field of cell p. When a landmark i is not visible, d_i is arbitrarily set to -0.1. Thus, newly recruited place cells encode the distances between the agent and visible landmarks. A limitation of this approach consists of the no sensory aliasing assumption (i.e., landmarks are assumed to provide unique visual cues), and the orientation independence of the simulated views.

7.1.6 Gaussier *et al.* (1998)

In this model, place recognition relies on the identification of at least two visual landmarks within a panoramic scene, and on the estimation of their allocentric azimuth relative to an absolute direction (e.g.,

[2]The topology of the maze is such that the agent can only move forward taking either left or right actions.

north) [118, 117]. The activity of a place cell i when the robot is at location (x, y) is given by

$$act_i = 1 - \frac{1}{\pi N_i} \sum_{k=1}^{N_i} V_{ik} \cdot f(|\Theta_{ik} - \theta_k|, v_k) \tag{7.1}$$

where: N_i is the number of landmarks; Θ_{ik} is the azimuth of landmark k as memorized when cell i has been created; θ_k is the azimuth of landmark k as measured from the current position (x, y); V_{ik} is equal to 1 if landmark k was visible when cell i has been created, 0 otherwise; v_k is equal to 1 if landmark k is visible from the current position (x, y), 0 otherwise; the angular difference $|\Theta_{ik} - \theta_k(x, y)|$ is computed modulo π; the function f is defined according to

$$f(\theta, v_k) = \begin{cases} \theta & v_k = 1 \\ \pi & v_k = 0 \end{cases}$$

Eq. 7.1 results in a place cell activity act_i that tends to 1 when all the azimuths θ_k relative to the agent's current position are similar to the azimuths Θ_{ik} learned when cell i has been created. At any time t, the recognized place is the one associated to the most active cell i^*, i.e., $i^* = \text{argmax}_i\,(act_i)$. During exploration, the simulated animal adds a new place cell i: *(i)* If it encounters an interesting place (e.g., a food source), *(ii)* if it is at a poorly represented location (i.e., $act_{i^*} < T$, where T is a recognition threshold), *(iii)* at the end of an obstacle avoidance behavior.

7.2 Path Integration: Previous Hypotheses

The hypothesis according to which animals might navigate by integrating internally-generated motion signals was first suggested by Darwin and Murphy in 1873 [81, 245]. Barlow, in 1964, reviews the early literature concerning inertial navigation, and elucidates the physical principles that might be the basis of this capability of animals [21].

The first model attempting to explain the mechanism underlying path integration was proposed by Jander in 1957 [155]. In this section, we give a brief overview of the models that have been postulated since the pioneer Jander's hypothesis. They have been classified according to the following taxonomy: *(i)* Mathematical approaches modeling path integration quantitatively [155, 224, 238, 28, 112]; *(ii)* Models based on the connectionist paradigm [199, 200, 141, 210, 363, 377]; *(iii)* Approaches focusing on the neurophysiological plausibility of the functional and anatomical aspects involved [207, 301, 319, 285].

7.2.1 Mathematical Models

7.2.1.1 Jander (1957)

Jander [155] postulates a model in which a time-weighted angular integration over time yields the representation of the current homing direction $\theta_h(t)$

$$\theta_h(t) = \frac{1}{t - t_0} \int_{t_0}^{t} \theta(t') \; dt' \tag{7.2}$$

where t_0 is the time at which the journey starts, and $\theta(t')$ is the instantaneous animal's bearing relative to an external reference direction ϕ. The distance of the animal to the starting point is not computed by the model. Note that, since the model measures the duration of each segment rather than its metric length, it assumes a constant translational speed of the animal. Below is an example of three-leg trajectory T (from point A to point B) to illustrate Eq. 7.2 in a discrete fashion.

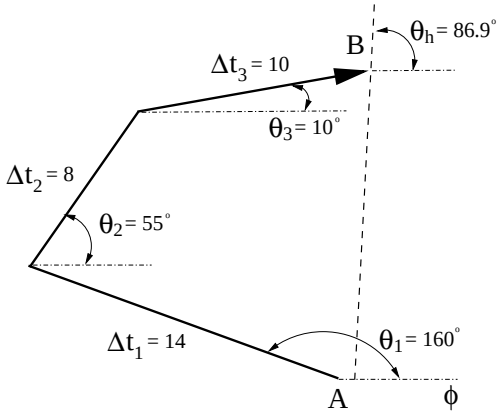

$$\begin{aligned} \theta_h &= \frac{1}{\sum_i \Delta t_i} \sum_i \theta_i \; \Delta t_i \\ &= \frac{\theta_1 \, \Delta t_1 + \theta_2 \, \Delta t_2 + \theta_3 \, \Delta t_3}{\Delta t_1 + \Delta t_2 + \Delta t_3} \\ &= 86.9^\circ \end{aligned}$$

According to Eq. 7.2, the animal has to memorize the duration Δt_i as well as its bearing θ_i *for all* segments $i \in T$ in order to estimate the final homing direction θ_h. That is, the process is not iterative, in the sense that the system needs to backtrack the entire history (from B to A) to compute the time-weighted average defined by Eq. 7.2.

Jander's model does not provide the geometrically correct solution to "close the polygon". In the above example, the correct homing direction would be $\theta_h = 84.4^\circ$, that is, the algorithm yields an error of 2.5°. Although the biological path integrator is also affected by systematic biases [314], Eq. 7.2 does not reproduce the type of errors observed in experiments with ants and other species [384].

7.2.1.2 Mittelstaedt and Mittelstaedt (1973)

Mittelstaedt and Mittelstaedt [222] put forward a model in which the homing vector $\vec{h}(t)$ (i.e., both direction and distance to home) is encoded by the animal's coordinates $x(t)$, $y(t)$ within a Cartesian allocentric frame

$$\begin{aligned} x(t) &= \int_{t_0}^{t} v(t') \cos(\theta(t'))\ dt' \\ y(t) &= \int_{t_0}^{t} v(t') \sin(\theta(t'))\ dt' \end{aligned} \tag{7.3}$$

where $v(t')$ and $\theta(t')$ are, respectively, the speed and the heading (relative to an arbitrary axis ϕ) of the animal during the time increment dt'.

This theory, termed *bicomponent* due to the trigonometric decomposition of the animal's movement (Eq. 7.3), was first introduced in 1962 by H. Mittelstaedt to model the orientation control system in insects [221]. Then, in 1973, it was employed to explain path integration in insects [222], and finally it was applied to vertebrates in 1980 [224].

Eq. 7.3 is the error-free trigonometric solution (ideal path integration) to compute the current animal's position by vector integration in allocentric space [201]. The homing direction is given by $\theta_h(t) = \arctan y(t)/x(t)$, and the distance to home by $l_h(t) = \sqrt{(x(t) - x(t_0))^2 + (y(t) - y(t_0))^2}$. Below is the same three-leg sample trajectory T used above, with the discrete solution for $\theta_h(t)$ and $l_h(t)$ according to Eq. 7.3:

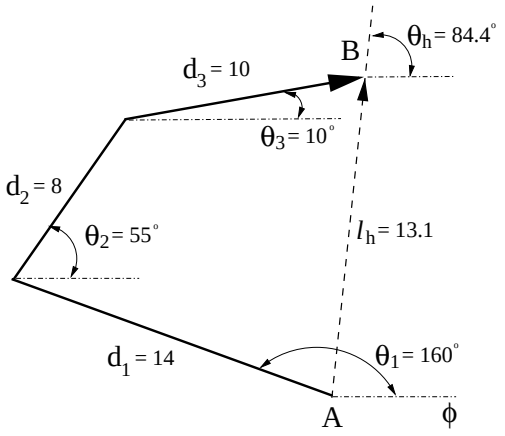

$$\begin{aligned} \theta_h &= \arctan \frac{\sum_i v_i \sin(\theta_i)\Delta t_i}{\sum_i v_i \cos(\theta_i)\Delta t_i} \\ &= \arctan \frac{\sum_i d_i \sin(\theta_i)}{\sum_i d_i \cos(\theta_i)} \\ &= 84.4^\circ \end{aligned}$$

$$l_h = |\vec{h}(t)| = 13.1$$

Thus, the model represents the correct geometrical solution to generate the homing vector $\vec{h}$. Systematic errors affecting the biological path integrator, as observed by behavioral experiments [201], are not accounted by the model. Rather, Mittelstaedt and Mittelstaedt assume that the drift in the homing process accumulates over time as a result of non-systematic errors (i.e., random noise), but the latter are not explicitly modeled by the system.

Finally, unlike Jander's model, Eq. 7.3 has an iterative solution: the homing vector $\vec{h}(t + \Delta t)$ can be computed entirely based on $\vec{h}(t)$, the vector at the previous time increment, without a complete backtrack through the animal's trajectory. Indeed,

$$\begin{aligned} x(t + \Delta t) &= x(t) + v(t) \cos(\theta(t))\, \Delta t \\ y(t + \Delta t) &= y(t) + v(t) \sin(\theta(t))\, \Delta t \end{aligned} \tag{7.4}$$

7.2.1.3 Müller and Wehner (1988)

Müller and Wehner [238] start from the hypothesis that the trigonometric vector decomposition proposed by Mittelstaedt and Mittelstaedt [222] requires computational capabilities that might be unlikely for simple organisms as ants. Thus, they propose an iterative method that provides an approximate solution by reformulating the original angular-integration hypothesis postulated by Jander.

The animal's trajectory is discretized in unit-length time steps. Let $\theta_h(t)$ and $l_h(t)$ be the allocentric homing direction and the distance from the starting point, respectively, at time t. Let δ denote the angular deviation (in degrees) between the animal's heading at time t and at time $t+1$, i.e., $\delta = \theta(t+1) - \theta_h(t)$. The iterative process to compute $l_h(t+1)$ and $\theta_h(t+1)$ is defined according to

$$\begin{aligned} l_h(t+1) &= l_h(t) + 1 - \frac{|\delta|}{90} \\ \theta_h(t+1) &= \theta_h(t) + k\frac{(180+\delta)(180-\delta)\,\delta}{l_h(t)} \end{aligned} \tag{7.5}$$

where $k = 4.009 \times 10^{-5} deg^{-2}$ is an empirically tuned constant to fit data concerning *cataglyphis fortis* ants. The following example illustrates the method.

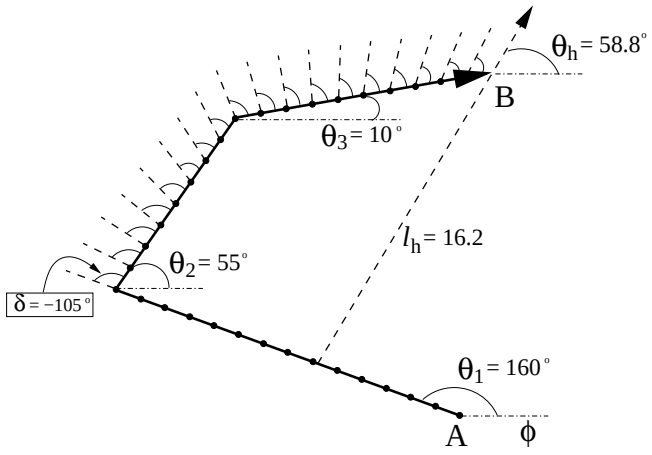

At each time step, the angular difference δ between the current animal's heading and the estimated homing direction is shown. The final homing vector $\vec{h}$, whose polar components are θ_h and l_h, approximates the geometrically correct solution. The direction error is approximately of $25.6°$, and the estimated A-B distance exceeds the correct value by an error of about 3.1.

The distance-weighted mean-direction algorithm proposed by Müller and Wehner (Eq. 7.5) successfully accounts for systematic biases in path integration. It fits several experimental data concerning homing capabilities of ants [238], as well as other insects and arthropods [382]. The model also provides a good approximation for the homing behavior of rodents, in particular of golden hamsters [314].

7.2.1.4 Benhamou *et al.* (1990)

Benhamou *et al.* [28] reconsider the trigonometric decomposition hypothesis by Mittelstaedt and Mittelstaedt [222] and reformulate it by *(i)* adding Gaussian noise to the system, *(ii)* encoding the translational and angular components of motion within an egocentric rather than allocentric coordinate system.

Let $l_h(t)$ be the distance to the starting point as estimated by animal, i.e., $l_h(t) = |\vec{h}(t)|$, and let $\theta_h(t)$ indicate the egocentric homing direction, i.e., the angle between the current animal's heading and the

starting location. Thus, the homing vector $\vec{h}(t)$ is represented within an animal-centered polar reference frame. The animal's trajectory is seen as a discretized sequence of movements. At each step, the animal either rotates by an angle α, or makes a translation of length d. After a rotation $\alpha(t)$ at time t and a translation $d(t+1)$ at time $t+1$, the egocentric polar representation of the animal's location is given by

$$
\begin{aligned}
l_h(t+1) &= \Big(l^2(t) + d^2(t+1) - 2\,l(t)\,d(t+1)\,\cos(\theta_h(t) - \alpha(t))\Big)^{\frac{1}{2}} \\
\theta_h(t+1) &= \arctan\left(\frac{\sin(\theta_h(t) - \alpha(t))}{\cos(\theta(t) - \alpha(t)) - d(t+1)/l_h(t)}\right) + k\,\pi
\end{aligned}
\tag{7.6}
$$

where $k = 0$ if the denominator is positive, and $k = 1$ otherwise.

Note that Eq. 7.6 provides a stochastic version of the geometrically correct trigonometric solution (i.e., systematic errors are dismissed by the model). Indeed, the model relies on the assumption that errors in the path integration process are not intrinsic in the updating computational process; rather, they are induced by inaccurate measurements of angular and translational components. Benhamou *et al.* suppose that Gaussian noise affects the estimate α of the angular displacement as well as the estimate d of the step length. Below, is the sample trajectory T and the homing vector as computed by Eq. 7.6:

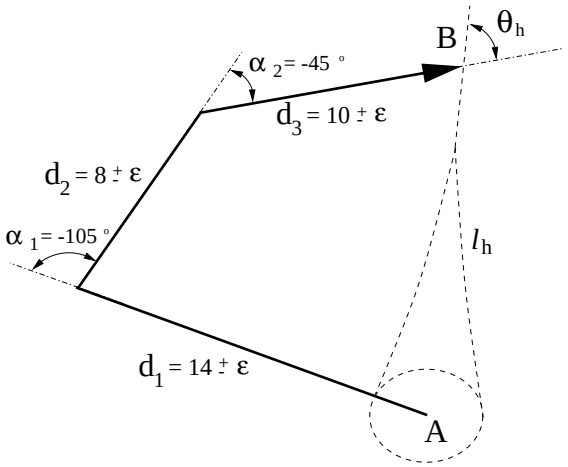

Benhamou *et al.* assume that translation estimation relies on idiothetic signals only, whereas rotations may be assessed based on both idiothetic and allothetic stimuli. Interestingly, the model predicts that *(i)* errors occurring only in the translation estimation do not affect the animal's homing performance, *(ii)* errors in the allothetic-based rotation assessment do not disrupt path integration, *(iii)* noise in the idiothetic-based estimation of angular movements strongly impairs the accuracy of the homing behavior. The prediction that linear errors have low influence compared to (idiothetic) angular errors might explain the fact that rodents seem to discount linear translation information when displaced passively [224], as well as the fact that they exhibit a poor capability of estimating the length of short, straight journeys [314]. Rodents might simply neglect linear information since the latter is not of eminent importance for their homing capabilities.

7.2.1.5 Fujita *et al.* (1990)

Fujita *et al.* [112] also suggest that a trigonometric decomposition within an egocentric framework underlies path integration. However, in contrast to the previous model, they propose an approximate solution to assess the components of motion, based on the assumption that computing the exact solution is beyond the capabilities of many organisms.

Fujita *et al.* first derive a set of equations providing the ideal solution to compute, at each time t, the egocentric homing vector $\vec{h}(t)$ in polar coordinates $(l_h(t),\, \theta_h(t))$. Those equations are the deterministic version of Eq. 7.6 employed by Benhamou *et al.*, and generate the vector $(l_h(t),\, \theta_h(t))$ recursively in terms of its previous components $(l_h(t-1),\, \theta_h(t-1))$.

As a second step, the authors proceed to a simplification of the above equations, by assuming that *(i)* the agent has traveled for a large distance from the origin before the first turn, *(ii)* the navigator does not approach the origin too closely during a winding route. Given that, the following solution is provided

$$\begin{aligned} l_h(t+1) &= l_h(t) - d(t)\, \cos(\theta_h(t)) \\ \theta_h(t+1) &= \theta_h(t) + d(t)\, \sin(\theta_h(t)/l_h(t)) - \alpha(t) \end{aligned} \tag{7.7}$$

where $d(t)$ and $\alpha(t)$ are the amplitude of the translation and of the rotation at time t, respectively. The authors observe that Eq. 7.7 constitutes the correct solution for $\vec{h}(t+1)$ in the case of step-size zero, that is, in the case of continuous integration of movements over time.

As a final step, Fujita *et al.* linearize Eq. 7.7 by approximating the sine/cosine functions as follows:

$$\sin(x) = \begin{cases} (-2/\pi)\, x - 2 & -3\pi/2 \leq x < -\pi/2 \\ (2/\pi)\, x & -\pi/2 \leq x < \pi/2 \\ (-2/\pi)\, x + 2 & \pi/2 \leq x < 3\pi/2 \end{cases} \tag{7.8}$$

$$\cos(x) = \begin{cases} (2/\pi)\, x + 1 & -\pi \leq x < 0 \\ (-2/\pi)\, x + 1 & 0 \leq x < \pi \\ (2/\pi)\, x - 3 & \pi \leq x < 2\pi \end{cases} \tag{7.9}$$

Fujita *et al.* report that their linearly approximate solution produces highly accurate estimates for the homing vector $\vec{h}(t)$. Nevertheless, they also observe that the model is not intended to emulate animals' path integration realistically. Indeed, their results do not fit experimental findings concerning ants' path integration capabilities.

7.2.2 Connectionist Models

7.2.2.1 McNaughton *et al.* (1991)

McNaughton *et al.* [210] suggest the first neural machinery capable to integrate internally generated signals over time. In their model, a linear associative network is employed to replace mathematical integrals of variables.

We have already discussed the method in Sec. 3.1, since the model focuses on the integration of angular velocity to establish an allocentric directional representation. Nonetheless, McNaughton *et al.* argue that the postulated mechanism might also be applied to maintain a representation of distance traveled. In other words, they suggest that their model might provide a neural basis for the more general process of path integration. In Sec. 7.2.3 we describe how this theory has been further developed by McNaughton and colleagues [207, 301].

7.2.2.2 Touretzky *et al.* (1993)

Touretzky *et al.* [363] start from O'Keefe's hypothesis (1991) [252] suggesting that rodents might encode the angle θ and the distance l to a location as the phase and the amplitude, respectively, of a sinusoidal

temporal function $f(t)$, namely a *phasor*[3]. Touretzky *et al.* propose a neural model to represent a phasor $f(t)$ by the sinusoidal activity of a population of neurons arranged in an array A of N elements. The continuous temporal signal $f(t)$ is encoded by the discretized representation provided by A:

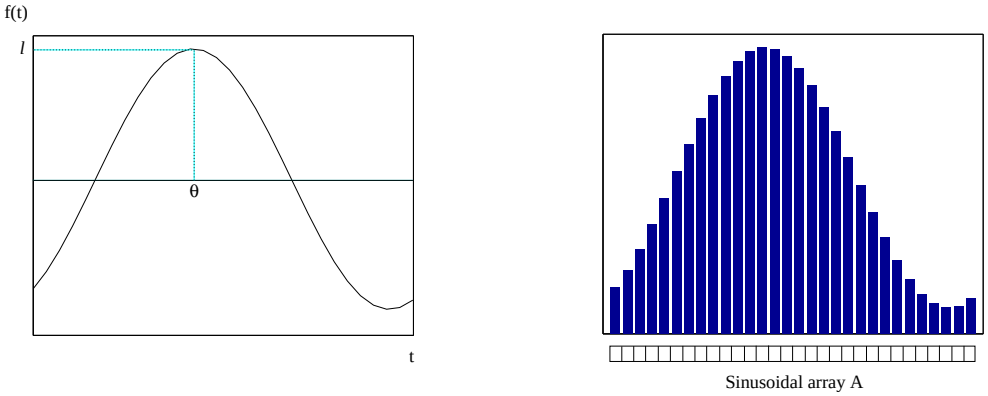

The value of each element $i \in A$ encodes the amplitude of the sinusoidal wave sampled at point $2\pi i/N$. When representing a vector $(l,\, \theta)$, the level of i is given by

$$f(l,\, \theta,\, i) = l \cdot \cos(\theta + 2\pi i/N) \tag{7.10}$$

The precision of the representation depends on the angular resolution $2\pi/N$. The authors report that $N = 24$ (i.e., a precision of $\pm 7.5°$) yields a sufficiently accurate approximation.

To encode the amplitude defined by Eq. 7.10, each element $i \in A$ consists of a population of noisy spiking neurons whose average activity provides the $f(l, \theta, i)$ value. Since neural firing can not be negative, the activity pattern is shifted by means of a baseline b in order to represent the negative part of the sinusoidal wave. Thus, the activity $r(i)$ of a neuron in the i_{th} element is defined according to

$$r(i) = b + k \cdot f(l, \theta, i) \tag{7.11}$$

where k is a constant gain factor. For their simulation, the authors take $b = 40\, spikes/sec$, which results in a firing range of $[0,\, 80]\, Hz$.

Touretzky *et al.* employ their sinusoidal array representation to manipulate vectors in a phasor form. Vector translation, in a phasor-based coordinate system, is simply achieved by addition of sinusoidal waves (e.g., $f(t) = f_1(t) + f_2(t)$). Therefore, in the sinusoidal array model, vector translation consists of simple element-wise addition of firing rates, e.g., $r(i) = r_1(i) + r_2(i) - b,\ \forall i \in A$ (the baseline b normalizes the activity). Similarly, vector subtraction is achieved by rotating the negated vector by $180°$, e.g., $r(i) = r_1(i) + r_2(i + N/2 \mathrm{mod} N) - b, \quad \forall i \in A$. Rotating a vector $f(t)$ by α radians in phasor coordinates corresponds to $f(t + \alpha)$. In the model, this operation is performed by rotating the sinusoidal array A by means of an additional module, namely a shifter circuitry.

In order to validate their system, Touretzky *et al.* have simulated two experiments with gerbils reported by Collett *et al.* [76]. In these tasks, an allocentric memory vector $\vec{v}_1$ describes the angle and distance of a landmark from a feeder location. A second vector $\vec{v}_2$ encodes the landmark position relative to the animal's current location. Thus, to solve its task the animals has to infer the goal position by vector subtraction within an allocentric directional framework, that is $\vec{v}_{goal} = \vec{v}_1 - \vec{v}_2$. The sinusoidal array model has been employed for that purpose.

[3]Phasor representation of vectors in polar coordinates: Any vector $\vec{v} = (l,\, \theta)$ can be temporally represented as a sinusoidal wave $f(t) = l \cdot \cos(\omega t + \theta)$, where l defines the wave amplitude, θ the phase, and ω the frequency. Function $f(t)$ is called phasor. The main advantage of the phasor representation is that vector manipulation (translation, rotation) is straightforward: Translation is accomplished by addition of sine waves, and rotation is achieved by phase shifting.

Note that, the model by Touretzky *et al.* does not focus on the mechanism for integrating movements over time (i.e., dead-reckoning). Rather, it provides a neural implementation suitable for manipulating vectors in allocentric space. Since path integration is a vector-based navigation process [380], we have introduced the model within this section.

7.2.2.3 Hartman and Wehner (1995)

Hartman and Wehner [141] put forth a modular neural architecture to emulate the navigation capabilities of *Cataglyphis* ants. Each neural component specifically addresses a computational aspect of dead-reckoning (e.g., integration of angular displacements over time), such that the entire system works as a neural path integrator.

The model is based on the idea of traveling activity patterns for representing a temporally changing variable $x(t)$: Given a population of neurons, the ensemble activity, at each time step, forms a stable state that encodes the current value of x. Such a blob of activity, then, travels over the neural population according to the variations of the variable. The neural implementation consists of a chain of mutually-coupled excitatory units. The stability of the traveling peak as well as its width are determined by two sets of inhibitory neurons projecting offset connections to the left and to the right, respectively. Below is a scheme of the model:

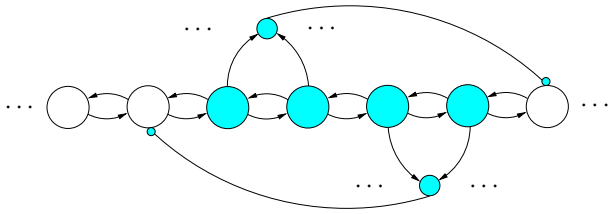

To shift the activity blob (grey neurons) one-step leftward, an extern current is injected into inhibitory neurons projecting to the right. Similarly, to induce a rightward shift, neurons inhibiting to the left are maximally activated.

Hartmann and Wehner employ traveling activity patterns to *(i)* assess the length of a path $s = \sum_i \Delta s_i$ by integrating linear displacements $\Delta s_i = v_i \, \Delta t_i$ over time (where v_i is the translational velocity), *(ii)* measure rotations $\alpha = \sum_i \Delta \alpha_i$ by summing angular increments $\Delta \alpha_i = \omega_i \, \Delta t_i$ over time (where ω_i is the angular velocity). The accuracy of the integration process depends on the number of units forming the chain, that is, on the discretization produced by the neural representation.

Let $\theta_h(t)$ and $\theta(t)$ be the homing direction and the animal heading at time t, respectively, with respect to an allocentric arbitrary reference ϕ. Let l_h denote the distance to the nest. Then, locomotion in direction θ by a linear path s results in a new animal's position encoded by $(l_h(t+1), \, \theta_h(t+1))$. As we have seen before, this new polar coordinates may be computed by either exact or approximate trigonometric decomposition.

The model involves a neural chain to integrate the nest distance l_h, and two cyclical neural chains encoding the homing angle θ_h and the direction of motion θ, respectively. The width of the blobs representing θ_h and θ is taken equal to $w(\theta_h) = w(\theta) = 90^\circ$. In order to perform goniometric computation (i.e., to assess the difference $\delta = \theta - \theta_h$), a third cyclical chain C^+ is considered whose ensemble activity encodes the overlap between the θ-peak and the θ_h-peak. In particular, the number of active neurons in C^+, i.e., $\#C^+$, is proportional to δ. That is, $\#C^+$ encodes a goniometric computation. Furthermore, $\#C^+(\delta)$ may be considered as a first-order approximation of the positive part of a cosine function $\cos(\delta)$:

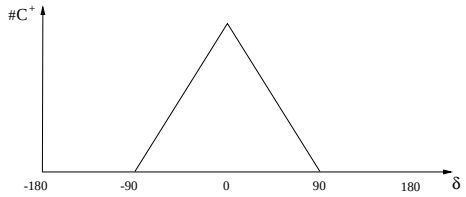

By rotating the θ_h chain relative to the θ chain properly, three more populations, C^-, S^+, and S^-, may be considered to provide a first-order approximation for the negative part of a cosine, and the positive and negative parts of a sine function, respectively.

Hartmann and Wehner show that by adjusting the network parameters properly (in particular the peak widths $w(\theta_h)$ and $w(\theta)$), the model reproduces the navigation behavior of *Cataglyphis* ants very accurately. In addition, they observe that parametric changes also enable their neural machinery to compute geometrically correct trigonometric decomposition, which demonstrates that trigonometry is computationally accessible by neural systems.

7.2.2.4 Maurer (1998)

Maurer [190, 200] proposes a neural network to realize path integration based on egocentered polar coordinates. The system consists of a multilayer neural network trained by back propagation [294] to predict the future animal's position, given the current coordinates and movement. The author argues that although back propagation is not neurophysiologically plausible, the resulting neural network may converge toward solutions comparable to those observed in biology. Below is the neural network used for a first version of the model

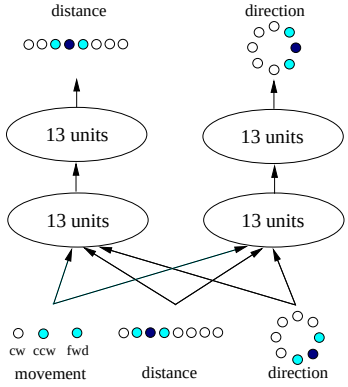

The network takes as input: *(i)* The activity of three units encoding, respectively, forward translations (fwd), clockwise rotations (cw), and counterclockwise rotations (ccw). These three units fire proportionally to the magnitude of the movement (either translation or rotation) being executed at the current step. *(ii)* Two sets of 8 units each coding for the current egocentric polar coordinates (distance and direction) of the animal relative to the nest. Each unit has a triangular activity tuned to one preferred coordinate value, and the responses of adjacent units overlap each other. In other words, polar coordinates are encoded by means of two coarse coding representations. The network provides as output the next animal's position. The network is actually folded, in the sense that its output is directly fed into its input. By this, the network forms a dynamic memory loop suitable for continuous path integration.

In order to validate the model, Maurer adopts simulated environments and simulated animals whose characteristics reflect those in experiments with hamsters [100, 314]. In particular, the model is tested in two different tasks. In the first, the animal undergoes three full revolutions around the feeder position (at the arena center) before eating, after which it returns to the nest. In the second, the hamster follows

an L-shaped outward journey from the nest to the food, and returns home after hoarding. Maurer's first model shows a plausible behavior with respect to the first experiment, but it does not fit experimental results concerning the L-shape experiment.

As a consequence, Maurer develops a new version of the model in which the neural network is trained to provide only the directional component of the polar representation of the animal's position. Below is the architecture of the second network:

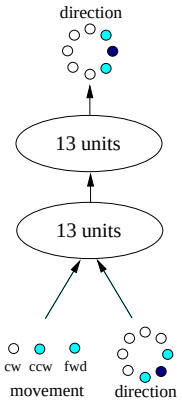

After training, the network has been tested with the same two experiments as above. In the three-revolution task, the network performed particularly well (even better than hamsters). In the L-shaped test the model reproduced experimental data quite properly. Maurer's findings suggest the counterintuitive hypothesis that distance representation plays a secondary role in path integration. This agrees with results reported by Benhamou *et al.* [28] showing that errors in the estimation of linear movements do not impair animals' homing capabilities.

7.2.3 Anatomical Hypotheses

Most of the models reviewed so far are rather theoretical, in the sense they focus on the mechanisms behind path integration without looking for plausible anatomical counterparts for their functional components. In this section, we move towards a somewhat different approach: Understanding the nature of path integration by identifying those brain areas that might provide the locus for such a navigational system. A general consensus emerges from the three theories discussed below: The hippocampal formation provides all the computational requirements necessary for path integration.

7.2.3.1 McNaughton *et al.* (1996)

McNaughton *et al.* [207] propose a theory in which the hippocampus works as an inertial path integrator. According to this thought, the space representation occurring in CA3-CA1 areas is primarily based on idiothetic signals. Exteroceptive information is a second order correlate, and is used to initialize and calibrate path integration by associative learning.

The core of the theory is the multi-chart hippocampal hypothesis: The CA3 synaptic matrix provides a large set of preconfigured quasi-independent reference frames, which permit spatial coding solely on the basis of internal information. Each of these reference frames can be imagined as a two-dimensional surface (a chart) consisting of a random subset of the total CA3 place cell population. Each cell is assigned a random location in the chart and the synaptic strength of CA3 collaterals defines metrical neighborhood relationships based on a two-dimensional Gaussian measure of distance. This is equivalent to consider place fields within a chart as vectors from an arbitrary point in an allocentric coordinate system.

Since hippocampal place cells show different representations across different environments [254, 177, 357], it must be assumed that many two-dimensional charts exist within the pyramidal cell layer, with only one being active at any time. Any CA3 cell would belong to many different maps, and would have as many different neighborhoods as many charts it belongs to. Since each map has associated a different vector representation determined by a precisely organized connection scheme, the following question arises: When (and how) is such synaptic organization established in CA3? The hypothesis that these connections might be learned through animal's experience is unlikely, because place fields appear almost immediately when the animal enters an unfamiliar environment [147, 394]. Thus, McNaughton and colleagues assume that the entire set of CA3 charts is prewired somehow during development, so that the animal disposes of a collection of maps to use in different environments. Another basic assumption concerns the fact that each chart is limited in size, in the sense that there is a maximal magnitude of the distance from the reference point that can be represented. Exceeding this limit is one of the factors that induce the selection of a new reference frame.

The cell activity within a chart is based on an attractor dynamics producing stable and localized firing patterns. At any time, one particular reference frame is active and supports the animal's spatial behavior. Which chart is selected depends on the collection of cells that are active: a chart with a well localized activity pattern will be more likely selected than a chart exhibiting a scattered activity. Under this assumption, a disoriented animal entering an unknown environment would select an arbitrary preconfigured chart. Then, external information (V) would be associated by correlational learning to locations in the current frame, and, on subsequent visits, would be responsible for recalling the appropriate chart. Below is a scheme to illustrate the model:

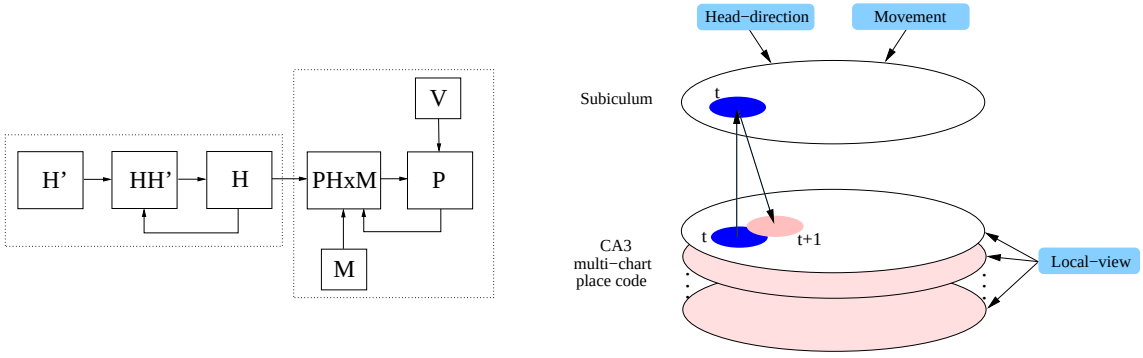

Within the active map, a blob of activity tracks the animal's position over time. The mechanism for shifting this peak over the network is a mere extension, from one to two dimensions, of the head-direction model postulated by McNaughton *et al.* in 1991 [210]. The path integrator circuit involves, beyond the CA3 layer (P), a population of PHxM cells coding for *(i)* the animal's current position (P), *(ii)* the animal's current heading (H), and *(iii)* the current self-movement information (M). The firing activity of each cell $i \in$ PHxM is correlated with all these three variables simultaneously. The authors observe that neurons in the subiculum, parasubiculum, and dorsal presubiculum have properties that fit the above PHxM theoretical cells [323]. In the model, PHxM cells and CA3 cells form a loop to update the spatial map: On the one hand, PHxM are told the animal's current position $\vec{p}(t)$ by CA3 place cells (via CA1). On the other hand, CA3 cells are informed about the future position $\vec{p}(t+1)$ by PHxM cells. This is obtained by means of prewired offset projections from PHxM to CA3. These offset connections are responsible for shifting the CA3 activity peak according to the current angular and linear motion components.

Samsonovich and McNaughton, in 1997 [301], have validated the above path integration model through numerical simulation. Their results fit several experimental data including doubling, reshaping, and vanishing of place fields in distorted environments [253], arising of place field directionality in a two-target shuttling task [132], rapid place field formation in new environments [147, 394], and slow place field rotation after animal disorientation [166].

A relevant question about the multi-chart model is: How many stable independent maps can be formed in a recurrent network such as CA3? Samsonovich [300] showed that the maximal number of uncorrelated attractor maps that can be stored in a Hopfield-like network is $0.0042 \cdot N$, where N is the number of network units. According to this, the rodent CA3 substrate (which contains approximatively $3 \cdot 10^5$ pyramidal cells [7]) may store about 1000 charts.

The authors identify the place module P with the CA3 area, since the recurrent nature of CA3 makes it a plausible candidate for both the attractor and the multi-chart properties, and the module PHxM with the subiculum (SC), since neurons in the latter are correlated to place, direction, and movement. Thus, the path integrator is formed by a CA3↔SC loop. However, as far as we know, these two structures are not directly interconnected. The subiculum receives place code information from the hippocampus via CA1, and projects efferents to the postsubiculum, the parasubiculum, the entorhinal cortex and other structures via the fornix, but not to the hippocampus [170, 171, 399, 400].

The model predicts that damaging the hippocampus would prevent rodents from path integrating. However, experiments by Alyan *et al.* [4] show that hippocampal rats have intact homing capabilities based on path integration.

To conclude, similar approaches suggesting a map-based path integrator scheme were postulated by Droulez and Berthoz in 1991 [90], and Zhang in 1996 [405]. However, these models were not explicitly simulated.

7.2.3.2 Sharp (1997)

Sharp [319, 320] adopts McNaughton and colleagues' idea of a path-integrator loop involving *(i)* an attractor network P coding for spatial locations, *(ii)* a set PHxM of cells encoding space, heading, and movement simultaneously, *(iii)* a connection scheme in which P tells PHxM the current location $\vec{p}(t)$, and PHxM sends back offset projections to inform P about the next animal's position $\vec{p}(t+1)$. However, in contrast to the multi-chart hypothesis by McNaughton *et al.*, Sharp suggests a unique preconfigured attractor network P that works as a universal place code used by the animal to fit all the environments it visits.

The model identifies the entorhinal cortex as the area in which the single-chart representation (P) takes place. Indeed, experimental findings suggest that entorhinal place cells have environment-independent activity patterns, and that they are not sensitive to cue card removal [278]. The model takes the subiculum as PHxM component of the path integrator, according to Sharp's experimental data suggesting that space, head orientation, and movements are simultaneously encoded by subicular activity [323]. In addition, Sharp showed that place cells in the subiculum are able to transfer a single, abstract spatial representation from one environment to another, and that they adapt such a universal map to fit the boundaries of the current environment [323, 319, 321]. In the model, Sharp assumes that the entorhinal cortex receives sensory information about the environmental size, and that this allows the system to center the entorhinal-subicular representation within the current environment, as well as to adjust the size of the universal map properly. Below is a functional scheme of the path integration model:

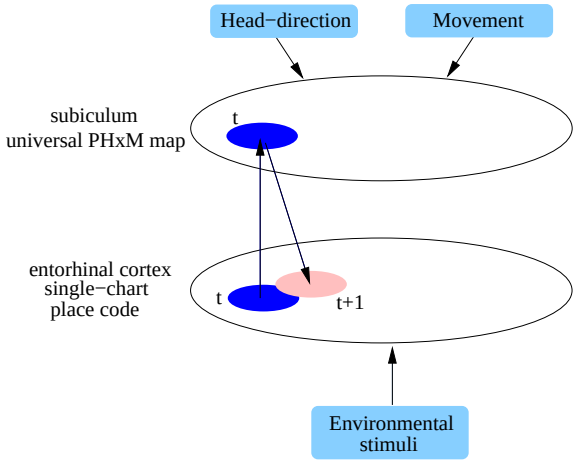

Sharp suggests that the path integration representation might converge onto the hippocampus proper to maintain a coherent place field activity despite ambiguous or absent exteroceptive inputs. Conversely, hippocampal place cells could help to reinitialize/stabilize the subicular-entorhinal representation over repeated episodes in the same environment.

In contrast to the hippocampal theory by McNaughton *et al.* [207, 301], this model predicts that lesions to the hippocampus should not impair path integration. As previously stated, this is consistent with experiments by Alyan *et al.* [4] in which hippocampal animals performed homing tasks in dark conditions as efficiently as control rats.

To conclude, the hypothesis proposed by Sharp has not been implemented neither numerically nor by a neural machinery, since the author does not report any simulation results validating the model.

7.2.3.3 Redish and Touretzky (1997)

Redish and Touretzky [285], as McNaughton *et al.* [207, 301], propose a path integration system relying on a two-dimensional attractor scheme. However, in agreement with Sharp [319, 320], they find the hypothesis of a multi-chart representation preconfigured in the hippocampus implausible and unnecessary complex. Therefore, they adopt the idea of a universal single chart updated by means of a path integrative process. The following scheme illustrates the model:

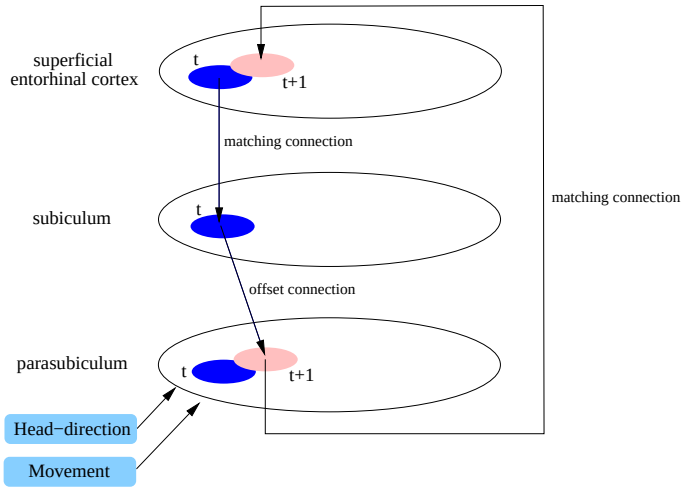

In the proposal by Redish and Touretzky, path integration is accomplished through a circuit involving three brain regions, namely the subiculum, the parasubiculum, and the superficial entorhinal cortex.

Experimental findings confirm the existence of the necessary anatomical connections for such a loop [171, 400, 138]. In the model, cells in the superficial entorhinal cortex project onto the subiculum via matching connections, the subiculum sends offset efferents to the parasubiculum, and the latter projects matching connections to the entorhinal cortex.

The authors stress the importance of head-direction as well as self-motion information in order to accomplish path integration. They emphasize the fact that both these inputs might enter the system via projections to the parasubiculum from the postsubiculum (head-direction) [138] as well as from the parietal and cingulate cortex (self-motion) [401]. However, in the model, these inputs are not utilized by the parasubiculum; rather, they are employed by the subiculum, which projects offset afferents to induce the position representation update from $\vec{p}(t)$ to $\vec{p}(t+1)$. In the model, the parasubiculum is told the new position $\vec{p}(t+1)$ from the subiculum, and it simply "closes the loop" by forwarding this information to the entorhinal cortex via matching connections. The authors disfavor Sharp's model [320] (in which the new location $\vec{p}(t+1)$ is directly sent from the subiculum to the entorhinal cortex) by arguing that there are no anatomical projections from subicular cells to the superficial layers of the entorhinal cortex. Indeed, they refer to early anatomical data by Kohler [170, 171] and Witter *et al.* [400] reporting that the subiculum projects only to deep layers of the entorhinal cortex. However, more recent studies have shown that the entorhinal cortex, in particular its medial area, does receive a prominent projection from the hippocampal formation (in particular from CA1 and subiculum) that distributes to the superficial layers I-III [397].

Finally, the above model predicts that *(i)* parasubicular cells should provide similar place field distributions across different environments (similar to subicular cells), *(ii)* lesions to the hippocampus should not disrupt path integration, *(iii)* damaging the subiculum, and/or the parasubiculum, and/or the superficial entorhinal cortex, should impair idiothetic-based homing behavior.

Chapter 8

Allothetic Space Representation: Processing Visual Information

Hippocampal place fields are determined by a combination of highly-processed multimodal sensory stimuli (e.g., visual, auditory, olfactory, and somatosensory cues) whose mutual relationships code for the animal's current location [256, 148, 92, 305]. Experiments on rodents suggest that vision plays an eminent role in determining place cell activity [101, 258, 240, 211, 166]. In this chapter, we focus on the visual pathway of our hippocampal model (Fig. 8.1), and we ask the question: How can place fields be established from real visual input [253]?

The vision-based localization problem implies extracting relevant information from *noisy* visual stimulation. We put forward a computational strategy to emulate the feature-extraction mechanism observed in the visual cortex. Moving up the visual pathway (from the retina to the lateral geniculate nucleus and then towards higher cortical areas), neurons become responsive to stimuli of increasing complexity, from orientation-sensitive cells (*simple cells*), to neurons responding to more complex patterns, such as faces [151, 143, 291, 364]. In our model, visual stimuli are interpreted by means of neurons that only respond to combinations of specific visual patterns. Due to this filtering process, properties like agent-landmark distance and egocentric orientation to visual cues can be measured *implicitly*, without explicit image processing. The activity of the neural filters propagates through the model yielding place cell firing.

Due to the high dimension of the visual input space, vision-based localization consists of *(i)* detecting a convenient low-dimensional representation of the continuous high-dimensional input space, *(ii)* learning the mapping function from the visual sensory space to points belonging to this representation. Since our animat moves on a two-dimensional space with a camera pointing in the motion direction, the high-dimensional visual space is not uniformly filled. Rather, all input data points lie on a low-dimensional surface embedded in a Euclidean space whose dimensionality is given by the total number of camera pixels. This low-dimensional description of the visual space is referred to as *view manifold* [109].

In the model, unsupervised Hebbian learning is used to detect the low-dimensional manifold representing the visual input space. The system consists of a multi-layer neural architecture modeling high-dimensional continuous inputs by means of overlapping place fields. Spatio-temporal properties of the environment are extracted from the visual stream to build the place field representation. Starting with no prior knowledge, the system grows incrementally and on-line as the agent interacts with the environment.

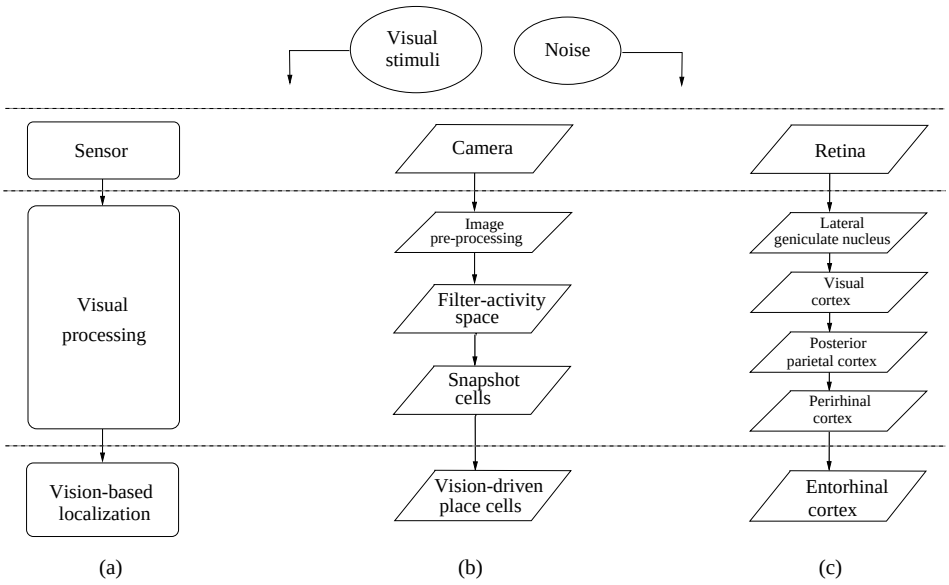

Figure 8.1: (a) Vision-based localization implies interpreting noisy visual information effectively. (b) Functional decomposition of the visual pathway in our hippocampal model. (c) The corresponding (highly simplified) anatomical decomposition.

In this chapter, we describe three robotic implementations that have been adopted to validate the visual pathway of the model. Our visual interpretation can be seen as a four-step processing:

(i) The first step consists of designing a set of filters to detect visual features. This phase strongly depends on the physical implementation and consists of choosing appropriate filters to interpret visual data.

(ii) As a second step, we map real images into the filter-activity space. This consists of computing the magnitude of the response of each filter for a given image.

(iii) Third, we map filter responses into neural activity. We take a population of hypothetical visual cells one synapse downstream from the filter layer. The activity of these neurons, termed *snapshot cells* (SnC), encodes the current visual input.

(iv) As a fourth step, we apply unsupervised Hebbian learning to achieve spatial coding. We consider a population of place cells one synapse downstream from the SnC layer. We suppose that the anatomical counterpart for this neural substrate is the superficial layer of the lateral entorhinal cortex (sLEC). Superficial layers of the entorhinal cortex receive spatial information about allothetic landmarks (e.g., visual cues) from the posterior parietal cortex, and contain location-sensitive neurons [278].

As we move up our visual pathway (from step *(i)* to *(iv)*), visual processing becomes independent with respect to the physical properties of the system. Neurons at the highest level (sLEC place cells) receive highly processed information, similar to biological entorhinal neurons that receive highly processed inputs from the posterior parietal cortex (Chapter 6).

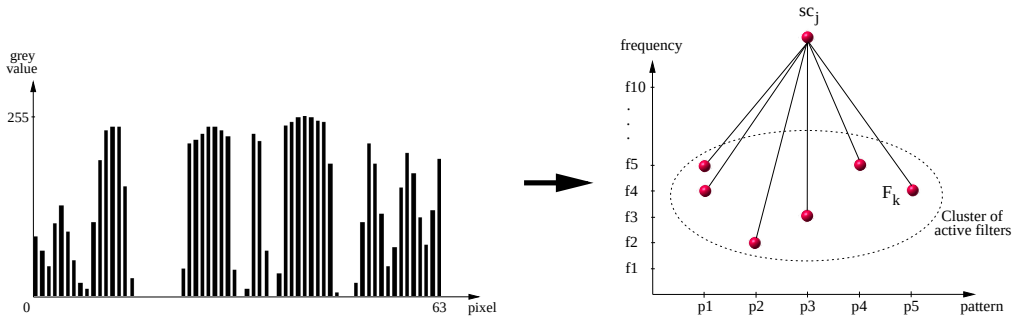

Figure 8.2: (a) A sample of linear image. (b) The filter-activity space. Along the x-axis are different Walsh-like filters, each of which responds to a specific pattern. Along the y-axis we vary the spatial frequency f_i of each pattern p_i. Each image is encoded by the cluster of filters maximally responding to that image.

8.1 Interpreting Visual Information

In this section we describe how stages *(i)*, *(ii)*, and *(iii)* can be accomplished in the case of both linear and two-dimensional visual information. Step *(iv)* is described in Sec. 8.2.

8.1.1 Processing Linear Visual Data

The linear-vision experimental setup consists of a high-walled square arena, with *randomly distributed* black and white stripes of variable width on the walls (Fig. 1.2 (a)). These stripes provide the visual input patterns to the system. Visual signals consist of 64-pixel one-dimensional images. Each pixel has a grey-level value within $[0,\ 255]$ (Fig. 8.2 (a)). The azimuthal visual range covers approximately $36°$. Image pre-processing consists of rescaling pixel values of input images within $[-1,\ 1]$. Thus, if $\mathbf{I}(x)$, with $0 \leq x \leq 63$, denotes an input image, the relation $-1 \leq \mathbf{I}(x) \leq 1\ \forall x \in [0,\ 63]$ holds.

We define several classes of *Walsh*-like binary filters [11]. Walsh filters are simple and permit effective and low-cost feature-detection in one-dimensional visual spaces. Each class corresponds to a specific visual pattern p_i, and the set of filters in that class corresponds to different spatial frequencies f_i for pattern p_i. This endows the system with a distance discrimination property. In total, we take 5 different classes of filters each of which contains filters at 10 different frequencies. As a consequence, a two-dimensional filter space including 50 distinct elements is established.

Let F_k be one of our Walsh filters, where $1 \leq k \leq 50$ is the index of the filter, and let l_k denote its length (i.e., number of pixels covered by the filter). In general, F_k can be seen as a simple binary function defined in $[0,\ l_k]$:

$$F_k(x) = \begin{cases} 1 & 0 \leq x \leq l_1 \\ -1 & l_1 < x \leq l_2 \\ 1 & l_2 < x \leq l_k \end{cases}$$

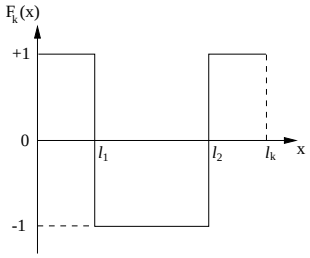

where $0 < l_1 < l_2 < l_k$, and $l_1 = l_k - l_2$.

The aim is to interpret incoming visual stimuli by mapping them into the above filter-activity space.

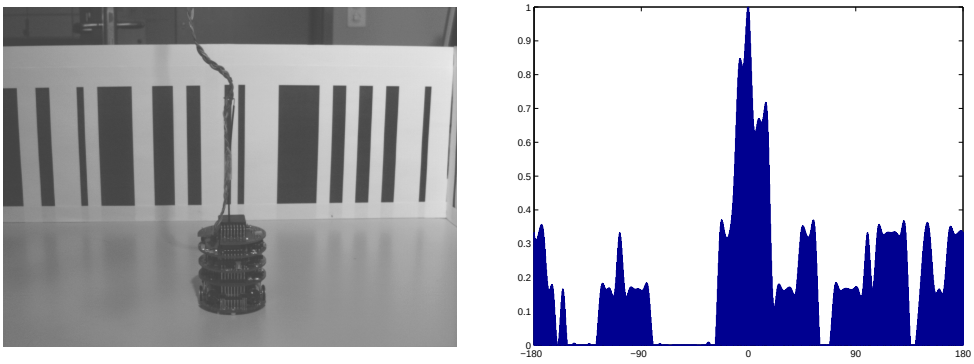

(a) (b)

Figure 8.3: (a) At time t_0, the robot takes a view and creates a snapshot cell sc_j. (b) Firing activity of cell sc_j as a function of the difference (in degrees) between the robot's current heading and its orientation at time t_0 (i.e., when it created sc_j).

Given an image $\mathbf{I}(x)$, the response a_k of filter F_k is computed by convolution

$$a_k = \max_n \Big\{ \sum_{x=0}^{l_k-1} F_k(x)\ \mathbf{I}(n+x) \Big\} \tag{8.1}$$

where $0 \leq n \leq 64 - l_k$. Since $\mathbf{I}(x) \in [-1,\ 1]\ \forall x$, and $F_k(x) = \pm 1\ \forall k, x$, the relationship $|a_k| \leq l_k$ holds.

Each filter F_k responds to a particular visual pattern. In order to model spatio-temporal relationships between visual cues by means of neural activity, we consider a layer of snapshot cells one synapse downstream of the filter layer. The idea is to represent each image $\mathbf{I}(x)$ by the cluster of maximally active filters, defined by Eq. 8.1. Let $C_k = 0.7 \cdot l_k$ be the threshold above which a filter F_k is considered as active. Given an image $\mathbf{I}(x)$, the set of active filters projects one layer forward to form a snapshot cell sc_j (Fig. 8.2 (b)):

$$sc_j = \{F_k \mid a_k \geq C_k\} \tag{8.2}$$

The firing activity r_j of a snapshot cell sc_j is given by

$$r_j = \frac{\sum_{k \in sc_j} \mathcal{H}(a_k - C_k)}{N_j} \tag{8.3}$$

where $\sum_{k \in sc_j}$ sums over all the N_j filters projecting to the cell sc_j, and $\mathcal{H}$ is the Heaviside function. The normalization has been chosen so that $0 \leq r_j \leq 1$.

Each snapshot cell sc_j receives afferent inputs from several visual-feature detectors. This allows the system to extract more complex environmental properties from visual information. As a consequence, the neural activity r_j can be employed to discriminate between views taken by the agent during exploration. In other words, the snapshot cell population works as a neural image classifier. However, snapshot cell activity does not provide allocentric spatial selectivity because it depends on the agent's gaze direction. Fig. 8.3 shows a result obtained by placing the robot perpendicularly to a wall and letting it take a snapshot at time t_0 (Fig. 8.3 (a)). Let sc_j denote the cell encoding that view. Fig. 8.3 (b) shows the firing activity r_j of cell sc_j as a function of the difference between the robot's current heading $\theta(t)$ and its heading $\theta(t_0)$ when it created the cell sc_j. That is, the robot turns on the spot taking a view each 4°, and the response of cell sc_j to each new image is measured.

8.1.2 Visual Processing based on Principal Component Analysis

The experimental setup used for the two-dimensional visual interpreter consists of an open-field square arena of approximately $80 \times 80\,cm$ within a standard laboratory background (Fig. 1.2 (b)). Visual input is provided by a black and white CCD camera having a view field of about 90° in the horizontal plane and about 60° in the vertical plane (Fig. 1.1 (c)). Image pre-processing involves histogram equalization[1] and resolution reduction (from 768×576 to 115×86 pixels). In addition, the 256 grey levels of original images are mapped into the $[-1, 1]$ range. Thus, working images $\mathbf{I}(x, y)$ have a resolution of $L_x \times L_y = 115 \times 86$ pixels, and the relation $-1 \leq \mathbf{I}(x, y) \leq 1$ holds.

In order to detect the low-dimensional manifold embedded in the $L_x \times L_y$ high-dimensional view space, we employ filter-based decomposition (similar to the approach adopted for the linear case). Thus, we take a class P of 10 two-dimensional filters f_i, each of which responds maximally to a specific visual pattern. Each filter has a size of 32×32 pixels[2]. To define the set of basic filters f_i, we consider the first 10 principal components extracted from real visual data.

Extracting principal components from natural images

Principal component analysis (PCA) (or Karhunen-Loeve transform) is a statistical method largely applied for dimensionality reduction [249, 160, 302]: PCA determines a linear transformation of the input data distribution such that high-dimensional inputs are projected onto a limited number of mutually-orthogonal linear descriptors (i.e., principal components), and such that the most information about the original data is conveyed.

Computing the principal components exactly corresponds to finding the eigenvectors of the correlation matrix of the input data. Since we want to compute principal components of 32×32 image pieces, we have an input space of 1024 variables, and then a correlation matrix with $1024 \times 1024 = 2^{20}$ entries, which is computationally too expensive. Therefore, we resort to Sanger's neural network technique [302, 140] to approximate principal components through off-line unsupervised learning. Sanger's algorithm permits the extraction of the first N principal components in sequence, and is a generalization of the neural network technique proposed by Oja [249, 250].

We let the robot collect 100 images $\mathbf{I}(x, y)$, and we build a training set by randomly selecting 10000 square samples $\mathbf{I}'(x, y)$ of 32×32 pixels from these images. Before training, samples $\mathbf{I}'$ are pre-processed by [140]: *(i)* subtracting the mean grey value (averaged over all 10000 patterns) from each pixel intensity, *(ii)* applying a Gaussian mask of width $\sigma = 7$ to reduce edge effects, *(iii)* normalizing all input vectors such that $\left(\sum_{xy} \mathbf{I}'(x, y)\right)^{1/2} = 1$.

We train a feed-forward neural network consisting of $32 \times 32 = 1024$ input units x_j and 10 output units y_i. Weights w_{ij} are initialized to small random values such that $\sum_j w_{ij}^2 \approx 1$. Output units are linearly activated according to $y_i = \sum_j w_{ij} x_j$. Learning is defined by [302]

$$\Delta w_{ij} = \eta\, y_i \Big(x_j - \sum_{k=1}^{i} w_{kj} y_k \Big) \tag{8.4}$$

The learning rate η is initialized to 2.0 and then decreased during training (in particular, it is halved every 1600 steps of learning).

[1] Histogram equalization is a form of image enhancement particularly useful when images suffer from poor contrast. The histograms of such images would have relatively narrow curves around a certain range of pixel values and not at the others. Histogram equalization consists of adjusting the ranges of pixel values such that each value has approximately the same number of image pixels.

[2] The size of the filters has been chosen empirically.

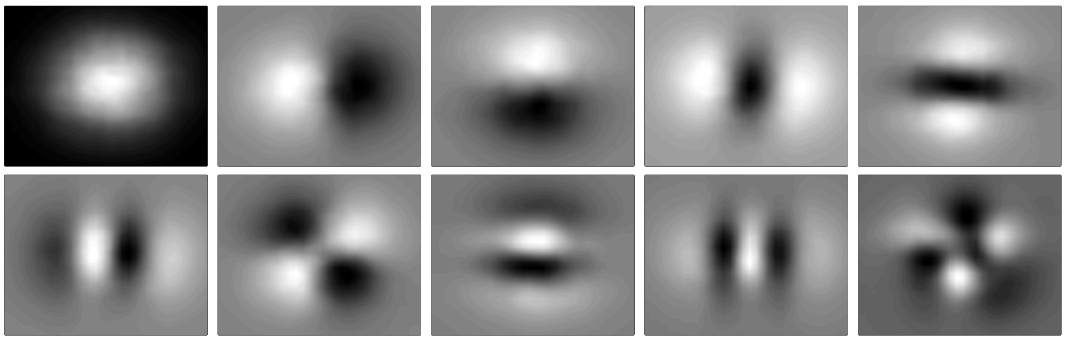

Figure 8.4: The 10 filters f_i corresponding to the first 10 principal components (numbered from left to right, top to bottom) learned by applying Sanger's neural network technique [302].

After training, each i_{th} weight vector $\vec{w}_i$ corresponds to the eigenvector of the input correlation matrix which has the i_{th} maximal eigenvalue [302]. Thus, the receptive field of the output unit y_i corresponds to the i_{th} principal component of input data. Fig. 8.4 shows the receptive fields of the 10 output units y_i after training. Neurons having receptive fields similar to those displayed in Fig. 8.4 have been experimentally observed in areas belonging to the biological visual pathway. Cells in the lateral geniculate nucleus (LGN) as well as retinal ganglion cells exhibit concentric center-surround antagonist receptive fields [364] similar to the first principal component. Simple cells in the primary visual cortex are also sensitive to on-off concentric stimuli as well as to oriented bars [106, 20] (similar to the other principal components in Fig. 8.4).

Modeling visual input

We take the 10 first principal components extracted by the above learning scheme to form our class P of basic 32×32 filters f_i. Then, for each filter $f_i \in P$, we consider 5 different spatial scales f_{is}[3]. Experimental findings on mammals suggest that the biological vision system uses filters with different spatial scales in the early visual-sensory processing [379].

Given a filter f_{is}, we compute its response a_{is} to an input image $\mathbf{I}$ by shifting the filter across the image and looking for the location of maximal matching between f_{is} and $\mathbf{I}$. That is,

$$a_{is} = \max_{x,y} \Big(\sum_{i=0}^{l_{is}-1} \sum_{j=0}^{l_{is}-1} f_{is}(i,j) \cdot \mathbf{I}(i+x,\, j+x) \Big) \tag{8.5}$$

where $-1 \leq \mathbf{I}(x,y) \leq 1$, $-1 \leq f_{is}(i,j) \leq 1$, $0 \leq x \leq L_x - l_{is}$, and $0 \leq y \leq L_y - l_{is}$.
The activity of a filter f_{is} is then normalized with respect to its potential maximal response

$$a_{is} = \frac{a_{is}}{\sum_{i=0}^{l_{is}-1} \sum_{j=0}^{l_{is}-1} |f_{is}(i,j)|} \tag{8.6}$$

Let (x^*, y^*), with $0 \leq x^* \leq L_x - l_{is}$ and $0 \leq y^* \leq L_x - l_{is}$, identify the location of best match. We segment images in 4 quadrants $\mathbf{I}_q$ and we utilize the spatial information (x^*, y^*) to characterize each filter f_{is} by the region of the image in which it has detected its preferred visual pattern. In other words, we interpret visual data by mapping images $\mathbf{I}$ into a three-dimensional filter-activity space $\mathcal{F}$ consisting

[3]Given the original size of a filter f_i, that is $l_i \times l_i = 32 \times 32$ pixels, we take 5 filters of size $l_{is} \times l_{is}$ where $l_{is} \in \{20,\, 12,\, 10,\, 8,\, 5\}$. Rescaling factors have been determined empirically.

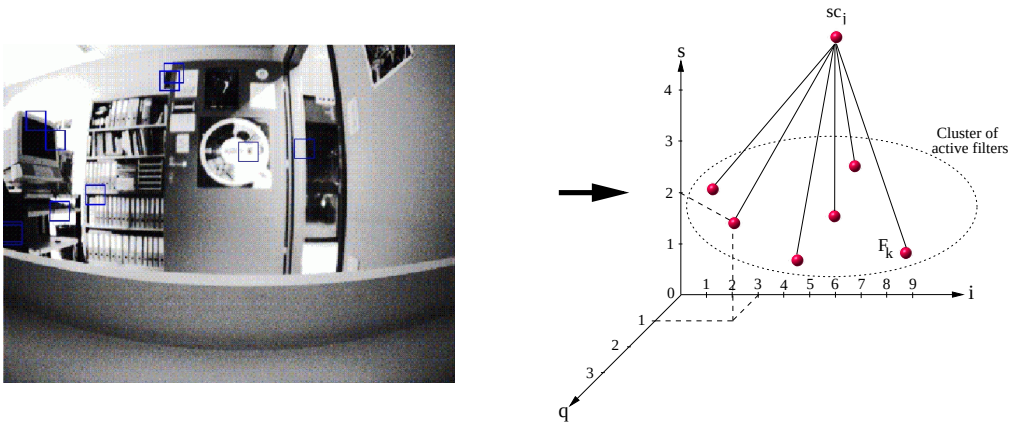

(a) (b)

Figure 8.5: Visual data are interpreted by mapping images into a filter-activity space. (a) A sample of image taken by the robot during exploration. Small black rectangles indicate features that have been detected by our small-scale filters. (b) The filter activity space is defined by taking 10 basic filters (i axis), considering 5 scales for each of them (s axis), and accounting for the quadrant of the image where each filter is maximally activated (q axis). Given an image **I**, the cluster of maximally active filters is used to generate a neural unit sc_j which is the internal representation of **I**.

of $F = 10 \times 5 \times 4 = 200$ elements f_{is}^q (Fig. 8.5). Let k be an index over $\mathcal{F}$, that is $0 \leq k \leq F$. Each element f_k identifies a unique neural filter $f_{is}^q \in \mathcal{F}$ responding to a specific localized pattern within **I**.

In order to detect more complex visual features, we consider a layer of snapshot cells one synapse downstream from the filter layer. As in Sec. 8.1.1, given an image $\mathbf{I}(x, y)$, the cluster of active filters is used to generate a neural unit sc_j which is the internal representation of **I**. Again, a threshold C_k is used to select maximally active filters, i.e., f_k such that $a_k \geq C_k$. Then, Eqs. 8.2 and 8.3 are used to recruit and drive, respectively, the cell $sc_j \in SnC$.

PCA-based filtering yields effective detection of visual properties and induces robust vision-based self-localization (see Sec. 8.2). However, processing real images according to Eq. 8.5 is computationally expensive. Given an image $\mathbf{I}(x, y)$, each filter f_{is} has to be shifted "pixel-by-pixel" across the image, and this process has to be repeated $\forall i$ and $\forall s$, with $0 \leq i \leq 9$, and $0 \leq s \leq 4$, respectively. For each input image, this yields a number n of multiplications given by $n = \sum_{i,s} l_{is}^2 \cdot (L_x - l_{is} + 1) \cdot (L_y - l_{is} + 1)$. On a SUN ultra 10, this results in a computation time of approximately $20\,s$ for processing a single image. In the next section we put forward an alternative approach that accounts for efficiency of the visual data interpretation.

8.1.3 Gabor-decomposition for Visual Feature Extraction

In this section, we describe a second approach to detect visual properties from two-dimensional video streams. We consider the same experimental setup as above: a $80 \times 80\,cm$ open-field arena, and a robot with a black and white camera on board. Visual scenes correspond to standard laboratory background. Working images $\mathbf{I}(x, y)$ are obtained from original data through histogram equalization, resolution reduction (from 768×576 to 422×316 pixels), and grey-value remapping into $[-1, 1]$.

We employ Gabor visual decomposition [113] to discover the mapping from the high-dimensional visual input to the low-dimensional view manifold [332]. Images are interpreted by means of a retinotopic strategy that selects a limited number of image-points to be sampled by the filtering process.

Defining a set of (modified) Gabor filters

Gabor filters [113] provide a suitable mathematical model for biological simple cells [82, 161]. They are generated by the two-dimensional Gabor function

$$G(x, y|\sigma_x, \sigma_y, \omega_0) = e^{-x^2/2\sigma_x^2} \cdot e^{-y^2/2\sigma_y^2} \cdot e^{i\omega_0 x} \tag{8.7}$$

This function produces a complex sinusoidal wave embedded in a Gaussian envelop centered at the origin. The angular frequency of the wave is ω_0, while σ_x and σ_y define the horizontal and vertical widths of the Gaussian window, respectively. The complex sinusoidal wave is parallel to the x axis, that is, the filter G responds selectively to vertical structures (e.g., vertical edges in the image). Therefore, to detect structures having different orientations and scales, we need to generate a family of two-dimensional Gabor filters by rotating and rescaling G.

In the frequency domain, the function G corresponds to a Gaussian $\hat{G}$ translated along the horizontal frequency axis ω_x by an offset ω_0. The horizontal and vertical frequency widths of function $\hat{G}$ are given by $\sigma_{\omega_x} = 1/\sigma_x$ and $\sigma_{\omega_y} = 1/\sigma_y$, respectively. Thus, the Fourier representation $\hat{G}$ of G is defined by

$$\hat{G}(\omega_x, \omega_y|\sigma_x, \sigma_y, \omega_0) = A \cdot e^{-(\omega_x-\omega_0)^2/2\sigma_{\omega_x}^2} \cdot e^{-\omega_y^2/2\sigma_{\omega_y}^2} \tag{8.8}$$

where A is a normalization factor.

In order to extract visual features effectively, we employ a family of *modified* Gabor filters [34, 332, 331]. A modified Gabor filter f_i, tuned to orientation ϕ_j and angular frequency $\omega_l = e^{\xi_l}$, corresponds to a Gaussian function in the *Log-polar* frequency plane rather than in the frequency domain itself, and is defined by the Fourier function

$$\hat{G}'(\xi, \phi) = A \cdot e^{-(\xi-\xi_j)^2/2\sigma_\xi^2} \cdot e^{-(\phi-\phi_l)^2/2\sigma_\phi^2} \tag{8.9}$$

where (ξ, ϕ) are coordinates in the Log-polar Fourier plane

$$(\xi, \phi) = \big(log||(\omega_x, \omega_y)||,\ \arctan(\omega_y/\omega_x)\big) \tag{8.10}$$

A key property of the Log-polar coordinate system is that translations along ϕ correspond to rotations in the image domain, while translations along ξ correspond to scaling the image. As a consequence, generating a family of modified Gabor filters by varying scale and orientation parameters, corresponds to arranging a set of identical Gaussian functions within a rectangular grid in the Log-polar frequency domain. This simplifies the designing process significantly (see [331] for further details). Another important benefit provided by modified Gabor filters is that they yield a more uniform coverage of the frequency plane. By contrast, pure Gabor wavelets tend to densely cover low-frequency regions of the Fourier plane, while provide coarse coverage of the high-frequency regions [331].

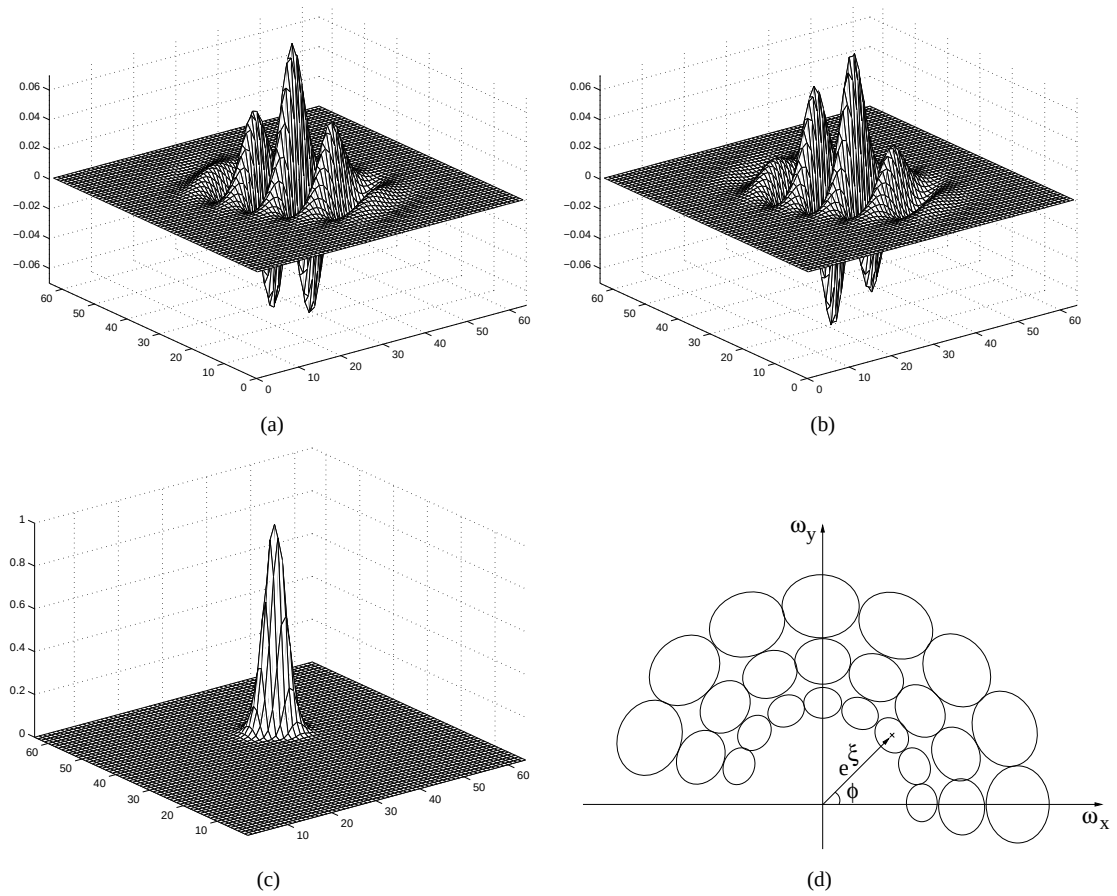

Figure 8.6: (a, b, c) An example of modified Gabor filter (taken from [331]): (a, b) show the real and imaginary part in the image domain of the filter, respectively, whereas (c) represents the filter in the frequency plane. (d) The set of modified Gabor filters, shown in the Fourier domain, employed for the visual feature detection task.

We employ a family of 24 modified Gabor filters, $\mathcal{F} = \{f_i(\omega_l, \phi_j) \,|\, 1 \leq l \leq 3, 1 \leq j \leq 8\}$, obtained by taking 3 distinct angular frequencies ω_1, ω_2, ω_3, and 8 orientations $\phi_1, \ldots, \phi_8$. The three angular frequencies ω_l have been determined by estimating three filter wavelengths λ_l suitable for our application, and then using the relation $\omega_l = 2\pi/\lambda_l$. The values $\lambda_1 = 8$, $\lambda_2 = 16$, and $\lambda_3 = 32$ have been chosen. The 8 orientations ϕ_j are evenly distributed over the range $[0, \pi]$, that is, $\phi_j = j\pi/8$. Figs. 8.6 (a, b) show the real and the imaginary waves, respectively, of a typical filter $f_i \in \mathcal{F}$ in the image domain. Fig. 8.6 (c) represents the filter f_i in the frequency plane. Finally, Fig. 8.6 (d) shows the entire set $\mathcal{F}$ of modified Gabor filters in the standard Fourier domain.

Retinotopic image sampling

Efficiency is a crucial issue in vision-based self-localization. We apply a *space-variant retinotopic sampling* [360] in order to perform visual feature extraction effectively (in contrast to the uniform sampling strategy employed in Sec. 8.1.2). We place a retinotopic grid on the image (Fig. 8.7), and we use the N_p points of this "retina" to sample visual data by means of Gabor decomposition [333]. The approach consists of having a high resolution of points only in a localized region of the view field (*fovea*), whereas peripheral areas are characterized by a low-resolution vision. Thus, the retinotopic graph is constructed

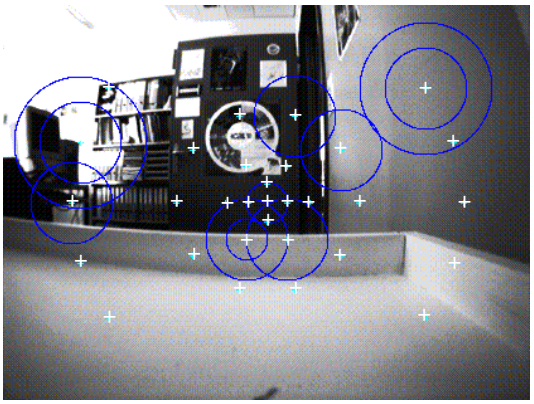

Figure 8.7: A visual scene acquired by the robot. The image resolution is 422×316 pixels. An example of retinotopic sampling grid (white crosses) employed to sample visual data by means of Gabor decomposition. In this example, the retina consists of $N_p = 31$ points arranged on $N_c = 5$ circles. Black circles represent maximally responding Gabor filters (the circle radius varies as a function of the filter's amplitude).

by arranging the N_p points on N_c concentric circles. The innermost circle (the center of the retina) has radius zero and coincides with the center of the image. The radii of the remaining $N_c - 1$ circles increase exponentially. On each circle, the retinal points are evenly distributed.

At each retinal point $\vec{g}$ we place the 24 modified Gabor filters $f_i \in \mathcal{F}$ with different orientations and amplitudes. This results in a population of overlapping Gaussian receptive fields that tend to cover the entire image continuously. The density of the coverage is higher at the center of the image and decreases by moving towards the peripheral regions of the image.

Given an image $\mathbf{I}(x, y)$, we take the *magnitude* of the responses of the Gabor filters to detect visual properties within video streams. Thus, we compute the magnitude of the response of all f_i filters for each retinal point $\vec{g}$

$$r_i(\vec{g}) = \left(\Big(\sum_{\vec{x}} Re(f_i(\vec{x})) \cdot \mathbf{I}(\vec{g} + \vec{x}) \Big)^2 + \Big(\sum_{\vec{x}} Im(f_i(\vec{x})) \cdot \mathbf{I}(\vec{g} + \vec{x}) \Big)^2 \right)^{\frac{1}{2}} \tag{8.11}$$

where $\vec{x}$ varies over the area occupied by the filter f_i in the spatial domain.

Due to the sparseness of the sampling graph, the retinotopic strategy allows the agent to perform Gabor-based feature extraction rather efficiently. Two factors influence computational costs: The structure of the retina (i.e., number of circles, and number of retinal points per circle), and the distribution of Gabor filters f_i over the retinal points (i.e., number of filters per point). We present experimental results obtained by employing a retina with $N_p = 17$ points distributed over $N_c = 4$ circles[4], and by placing all the 24 modified Gabor filters at each retinal point. The resulting computation time for processing an image on a SUN ultra 10 is approximately 1.25 seconds.

Modeling visual input by neural activity

The third step in the visual pathway of our model consists of interpreting visual cues by means of neural activity. As in Secs. 8.1.1 and 8.1.2, we take the responses of our visual filters as neural afferents to a population of snapshot cells (SnC) one synapse downstream of the filter layer. Note that our Gabor-based

[4]The number of points per circle (from the innermost to the outermost) is as follows: 1, 4, 6, 6.

retinotopic system may be abstracted as a three-dimensional filter-activity space: The angular frequency and the orientation of each Gabor filter provide the two first dimensions, while the spatial distribution of the retinal points provides the third one (i.e., the sampled image location). Thus, we model each image $\mathbf{I}(x, y)$ by mapping its representation $\mathbf{I}_f$ in the filter-activity space into snapshot cell activity.

Let k be an index over all $K = 24 \times 17$ filters forming the retinotopic grid. Given an image $\mathbf{I}$, a snapshot cell $sc_j \in$ SnC is recruited to receive afferents from all f_k filters. Synaptic connections w_{jk} from filters f_k to cell sc_j are initialized according to

$$w_{jk} = r_k \quad \forall k \tag{8.12}$$

Synaptic connections to cell sc_j work as the long-term representation of the filter activity associated to image $\mathbf{I}$. If, at a later point, the robot sees an image $\mathbf{I}'$, the firing activity r_j of cell $sc_j \in$ SnC is given by

$$r_j = e^{-(\frac{1}{K}\sum_k |r_k - w_{sk}|)^2/2\sigma^2} \tag{8.13}$$

where r_k are the Gabor filter responses to image $\mathbf{I}'$. Eq. 8.13 defines a radial basis function in the filter space that measures the similarity of the current image to the image stored in the weights w_{jk}. The width σ determines the discrimination capacity of the system for visual scene recognition[5].

In contrast to the strategy applied in Secs. 8.1.1 and 8.1.2 to map filter activity into snapshot cell activity (Eqs. 8.2, 8.3), here filter responses are not thresholded. Rather, filter activity is interpreted as a continuum by means of the radial basis function defined by Eq. 8.13, and each cell sc_j receives afferents from all filters f_k.

[5]In the current implementation σ has been empirically set to 0.07.

8.2 Allothetic Space Representation: Place Fields in the Superficial Layer of the Lateral Entorhinal Cortex

In this section, we complete the visual pathway of the model (step *(iv)*). Snapshot cell activity (resulting from steps *(i)*, *(ii)*, and *(iii)*) depends on the agent's gaze direction (Fig. 8.3), and does not code for spatial locations. We apply unsupervised Hebbian learning to achieve allocentric spatial coding one synapse downstream from the SnC layer (Fig. 8.1). We assume that this neural substrate corresponds to the superficial layer of the lateral entorhinal cortex, and we call these neurons sLEC cells. Each cell $i \in$ sLEC receives afferents from a set of snapshot cells whose activities code for visual features of the environment. As a consequence, the activity r_i of a unit $i \in$ sLEC depends on the combination of multiple visual cues, which makes cell i location sensitive (i.e., a place cell). This results in a sLEC population activity providing allocentric spatial coding.

8.2.1 Unsupervised Growing Network Scheme

Every time the agent is at a new location[6], all simultaneously active snapshot cells are connected to a newly recruited sLEC cell. Let i and j be indices for sLEC cells and snapshot cells, respectively. If r_j is the firing activity of a snapshot cell j, then a connection w_{ij}^{new} is created such that

$$w_{ij}^{new} = \mathcal{H}(r_j - \epsilon) \cdot \mathrm{rnd}_{0,1} \tag{8.14}$$

where $\mathcal{H}$ is the Heaviside function, $\epsilon = 0.75$ is the activity threshold above which a snapshot cell is considered to be active, and $\mathrm{rnd}_{0,1}$ means that each new synapse w_{ij}^{new} is initialized by a random weight in $(0, 1)$.

The firing rate r_i of a place cell $i \in$ sLEC is computed by linearly averaging the activity over its SnC afferents

$$r_i = \sum_j w_{ij} r_j / \sum_j w_{ij} \tag{8.15}$$

where j varies over all snapshot cells that have been connected to cell i according to Eq. 8.14. The relation $0 \leq r_i \leq 1$ holds.

Once connections w_{ij} are established, their synaptic strength is changed by Hebbian learning

$$\Delta w_{ij} = r_j \, (r_i - w_{ij}) \tag{8.16}$$

The rationale behind Eq. 8.16 is to *(i)* induce correlational LTP whenever pre- and postsynaptic neurons are simultaneously active, *(ii)* depress the connection w_{ij} (LTD) whenever the presynaptic snapshot cell j is active while the postsynaptic sLEC cell i is not. For instance, Eq. 8.16 tends to reduce the strength of the connection from an "always activated" SnC cell j (i.e., a cell that does not discriminate visual information effectively) to a sLEC cell i. Note also that Eq. 8.16 keeps weights $w_{ij} \leq 1$.

We call the learning scheme defined by Eqs. 8.14, 8.15, and 8.16 an *unsupervised growing network scheme* (see, e.g., [110]). When the animat first enters a novel environment, it has no prior knowledge, that is, there is no place field representation at all. The population of place cells grows incrementally as a result of the animat interaction with the environment, that is, through exploration (see Sec. 10.2 for a description of the adopted exploration strategy).

[6] A simple definition of "location familiarity", based on the number of place cells active at a given position, will be given in Chapter 10.

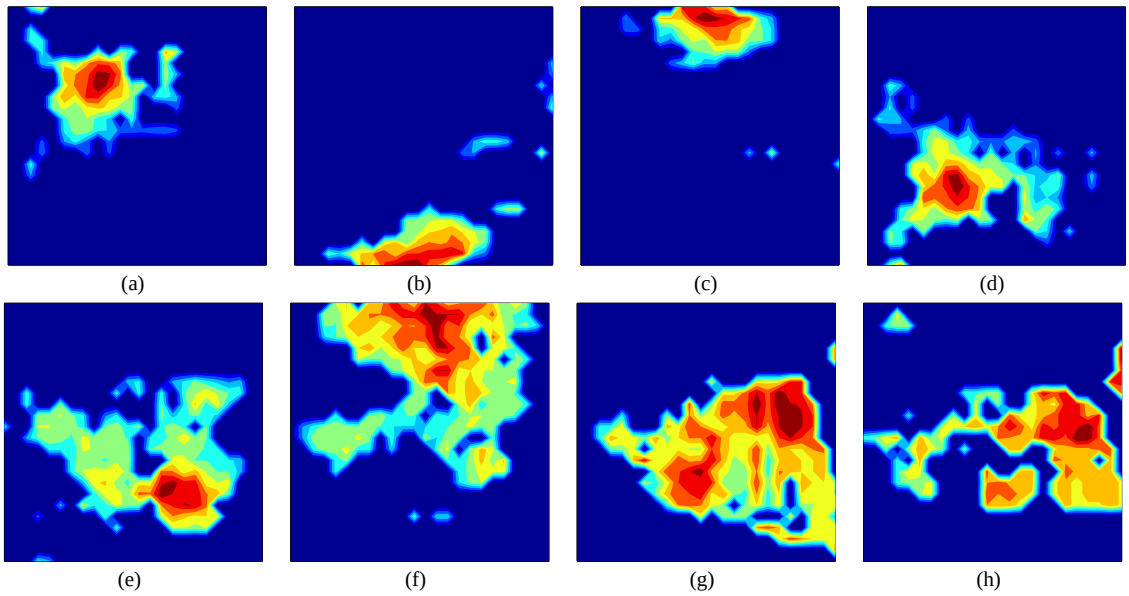

Figure 8.8: Some samples of sLEC place fields obtained with the linear-vision experimental setup. The square box represents the arena. The firing rate of a cell is maximal (center of the place field) when the agent is at the red location. The dark-blue region indicates the area in which the cell remains silent.

8.2.2 Recording Single sLEC Cell Activity

At each visited location, the robot takes four views v_1, v_2, v_3, v_4 [62], such that $v_{i+1} = v_i + 90°$. For simplicity, we shall refer to these gaze directions as north, east, south, and west, but they may actually be associated to an arbitrary allocentric reference frame. To take these views the animat strongly relies on the compass information provided by the head-direction system (Chapter 4).

For each visited location the robot creates four snapshot cells (one for each view v_i), which are bound together to form a quasi-panoramic description of the place[7]. This information is then used for driving neurons in the sLEC layer (Eqs. 8.14, 8.15, and 10.1). This results in *non-directional* sLEC place cell activity, consistently with experimental data concerning place fields recorded from freely-moving rats in open-field environments [278]. On the other hand, when recorded in a linear track maze, place cells tend to be directional [209]. In a linear track the animal always runs in the same direction. If we would model this by taking a *single* view only, then we would get directionality.

We have validated the above unsupervised growing network scheme in all three cases of visual data interpretation described in Sec. 8.1. Here, we present results obtained by recording single place cell activity from the sLEC layer of the model during robot exploration.

Fig. 8.8 shows eight samples of sLEC place fields obtained with the linear-vision experimental setup. Figs. 8.8 (a-e) show five examples of cells that are maximally activated only if the agent is in a rather localized region of the environment. Thus, the robot can use the center of the place field for self-localization. Fig. 8.8 (f) shows a cell whose receptive field is not well-localized, but still single-peak. On the other hand, Figs. 8.8 (g, h) present two examples of multiple-peak place field. The activities of these sLEC cells encode ambiguous visual inputs, that is, their receptive fields identify different spatial locations that provide similar visual stimulation. About 30% of the recorded sLEC cells are of type (g, h) when using linear vision.

[7]Notice that rats have a wide angle vision: $320 - 360°$ in the azimuthal plane depending on head angle [152].

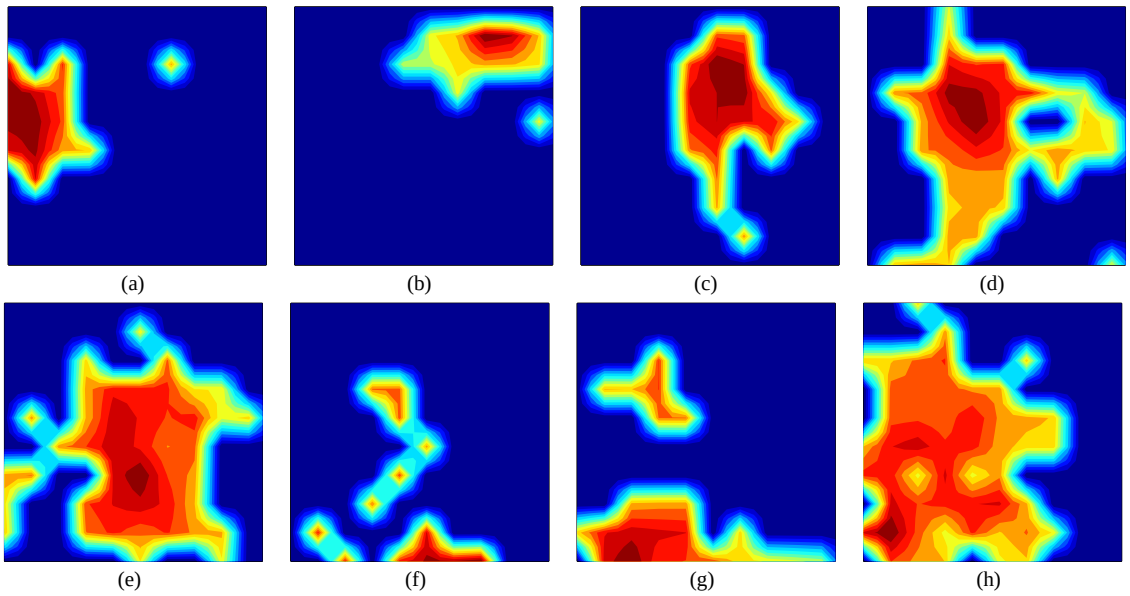

Figure 8.9: Some place fields recorded from the sLEC layer of the model in the case of PCA-based visual processing.

Fig. 8.9 presents eight sLEC place fields recorded when adopting PCA-based visual interpretation. Again, most of the recorded place cells exhibit single-peak localized fields (e.g., Figs. 8.9 (a, b)). However, due to visual aliasing, some cells have multiple subfields, that is, they do not discriminate spatial locations effectively (e.g., Figs. 8.9 (g, h)). About 25% of recorded sLEC cells present unclean location selectivity due to multi-peak activity. Note that, with respect to sLEC place cells created in the linear-vision case, receptive fields of Fig. 8.9 are more expanded, that is, place cells tend to be maximally active over larger spatial regions.

Finally, Fig. 8.10 shows eight sLEC receptive fields obtained by applying Gabor-based visual feature extraction. More than 90% of recorded cells showed clean location-correlated firing (e.g., Fig. 8.10 (a)). Still, some sLEC cells encode ambiguous visual information which results in non-localized receptive fields (e.g., Fig. 8.10 (f)). A general property of our Gabor-based sLEC place cells is that they tend to fire at appreciable rate throughout the entire arena surface (Fig. 8.10 (h)), reducing the contrast between low and high activity regions. However, despite this noisy background activity, their maximal firing rate occurs only over restricted areas of the environment.

8.2.3 Interpreting Place Cell Ensemble Activity by Population Vector Decoding

The proposed model yields an allothetic spatial representation consisting of a large number of localized overlapping place fields. The rationale behind such a redundant approach is two-fold: *(i)* To cover space uniformly in order to generate a *continuous* coarse coding representation (similar to a dense family of overlapping basis functions); *(ii)* To use the place cell population activity, rather than the single cell activity, for the self-localization task.

To locate itself, our animat utilizes the information provided by all sLEC cells, i.e., $\mathcal{R}^{sLEC}(t) = \{\, r_i(t) \,|\, \forall i \in sLEC\}$. Looking at the population activity rather than at the single cell activity, allows the system to compensate for misleading sLEC place cell activity (e.g., Fig. 8.10 (g)). This enhances the stability and the robustness of the vision-based self-localization process.

Decoding the ensemble place cell activity for position reconstruction is neurophysiologically plau-

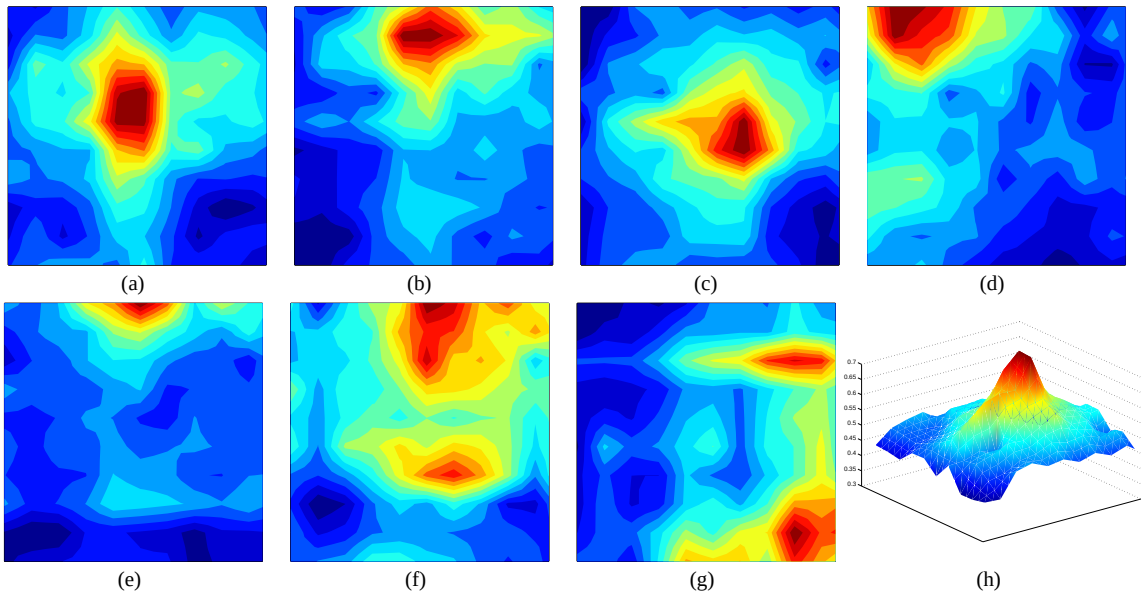

Figure 8.10: Samples of sLEC receptive fields obtained by means of Gabor-based visual feature extraction.

sible. Hippocampal place cells seem indeed to transmit an ensemble code to identify spatial locations [394, 406]. Population decoding has been also applied to interpret neuronal activity patterns in several other brain areas: The primate motor and premotor cortices [121, 120, 67, 299], the parietal cortex [164, 339], the cerebellum [105], the primary visual cortex [123], the superior colliculus [124, 179], and the inferotemporal cortex [404].

In order to interpret the ensemble firing $\mathcal{R}^{sLEC}(t)$ as a spatial location, we compute the population vector $\vec{p}(t)$ by averaging sLEC activity [121, 243, 1, 61]

$$\vec{p}(t) = \frac{\sum_i \vec{x}_i \, r_i(t)}{\sum_i r_i(t)} \tag{8.17}$$

where $\vec{x}_i$ is the center of the place field of neuron $i \in$ sLEC. Let $\vec{s}(t)$ be the true agent's location at time t. According to Eq. 8.17, the system estimates $\vec{s}(t)$ by computing the center of mass $\vec{p}(t)$ of the sLEC network activity. The encoded position $\vec{p}$ is near, but not necessarily identical, to the true location $\vec{s}$. The approximation $\vec{p} \approx \vec{s}$ is good for large neural populations covering the environment densely and uniformly [299].

Note that the place field center $\vec{x}_i$ has been made explicit for interpreting and monitoring purposes only. Associated with each place cell i is a vector $\vec{x}_i$ that represents the estimated location of the robot (based on dead-reckoning) when it creates the cell i. While the vector $\vec{x}_i$ is used in Eq. 8.17 for decoding the observed population activity, knowledge of $\vec{x}_i$ is not necessary for spatial learning and navigation.

Fig. 8.11 (a) shows an example of sLEC place cell distribution created by the robot after spatial learning (in the case of linear visual information). The square box represents the experimental arena. Each dot represents a place cell $i \in$ sLEC, and the position of the dot represents the center $\vec{x}_i$ of the place field. The two-dimensional space is uniformly and densely covered by the population of sLEC cells forming the allothetic space representation[8]. In this experiment the animat, starting from an empty

[8]Note that place cells are not topographically arranged within the sLEC layer of the model. That is, two cells i and j coding for two adjacent locations $\vec{x}_i$ and $\vec{x}_j$, respectively, are not neighboring neurons in the sLEC network. In Fig. 8.11 (a) adjacency of physical locations is preserved at the level of sLEC cells for interpreting purposes only.

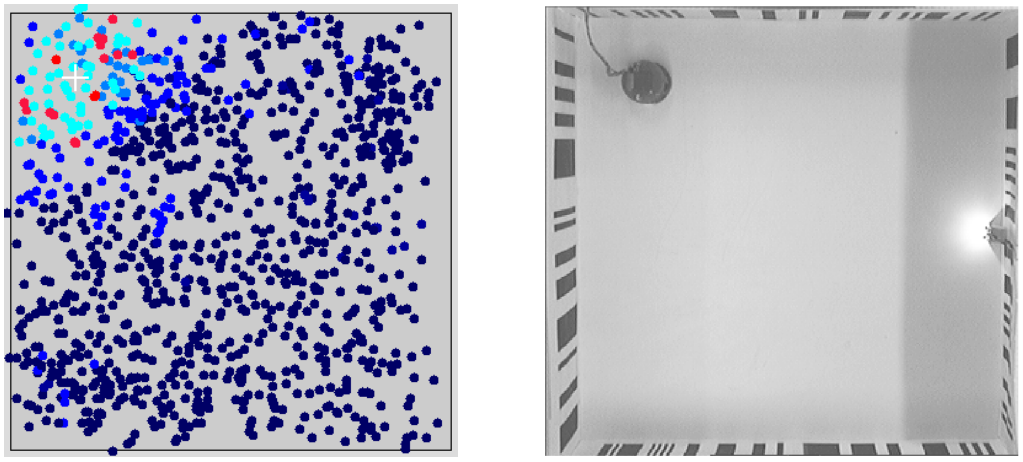

(a) (b)

Figure 8.11: Allothetic vision-based space representation. (a) An example of sLEC population created by the robot after learning. The square box represents the experimental arena. Each dot represents a place cell $i \in$ sLEC, and the position of the dot represents the center of the place field $\vec{x}_i$. The white cross represents the center of mass of the ensemble activity computed by population vector decoding. (b) The location of the robot corresponding to the sLEC ensemble activity of (a).

population, has created about 600 vision-driven sLEC cells. The ensemble sLEC cell activity of Fig. 8.11 (a) codes for the robot's location shown in Fig. 8.11 (b). The center of mass of the network activity (computed by Eq. 8.17) is represented by the white cross.

8.2.4 Accuracy of the Allothetic Spatial Representation

The sLEC activity blob moves over the two-dimensional space tracking robot's displacements. The accuracy of the representation is not uniformly distributed over the arena surface: The mapping from the visual input space to the two-dimensional view-manifold reflects the reliability of local visual stimulations, such that locations characterized by ambiguous local views will be poorly encoded by sLEC firing activity. Fig. 8.12 (a) shows an example of inadequate sLEC ensemble activity. The robot is approximately at location A. However, sLEC exhibits a multi-blob ensemble activity and population vector decoding (Eq. 8.17) reconstructs a position $\vec{p}$ that is totally inaccurate (white cross).

The sLEC allothetic space representation does not form the final hippocampal spatial map, rather it is combined along with path integration to yield a stable CA3-CA1 space representation (Chapter 10). Nevertheless, sLEC activity plays a primary role in determining the system's dynamics, and it is used to calibrate the idiothetic space representation occasionally. As a consequence, the agent needs a criterion to evaluate the reliability of the vision-based space coding.

As a first step, we have adopted a mere algorithmic technique to assess sLEC population activity. As a first approximation, one way to penalize representations of type Fig. 8.12 (a), consists of measuring the dispersion σ of the sLEC population activity around the center of mass $\vec{p}$. According to this technique, the robot may assess sLEC ensemble activity over time by employing a threshold Σ to simply neglect improper sLEC representations, that is, those activity patterns characterized by a dispersion $\sigma > \Sigma$ (Fig. 8.12 (b)). In a future step, we will implement this technique to assess the ensemble sLEC activity by means of a neural network being included in the model.

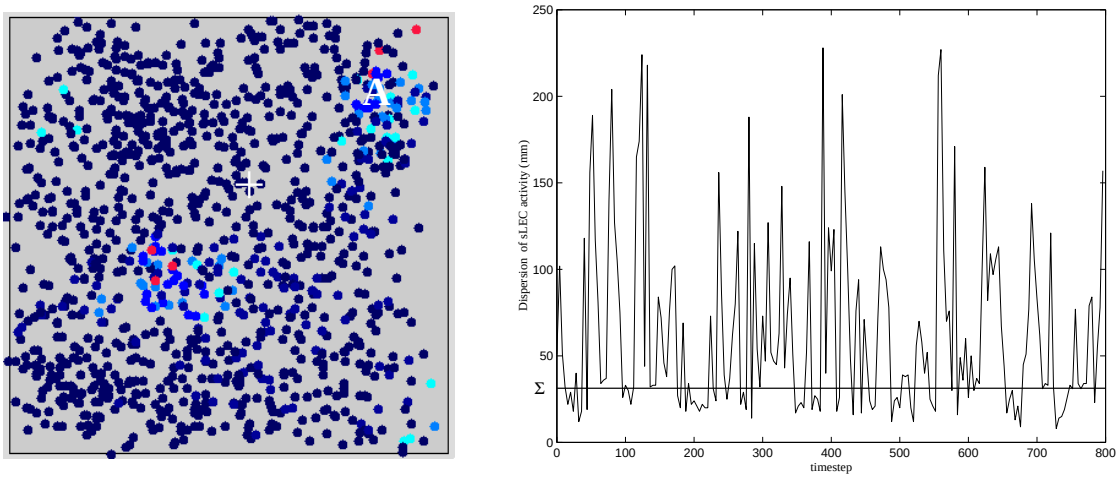

(a) (b)

Figure 8.12: (a) An example of multi-blob population activity due to ambiguous visual stimuli. Location A is the actual animat's position, whereas the white cross represents the position reconstructed from sLEC vision-based activity. (b) The reliability of the positional information encoded by sLEC activity may be assessed by looking at the dispersion σ (around the center of mass) over time, and a threshold Σ may be used to select proper sLEC representations.

8.3 Discussion

In this chapter, we have described the visual pathway of our hippocampal model. Starting from raw visual information, we build an allothetic space representation based on place-sensitive neurons. In the first part of the chapter (Sec. 8.1), we have focused on the interpretation of visual data. We put forth a mechanism to detect the view manifold that provides a low-dimensional representation of the environmental spatio-temporal properties carried by the visual signal. Incoming stimuli are interpreted by means of neurons that only respond to combinations of specific visual cues. This yields a neural activity in which metric information (e.g., distance to visual cues), landmark recognition, and interrelations between landmarks are *implicitly* encoded. This is in contrast to other hippocampal models in which metric information is directly used as input for the system [317, 62, 377, 286, 369].

In order to validate our approach, we have tested it on a robotic platform by means of three distinct implementations: *(i)* A system based on Walsh-like filters for linear image processing, *(ii)* a PCA-based module to interpret two-dimensional video signals, and *(iii)* a retinotopic Gabor-based approach suitable for modeling biological low-level visual processing.

In the second part of this chapter (Sec. 8.2), we have proposed an unsupervised learning scheme to generate location-sensitive neurons (place cells) relying on processed visual information. As a prediction, we assume that these cells belong to the superficial layer of the lateral entorhinal cortex (sLEC). Indeed: *(i)* They receive processed allothetic information, as neurons in the superficial lateral entorhinal cortex do from the parietal lobe via the perirhinal cortex [278, 60]; *(ii)* They exhibit spatial correlates similar to cells observed in the rat entorhinal cortex. Actually, available recording data from the entorhinal cortex concern the superficial layer of its medial part (sMEC) [278]. The only available finding concerning LEC comes from lesion studies: Animals with lesioned LEC exhibit deficits in an odor-to-place matching task [262]. This result is consistent with the hypothesis that LEC (in particular its superficial layer) is involved in encoding allothetic signals (e.g., odor). However, it does not tell us anything about the presence of place cells in sLEC. To our knowledge, not any recordings have been done from LEC cells [283]. Thus, the model predicts the presence of spatially tuned neurons in the sLEC that are driven by multimodal

external cues, and in particular by vision.

According to experimental findings about place cells recorded from freely-moving rodents in open-field arenas [239], our sLEC cells have non directional place fields. In order to obtain direction independent sLEC place fields, we adopt a technique similar to the approach proposed by Burgess *et al.* [62], in which, for every visited location, four distinct views (e.g., north, east, south, and west) are bound together to form a non-directional local view. In our model, for each view taken, the animat creates a snapshot cell coding for that view and whose activity is viewpoint dependent. Unsupervised correlational learning is then applied to combine the response of several snapshot cells for driving a single sLEC cell. In this process, the compass information provided by the head-direction module (Chapter 4) plays an important role. As a result of learning, sLEC cells are spatially tuned with respect to an allocentric frame of reference. Interestingly, experimental data suggest that viewpoint-dependent spatial coding occurs in the parietal lobe, whereas neurons in the hippocampal formation tend to provide allocentric representations [60, 26]. Thus, we assume that *(i)* our hypothetical snapshot cells might find their biological counterpart in neurons in the parietal cortex, *(ii)* unsupervised correlational learning might be responsible for the translation from gaze-dependent snapshot cell activity into allocentric spatial activity of sLEC cells, and *(iii)* head-direction cells should be important for this translation process.

A direct consequence of adopting unsupervised correlational learning is that sLEC place fields are established quickly (at the first visit of the place) and subsequently tuned to achieve stable coding. This is consistent with recordings from real hippocampal place cells [147, 395]. After learning, the model develops a space representation consisting of a large number of overlapping sLEC place fields covering the environment uniformly. This is compatible with entorhinal cells recorded from rats, whose field centers evenly cover the surface of the recording apparatus [278].

In the model, redundancy in the place field representation is considered a crucial issue to yield robustness. To interpret the ensemble sLEC cell activity as spatial locations we apply a population vector scheme [121, 394, 62], a decoding mechanism that reduces the effect of noise in the response of single neurons. The biological plausibility of reconstruction methods such as population coding has been recently demonstrated by Zhang *et al.* [406].

A critical aspect of our vision-based localization system is the need for rather rich visual stimulation. Future research will focus on the adaptability of the feature extraction process. For instance, in the case of a high-walled arena without visual cues on the walls, the animat might resort to a more geometric approach in which agent-wall distance would be explicitly considered [59] as well as the geometric properties of the environment (e.g., metric relations between walls) [276].

To conclude, an intrinsic limitation that occurs when performing vision-based spatial learning is the sensory aliasing problem [193, 109]. In other words, when relying on visual information only, the spatial representation might not fulfill the Markov hypothesis [202]. Indeed, since distinct spatial locations may provide identical visual information, the low-dimensional view manifold may be singular [178, 345, 308]. The presence of multi-peak place fields in our sLEC population reflects this singularity property of the view manifold. As already mentioned, in our hippocampal model we consider, along with the visual pathway, an idiothetic representation encoding self-motion information (i.e., path integration). This idiothetic signal provides an internally generated context that enables the system to disambiguate visual input. Path integration is the topic of the next chapter.

Chapter 9

Idiothetic Space Representation: Path Integration

Neurons in the hippocampal formation exhibit stable place fields even in the absence of allothetic cues, i.e., either after visual cue removal, or in the dark [251, 258, 212, 277, 278, 194] (Chapter 6). This suggests that, along with allothetic signals, the dynamics of place cells relies on some internally-generated motion information.

The ability to navigate relying on self-motion information is referred to as path integration [223, 99]. More precisely, path integration (also termed dead-reckoning[1] [114, 272], or vector navigation [380]) enables a navigating agent to infer its current position relative to a departure point from its own movement [201, 98]. This implies integrating the translational and angular motion components over time in order to determine the current distance and orientation with respect to the starting point [314]. To assess the angular component of motion continuously, the information provided by head-direction cells (Part II) is determinant.

Path integration does not consists of *(i)* memorizing the sequence of movements since the beginning of the outward journey until the current position, and *(ii)* performing path reversal to return home. Rather, the continuous integration of translations and rotations over time generates a homing vector $\vec{h}$ leading the animal *directly* to the departure location [98] (Fig. 9.1 (a)).

The capability of homing by path integration has been observed in several living species through a large body of behavioral studies. Experiments suggest that arthropods [225], ants [382, 77, 238], spiders [131, 316], bees [35], birds [223], rodents [361, 224, 223, 273, 100, 314, 3], dogs [21, 30, 313], cats [30], and humans [30, 304, 111, 102], path integrate when navigating without external references. Fig. 9.1 (b) shows an example of hamsters' homing trajectory reported by Etienne and colleagues [98]. In particular, they have shown that hamsters do resort to vector addition in absolute coordinates to perform goal-oriented navigation independently of external landmarks [99].

In mammals, assessing translational and angular displacements relies on vestibular, somatosensory, and motor command efference copy information. Vestibular signals primarily influence rotation estimation, whereas somatosensory and efference copies are predominant to measure translations [125].

Dead-reckoning allows an agent to navigate in totally unfamiliar environments since its very first

[1] Deduced reckoning.

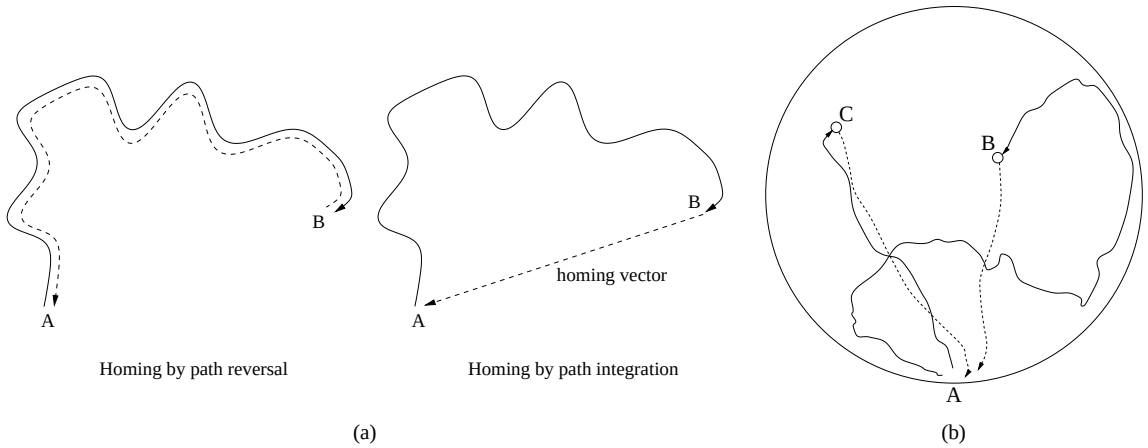

Figure 9.1: (a) Difference between path reversal (i.e., inverting the sequence of movements performed from a starting point A to a current location B, and path integration (i.e., returning to location A by a direct homing vector). Solid lines represent the outward journey, whereas dashed lines indicate the return path. (b) Homing behavior of two hamsters. After being guided by a bait from the nest location A to points B e C, respectively, within the arena (solid lines), animals return home by an approximately direct path (dashed lines). The experiment was performed in the dark, with a circular arena of radius approximately equal to 1 m. Data courtesy A. Etienne [98].

exploring excursion [136]. To path integrate, indeed, animals do not have to learn environmental properties, because they solely rely on their internal references. This makes path integration a sort of basic underlying mechanism for navigation. Indeed, it is functionally available *(i)* in all type of environments (i.e., with or without reliable landmarks), *(ii)* as soon as the animal enters an unknown environment, *(iii)* also for "naive" navigators (e.g., young animals), that is, agents that can not exploit their interaction with the external world effectively. In addition, since dead-reckoning endows animals with vector addition capabilities in allocentric space [99], it might play an important role in determining short-cut finding behavior. Preliminary behavioral results on hamsters support this hypothesis [97].

A limitation of path integration is its vulnerability to cumulative drift over time. This holds for both biological [98] and artificial navigating systems [48]. Indeed, the idiothetic-based dynamics that consists of integrating translational and rotational signals instantaneously, induces systematic as well as non-systematic errors that quickly disrupt the position representation [238, 314, 201, 136]. As a consequence, path integration needs to be reset episodically in order to maintain a consistent representation over time. When available, stable allothetic cues (e.g., visual fixes) may be used to accomplish such a calibration process [210, 207, 98].

In this chapter, we focus on the idiothetic pathway of our hippocampal model. Internal movement-related signals are used to drive a population of place cells, PiC, that we suppose to be located in the superficial layer of the medial entorhinal cortex (sMEC). The model also involves the subicular area (SC) of the hippocampal formation. For an overview of previously postulated hypotheses modeling path integration, see Sec. 7.2.

9.1 Path Integration-based Place Coding

The path integrator providing the idiothetic pathway of our hippocampal model must satisfy the following requirements: *(i)* To encode the current agent's position $\vec{p}(t)$ with respect to a departure point $\vec{p}(t_0)$ (where t_0 is the time at which the journey starts), i.e., to provide the necessary information to compute the current homing vector $\vec{h}(t)$. *(ii)* To update the representation $\vec{p}(t)$ of the agent's position by integrating

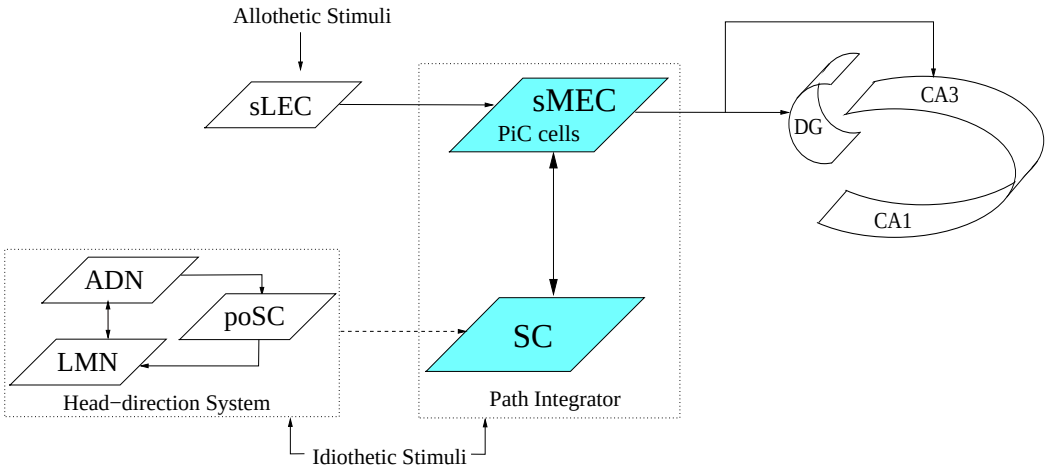

Figure 9.2: A functional overview of the idiothetic pathway (grey areas) of the model. Path integration occurs via the interaction between cells in the superficial medial entorhinal cortex (sMEC) and cells in the subiculum (SC). SC in the model receives directional information from the head-direction system (Part II). sMEC forms the output of the path integrator and projects onto the hippocampus proper. sMEC receives afferent information from sLEC which, in our model, is the locus of the allothetic space representation. The latter is used to occasionally recalibrate the sMEC idiothetic map.

inertial signals over time. Thus, the path integrator has to receive directional information from the head-direction system, and self-motion information from the proprioceptive and vestibular systems. *(iii)* To convey its idiothetic representation to the hippocampus proper, in which the CA3-CA1 place field representation takes place. *(iv)* To receive afferent information from an allothetic-based place coding area in order to recalibrate the idiothetic representation.

Fig. 9.2 shows the functional architecture of our path integrator as well as its main afferents and efferents. The system satisfies the four above requirements:

(i) The idiothetic place representation occurs in a population of neurons that we call path-integration place cells (PiC). The ensemble activity $\mathcal{R}^{PiC}(t) = \{r_i(t) \,|\, i \in PiC\}$ codes for the robot's current position $\vec{p}(t) = (x(t), y(t))$ within an allocentric frame of reference centered at the starting location $\vec{p}(t_0)$. $\mathcal{R}^{PiC}(t)$ depends on internal information only. The activity of the PiC cells is environment-independent, that is, PiC place fields do not change from environment to environment. We assume that the anatomical locus for our PiC cells is the superficial layer of the medial entorhinal cortex (sMEC). Indeed, neurophysiological data suggest that the place-field topology of location-selective neurons in sMEC does not change across different environments [278].

(ii) Similar to previous anatomical hypothesis [207, 319, 285], we assume that subicular cells (SC) have a primary role in path integration. Indeed, SC cells have the property of encoding both location and direction of the animal [323], and have a firing activity that is modulated by self-motion information [207]. Moreover, similar to sMEC, SC has the capability of transferring a single place field pattern from one environment to the next. As shown in Fig. 9.2, we assume that integrating self-movement signals occurs through the interaction between sMEC (PiC) and SC cells. In particular, we suppose that SC cells are responsible for shifting the PiC activity blob over time. Anatomical interconnections exist between the subiculum and layers I-III of the medial entorhinal cortex [397].

(iii) In the model, PiC cells form the output of the path integrator. Since the PiC representation takes place in sMEC and the latter projects anatomically onto the hippocampus proper (via the per-

forant path), the third requirement is satisfied. In Chapter 10 we will describe how sMEC→CA3 connections can be learned during environment exploration.

(iv) In order to prevent the path integration error from diverging over time, we employ the vision-based space code established in the sLEC layer (Chapter 8) to calibrate the PiC representation. Anatomical data show that the lateral entorhinal cortex (LEC) projects to the medial entorhinal area (MEC) [278]. We will describe in Chapter 10 how path integration calibration is achieved in the model.

9.1.1 Robotic Implementation

In the current system, we only employ a neural network to implement sMEC (PiC) cells. On the other hand, we do not propose a neural implementation for the SC layer shown in Fig. 9.2. Rather, to update PiC activity, translational velocity $v(t)$ provided by the odometer is algorithmically combined with the allocentric directional signal $\theta(t)$ provided by our head-direction module. That is, similar to Mittelstaedt and Mittelstaedt (1973) [222], we compute the agent's current location $\vec{p}(t) = (x(t), y(t))$ as

$$\begin{aligned} x(t) &= \int_{t_0}^{t} v(t') \cos(\theta(t'))\ dt' \\ y(t) &= \int_{t_0}^{t} v(t') \sin(\theta(t'))\ dt' \end{aligned} \tag{9.1}$$

where t_0 is the time at which the agent enters the arena. Thus, the position $\vec{p}(t)$ is computed within an allocentric framework whose reference point is the entry position $\vec{p}(t_0)$, the reference orientation ϕ is provided by the head-direction system, and the metric is provided by the odometer.

We take a population PiC $= \{\, i \mid 1 \leq i \leq N \,\}$ of $N = 1000$ cells. Our PiC cell assembly could be interpreted either as *one* of the charts of the multi-chart path integrator by McNaughton *et al.* [207], or as the universal place map proposed by Sharp [319]. The advantage of such a single-chart approach is that only one map has to be preconfigured in advance. By contrast, McNaughton *et al.* assume a large number of preconfigured charts (and for each of them the necessary prewired connections for the update mechanism) from which one chart is extracted to be used as current map.

PiC neural activity relies on preconfigured metric interrelations: *(i)* One cell $i_o \in$ PiC encodes the origin $(0, 0)$ of the abstract frame of reference $\mathcal{S}$ the PiC firing relies on. Every time the robot enters a novel environment, cell i_o is recruited to encode the entry position $\vec{p}(t_0)$. *(ii)* Each cell $i \in$ PiC, with $i \neq i_o$, has one preferred firing location $\vec{p}_i$ relative to the origin identified by i_o. PiC preferred positions are evenly distributed over the two-dimensional abstract space $\mathcal{S}$. In the current implementation, we adopt a fixed spatial resolution of about $10\,cm$.

Note that, since the abstract space $\mathcal{S}$ is mapped onto the physical space $\mathcal{S}'$ depending on the entry position $\vec{p}(t_0) \in \mathcal{S}'$, a *novel* environment may be encoded by two distinct PiC firing patterns if two explorations start at points $\vec{p}^{\,1}(t_0) \neq \vec{p}^{\,2}(t_0)$. In other words, PiC cells have preconfigured metric relations within the abstract allocentric space $\mathcal{S}$, but not with respect to a physical absolute framework $\mathcal{S}'$. As we will see in Chapter 10, correlational learning can be employed to establish a stable $\mathcal{S} \rightarrow \mathcal{S}'$ mapping, such that in a *familiar* environment PiC cells maintain similar firing patterns across different sessions.

As the robot moves, PiC cell activity changes according to translational self-motion signals (supplied by the odometer) and to the current heading of the robot as estimated by the directional system. The firing rate $r_i(t)$ of a PiC cell i at time t is taken as a Gaussian

$$r_i(t) = \exp\Big(- \big(\vec{p}_{dr}(t) - \vec{p}_i\big)^2 / 2\sigma^2 \Big) \tag{9.2}$$

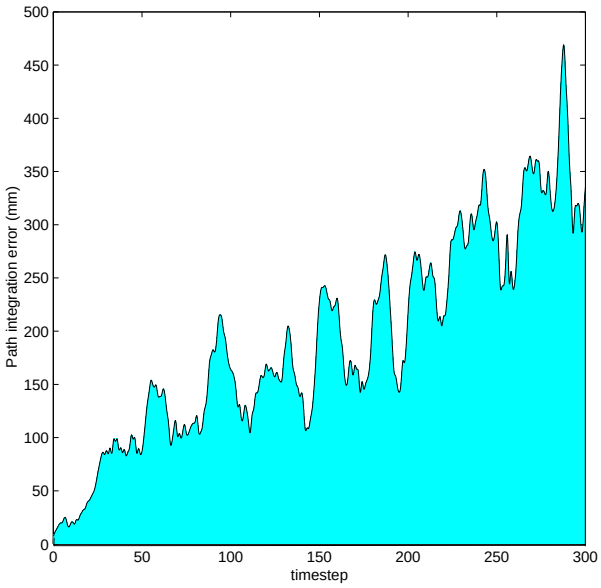

Figure 9.3: The mean path integration error over time. Idiothetic place coding is affected by cumulative shift.

where $\vec{p}_{dr}(t)$ is the current robot position (relative to the starting point $\vec{p}(t_0)$) estimated by dead-reckoning, $\vec{p}_i$ is the preferred firing location (center of the receptive field) of cell i, and $\sigma \approx 10\,cm$ defines the width of the field. According to Eq. 9.2, PiC cells have quite large firing fields and tend to be rarely silent within a $80 \times 80\,cm$ square arena (Fig. 1.2 (b)). This is consistent with the fact that neurons recorded from sMEC fire considerably over the entire environment [278].

9.1.1.1 Position Tracking Error Over Time

In order to test the idiothetic place representation provided by the above system, we run a series of n experiments all starting at $t = 0$. At each time, the robot rotates by a random angle within $[0°, 360°]$, and moves linearly by a random step within $[0, 5]\,cm$. Population vector decoding [121, 243, 241, 394, 299] is applied to estimate the robot's position $\vec{p}\,'(t)$ according to the ensemble PiC activity, that is, $\vec{p}\,'(t) = \big(\sum_i r_i(t)\,\vec{p}_i\big)/\sum_i r_i(t)\big)$. Vision-information is not made available to the robot (to simulate experiments with animals in complete darkness).

At each step t, we measure the difference $e(t)$ between the actual robot's position $\vec{p}(t) = (x(t), y(t))$ and the estimate $\vec{p}\,'(t) = (x'(t), y'(t))$. The mean tracking error at time t is defined as

$$e(t) = \frac{1}{n}\sum_{i=1}^{n}\big((x' - x)^2 + (y' - y)^2\big)^{\frac{1}{2}} \tag{9.3}$$

The experiment includes $n = 5$ trials, each of which consists of 300 steps. At the beginning of each session the robot is placed inside the arena at the same starting position $\vec{p}(t_0)$, and with the same initial arbitrary heading. Fig. 9.3 shows the mean path integration error $e(t)$. As expected, the idiothetic representation is affected by a cumulative shift over time. As a consequence, in order to make path integration useful for effective space coding, we need a mechanism to bound dead-reckoning error by occasionally calibrating the idiothetic representation. Allothetic information can be used to achieve this aim. In Chapter 10 we will take the visually-driven activity of sLEC place cells as the signal for calibration.

9.1.2 Having a Finite Idiothetic Representation: The Edge-Effect Problem

Our path integrator model (as well as those by McNaughton *et al.* [207], Sharp [319], and Redish and Touretzky [285]) has an intrinsic limitation due to the following factors: *(i)* The idiothetic place map is preconfigured rather than learned; *(ii)* The representation is finite, i.e., the finite PiC ensemble can only encode a confined physical space. This gives rise to the edge-effect problem [283]: What happens when the navigating agent (either a robot or a rat) crosses one of the edges of the place map?

Neurophysiological data by Sharp [320] suggest that subicular cells have the ability to transfer a single place code across environments with different shapes and sizes. In particular, subicular place cells seem to be able to adjust the size of their place field in order to fit the boundaries of the current environment. Based on these findings, we might imagine a scalable idiothetic representation such that both the size of place fields and their mutual metric relations would change to fit larger environments. Although this hypothesis might contribute to solving the edge-effect problem, it does not provide the full solution. Indeed, there would still be a limit to the scalability of the map, and the accuracy of the place representation would decrease as place fields would become too large and distant [283].

An alternative solution consists of resorting to multiple place cell representations (i.e., multiple hippocampal maps [256], active place cell subsets [240], multi-charts [207], multiple reference frames [362, 285]). The basic idea is that the animal might be able to do map transitions appropriately, such that it would always have a valid place representation associated to the current context. Although we do not address the issue of multiple hippocampal maps [256, 240, 207, 362], we assume that the robot would switch to a different map as soon as the current PiC representation would be no longer adequate.

9.2 Discussion

In this chapter we have described the idiothetic pathway of our hippocampal model. The aim is to provide the animat with a place code system relying on inertial stimulation only. This is consistent with neurophysiological data showing that place fields do persist even in the absence of external cues (Chapter 6).

We consider the superficial layer of the medial entorhinal cortex (sMEC) and the subiculum (SC) as primary structures for path integration (Fig. 9.2). This is based on the experimental evidence that sMEC and SC seem to be capable to transfer the same place field topology across different environments, e.g., from a square box to a cylindrical arena (Chapter 6). In the model, sMEC and SC interact with each other to integrate self-movement signals. A fundamental contribution to build the idiothetic space representation comes from the head-direction cells (Chapter 4). We assume that SC cells receive (either directly or indirectly) the signal encoding the current allocentric heading of the animal (e.g., from the postsubiculum (poSC)).

Similar to McNaughton *et al.* [207], we take the subiculum (SC) as a primary component to update path integration. On the other hand, we do not take the hippocampus proper as a part of our path integrator. Rather, the idiothetic representation is maintained outside CA3-CA1 (in particular in sMEC) and converges onto the hippocampus via the perforant path. Also, we adopt a single preconfigured chart, in contrast to the multi-chart hypothesis by McNaughton and colleagues. In that, our approach is similar to the model by Redish and Touretzky [285]. However, their system identifies the parasubiculum (paSC) as a crucial part of the path integrator. Our model does not involve paSC. Rather, paSC might be responsible for feeding SC with the directional signal from the postsubiculum (poSC). Similar to Sharp [319], we use a single chart and we take the subiculum and the entorhinal cortex as primary components for the path integrator. However, Sharp assumes that EC receives information about the size of the current environment and that this knowledge is used to center the universal map within the

physical space as well as to adapt its size appropriately. Our animat does not estimate the size of a novel environment before starting exploration. Also, our idiothetic representation is not universally mapped onto the physical space (i.e., it depends on the starting point). Another important difference concerns Sharp's hypothesis about reinitializing/correcting the path integrator. Sharp suggests that the hippocampus proper sends the calibrating signal to SC. By contrast, in our model, the allothetic sLEC map is responsible for calibrating the idiothetic sMEC representation, via the sLEC→sMEC projections (Chapter 10). Since the hippocampus proper combines both idiothetic and allothetic representations, using the CA3-CA1 signal to reset path integration would not be effective: an erroneous path integration would affect CA3-CA1 firing, which should not be used for calibration. On the other hand, the map formed in our sLEC layer is both allothetic and totally independent from path integration.

Some predictions can be made based on the path integrator model presented in this chapter: *(i)* Place cells in sMEC mainly depend on idiothetic information; *(ii)* The sMEC-SC loop provides the anatomical locus for path integration; *(iii)* Lesions to the interconnections between sMEC and SC should impair the animal's ability to path integrate; *(iv)* Lesions to SC should impair sMEC firing as well as disrupting path integration; *(v)* Damaging sMEC should reduce the influence of inertial stimuli upon CA3-CA1 cells (e.g., animals with lesioned sMEC should not exhibit clean hippocampal place fields in complete darkness); *(vi)* In contrast to the theory by McNaughton *et al.* [207] and similar to Sharp [319, 320] and Redish and Touretzky [285], our model postulates that lesions to the hippocampus proper should not disrupt path integration. Alyan and colleagues (1997) [4] have explicitly tested this hypothesis. Preliminary results suggest that, when solving homing tasks, hippocampal rats exhibit intact path integration.

There are some aspects we need to take into account in the future in order to extend our path integrator: *(i)* We need to implement the entire system by a neural machinery. At the moment, only the sMEC layer (i.e., PiC) is a cell assembly. Similar to our neural approach to integrate angular velocity (Chapter 4), we will consider a SC network in which for each preferred location $\vec{p}_i$ there are several neurons coding for different robot's headings θ_i. That is, each unit $i \in$ SC will respond to the combined signal $(\vec{p}_i, \theta_i)$. Moreover, according to experimental results, the activity r_i of cell i should be modulated by the translational velocity v. Finally, a sMEC↔SC connection scheme based on matching and offset projections might produce the activity shift necessary for tracking robot locomotion. *(ii)* As previously stated, in the current model we do not address the problem of switching between distinct idiothetic representations to cope with the edge-effect problem. Indeed, the neural PiC population is large enough to fit both our experimental arenas (i.e., a $60 \times 60\,cm$ square box, and a $80 \times 80\,cm$ open environment). In the future, a mechanism to endow the animat with multiple mapping capabilities should be included. *(iii)* Our path integrator accounts for non-systematic errors, whereas it does not model systematic biases (Eq. 9.1 is the correct trigonometric solution to generate the homing vector $\vec{h}(t)$) [314].

To conclude, we have stressed the fact that the idiothetic representation alone cannot provide stable self-localization capabilities (Fig. 9.3). This holds both for animals (e.g., [98]) and mobile artifacts (e.g., [48, 47]). Employing our vision-based place code to correct dead-reckoning is one of the topics of the next chapter.

Chapter 10

Place Cells in the Hippocampus Proper

In the previous two chapters, the allothetic and the idiothetic pathways of the model have been presented. The former allows the agent to extract spatio-temporal properties of the environment from visual signals. This results in place fields that we suppose being anatomically located in the superficial layers of the lateral entorhinal cortex (sLEC). An intrinsic limitation of vision-based spatial coding is its vulnerability to sensory aliasing (Sec. 8.2) [178, 345, 193, 109]. The idiothetic pathway (i.e., path integration) provides the agent with a basic mechanism to behave independently from environmental conditions (e.g., in complete darkness). We take the superficial layers of the medial entorhinal cortex (sMEC) as output for our path integrator (which also involves the subiculum). Self-localization based on idiothetic signals suffers from an instability problem: the representation is not coherent over time [98, 48] (Sec. 9.1).

In this chapter, allothetic and idiothetic representations are combined to form a population of CA3-CA1 place fields. Correlational learning is applied to combine these two types of information based on the agent-environment interaction. This induces a mutual benefit in the sense that path integration may disambiguate visual singularities and, conversely, visual cues may be used for resetting the path integrator. This process is done on-line during the development of the hippocampal space representation (i.e., exploration).

Fig. 10.1 (a) illustrates a functional view of the entire model. Hippocampal place cells and head-direction cells are strongly coupled and interact with each other to form a unitary spatial learning system. Consistently with experimental data [166], inhibiting either one of these two modules would critically impair the performance of the other and would disrupt spatial learning. Both systems rely on extrinsic as well as intrinsic signals to maintain stable internal representations. Finally, the combined output of both systems is used for motor commands to achieve goal-oriented navigation (Part IV). Fig. 10.1 (b) shows a more detailed view of the model as well as its anatomical counterparts.

In Sec. 10.1 we establish CA3-CA1 place fields by combining sLEC and sMEC place codes. In Sec. 10.2 we show how path integration can be recalibrated based on vision during exploration. Finally, in Sec. 10.3 we study the interrelation between allothetic and idiothetic cues in controlling place cells.

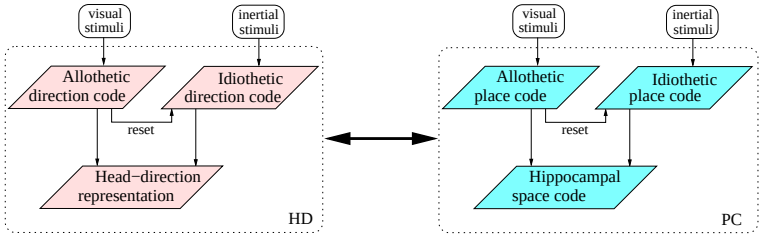

(a)

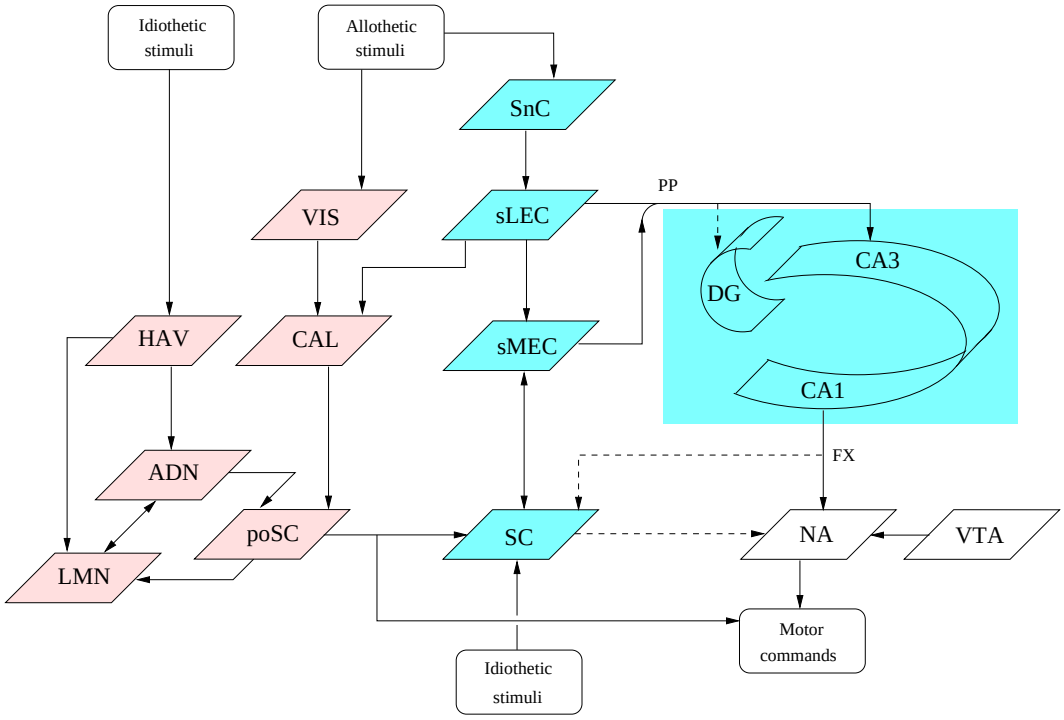

(b)

Figure 10.1: (a) Functional view of the entire spatial learning system. Place coding areas (hippocampal place cells) and direction coding areas (head-direction cells) interact with each other to maintain stable internal representations. (b) Representation of the constituents of the model with their anatomical counterparts. Glossary: LMN: lateral mammillary nuclei, ADN: anterodorsal thalamic nucleus, poSC: postsubiculum, HAV: hypothetical head angular velocity cells, CAL: hypothetical direction calibration cells, VIS: hypothetical visual cells, SC: subiculum, sMEC: superficial medial entorhinal cortex, sLEC: superficial lateral entorhinal cortex, SnC: hypothetical snapshot cells, DG: dentate gyrus, NA: nucleus accumbens, VTA: ventral tegmental area, FX: fornix, PP: perforant path.

10.1 CA3-CA1 Place Fields: Combining Allothetic and Idiothetic Representations

In the model, allothetic and idiothetic representations converge onto the hippocampus proper to form a spatial representation based on CA3-CA1 place fields. Both sLEC and sMEC cells project to CA3-CA1 units by means of downstream synapses that are established according to a correlational learning rule. Anatomically, these projections would form the perforant path from the superficial layers of the entorhinal cortex to the hippocampus (Chapter 6).

As mentioned in Sec. 8.2, when the robot enters a novel environment it does not have any prior knowledge. Place cells are recruited incrementally as exploration proceeds. In order to cover the space uniformly by a population of place fields, the agent needs a mechanism to assess "familiarity" with respect to spatial locations. In the model, whenever the robot visits a place $\vec{p}(t)$, it first checks whether there are already CA3-CA1 cells i coding for that place. A new subset of place cells is recruited only if

$$\sum_i \mathcal{H}(r_i(t) - \epsilon) < A \tag{10.1}$$

where $r_i(t)$ is the firing activity of cell i at time t, and $\mathcal{H}$ is the Heaviside function. Therefore, the population of hippocampal place cells grows only if the number of CA3-CA1 units active at location $\vec{p}(t)$ does not exceed a threshold A [286, 117]. Eq. 10.1, which is a mere algorithmic implementation, enables the system to control the redundancy level in the space representation.

For each new location, the system recruits a new subset of n CA3-CA1 place cells i[1]. New connections are then formed from all simultaneously active cells in sLEC and sMEC to the new place cells. Then, during the agent-environment interaction, Hebbian learning is used to establish the weight of the efferents from sLEC and sMEC to the hippocampus proper. If i and j represent a postsynaptic CA3-CA1 cell and a presynaptic cell in sLEC or sMEC, respectively, the weight w_{ij} is changed according to

$$\Delta w_{ij} = \eta\, r_i\, r_j\, (1 - w_{ij}) \tag{10.2}$$

where the learning rate η is taken equal to 1. The firing rate r_i of a cell $i \in$ CA3-CA1 is simply a weighted average of the activity of its presynaptic cells

$$r_i = \sum_j w_{ij}\, r_j \Big/ \sum_j w_{ij} \tag{10.3}$$

where j varies over sLEC and sMEC cells. The relation $0 \leq r_i \leq 1$ holds.

The above technique to generate CA3-CA1 place fields yields a stable representation in the hippocampus proper. The reason for this stability is two-fold: *(i)* Each CA3-CA1 place cell receives afferents from several entorhinal cells with similar firing patterns, and it averages their activities. This results in a cleaner positional signal with respect, for instance, to the noisy activity recorded in sLEC (Fig. 8.10) [278]. *(ii)* CA3-CA1 activity relies on the combination of internal and external stimuli. This allows the system to eliminate the sensory aliasing problem affecting pure allothetic representations. Also, in the dark, CA3-CA1 firing can be supported by the input provided by the sMEC path integration signal.

[1]The number n of newly recruited cells is a fixed parameter of the model that influences the redundancy of the place field space representation. In the current implementation, $n = 2$.

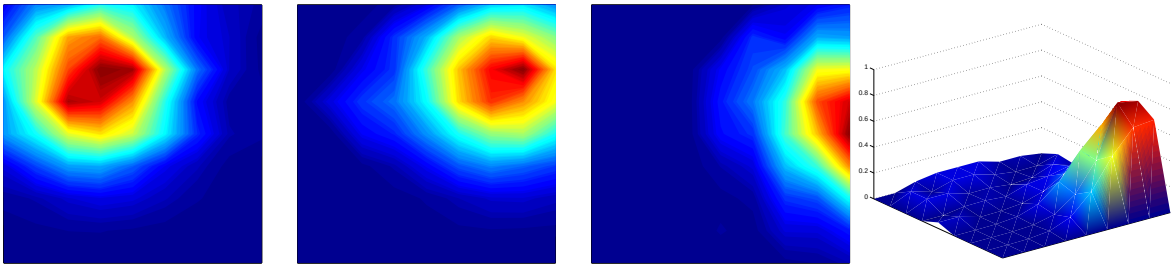

(a) (b) (c) (d)

Figure 10.2: Samples of CA3-CA1 place fields recorded during experiments using linear vision.

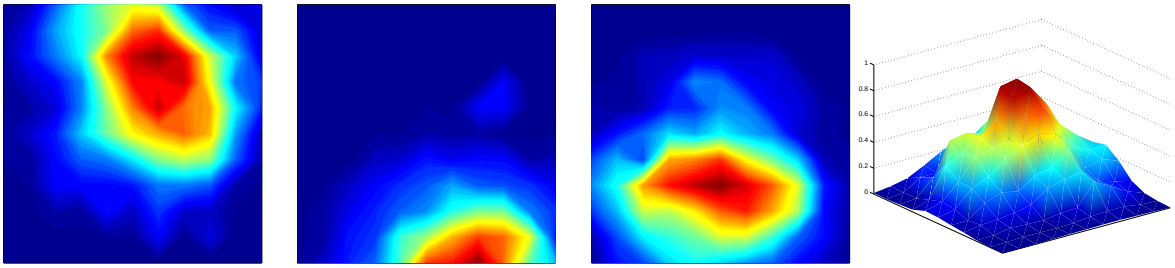

(a) (b) (c) (d)

Figure 10.3: Samples of CA3-CA1 place fields recorded during experiments using PCA-based vision interpretation.

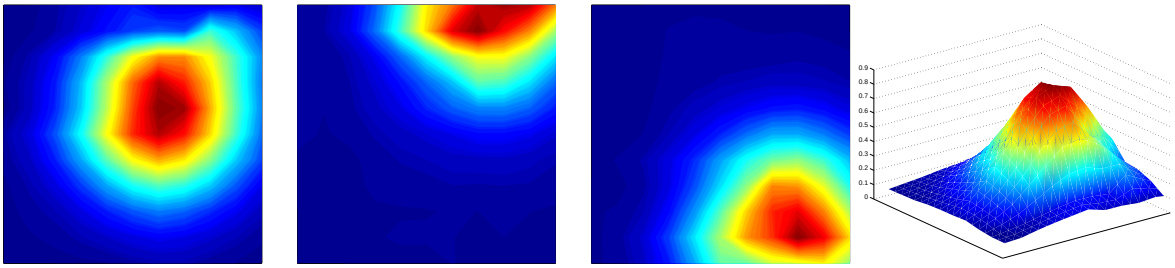

(a) (b) (c) (d)

Figure 10.4: Samples of CA3-CA1 place fields recorded during experiments using Gabor-based visual feature extraction.

10.1.1 Single-Cell Recordings

Figs. 10.2, 10.3, and 10.4 illustrate some typical CA3-CA1 place fields recorded in experiments with linear vision, PCA-based image interpretation, and Gabor-based feature extraction, respectively. The model generates biologically plausible CA3-CA1 place cells in all three cases. Receptive fields are clean and less noisy than those recorded in the sLEC layer of our model (Chapter 8). Most of the recorded CA3-CA1 cells do not exhibit multiple subfields (e.g., 97% in the linear vision case). This is mainly due to the contribution of the idiothetic representation. Such a clean spatial selectivity at the level of the hippocampus proper is compatible with experimental single-unit recordings showing that about 5% of hippocampal cells have multiple subfields within a single environment [281].

10.1.2 Multiple-Cell Recordings

As already mentioned in Chapter 8, the purpose is to cover the environment by a large population of localized overlapping place fields. This yields a dense family of overlapping basis functions that can be

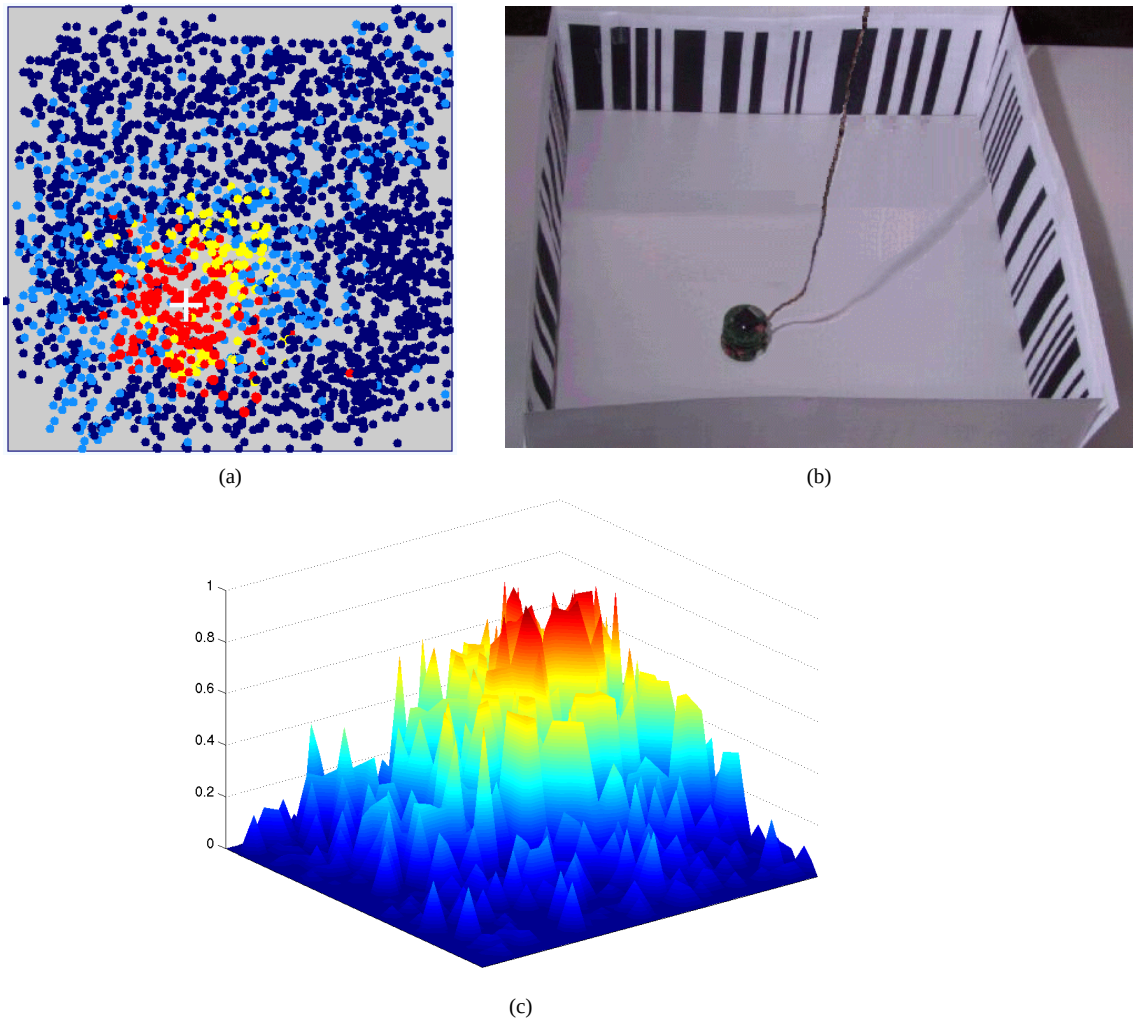

Figure 10.5: (a) Sample of learned CA3-CA1 population. Each dot denotes the center of a place field. The white cross represents the center of mass of the ensemble activity computed by population vector decoding. (b) The $60 \times 60\,cm$ experimental arena with the Khepera robot inside. The location of the robot corresponds to the CA3-CA1 ensemble activity of (a). (c) An example of population activity coding for the upper-right corner of the arena.

used for the self-localization task. Redundancy helps in terms of stability and robustness of the place code. Let $\mathcal{R}^{CA}(t) = \{\, r_i(t) \,|\, \forall i \in CA3{-}CA1\}$ be the ensemble CA3-CA1 activity at time t. As usual, we employ population vector decoding (Sec. 8.2.3) to reconstruct the agent's current position based on $\mathcal{R}^{CA}(t)$ [394, 406].

Fig. 10.5 (a) shows an example of CA3-CA1 population created by the robot after spatial learning based on linear visual information. Each dot is a place cell $i \in$ CA3-CA1, and its position denotes the center $\vec{x}_i$ of the place field. Note that, place cells do not need to be tied to the physical space to enable spatial navigation of the system (Part IV). Rather, we use this information in order to monitor the self-localization process. Also, CA3-CA1 cells are not topographically arranged (i.e., cells $i, j \in$ CA3-CA1 coding for two adjacent locations $\vec{x}_i$ and $\vec{x}_j$, respectively, are not adjacent neurons of the network). However, in Fig. 10.5 topology is preserved for interpreting purpose. In this experiment, the robot recruited approximately 1000 CA3-CA1 place cells (starting from an empty population) covering the environment uniformly and densely. The firing pattern of Fig. 10.5 (a) has been recorded with the

robot at the location shown in Fig. 10.5 (b). The estimated position (computed by population vector coding) is represented by the white cross (i.e., the center of mass of the CA3-CA1 ensemble activity). Finally, Fig. 10.5 (c) is a three-dimensional representation of CA3-CA1 population activity recorded with the robot being at the upper-right corner of the arena. The diagram has been obtained by discretizing the two-dimensional environmental space in 40×40 subregions, and then plotting the mean ensemble activity for each subregion.

10.2 Exploring Behavior and Path Integration Calibration

During locomotor behavior, the robot moves in discrete time steps Δt that determine the frequency at which it senses the world, interprets sensory inputs, and takes an action. We assume that each step Δt corresponds to one theta cycle of approximately $0.1\,s$ (the real movement of the robot is, of course, slower than this) [317, 59]. Thus, place cell activity is updated with a frequency of $10\,Hz$, which simulates the fact that hippocampal processing, during animal's motion, is timed by a sinusoidal EEG signal of $7 - 12\,Hz$, namely the theta rhythm (Sec. 6.4).

When entering a novel environment, the robot needs to explore it in order to establish a CA3-CA1 place field representation. Since the environment is unfamiliar, the robot starts by relying upon path integration only. The entry location becomes the reference point (home) relative to which the sMEC idiothetic representation is built. Then, exploration is performed to *(i)* extract the environmental visual properties in order to learn the sLEC allothetic code, *(ii)* combine local views encoded by sLEC activity with the spatial framework provided by the path integrator. This results in a growing CA3-CA1 population whose dynamics relies on two coherent allothetic and idiothetic space codes.

The robot adopts an exploration strategy that emulates the exploratory behavior of animals [180, 98, 78]. As exploration proceeds, new place cells are recruited by the system. At the very beginning, exploration consists of short return trips (e.g., narrow loops) centered in the home location and directed towards the principal radial directions (e.g., north, east, and so on). This overall behavior relies on the head-direction system and allows the robot to explore the space around the home base exhaustively. Afterwards, the robot switches to a more open-field exploration strategy: It starts moving in a random direction recruiting a new subset of CA3-CA1 place units for each new location encountered (i.e., spatial learning). After a while, the agent "feels" the need to re-calibrate its path integrator. Note that we do not propose a specific uncertainty model for the dead-reckoning system. We simply assume that the "need of calibration" grows monotonically as some function $n(t)$ of time t. When, after a time t_{cal}, $n(t)$ overcomes a fixed threshold n_{cal}, the animat stops creating place cells and starts following the homing vector to return to the starting location [98, 78]. As soon as it finds a previously visited location (not necessarily the home location), it tries to use the learned allothetic representation to re-align its path integrator. Fig. 10.6 (a) shows an example of open-field exploring excursion.

We take the vision-based sLEC activity as calibrating signal. Let $\vec{p}(t)$ be the center of mass of the sLEC ensemble activity computed by population vector coding at time t (Sec. 8.2.3). Let $\sigma(t)$ denote the variance of the sLEC activity around $\vec{p}(t)$. As described in Sec. 8.2.4, we take a fixed variance threshold Σ to evaluate the reliability of the sLEC cell activity. We further assume that only if $\sigma \leq \Sigma$ the signal $\vec{p}$ is suitable for re-calibrating the robot (Fig. 8.12 (b)). More precisely, we define a weight coefficient

$$\alpha(t) = \begin{cases} 1 - \frac{\sigma(t)}{\Sigma} & \sigma(t) \leq \Sigma \\ 0 & \text{otherwise} \end{cases} \tag{10.4}$$

and then we use it to compute the calibrated robot position $\vec{p}^{\,*}(t)$ according to

$$\vec{p}^{\,*}(t) = \alpha(t)\ \vec{p}(t) + (1 - \alpha(t))\ \vec{p}^{\,\prime}(t) \tag{10.5}$$

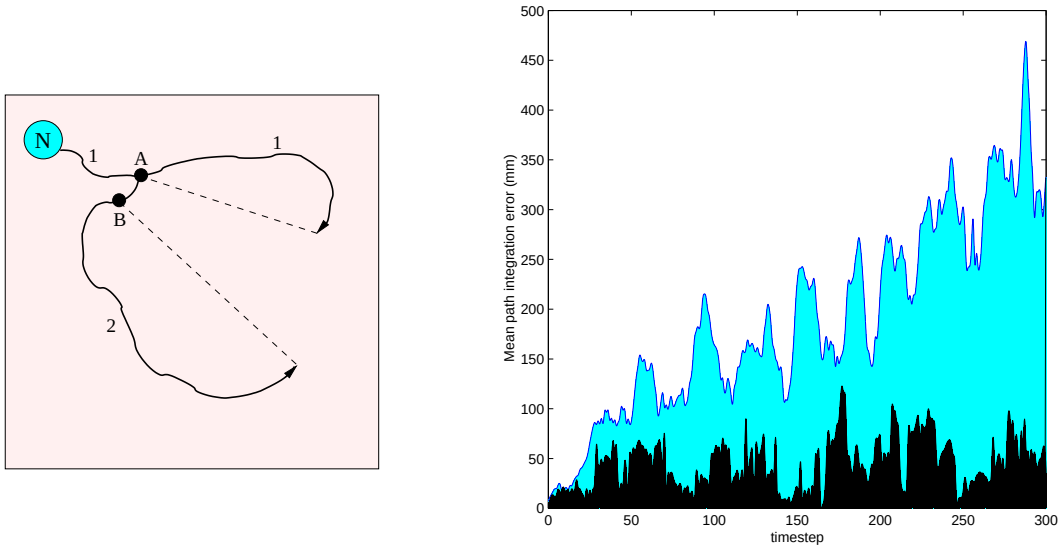

(a) (b)

Figure 10.6: (a) Sample of exploring excursion from a starting point (N). During exploration (solid trajectory n.1) the animat recruits new place cells to model spatial locations. After a while the robot stops exploring and follows its homing vector (dashed line). During homing, spatial learning does not occur (i.e., the robot does not recruit new place cells). Exploration is resumed as soon as a previously visited position A is found in which the path integrator can be calibrated based on visual information. The process is then iterated (excursion n.2) to propagate exploration over the entire environment by keeping the dead-reckoning error bounded. (b) Uncalibrated (lightgrey curve) and calibrated (black curve) mean path integration error.

where $\vec{p}'(t)$ is the position estimated by dead-reckoning at time t. Eq. 10.5 is an algorithmic implementation. In the future, we will calibrate path integration by applying associative learning to correlate sLEC activity with sMEC activity.

Once the robot has calibrated itself, exploration is resumed and the robot starts creating new place cells (Fig. 10.6 (a)). This technique allows the agent to propagate exploration over the entire environment by keeping the dead-reckoning error bounded. This means that allothetic (sLEC) and idiothetic (sMEC) representations are maintained mutually consistent over time. As in Sec. 9.1.1.1, we run a series of n exploring trials all starting at $t = 0$. At each step t, we measure the difference $e(t)$ between the robot's true position $\vec{p}(t)$ and the estimation $\vec{p}'(t)$. The experiment involves $n = 5$ trials, each of which includes 300 exploration steps. At the beginning of each trial the robot is placed at the same starting location $\vec{p}(t_0)$ (home). Fig. 10.6 (b) shows the calibrated and the uncalibrated mean error e. While the uncalibrated system exhibits a cumulating shift over time, the calibrated error is bounded and has an average value of about $4.4\,cm$.

Note that, during the homing behavior, the robot might reach the starting location without having re-aligned its path integration, i.e., without having found a location where sLEC activity is suitable for calibration[2]. In this case, the robot would resort to a spiral searching behavior centered around the home location. After having found a calibration point, the open-field exploring behavior would be resumed.

Another important concern is that an initially unfamiliar environment has to be explored uniformly in order to model all locations adequately. The robot adopts a simple active-exploration strategy that helps to cover the environment uniformly. At each time step t, it updates its direction of motion ϕ based on CA3-CA1 activity. If a relatively large number of neurons are currently active, it means that a well known region of the environment is being visited. Then, a small directional change, $\Delta\phi_s$, will

[2]This case has never occurred in our experiments.

increase the probability of leaving that area. Conversely, a large variability of the robot's direction, $\Delta\phi_l$, is associated to low CA3-CA1 place cell activity, which results in a thorough exploration of that region. In our experiments $\Delta\phi_s$ and $\Delta\phi_l$ are randomly drawn from $[-5°, +5°]$ and $[-60°, +60°]$, respectively.

10.2.1 Entering a familiar environment

In this thesis, we do not address the problem of consolidating and recalling hippocampal maps [283]. Rather, we simply assume that re-entering a familiar environment results in the attempt to retrieve the previously learned hippocampal chart. In particular, the agent needs to re-align its internal representations (i.e., head-direction, allothetic space code, and idiothetic chart) in order to re-activate a coherent description of the environment. Such a re-instantiating process relies on the coupling between the different components of the model established during training by associative learning. In a highly simplified view (assuming a constellation of external cues that can just rotate as whole between sessions) we can distinguish four different situations:

1. At the beginning of the trial, the agent is placed at the nest location[3] without undergoing disorientation before entering the arena (i.e., consistency is maintained between the animal's directional framework outside and inside the experimental arena). In this condition, neither the head-direction system nor the path integrator need to be re-aligned.

2. The animat is placed at the nest location but it is disoriented before entering the arena. In this case, the agent needs to polarize its directional sense by using a visual reference (e.g., a light source) experienced as stable during training (see Part II, in particular Chapter 5).

3. The animat is entered the arena at a random location $\vec{p}(t_0)$ without undergoing disorientation before starting the session. In this case, the robot has to re-align the idiothetic space chart (sMEC) with respect to the allothetic map (sLEC), i.e., to reset its path integrator according to visual information (Sec. 10.2).

4. At the beginning of the trial, the animat is placed at a random location $\vec{p}(t_0)$ and undergoes disorientation. Therefore, both head-direction cells and path integration need to be re-aligned. To do this, the agent has to retrieve a clean positional signal at the level of the allothetic space code, i.e., a clean blob in the sLEC ensemble activity. Since sLEC firing relies on head-direction information, the only thing the agent can do is to resort to a random search in the visual input space by taking views in random directions by turning on the spot. Given that the animat's entry position $\vec{p}(t_0)$ is a location that has been experienced during training, this search should converge to an activation of the subset of sLEC place cells coding for $\vec{p}(t_0)$. Once the allothetic representation has been restored, the head-direction system can be calibrated according to the technique described in Chapter 4. Also, path integration can be reset as described in Sec. 10.2.

Failing to retrieve a coherent representation of the environment (i.e., to mutually re-align head-direction and place cells, allothetic and idiothetic charts) results in a failure in "recognizing" a previously experienced arena. Then, a between-session remapping (i.e., the development of a new hippocampal representation of the environment) is likely to occur [258, 50, 166, 25]. Since the retrieving procedure (points 2, 3, and 4) relies on LTP correlational learning, impairing this mechanism would result in unstable representations between separate sessions within a same environment. This is consistent with experimental findings showing that old animals with deficient LTP exhibit stable hippocampal maps within sessions, but unstable mapping between separate runs [25, 283].

[3] We assume that the nest is a location (e.g., a home box [99]) that the agent can easily recognize across sessions and that provides the reference point for the path integration framework.

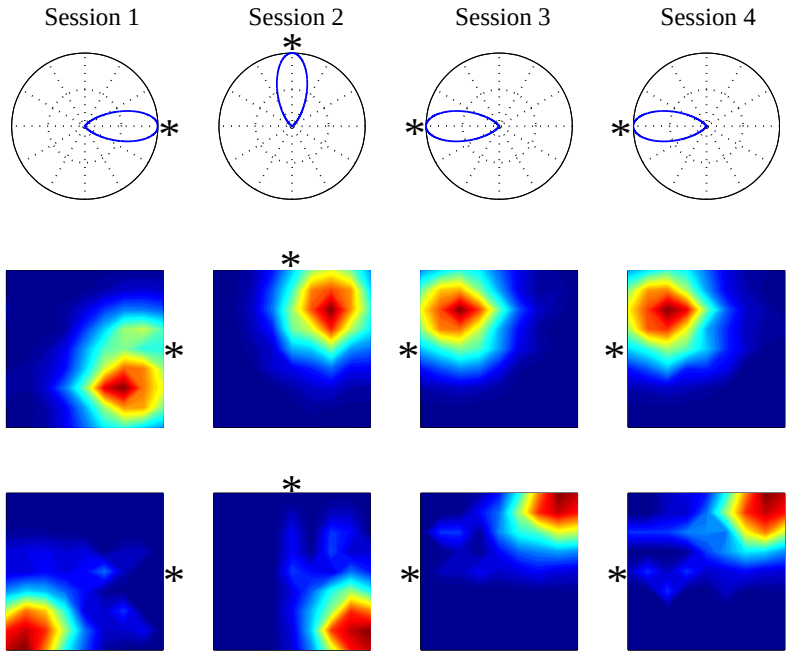

Figure 10.7: A poSC head-direction cell and two CA3-CA1 place cells recorded from the robot trained under non disorienting conditions. At the beginning of each probe trial, the robot undergoes disorientation. The head-direction cell as well as the two place cells are controlled by the light cue (whose position is indicated by the '*' next to each diagram).

10.3 Studying the Interaction between External and Internal Cues

In Chapter 5 we have studied the influence of external directional cues upon the head-direction system. In particular, we have done some experiments inspired by those performed by Knierim *et al.* on rats [166] (Sec. 5.1.3). Consistent with neurophysiological data, we have shown that "disorienting" our robot before each training session reduces the control exerted by an external landmark (e.g., a light source) upon head-direction cells during probe trials (Fig. 5.4). By contrast, the robot trained under non disorienting conditions is able to re-align its direction representation across test trials based on the external directional reference (Fig. 5.3).

In this section, we present results obtained by recording our CA3-CA1 place cells in a similar experiment[4]. The objective being to demonstrate that our head-direction and place cells are two coupled systems, so that if the direction representation rotates by an angle $\Delta\theta$, a comparable rotation will be observed in the hippocampal place code.

We run 2 training series of 10 trials each. A single trial lasts about 15 minutes. At the beginning of each trial the robot is placed at the center of the arena (i.e., nest position). First, the animat turns on the spot and learns to associate its inertial directional sense to the external reference provided by a light source. Then, random exploration is triggered and spatial learning takes place. In the first series of experiments, the robot is trained under non disorienting conditions. In the second series, the robot's direction representation is initialized randomly (i.e., the animat is disoriented) before each training trial.

After training, we run 2 probe series of 4 trials each. In the first series, we record from the robot trained under non disorienting conditions. In the second, we take the robot that underwent disorientation. We record poSC head-direction cells and CA3-CA1 place cells of the model.

[4]For this experiment we have utilized our linear-vision setup.

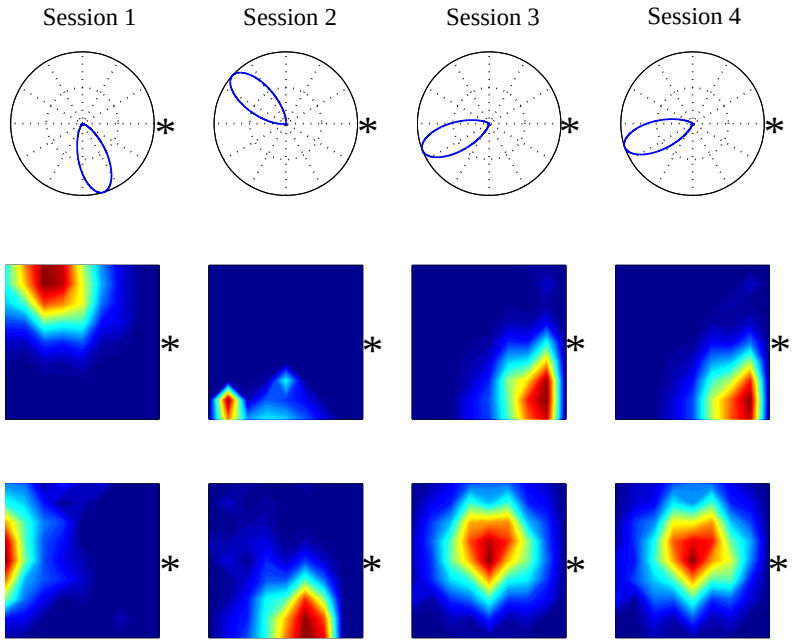

Figure 10.8: Effects of disorienting training on poSC head-direction and CA3-CA1 place cells. At the beginning of the four recording trials, the robot fails to calibrate its directional system with respect to the allothetic cue. The light (whose position is indicated by the '*' next to each diagram) does not influence the dynamics of the direction and place systems. Note that head-direction and place cells are strongly coupled and always rotate consistently with each other.

At the beginning of each recording trial, the robot is "disoriented" and placed inside the arena at the nest location (i.e., case 2 in Sec. 10.2.1). Then, it first tries to re-align its direction representation relative to the external cue. This process usually succeeds for the robot that did not undergo disorienting training. Indeed, with respect to its inertial system, the light has been perceived as stable during training. By contrast, the robot trained under disorienting conditions usually fails to recalibrate its head-direction cells, and keeps its directional sense as randomly initialized at the beginning of the trial.

Fig. 10.7 shows some results obtained during the first recording series. The tuning curve of one poSC direction cell as well as the place fields of two CA3-CA1 neurons are displayed. The light source exerts a strong control upon both directional and place coding: the head-direction cell and the place cells rotate their firing patterns following the light reference. Importantly, head-direction and place cells always rotate consistently with each other, forming a unitary system.

Fig. 10.8 illustrates results obtained with the robot trained under disorienting conditions. The external cue does not influence the directional and place representations: despite its constant position, recorded cells vary their receptive fields across trials. Since the robot fails to re-align the components of its spatial learning system, between-session hippocampal remapping occurs (even if the robot enters a familiar environment). Note that, recorded cells maintain similar receptive fields during sessions 3 and 4 (i.e., remapping does not occur). This is solely due to the fact that at the beginning of these trials the robot's directional system has been initialized to two similar random configurations.

Results shown in Figs. 10.7 and 10.8 are consistent to those reported by Knierim *et al.* when recording from freely-moving rats [166].

10.4 Discussion

In this chapter, the different components of our spatial learning system (i.e., head-direction cells, vision-based place code, and path integration) are combined to form a unitary system. We stress the importance of integrating external and internal signals to drive place cell activity in CA3-CA1. The idiothetic pathway provides a spatial framework suitable to compensate for unreliable visual information. Conversely, visual cues allow the path integrator to maintain a coherent representation over time. This closed sensory loop results in a stable representation in our hippocampus proper. We also stress the importance of the coupling between head-direction and place cells. Their mutual interaction is a crucial issue with respect to the self-localization task. Consistently with experimental results, disrupting our poSC layer, for instance, would severely impair spatial learning [352].

The system relies upon the interaction between separate maps (i.e., direction representation, allothetic and idiothetic space codes), so that entering a familiar environment results in retrieving a coherent description of the world by realigning the above representations. Such a recalling mechanism relies on experience-dependent LTP coupling. Failing to retrieve a consistent internal representation, leads to hippocampal remapping. Thus, the model is consistent with the fact the LTP-deficient animals (e.g., old rats) may develop a new hippocampal representation upon arrival into a familiar environment.

Our place cell assembly grows incrementally during exploration. Redundancy in the place code is considered as a crucial property to yield robustness. The model produces a spatial representation consisting of a large population of localized overlapping place fields covering the environment uniformly and densely. Population vector coding is employed to interpret the ensemble CA3-CA1 activity [394].

The adopted working hypothesis is compatible with Tolman's latent learning concept [361]: The hippocampal space representation is built in the absence of any specific reward, and a sort of "curiosity" solely motivates exploration [367]. Nevertheless, neurophysiological findings show that place cells can become strongly correlated to reward locations [209, 391, 51, 239] (Chapter 6). A way to understand this consists of considering the dependency of the hippocampal representation on multivariate sensory inputs: Among several sources of information, the animal might use the reward signal received at location $\vec{p}$ to tune up one pyramidal cell i coding for $\vec{p}_i \approx \vec{p}$. If the reward information becomes predominant with respect to the other sensory inputs (e.g., vision), then cell i might become mostly correlated to the reward signal. Then, moving the rewarding location would result in a shift of the place field of i.

The model captures some properties of biological place cells (Chapter 6): It produces plausible CA3-CA1 place fields [254, 243, 242], anatomical topology does not occur in the hippocampal ensemble [254, 177, 357, 240], place cell firing patterns are uncorrelated across different environments [177, 243, 240, 357, 278], rapid learning of place fields occurs [147, 394, 18], asymmetry in the CA3-CA1 firing persists even in visually symmetric environments [322, 324]. Due to the contribution of the idiothetic pathway, cells in the CA3-CA1 layer of the model exhibit clean location selectivity even when allothetic cues are removed or in complete darkness [254, 148, 271, 240, 258, 212, 278]. In addition, since our idiothetic representation depends on the entry position, the latter plays an important role in determining the influence of visual cues upon place field locations [324]. Also, according to experimental findings [32, 166], the more stable an allothetic cue (e.g., a light cue) is perceived by the robot, the more it will control the dynamics of our CA3-CA1 place cells.

The model is compatible with anatomical lesion data showing that damaging the hippocampus proper disrupts spatial cognition [236, 237], whereas does not affect path integration [4]. Lesions to the entorhinal layer of our model (i.e., sLEC and sMEC substrates) would seriously impair spatial selectivity of pyramidal cells in CA3-CA1 (consistent to [220, 261, 307, 127, 278]). The model predicts the damaging sMEC would diminish the influence of idiothetic signals upon CA3-CA1 firing, whereas lesioning sLEC would impair the influence of visual information. Finally, since our path integrator is reset by taking

the allothetic sLEC code as calibrating signal, the model predicts that cutting off the sLEC → sMEC projections would impair the animal's capability of maintaining a coherent idiothetic space code.

Much has to be done to account for several other important properties of biological place cells. For instance, the current system does not model the phase relationship between the theta rhythm and hippocampal place cell firing [257, 330]. Another property that is not captured by the model is the experience-dependent reshaping of CA1 place fields observed when animals experience several times a route [215, 216]. Also, the role of the dentate gyrus (DG) is currently neglected by the model. To conclude, experiments involving environment manipulations (e.g., shrinking and stretching the experimental arena) as well as tests in changing environmental conditions (e.g., changing light) need to be performed to replicate several neurophysiological data (see Chapter 14 for a discussion concerning future work).

Part IV

Spatial Behavior: Goal-oriented Navigation

Foreword

The spatial behavior of a rat is likely to be the result of a continuous combination of different navigation modalities [114, 274]. Which strategy predominates at any time depends on the sensory information available to the animal [287]. A simple stimulus-response strategy can be used to finding a food source when the latter is visible or it smells (i.e., taxon navigation [256, 309]). Applying a specific motor sequence can be used to solve tasks in which the target is identified by a succession of cues (i.e., praxic navigation [272]). However, in more complex navigation tasks (e.g., when the target cannot be directly perceived by the animal, as the hidden-escape platform in the water maze [233]), goal-oriented navigation needs to be supported by spatial learning (i.e., locale navigation [256]). As mentioned early in this dissertation (Part III), the hippocampus provides a neural support for such a cognitive spatial behavior. Lesion data indicate that the hippocampus is important for locale navigation but not for taxon navigation [236, 237, 267] (Chapter 6).

The hippocampal place cell activity encodes locations but does not provide information about how to navigate from one place to another. In this part of the thesis we focus on locale navigation and we ask the following question: How can target-directed cognitive behavior be achieved based on hippocampal place fields?

- In Chapter 11 we discuss some general issues concerning navigation capabilities of animals and we review experimental data on rodents.

- In Chapter 12 some previous models of hippocampal-based navigation are reviewed. We also present results obtained by implementing: *(i)* The model by Abbott, Blum, and Gerstner [1, 46, 122] using our place field substrate as underlying space representation (Chapter 10); *(ii)* The navigation part of the model by Burgess, Recce, and O'Keefe (results concerning the implementation of the space-code part have been reported in Sec. 7.1.2).

Contributions: We propose an action learning scheme that allows the agent to acquire navigational maps based on rewarding signals. Our place fields provide a coarse coding representation which is suitable for applying reinforcement learning in continuous space [343]. Furthermore, since the established place code is goal-independent, it provides a basis onto which navigation to a newly added target (or to several distinct targets) can be quickly learned, i.e., latent learning (the navigation model described in this part of the thesis has been published in [14, 16]):

- In Chapter 13 we describe the reward-based learning scheme and we present results obtained by validating the model with the Khepera robot.

Chapter 11

Goal-oriented Behavior in Animals

In Chapter 6, we have discussed neurophysiological as well as anatomical-lesion data supporting the hypothesis that the hippocampus endows animals with cognitive spatial behavior. In this chapter, we consider behavioral studies, that is, experiments to investigate, through behavior analysis, the ability of animals to reach interesting places. Returning home, finding a food source, avoiding dangerous areas: all these tasks demand effective navigation skills, and can be reproduced in controlled environmental conditions[1].

In particular, we review experiments concerning open-field navigation [233, 236, 76, 32, 33, 132, 337]. We do not provide an exhaustive experimental review[2]. Rather, we provide some elements that are relevant to the concepts treated in this part of the thesis. To this extend, for instance, we briefly describe some properties of the nucleus accumbens and of dopaminergic neurons which will support our discussion about reward-based learning.

11.1 Goal-oriented Navigation in the Water Maze

The water maze experimental setup is one of the most utilized for studying hippocampally dependent navigation. Since it was first introduced by Morris in 1981, a huge amount of data has been collected about rodents' capability of learning this task [233, 236, 307, 237, 385, 25, 337]. In particular, many lesion studies have been done by employing this open-field environment (Chapter 6).

As shown in Fig. 6.5 (a), the Morris water maze is a large circular pool (e.g., $1-2$ m of diameter) containing opaque liquid (e.g., milk mixed with water). The navigation task consists of reaching an escape platform (e.g., 0.1 m of diameter) located somewhere within the pool. The animal must learn to navigate to the platform from any location. The task is termed *hidden escape-platform* water maze when the platform is submerged just below the surface. Alternatively, it is termed *visible escape-platform* water maze if the platform is above the water surface. Finally, a *cued-platform* water maze is a task in which a local landmark (e.g., a ball hanging above the platform) indicates the target location. In the following, we will concentrate on the hidden escape-platform task. Two major protocols have been defined:

[1] A controlled environment is an experimental setup in which most of the variables that can influence the animal behavior (e.g., light, noise, odor, visual cues) are controlled by the experimenter.

[2] Refer to [361, 256, 114, 282, 371] for more comprehensive experimental overviews.

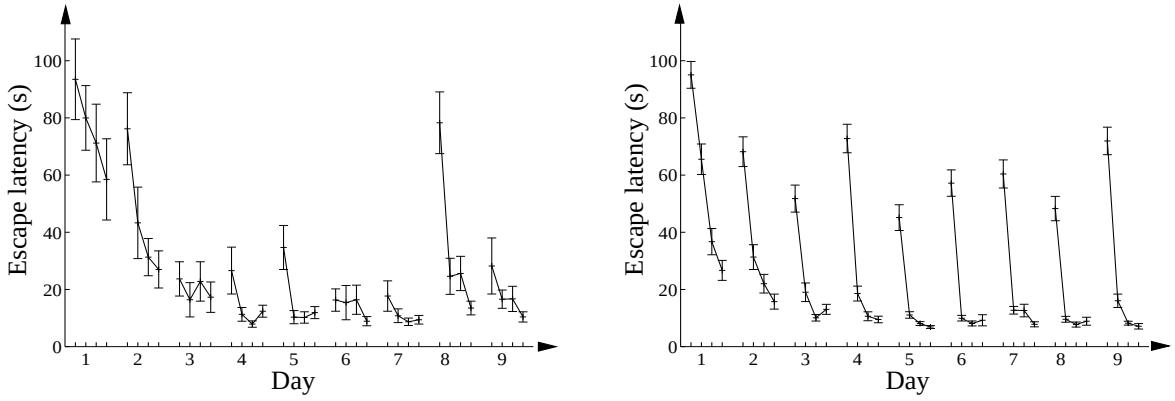

(a) (b)

Figure 11.1: Experimental data showing the performance of rats on solving the reference memory water maze RMW, (a), and the delayed matching-to-place water maze DMP, (b). Animals are given four trials per day during consecutive days. For each trial, the time needed by the animal to find the hidden platform is measured (i.e., escape latency). Data courtesy Foster, Morris and Dayan [337, 107].

(i) The reference memory water maze task (RMW) [233]: The platform has a fixed location during the entire training. Rats are given several learning trials per day (e.g., four), for several consecutive days. At the beginning of each trial, the rat is entered the pool at one of four starting locations (e.g., north, south, west, or east edge). Using different entry positions over trials reduces the possibility of solving the task by praxic navigation (i.e., by simply re-executing a sequence of locomotor actions that worked successfully previously). As shown in Fig. 11.1 (a), normal rats can learn the RMW task in a few trials. Performance are measured in terms of mean escape latency (i.e., the average time needed by the animal to find the platform). In this experiment, rats received four training trials per day. The diagram shows that they are able to navigate directly to the platform after less than 20 trials (i.e., escape latencies become very short during days $5-7$). However, if the platform is subsequentially relocated, animals need some re-training before exhibiting direct navigation trajectories to the new target location. Nonetheless, learning the new task occurs rapidly (i.e., escape latencies during days $8-9$).

(ii) The delayed matching-to-place water maze (DMP) [337]: The hidden platform is maintained in the same position throughout the day, but it is given a new location at the beginning of each day. Fig. 11.1 (b) shows the performance of rats when solving the DMP task. Interestingly, animals exhibit one-trial learning after a few days of training. Indeed, as shown in Fig. 11.1 (b), they can swim directly to the platform on the second trial of day 6. Thus, rats are able not only to avoid interference between consecutive training days, but they can also generalize from previous experience to improve navigation during later days [107].

11.2 Navigation based on Local-landmark Information

Collett and colleagues, in 1986, performed a series of behavioral experiments to investigate the ability of animals to use local-landmark information (within an otherwise impoverished open-field environment) to achieve goal-directed navigation [76]. In particular, they were interested in understanding *(i)* what a rodent can learn about the geometrical relationships between a goal and nearby visual cues, *(ii)* how it uses this information to navigate to a target. They trained gerbils to find a feeder location identified by an array of cylindrical landmarks. In this experiment, the animal's entry position as well as the location

of the landmark array changed randomly across training trials. However, the food source remained at the same position relative to the landmark array. After training, probe tests were run by removing the food source and measuring the time spent by animals at each location of the environment.

Experimental results reveal that gerbils are able to learn the spatial relationship between the target location and either a single local cue or an array of local landmarks [76, 282]:

(i) Single landmark: Gerbils can easily learn to find a food source located at a specific bearing and distance to a single cylinder (Fig. 11.2 (a)).

(ii) Two landmarks: When trained to search for food at a position identified by two local landmarks, gerbils are able to combine the bearing-distance information relative to both cues (Fig. 11.2 (b)). If probe trials are run with only one cylinder, animals search alternatively in two locations, each of which located at the learned distance and bearing from one of the experienced landmarks (Fig. 11.2 (c)). If, after training, the two-landmark array is stretched (e.g., the distance between the two cylinders is doubled), gerbils concentrate their search on two locations, each at the learned distance-bearing from the corresponding landmark (Fig. 11.2 (d)).

(iii) Three landmarks: Animals are able to find a target located at the center of a configuration of three landmarks (Fig. 11.2 (e)). Adding one more cylinder during probe trials does not modify the location where gerbils expect to find food, i.e., the fourth landmark is ignored, (Fig. 11.2 (f)). Removing one cylinder during test does not impair animals' searching behavior (Fig. 11.2 (g)). If only one landmark is used during probe sessions, gerbils search at three distinct locations equally distant from the cylinder but with bearings relative to each experienced landmark (Fig. 11.2 (h)). Finally, if one element of the three-landmark array is relocated with respect to the original configuration (e.g., the triangular array is stretched), gerbils continue to search at the correct location by using information provided by the two unchanged cues (Fig. 11.2 (i)).

Given the above findings, Collett and colleagues concluded that gerbils are able to establish a complete representation of the environment which allows them to store (or to compute from) the geometric arrangement of landmarks and goal [76].

11.3 The Role of Dopamine in Reward-based Learning

Neurophysiological data show that dopamine neurons in the ventral tegmental area (VTA) encode the difference between expected and actual occurrence of reward-related stimuli [189, 312]. Indeed, the response of these neurons is a function of the unpredictability of incoming reward (or reward-predicting) stimuli: They respond positively to rewards which occur unpredictably. Instead, they remain silent if a fully predicted stimulus arrives. Conversely, when a predicted reward fails to occur, dopamine neurons respond negatively[3] exactly at the time at which the reward is expected.

The response of dopamine neurons occurs for specific rewards. For instance, when the animal touches a small morsel of hidden food. However, if a similarly shaped non-food object is given to the animal, there is no dopamine response. A further property is that dopamine responses are independent of the direction of the stimulus relative to the body axis. Also, reward-related stimulation activates most of dopamine neurons in parallel. This suggests that these neurons respond as a population rather than exhibiting distinct response profiles among each other. Finally, dopamine neurons respond mostly to unconditioned rewards as well as to conditioned reward-predicting stimuli. They also respond to novel

[3] A positive (negative) response occurs when neurons fire more (less) action-potentials with respect to their background activity. Similarly, a silent dopamine neuron is one that does not vary its activity when a rewarding stimulus arrives.

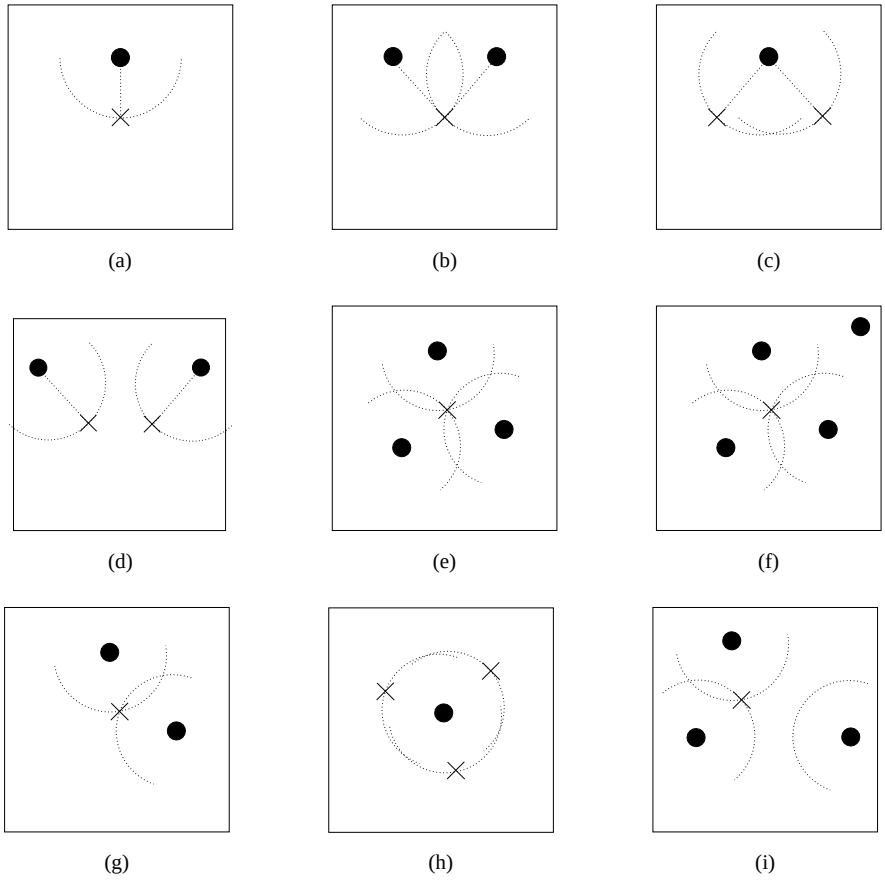

Figure 11.2: Experimental results showing the performance of gerbils in searching for a food source at a position identified by a landmark array. Filled black circles indicate cylindrical landmarks, whereas black crosses indicate the locations mostly visited by animals during probe trials. Adapted from Collett *et al.* [76].

or arousing stimuli. A novel stimulus could be a potential reward or a reward predictor and it can be considered as a reward-related event [310].

Dopamine neurons project to the frontal cortex and to the striatum (i.e., caudate nucleus, putmen, and ventral striatum including nucleus accumbens). These structures seem to be involved in processing information concerning behavioral actions (e.g., goal-oriented movements). In particular, neurons in the ventral part of the striatum are primarily activated in relation to the expectation of rewards [311]. Thus, striatal cells might be informed by dopamine neurons about the occurrence of unconditioned or conditioned reward-related events. The presence of dopamine-dependent plasticity in the striatum suggests that dopamine responses might be involved in plasticity changes yielding reward-based learning [310]. Indeed, since the activity of dopamine neurons is a function of the unpredictability of a reward, dopaminergic activity might work as a prediction error signal (i.e., the difference $R(t) - R'(t)$ between the actual reward $R(t)$ at time t and the predicted reward $R'(t)$) suitable for learning. Then, striatal plasticity may be exploited to modify synaptic efficacy based on the dopamine prediction error [312].

11.4 The Nucleus Accumbens

As mentioned in Chapter 6, the nucleus accumbens (NA) seems to be involved in the neural mechanisms underlying cognitive spatial behavior. Experiments show that NA lesions disrupt the performance of rats

on learning the hidden water maze, but not the visible-platform task [341].

In particular, the nucleus accumbens seems to play an important role in spatial locomotion based on reward-dependent learning [387, 55]. Indeed:

(i) NA receives space coding information from the hippocampal formation via the fornix fiber bundle [344, 8];

(ii) NA receives dopaminergic inputs from the brainstem (via the ventral tegmental area) which are likely to transmit rewarding brain stimulation [188, 315];

(iii) NA may influence locomotion through its efferent projections to pallidal, hypothalamic, and mesencephalic structures [139, 75]. Furthermore, when testing hippocampal lesioned rats with disrupted dopaminergic input to the striatum (including NA), Whishaw and Mittleman [387] found that combined hippocampal-accumbens lesions yield a decrease of the animal's locomotor behavior [387].

Therefore, the nucleus accumbens can be thought of as a locomotor-related structure in which spatial information as well as reward-based signals converge [55]. This suggests that NA might be strongly implicated in target-dependent action learning.

Chapter 12

Modeling the Role of Hippocampus in Navigation: State of the Art

Since the first experimental findings have shown the hippocampal role in rodent cognitive behavior, several models have been proposed to achieve hippocampally dependent navigation [62, 377, 55, 308, 46, 122, 116, 368, 107]. Some of these models assume the existence of an 'ideal' place cell representation [377, 55, 46, 122, 368, 107], some others are also concerned with establishing place fields from sensory inputs [62, 308, 117]. Despite their differences, all these approaches share the idea that complex navigation (i.e., not merely stimulus-response mapping) can be effectively accomplished based on hippocampal space coding.

In this chapter, we review some of the above navigation models[1]. We also present simulation results obtained by implementing two of these approaches: The model by Burgess, Recce, and O'Keefe [62], and the model by Abbott, Blum, and Gerstner [1, 46, 122] (Part III).

12.1 Burgess, Recce, and O'Keefe (1994)*

This model postulates a goal-memory system in which each target is represented by a set of goal cells (*GC*) one synapse downstream of hippocampal place cells [62, 61]. GC activity encodes the animal's position with respect to the goal (e.g., cell GC_N will fire maximally whenever the rat is north of the target). In particular, each goal cell is assumed to have a conic firing distribution over the entire environment peaked close to the goal (e.g., GC_N has a firing map peaked to the north of the target). As a consequence, since GC activity decreases as the distance to the goal increases, the animal can estimate both its direction and its distance to the target. In order to approximate the direction of the animal to the goal, population vector coding is applied to interpret the ensemble goal cell activity.

One-shot Hebbian learning is applied to modify the connections from place cells to goal cells whenever the animal receives the corresponding reward. In particular, the reinforcement signal is gated by the head-direction system before reaching the goal cell population, such that, for instance, cell GC_N receives

[1]Those models that do not assume the existence of place fields but that create them from sensory inputs, have already been introduced in Chapter 7. Here, we only discuss their navigation part. A label '*' is used to distinguish these models from those that do assume an ideal place field population.

an excitatory input only when the rat is at the target location and faces north. Whenever a goal cell GC_i receives the reward signal, connections from the currently active place cells to GC_i are "switched on". Below is scheme of the model:

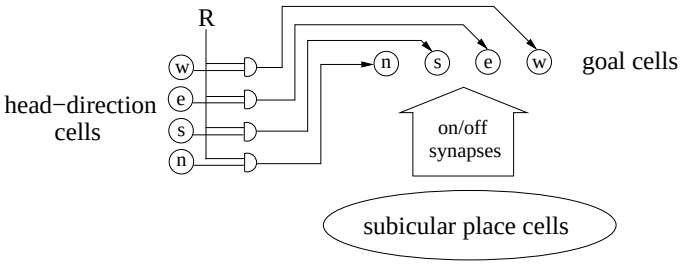

The authors assume that *(i)* the reinforcement signal R always arrives at a late phase of each theta cycle, *(ii)* whenever the rat finds a rewarding location, it turns on the spot and looks in several directions. Since learning occurs at a late phase of theta, when the rat is at the goal location and faces north, for instance, connections to cell GC_N will be created from those place cells having place fields centered to the north of the goal[2]. Thus, after learning, whenever the rat is north with respect to the goal, cell GC_N will fire more strongly than each other cell GC_i. This will provide the animal with the direction to be taken in order to navigate towards the goal.

An intrinsic limitation of the hypothesis by Burgess, Recce, and O'Keefe is that only those place cells whose place field contains the target are correlated to the goal cells by Hebbian learning. This results in goal cells of limited attraction radius which impairs agent's navigation at large distances from the target. For this reason, the authors postulate that place cells driving GC activity are not located in the hippocampus proper but in the subiculum. Indeed, subicular cells have broader place fields than CA3-CA1 cells (Chapter 6). However, this does not solve the problem of loosing navigation information at very large distances from the target (i.e., for distances larger than the size of subicular place fields).

In Fig. 12.1 we show some simulation results obtained by implementing the above navigation model[3]. Figs. 12.1 (a, b, c) show, respectively, the simulated environment with the target (darkgrey square) at the center, a navigational map acquired after training, and an example of goal-directed behavior performed by the agent (lightgrey circle) based on the learned map. Fig. 12.1 (d) shows the same environment with the goal located near to the bottom-right corner. As a consequence, the maximal distance the robot has to travel to reach the goal is larger than in the previous experiment. Figs. 12.1 (e, f) show, respectively, a successful and a failing goal-oriented path. As previously discussed, the above navigation model suffers from a "distal reward problem" [107]. Indeed, in Fig. 12.1 (f) the simulated agent becomes lost at the upper-left corner (where the distance to the goal is maximal) since it does not receive any directional signal from the population of goal cells.

12.2 Wan, Redish, and Touretzky (1994, 1997)*

Hippocampal-dependent navigation, in the model by Wan, Redish, and Touretzky [377, 286], relies upon vector subtraction within an allocentric coordinate system. As we have seen in Sec. 7.1.3, the authors put forward a theory of rodent navigation in which place cell activity depends on both local views (i.e., allothetic-based place recognition) and metric information relative to a reference point (i.e., path integrator coordinates). On entering a familiar environment, the agent uses visual cues to activate the place code for the current location $\vec{p}(t)$, and it recalls the path integrator coordinates (x_p, y_p) associated

[2] As seen in Sec. 6.4, a place cell i firing at a late phase of theta tends to have a place field centered ahead of the animal.

[3] The part of the model concerning the acquisition of the space representation has also been implemented (Sec. 7.1.2).

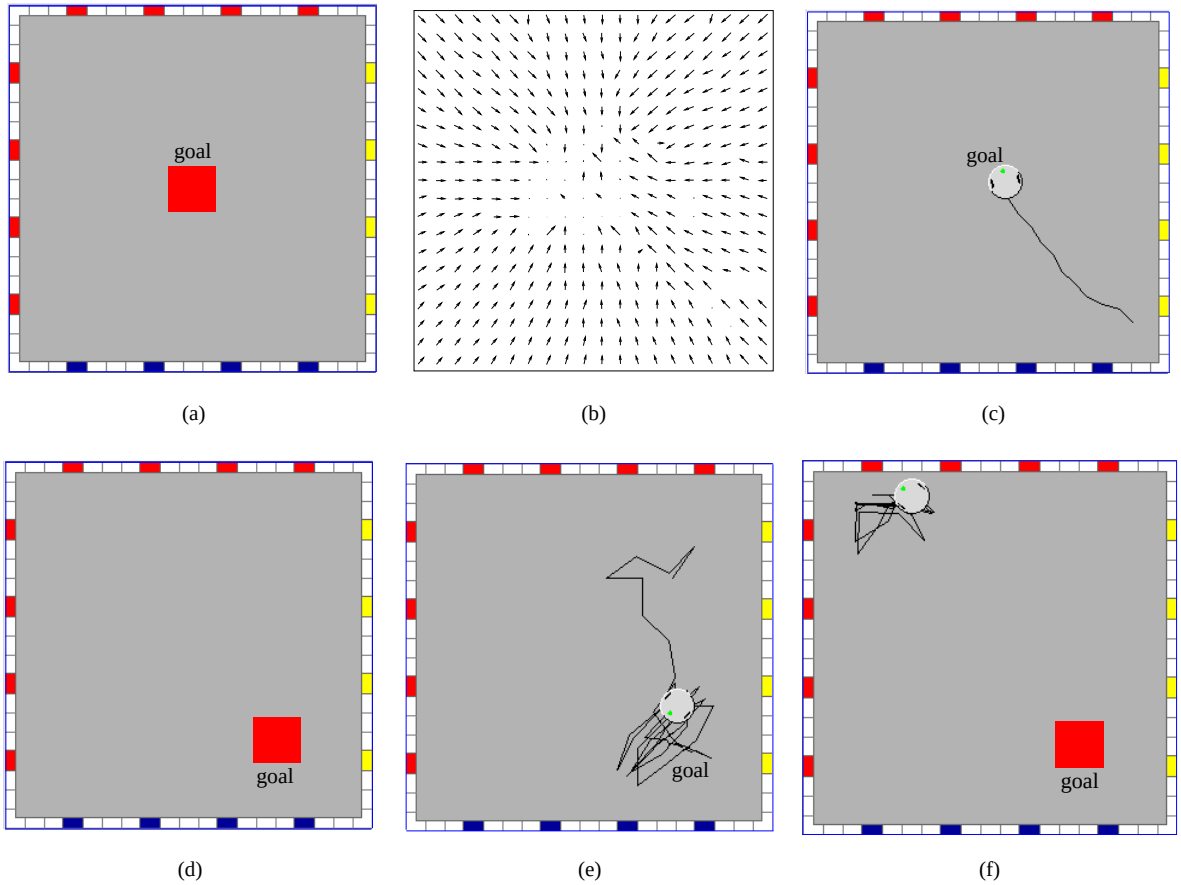

Figure 12.1: Some results obtained by implementing the navigation model by Burgess, Recce, and O'Keefe [62].

with place $\vec{p}$. Once the animal has reconstructed its position within the metric reference frame, it can estimate the position of the goal within this coordinate system, $(\overline{x}_g, \overline{y}_g)$[4], and it can compute the direction $\vec{d}_g$ leading to the goal by vector subtraction, i.e., $\vec{d}_g = \binom{\overline{x}_g}{\overline{y}_g} - \binom{x_p}{y_p}$. Note that this does not provide explicit obstacle avoidance information.

The authors model behavioral data from Collett and colleagues [76], in which gerbils have to find a food location identified by an array of local landmarks in an otherwise impoverished environment. Computer simulations reported by Wan, Redish, and Touretzky show goal-position estimates that are consistent with experimental data. The model, however, does not produce explicit locomotion behavior towards a rewarding location. Rather, the simulated gerbil is placed at 100 randomly-chosen locations in the environment. For each trial i, $1 \leq i \leq 100$, the model reconstructs its entry position (x_p^i, y_p^i), and from that it simply provides an estimate of the goal location $(\overline{x}_g, \overline{y}_g)$.

12.3 Brown and Sharp (1995)*

Similar to the model by Burgess, Recce, and O'Keefe, Brown and Sharp propose a neural system in which place and direction codes are mapped into actions based on reward signals [55]. An important difference, however, is that they model motion directions egocentrically (e.g., turn left), rather than allocentrically (e.g., the goal is to the north).

[4]Given that the goal position (x_g, y_g) has been learned by the agent during training.

The authors concentrate on the nucleus accumbens (NA) and model it as two separate clusters of 60 output units each: one cluster is associated with left-turns, the other with right-turns. Each cell $i \in$ NA can be considered as a motor cell[5]. Each cluster also contains 60 inhibitory interneurons, each of which projects to a unique subset of 59 motor cells. Thus, whenever an interneuron becomes active, it inhibits all the motor units within its cluster but one. Interneurons are activated by hippocampal place cells. In particular, Brown and Sharp consider 60 place units in total. For each spatial location there is only one place cell firing (i.e., winner-take-all scheme). Each place cell projects to a unique interneuron within each cluster. Thus, for each position, the activation of one place cell mediates the inhibition of all but two motor units, i_l and i_r, one unit for each NA cluster.

Which of the two units i_l and i_r will actually fire is determined by the head-direction cells (HD): each cell $i \in$ NA receives one plastic excitatory afferent from each directional cell j and is activated according to $a_i = \sum_j w_{ij} r_j$. Thus, of the two motor neurons that have not been inhibited i_l and i_r, the one fires which has the highest activation (e.g., i_l, if $a_{i_l} > a_{i_r}$). The winner cell (e.g., i_l) will determine the agent's motion direction[6] (e.g., left). Below is a diagram of the model:

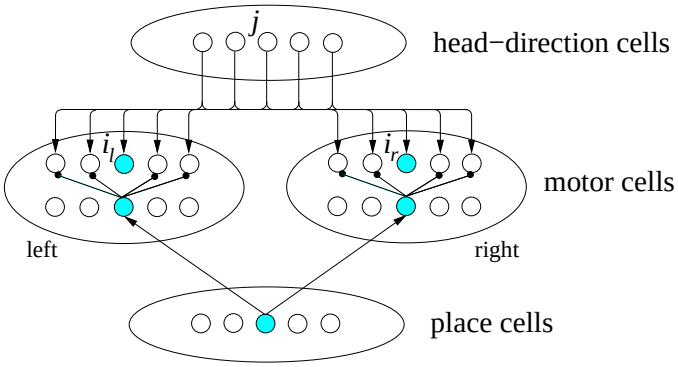

Synaptic weights w_{ij} from HD cells to NA motor cells are initialized randomly and changed whenever the agent reaches the target. In order to back-propagate the reward-triggered synaptic change (i.e., to update those HD→NA projections responsible for the actions taken during the agent's trajectory), Brown and Sharp use a memory trace: those synapses which led to more recently correct movements are strengthened the most, while synapses associated to early good actions are strengthened less.

The above model has been tested by computer simulations to learn the hidden-platform water maze (in particular, the reference memory task (RMW) as defined in Chapter 11). Results are comparable to experimental data [233].

12.4 Schölkopf and Mallot (1995)*

In the models by Burgess, Recce, and O'Keefe [62, 61] and Brown and Sharp [55], goal-oriented navigation is achieved by mapping places into local actions. Alternatively, navigation can be accomplished by path planning based on a topological map. The model by Schölkopf and Mallot [308] adopts this view.

As seen in Chapter 7, the central idea of the authors is that a topological view graph (rather that a place graph) is sufficient for navigation. Each node of the graph corresponds to a given view, while edges indicate temporal coherence between views: Two nodes i and j are connected if and only if the

[5]Note that, since actions are encoded within an egocentric framework it is appropriate to speak about motor units; real motor neurons would be neurons in the motor cortex or in the spinal cord.

[6]The agent rotates by an angle of 11° and it moves by a small distance.

corresponding views v_i and v_j are temporally sequential. Furthermore, each edge $i \rightarrow j$ is labeled by the egocentric movement (i.e., go left, right, back) leading the agent from view v_i to the subsequent view v_j.

Given this view-graph representation, path planning is employed to determine the minimal-length sequence of movements to go from the current view v_i to the desired view v_{goal}. In particular, path planning is performed by adopting a breadth-first search algorithm that can be seen as "mental exploration": *(i)* The goal view v_{goal} is "imagined" in order to find which node in the graph it corresponds to; *(ii)* The lengths of all possible paths from the current node v_i to v_{goal} produced by "imaging" all possible movements from v_i are evaluated; *(iii)* The movement μ producing the optimal path is selected; *(iv)* The process is iterated from the new current view v_j reached by executing μ from v_i.

In other words, this planning algorithm "explores mentally" all paths that have been experienced during exploration (as well as novel combinations of them) in order to find the optimal trajectory between arbitrary points in the representation. As reported by the authors, the necessary concepts to implement the above path planning mechanism by a neural machinery have not been explicitly implemented in the model.

12.5 Abbott, Blum, and Gerstner (1996, 1997)

Abbott and Blum [1] put forward a theory concerning the functional understanding of NMDA-dependent LTP (Chapter 6). They address the following question: How does LTP-synaptic enhancement affect the behavioral read-out of a neural population? As a general outcome, they postulate that LTP enables a neural network encoding a quantity $\mathcal{A}$ to generate an experienced-based prediction of future $\mathcal{A}$ values. In particular, they argue that LTP-dependent training induces a shift in the ensemble code with important behavioral implications.

To validate their theory, Blum and Abbott [46] examine the effect of LTP on the ensemble place cell coding. They show that the shift arising from collateral synaptic potentiation provides a useful information to support goal-oriented behavior. The navigation model relies on two experimentally supported elements: *(i)* NMDA-mediated LTP is temporally asymmetric (i.e., only if presynaptic firing precedes postsynaptic activity by less than about $200\,ms$ LTP occurs, otherwise either there is no LTP or LTD takes place [183, 85, 1]); *(ii)* Population vector decoding can be applied to interpret place cell ensemble activity as physical locations [394].

The authors focus on CA3 collaterals and change their synaptic efficacy according to a reward-induced LTP. After training (consisting of a set of exploratory paths leading the agent to a target location), CA3 ensemble activity no longer encodes the agent's actual position $\vec{p}_1$, but a different location $\vec{p}_2$ nearby. This experience-induced shift in the spatial information encoded by the ensemble CA3 firing, $\Delta\vec{p} = \vec{p}_2 - \vec{p}_1$, provides the local action to reach the target position. Blum and Abbott report that applying this navigation mechanism to solve the hidden-platform water maze (in particular, its reference memory version (RMW)), produces a latency decrease over trial consistent with experimental data [236].

Gerstner and Abbott [122] extend the above approach in order to achieve navigation *(i)* in the presence of obstacles, *(ii)* for multiple target locations. With respect to this latter issue, they introduce a target-dependent modulatory mechanism influencing place cell activity. After learning (involving multiple reward locations), this mechanism allows place cells to provide multiple navigational maps corresponding to different targets. Furthermore, interpolating maps to learned targets provides directions to novel unexperienced goals.

The model by Abbott, Blum, and Gerstner predicts that individual place fields should shift as a result of learning. As mentioned in Chapter 6, recent experimental data suggest that an experience-dependent place field shifting does actually occur when an animal takes several times a route [215, 216].

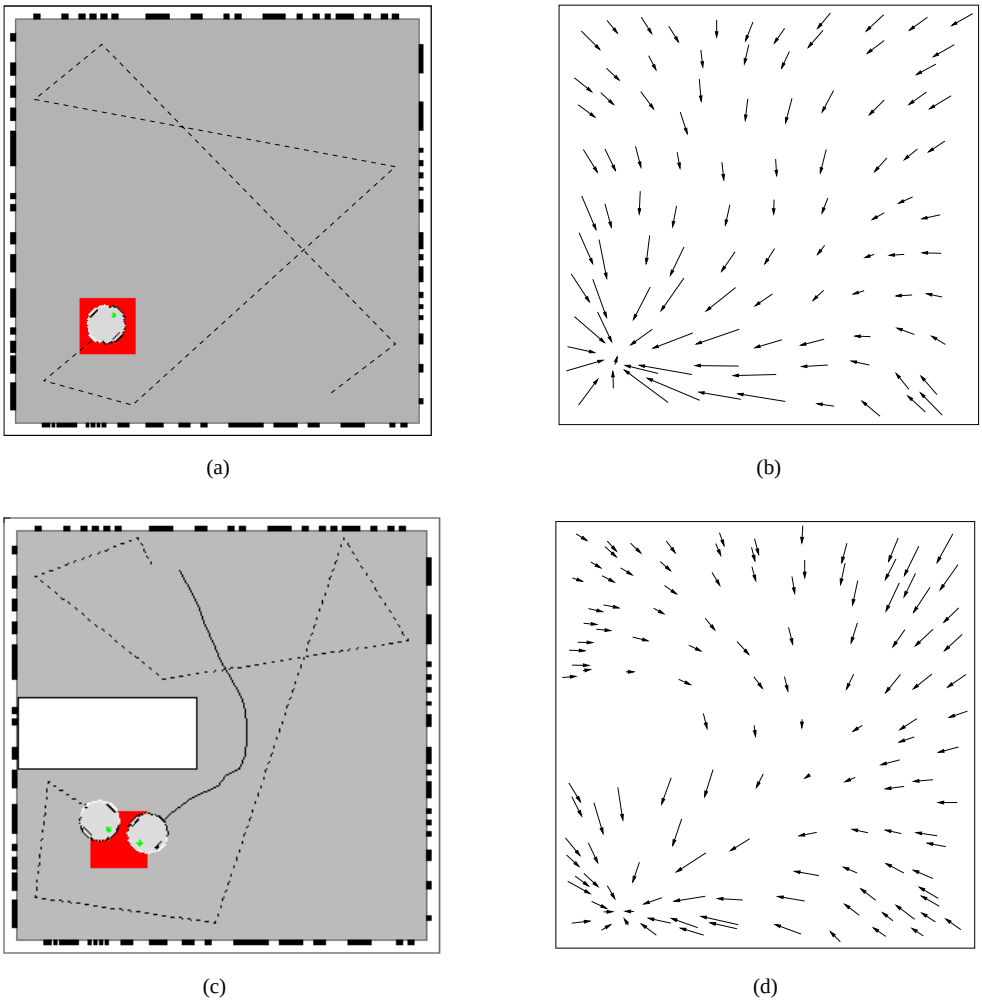

Figure 12.2: Some results obtained by implementing the navigation model by Abbott, Blum, and Gerstner [1, 46, 122]. Our incrementally learned place field representation has been used as basis to support the proposed action learning mechanism. The navigational map shown in (b) has been learned after 5 training trials, whereas the map in (d) has been acquired after 10 trials.

A limitation of the above theory is that place cells need to be tied to physical locations within a metric coordinate system (as in Wan, Redish, and Touretzky [377, 286]). This is indeed necessary to derive the local directions of movement encoded by the shift of the center of mass of the ensemble hippocampal activity. A second assumption of the model is that hippocampal place cells need to exhibit a dual space code. Indeed, they have to simultaneously represent both the true animal's location and the shifted position in order to extract the vector difference providing the navigational map.

The theory by Abbott, Blum, and Gerstner does not account for place field establishment based on sensory information. Rather, it focuses on how LTP can affect the activity of an ideal population of place cells consisting of evenly distributed Gaussian-tuned units. In an early stage of this thesis work, we have employed our incrementally learned place field substrate (Chapter 10) as a basis to implement the above navigation model [13]. This completes the system in the sense that the underlying space code supporting action learning is not manually built, but it is learned by modeling spatio-temporal features of the environment. Fig. 12.2 shows two examples of navigational maps obtained by computer simulation. During training, the simulated agent (lightgrey circle) moves along straight exploratory paths avoiding collisions with walls and obstacles (white rectangle). Every exploratory path starts at a random position

and with a random orientation and ends when the target (darkgrey square) is reached. Dashed lines shown in Fig. 12.2 are examples of training paths. The two vector field representations, Figs. 12.2 (b, d), show the two navigational maps generated after 5 and 10 training trials, respectively. Arrows indicate local directions as suggested by the shifted population vector for each sampled position[7]. Finally, solid trajectories are examples of paths taken by the agent according the encoded navigation map.

12.6 Trullier and Meyer (1997)*

As described in Sec. 7.1.5, Trullier and Meyer [370, 368, 369] consider the rat hippocampus (in particular, its CA3 region) as an hetero-associative memory in the spatial domain: place cells form a recurrent network that learns the spatial relationships between neighboring locations. Thus, CA3 provides a directed graph in which nodes are places and links indicate adjacency between places. In addition, similar to Schölkopf and Mallot [308], links are labeled by head-direction information. This results in a topological representation suitable to support navigation.

In the model, goal-oriented behavior is achieved by combining reward-signal propagation with the existence of goal cells. Whenever the agent finds a target, it recruits a set of goal cells and it looks in all directions [62]. For each orientation θ_i, the system induces a firing activity in all CA3 cells encoding locations in that gaze direction (by exploiting the fact that links in the CA3 graph are gated by head-direction information), regardless of the distance from the goal. Then, Hebbian learning is applied to correlate active CA3 place cells to the goal cell corresponding to θ_i. After learning, activating one of these CA3 cells will result in exciting the appropriate goal cell which will indicate where the animat is relative to the goal.

However, the above propagation mechanism does not bypass obstacles. Thus, in the presence of barriers, the activity of a set of goal cells does not cover the entire environment. Let S_g denote the surface of the environment in which goal cells can guide the agent towards the target. In order to overcome this problem, Trullier and Meyer postulate the existence of subgoal cells: neurons that allow the animat to navigate from places $\vec{p}$ that are not covered by goal cell activity, i.e., $\vec{p} \notin S_g$, to places $\vec{p}' \in S_g$. Whenever the agent is at a location $\vec{p} \notin S_g$, it starts wandering until when it finds a place $\vec{p}' \in S_g$. At that time, it recruits a new set of subgoals and it looks in all directions to update the connections from CA3 to this new set of subgoal cells. After learning, whenever the animat will be in a position $\vec{p} \notin S_g$, an appropriate subgoal cell will be activated in order to guide it towards a place $\vec{p}' \in S_g$. Thus, goal-oriented information is stored which allows the agent to skirt around obstacles (if the environment contains several barriers, the agent has to recruit multiple sets of subgoals).

Trullier and Meyer report computer simulations showing that the model allows an animat to perform goal-oriented navigation in presence of obstacles. However, as pointed out by the authors, the system does not cope with dynamically changing environments (e.g., moving, appearing, or disappearing obstacles).

12.7 Gaussier *et al.* (1998)*

In the model by Gaussier and colleagues [117] cognitive navigation is seen as a temporal sequence process (similar to Schölkopf and Mallot [308], and Trullier and Meyer [370]). A goal map layer is considered one synapse downstream of the place recognition level (the latter has been described in Sec. 7.1.6). The goal map provides a topological representation of the environment in which places (nodes) are interconnected according to their experienced temporal proximity. In particular, the weight w_{ij} of the link

[7] The vector field representations have been obtained by rastering randomly over the obstacle-free areas of the environment.

between two successively visited places j and i is increased by Hebbian learning. Thus, $w_{ij} = 0$ when there is no path from j to i, whereas $0 < w_{ij} \leq 1$ when i is directly reachable from j. Below is a scheme illustrating the architecture of the system:

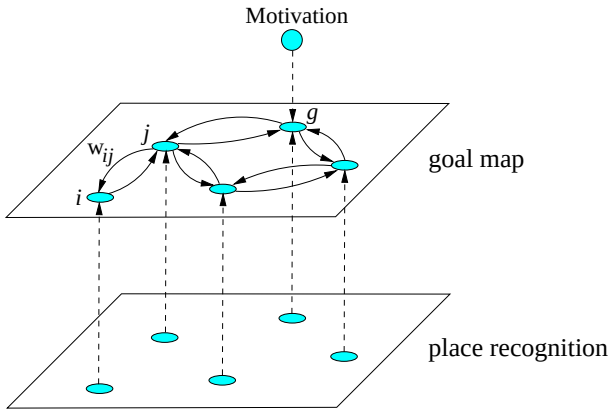

Whenever a motivation activates one unit g of the goal map, it is propagated to the other neurons through the synaptic weights w_{ij}. If r_i, with $0 \leq r_i \leq 1$, denotes the activation of a cell i in the goal map, the motivation-propagation algorithm is: *(i)* $r_g = 1$, *(ii)* $r_i = \max_j(w_{ij} \cdot r_j), \forall i$. As a consequence, a motivation has the effect of activating neurons in the goal map according to their distance to the goal neuron g. Navigation can then be achieved by following the gradient of the neural activity over the goal map. Note that the algorithm allows the agent to find the shortest path in the graph.

The authors present computer simulation results in which an agent has to find food and water sources within an environment with obstacles. Although they claim that the above method can cope with moving goals and dynamic environments, they do not report any result concerning those experimental situations.

12.8 Foster, Morris, and Dayan (2000)

This model postulates an actor-critic temporal difference (TD) learning mechanism to achieve reward-based hippocampally dependent navigation [107, 83]. The system consists of three parts: *(i)* An input layer of place cells. The authors assume an ideal population of Gaussian units i encoding the agent's current position $\vec{p}(t)$. *(ii)* A critic unit c one synapse downstream of hippocampal place cells. The critic receives afferent projections w_{ci} from all place cells i. *(iii)* An actor network consisting of eight action cells a_j one synapse downstream of place cells. Each action unit a_j, $1 \leq j \leq 8$, codes for a specific direction of movement (e.g., north, north-west, and so on) and receives input connections w_{ji} from all place units i.

Given this architecture, TD-learning is applied to update weights w_{ci} and w_{ji} from place cells i to the critic c and to the actors a_j, respectively. During reward-based learning, the actor continually produces actions taking the agent around the environment, while the critic continually criticizes the selected actions. This results in an actor adapting its policy according to the critic's information, and in a critic that also adapts to the changing actor[8] [343].

The authors focus on experimental findings concerning rodent navigation in the hidden-platform water maze. In particular, they run computer simulations to solve both the fixed-platform task (i.e., the reference-memory water maze (RMW)), and the moving-platform task (i.e., the delayed matching-to-place water maze (DMP)[9]) [337]. Results show that the above actor-critic model allows the simulated

[8]The critic network has to learn a value function $V(\vec{p})$, over locations $\vec{p}$, which is used in TD-learning to evaluate the action-selection policy currently adopted by the actor [343].

[9]See Chapter 11.

animal to learn the RMW task in about 10 trials (consistently with rodent navigation data). On the other hand, the model fails to reproduce the performance of rats in the DMP task: the system does not adapt to moving platforms over trial because of the interference from previously learned policies. Rats, instead, can avoid interference from previous training as well as generalize from experience on previous days to improve performance on later training [337] (Chapter 11).

In order to overcome this limitation, a further component is added to the model which learns allocentric coordinates (i.e., independent of the animal's entry position). The coordinate system relies on the hippocampal place cell representation and consists of three elements: *(i)* A coordinate representation (x_p, y_p) of the agent's current location $\vec{p}$. This is implemented by two units x and y one synapse downstream of hippocampal place cells. Let w_{xi} and w_{yi} be the projection weights from all place cells i to units x and y, respectively. *(ii)* A goal coordinate memory (x_g, y_g) providing the location at which the platform was lastly found. *(iii)* A mechanism that performs vector subtraction algorithmically to derive the direction θ_g from the current location (x_p, y_p) to the target (x_g, y_g).

As the agent moves, TD-learning is applied to update connections w_{xi} and w_{yi} from place cells to coordinate cells, that is, to acquire the allocentric coordinate system based on relative self-motion. For every move, the agent has to estimate the difference between its coordinates at the end and at the begin of the move. This difference should correspond exactly to the relative self-motion during the move. This consistency condition is used as the basis for TD-learning. The authors report that learning a coordinate system takes about 16 trials.

The output of the coordinate system is the direction θ_g to be followed to navigate to the goal. This information is integrated within the actor-critic system by considering an abstract action cell a_{coord}. This unit receives reward by the critic depending on its performance in a similar manner to actor cells a_j. This results in a coordinate system which takes the control gradually as it gives increasingly accurate information about where both the animal and the goal are located. Below is a scheme of the model:

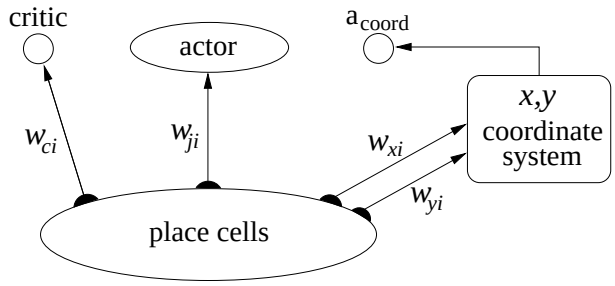

The combined coordinate and actor-critic model has been tested in both the fixed-platform water maze (RMW) and the moving-platform task (DMP). In both cases, simulation results are consistent with experimental data [337].

An intrinsic limitation of the above model is that it relies upon an ideal place cell representation in the sense that *(i)* it is not built from sensory inputs, *(ii)* it is assumed to be coherent over time. In particular, using place cells to learn metric coordinates assumes a completely stable allocentric place code. Also, as reported by the authors, head-direction information is assumed to be consistent over trial (e.g., the simulated agent never undergoes disorientation between training trials). Finally, similar to Wan, Redish, and Touretzky [377, 286], the coordinate model relies upon an explicit memory of the current goal coordinates (x_g, y_g) in order to derive the bearing information by means of vector subtraction.

Chapter 13

Modeling Hippocampus-based Navigation: Action Learning in Continuous Space

The proposed hippocampal model enables the agent to localize itself within an unfamiliar environment based on available sensory information (i.e., spatial learning). However, to accomplish its functional role in spatial behavior (i.e., cognitive navigation), the model must also incorporate the knowledge about relationships between the environment, its obstacles and reward locations. That is, place cell activity has to be used to guide navigation.

Fig. 13.1 (a) shows a functional view of the entire system. Spatial learning is achieved by acquiring place and direction coding incrementally. Both of these modules rely on allothetic as well as idiothetic stimuli, and interact with each other to maintain consistent representations over time. One synapse downstream of hippocampal place cells, we have the action learning module. We focus on a specific neural pathway, namely the fornix projection, connecting the hippocampus to the nucleus accumbens (NA). As previously mentioned, the latter is an extra-hippocampal structure involved in goal-oriented locomotor behavior (Chapter 11). We do not propose a model for the nucleus accumbens [55]. Rather, we simply assume that CA1 place cells project efferents to a population NA $= \{a_i \,|\, 1 \leq i \leq N\}$ of *locomotor action neurons* a_i whose anatomical locus is the nucleus accumbens. Each cell a_i provides an allocentric directional motor command (e.g., go north).

Given the above system, the navigation problem is: How can we establish a mapping function $\mathcal{M} : \mathcal{R} \rightarrow \mathcal{A}$ from the place cell activity space $\mathcal{R}$ to the action space $\mathcal{A}$ in order to achieve goal-directed behavior? We employ reinforcement learning [343] to acquire $\mathcal{M}$ through continuous agent-environment interaction [62, 55, 107]. Synaptic efficacy between CA1 and NA cells is changed as a function of reward-related signals. This results in an ensemble action cell activity providing, at each time step, the correct action to navigate to the target.

Within the reinforcement learning paradigm, temporal difference (TD) learning relies on sound mathematical foundation and represents a well understood approach [342, 84]. In particular, we utilize Q-learning [378], a TD-based learning technique.

Problems with high-dimensional *continuous* state space are a critical issue in reinforcement learning [303]. In these cases, optimizing the *value function* (i.e., deriving the function predicting the optimal long-term reward for a given state when selecting a specific action and following an optimal policy

thenceforth) means to learn the optimal mapping over a very large number of state-action pairs. Since exploring the whole state-action space is infeasible, a key property in these problems is the ability of the agent to estimate the expected value function for never experienced state-action pairs (i.e., generalization property) [303, 343, 187, 297].

Combining reinforcement learning with function approximation methods permits to learn an action-value function over a continuous location space, which endows the system with spatial generalization capabilities. Thus, the animat may be able to associate appropriate actions to spatial positions that it has never seen before. Our overlapping localized CA3-CA1 place fields provide a natural set of basis functions that can be used to learn a parameterized form of the action-value Q-function. Note that, in contrast to other approaches, we do not have to choose parameters like width and location of the basis functions. Rather, the basis functions are created automatically during the phase of unsupervised learning.

13.1 Learning Navigational Maps: Reward-based Action Selection in Continuous Space

Fig. 13.1 (b) shows our action learning system. For each type of target (e.g., food, water), a discrete set of four NA action cells $\mathcal{A} = \{north, south, west, east\}$ is recruited. Each cell $a \in \mathcal{A}$ receives afferent projections $\vec{w}^a = (w_1^a, \ldots, w_n^a)$ from all n CA1 place cells. Therefore, the navigation task consists of modifying the four synaptic vectors $\vec{w}^a$ to learn the location-to-action mapping from each agent's position $\vec{p}(t)$ (i.e., *state*) to the correct local action $a(t)$. This results in learning a set of four Q-value functions $Q_w(\vec{p}, a)$.

Let r_i and r_a denote the activity of place cells i and of action cells a, respectively. Each state $\vec{p}(t)$ is encoded by the ensemble place cell activity $\vec{r}(\vec{p}) = (r_1(\vec{p}), r_2(\vec{p}), \ldots, r_n(\vec{p}))$, where n is the number of place cells. Thus, the population of place cells can be thought of as a continuous-valued coarse coding representation. The activity r_a of an action cell a depends linearly on the place cell activity $\vec{r}(\vec{p})$ and on the synaptic weights $\vec{w}^a$:

$$r_a(\vec{p}) = Q_w(\vec{p}, a) = (\vec{w}^a)^T \, \vec{r}(\vec{p}) = \sum_{i=1}^{n} w_i^a \, r_i(\vec{p}) \tag{13.1}$$

In order to update $\vec{w}^a$ to approximate the optimal goal-oriented action-value function $Q_w^*(\vec{p}, a)$, we employ the linear gradient-descent version of Watkins' Q-learning algorithm [343]. Given a position $\vec{p}$, we interpret the action cell activity $r_a(\vec{p})$ as the "expected gain" when taking action a at location $\vec{p}$ of the environment.

During training, the robot behaves in order to either consolidate goal-directed paths (*exploitation*) or find novel routes (*exploration*). This exploitation-exploration trade-off is determined by an ϵ-greedy action selection policy, with $0 \leq \epsilon \leq 1$ [343]. At each time t, the robot takes the "optimal" action a_t^* with probability $1 - \epsilon$ (exploitation)

$$a_t^* = \text{argmax}_a \, r_a(\vec{p}_t) \tag{13.2}$$

or, it might resort to uniform random action selection with probability equal to ϵ (exploration). At each step Δt, the synaptic efficacy of projections $\vec{w}^a$ changes according to [343]

$$\Delta \vec{w}^a = \alpha \, \delta_t \, \vec{e}_t \tag{13.3}$$

The terms in Eq. 13.3 have the following interpretation:

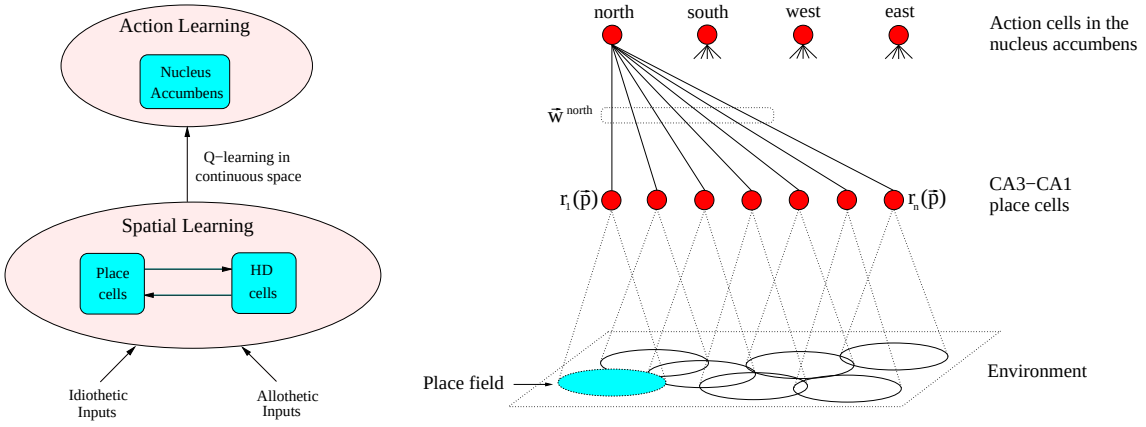

Figure 13.1: (a) A functional view of the entire system. The space learning module presented in the previous chapters provides a neural basis for action learning. We focus on the connection from CA1 cells to neurons in the nucleus accumbens (NA) in order to establish navigation maps. (b) Our Q-learning scheme. The overlapping localized CA3-CA1 place fields are used as basis functions. A discrete set of actions $\mathcal{A} = \{north, south, west, east\}$ is considered, which results in four Q-functions $Q(\vec{p}, north), Q(\vec{p}, south), Q(\vec{p}, west), Q(\vec{p}, east)$ to be approximated.

(i) The factor α, $0 \leq \alpha \leq 1$, is a constant learning rate.

(ii) The term δ_t is the *prediction error* defined by

$$\delta_t = R_{t+1} + \gamma \max_a r_a(\vec{p}_{t+1}) - r_a(\vec{p}_t) \tag{13.4}$$

where R_{t+1} is the actual reward delivered by an internal brain signal, and γ, $0 \leq \gamma \leq 1$, is a constant discount factor. The temporal difference δ_t estimates the error between the expected and the actual reward when, given the location $\vec{p}_t$ at time t, the agent takes action a and reaches location $\vec{p}_{t+1}$ at time $t+1$. Training trials allow the system to minimize this error signal. Thus, asymptotically $\delta_t \approx 0$, which means that, given a state-action pair, the deviation between predicted and actual reward tends to zero.

Neurophysiological data show that the activity of dopamine neurons in mammalian midbrain encodes the difference between expected and actual occurrence of reward stimuli [312]. In particular, the more reliably a reward is predicted, the closer is the rate of a dopamine neuron to its baseline activity. Thus, the temporal difference error δ_t used to update our synaptic weights $\vec{w}^a$ may be thought of as a dopamine-like teaching signal (Sec. 11.3).

(iii) During training paths, Eq. 13.3 allows the robot to memorize action sequences. Since recently taken actions are more relevant than earlier ones, we need a memory trace mechanism to weight actions as a function of their occurrence time. The vector $\vec{e}_t$, called *eligibility trace*, provides such a mechanism [343]. The update of the eligibility trace depends on whether the robot selects an exploratory or an exploiting action. Specifically, the vector $\vec{e}_t$ changes according to

$$\vec{e}_t = \vec{r}(\vec{p}_t) + \begin{cases} \gamma\lambda\vec{e}_{t-1} & \text{if exploiting} \\ 0 & \text{if exploring} \end{cases} \tag{13.5}$$

where λ, $0 \leq \lambda \leq 1$, is a trace-decay parameter [343], and $\vec{r}(\vec{p}_t)$ is the place cell vector activity. We start with $\vec{e}_0 = \vec{0}$.

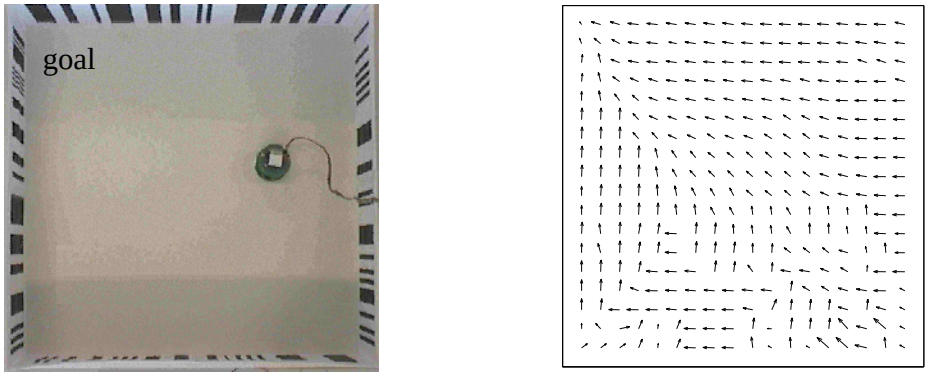

(a) (b)

Figure 13.2: (a) A view of the environment with a target location. (b) Vector field representation of the learned navigational map after 5 training trials.

13.2 Behavioral Experiments in Open-field Environments

Given our experimental setup, we define a specific target region (e.g., a feeding location) within the environment. We apply the above reward-based learning scheme to build up a navigational strategy leading the robot towards the target from any location, while avoiding obstacles.

To study robot behavior, we adopt a similar protocol as employed for experiments with rats. Training involves a series of consecutive sessions each of which consists of a sequence of trials (e.g., five) beginning at a random location and ending when the robot reaches the target or after a timeout (e.g., 120 time steps). At the beginning of each trial the robot retrieves its starting location on the hippocampal chart based on the allothetic (visually-driven) place field representation (Chapter 10) [207, 205].

During learning, we take a discrete set of actions $\mathcal{A} = \{north, south, west, east\}$. However, after learning, population vector coding is applied to map $\mathcal{A}$ into a continuous action space $\mathcal{A}'$ by averaging the ensemble action cell activity. Given a position $\vec{p}$ of the robot, the action

$$a'(\vec{p}) \propto \begin{pmatrix} \cos\phi \\ \sin\phi \end{pmatrix} \tag{13.6}$$

is a direction in the environment encoded by the NA action cell activity

$$a'(\vec{p}) = \frac{\sum_{a\in\mathcal{A}} a\, r_a(\vec{p})}{\sum_{a\in\mathcal{A}} r_a(\vec{p})} \tag{13.7}$$

where $a_n = \binom{0}{1}$, $a_s = \binom{0}{-1}$, $a_w = \binom{-1}{0}$, and $a_e = \binom{1}{0}$ are the four principal directions. Eq. 13.7 results in smooth trajectories.

Experiments have been carried out with a learning rate $\alpha = 0.1$, a discount factor $\gamma = 1.0$, and a decay factor $\lambda = 0.9$. The reward-signal function $R(\vec{p})$ is defined by

$$R(\vec{p}) = \begin{cases} 1 & \text{if } \vec{p} = \text{target state} \\ -0.5 & \text{if } \vec{p} = \text{collision state} \\ 0 & \text{otherwise} \end{cases} \tag{13.8}$$

where collision means contact with walls or obstacles.

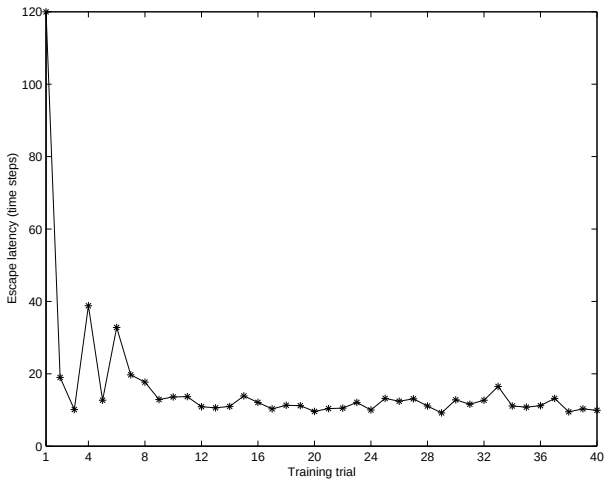

Figure 13.3: Escape latency, in number of time steps, as a function of training trials.

We adopt a dynamically changing ϵ-probability. The idea is to increase the probability of exploring novel routes as the time to reach the target increases. The ϵ parameter is defined by the exponential function

$$\epsilon(t) = \frac{\exp(\beta\, t) + k_1}{k_2} \tag{13.9}$$

where $\beta = 0.068$, $k_1 = 100$, $k_2 = 1000$, and where $t = 0, 1, 2, \ldots$ are discrete time steps. If we consider the dynamic of ϵ over a time window of 100 time steps, at $t = 0$ the robot behaves according to a value $\epsilon = 0.101$ (i.e., enhancing exploitation), and at $t = 100$ it behaves according to a value $\epsilon = 1.0$ (i.e., enhancing exploration). If at the end of the time, $t = 100$, the target is not reached yet, exploration is further enhanced by keeping a fixed $\epsilon = 1.0$ for another 100 time steps. Then, exploitation is resumed by setting $t = 0$ and $\epsilon = 0.101$. Moreover, every time the target is reached the time window is re-initialized as well, and ϵ is set equal to 0.101. These are heuristic methods to ensure a sufficient amount of exploration.

13.2.1 Single Target Experiments

Fig. 13.2 (a) is a view of the experimental arena as seen by the camera above the environment[1]. The target area is about 2.5 times the area occupied by the robot. During training, the target is maintained at a fixed position. Therefore, this experiment is similar to the reference memory task (RMW) performed with rats in the water maze (Sec. 11.1). Fig. 13.2 (b) shows the navigational map learned by the robot after 5 training trials. The vector field representation has been obtained by rastering uniformly over the whole environment. Each arrow indicates the local action encoded by the ensemble action cell activity after learning (according to Eq. 13.7). Many of sampled locations were not visited by the robot during training, which confirms the generalization capabilities of the method. That is, the robot was able to associate appropriate goal-directed actions to never experienced spatial states. Finally, Fig. 13.3 shows the mean escape latency (measured in time steps) as a function of training trials. Our model captures the learning capabilities of rats when solving the RMW task: escape latencies decrease over time and reach asymptotically low values after few training trials (similar to Fig. 11.1 (a)).

[1]Results concerning goal-oriented navigation have been mainly obtained with the experimental setup involving the linear vision system (Chapter 8, Sec. 8.1.1).

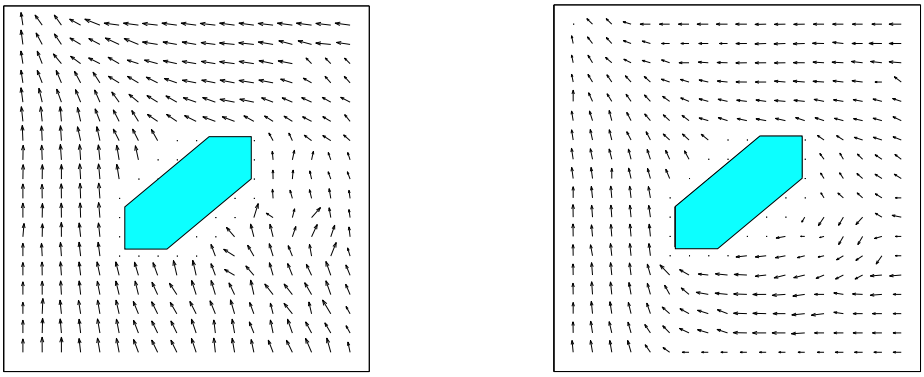

(a) (b)

Figure 13.4: Navigation tasks in the presence of obstacles (grey objects). (a) Vector field representation after 20 training trials. (b) The same experiment as (a), but the map is the result of 80 training trials.

In a further series of experiments, obstacles have been added within the arena. Figs. 13.4 and 13.5 show that our navigation model enables the robot to cope with the presence of barriers (grey objects). The vector field representation of Fig. 13.4 (a) has been acquired after 20 training paths. It contains appropriate goal memory information, whereas it does not provide obstacle avoidance accurately. Note that, this does not really impair the robot's target-directed behavior, since obstacle avoidance is supported by a low-level reactive module driven by infrared sensors. Fig. 13.4 (b) displays a navigational map learned after 80 trials for the same task. Due to longer training, the map provides both goal-oriented and obstacle avoidance behavior. Fig. 13.5 (a) shows a map learned after 20 training paths in an experiment in which the target location is located at the bottom-left corner of the arena. Finally, the map in Fig. 13.5 (b) concerns an experiment with two obstacles and the goal at the same bottom-left location. The map has been learned after 20 training trials.

13.2.2 Moving a Learned Target

As discussed in Sec. 11.1, if rats are trained to solve the reference memory task (RMW) for several days and the escape platform is subsequentially relocated, they need some re-training to adapt their navigation behavior. Nevertheless, they exhibit direct navigation paths to the new target after a few re-training trials (Fig. 11.1).

This experiment concerns relocating a previously learned target (in particular, the one used for the experiment of Fig. 13.5 (b)), and enabling the robot to adapt its navigational behavior consequently. Fig. 13.6 (a) shows an example of robot's trajectory after it has learned the target at the bottom-left location. The animat is attracted by the previous feeder position and concentrates its search in the neighboring positions. Therefore, some mechanism is needed to allow the robot to find the new target in the upper-right corner of the arena. Note that, due to the exploration component of the Q-learning algorithm (Eq. 13.9), the animat might occasionally find the new goal and then adapt its action-selection policy over time. However, it would take a large number of re-training trials in order to achieve direct navigation towards a new target placed in an opposite position with respect to the learned one.

The idea is to endow the system with an internal *reward-expectation mechanism*. During training, the robot learns to correlate the place cell activity to the positive reward signal, $R = 1$, received at the target location. This is done by taking a neuron d, that we call the *reward-expectation cell*, one synapse downstream from the place cell layer (Fig. 13.7 (a)). Let i be an index over the CA3-CA1 cell population.

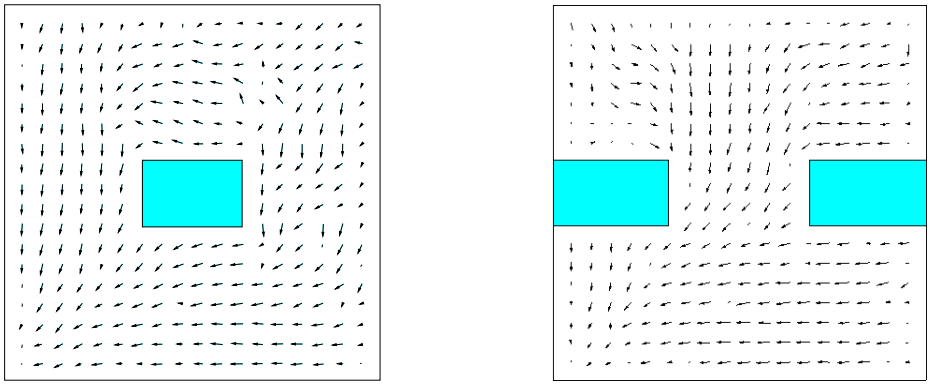

(a) (b)

Figure 13.5: Navigation tasks in the presence of obstacles (grey objects). (a) The goal is at the bottom-left corner. The map has been learned after 20 training trials. (b) The same learning task as (a) but with two obstacles. The map has been learned after 20 trials.

Projections w_{di} from place cells i to cell d are inhibitory, and are initialized to random values within the interval $[-0.1, 0]$. Cell d receives as input also the external rewarding stimulus R. The activity r_d of cell d is non linear and it is defined by

$$r_d = \begin{cases} f(\sum_i w_{di} r_i) + R & \text{if } R \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{13.10}$$

where $f(x) = \tanh(x)$. Thus, the activity of cell d depends on both the external reward R and the CA3-CA1 network activity.

In order to learn the desired correlation between the event "positive reward" and the place cell activity, we apply Hebbian learning to modify the inhibitory weights w_{di}:

$$\Delta w_{di} = r_i \cdot r_d (w_{di} - 1) \tag{13.11}$$

The more correlated the activities r_i and r_d, the more inhibitory the synapse w_{di}.

As a consequence, before associating the external reward signal with the internal spatial representation, the cell d responds maximally whenever the robot receives a positive $R = 1$. Indeed, since weights w_{di} are initially close to zero, the activity $r_d \approx R = 1$ (according to Eq. 13.10). As training proceeds, the robot starts predicting the external stimulus R by learning synapses w_{di}. Thus, whenever the robot is near the target location, cell d receives a strong inhibitory input $\sum_i w_{di} r_i$ which compensates for the excitatory reward R. Therefore, when R is fully predicted, even when a reward $R = 1$ occurs, cell d remains silent. On the other hand, if the fully predicted reward fails to occur (i.e., the learned target has been moved away), the activity of cell d is strongly depressed ($r_d \approx -1$), and an internal negative reward is generated. When the number of negative internal rewards exceeds a fixed threshold D (e.g., $D = 10$), the robot "gives up" searching the previous target location and starts looking for a new goal. Thus, after a re-training period, it will start navigate directly to the relocated target. Fig. 13.6 (b) presents the navigational map obtained after 20 re-training trials. Fig. 13.7 (b) shows the escape latency during training before target relocation (until trial 40), and after (trials $41 - 60$). This result is consistent with the capability of rats of quickly re-adapting their behavior relative to an escape platform which has been relocated after several days of training.

The reward-expectation cell d finds its neuro-physiological counterpart in dopamine neurons observed in mammalian midbrain. The response of these neurons encodes the unpredictability of incoming

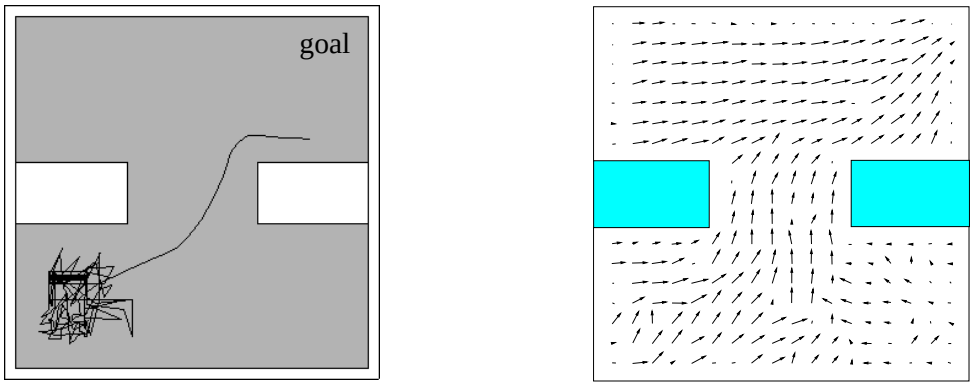

(a) (b)

Figure 13.6: (a) The robot uses the navigation map of Fig. 13.5 (b) to navigate to the previously learned target. However, the latter has been moved to a new location (the solid line is the robot's trajectory). (b) The re-adapted navigational map corresponding to the new rewarding location.

stimuli [312]. Indeed, they respond positively to external rewards which occur unpredictably, but they remain silent if a fully predicted stimulus arrives. Moreover, when a fully-expected reward fails to occur, dopamine neurons respond negatively exactly at the time at which the reward is expected [312] (Sec. 11.3)[2].

13.2.3 Multiple Target Experiment

The learning scheme of Sec. 13.1 also applies to multi-target tasks. Let $\{T_1, \ldots, T_m\}$ be a set of distinct target types (e.g., T_1 could be a food location, T_2 a water location, and so on). Whenever the agent finds a target T_i, it recruits a new set of action cells $\mathcal{A}^{T_i}$. Let $r_a^{T_i}$ be the activity of cell $a \in$ NA associated to target T_i, $\vec{w}^{a,T_i}$ denote the synaptic projections from CA1 place cells to cell a, and $\{R^{T_1}, \ldots, R^{T_m}\}$ be the set of reward signals. We employ the above Q-learning algorithm to approximate the set of functions $r_a^{T_i}(\vec{p})$.

In this experiment we consider two distinct types of rewards T_1 (e.g., food) and T_2 (e.g., water). Fig. 13.8 (a) shows the two target locations within the environment. The agent is initially trained to reach the feeder location T_1. That is, its primary task is to approximate the $r_a^{T_1}(\vec{p})$ functions. Fig. 13.8 (b) illustrates the navigation map learned by the animat after about 30 trials.

When searching for food, it might happen that the agent encounters the water location and receives a positive reward signal with respect to T_2, $R^{T_2} = 1$. This information can be exploited by the robot by adjusting the weights $\vec{w}^{T_2}$. Thus, even if T_2 is not the current target, the robot can partially learn a navigational map leading to it. Fig. 13.8 (c) shows the knowledge about the water location T_2 acquired by the robot while learning the optimal policy to reach the food T_1. Therefore, when the robot decides to focus on water (i.e., to approximate $r_a^{T_2}(\vec{p})$), it does not start from zero knowledge. This results in a shorter training time for T_2, and accelerates the robot's progress. Fig. 13.8 (d) presents the navigational map learned by the robot after about 10 training trials when looking for water.

[2]Note that, instead of Eq. 13.10, we could have also used the prediction error δ_t defined in Eq. 13.4 to monitor an unexpected target location.

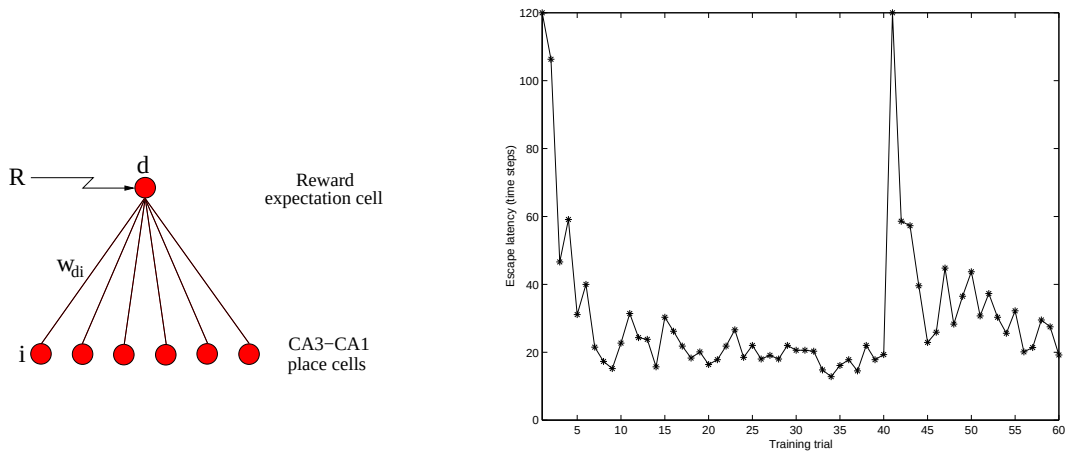

(a) (b)

Figure 13.7: (a) The internal *reward-expectation mechanism*. Cell d has a dopaminergic-like activity and sends an internal negative reward whenever a fully predicted reward fails to occur. (b) Escape latency during training (until trial 40) and after relocating the target (trials 41-60).

13.3 Discussion

Action learning, in our navigation model, relies on the reward-dependent update of hippocampal efferents to the nucleus accumbens (NA) [55]. In the system, CA1 cells project to a population of locomotor action cells in NA whose ensemble activity is used to guide navigation. Thus, solving the action learning task results in establishing a mapping function from the continuous space of physical locations to the activity space of action cells. We employ Q-learning [378] (a temporal-difference learning technique [342, 84, 343]) to establish this function based on the agent's experience. The robot interacts with the environment, and reward-related stimuli elicit the synaptic changes of the CA1 $\rightarrow$ NA projections to adapt the action-selection policy to the task. After training, the hippocampally-dependent NA activity provides a map to support goal-oriented behavior and obstacle avoidance.

During recent years, reinforcement learning has become a mathematically well-defined learning paradigm with numerous applications [343, 84]. Algorithms like temporal difference learning TD [342], Q-learning [378], or SARSA [343] are easy to implement, in particular for low-dimensional discrete problems. In real world applications, input data (e.g., sensory data for an autonomous robot) are rather high-dimensional and continuous. The most important practical issue for applications of reinforcement learning in these cases is probably the construction of a suitable representation of the input space. We have shown that our incrementally learned place fields provide a coarse coding representation suitable for applying reinforcement learning in continuous space. Our representation also solves the problem of partially hidden states [202]. The standard reinforcement paradigm relies on the Markov hypothesis that the present state of the system is known to the learner [343]. However, in real-world applications the Markov hypothesis is not always fulfilled (e.g., when the current state is identified based on visual input only). Since our hippocampal place fields integrate allothetic and idiothetic information, sensory aliasing does not occur, and therefore the current state is always fully known to the system. Another limitation of standard reinforcement techniques is that learning takes a long training time when applied directly on high-dimensional input spaces [343]. We have shown that by means of an appropriate state space representation, based on localized overlapping place fields, the robot can learn goal-oriented behavior after few training trials (e.g., 5 as shown in Fig. 13.2).

Recently, Schultz and colleagues [311, 312] have demonstrated that TD reinforcement learning has

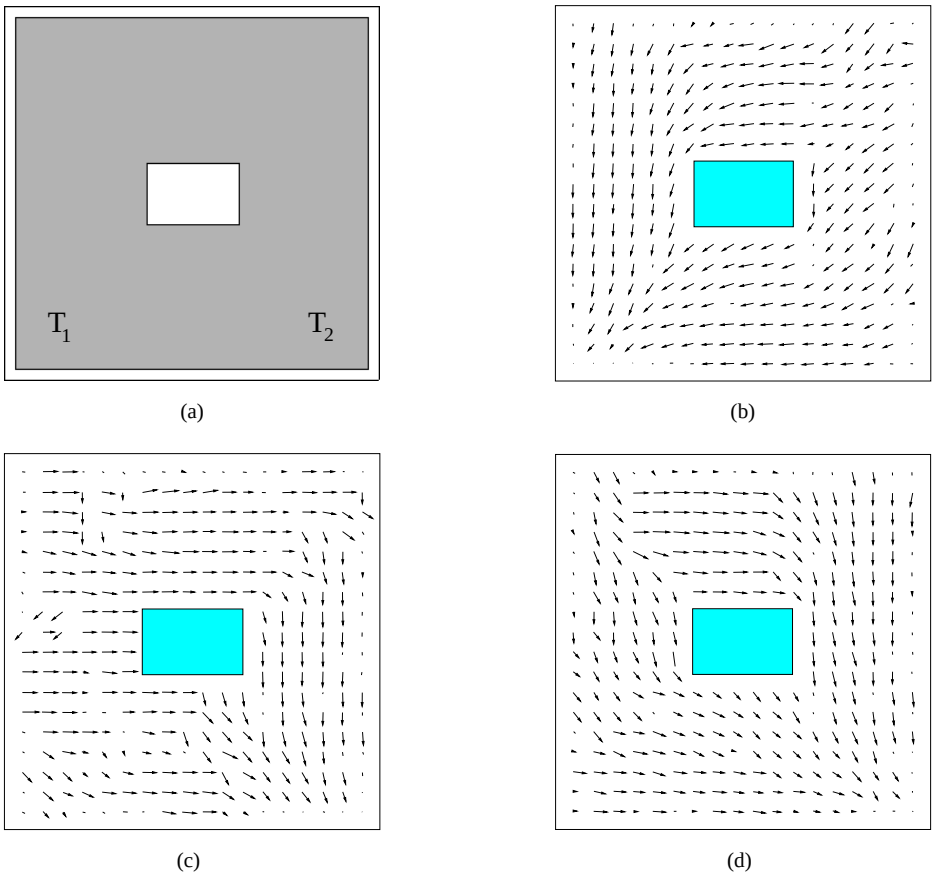

Figure 13.8: (a) The arena with two distinct target types T_1 (e.g., food) and T_2 (e.g., water). The white rectangle is an obstacle. (b) The navigation map corresponding to the food location T_1. (c) The partial navigation map corresponding to the water T_2 learned by the robot when focusing on food T_1. (d) The final map acquired by the robot when focusing on water T_2.

also a sound neurophysiological foundation: midbrain dopamine neurons fire as a function of the unpredictability of a reward and, then, they might provide a prediction error signal similar to the teaching signal used in TD learning. Dopaminergic inputs reach the ventral striatum (including the nucleus accumbens) via the ventral tegmental area (Chapter 11).

The navigation model has been validated by a series of experiments which can been thought of as instances of the reference memory task (RMW) in the water maze. The system captures the ability of rats to acquire navigational map rapidly (less than 10 trials) and effectively (generalization property). In one of these experiments, the target location is changed after training. We postulate a mechanism based on a dopamine-like signal in order to let the robot "give up" using the learned map and start approximating a new action-selection policy for the new goal. This results in rapid adaptation of the goal-oriented behavior which is consistent with experimental data. The system has been also tested in a two-target learning task (e.g., food and water). A nice property is that even if the robot is focusing on a specific target T_1, it can exploit exploration in order to learn navigation towards target T_2. Thus, when the agent focuses on T_2, it does not start from zero and it has only to refine the partial map previously established. This results in fast learning of the optimal policy for T_2.

An important issue we have discussed in Part III is latent learning, i.e., the fact that the rodents establish a representation of the environment even in the absence of explicit rewards [361, 256]. In this chapter, we have shown that having a target-independent space representation (i.e., our CA3-CA1 place

fields) enables the animat to learn target-oriented navigation very quickly. That is, goal-directed behavior is supported by the stable activity in our place cell substrate. This is consistent with the concept of latent learning.

The model predicts the presence of locomotor action neurons in the nucleus accumbens (NA) driven by hippocampal place cells. This is consistent with lesion studies. Indeed, damaging either the fimbria-fornix (connecting the hippocampal formation to NA) or the nucleus accumbens impairs animals' performance to learn the hidden platform maze [93, 341, 385] (Chapter 6). The model also predicts that blocking synaptic plasticity in the hippocampus proper after a place field representation has been established should not impair action learning in a hidden platform navigation task [107]. Indeed, since NA locomotion units are extrahippocampal, as long as CA1 $\rightarrow$ NA projections are plastic goal-oriented behavior can be achieved.

The fact that we locate action selection outside the hippocampal formation is an important difference with respect to models by Abbott, Blum, and Gerstner [1, 46, 122], and by Trullier and Meyer [370, 368, 369]. Indeed, those approaches postulate that navigation information is stored in the recurrent connections between hippocampal place cells (e.g., CA3 collaterals). Also, in contrast to models by Schölkopf and Mallot [308], and Gaussier *et al.* [117], our navigation model does not rely on a topological graph supporting path planning. Our approach is similar to those by Burgess, Recce, and O'Keefe [62, 61], Brown and Sharp [55], and Foster, Morris, and Dayan [107, 83] in that we also postulate that goal-oriented behavior is achieved by mapping place cell activity into actions based on reward signals.

However, goal cells (GC) by Burgess, Recce, and O'Keefe [62, 61] do not depend on a *local* place-to-action function as our NA action cells do. This results in a GC activity that encodes the position of the animal relative to the goal, but that does not account for obstacle avoidance. Also, our action selection mechanism does not suffer from the distal reward problem, i.e., the fact that places that are far away from the target may not be associated to any specific goal-oriented action. Finally, Burgess, Recce, and O'Keefe do not propose any re-learning mechanism to cope with targets whose location might change over time (e.g., DMP water maze task).

With respect to the model by Brown and Sharp, we do not encode motion directions egocentrically (e.g., turn left), but allocentrically (e.g., locomote to north). Also, they use an explicit memory trace mechanism to overcome the distal reward problem, whereas our system relies on a TD learning method [343]. Another difference is that we do not propose an explicit model for the nucleus accumbens as Brown and Sharp do. Rather, we simply assume the presence of action neurons in that specific brain region. Finally, the model by Brown and Sharp is adequate for RMW-like problems, but not for DMP learning tasks [337].

Our approach is similar to that by Foster, Morris, and Dayan [107, 83] in that we also apply a hippocampally dependent temporal difference learning for navigation. However, they postulate an actor-critic architecture to implement reward-based synaptic modification, whereas in our system we implement a Q-learning scheme in continuous space. Also, their model relies on an ideal place field representation which is assumed to be stable over time. By contrast, we build our population of CA3-CA1 cells from real sensory information and we get stability by combining visual and self-motion cues. Finally, similar to Wan, Redish, and Touretzky [377, 286], Foster, Morris, and Dayan resort to an allocentric metric coordinate system in order to enable the system to solve complex navigation tasks (e.g., the DMP water maze). The coordinate system relies on the hippocampal place cell representation, has an explicit memory of the current goal coordinates (x_g, y_g), and derives the current direction to the target θ_g by performing vector subtraction algorithmically. In our model, path integration (encoding metric information) is independent from CA3-CA1 representation and is not used for action selection.

Part V

Conclusions and Appendices

Chapter 14

Conclusions

At the beginning of this dissertation we established two milestones for this thesis work: *(i)* Understanding and modeling the neural mechanisms underlying biological spatial learning; *(ii)* Designing an autonomous navigation system based on the information provided by a neural representation of the agent-environment interaction. It's time to ask the crucial question: Have we accomplished these two objectives? In this chapter, we summarize the achievements of this work and we give its overall contributions. We also analyze the limitations of the current model, and we state our long-term future direction of research.

14.1 Contributions

In this work we have developed a modular neural model that captures several neurophysiological findings concerning biological head-direction cells and hippocampal place cells. The model involves two principal neural substrates providing directional and place coding. These two modules are strongly coupled and interact with each other to form a unitary spatial learning system. Inhibiting either one of the two systems critically impairs the performance of the other and disrupts the spatial learning capability of the entire system. The model stresses the importance of combining idiothetic (e.g., vestibular and proprioceptive) and allothetic (e.g., visual) signals to generate and maintain stable direction and place representations. Indeed, the dynamics of both neural systems relies on vision as well as self-motion stimuli. The representation of direction and position by head-direction cells and place cells, respectively, provides a suitable basis for achieving goal-oriented navigation using reinforcement learning. Fig. 14.1 is a scheme of the overall system's architecture.

In order to validate the model experimentally, we have implemented it on a real mobile Khepera robot. The experimental setup consists of a square arena within which the robot can freely move. The robot's behavior is monitored by means of a camera above the arena. The robot's sensory system consists of eight infrared sensors to detect obstacles and measure ambient light, a light-detector, an on-board camera for vision-based self-localization, and an odometer for sensing internal self-motion signals. We stress the importance of continuous interaction between the agent and the environment. This results in an incremental and dynamic development of the navigation system, and enables the agent to adapt its lifelong behavior according to situations that it has never experienced before.

The three main subsystems, i.e., head-direction cells, place cells, and goal-directed navigation, have

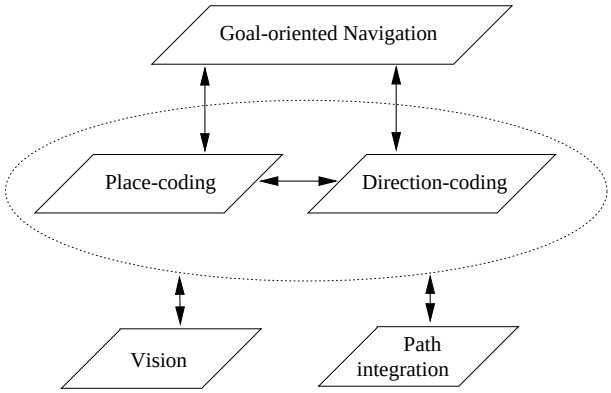

Figure 14.1: Modular architecture of the current system.

been described in the three central parts of this dissertation, Part II, III, and IV, respectively. Below, we give a brief description of their basic principles.

Head-direction Cells

The directional system models biological head-direction cells observed in three structures of the rat brain, namely the anterodorsal thalamic nucleus (ADN), the lateral mammillary nuclei (LMN), and the postsubiculum (poSC) [279, 353, 347, 388, 229, 71, 41]. Idiothetic signals strongly influence directional selectivity since ADN and LMN neurons are primarily involved in integrating angular velocity over time. On the other hand, allothetic stimuli influence poSC neurons and can therefore be used to occasionally modify the system's dynamics to calibrate head-direction cells. poSC forms the output of the entire directional system. That is, the population activity of poSC cells is used to estimate the robot's current allocentric heading. Population vector coding [121, 394] is employed to interpret the ensemble poSC pattern of activity.

Hippocampal Place Cells

The place-coding system is based on a computational model of the rat Hippocampus. Unsupervised Hebbian learning is applied to correlate visual information and path integration. This yields a stable space representation in which unreliable visual data are compensated for by internal self-motion signals, and reliable visual fixes are used to occasionally calibrate the path integrator. Hippocampal place fields [255, 256] are determined by modeling high-dimensional continuous sensory input by means of localized overlapping place fields. Starting with no prior knowledge, the system grows incrementally based on the agent-environment interaction. The anatomical areas of the hippocampal formation which are primarily involved in the model are *(i)* the superficial layers of the lateral entorhinal cortex (sLEC) in which we assume the allothetic representation is formed, *(ii)* the superficial layers of the medial entorhinal cortex in which we suppose a representation induced by path integration takes place, *(iii)* the subiculum (SC) which plays an important role in path integration, and *(iv)* the hippocampal CA3-CA1 regions in which the "final" place field representation is formed.

Navigation and Action Learning

The proposed spatial representation results in an incrementally learned coarse coding representation suitable for applying reinforcement learning to map continuous input state spaces into agent's actions. In particular, a temporal difference (TD) learning technique, namely Q-learning, is applied to acquire navigational maps. CA3-CA1 place fields work as a basis function approximator that is utilized to

learn a parameterized form of the Q-function [378, 343]. The navigation system involves the nucleus accumbens (NA), an extra-hippocampal structure that seems to be involved in reward-based goal memory and in locomotor behavior [55, 282]. In the model, CA3-CA1 place cells drive a population of locomotor action neurons in NA. The synaptic efficacy between CA3-CA1 cells and action cells is modified as a function of dopaminergic target-related reward signals. This results in an ensemble activity of the action neurons that provides a navigational map to support spatial behavior.

To conclude, we think that the present study has produced some interesting insights about the neurophysiological processes involved in spatial cognition (e.g., functional and anatomical predictions[1]), in particular for the mechanisms underlying hippocampal place cells and head-direction cells. From a robotic point of view, we have endowed an artificial agent with animal-like exploration and self-localization capabilities (e.g., bio-inspired vision-based place recognition). Furthermore, the robot has learned to accomplish effective target-oriented navigation based on its continuous interaction with the environment.

14.2 Limitations and Future Work

The attempt to answer a complex basic question most often gives rise to many other complex basic questions. Since this is the essence of scientific research, hopefully, the issues listed below are just a subset of the potential future directions sprung from this thesis.

Throughout this dissertation we have already discussed several important points that need to be addressed in order to ameliorate the soundness and the biological plausibility of our approach[1]. Our future research will focus on the following issues:

Environment manipulations. The existing model has been validated in rather stable environments. In the future, data analysis will be focused on the dynamics of the robot behavior using the same methodology as employed for experiments with rats. In particular, we will evaluate our hippocampal model through experiments concerning environment manipulations (e.g., shrinking and stretching the arena, changing light conditions). We will also address the two following issues:

- Open-field *vs* linear-track mazes. Biological place cells tend to be direction independent when the animal moves freely within an open-field arena. Alternatively, place fields become directional in the case of highly structured or linear mazes. We will investigate whether and how place fields of our model can show a similar property.

- Rich *vs* poor visual information. In principle, three distinct cue cards in an otherwise impoverished environment provide enough information for precise localization. Our current allothetic pathway relies on rich visual stimulation (e.g., standard laboratory background). We will validate the performance of our allothetic self-localization mechanism in more visually deprived environments.

Redundant vs sparse topological representation. The current model stresses the importance of having a redundant space code. The environment is covered by a large population of overlapping place fields. An interesting issue concerns investigating more sparse place coding systems and comparing their self-localization capability with that of the current system with respect to neurophysiological data.

Active search within the visual input space. In the current model, visual data are interpreted by means of filters (e.g., Gabor filters) which are either moved over the whole visual scene to detect specific patterns, or are evaluated in fixed image locations (i.e., retinotopic sampling). In order to optimize the

[1] We refer the reader to the following Discussion sections: Secs. 4.4, 5.3, 8.3, 9.2, 10.4, 13.3.

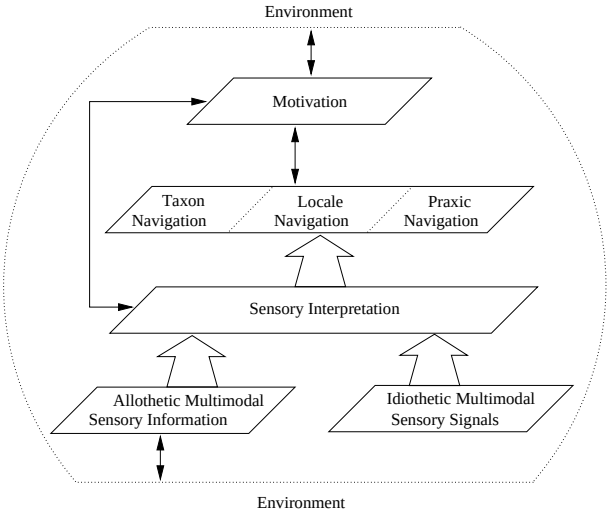

Figure 14.2: The extended system.

interpretation of visual data, we might switch to a somewhat different strategy that can be compared to saccadic search. That is, the system might perform "jumps", namely saccades, between the points of interest in the visual scene. The idea is to let the agent search for interesting cues within the image by shifting its focus of attention and use these salient landmarks for self-localization. The concept of "interesting" will be probably defined based on a saliency measure in the Gabor feature space.

Exploration strategies. What are suitable exploration strategies to optimize the interaction between a navigating agent and its environment? When an agent (rat or robot) first enters an unfamiliar environment, it does not have an internal spatial map yet. How do animals explore their environment in order to optimize spatial learning? How does the exploration strategy depend on the available sensory information? When establishing the space representation, the current system directs exploration towards unknown regions, based on the activity of existing place cells. In order to replace this strategy by a more environment-dependent exploration, we will employ the above concept of "interesting" landmark to modulate the robot's curiosity and to direct exploration.

Multiple environments: Consolidating and recalling hippocampal maps. In the current model, the problem of long-term memorization of learned space representations is not addressed. This interesting issue relates to the role of the hippocampus in memory as well as to the interaction between the hippocampus and neocortical areas (in which long-term consolidation is likely to take place). We will also address this issue from the point of view of optimizing the use of memory. Indeed, reusing previously learned maps is necessary to avoid the explosion of time and memory resources that would arise when the agent has to cope with multiple environments. In particular, the reuse of hippocampal codes devoted to previously experiences, requires the recognition of the environment from contextual cues and the recall of the appropriate representation.

Real-time acting. In the current model, the animat does not really behave in real-time. Rather, it moves in discrete time steps Δt which determine the frequency at which the robot interacts with the environment and takes an action. For a more realistic bio-inspired behavior, the animat should act in a continuous manner. Due to the non neglectable time needed for visual interpretation (e.g., about $1\,s$ to map two-dimensional images into neural activity), a possible solution might be to let the robot move also during signal processing. During this "thinking-time", the reactive controller could maintain the agent in a safety condition (e.g., providing obstacle avoidance).

Multimodal sensory information. The current model integrates external (vision) and internal (path

integration) signals. LTP learning is applied to correlate these two types of information. However, their relative contribution does not vary explicitly as a function of environmental conditions (e.g., changing light). The idea is *(i)* to extend the current system by taking into account other allothetic non-visual sensory inputs (e.g., odor or sound), *(ii)* to include a low-level sensory-gating network to modulate the contribution of each sensor modality (e.g., in the dark path integration and odor should play an important role for self-localization). Such a neural substrate for sensory selective attention might find its neurological counterpart in the thalamus, and should adapt the relative priority between sensory signals according to the current agent-environment context.

Motivation-based action selection. At a high control level, action selection can be thought of as based on the agent's current motivational state. Motivation derives from internal states (e.g., hunger, fear, curiosity) as well as from external stimulation (e.g., changing environmental situations, conflict conditions). Interestingly, the relationship between motivation and perception forms a closed loop. On the one hand, motivation is driven by perceptual signals deriving from the agent-environment interaction. On the other hand, sensory information is interpreted with respect to the current motivational state of the agent (i.e., selective attention). We will include a high-level controller to model how motivation may influence action selection[2].

Integrating different navigation strategies. This thesis has focused on cognitive (i.e., locale) navigation. The animat establishes an internal model of the spatio-temporal properties of the environment, and uses such a space representation to perform navigation. Nevertheless, animals are able to adapt their navigation strategy to the complexity of the spatial task to be solved. Adopting the optimal navigation technique is crucial for minimizing the use of resources like time and memory. In particular, animals combine taxon, praxic, and locale navigation. In the future we will extend the current model to incorporate taxon and praxic behavior modalities. A sensory evaluation process might emulate the animals' capability of adapting the complexity of their navigation strategy to the current task-environment context. For instance, the agent might resort to simple reflex-behavior to follow a smell trace, whereas it might adopt locale navigation to perform hoarding trips in complex mazes. Fig. 14.2 shows a diagram of the extended system.

Neuroethological and robotic experiments. A crucial issue of our future work will be the study of the above questions in parallel with both animals and robots. A new series of rat experiments and robotic experiments will be performed in parallel by employing equivalent experimental setups (or at least setups with the same characteristic features). Such a two-pronged experimental approach should yield a better understanding of the behavioral and cognitive mechanisms underlying intelligent navigation[3].

Validating the hypotheses/predictions derived by the model. Each component of our system (i.e., head-direction cells, path integration, visual-based space code, place cells in the hippocampus proper, and goal-oriented behavior) relies on functional as well as anatomical assumptions, and predicts that specific capabilities should be impaired by performing lesions at the level both of brain structures and of their anatomical interconnections. Some of these predictions will be tested by designing specific recording experiments with rats[4].

[2]This work will be done in collaboration with J.-A. Meyer, AnimatLab, Laboratoire d'Informatique de Paris 6, France.

[3]This work will be done in collaboration with C. Brandner of the Lab. of Neuroethology, Inst. of Physiology, Fac. of Medicine, University of Lausanne, Switzerland.

[4]This work will be done in collaboration with A. Berthoz and S. Wiener of the Lab. de Physiologie de la Perception et de l'Action (LPPA), Collège de France, Paris, France.

Appendix A

The Hippocampal Formation: Anatomical Images

The hippocampal formation is one of the most studied brain regions. This part of the limbic system has been so named by Arantius (in 1587) [185, 367] due to its characteristic three-dimensional shape. Below is an image of the rat hippocampi (adapted from Amaral and Witter (1989) [8]):

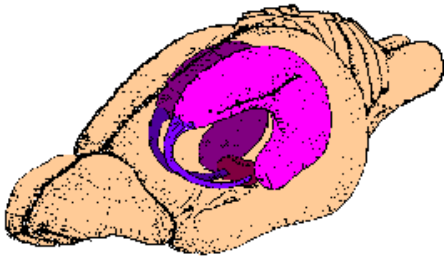

As clearly indicated by the above image, in rats the hippocampi occupy a large volume of the cerebral hemispheres. By contrast, in primates and human (whose brain volume is mostly occupied by neocortical areas) the hippocampi are buried deep within the temporal lobes. Below is a medial view of the right human brain showing the hippocampal formation and its neighboring areas (adapted from Burgess, Jeffery, and O'Keefe (1999) [60]):

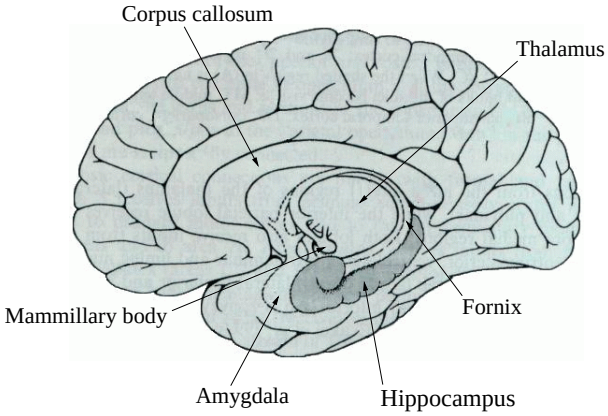

In this appendix we focus on the rat hippocampus and we show the anatomical locus of some of the regions involved by the model presented in this thesis. The following image illustrates a section of the rat brain (bregma -2.12 mm) and focuses on the left hippocampus to identify its main structures: The dentate gyrus (DG), the cornu ammonis (CA1-CA3), and the subiculum (SC).

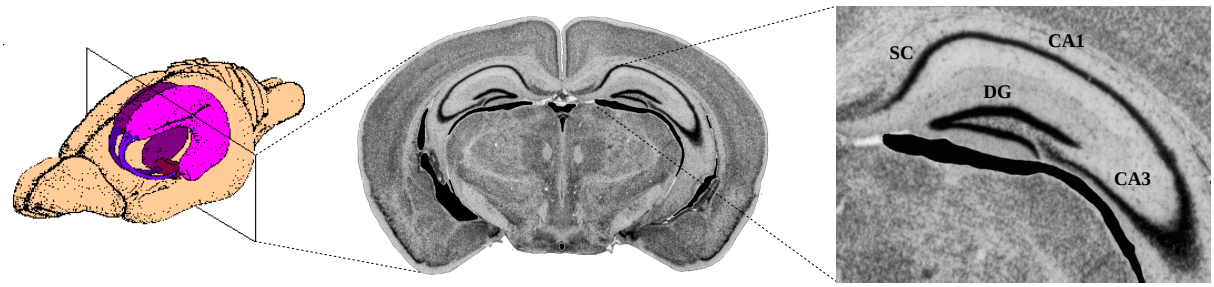

The following image is a more detailed representation of the hippocampal formation (adapted from Brown and Zador (1979) [57]):

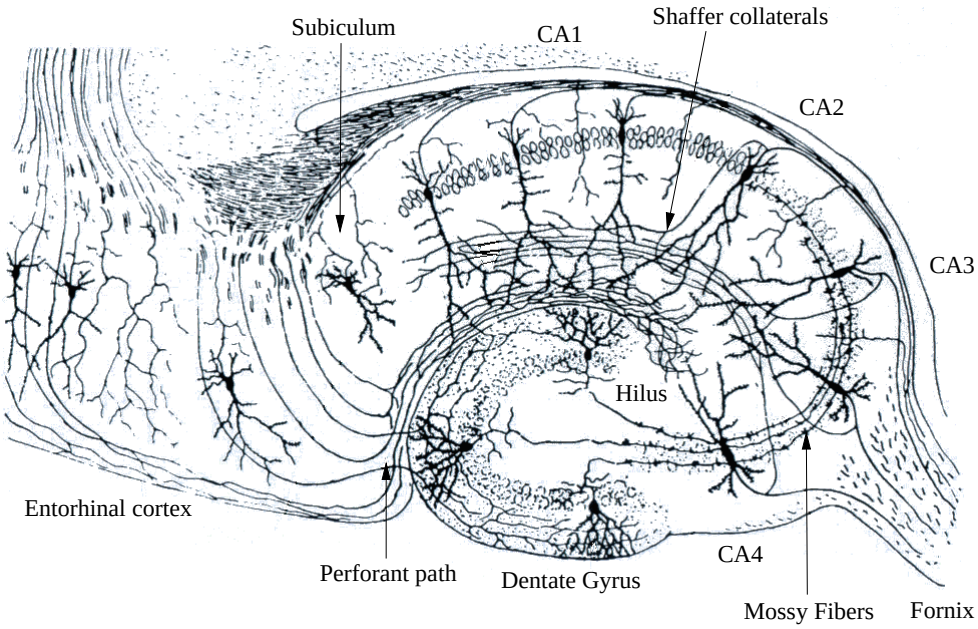

Appendix B

Neural Networks for Egocentric Bearing Estimation

B.1 Introduction

The task is to allow the robot to estimate its egocentric bearing α relative to a light source L. The robot senses light by means of eight infrared sensors and one light-detector. This results in a sensory reading vector $\vec{s}(t) = (s_1(t), \ldots, s_9(t))$. In order to interpret this noisy information, we train a feed-forward neural network $\mathcal{N}$ to learn the mapping function

$$\mathcal{M} : \mathcal{S} \rightarrow \mathcal{R}^{VIS}(\alpha) \tag{B.1}$$

where $\mathcal{S}$ is the input sensory space (i.e., $\vec{s}(t) \in \mathcal{S}$), and $\mathcal{R}^{VIS}(\alpha)$ is the population activity of our visual bearing cells VIS (Chapter 4). We let the robot approximate $\mathcal{M}$ through supervised learning.

B.2 Supervised Training

The input of the network $\mathcal{N}$ consists of the sensory vector $\vec{s}(t) = (s_1(t), \ldots, s_9(t))$. We take one hidden layer of 30 units. The output layer is our population of $S = 180$ VIS cells. The activity r_i of a neuron i in layer l is given by $f(\sum_j w_{ij} r_j)$, where j varies over all units in layer $l-1$, and f is the sigmoid function. Network $\mathcal{N}$ is trained by using gradient descent back-propagation[1] [294]. The error function being minimized is

$$E = \frac{1}{2N} \sum_{i=1}^{N} \sum_{v=1}^{S} (r_v(\vec{s}_i) - t_v(\vec{s}_i))^2 \tag{B.2}$$

where r_v and t_v are the actual and the desired value, respectively, of the v_{th} output unit, and N is the number of training examples. Training patterns are of the form $(\vec{s}_i, \vec{t}(\vec{s}_i))$, where $\vec{s}_i = (s_1, \ldots, s_9)$ is the sensory input vector, and $\vec{t}(\vec{s}_i) = (t_1(\vec{s}_i), \ldots, t_S(\vec{s}_i))$ is the corresponding target vector.

Vector $\vec{t}(\vec{s}_i)$ is the desired VIS activity encoding the bearing α which corresponds to $\vec{s}_i$. That is, if we apply population vector coding to interpret $\vec{t}(\vec{s}_i)$ we obtain α. The ensemble activity $\vec{t}(\vec{s}_i)$ is taken as a Gaussian centered on the mean $\mu = \alpha$ and with standard deviation $\sigma = 20°$. Thus, for a given α, the target activity of the v_{th} output cell is $t_v(\vec{s}_i) = \exp(-d_v^2/2\sigma^2)$, with

$$d_v = \begin{cases} \min(|\,\alpha_v - \alpha\,|,\ |\,\alpha - \alpha_v + S\,|) & \alpha_v > \alpha \\ \min(|\,\alpha_v - \alpha\,|,\ |\,\alpha_v - \alpha + S\,|) & otherwise \end{cases} \tag{B.3}$$

[1] The learning rate η is taken equal to 0.05, and the momentum term $\alpha = 0.6$.

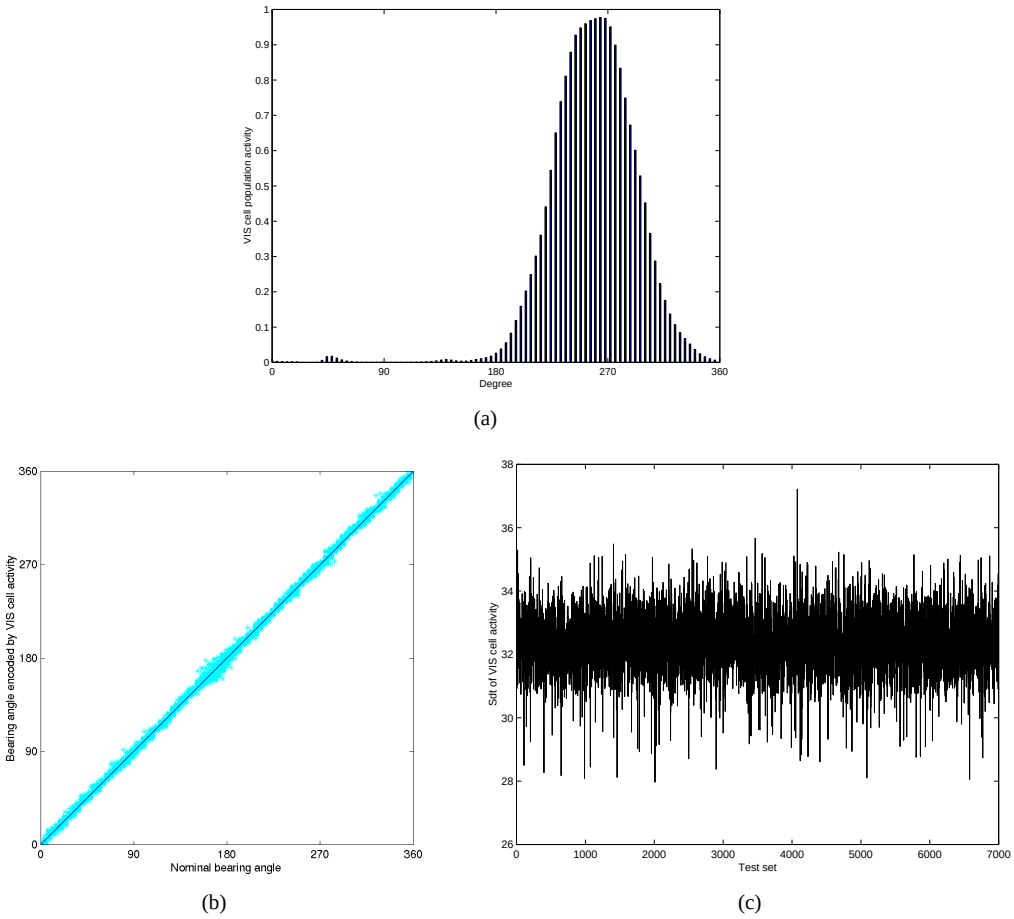

Figure B.1: (a) A sample of the output of the network after training. The output encodes a bearing angle $\overline{\alpha}$ of approximately $260°$. (b) The behavior of the network during test. The solid black line is the desired behavior, whereas grey dots represent the actual behavior. (c) The standard deviation of the network's output around the mean μ during test.

where α_v is the preferred bearing of cell $v \in$ VIS. To build the training data set, the robot is located at different distances from the light L. At each location the robot rotates by steps of $2°$. At each step, it builds a new example $(\vec{s}_i, \vec{t}(\vec{s}_i))$. After a complete turn, the robot is relocated at a different distance from L. A data set containing 28000 examples has been built by iterating this process. We take a training set of $N = 14000$ examples, a validation set of 7000 examples, and a test set of 7000 examples.

B.3 Results

Fig. B.1 (a) represents a typical output distribution for the trained network $\mathcal{N}$. According to the population decoding scheme, the ensemble activity corresponds to an angle $\overline{\alpha}$ of approximately $260°$. In order to assess the performance $\mathcal{P}$ of the trained network $\mathcal{N}$ we compute

$$\mathcal{P} = \frac{1}{M} \sum_{i=1}^{M} \mathcal{H}(\beta - |\,\overline{\alpha} - \alpha\,|) \tag{B.4}$$

where $M = 7000$ is the size of the test set, β is a fixed threshold, and $\mathcal{H}$ is the Heaviside function. The performance $\mathcal{P}$ is approximately 99% for $\beta = 10°$, and approximately 95.5% for $\beta = 5°$.

Fig. B.1 (b) shows the behavior of $\mathcal{N}$ when sampled over the test set. The solid black line represents the desired behavior, whereas grey dots represent the network's actual behavior. The dispersion of the actual output around the nominal behavior is rather uniform and narrow. Fig. B.1 (c) shows the standard deviation of the output distribution around the mean firing activity μ during the same experiment.

Appendix C

Glossary

ADN: Anterodorsal nucleus.

AT: Anterior thalamus.

CA3-CA1: Brain regions forming the hippocampus proper.

DG: Dentate gyrus.

EC: Entorhinal cortex.

HD: Head direction.

LDN: Laterodorsal nucleus.

LEC: Lateral entorhinal cortex.

LGN: Lateral geniculate nucleus.

LMN: Lateral mammillary nuclei.

LTD: Long-term depression.

LTP: Long-term potentiation.

MB: Mammillary bodies.

MEC: Medial entorhinal cortex.

PC: Posterior cingulate cortex.

PP: Posterior parietal cortex.

paHI: Parahippocampal cortex.

paSC: Parasubiculum.

peRH: Perirhinal cortex.

poSC: Postsubiculum.

prSC: Presubiculum.

SC: Subiculum.

sLEC: Superficial layers of the lateral entorhinal cortex.

sMEC: Superficial layers of the medial entorhinal cortex.

VTA: Ventral tegmental area.

Acetylcholine: A neurotransmitter.

Afferent signals: Incoming signals.

Allocentric: Relative to an external absolute frame of reference.

Allothetic stimuli: External stimuli such as visual, olfactory, auditory, and tactile signals.

Amygdala: A brain region involved in emotional and motivational behavior.

Anterior thalamus: A brain region possibly involved in the rodent directional system.

Anterodorsal nucleus: A brain region of the anterior thalamus possibly involved in head-direction representation.

Back-propagation: Gradient-descent learning algorithm employed for training artificial neural networks.

Caudate nucleus (dorsal striatum): A brain region possibly involved in the rodent directional system.

Cholinergic: From acetylcholine, a neurotransmitter.

Dead-reckoning: Also termed *path integration* and *vector navigation*, is the process that allows a navigating system to infer its position with respect to a departure point (i.e., *homing vector*) based on purely inertial self-motion signals.

Declarative memory: Memory of facts, names, and episodes.

Deep entorhinal cortex: Layers IV-VI of the entorhinal region.

Dentate gyrus: Brain region, belonging to the hippocampal formation, possibly involved in space representation.

Dopamine: A neurotransmitter.

Dorsal presubiculum (postsubiculum): A structure, within the hippocampal formation, possibly involved in the head-direction system.

Dorsal striatum (caudate nucleus): A brain region possibly involved in the rodent directional system.

EEG: ElectroEncephaloGraph.

Efferent signals: Outcoming signals.

Egocentric: Relative to an observer-centered frame of reference.

Entorhinal cortex: Six-layered cortical structure providing most of the cortical inputs to the hippocampal formation. It is divided into superficial and deep areas. In rat, it is also divided into lateral and medial regions.

Episodic memory: Memory of events (see declarative memory).

Fornix: Fiber bundle connecting the hippocampus and the subiculum to subcortical structures.

GABA: A neurotransmitter, generally inhibitory.

Glutamate: A neurotransmitter, generally excitatory.

Head-direction cell: A neuron whose activity encodes the rodent's allocentric headings in the horizontal plane.

Hebbian learning: A local learning mechanism based on the correlation between the activity of the pre- and the post-synaptic neuron.

Hippocampal formation: The group of brain regions including the hippocampus proper and its adjacent regions (e.g., the dentate gyrus).

Hippocampus proper: The hippocampal region including CA3 and CA1 areas.

Idiothetic stimuli: Internally-generated information such as vestibular and proprioceptive signals.

Latent learning: Form of learning occurring without any explicit reward (i.e., food).

Lateral entorhinal cortex: A portion of the entorhinal cortex.

Lateral geniculate nucleus: A brain region belonging to the visual pathway (it receives topographically organized input from both retinae and projects to the visual cortex).

Lateral mammillary nuclei: A brain region possibly involved in head-direction representation.

Laterodorsal nucleus: A thalamic brain region possibly involved in head-direction representation.

Limbic system (medial temporal lobe): A brain area grouping the olfactory cortex, the amygdala, and the hippocampus.

Local-view: Usually, it identifies a spatial representation based on spatio-temporal relationships between an agent and salient allothetic cues.

Long-term potentiation: A basic learning mechanism in the brain that consists of increasing the synaptic strength of the connection between two neurons whenever they fire synchronously.

Mammillary bodies: A brain region where the fornix (projecting from the hippocampus) ends.

Medial entorhinal cortex: A portion of the entorhinal cortex.

NMDA (N-methyl-D-aspartate): A neuroreceptor involved in LTP.

Parahippocampal cortex: A primate brain region. The analogous in rats is the postrhinal cortex.

Parasubiculum: A brain region belonging to the subicular cortex.

Path integration: Also termed *dead-reckoning* and *vector navigation*, is the process that allows a navigating system to infer its position with respect to a departure point (i.e., *homing vector*) based on purely inertial self-motion signals.

Perforant path: Main neural pathway from the superficial entorhinal cortex to the hippocampus.

Perirhinal cortex: A cortical area possibly involved in novelty detection.

Place cell: Spatially tuned neuron whose activity codes for a localized environmental area (its place field).

Place field: Spatial region in which a place cell exhibits its maximal firing activity.

Posterior cingulate cortex (retrosplenial cortex): A cortical area possibly involved in spatial orientation.

Posterior parietal cortex: Neocortical area possibly involved in spatial cognition.

Postrhinal cortex: A rat brain region. The analogous in primate is the parahippocampal cortex.

Postsubiculum (dorsal presubiculum): A structure, within the hippocampal formation, possibly involved in the head-direction system.

Presubiculum: A brain region belonging to the subicular cortex.

Procedural memory: Memory of skills learned through practice.

Retrosplenial cortex (posterior cingulate cortex): A cortical area possibly involved in spatial orientation.

Septal nuclei: A brain region providing cholinergic and GABA-ergic input to the hippocampus.

Simple cell: Oriented-bar sensitive neuron in the primary visual cortex (V1).

Subiculum: A brain region belonging to the hippocampal formation.

Superficial entorhinal cortex: Layers II-III of the entorhinal region.

Theta rhythm: EEG sinusoidal signal in the hippocampal formation with a frequency of $7 - 12\,Hz$.

Vector Navigation: Also termed *path integration* and *dead-reckoning*, is the process that allows a navigating system to infer its position with respect to a departure point (i.e., *homing vector*) based on purely inertial self-motion signals.

Ventral tegmental area: A structure in the midbrain.

Vestibular system: Has the purpose of keeping tabs on the position and motion of the head in space.

Bibliography

[1] L.F. Abbott and K.I. Blum. Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex*, 6:406–416, 1996.

[2] E.D. Adrian. The impulses produced by sensory nerve endings. *Journal of Physiology (London)*, 61:49–72, 1926.

[3] S. Alyan and R. Jander. Short-range homing in the house mouse, *Mus musculus*: Stage in the learning of directions. *Animal Behaviour*, 48:285–298, 1994.

[4] S.H. Alyan, B.M. Paul, E. Ellesworth, R.D. White, and B.L. McNaughton. Is the hippocampus required for path integration? *Society for Neuroscience Abstracts*, 23:504, 1997.

[5] D.G. Amaral. Emerging principles of intrinsic hippocampal organization. *Current Opinion in Neurobiology*, 3:225–229, 1993.

[6] D.G. Amaral, C. Dolorfo, and P. Alvarez-Royo. Organization of Ca1 projections to the subiculum: a PHA-L analysis in the rat. *Hippocampus*, 1:415–435, 1991.

[7] D.G. Amaral, N. Ishizuka, and B. Claiborne. Neurons, numbers and the hippocampal network. *Progress in Brain Research*, 83:1–11, 1990.

[8] D.G. Amaral and M.P. Witter. The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, 31(3):571–591, 1989.

[9] S.-I. Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27:77–87, 1977.

[10] P. Andersen, T.V.P. Bliss, and K.K. Skrede. Lamellar organization of hippocampal excitatory path ways in the rat. *Experimental Brain Research*, 13:222–238, 1971.

[11] H.C. Andrews. *Computer techniques in image processing*. Academic Press, New York, 1970.

[12] A. Arleo, J. del R. Millán, and D. Floreano. Efficient learning of variable-resolution cognitive maps for autonomous indoor navigation. *IEEE Transactions on Robotics and Automation*, 15(6):990–1000, 1999.

[13] A. Arleo and W. Gerstner. A vision-driven model of hippocampal place cells and temporally asymmetric LTP-induction for action learning. In *Artificial Neural Networks - ICANN'99. 9th International Conference*, Edinburgh, UK, 1999. IEE, London.

[14] A. Arleo and W. Gerstner. Neuro-mimetic navigation systems: A computational model of the rat Hippocampus. In *Proceedings of the Conference on Situated Artificial Intelligence (IAS99)*, pages 193–211, Paris, 1999. Hermès.

[15] A. Arleo and W. Gerstner. Modeling rodent head-direction cells and place cells for spatial learning in biomimetic robotics. In J.-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S.W. Wilson, editors, *From Animals to Animats VI*, pages 236–245, Cambridge MA, 2000. MIT Press.

[16] A. Arleo and W. Gerstner. Spatial cognition and neuro-mimetic navigation: A model of hippocampal place cell activity. *Biological Cybernetics, Special Issue on Navigation in Biological and Artificial Systems*, 83:287–299, 2000.

[17] A. Arleo and W. Gerstner. Spatial orientation in navigating agents: Modeling head-direction cells. *Neurocomputing*, (to appear), 2001.

[18] K.B. Austin, L.H. White, and M.L. Shapiro. Short- and long-term effects of experience on hippocampal place fields. *Society for Neuroscience Abstracts*, 19:263, 1993.

[19] M.E. Bach, R.D. Hawkins, M. Osman, E.R. Kandel, and M. Mayford. Impairment of spatial but not contextual memory in CaMKII mutant mice with a selective loss of hippocampal LTP in the range of the theta frequency. *Cell*, 81:905–915, 1995.

[20] R.J. Baddeley and P.J.B. Hancock. A statistical analysis of natural images matches psychophysically derived orientation tuning curves. *Proceedings of the Royal Society of London, series B*, 246:219–223, 1991.

[21] J.S. Barlow. Inertial navigation as a basis for animal navigation. *Journal of Theoretical Biology*, 6:76–117, 1964.

[22] C.A. Barnes. Memory deficits associated with senscence: A neurophysiological and behavioral study in the rat. *Journal of Comparative and Physiological Psychology*, 93:74–104, 1979.

[23] C.A. Barnes, B.L. McNaughton, S.J.Y. Mizumori, B.W. Leonard, and L.H. Lin. Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. In J. Storm-Mathisen, J. Zimmer, and O.P. Ottersen, editors, *Understanding the Brain through the Hippocampus: The Hippocampal Region as a Model for Studying Brain Structure and Function*, volume 83 of Progress in Brain Research, pages 287–300. Elsevier Science Publishers B.V., New York, 1990.

[24] C.A. Barnes, L. Nadel, and W.K. Honig. Spatial memory deficit in senescent rats. *Canadian Journal of Psychology*, 34(1):29–39, 1980.

[25] C.A. Barnes, M.S. Suster, J. Shen, and B.L. McNaughton. Multistability of cognitive maps in the hippocampus of old rats. *Nature*, 388(6639):272–275, 1997.

[26] S. Becker and N. Burgess. Modelling spatial recall, mental imagery and neglect. In *Advances in Neural Processing Systems 13*. MIT Press, 2001.

[27] R.D. Beer. *Intelligence as adaptive behavior: An experiment in computational neuroethology*. Academic Press, San Diego, CA, 1990.

[28] S. Benhamou, J.-P. Sauvé, and P. Bovet. Spatial memory in large scale movements: Efficiency and limitation of the egocentric coding process. *Journal of Theoretical Biology*, 145:1–12, 1990.

[29] M. Bentivoglio, K. Kultas-Ilinsky, and I. Illinsky. Limbic thalamus: Structure, intrinsic organisation, and connections. In B.A. Vogt and M. Gabriel, editors, *Neurobiology of Cingulate Cortex and Limbic Thalamus: A Comprehensive Handbook*, pages 71–122. Birkhauser, Boston, 1993.

[30] J.S. Beritoff. *Neural Mechanisms of Higher Vertebrates*. Boston: Little Brown, 1965.

[31] A. Berthoz. Hippocampal and parietal contribution topokinetic and topographic memory. In K.J. Jeffery N. Burgess and J. O'Keefe, editors, *The Hippocampal and Parietal Foundations of Spatial Cognition*, chapter 20, pages 381–403. Oxford University Press, 1999.

[32] R. Biegler and R.G.M. Morris. Landmark stability is a prerequisite for spatial but not discrimination learning. *Nature*, 361:631–633, 1993.

[33] R. Biegler and R.G.M. Morris. Landmark stability: Studies exploring whether the perceived stability of the environment influences spatial representation. *Journal of Experimental Biology*, 199(1):187–193, 1996.

[34] J. Bigün. Speed, frequency, and orientation tuned 3-D Gabor filter banks and their design. In *Proceedings of the International Conference on Pattern Recognition*, volume C, pages 184–187. IEEE Computer Society, 1994.

[35] A.R. Bisetzky. Die Tänze der Bienen nach einem Fussweg zum Futterplatz. *Z. verg. Physiol.*, 40:264–288, 1957.

[36] C.M. Bishop. *Neural networks for pattern recognition*. Oxford University Press, Oxford, 1995.

[37] H.T. Blair. A thalamocortical circuit for computing directional heading in the rat. In D.S. Touretzky, M.C. Mozer, and M.E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*, pages 152–158. MIT Press, 1996.

[38] H.T. Blair, J. Cho, and P.E. Sharp. Role of the lateral mammillary nucleus in the rat head direction circuit: A combined single unit recording and lesion study. *Neuron*, 21:1387–1397, 1998.

[39] H.T. Blair, B.W. Lipscomb, and P.E. Sharp. Temporal tuning of anticipatory head-direction cells in the anterior thalamus of the rat. *Society for Neuroscience Abstracts*, 22:913, 1996.

[40] H.T. Blair, B.W. Lipscomb, and P.E. Sharp. Anticipatory time intervals of head-direction cells in the anterior thalamus of the rat: Implications for path integration in the head-direction circuit. *Journal of Neurophysiology*, pages 145–159, 1997.

[41] H.T. Blair and P.E. Sharp. Anticipatory head direction signals in anterior thalamus: Evidence for a thalamocortical circuit that integrates angular head motion to compute head direction. *Journal of Neuroscience*, 15(9):6260–6270, 1995.

[42] W. Bles, J.M.B.V. de Jong, and G. De Wit. Somatosensory compensation for loss of labyrinthine function. *Acta Otolaryngology*, 97:213–221, 1984.

[43] T.V.P. Bliss and G.L. Collingridge. Long term potentiation: A synaptic model of memory. *Nature*, 361:31–39, 1993.

[44] T.V.P. Bliss and T. Lomo. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*, 232:331–356, 1973.

[45] T.V.P. Bliss and M.A. Lynch. Long-term potentiation of synaptic transmission in the hippocampus: Properties and mechanisms. In P.W. Landfield and S.A. Deadwyler, editors, *Long-Term Potentiation: From Biophysics to Behavior*, pages 3–72. Liss, New York, 1988.

[46] K.I. Blum and L.F. Abbott. A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, 8:85–93, 1996.

[47] J. Borenstein, H.R. Everett, and L. Feng. Where am I? Systems and methods for mobile robot positioning. Technical report, University of Michigan, 1996.

[48] J. Borenstein and L. Feng. Correction of systematic odometry errors in mobile robots. In *Proceedings of the International Conference on Intelligent Robots and Systems*, pages 569–574, 1995.

[49] A. Bose, V. Booth, and M. Recce. A temporal mechanism for generating the phase precession of hippocampal place cells. *Journal of Computational Neuroscience*, 9:5–30, 2000.

[50] E. Bostock, R.U. Muller, and J.L. Kubie. Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, 1(2):193–206, 1991.

[51] C.R. Breese, R.E. Hampson, and S.A. Deadwyler. Hippocampal place cells: Stereotypy and plasticity. *Journal of Neuroscience*, 9(4):1097–1111, 1989.

[52] R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Robotics and Automation*, RA-2:14–23, 1986.

[53] R.A. Brooks. Elephants don't play chess. *Robotics and Autonomous Systems*, 6:3–15, 1990.

[54] R.A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

[55] M.A. Brown and P.E. Sharp. Simulation of spatial-learning in the Morris water maze by a neural network model of the hippocampal-formation and nucleus accumbens. *Hippocampus*, 5:171–188, 1995.

[56] T.H. Brown, E.W. Kairiss, and C.L. Keenan. Hebbian synapses: Biophysical mechanisms and algorithms. *Annual Review of Neuroscience*, 13:475–511, 1990.

[57] T.H. Brown and A.M. Zador. Hippocampus. In G.M. Shepherd, editor, *The synaptic organization of the brain*, chapter 11. Oxford University Press, 1979.

[58] T.H. Brown, A.M. Zador, Z.F. Mainen, and B.J. Clairborne. Hebbian modifications in hippocampal neurons. In M. Boudry and J.L. Davis, editors, *Long-Term Potentiation: A Debate of Current Issues*, pages 357–389. MIT Press, Cambridge, MA, 1991.

[59] N. Burgess, J.G. Donnett, K.J. Jeffery, and J. O'Keefe. Robotic and neural simulation of the Hippocampus and rat. *Phil. Trans. of the Royal Society [Biol]*, 352, 1997.

[60] N. Burgess, K.J. Jeffery, and J. O'Keefe. Integrating hippocampal and parietal functions: a spatial point of view. In K.J. Jeffery N. Burgess and J. O'Keefe, editors, *The Hippocampal and Parietal Foundations of Spatial Cognition*, chapter 1, pages 3–29. Oxford University Press, 1999.

[61] N. Burgess and J. O'Keefe. Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*, 6:749–762, 1996.

[62] N. Burgess, M. Recce, and J. O'Keefe. A model of hippocampal function. *Neural Networks*, 7:1065–1081, 1994.

[63] R.D. Burwell, M.P. Witter, and D.G. Amaral. Perirhinal and postrhinal cortices of the rat: A review of the neuroanatomical literature and comparison with findings from the monkey brain. *Hippocampus*, 5:390–408, 1995.

[64] G. Buzsáki. Feed-forward inhibition in the hippocampal formation. *Progress in Neurobiology*, 22:131–153, 1984.

[65] G. Buzsáki. The hippocampo-neocortical dialogue. *Cerebral Cortex*, 6:81–92, 1996.

[66] G. Buzsáki, L.S. Chen, and F.H. Gage. Spatial organization of physiological activity in the hippocampal region: Relevance to memory formation. *Progress in Brain Research*, 83:257–268, 1990.

[67] R. Caminiti, P.B. Johnson, C. Galli, S. Ferraina, and Y. Burnod. Making arm movements within different parts of space: The premotor and motor cortical representations of a coordinate system for reaching to visual targets. *Journal of Neuroscience*, 11:1182–1197, 1991.

[68] R. Chatila and J.P. Laumond. Position referencing and consistent world modeling for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1985.

[69] L.L. Chen. *Head-Directional Information Processing in the Rat Posterior Cortical Areas*. PhD thesis, University of Colorado, 1991.

[70] L.L. Chen, L. Lin, C.A. Barnes, and B.L. McNaughton. Head-direction cells in the rat posterior cortex. II. Contributions of visual and idiothetic information to the directional firing. *Experimental Brain Research*, 101:23–34, 1994.

[71] L.L. Chen, L. Lin, E.J. Green, C.A. Barnes, and B.L. McNaughton. Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Experimental Brain Research*, 101:8–23, 1994.

[72] D. Cliff, P. Husband, J.A. Meyer, and S.W. Wilson. *From animals to animats 3 - SAB94*. MIT Press-Bradford Books, Cambridge, Massachusetts, 1994.

[73] D.T. Cliff. Computational neuroethology: a provisional manifesto. In J. A. Meyer and S. W. Wilson, editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, Cambridge, Massachusetts, 1991. MIT Press-Bradford Books.

[74] D.T. Cliff and S.G. Bullock. Adding "foveal vision" to Wilson's animat. *Adaptive Behaviour*, 2:49–72, 1993.

[75] L.M. Colle and R.A. Wise. Circling induced by intra-accumbens amphetamine injections. *Psychopharmacology*, 105:157–161, 1991.

[76] T.S. Collett, B.A. Cartwright, and B.A. Smith. Landmark learning and visuospatial memories in gerbils. *Journal of Comparative Physiology A*, 158:835–851, 1986.

[77] T.S. Collett, E. Dillman, A. Giger, and R. Wehner. Visual landmarks and route following in desert ants. *Journal of Comparative Physiology*, 170:435–442, 1992.

[78] T.S. Collett and J. Zeil. Places and landmarks: An arthropod perspective. In Sue Healy, editor, *Spatial Representation in Animals*, chapter 2, pages 18–53. Oxford University Press, 1998.

[79] A. Cressant, R.U. Muller, and B. Poucet. Failure of centrally placed objects to control firing fields of hippocampal place cells. *Journal of Neuroscience*, 17(7):2531–2542, 1997.

[80] F.H. Lopes da Silva, M.P. Witter, P.H. Boeijinga, and A.H.M. Lohman. Anatomical organization and physiology of the limbic cortex. *Physiological Reviews*, 70:453–511, 1990.

[81] C. Darwin. Origin of certain instincts. *Nature*, 7:417–418, 1873.

[82] J.G. Daugman. Two-dimensional spectral analysis of cortical receptive field profiles. *Vision Research*, 20:847–856, 1980.

[83] P. Dayan. Navigating through temporal difference. In R.P. Lippmann, J.E. Moody, and D.S. Touretzky, editors, *Advances in Neural Information Processing Systems 3*, pages 464–470. Morgan Kaufmann, San Mateo, CA, 1991.

[84] P. Dayan and T. Sejnowski. TD(λ) converges with probability 1. *Machine Learning*, 14:295–301, 1994.

[85] D. Debanne, B.H. Gahwiler, and S.M. Thompson. Asynchronous pre- and postsynaptic activity induces associative long-term depression in area CA1 of the rat hippocampus *in vitro*. *Proceedings Natl. Acad. Science USA*, 91:1148–1152, 1994.

[86] J. del R. Millán. Rapid, safe, and incremental learning of navigation strategies. *IEEE Transactions on Systems, Man and Cybernetics—Part B*, 26:408–420, 1996.

[87] M. Dorigo. Editorial introduction to the special issue on learning autonomous robots. *IEEE Transactions on Systems, Man and Cybernetics-Part B*, 26:361–364, 1993.

[88] M. Dorigo and M. Colombetti. *Robot shaping*. MIT Press-Bradford Books, Cambridge, Massachusetts, 1998.

[89] M. Dorigo and U. Schnepf. Genetic-based machine learning and behaviour based robotics. *IEEE Transactions on Systems, Man and Cybernetics*, 23:141–154, 1993.

[90] J. Droulez and A. Berthoz. The concept of dynamic memory in sensorimotor control. In D.R. Humphrey and H.-J. Freund, editors, *Motor Control: Concepts and Issues*, pages 137–161. John Wiley Sons, 1991.

[91] H. Eichenbaum and N.J. Cohen. Representation in the hippocampus: What do hippocampal neurons code? *Trends in Neuroscience*, 11(6):244–248, 1988.

[92] H. Eichenbaum, M. Kuperstein, A. Fagan, and J. Nagode. Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *Journal of Neuroscience*, 7(3):716–732, 1987.

[93] H. Eichenbaum, C. Stewart, and R.G.M. Morris. Hippocampal representation in place learning. *Journal of Neuroscience*, 10(11):3531–3542, 1990.

[94] A. Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, 3:249–265, 1987.

[95] A. Elfes. *Occupancy grids: A probabilistic framework for robot perception and navigation*. PhD thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 1989.

[96] B. Ermentrout and J. Cowan. A mathematical theory of visual hallucination patterns. *Biological Cybernetics*, 34:137–150, 1979.

[97] A.S. Etienne. Mammalian navigation, Neural Models and Biorobotics. *Connection Science*, 10:271–289, 1998.

[98] A.S. Etienne, J. Berlie, J. Georgakopoulos, and R. Maurer. Role of dead reckoning in navigation. In Sue Healy, editor, *Spatial Representation in Animals*, chapter 3, pages 54–68. Oxford University Press, 1998.

[99] A.S. Etienne, R. Maurer, J. Berlie, B. Reverdin, T. Rowe, J. Georgakopoulos, and V. Séguinot. Navigation through vector addition. *Nature*, 396:161–164, 1998.

[100] A.S. Etienne, R. Maurer, and F. Saucy. Limitations in the assessment of path dependent information. *Behavior*, 106(1/2):81–111, 1988.

[101] A.S. Etienne, R. Maurer, F. Saucy, and E. Teroni. Short-distance homing in the golden hamster after a passive outward journey. *Animal Behavior*, 34:696–715, 1986.

[102] A.S. Etienne, R. Maurer, and V. Séguinot. Path integration in mammals and its interaction with visual landmarks. *Journal of Experimental Biology*, 199:201–209, 1996.

[103] A.S. Etienne, E. Teroni, C. Hurni, and V. Portenier. The effect of a single light cue on homing behaviour of the golden hamster. *Animal Behaviour*, 39:17–41, 1990.

[104] D. Floreano and F. Mondada. Evolution of plastic neurocontrollers for situated agents. In P. Maes, M. Mataric, J. Meyer, J. Pollack, H. Roitblat, and S. Wilson, editors, *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, Cambridge, Massachusetts, 1996. MIT Press-Bradford Books.

[105] P.A. Fortier, J.F. Kalaska, and A.M. Smith. Cerebellar neuronal activity related to whole-arm reaching movements in the monkey. *Journal of Neurophysiology*, 62:198–211, 1989.

[106] D.H. Foster and P.A. Ward. Asymmetries in oriented-line detection indicate two orthogonal filters in early vision. *Proceedings of the Royal Society of London, series B*, 243:75–81, 1991.

[107] D.J. Foster, R.G.M. Morris, and P. Dayan. A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus*, 10(1):1–16, 2000.

[108] T.C. Foster, C.A. Castro, and B.L. McNaughton. Spatial selectivity of rat hippocampal neurons: dependence on preparedness for movement. *Science*, 244:1580–1582, 1989.

[109] M.O. Franz, B. Schölkopf, H.A. Mallot, and H.H. Bülthoff. Learning View Graphs for Robot Navigation. *Autonomous Robots*, 5:111–125, 1998.

[110] B. Fritzke. Growing cell structures –A self-organizing network for unsupervised and supervised learning. *Neural Networks*, 7(9):1441–1460, 1994.

[111] N. Fujita, R.L. Klatzky, J.M. Loomis, and R.G. Golledge. The encoding-error model of pathway completion without vision. *Geogr. Analysis*, 25:295–314, 1993.

[112] N. Fujita, J.M. Loomis, R.L. Klatzky, and R.G. Golledge. A minimal representation for dead-reckoning navigation: Updating the homing vector. *Geogr. Analysis*, 22:326–335, 1990.

[113] D. Gabor. Theory of communication. *Journal of the IEE*, 93:429–457, 1946.

[114] C.R. Gallistel. *The organization of learning*. MIT Press, Cambridge, MA, 1990.

[115] P. Gaussier. Special issue on animat approach to control autonomous robots interacting with an unknown world. *Robotics and Autonomous Systems*, 16, 1995.

[116] P. Gaussier, C. Joulain, A. Revel, S. Zrehen, and J.P. Banquet. Animat navigation: about the interest of hippocampus studies. In *Fourth European Conference on Artificial Life*, pages 299–308, 1997.

[117] P. Gaussier, S. Leprêtre, C. Joulain, A. Revel, M. Quoy, and J.P. Banquet. Animal and robot learning: Experiments and models about visual navigation. In *7th European Workshop on Learning Robots*, Edinburgh, UK, 1998.

[118] P. Gaussier and S. Zrehen. PerAc: A neural architecture to control artificial animals. *Robotics and Autonomous Systems*, 16(2-4):291–320, 1995.

[119] V.V. Gavrilov, S.I. Wiener, and A. Berthoz. Whole body rotations enhance hippocampal theta rhythm slow activity in awake rats passively transported on a mobile robot. *Annals of the New York Academy of Sciences*, 781:385–398, 1996.

[120] A.P. Georgopoulos, R.E. Kettner, and A. Schwartz. Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Neuroscience*, 8:2928–2937, 1988.

[121] A.P. Georgopoulos, A. Schwartz, and R.E. Kettner. Neuronal population coding of movement direction. *Science*, 233:1416–1419, 1986.

[122] W. Gerstner and L.F. Abbott. Learning navigational maps through potentiation and modulation of hippocampal place cells. *Journal of Computational Neuroscience*, 4:79–94, 1997.

[123] C.D. Gilbert and T.N. Wiesel. The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research*, 30:1689–1701, 1990.

[124] J.A.M. Van Gisbergen, A.J. Van Opstal, and A.M.M. Tax. Collicular ensemble coding of saccades based on vector summation. *Neuroscience*, 21:541–555, 1987.

[125] S. Glasauer, M.-A. Amorim, E. Vitte, and A. Berthoz. Goal-directed linear locomotion in normal and labyrinthine-defective subjects. *Experimental Brain Research*, 98:323–335, 1994.

[126] E.J. Golob and J.S. Taube. Head direction cells recorded from the postsubiculum in animals with lesions of the lateral dorsal thalamic nucleus. *Society for Neuroscience Abstracts*, 20:805, 1994.

[127] C.R. Goodlett, J.M. Nichols, R.W. Halloran, and J.R. West. Long-term deficits in water maze spatial conditional alternation performance following retrohippocampal lesions in rats. *Behavioral Brain Research*, 32:63–67, 1989.

[128] J.P. Goodridge and J.S. Taube. Lesions of the anterior thalamic nucleus disrupt head direction cell firing in the dorsal presubiculum. *Society for Neuroscience Abstracts*, 19:796, 1993.

[129] J.P. Goodridge and J.S. Taube. The effect of lesions of the postsubiculum on head direction cell firing in the anterior thalamic nuclei. *Society for Neuroscience Abstracts*, 20:805, 1994.

[130] J.P. Goodridge and J.S. Taube. Preferential use of landmark navigational system by head direction cells in rats. *Behavioral Neuroscience*, 109:49–61, 1995.

[131] P. Görner and B. Claas. Homing behavior and orientation in the funne-web spider, *Agelema labyrinthica* Clerck. In F. Barth, editor, *Neurobiology of Arachnids*. Springer, Berlin, 1985.

[132] K.M. Gothard, W.E. Skaggs, and B.L. McNaughton. Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. *Journal of Neuroscience*, 16(24):8027–8040, 1996.

[133] K.M. Gothard, W.E. Skaggs, K.M. Moore, and B.L. McNaughton. Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *Journal of Neuroscience*, 16(2):823–835, 1996.

[134] J.A. Gray. *The Neuropsychology of Anxiety: An Enquiry into the Functions of the Septo-Hippocampal System*. Oxford University Press, New York, 1982.

[135] J. Green and A. Arduini. Hippocampal electrical activity in arousal. *Journal of Neurophysiology*, 17:533–557, 1954.

[136] A.S. Griffin and A.S. Etienne. Updating the path integrator through a visual fix. *Psychobiology*, 26(3):240–248, 1998.

[137] T. Van Groen and J.M Wyss. Extrinsic projections from area CA1 of the rat hippocampus: Olfactory, cortical, subcortical, and bilateral hippocampal formation projections. *Journal of Comparative Neurology*, 302:515–528, 1990.

[138] T. Van Groen and M.J. Wyss. The postsubicular cortex in rat: Characterization of the fourth region of the subicular cortex and its connections. *Brain Research*, 529:165–177, 1990.

[139] H.J. Groenewegen and F.T. Russchen. Organization of the efferent projections of the nucleus accumbens to pallidal, hypothalamic and mesencephalic structures. A tracing and immunohistochemical study in the cat. *Journal of Comparative Neurology*, 223:347–367, 1984.

[140] P.J.B. Hancock, R.J. Baddeley, and L.S. Smith. The principal components of natural images. *Network: Computation in Neural Systems*, 3(1):61–70, 1992.

[141] G. Hartmann and R. Wehner. The ant's path integration system: a neural architecture. *Biological Cybernetics*, 73(6):483–497, 1995.

[142] M.E. Hasselmo and J.M. Bower. Acetylcholine and memory. *Trends in Neurosciences*, 16(6):218–222, 1993.

[143] M.E. Hasselmo, E.T. Rolls, and G.C. Baylis. The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. *Experimental Brain Research*, 32:203–218, 1989.

[144] S. Haykin. *Neural Networks. A Comprehensive Foundation*. MacMillan, Englewood Cliffs, NJd, 1994.

[145] D.O. Hebb. *The organization of behaviour: A neuropsychological theory*. J. Wiley, New York, 1949.

[146] J. Hertz, A. Krogh, and R.G. Palmer. *Introduction to the theory of neural computation*. Addison-Wesley, Redwood City, CA, 1991.

[147] A.J. Hill. First occurrence of hippocampal spatial firing in a new environment. *Experimental Neurology*, 62:282–297, 1978.

[148] A.J. Hill and P.J. Best. Effects of deafness and blindness on the spatial correlates of hippocampal unit activity in the rat. *Experimental neurology*, 74:204–217, 1981.

[149] R. Hirsh. The hippocampus and contextual retrieval of information from memory: A theory. *Behavioral Biology*, 12:421–444, 1974.

[150] R. Hirsh, B. Leber, and K. Gillman. Fornix fibers and motivational states as controllers of behavior: A study stimulated by the contextual retrieval theory. *Behavioral Biology*, 22:463–478, 1978.

[151] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160:106–154, 1962.

[152] A. Hughes. A schematic eye for the rat. *Visual Resources*, 19:569–588, 1977.

[153] R. Insausti, M.T. Herrero, and M.P. Witter. Entorhinal cortex of the rat: Cytoarchitectonic subdivision and the origin and distribution of cortical efferents. *Hippocampus*, 7:146–83, 1997.

[154] R.L. Isaacson. *The Limbic System*. Plenum Press, New York, 1974.

[155] R. Jander. The optical directional orientation of the red ant (*Formica rufa* L.). *Zeitschrift für vergleichende Physiologie*, 40:162–238, 1957.

[156] L.E. Jarrard. Selective hippocampal lesions and behavior: Effects of kainic acid lesions on performance of place and cue tasks. *Behavioral Neuroscience*, 97:873–889, 1983.

[157] L.E. Jarrard. On the role of the hippocampus in learning and memory in the rat. *Behavioral Neural Biology*, 60(1):9–26, 1993.

[158] K.J. Jeffery. Learning of landmark stability and instability by hippocampal place cells. *Neuropharmacology*, 37:677–687, 1998.

[159] K.J. Jeffery, J.G. Donnett, N. Burgess, and J. O'Keefe. Directional control of hippocampal place fields. *Experimental Brain Research*, 117:131–142, 1997.

[160] I.T. Jolliffe. *Principal Component Analysis*. Springer-Verlag, New York, 1986.

[161] J.P. Jones and L.A. Palmer. An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat strate cortex. *Journal of Neurophysiology*, 58(6):1233–1258, 1987.

[162] M.W. Jung and B.L. McNaughton. Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*, 3(2):165–182, 1993.

[163] M.W. Jung, S.I. Wiener, and B.L. McNaughton. Comparison of spatial firing characteristics of the dorsal and ventral hippocampus of the rat. *Journal of Neuroscience*, 14(12):7347–7356, 1994.

[164] J.F. Kalaska, R. Caminiti, and A.P. Georgopoulos. Cortical mechanisms related to the direction of two-dimensional arm movements: Relations in parietal area 5 and comparison with motor cortex. *Experimental Brain Research*, 51:247–260, 1983.

[165] K. Kishimoto and S.-I. Amari. Existence and stability of local excitations in homogeneous neural fields. *Journal of Mathematical Biology*, 7(4):303–318, 1979.

[166] J.J. Knierim, H.S. Kudrimoti, and B.L. McNaughton. Place cells, head direction cells, and the learning of landmark stability. *Journal of Neuroscience*, 15:1648–1659, 1995.

[167] J.J. Knierim, H.S. Kudrimoti, and B.L. McNaughton. Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells. *Journal of Neurophysiology*, 80:425–446, 1998.

[168] J.J. Knierim, H.S. Kudrimoti, W.E. Skaggs, and B.L. McNaughton. The interaction between vestibular cues and visual landmark learning in spatial navigation. In T. Ono, B.L. McNaughton, S. Molotchnikoff, E. Rolls, and H. Nishijo, editors, *Perception, Memory, and Emotion: Frontiers in Neuroscience*, pages 343–357. Oxford: Pergamon Press, 1996.

[169] T. Kobayashi, H. Nishijo, M. Fukuda, J. Bures, and T. Ono. Task-dependent representation in rat hippocampal place neurons. *Journal of Neurophysiology*, 78:597–613, 1997.

[170] C. Kohler. Intrinsic connections of the retrohippocampal region in the rat brain. II. The medial entorhinal area. *Journal of Comparative Neurology*, 246:149–169, 1986.

[171] C. Kohler. Intrinsic connections of the retrohippocampal region in the rat brain. III. The lateral entorhinal area. *Journal of Comparative Neurology*, 271:208–228, 1988.

[172] T. Kohonen. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69, 1982.

[173] T. Kohonen. *Self-organization and associative memories*. Springer-Verlag, Berlin, 1984.

[174] B. Kolb. Posterior parietal and temporal association cortex. In B. Kolb and R.C. Tees, editors, *The Cerebral Cortex of the Rat*, chapter 19, pages 459–471. MIT Press, Cambridge MA, 1990.

[175] K.C. Kosel, G.W. Van Hoesen, and J.R. West. Olfactory bulb projections to the parahippocampal area of the rat. *Journal of Comparative Neurology*, 198:467–482, 1981.

[176] B. Kroese. Special issue on reinforcement learning and robotics. *Robotics and Autonomous Systems*, 15, 1995.

[177] J.L. Kubie and J.B. Ranck. Sensory-behavioral correlates in individual hippocampus neurons in three situations: Space and context. In W. Seifert, editor, *Neurobiology of the Hippocampus*, pages 433–447. Academic Press, New York, 1983.

[178] B.J. Kuipers and Y.T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems*, 8:47–63, 1991.

[179] C. Lee, W.H. Rohrer, and D.L. Sparks. Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature*, 332:357–360, 1988.

[180] B. Leonard and B.L. McNaughton. Spatial representation in the rat: Conceptual, behavioral, and neurophysiological perspectives. In R.P. Kesner and D.S. Olton, editors, *Neurobiology of Comparative Cognition*, pages 363–422. Erlbaum, Hillsdale, NJ, 1990.

[181] C.L. Leonhard, R.W. Stackman, and J.S. Taube. Head direction cells recorded from the lateral mammillary nuclei in rats. *Society for Neuroscience Abstracts*, 22:1873, 1996.

[182] W.B. Levy. A computational approach to hippocampal function. In R.D. Hawkins and G.H. Bower, editors, *Computational Models of Learning in Simple Neural Systems*, pages 243–305. Academic Press, San Diego, CA, 1989.

[183] W.B. Levy and D. Steward. Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Journal of Neuroscience*, 8:791–797, 1983.

[184] W.B. Levy, X. Wu, and R.A. Baxter. Unification of hippocampal function via computational/encoding considerations. *International Journal of Neural Systems*, 6:71–80, 1995.

[185] F.T. Lewis. The significance of the term *hippocampus*. *Journal of Comparative Neurology*, 35:213–230, 1923.

[186] G. Li. *Towards on-line learning agent for autonomous navigation*. PhD thesis, Chalmers University of Technology, Göteborg, Sweden, 1999.

[187] L. Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning*, 8:293–321, 1992.

[188] O. Lindvall. Dopamine pathways in the rat brain. In A.S. Horn, J. Korb, and B.H.C. Westerink, editors, *The neurobiology of dopamine*, pages 109–120. New York, Academic Press, 1979.

[189] T. Ljungberg, P. Apicella, and W. Schultz. Responses of monkey midbrain dopamine neurons during delayed alternation performance. *Brain Research*, 586:337–341, 1991.

[190] P. Maes. Behaviour-based artificial intelligence. In J. A. Meyer H. L. Roiblat and S. W. Wilson, editors, *From Animals to Animats II: Proceedings of the Second International Conference on Simulation of Adaptive Behaviour*, Cambridge, Massachusetts, 1993. MIT Press-Bradford Books.

[191] P. Maes, M. Mataric, and J.A. Meyer. *From animals to animats 4 - SAB96*. MIT Press-Bradford Books, Cambridge, Massachusetts, 1996.

[192] E.A. Maguire. Hippocampal and parietal involvement in human topographical memory: Evidence from functional neuroimaging. In K.J. Jeffery N. Burgess and J. O'Keefe, editors, *he Hippocampal and Parietal Foundations of Spatial Cognition*, chapter 21, pages 404–415. Oxford University Press, 1999.

[193] H.A. Mallot, M.O. Franz, B. Schölkopf, and H.H. Bülthoff. The view–graph approach to visual navigation and spatial memory. In W. Gerstner, A. Germond, M. Hasler, and J.D. Nicoud, editors, *Artificial Neural Networks - ICANN'97. 7th International Conference*, pages 751–756, Lausanne, Switzerland, 1997. Springer Verlag.

[194] E.J. Markus, C.A. Barnes, B.L. McNaughton, V.L. Gladden, and W.E. Skaggs. Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus*, 4:410–421, 1994.

[195] E.J. Markus, Y. Qin, B. Leonard, W.E. Skaggs, B.L. McNaughton, and C.A. Barnes. Interactions between location and task affect the spatial and direction firing of hippocampal neurons. *Journal of Neuroscience*, 15:7079–7094, 1995.

[196] D. Marr. Simple memory: A theory for archicortex. *Phil. Trans. of the Royal Society [Biol]*, 176:23–81, 1971.

[197] M.J. Mataric. Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation*, 8:304–312, 1992.

[198] B.L. Matthews, K.A. Campbell, and S.A. Deadwyler. Rotational stimulation disrupts spatial learning in fornix-lesioned rats. *Behavioral Neuroscience*, 102:35–42, 1988.

[199] R. Maurer. *L'intégration du chemin ou navigation à l'estime chez l'animal: Formalisation de modèles neuromimétiques*. PhD thesis, University of Geneva, Switzerland, 1993.

[200] R. Maurer. A connectionist model of path integration with and without a representation of distance to the starting point. *Psychobiology*, 26(1):21–35, 1998.

[201] R. Maurer and V. Seguinot. What is modelling for? A critical review of the models of path integration. *Journal of Theoretical Biology*, 175:457–475, 1995.

[202] R.A. McCallum. Hidden state and reinforcement learning with instance-based state identification. *IEEE Systems, Man, and Cybernetics*, 26(3):464–473, 1996.

[203] R.J. McDonald and N.M. White. Parallel information processing in the water maze: Evidence for independent memory systems involving dorsal striatum and hippocampus. *Behavioral and Neural Biology*, 61:260–270, 1994.

[204] D.J. McFarland and T. Boesser. *Intelligent behavior in animals and robots*. MIT Press-Bradford Books, Cambridge, Massachusetts, 1993.

[205] B.L. McNaughton. Neural mechanisms for spatial computation and information storage. In L. Nadel, L. Cooper, P. Culicover, and R.M. Harnish, editors, *Neural Connections, Mental Computation*, chapter 9, pages 285–350. MIT Press, Cambridge, MA, 1989.

[206] B.L. McNaughton. The mechanism of expression of long-term enhancement of hippocampal synapses: Current issues and theoretical implications. *Annual Review of Physiology*, 55:375–396, 1993.

[207] B.L. McNaughton, C.A. Barnes, J.L. Gerrard, K. Gothard, M.W. Jung, J.J. Knierim, H. Kudrimoti, Y. Qin, W.E. Skaggs, M. Suster, and K.L. Weaver. Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *Journal of Experimental Biology*, 199:173–185, 1996.

[208] B.L. McNaughton, C.A. Barnes, J. Meltzer, and R.J. Sutherland. Hippocampal granule cells are necessary for normal spatial learning but not for spatially-selective pyramidal cell discharge. *Experimental Brain Research*, 76:485–496, 1989.

[209] B.L. McNaughton, C.A. Barnes, and J. O'Keefe. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 52:41–49, 1983.

[210] B.L. McNaughton, L.L. Chen, and E.J. Markus. Dead reckoning, landmark learning, and the sense of direction: A neurophysiological and computational hypothesis. *Journal of Cognitive Neuroscience*, 3:190, 1991.

[211] B.L. McNaughton, J.J. Knierim, and M.A. Wilson. Vector encoding and the vestibular foundations of spatial cognition: Neurophysiological and computational mechanisms. In M. Gazzaniga, editor, *The Cognitive Neuroscience*, chapter 37, pages 585–595. MIT Press, Boston, 1994.

[212] B.L. McNaughton, B. Leonard, and L. Chen. Cortical-hippocampal interactions and cognitive mapping: A hypothesis based on reintegration of the parietal and inferotemporal pathways for visual processing. *Psychobiology*, 17(3):230–235, 1989.

[213] B.L. McNaughton, S.J.Y. Mizumori, C.A. Barnes, B.J. Leonard, M. Marquis, and E.J. Green. Cortical representation of motion during unrestrained spatial navigation in the rat. *Cerebral Cortex*, 4:27–39, 1994.

[214] B.L. McNaughton and R.G.M. Morris. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, 10(10):408–415, 1987.

[215] M.R. Mehta, C.A. Barnes, and B.L. McNaughton. Experience-dependent, asymmetric expansion of hippocampal place fields. In *Proceedings Natl. Acad. Science USA*, volume 94, pages 8918–8921, 1997.

[216] M.R. Mehta, M.C. Quirk, and M.A. Wilson. Experience-dependent, asymmetric shape of hippocampal receptive fields. *Neuron*, 25:707–715, 2000.

[217] J.A. Meyer, H.L. Roitblat, and S.W. Wilson. *From animals to animats 2 - SAB92*. MIT Press-Bradford Books, Cambridge, Massachusetts, 1993.

[218] J.A. Meyer and S.W. Wilson. *From animals to animats - SAB90*. MIT Press-Bradford Books, Cambridge, Massachusetts, 1991.

[219] R. Miller. *Cortico-Hippocampal interplay and the representation of contexts in the brain*. Springer-Verlag, 1991.

[220] V.M. Miller and P.J. Best. Spatial correlates of hippocampal unit activity are altered by lesions of the fornix and entorhinal cortex. *Brain Research*, 194:311–323, 1980.

[221] H. Mittelstaedt. Control systems of orientation in insects. *Annual Review of Entemology*, 7:177–198, 1962.

[222] H. Mittelstaedt and M.L. Mittelstaedt. Mechanismen der orientierung ohne richtende aussenreize. *Fortschritte der Zoologie*, 21:46–58, 1973.

[223] H. Mittelstaedt and M.L. Mittelstaedt. Homing by path integration. In F. Papi and H.G. Wallraff, editors, *Avian navigation*. Springer, Berlin Heidelberg, 1982.

[224] M.L. Mittelstaedt and H. Mittelstaedt. Homing by path integration in a mammal. *Naturwissenschaften*, 67:566–567, 1980.

[225] H. Mitterlstaed. The role of multimodal convergence in homing by path integration. *Fortschritte der Zoologie*, 28:197–212, 1983.

[226] S.J.Y. Mizumori and B.G. Cooper. Spatial representations of dorsal caudate neurons of freely-behaving rats. *Society for Neuroscience Abstracts*, 21:1929, 1995.

[227] S.J.Y. Mizumori, A.M. Lavoie, and A. Kalyani. Redistribution of spatial representation in the hippocampus of aged rats performing a spatial memory task. *Behavioral Neuroscience*, 110(5):1006–1016, 1996.

[228] S.J.Y. Mizumori, K.E. Ward, and A.M. Lavoie. Medial septal modulation of entorhinal single unit activity in anesthetized and freely moving rats. *Brain Research*, 570:188–197, 1992.

[229] S.J.Y. Mizumori and J.D. Williams. Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *Journal of Neuroscience*, 13:4015–4028, 1993.

[230] F. Mondada, E. Franzi, and P. Ienne. Mobile robot miniaturization: A tool for investigation in control algorithms. In T. Yoshikawa and F. Miyazaki, editors, *Proceedings of the Third International Symposium on Experimental Robotics*, pages 501–513, Tokyo, 1993. Springer Verlag.

[231] H.P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, 9:61–74, 1988.

[232] H.P. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 116–121, 1985.

[233] R.G.M. Morris. Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12:239–260, 1981.

[234] R.G.M. Morris, E. Anderson, G.S. Lynch, and M. Baudry. Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antogonist, AP5. *Nature*, 319:774–776, 1986.

[235] R.G.M. Morris and U. Frey. Hippocampal synaptic plasticity: Role in spatial learning or the automatic recording of attended experience? In K.J. Jeffery N. Burgess and J. O'Keefe, editors, *The Hippocampal and Parietal Foundations of Spatial Cognition*, chapter 12, pages 220–246. Oxford University Press, 1999.

[236] R.G.M. Morris, P. Garrud, J.N.P. Rawlins, and J. O'Keefe. Place navigation impaired in rats with hippocampal lesions. *Nature*, 297:681–683, 1982.

[237] R.G.M. Morris, F. Schenk, F. Tweedie, and L.E. Jarrard. Ibotenate lesions of the hippocampus and/or subiculum: Dissociating components of allocentric spatial learning. *European Journal of Neuroscience*, 2:1016–1028, 1990.

[238] M. Müller and R. Wehner. Path integration in desert ants, *Cataglyphis fortis*. *Proc. Acad. Sci. U.S.A.*, 85:5287–5290, 1988.

[239] R.U. Muller, E. Bostock, J.S. Taube, and J.L. Kubie. On the directional firing properties of hippocampal place cells. *Journal of Neuroscience*, 14(12):7235–7251, 1994.

[240] R.U. Muller and J.L. Kubie. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7:1951–1968, 1987.

[241] R.U. Muller and J.L. Kubie. The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, 9:4101–4110, 1989.

[242] R.U. Muller, J.L. Kubie, E.M. Bostock, J.S. Taube, and G.J. Quirk. Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In J. Paillard, editor, *Brain and Space*, chapter 17, pages 296–333. Oxford University Press, New York, 1991.

[243] R.U. Muller, J.L. Kubie, and J.J.B. Ranck. Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *Journal of Neuroscience*, 7:1935–1950, 1987.

[244] R.U. Muller, J.L. Kubie, and R. Saypoff. The hippocampus as a cognitive graph. *Hippocampus*, 1(3):243–246, 1991.

[245] J.J. Murphy. A mechanical analogy. *Nature*, 7:483, 1873.

[246] L. Nadel and J. Willner. Context and conditioning: A place for space. *Physiological Psychology*, 8(2):218–228, 1980.

[247] L. Nadel, J. Willner, and E.M. Kurz. Cognitive maps and environmental context. In P. Balsam and A. Tomie, editors, *Context and Learning*, pages 385–406. Erlbaum, Hillsdale, NJ, 1985.

[248] S. Nolfi and D. Floreano. *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-organizing Machines*. MIT Press, Cambridge, MA, 2000.

[249] E. Oja. A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15:267–273, 1982.

[250] E. Oja. Neural networks, principal components, and subspaces. *International Journal of Neural Systems*, 1:61–68, 1989.

[251] J. O'Keefe. Place unit in the hippocampus of the freely moving rat. *Experimental Neurology*, 51:78–109, 1976.

[252] J. O'Keefe. An allocentric spatial model for the hippocampal cognitive map. *Hippocampus*, 1(3):230–235, 1991.

[253] J. O'Keefe and N. Burgess. Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381:425–428, 1996.

[254] J. O'Keefe and D.H. Conway. Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental Brain Research*, 31:573–590, 1978.

[255] J. O'Keefe and J. Dostrovsky. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 34:171–175, 1971.

[256] J. O'Keefe and L. Nadel. *The Hippocampus as a cognitive map*. Clarendon Press, Oxford, 1978.

[257] J. O'Keefe and M. Recce. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3:317–330, 1993.

[258] J. O'Keefe and A. Speakman. Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, 68:1–27, 1987.

[259] D.S. Olton, J.T. Becker, and G.E. Handelmann. Hippocampus, space, and memory. *Behavioral and Brain Sciences*, 2:313–322, 1979.

[260] D.S. Olton, M. Branch, and P.J. Best. Spatial correlates of hippocampal unit activity. *Experimental Neurology*, 58:387–409, 1978.

[261] D.S. Olton, J.A. Walker, and W.A. Wolf. A disconnection analysis of hippocampal function. *Brain Research*, 233:241–253, 1982.

[262] T. Otto, C. Ding, G. Cousens, and K. Schiller. Effects of lateral vs. medial entorhinal cortex aspiration on the acquisition of odor-place associations. *Society for Neuroscience Abstracts*, 22:1120, 1996.

[263] T. Otto and H. Eichenbaum. Neuronal activity in the hippocampus during delayed non-match to sample performance in rats: Evidence for hippocampal processing in recognition memory. *Hippocampus*, 2(3):323–334, 1992.

[264] M.G. Packard and J.L. McGaugh. Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience*, 106(3):439–446, 1992.

[265] J.W. Papez. A proposed mechanism of emotion. *Archives of Neurology and Psychiatry*, 38:728–744, 1937.

[266] C. Pavlides and J. Winson. Influence of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *Journal of Neuroscience*, 9(8):2907–2918, 1989.

[267] J.M. Pearce, A.D.L. Roberts, and M. Good. Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. *Nature*, 396:75–77, 1998.

[268] R. Pfeifer. Cognition-perspectives from autonomous agents. *Robotics and Autonomous Systems*, 15:47–70, 1995.

[269] R. Pfeifer. Building "fungus eaters": Design principles of autonomous agents. In M. Mataric and J. A. Meyer, editors, *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, Cambridge, Massachusetts, 1996. MIT Press-Bradford Books.

[270] R. Pfeifer and C. Scheier. *Understanding Intelligence*. MIT Press, Cambridge, MA, 1999.

[271] R.M. Pico, L.K. Gerbrandt, M. Pondel, and G. Ivy. During stepwise cue deletion, rat place behaviors correlate with place unit responses. *Brain Research*, 330:369–373, 1985.

[272] M. Potegal. Vestibular and neostriatal contributions to spatial orientation. In M. Potegal, editor, *Spatial abilities (Development and physiological foundations)*, pages 361–387. Academic Press, New York, 1982.

[273] M. Potegal. The vestibular navigation hypothesis: A progress report. In P. Ellen and C. Thinus-Blanc, editors, *Cognitive Processes and Spatial Orentation in Animal and Man*, volume 2. Dordrecht: Martinus Nijhoff, 1987.

[274] B. Poucet. Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review*, 100:163–182, 1993.

[275] B. Poucet, C. Thinus-Blanc, and R.U. Muller. Place cells in the ventral hippocampus of rats. *NeuroReport*, 5(16):2045–2048, 1994.

[276] T.J. Prescott. Spatial learning and representation in animats. In J. Meyer, H.L. Roiblat, and S.W. Wilson, editors, *Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*, pages 164–173, Cambridge, MA, 1994. MIT Press-Bradford Books.

[277] G.J. Quirk, R.U. Muller, and J.L. Kubie. The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience*, 10(6):2008–2017, 1990.

[278] G.J. Quirk, R.U. Muller, J.L. Kubie, and J.B.Jr. Ranck. The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *Journal of Neuroscience*, 12(5):1945–1963, 1992.

[279] J.B.Jr. Ranck. Head-direction cells in the deep cell layers of dorsal presubiculum in freely moving rats. *Society for Neuroscience Abstracts*, 10:599, 1984.

[280] M. Recce and K.D. Harris. Memory for places: a navigational model in support of Marr's theory of hippocampal function. *Hippocampus*, 6:735–748, 1996.

[281] M. Recce, A. Speakman, and J. O'Keefe. Place fields of single hippocampal cells are smaller and more spatially localized than you thought. *Society of Neuroscience Abstracts*, 17:484, 1991.

[282] A.D. Redish. *Beyond the cognitive map*. PhD thesis, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1997.

[283] A.D. Redish. *Beyond the Cognitive Map, From Place Cells to Episodic Memory*. MIT Press-Bradford Books, London, 1999.

[284] A.D. Redish, A.N. Elga, and D.S. Touretzky. A coupled attractor model of the rodent head direction system. *Network*, 7(4):671–685, 1996.

[285] A.D. Redish and D.S. Touretzky. Cognitive maps beyond the hippocampus. *Hippocampus*, 7(1):15–35, 1997.

[286] A.D. Redish and D.S. Touretzky. Navigating with landmarks: Computing goal locations from place codes. In K. Ikeuchi and M. Veloso, editors, *Symbolic Visual Learning*, chapter 12, pages 325–351. Oxford University Press, 1997.

[287] F. Restle. Discrimination of cues in mazes: A resolution of the 'place-vs-response' question. *Psychological Review*, 64(4):217–228, 1957.

[288] E.T. Rolls. Functions of neuronal networks in the hippocampus and cerebral cortex in memory. In R.M.J. Cotterill, editor, *Models of brain function*, pages 15–33. Cambridge University Press, 1989.

[289] E.T. Rolls. A model of the operation of the hippocampus and entorhinal cortex in memory. *International Journal of Neural Systems*, 6:51–70, 1995.

[290] E.T. Rolls and S. O'Mara. Neurophysiological and theoretical analysis of how the hippocampus functions in memory. In T. Ono, L.R. Squire, M. Raichle, D. Perrett, and M. Fukuda, editors, *Brain mechanisms of perception: From neuron to behavior*, pages 276–300. Oxford University Press, New York, 1993.

[291] E.T. Rolls and M.J. Tovée. Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *Journal of Neurophysiology*, 73:713–726, 1995.

[292] A. Rotenberg, M. Mayford, R.D. Hawkins, E.R. Kandel, and R.U. Muller. Mice expressing activated CaMKII lack low frequency LTP and do not form stable place cells in the CA1 region of the hippocampus. *Cell*, 87:1351–1361, 1996.

[293] A. Rotenberg and R.U. Muller. Variable place cell coupling to a continuously viewed stimulus: Evidence that the hippocampus acts as a perceptual system. *Royal Society Philosophical Transactions: Biological Sciences*, 352(1360):1505–1514, 1997.

[294] D.E. Rumelhart, G.E. Hinton, and R.J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986.

[295] D.E. Rumelhart and D. Zipser. Feature discovery by competitive learning. *Cognitive Science*, 9:75–112, 1985.

[296] D.E. Rumelhart and D. Zipser. Feature discovery by competitive learning. In D.E. Rumelhart and J.L. McClelland, editors, *Parallel distributed processing 1*, Cambridge, Massachusetts, 1986. MIT Press.

[297] G.A. Rummery and M. Niranjan. On-line Q-learning using connectionist systems. Technical Report CUED/F-INFENG/TR 166, Engineering Department, Cambridge University, 1994.

[298] Y. Sakurai. Involvement of auditory cortical and hippocampal neurons in auditory working memory and reference memory in the rat. *Journal of Neuroscience*, 14(5):2606–2623, 1994.

[299] E. Salinas and L.F. Abbott. Vector reconstruction from firing rates. *Journal of Computational Science*, 1:89–107, 1994.

[300] A. Samsonovich. *Attractor-map theory of the hippocampal representation of space*. PhD thesis, The University of Arizona, 1997.

[301] A. Samsonovich and B.L. McNaughton. Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, 17(15):5900–5920, 1997.

[302] T.D. Sanger. Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural Networks*, 2:459–473, 1989.

[303] J.C. Santamaría, R.S. Sutton, and A. Ram. Experiments with reinforcement learning in problems with continuous state and action spaces. *Adaptive Behavior*, 6(2):163–218, 1998.

[304] J.P. Sauvé. *L'orientation spatiale: formalisation d'un modèle de mémorisation égocentrée et expérimentation chez l'homme*. PhD thesis, Université d'Aix-Marseille II, 1989.

[305] E. Save, A. Cressant, C. Thinus-Blanc, and B. Poucet. Early visual deprivation does not prevent hippocampal place cell firing in the rat. *Society for Neuroscience Abstracts*, 22:912, 1996.

[306] E. Save, A. Cressant, C. Thinus-Blanc, and B. Poucet. Spatial firing of hippocampal place cells in blind rats. *Journal of Neuroscience*, 18(5):1818–1826, 1998.

[307] F. Schenk and R.G.M. Morris. Dissociation between components of spatial memory in rats after recovery from the effects of retrohippocampal lesions. *Experimental Brain Research*, 58:11–28, 1985.

[308] B. Schölkopf and H.A. Mallot. View-based cognitive mapping and path planning. *Adaptive Behavior*, 3:311–348, 1995.

[309] H. Schöne. *Spatial Orientation*. Princeton, NJ, 1984.

[310] W. Schultz. Reward responses of dopamine neurons: A biological reinforcement signal. In W. Gerstner, A. Germond, M. Hasler, and J.D. Nicoud, editors, *Artificial Neural Networks - ICANN'97. 7th International Conference*, pages 3–12, Lausanne, Switzerland, 1997. Springer Verlag.

[311] W. Schultz, P. Apicella, E. Scarnati, and T. Ljungberg. Neuronal activity in monkey ventral striatum related to the expectation of reward. *Journal of Neuroscience*, 12:4595–4610, 1992.

[312] W. Schultz, P. Dayan, and R.R. Montague. A neural substrate of prediction and reward. *Science*, 275:1593–1599, 1997.

[313] V. Séguinot, J. Cattet, and S. Benhamou. Path integration in dogs. *Animal Behavior*, 55:787–797, 1998.

[314] V. Séguinot, R. Maurer, and A.S. Etienne. Dead reckoning in a small mammal: The evaluation of distance. *Journal of Comparative Physiology A*, 173:103–113, 1993.

[315] S.R. Sesack and V.M. Pickel. In the rat medial nucleus accumbens, hippocampal and catecholaminergic terminals converge on spiny neurons and are in apposition to each other. *Brain Research*, 527(2):266–279, 1990.

[316] E.-A. Seyfarth, R. Hergenröder, H. Ebbes, and F.G. Barth. Idiothetic orientation of a wandering spider: Compensation of detours and estimates of goal distance. *Behavioral Ecology and Sociobiology*, 11:139–148, 1982.

[317] P.E. Sharp. Computer simulation of hippocampal place cells. *Psychobiology*, 19(2):103–115, 1991.

[318] P.E. Sharp. Multiple spatial/behavioral correlates for cells in the rat postsubiculum: Multiple regression analysis and comparison to other hippocampal areas. *Cerebral Cortex*, 6(2):238–259, 1996.

[319] P.E. Sharp. Subicular cells generate similar spatial firing patterns in two geometrically and visually distinctive environments: Comparison with hippocampal place cells. *Behavioral and Brain Research*, 85:71–92, 1997.

[320] P.E. Sharp. Complimentary roles for hippocampal versus subicular/entorhinal place cells in coding place, context, and events. *Hippocampus*, 9:432–443, 1999.

[321] P.E. Sharp. Subicular place cells expand/contract their spatial firing pattern to fit the size of the environment in an open field, but not inthe presence of barriers: Comparison with hippocampal place cells. *Behavioral Neuroscience*, ??:??, 1999.

[322] P.E. Sharp, H.T. Blair, D. Etkin, and D.B. Tzanetos. Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. *Journal of Neuroscience*, 15:173–189, 1995.

[323] P.E. Sharp and C. Green. Spatial correlates of firing patterns of single cells in the subiculum of freely moving rat. *Journal of Neuroscience*, 14(4):2339–2356, 1994.

[324] P.E. Sharp, J.L. Kubie, and R.U. Muller. Firing properties of hippocampal neurons in a visually symmetrical environment: Contributions of multiple sensory cues and mnemonic processes. *Journal of Neuroscience*, 10(9):3093–3105, 1990.

[325] G.M. Shepherd and C. Koch. Introduction to synaptic circuits. In G.M. Shepherd, editor, *The Synaptic Organization of the Brain*, pages 3–31. Oxford University Press, New York, 1990.

[326] A.J. Silva, R. Paylor, J.M. Wehner, and S. Tonegawa. Impaired spatial learning in α-calcium-calmodulin kinase II mutant mice. *Science*, 257:206–211, 1992.

[327] W.E. Skaggs, J.J. Knierim, H.S. Kudrimoti, and B.L. McNaughton. A model of the neural basis of the rat's sense of direction. In G. Tesauro, D.S. Touretzky, and T.K. Leen, editors, *Advances in Neural Information Processing Systems 7*, pages 173–180, Cambridge, MA, 1995. MIT Press.

[328] W.E. Skaggs and B.L. McNaughton. The hippocampal theta rhythm and memory for temporal sequences of events. *International Journal of Neural Systems*, 6:101–105, 1995.

[329] W.E. Skaggs and B.L. McNaughton. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, 271:1870–1873, 1996.

[330] W.E. Skaggs, B.L. McNaughton, M.A. Wilson, and C.A. Barnes. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2):149–173, 1996.

[331] F. Smeraldi. *Attention-driven pattern recognition*. PhD thesis, Swiss Federal Institute of Technology, EPFL, Lausanne, Switzerland, 2000.

[332] F. Smeraldi and J. Bigün. Retinal vision applied to facial features detection and face authentication. *IEEE-PAMI Transactions on Pattern Analysis and Machine Intelligence (submitted)*, 1999.

[333] F. Smeraldi, J. Bigün, and W. Gerstner. Face space dimensionality in the context of face authentication. *IEEE Transactions on Pattern Analysis and Machine Intelligence (submitted)*, 1999.

[334] A. Speakman and J. O'Keefe. Hippocampal complex spike cells do not change their place fields if the goal is moved within a cue controlled environment. *European Journal of Neuroscience*, 2:544–555, 1990.

[335] R.W. Stackman and J.S. Taube. Influence of vestibular system lesions upon anterior thalamic head-direction cell activity. *Society for Neuroscience Abstracts*, 21:945, 1995.

[336] R.W. Stackman and J.S. Taube. Firing properties of rat lateral mammillary single units: Head direction, head pitch, and angular head velocity. *Journal of Neuroscience*, 18(21):9020–9037, 1998.

[337] R.J. Steele and R.G.M. Morris. Delay-dependent impairment of a matching-to-placetask with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus*, 9:118–136, 1999.

[338] L. Steels. Building agents out of autonomous behavior systems. In L. Steels and R. Brooks, editors, *The "artificial life" route to "artificial intelligence". Building situated embodied agents*, New Haven, 1993. Lawrence Erlbaum.

[339] M.A. Steinmets, B.C. Motter, C.J. Duffy, and V.B. Mountcastle. Functional properties of parietal visual neurons: Radial organization of directionalities with the visual field. *Journal of Neuroscience*, 7:177–191, 1987.

[340] R.J. Sutherland and J.M. Hoesing. Posterior cingulate cortex and spatial memory: A microlimnology analysis. In B.A. Vogt and M. Gabriel, editors, *Neurobiology of Cingulate Cortex and Limbic Thalamus: A Comprehensive Handbook*, pages 461–477. Birkhauser, Boston, 1993.

[341] R.J. Sutherland and A.J. Rodriguez. The role of the fornix/fimbria and some related subcortical structures in place learning and memory. *Behavioral and Brain Research*, 32:265–277, 1990.

[342] R.S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3:9–44, 1988.

[343] R.S. Sutton and A.G. Barto. *Reinforcement learning, an introduction*. MIT Press-Bradford Books, Cambridge, Massachusetts, 1998.

[344] L.W. Swanson and W.M. Cowan. An autoradiographic study of the organization of the efferent connections of the hippocampal formation in the rat. *Journal of Comparative Neurology*, 172:49–84, 1977.

[345] J. Tani and N. Fukumura. Learning goal-directed sensory-based navigation of a mobile robot. *Neural Networks*, 7:553–563, 1994.

[346] H. Tanila, P. Sipilä, M. Shapiro, and H. Eichenbaum. Brain aging: Impaired coding of novel environmental cues. *Journal of Neuroscience*, 17(13):5167–5174, 1997.

[347] J.S. Taube. Qualitative analysis of head-direction cells recorded in the rat anterior thalamus. *Society for Neuroscience Abstracts*, 18:708, 1992.

[348] J.S. Taube. Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, 15(1):1953–1971, 1995.

[349] J.S. Taube. Place cells recorded in the parasubiculum of freely moving rats. *Hippocampus*, 5(6):569–583, 1996.

[350] J.S. Taube and H.L. Burton. Head direction cell activity monitored in a novel environment and during a cue conflict situation. *Journal of Neurophysiology*, 74:1953–1971, 1995.

[351] J.S. Taube, J.P. Goodridge, E.J. Golob, P.A. Dudchenko, and R.W. Stackman. Processing the head direction cell signal: A review and commentary. *Brain Research Bulletin*, 40(5/6):477–486, 1996.

[352] J.S. Taube, J.P. Klessak, and C.W. Cotman. Lesions of the rat postsubiculum impair performance on spatial tasks. *Behavioral and Neural Biology*, 5:131–143, 1992.

[353] J.S. Taube, R.I. Muller, and J.B.Jr. Ranck. Head direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, 10:420–435, 1990.

[354] J.S. Taube, R.I. Muller, and J.B.Jr. Ranck. Head direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *Journal of Neuroscience*, 10:436–447, 1990.

[355] J.S. Taube and R.U. Muller. Head direction cell activity in the anterior thala- mic nuclei, but not the postsubiculum, predicts the animal's future directional heading. *Society for Neuroscience Abstracts*, 21:946, 1995.

[356] C.J. Taylor and D.J. Kriegman. Exploration strategies for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 2, pages 248–253, Los Alamitos, CA, 1993. IEEE Computer Society Press.

[357] L.T. Thompson and P.J. Best. Place cells and silent cells in the hippocampus of freely-behaving rats. *Journal of Neuroscience*, 9(7):2382–2390, 1989.

[358] S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. *Nature*, 381:520–522, 1996.

[359] S. Thrun. Learning maps for indoor mobile robot navigation. *Artificial Intelligence*, 99:21–71, 1998.

[360] M. Tistarelli and G. Sandini. On the advantages of log-polar mapping for direct estimation of time-to-impact from optical flow. *IEEE Transactions on Pattern Analysis and Machine Recognition*, 15(4):401–410, 1993.

[361] E.C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55:189–208, 1948.

[362] D.S. Touretzky and A.D. Redish. A theory of rodent navigation based on interacting representations of space. *Hippocampus*, 6(3):247–270, 1996.

[363] D.S. Touretzky, A.D. Redish, and H.S. Wan. Neural representation of space using sinusoidal arrays. *Neural Computation*, 5:869–884, 1993.

[364] M.J. Tovée. *An introduction to the visual system*. Cambridge University Press, 1996.

[365] A. Treves and E.T. Rolls. A computational analysis of the role of the hippocampus in memory. *Hippocampus*, 4:373–391, 1994.

[366] A. Treves, W.E. Skaggs, and C.A. Barnes. How much of the Hippocampus can be explained by functional constraints? *Hippocampus*, 6:666–674, 1996.

[367] O. Trullier. *Elaboration et traitment des représentations spatiales servant à la navigation chez le rat*. PhD thesis, Université Paris VI, Paris, 1998.

[368] O. Trullier and J.-A. Meyer. Animat navigation using a cognitive graph. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S. Wilson, editors, *From Animals to Animats V*, pages 213–222, Cambridge, MA, 1998. MIT Press.

[369] O. Trullier and J.-A. Meyer. Animat navigation using a cognitive graph. *Biological Cybernetics*, 83:271–285, 2000.

[370] O. Trullier and J.A. Meyer. Place sequence learning for navigation. In W. Gerstner, A. Germond, M. Hasler, and J.D. Nicoud, editors, *Artificial Neural Networks - ICANN'97. 7th International Conference*, pages 757–762, Lausanne, Switzerland, 1997. Springer Verlag.

[371] O. Trullier, S.I. Wiener, A. Berthoz, and J.A. Meyer. Biologically-based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51:483–544, 1997.

[372] J.Z. Tsien, P.T. Huerta, and S. Tonegawa. The essential role of hippocampal CA1 NMDA receptor dependent synaptic plasticity in spatial memory. *Cell*, 87:1327–1338, 1996.

[373] M. Tsodyks and T. Sejnowski. Associative memory and hippocampal place cells. *International Journal of Neural Systems*, 6:81–86, 1995.

[374] L.G. Ungerleider and M. Mishkin. Two cortical visual systems. In D.J. Ingle, M.A. Goodale, and R.J.W. Mansfield, editors, *Analysis of Visual Behavior*. MIT Press, Cambridge, MA, 1982.

[375] C.H. Vanderwolf. Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalography and Clinical Neurophysiology*, 26:407–418, 1969.

[376] G. Vanni-Mercier and M. Magnin. Single neuron activity related to natural vestibular stimulation in the cat's visual cortex. *Experimental Brain Research*, 45:451–455, 1982.

[377] H.S. Wan, D.S. Touretzky, and A.D. Redish. Towards a computational theory of rat navigation. In M. Mozer, P. Smolensky, D.S. Touretzky, J.Elman, and A. Weigend, editors, *Proceedings of the 1993 Connectionist Models Summer School*, pages 11–19, Hillsdale, NJ, Erlbaum, 1994.

[378] C.J.C.H. Watkins. *Learning from delayed rewards*. PhD thesis, University of Cambridge, England, 1989.

[379] R. Watt. *Understanding Vision*. Academic Press, London, 1991.

[380] R. Wehner. Himmelsnavigation bei insekten. Neurophysiologie und Verhalten. *Neujahrsblatt der Naturforschenden Gesellschaft Zürich*, 184:1–132, 1982.

[381] R. Wehner. Spatial organization of foraging behavior in individually searching desert ants, *Cataglyphis* (Sahara Desert) and *Ocymyrmex* (Namib Desert). In J.M Pasteels and J.-L. Deneubourg, editors, *From individual to collective behavior in social insects*, pages 15–42. Birkhäuser, Basel, 1987.

[382] R. Wehner. Arthropods. In F. Papi, editor, *Animal Homing*, pages 45–144. Chapman & Hall, London, 1992.

[383] R. Wehner and S. Rossel. The bee's celestial compass–a case of study in behavioural neurobiology. *Fortschritte der Zoologie*, 31:11–53, 1985.

[384] R. Wehner and S. Wehner. Path integration in desert ants: Approching a long-standing puzzle ininsect navigation. *Monitore Zool. Ital.*, 20(3):309–331, 1986.

[385] I.Q. Whishaw, J.C. Cassel, and L.E. Jarrard. Rats with fimbria-fornix lesions display a place response in a swimming pool: A dissociation between getting there and knowing where. *Journal of Neuroscience*, 15(8):5779–5788, 1995.

[386] I.Q. Whishaw and L.E. Jarrard. Evidence for extrahippocampal involvement in place learning and hippocampal involvement in path integration. *Hippocampus*, 6:513–524, 1996.

[387] I.Q. Whishaw and G. Mittleman. Hippocampal modulation of nucleus accumbens: Behavioral evidence from amphetamine-induced activity profiles. *Behavioral and Neural Biology*, 55:289–306, 1991.

[388] S.I. Wiener. Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *Journal of Neuroscience*, 13:3802–3817, 1993.

[389] S.I. Wiener. Spatial, behavioral and sensory correlates of hippocampal CA1 complex spike cell activity: Implications for information processing functions. *Progress in Neurobiology*, 49:335–361, 1996.

[390] S.I. Wiener, V.A. Korshunov, R. Garcia, and A. Berthoz. Inertial, substratal and landmark cue control of hippocampal CA1 place cell activity. *Eur. Journal of Neuroscience*, 7:2206–2219, 1995.

[391] S.I. Wiener, C.A. Paul, and H. Eichenbaum. Spatial and behavioral correlates of hippocampal neuronal activity. *Journal of Neuroscience*, 9(8):2737–2783, 1989.

[392] D.J. Willshaw and J.T. Buckingham. An assessment of Marr's theory of hippocampus as a temporary memory store. *Phil. Trans. of the Royal Society [Biol]*, 329:205–215, 1990.

[393] H.R. Wilson and J.D. Cowan. A mathematical theory of the functional dynamics of cortical and thalamic tissue. *Kybernetik*, 13:55–80, 1973.

[394] M.A. Wilson and B.L. McNaughton. Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055–1058, 1993.

[395] M.A. Wilson and B.L. McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265:676–679, 1994.

[396] J. Winson. Loss of hippocampal theta rhythm results in spatial memory deficits in the rat. *Science*, 201:160–163, 1978.

[397] M.P. Witter. Organization of the entorhinal-hippocampal system: A review of current anatomical data. *Hippocampus*, 3:33–44, 1993.

[398] M.P. Witter and H.J. Groenewegen. Laminar origin and septotemporal distribution of entorhinal and perirhinal projections to the hippocampus in the cat. *Journal of Comparative Neurology*, 224:371–385, 1984.

[399] M.P. Witter, H.J. Groenewegen, F.H. Lopes da Silva, and A.H.M. Lohman. Functional organization of the extrinsic and intrinsic circuitry of the parahippocampal region. *Progress in Neurobiology*, 33:161–253, 1989.

[400] M.P. Witter, R.H. Ostendorf, and H.J. Groenewegen. Heterogeneity in the dorsal subiculum of the rat. Distinc neuronal zones project to different cortical and subcortical targets. *European Journal of Neuroscience*, 2:718–725, 1990.

[401] J.M. Wyss and T. Van Groen. Connections between the retrosplenial cortex and the hippocampal formation in the rat: A review. *Hippocampus*, 2(1):1–12, 1992.

[402] M.F. Yeckel and T.W. Berger. Feedforward excitation of the hippocampus by afferents from the entorhinal cortex: Redefinition of the role of the trisynaptic pathway. In *Proceedings of the National Academy of Sciences*, volume 87, pages 5832–5836, USA, 1990.

[403] B.J. Young, G.D. Fox, and H. Eichenbaum. Correlates of hippocampal complex-spike cell activity in rats performing a nonspatial radial maze task. *Journal of Neuroscience*, 14(11):6553–6563, 1994.

[404] M.P. Young and S. Yamane. Sparse population coding of faces in the inferotemporal cortex. *Science*, 256:1327–1331, 1992.

[405] K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, 16(6):2112–2126, 1996.

[406] K. Zhang, I. Ginzburg, B.L. McNaughton, and T.J. Sejnowski. Interpreting neuronal population activity by reconstruction: A unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, 79:1017–1044, 1998.

[407] U.R. Zimmer. Robust world-modeling and navigation in a real world. *Neurocomputing*, 13:2–4, 1996.

Curriculum Vitæ

Angelo ARLEO
Chemin des Clos, 53
1024 Ecublens
Tel. +41 21 691 7168
angelo.arleo@epfl.ch
Born 17th August 1969 in S. Lombardo, Varese, Italy

Summary

Research interests in spatial learning and navigation in biological and neuro-mimetic systems, autonomous mobile robotics, computer vision, artificial neural networks.

Education

1997–2000 **Ph.D. in Computer Science**. Centre for Neuro-Mimetic Systems (MANTRA), Microprocessors and Interface Lab (LAMI), Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland.
Focus: Modeling mechanisms underlying animals' spatial learning.

1996 **Degree project on Robotics and System Control**. Institute for Systems, Informatics and Safety. Joint Research Centre of the European Union, Ispra, Italy.
Focus: Connectionist models, adaptive robotics, and artificial intelligence.

1991–1996 **Master's Degree (with distinction) in Computer Science**. University of Mathematical Science of Milan, Italy.

Professional Experience

1997–2000 (sept.) (dec.) Research and teaching assistant at the Centre for Neuro-Mimetic Systems at the EPFL, Lausanne, Switzerland.
Functions: Member of the Robot Learning Group. Supervisor for ungraduate student projects. Assistant of the neural networks course. Web-master for the lab web-page.

1997–1997 (jan.) (aug.) Software Engineer at Nomadic Technologies Inc., Robots Manufacturer, Palo Alto, California, USA.
Functions: Designing a communication protocol between a mobile robot and a laser system. Developing a control graphical interface for a mobile robotic platform.

1995–1996 (sept.) (dec.) Research Assistant at the Institute for Systems, Informatics and Safety, Joint Research Centre of the EU, Ispra, Italy.
Functions: Designing a map-learning neural system for robot indoor navigation.

1988–1991 (oct.) (oct.) Design Engineer at Agusta S.p.A., Aeronautics Industry, Varese, Italy.
Functions: Designing electric-avionic circuits, helicopter area.

Languages

Italian: Native.
English: Read, write, and speak fluently.
French: Read, write, and speak fluently.
Spanish: Basic knowledge.

Computer Skills

Languages: C/C++, LISP, Prolog, Pascal, HTML, XMotif, Assembler 80x86.
Operating Systems: UNIX/Linux, DOS, Windows 95/NT.
Software tools: Matlab, CADAM.

Miscellaneous Activities

- Reviewer for the Journal of Biological Cybernetics.
- Member of the Program Committees of GECCO'99 (Orlando, Florida).

Publications

Journal Papers

A. Arleo and W. Gerstner. Spatial cognition and neuro-mimetic navigation: A model of hippocampal place cell activity. *Biological Cybernetics, Special Issue on Navigation in Biological and Artificial Systems*, 83: 287–299, 2000.

A. Arleo, J. del R. Millán, and D. Floreano. Efficient learning of variable-resolution cognitive maps for autonomous indoor navigation. *IEEE Transactions on Robotics and Automation*, 15(6): 990-1000, 1999.

A. Arleo and W. Gerstner. Spatial orientation in navigating agents: Modeling head-direction cells. *Neurocomputing (to appear)*, 2001.

Peer-reviewed Conference Papers

A. Arleo and W. Gerstner. Place Cells and Spatial Navigation based on Vision, Path Integration, and Reinforcement Learning. In *Advances in Neural Information Processing Systems 13*, MIT-Press, Denver, December 2000.

A. Arleo and W. Gerstner. Modeling rodent head-direction cells and place cells for spatial learning in bio-mimetic robotics. In J.-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S.W. Wilson, editors, *Proceedings of the 6th Int. Conf. on Simulation of Adaptive Behavior, from Animals to Animats 6 (SAB2000)*, pages 236-245, MIT Press, Cambridge MA, 2000.

A. Arleo and W. Gerstner. A vision-driven model of hippocampal place cells and temporally asymmetric LTP-induction for action learning. In *Proceedings of the 9th Int. Conf. on Artificial Neural Networks (ICANN'99)*, vol. 1, pages 132-137, IEE, London, 1999.

A. Arleo and W. Gerstner. Neuro-mimetic navigation systems: A computational model of the hippocampus. In A. Drogoul and J.-A. Meyer, editors, *Proceedings of the Conference on Situated Artificial Intelligence (IAS99)*, pages 193-211, Hermès, Paris, 1999.

J. del R. Millán and A. Arleo. Neural network learning of variable grid-based maps for the autonomous navigation of robots. In *Proceedings of the IEEE Int. Symposium on Computational Intelligence in Robotics and Automation*, pages 40-45, Monterey, CA, 1997.

Other Conference Papers

A. Arleo and W. Gerstner. Spatial models and autonomous navigation in neuro-mimetic systems. In *Third International Conference on Cognitive and Neural Systems*, Boston, MA, 1999.

A. Arleo and D. Floreano and W. Gerstner. Modélisation de l'hippocampe: Représentation spatiale et navigation des systèmes autonomes, In *10ème Journées des Jeunes Chercheurs en Robotique*, Amiens, France, 1998.